# ADAPTIVE MULTIPROCESSOR REAL-TIME SYSTEMS 

Aaron D. Block

A dissertation submitted to the faculty of the University of North Carolina at Chapel Hill in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Computer Science.

Chapel Hill
2008

Approved by:
James H. Anderson
Tarek Abdelzaher
Sanjoy Baruah
Gary Bishop
Kevin Jeffay
Stephen Quint
(c) 2008

Aaron D. Block

## ALL RIGHTS RESERVED

## ABSTRACT

## AARON D. BLOCK: Adaptive Multiprocessor Real-Time Systems <br> (Under the direction of James H. Anderson)

Over the past few years, as multicore technology has become cost-effective, multiprocessor systems have become increasingly prevalent. The growing availability of such systems has spurred the development of computationally-intensive applications for which single-processor designs are insufficient. Many such applications have timing constraints; such timing constraints are often not static, but may change in response to both external and internal stimuli. Examples of such applications include tracking systems and many multimedia applications. Motivated by these observations, this dissertation proposes several different adaptive scheduling algorithms that are capable of guaranteeing flexible timing constraints on multiprocessor platforms.

Under traditional task models (e.g., periodic, sporadic, etc.), the schedulability of a system is based on each task's worst-case execution time (WCET), which defines the maximum amount of time that each of its jobs can execute. The disadvantage of using WCETs is that systems may be deemed unschedulable even if they would function correctly most of the time when deployed. Adaptive real-time scheduling algorithms allow the timing constraints of applications to be adjusted based upon runtime conditions, instead of always using fixed timing constraints based upon WCETs. While there is a substantial body of prior work on scheduling applications with static timing constraints on multiprocessor systems, prior to this dissertation, no adaptive multiprocessor scheduling approach existed that is capable of ensuring bounded "error" (where error is measured by comparison to an ideal allocation).

In this dissertation, this limitation is addressed by proposing five different multiprocessor scheduling algorithms that allow a task's timing constraints to change at runtime. The five proposed adaptive algorithms are based on different non-adaptive multiprocessor scheduling algorithms that place different restrictions on task migrations and preemptions. The relative
advantages of these algorithms are compared by simulating both the Whisper human tracking system and the Virtual Exposure Camera (VEC), both of which were developed at The University of North Carolina at Chapel Hill. In addition, a feedback-based adaptive framework is proposed that not only allows timing constraints to adapt at runtime, but also detects which adaptions are needed. An implementation of this adaptive framework on a real-time multiprocessor testbed is discussed and its performance is evaluated by using the core operations of both Whisper and VEC. From this dissertation, it can be concluded that feedback and optimization techniques can be used to determine at runtime which adaptions are needed. Moreover, the accuracy of an adaptive algorithm can be improved by allowing more frequent task migrations and preemptions; however, this accuracy comes at the expense of higher migration and preemption costs, which impacts average-case performance. Thus, there is a tradeoff between accuracy and average-case performance that depends on the frequency of task migrations/preemptions and their cost.

To my wife, parents, and brother,
without whom I would only be three sevenths of the man I am today.

## ACKNOWLEDGMENTS

No document this large can be written without help from many people. I would first like to thank my advisor, James Anderson, who taught me how to research and write over six dyslexia-filled years. I cannot imagine a better advisor.

I would also like to than the rest of my committee: Tarek Abdelzaher, Sanjoy Baruah, Gary Bishop, Kevin Jeffay, and Stephen Quint. I am lucky to have such a qualified committee, and I deeply appreciate all of the help and feedback they have provided me over the years.

I am indebted to the many collegues with whom I have published over the years: Gary Bishop, Stephen Quint, Uma Devi, Björn Brandenburg, John Calandrino, Hennadiy Leontyev, Anand Srinivasan, Warren Davis, Russell Hamm, Sarah Knoop, and Peter Schwarz. I will always remember fondly the many crushed-red-pepper-pizza filled nights that I spent with you. Also, I owe many thanks to my other real-time colleagues: Shelby Funk, Phil Holman, and Nathan Fisher. I never wrote a paper with you, but I always wished I had.

So many of my friends (with whom I never published) helped to contribute to this dissertation and my life in general that it is hard to name them all, but here it goes (I apologize to any friends I may have omitted): Alex Higbee, Joel Heires, Craig Morris, Mike West, Austin Parker, the entire Haverford fencing team, Joanna Grayer, Julia Diepold, Jack and Shilpa McManus, Peter and Lori Adler, Galvin Chow, Chas Budnick, Harrison Breuer, Carl Knutson, Matthew Gjenvic, Jesse Milnes, Luv Kohli, Brian Eastwood, Chris VanderKnyff, Jeremy Wendt, Russell Gayle, Stephen Oliver, Sasa Junuzovic, Sandra Neely, Keith Lee, Jeff Terrell, and Avneesh Sud. Last, but not least, I want to thank my groomsmen, Adam Ruder, John Stevens "Terry" McMahon III, and Eric Bennett. You guys have been there for me in the best of times and the worst of times. I owe you more than I can imagine.

I would also like to thank my in-laws. Terri, I cannot conceptualize a more fun sister. Also, you are probably the best person I have ever known at keeping a surprise a secret. Brian
and Brenda, you guys have been there for me nearly every Friday night for the past six years. Thank you so much for always being there for me to kvetch about my week.

Next I would like to thank my brother. Stefan, you are the Platonic form of a brother to which all other brothers aspire in an attempt to reach that ideal. I want to apologize for any references to younger brothers as the PEDF scheduling algorithm; I felt they were necessary for the literary quality of the document. (An alternative joke would be "Along side this processor there is another, there are places where you can adapt.")

Penultimately, I would like to thank my parents, who (aside from teaching me transition words) taught me everything I know. Thank you so much for helping me grow from a nerdy little boy into a relatively less nerdy man. It is possible to mathematically prove that you are the best parents I could have.

Finally, I want to thank my wife. Nicki, you are the most wonderful wife I could have asked for. So, to quote Jerry Maguire, as I did in my wedding speech,"Show me the money." Wait, that doesn't sound right. What I want to say is that you enrich my life every day and always bring a smile to my face-I love you.

That's it. Enjoy the dissertation.

## TABLE OF CONTENTS

LIST OF TABLES ..... xiv
LIST OF FIGURES ..... xv
LIST OF ABBREVIATIONS ..... xx
1 INTRODUCTION ..... 1
1.1 Applications ..... 2
1.1.1 Whisper ..... 2
1.1.2 Virtual Exposure Camera ..... 3
1.2 Real-Time Systems ..... 3
1.2.1 Uniprocessor Systems ..... 6
1.2.2 Multiprocessor Partitioned Scheduling ..... 7
1.2.3 Restricted Global Multiprocessor Scheduling ..... 9
1.2.4 Unrestricted Global Multiprocessor Scheduling ..... 11
1.2.5 Impact of Migrations and Preemptions ..... 12
1.2.6 Research Needed ..... 14
1.3 Adaptivity ..... 16
1.3.1 Leave/join Reweighting ..... 16
1.3.2 Rate-Based Earliest Deadline ..... 16
1.3.3 Proportional Share Scheduling ..... 17
1.3.4 Adaptive Feedback-Based Frameworks ..... 18
1.4 Thesis Statement ..... 19
1.5 Contributions ..... 20
1.5.1 Adaptable Task Model and Reweighting Algorithms ..... 20
1.5.2 Reweighting Rules. ..... 21
1.5.3 Evaluating Algorithms ..... 22
1.5.4 AGEDF ..... 27
1.5.5 AGEDF Implementation ..... 28
1.6 Organization ..... 29
2 PRIOR WORK ..... 30
2.1 Leave/Join Reweighting ..... 30
2.2 Rate-Based Earliest Deadline ..... 31
2.3 Earliest Eligible Virtual Deadline First ..... 33
2.4 Feedback-Control Theory ..... 36
2.4.1 Basics of Feedback Theory ..... 36
2.4.2 Feedback Characteristics ..... 40
2.4.3 Controllers ..... 43
2.4.4 Disturbance ..... 45
2.4.5 Feedback Theory For a Predictor ..... 46
2.5 The FCS Framework ..... 46
2.5.1 The FCS Framework's Architecture ..... 47
2.5.2 Feedback in the Controller and QoS Actuator ..... 48
2.5.3 Assumptions of the FCS Framework ..... 54
2.5.4 Limitations of the FCS Framework ..... 55
2.6 The Constant Bandwidth Server Feedback Scheduler ..... 57
2.6.1 Constant Bandwidth Server ..... 58
2.6.2 Feedback Framework ..... 58
2.6.3 Stability ..... 60
2.6.4 Scheduling Error Assumption ..... 62
2.6.5 Limitations ..... 62
2.7 Conclusion ..... 62
3 GEDF and NP-GEDF ..... 63
3.1 Adaptable Sporadic Task System ..... 63
3.2 The SW Scheduling Algorithm and Deviance ..... 67
3.3 Modifications ..... 70
3.4 Task Reweighting ..... 70
3.4.1 Reweighting Under GEDF ..... 71
3.4.2 Modifications for NP-GEDF ..... 78
3.5 Tardiness and Drift Bounds ..... 79
3.5.1 Tardiness Bounds ..... 79
3.5.2 Additional Theoretical Algorithms ..... 86
3.5.3 Drift ..... 91
3.6 Conclusion ..... 97
4 PEDF and NP-PEDF ..... 99
4.1 Preliminaries ..... 99
4.2 A Limitation of Partitioning Schemes ..... 100
4.3 Partitioning and Repartitioning ..... 103
4.4 Allowing Guaranteed and Desired Weights to Differ ..... 104
4.4.1 Determining Guaranteed Weights ..... 104
4.4.2 The Adaptable Sporadic Task Model, Revisted ..... 108
4.4.3 Modifying the SW Algorithm ..... 112
4.5 Changing Desired Weights ..... 113
4.5.1 Changing Desired Weights in PEDF ..... 113
4.5.2 Modifications for NP-PEDF ..... 122
4.6 Resetting Rules ..... 124
4.7 Scheduling Correctness ..... 126
4.7.1 The CSW Algorithm ..... 126
4.7.2 Overlap ..... 128
4.7.3 Lag ..... 128
4.7.4 Higher-Priority Jobs ..... 130
4.7.5 PEDF Correctness ..... 131
4.7.6 NP-PEDF Correctness ..... 134
4.8 The IDEAL and PT Algorithms ..... 135
4.9 Drift ..... 140
4.9.1 Partial Drift ..... 143
4.9.2 Relationship Between PT and IDEAL ..... 151
4.9.3 Calculating Drift ..... 164
4.9.4 Incorporating Resets ..... 166
4.9.5 Total Drift Incurred ..... 167
4.9.6 Modifications for NP-PEDF ..... 168
4.10 Adjusting PEDF for Use with any Metric ..... 171
4.11 Time Complexity ..... 173
4.12 Conclusion ..... 173
$5 \mathrm{PD}^{2}$ ..... 175
5.1 Preliminaries ..... 175
5.1.1 Periodic Pfair Scheduling ..... 175
5.1.2 The Intra-Sporadic Task Model ..... 178
5.1.3 The $\mathrm{PD}^{2}$ Algorithm ..... 179
5.1.4 IS Ideal Schedule ..... 182
5.1.5 Dynamic Task Systems ..... 183
5.2 Adaptable Task Model ..... 184
5.3 SW Scheduling Algorithm ..... 189
5.4 Reweighting Rules ..... 190
5.4.1 Positive- and Negative-Changeable ..... 191
5.4.2 Heavy-Changeable ..... 194
5.5 Scheduling Correctness ..... 199
5.5.1 The AGIS Task Model ..... 205
5.5.2 Displacements ..... 207
5.5.3 Reweighting Properties ..... 210
5.5.4 Correctness Proof ..... 213
5.6 Drift ..... 230
5.6.1 PD-LJ is Not Fine-Grained ..... 232
5.6.2 All EPDF Scheduling Algorithms Incur Drift ..... 233
5.6.3 PD-PNH is fine-grained ..... 234
5.7 Lost Utilization ..... 243
5.8 Conclusion ..... 246
6 AGEDF ..... 247
6.1 Adaptable Service Level Tasks ..... 247
6.2 The AGEDF Scheduling Algorithm ..... 250
6.2.1 The Feedback Predictor ..... 251
6.2.2 Optimization ..... 258
6.2.3 Reweighting ..... 260
6.2.4 User-Defined Threshold ..... 262
6.3 Conclusion ..... 263
7 IMPLEMENTATION and EXPERIMENTS ..... 264
7.1 Whisper ..... 265
7.1.1 Correlation ..... 266
7.1.2 Kalman Filter ..... 268
7.1.3 Occlusion ..... 270
7.1.4 Real-time Characteristics ..... 271
7.2 VEC ..... 271
7.2.1 Bilateral Filter ..... 272
7.2.2 A Few Observations ..... 275
7.2.3 VEC's Algorithm ..... 276
7.2.4 Real-time Characteristics ..... 278
7.3 LITMUS ${ }^{\text {RT }}$ ..... 279
7.3.1 The Design of LITMUS ${ }^{\text {RT }}$ ..... 280
7.3.2 Core Infrastructure ..... 281
7.3.3 Scheduler Plugins ..... 282
7.3.4 System Call API ..... 283
7.4 Comparison ..... 283
7.4.1 Whisper Experiments ..... 285
7.4.2 VEC Experiments ..... 289
7.5 AGEDF Implementation and Evaluation ..... 290
7.5.1 Implementation ..... 290
7.5.2 Evaluation ..... 293
7.6 Conclusion ..... 316
8 CONCLUSION AND FUTURE WORK ..... 317
8.1 Summary of Results ..... 317
8.2 Other Related Work ..... 320
8.3 Future Work ..... 322
BIBLIOGRAPHY ..... 324

## LIST OF TABLES

1.1 Summary of algorithms and their properties. ..... 15
1.2 Empirical performance of the algorithms under different conditions. ..... 15
1.3 Summary of worst-case results for reweighting systems ..... 24
3.1 Summary of notation used in this chapter. ..... 64
4.1 Summary of notation used in this chapter. ..... 101
4.2 The MAOE, AAOE, MROE, and AROE metrics. ..... 104
4.3 Guaranteed weight values for the MAOE, AAOE, MROE, and AROE metrics. ..... 105
4.4 Summary of properties used in Section 4.7. ..... 126
4.5 Summary of properties used in Section 4.9. ..... 140
5.1 Brief description of the notation used in this chapter. ..... 176
5.2 Summary of properties used in Section 5.5. ..... 200

## LIST OF FIGURES

1.1 A one-processor example with three sporadic tasks ..... 4
1.2 A one-processor example of an ideal and EDF schedule. ..... 7
1.3 A two-processor example of PEDF and NP-PEDF. ..... 8
1.4 Partitioning three tasks with weight $2 / 3$ on two processors. ..... 9
1.5 A two-processor example of GEDF and NP-GEDF ..... 10
1.6 A two-processor example of $\mathrm{PD}^{2}$. ..... 12
1.7 A one-processor example of leave/join reweighting. ..... 17
1.8 Several one-processor examples of this dissertation's reweighting rules. ..... 23
1.9 A one-processor example of why leaves must be delayed. ..... 27
1.10 The AGEDF system. ..... 27
2.1 A one-processor example of leave/join reweighting. ..... 31
2.2 A one-processor example of RBED. ..... 33
2.3 A one-processor example of EEVDF. ..... 35
2.4 A simple feedback-control loop. ..... 37
2.5 Example system response to a step input. ..... 38
2.6 An example feedback-control loop. ..... 39
2.7 A feedback-control loop with a disturbance. ..... 45
2.8 The design of the FCS. ..... 48
2.9 The feedback loops for the FCS framework. ..... 52
2.10 An example one of the FCS framework's assumptions. ..... 56
2.11 An example one of the FCS framework's limitations. ..... 57
2.12 A one-processor example of the CBS ..... 59
2.13 The adaptive reservation-based feedback design. ..... 59
3.1 A one-processor example of the adaptable sporadic task model. ..... 65
3.2 An example of the GEDF, IDEAL, CSW, and SW scheduling algorithms. ..... 68
3.3 An example of the GEDF, IDEAL, CSW, and SW scheduling algorithms. ..... 69
3.4 A one-processor example of reweighting via Case (i) of Rule P under GEDF. ..... 73
3.5 A one-processor example of reweighting via Case (ii) of Rule P under GEDF. ..... 73
3.6 A one-processor example of reweighting via Case (i) of Rule $N$ under GEDF. . ..... 75
3.7 A one-processor example of reweighting via Case (ii) of Rule N under GEDF. ..... 75
3.8 A one-processor example of canceling a reweighting event. ..... 77
3.9 A one-processor example of NP-GEDF. ..... 78
3.10 A one-processor example of task classes. ..... 81
3.11 A one-processor example of task classes. ..... 84
3.12 An example of the GEDF, IDEAL, CSW, and SW scheduling algorithms. ..... 87
3.13 A continuation of Figure 3.12. ..... 88
3.14 An example of the EDF, IDEAL, CSW, and SW scheduling algorithms. ..... 90
3.15 A partial schedule that illustrates drift when tasks miss deadlines. ..... 95
3.16 A one-processor example of drift in NP-GEDF. ..... 97
3.17 A one-processor example of drift in NP-GEDF ..... 98
4.1 An example of the MAOE, MROE, and AROE metrics. ..... 106
4.2 A one-processor example of deadlines when minimizing the MROE. ..... 112
4.3 A one-processor example of the PEDF and SW scheduling algorithms. ..... 114
4.4 A one-processor example of the PEDF and SW scheduling algorithms. ..... 115
4.5 A one-processor example of reweighting via Case (i) of Rule P under PEDF. ..... 119
4.6 A one-processor example of reweighting via Case (ii) of Rule P under PEDF. ..... 119
4.7 A one-processor example of reweighting via Case (i) of Rule N under PEDF. ..... 121
4.8 A one-processor example of reweighting via Case (ii) of Rule N under PEDF. ..... 121
4.9 A one-processor example of canceling a reweighting event ..... 122
4.10 A one-processor example of NP-PEDF reweighting. ..... 123
4.11 A two-processor example of resetting under PEDF. ..... 125
4.12 An illustration of $t_{1}$ and $t_{d}$ ..... 132
4.13 An illustration of $t_{z}$ and $\epsilon_{z}$. ..... 133
4.14 An example illustrating the the NP-PEDF tardiness bound. ..... 135
4.15 An example of the PEDF, IDEAL, CSW, SW, PT scheduling algorithms. ..... 137
4.16 A continuation of Figure 4.15. ..... 138
4.17 An example of the PEDF, IDEAL, CSW, SW, and PT scheduling algorithms. ..... 139
4.18 A one-processor example of the task decomposition in Lemma 4.4. ..... 146
4.19 A one-processor example of multiple reweighting events. ..... 152
4.20 Decomposition of a task for Lemma 4.7. ..... 155
4.21 Decomposition of a task for Lemma 4.9. ..... 160
4.22 A one-processor example of drift in NP-PEDF ..... 169
4.23 A one-processor example of drift in NP-PEDF. ..... 170
5.1 The behavior of a periodic and IS task. ..... 178
5.2 The behavior of a periodic heavy task. ..... 181
5.3 Pseudo-code defining $\mathrm{A}\left(\mathcal{I}_{\mathrm{IS}}, T_{i}^{[j]}, t\right)$. ..... 182
5.4 Per-slot allocations for AIS system. ..... 185
5.5 An illustration of freeing the capacity. ..... 187
5.6 A one-processor $\mathrm{PD}^{2}$ schedule of two tasks. ..... 188
5.7 Pseudo-code defining the $\mathrm{A}\left(\mathcal{S} W, T_{i}^{[j]}, t\right)$ ..... 189
5.8 An illustration of Rule P under $\mathrm{PD}^{2}$. ..... 192
5.9 An illustration of Rule N under $\mathrm{PD}^{2}$. ..... 193
5.10 An example of instant reweighting causing a heavy task to miss its deadline. ..... 194
5.11 An example of increasing the weight of a heavy-changeable task. ..... 196
5.12 An example of decreasing the weight of a heavy-changeable task. ..... 197
5.13 An example of decreasing the weight of a heavy-changeable task. ..... 198
5.14 An illustration of the AGIS task model. ..... 205
5.15 The behavior of a periodic heavy task in an AGIS system. ..... 206
5.16 The reweighting rules in an AGIS system. ..... 207
5.17 An example of a heavy-changeable task in an AGIS system. ..... 208
5.18 An illustration of displacements. ..... 209
5.19 Illustrating the sets $A, B$, and $I$. ..... 215
5.20 An illustration of displacements caused by remove $\mathbf{X}^{(1)}$. ..... 216
5.21 An illustration of Corollary 5.2. ..... 220
5.22 An illustration of the proof of Lemma 5.13. ..... 222
5.23 An illustration of $T_{x}$ Lemma 5.14. ..... 225
5.24 An example of the IDEAL algorithm under $\mathrm{PD}^{2}$. ..... 232
5.25 An illustration of why PD-LJ is coarse-grained. ..... 233
5.26 An illustration that all EPDF algorithms incur drift. ..... 234
5.27 An illustration of Rule P under $\mathrm{PD}^{2}$. ..... 236
5.28 An illustration of Rule N under $\mathrm{PD}^{2}$. ..... 237
5.29 An illustration of Rule N under $\mathrm{PD}^{2}$. ..... 238
6.1 Estimated weight vs. importance value/service level. ..... 249
6.2 The AGEDF scheduling algorithm. ..... 251
6.3 A feedback system's response to a step input. ..... 252
6.4 Estimated weight vs. importance value/service level. ..... 259
6.5 An illustration of changing the code segment. ..... 262
7.1 The Whisper system ..... 265
7.2 Pseudo-code defining correlation. ..... 266
7.3 An illustration of correlation. ..... 267
7.4 An illustration of the Kalman Filter. ..... 269
7.5 The core loop in Whisper. ..... 270
7.6 An illustration of an occluding object. ..... 270
7.7 Example of different methods for removing noise from a pixel. ..... 273
7.8 The flow diagram of the VEC. ..... 277
7.9 The VEC system divided into tasks. ..... 279
7.10 The simulated Whisper system ..... 286
7.11 Whisper simulations. ..... 287
7.12 Whisper simulations continued. ..... 288
7.13 The simulated VEC system. ..... 289
7.14 VEC simulations. ..... 291
7.15 VEC simulations, continued. ..... 292
7.16 The simulated Whisper system. ..... 294
7.17 The actual weight of a Whisper task at three different service levels. ..... 296
7.18 AGEDF running Whisper Continued. ..... 297
7.19 AGEDF running Whisper. ..... 298
7.20 AGEDF running Whisper. ..... 299
7.21 AGEDF running Whisper. ..... 300
7.22 AGEDF running Whisper. ..... 301
7.23 AGEDF running Whisper. ..... 302
7.24 AGEDF running Whisper. ..... 303
7.25 The simulated VEC system. ..... 304
7.26 The actual weight of a VEC task at three different service levels. ..... 305
7.27 The actual weight of a VEC task at three different service levels. ..... 306
7.28 AGEDF running VEC. ..... 307
7.29 AGEDF running VEC. ..... 308
7.30 AGEDF running VEC. ..... 309
7.31 AGEDF running VEC. ..... 310
7.32 AGEDF running VEC. ..... 311
7.33 AGEDF running VEC. ..... 312
7.34 AGEDF running VEC. ..... 313
7.35 AGEDF running VEC. ..... 314
7.36 AGEDF running VEC. ..... 315

## LIST OF ABBREVIATIONS

| AIS | Adaptable Intra-Sporadic Task Model |
| :--- | :--- |
| AGIS | Adaptable Generalized Intra-Sporadic Task Model |
| AGEDF | Adaptive GEDF |
| AAOE | Average Absolute Overall Error |
| AROE | Average Relative Overall Error |
| EDF | Earliest-Deadline-First |
| EPDF | Earliest-Pseudo-Deadline-First |
| EEVDF | Earliest-Eligible-Virtual-Deadline-First |
| FF | Fairness Factor |
| FCS | Feedback-Control Real-Time Scheduling |
| FMLP | Flexible Multiprocessor Locking Protocol |
| GEDF | Global EDF |
| IS | Intra-Sporadic |
| MAOE | Maximal Absolute Overall Error |
| MROE | Maximal Relative Overall Error |
| NP-GEDF | Non-Preemptable GEDF |
| NP-PEDF | Non-Preemptable PEDF |
| PEDF | Partitioned EDF |
| RBED | Rate-Based Earliest-Deadline |
| RTOS | Real-Time Operating System |
| RAD | Reasonable Allocation Decreasing |
| UA | Average Under Allocation |

## CHAPTER 1

## INTRODUCTION

The goal of this dissertation is to extend research on multiprocessor real-time systems in order to enable such systems to adapt tasks' processor shares - a process called reweighting -in response to both external and internal stimuli. The particular focus of this work is on adaptive systems that are deployed in environments in which tasks may frequently require significant share changes. Such environments are commonplace in computationally-intensive multimedia applications. Prior to the research in this dissertation, no multiprocessor reweighting algorithms had been proposed that could change task shares with bounded overhead. In this dissertation, we extend prior work on uniprocessor and multiprocessor systems to construct reweighting algorithms with minimal overhead for several different types of multiprocessor systems. Furthermore, we examine how feedback and optimization techniques can be use to determine, at run time, which reweighting events are needed. Finally, we evaluate the proposed adaptive scheduling algorithms by using two multimedia applications developed at The University of North Carolina at Chapel Hill: the Whisper human tracking system (Vallidis, 2002) and the Virtual Exposure Camera (VEC) night vision system (Bennett and McMillan, 2005).

To motivate the need for adaptive real-time scheduling, we begin this chapter with a brief overview of Whisper and VEC. Next, we discuss the core real-time concepts that are relevant to this dissertation. We then review prior work on adaptive real-time systems and state the thesis of this dissertation. We conclude this chapter by summarizing this dissertation's contributions and providing an outline for the remainder of the dissertation.

### 1.1 Applications

Brief descriptions of Whisper and VEC are provided below; more detailed descriptions will be given later in Chapter 7, where our experimental results are given. Before discussing Whisper and VEC, it is important to point out that, since each is a multimedia application, both must ensure certain timing constraints to provide an acceptable user experience and thus are examples of real-time application. One of the questions we will answer in this dissertation is how the timing constraints of real-time applications can be changed at run time while still providing an acceptable quality-of-service (QoS).

### 1.1.1 Whisper

As mentioned above, Whisper performs full-body tracking in virtual environments (Vallidis, 2002). Whisper tracks users via an array of wall- and ceiling-mounted microphones that detect signals (i.e., white noise) emitted from speakers attached to each user's hands, feet, and head. Specifically, by calculating the time-shift in the signal for each microphone-speaker pair, Whisper is able to triangulate each speaker's position. The amount of time required to calculate the distance between a microphone-speaker pair is indirectly related to the signal-to-noise ratio. As the distance between a microphone-speaker pair increases, the signal-tonoise ratio decreases, which increases the amount of time required to calculate this distance. Also, other factors, like ambient noise, further degrade the signal-to-noise ratio causing total computation time to increase. Because the computational cost associated with calculating the distance between a microphone speaker-pair can change at run time, the tasks comprising Whisper must be scheduled using algorithms that either allow task parameters to adapt or allow task shares to be defined based on worst-case scenarios. Unfortunately, provisioning the system based on worst-case scenarios may not be a viable option because there exist scenarios for which no reasonable computational platform can correctly track a user (e.g., a room with a 100 dB of ambient noise). While using adaptive techniques reduces the resources required to correctly track a user (relative to over-provisioning), the workload is still intensive enough to necessitate a multiprocessor system.

### 1.1.2 Virtual Exposure Camera

The second application considered in this dissertation is the VEC video-enhancement system (Bennett and McMillan, 2005). ${ }^{1}$ VEC is capable of improving the quality of an underexposed video feed so that objects that are indistinguishable from the background become clear and in full color. In VEC, darker objects require more computation to correct. Thus, as dark objects move in the video, the processor shares of the tasks assigned to process different areas of the video will change. Like all multimedia applications, to create an acceptable user experience, VEC must update the corrected image at a regular rate. VEC will eventually be deployed in a full-color night vision system, so tasks will need to change shares as fast as a user's head can turn. In the planned configuration, a multicore platform consisting of approximately ten processing cores will be used.

### 1.2 Real-Time Systems

The distinguishing characteristic of a real-time system is the need to satisfy timing constraints. The timing constraints of recurrent applications (e.g., Whisper and VEC) can be represented using the sporadic task model. In this model, each piece of sequential recurrent code is called a task. Each invocation of such a task is called a $j o b$. We denote the $i^{\text {th }}$ task of a set of tasks $T$ as $T_{i}$ (where tasks are ordered by some arbitrary method), and denote the $j^{\text {th }}$ job of the task $T_{i}$ as $T_{i}^{j}$ (where jobs are ordered by the sequence in which they are invoked). Associated with a sporadic task is a worst-case execution time (WCET), denoted $\mathrm{e}\left(T_{i}\right)$, and a period, denoted $\mathrm{p}\left(T_{i}\right)$. The WCET denotes the maximum amount of time any job of the task requires; the period denotes the minimum separation between consecutive job invocations and defines a relative "deadline" for each job. The time at which a job is invoked is called its release time, denoted $\mathrm{r}\left(T_{i}^{j}\right)$, and the (absolute) time by which a job must complete is called its deadline, denoted $\mathrm{d}\left(T_{i}^{j}\right)$. The weight of a task $T_{i}$, denoted $\mathrm{wt}\left(T_{i}\right)$, is the fraction of a process it requires to be correctly scheduled and is defined as $\mathrm{e}\left(T_{i}\right) / \mathrm{p}\left(T_{i}\right)$. For shorthand, we will use $T_{i}:(e, p)$ to denote a task $T_{i}$ with a WCET of $e$ and a period of $p$.

[^0]

Figure 1.1: A one-processor example with three sporadic tasks.
The release and deadline of the job $T_{i}^{j}$ of a sporadic task $T_{i}$ can be specified as

$$
\begin{aligned}
\mathrm{r}\left(T_{i}^{1}\right) & =\theta\left(T_{i}^{1}\right) \\
\mathrm{r}\left(T_{i}^{j}\right) & =\mathrm{d}\left(T_{i}^{j-1}\right)+\theta\left(T_{i}^{j}\right), j>1 \\
\mathrm{~d}\left(T_{i}^{j}\right) & =\mathrm{r}\left(T_{i}^{j}\right)+\mathrm{p}\left(T_{i}\right), j \geq 1
\end{aligned}
$$

where $\theta\left(T_{i}^{j}\right) \geq 0$ for $j \geq 1$. $\theta\left(T_{i}^{j}\right)$ denotes the sporadic separation between job releases. If $\theta\left(T_{i}^{j+1}\right)=0$, then $T_{i}^{j+1}$ is released at $T_{i}^{j}$ s deadline.

Example (Figure 1.1). Consider the example in Figure 1.1, which depicts a one-processor system with three tasks: $T_{1}:(2,7)$, which has a sporadic separation of one time unit between $T_{1}^{1}$ and $T_{1}^{2} ; T_{2}:(1,7 / 3)$; and $T_{3}:(4,14)$. The grey boxes denote the time at which the associated job is scheduled. Down arrows denote a job release. Up arrows denote a job deadline. Up-and-down arrows denote that a job's deadline and its successor's release occur at the same time. Similar notation will be used in later figures.

The actual execution time of job $T_{j}^{j}$, denoted $\operatorname{Ae}\left(T_{i}^{j}\right)$, is the amount of time for which $T_{i}^{j}$ is scheduled; this value is upper-bounded by e $\left(T_{i}\right)$. Depending on the scenario, this value may or may not be known before the job finishes execution. To facilitate further discussion, a few additional terms need to be defined.

Definition 1.1 (Window, Active, and Inactive). If $T_{i}^{j}$ is a job in the task system $T$, then the window of $T_{i}^{j}$ defined as the range $\left[\mathrm{r}\left(T_{i}^{j}\right), \mathrm{d}\left(T_{i}^{j}\right)\right)$. Furthermore, the job $T_{i}^{j}$ is active correct.
at time $t$ iff $t$ is in $T_{i}^{j}$,s window (i.e., $t \in\left[\mathrm{r}\left(T_{i}^{j}\right), \mathrm{d}\left(T_{i}^{j}\right)\right)$ ), and inactive otherwise. We use $\mathrm{ACT}(t)$ to denote the set of active jobs at time $t$.

For example, in Figure 1.1, $T_{1}^{1}$ is active over the range $[1,7)$ and is inactive at every other time.

Definition 1.2 (Completed). If $\mathcal{S}$ is a schedule of the task system $T$, then a job $T_{i}^{j} \in T$ is said to have completed by time $t$ in $\mathcal{S}$ iff $T_{i}^{j}$ has executed for $\operatorname{Ae}\left(T_{i}\right)$ time units by $t$ in $\mathcal{S}$. Similarly, a task $T_{i}$ is said to be complete at time $t$ iff at time $t$ every job of $T_{i}$ that was released by $t$ has completed.

For example, in Figure 1.1, $T_{1}^{1}$ is complete at and after time 4. Also, at time $10 / 3, T_{2}$ is complete since both $T_{2}^{1}$ and $T_{2}^{2}$ are complete by time $10 / 3$.

Definition 1.3 (Pending and Ready). For an arbitrary scheduling algorithm $\mathcal{A}$, if $\mathcal{S}$ is a schedule of the task system $T$ under $\mathcal{A}$, then a job $T_{i}^{j}$ is said to be pending at time $t$ in $\mathcal{S}$ if $\mathrm{r}\left(T_{i}^{j}\right) \leq t$ and $T_{i}^{j}$ is incomplete at $t$ in $\mathcal{S}$. Note that a job can be both pending and inactive, if it misses its deadline. A pending job $T_{i}^{j}$ is said to be ready at time $t$ in $\mathcal{S}$ if all prior jobs of task $T_{i}$ have completed by $t$. A job $T_{i}^{j}$ can be pending but not ready if $T_{i}^{j-1}$ is incomplete at $\mathrm{r}\left(T_{i}^{j}\right)$. (Such a scenario may occur in some multiprocessor algorithms.)

For example, in Figure 1.1, $T_{1}^{1}$ is pending until time 4, and $T_{3}^{1}$ is pending until time 11.
Let $\mathcal{A}$ be an arbitrary scheduling algorithm, $\tau$ be an arbitrary task system, and $\mathcal{S}$ denote schedule of $\tau$ generated by $\mathcal{A}$. Then, we use $\mathrm{A}\left(\mathcal{S}, T_{i}^{j}, t_{1}, t_{2}\right)$ denote the total time allocated to $T_{i}^{j}$ in $\mathcal{S}$ in $\left[t_{1}, t_{2}\right)$. Similarly, we use $\mathrm{A}\left(\mathcal{S}, T_{i}, t_{1}, t_{2}\right)$ and $\mathrm{A}\left(\mathcal{S}, \tau, t_{1}, t_{2}\right)$, respectively, to denote the total time allocated to all jobs of $T_{i}$ in $\mathcal{S}$ and all tasks of $\tau$ in $\mathcal{S}$, over the interval $\left[t_{1}, t_{2}\right)$. We say that the value of $\mathrm{A}\left(\mathcal{S}, T_{i}^{j}, 0, t\right)$ is the amount that $T_{i}^{j}$ has executed by $t$. For example in Figure 1.1, $\mathrm{A}\left(\mathcal{S}, T_{1}^{1}, 0,2\right)=1$ and $\mathrm{A}\left(\mathcal{S}, T_{1}^{1}, 1,4\right)=2$.

Depending on the consequences of missing a deadline, real-time systems can be classified as either "hard" or "soft." A system is a hard real-time (HRT) system if missing any deadline implies that the system fails. In contrast, in a soft real-time (SRT) system, deadlines may be missed. Examples of HRT systems include avionics and automotive applications. Examples of SRT systems include multimedia and virtual-reality applications. While SRT systems may
miss an occasional deadline, it is still possible for such systems to "fail;" however, there is no single notion of a "correct" SRT system. Some possible notions of SRT correctness include: bounded deadline tardiness (i.e., all jobs complete within some bound of their deadline) (Devi and Anderson, 2008); a specified percentage of deadlines must be met (Lu et al., 2002); and $m$ out of every $k$ consecutive jobs of each task complete before their deadline (Hamadoui and Ramanathan, 1995). In this dissertation, we are primarily concerned with HRT systems and SRT systems with bounded deadline tardiness. For HRT systems, we say that a given task set is schedulable if it is possible to guarantee that no single task will miss its deadline; otherwise we say that it is unschedulable. Similarly, for SRT systems, we say that a given task set is schedulable if it is possible to guarantee that every task has bounded deadline tardiness, and is unschedulable otherwise. Moreover, for many types of system, we can determine if a given task set is schedulable using a scheduability test, i.e., a set of conditions that, when satisfied by the task set, imply that it is schedulable.

### 1.2.1 Uniprocessor Systems

The weight of a task can be used to define an ideal schedule, in which, at each instant, each task is allocated a fraction of a processor equal to its weight. While the ideal schedule represents the most equitable allocation of resources possible, it is infeasible to implement since it requires tasks to be preempted and swapped at arbitrarily small intervals. For a uniprocessor system, a more realistic scheduling algorithm is the earliest-deadline-first (EDF) algorithm, which schedules jobs based on their deadlines, with earlier deadlines having higher priority. On a uniprocessor system, EDF can guarantee that every job completes before its deadline if the total weight of all tasks is at most one, the total available utilization.

Example (Figure 1.2). Figure 1.2 depicts a one-processor example of an ideal and EDF schedule of a system with four tasks: $T_{1}:(1,2) ; T_{2}:(2,8) ; T_{3}:(1,8) ;$ and $T_{4}:(1,8) T_{1}$. The numbers in each box denote the fraction of the processor consumed by the associated task. Insets (a) through (c) depict, respectively, an ideal schedule, an EDF schedule, and the actual and ideal allocations for $T_{1}$.

(a)

(b)

(c)

Figure 1.2: A one-processor example of an (a) ideal and (b) EDF schedule, and (c) $T_{1}$ 's allocation in both.

### 1.2.2 Multiprocessor Partitioned Scheduling

Most multiprocessor scheduling algorithms can be classified as either partitioned or global. Under partitioned algorithms, each task is permanently assigned to a specific processor and each processor independently schedules its assigned tasks using a uniprocessor scheduling algorithm. Alternatively, under global algorithms, a task may migrate among processors. The advantage of partitioned approaches over global approaches is that they have lower migration/preemption costs. This is because, under partitioned approaches, tasks maintain cache affinity for longer durations of time due to fewer task migrations than in global approaches. The disadvantage of partitioned approaches is that such systems have inferior scheduability conditions when compared to global approaches (as we shall see).

This section discusses the specifics of two partitioned scheduling algorithms: preemptive and nonpreemptive partitioned EDF (abbreviated as PEDF and NP-PEDF, respectively). Under PEDF and NP-PEDF each processor is scheduled independently using the EDF scheduling algorithm. The difference between them is that, under PEDF, a job can be preempted, and under NP-PEDF, a job cannot be preempted.

Example (Figure 1.3). Consider the example in Figure 1.3, which depicts a two-processor system with six tasks.: $T_{1}:(3,12) ; T_{2}:(1,6) ; T_{3}:(3,6) ; T_{4}:(1,4) ; T_{5}:(1,4) ;$ and $T_{6}:(5,12)$. Tasks $T_{1}, T_{3}$, and $T_{5}$ are assigned to Processor 1, and tasks $T_{2}, T_{4}$, and $T_{6}$ are assigned to Processor 2. Inset (a) depicts a PEDF schedule, and (b) depicts an NP-PEDF schedule.


Figure 1.3: Two-processor (a) PEDF and (b) NP-PEDF schedules.

Notice that, since $T_{5}$ is assigned to Processor 1, Processor 2 is idle over the range [10, 11) even though $T_{5}$ has work to be completed. Also note that the difference between Figure 1.3(a) and Figure 1.3(b) is that in Figure 1.3(b) once a job begins executing it continues to do so until it completes. This behavior is illustrated by $T_{6}^{1}$, which has a contiguous execution in Figure 1.3(b) but not in Figure 1.3(a).

Before continuing, there is one subtlety that must be discussed. Throughout this dissertation, whenever we discuss partitioned scheduling algorithms, we will make a distinction between the guaranteed weight of a task and the desired weight of a task (given by $\mathrm{e}\left(T_{i}\right) / \mathrm{p}\left(T_{i}\right)$ ). This distinction is important because under partitioned algorithms, it is possible for a processor to be over-allocated, i.e., the processor is assigned tasks with a total weight exceeding one. When a processor is over-allocated, there are two options: reject one or more tasks; or reduce the shares of the tasks on that processor. For example, in a two-processor system with three tasks each of weight $2 / 3$ (depicted in Figure 1.4), either one of the three tasks must be rejected or the shares of the tasks on the over-allocated processor must be reduced. Both of these options guarantee that the shares of tasks do not overutilize either of the processors even though the weights do. While both variants have their relative advantages, for adaptive systems, the latter option is likely a better option. (A thorough discussion of this is issue given in Section 4.2.) It is important to note that for the majority of the algorithms considered in


Figure 1.4: (a) A two-processor system with three tasks each with a desired weight of $2 / 3$. (b) The guaranteed weights of tasks in (a) when no task is rejected. (c) The guaranteed weights of tasks in (a) when $T_{2}$ is rejected.
this dissertation, a task's guaranteed weight equals its desired weight. Thus, for brevity, we will only use the terms "desired weight" and "guaranteed weight" when discussing algorithms where the two may differ.

### 1.2.3 Restricted Global Multiprocessor Scheduling

In global scheduling algorithms, tasks are scheduled from a single priority queue and may migrate among processors. Global algorithms can be further classified as either restricted and unrestricted. A scheduling algorithm is considered to be restricted if the scheduling priority of each job (for any given schedule) does not change once it has been released. A scheduling algorithm is considered to be unrestricted if there exists a schedule in which some job changes its priority after it is released. In this section, we discuss two restricted global scheduling algorithms, preemptive global-EDF (GEDF) and non-preemptive global-EDF (NP-GEDF); unrestricted algorithms are considered in Section 1.2.4. Under both GEDF and NP-GEDF, tasks are scheduled from a single priority queue on an EDF basis. As with partitioned algorithms, the only difference between GEDF and NP-GEDF is that jobs can be preempted in GEDF and cannot be preempted in NP-GEDF.

Example (Figure 1.5). Consider the example in Figure 1.5, which pertains to a twoprocessor system with five tasks: $T_{1}:(2,7) ; T_{2}:(1,7) ; T_{3}:(1,7) ; T_{4}:(3,7)$; and $T_{5}:(3,7)$. Inset (a) depicts a GEDF schedule and inset (b) depicts an NP-GEDF schedule. In (a), $T_{5}^{1}$ misses a deadline by one time unit at time 7 , and $T_{5}^{2}$ misses a deadline by one time unit at time 14 . In (b), $T_{5}^{2}$ misses a deadline by two time unit at time 14 .

Since tasks are not assigned to processors for prolonged periods of time in these algorithms,


Figure 1.5: Two-processor (a) GEDF and (b) NP-GEDF schedules.
it is not possible for a processor to be over-allocated in the long term, provided that total utilization is at most $m$, the number of processors. However, short-term over-allocations that cause deadline misses are possible. Nonetheless, as shown by Devi and Anderson (Devi and Anderson, 2008), such misses are only by bounded amounts. The amount of time by which a job misses its deadline is called its tardiness. For example, in Figure 1.5(a), $T_{5}^{1}$ misses a deadline at time 7, and in both Figure 1.5(a) and Figure 1.5(b), $T_{5}^{2}$ misses a deadline at time 14. As shown in (Devi and Anderson, 2008), under GEDF, the maximal tardiness of any job of task a $T_{i}$ is bounded by a formula given shortly; however, before presenting their equation, we briefly introduce a few needed terms. Let $\mathrm{e}_{\text {min }}$ denote the minimum execution time of any task in the system $\tau$. Define $\mathrm{WT}(\tau)$ as

$$
\mathrm{WT}(\tau)=\sum_{T_{i} \in \mathcal{T}} \mathrm{wt}\left(T_{i}\right) .
$$

Additionally, let $\max _{\mathrm{e}}(k)$ and $\max _{\mathrm{wt}}(k)$ denote, respectively, the $k^{\text {th }}$ largest execution time and weight of any task. Finally, let the value $\Gamma$ denote

$$
\Gamma=\left\{\begin{array}{lr}
\mathrm{WT}(\tau)-1, & \mathrm{WT}(\tau) \text { is integral } \\
\lfloor\mathrm{WT}(\tau)\rfloor, & \text { otherwise }
\end{array}\right.
$$

Using these terms, the maximal tardiness (as given in (Devi and Anderson, 2008)) for any
job of a task $T_{i}$ is given as

$$
\begin{equation*}
\frac{\sum_{k=1}^{\Gamma} \max _{\mathrm{e}}(k)-\mathrm{e}_{\min }}{m-\sum_{k=1}^{\Gamma-1} \max _{\mathrm{wt}}(k)}+\mathrm{e}\left(T_{i}\right), \tag{1.1}
\end{equation*}
$$

provided $\mathrm{WT}(\tau) \leq m$, where $m$ is the number of processors. Since tasks cannot be preempted in NP-GEDF, its tardiness bound is slightly larger than (1.1). (See (Devi and Anderson, 2008) for details.)

### 1.2.4 Unrestricted Global Multiprocessor Scheduling

The final multiprocessor scheduling algorithm we consider in this dissertation is the unrestricted global Pfair algorithm PD ${ }^{2}$ (Srinivasan and Anderson, 2005). PD ${ }^{2}$ schedules a task by breaking it into a sequence in which of subtasks, each of which represents one time unit of execution. Each such time unit is called a quantum. The $j^{\text {th }}$ subtask of the task $T_{i}$ is denoted as $T_{i}^{[j]}$, where subtasks are ordered by the sequence in which they are invoked. Associated with each subtask is a pseudo-release and pseudo-deadline (often called "release" and "deadline" for brevity), which are defined as

$$
\begin{aligned}
\mathrm{r}\left(T_{i}^{[1]}\right) & =\theta\left(T_{i}^{[1]}\right) \\
\mathrm{r}\left(T_{i}^{[j]}\right) & =\mathrm{d}\left(T_{i}^{[j-1]}\right)+\left\lfloor\frac{j-1}{\mathrm{wt}\left(T_{i}\right)}\right\rfloor-\left\lceil\frac{j-1}{\mathrm{wt}\left(T_{i}\right)}\right\rceil+\theta\left(T_{i}^{j}\right), j>1 \\
\mathrm{~d}\left(T_{i}^{[j]}\right) & =\mathrm{r}\left(T_{i}^{[j]}\right)-\left\lfloor\frac{j-1}{\mathrm{wt}\left(T_{i}\right)}\right\rfloor+\left\lceil\frac{j}{\mathrm{wt}\left(T_{i}\right)}\right\rceil, j \geq 1
\end{aligned}
$$

where $\theta\left(T_{i}^{[j]}\right) \geq 0$ for $j \geq 1 . \theta\left(T_{i}^{[j]}\right)$ denotes the "sporadic" separation between subtask releases (such separations are called "intra-sporadic"). PD ${ }^{2}$ schedules subtasks on an earliest-pseudo-deadline-first basis with two tie-breaking rules, which are used in the event of a deadline tie. (A more thorough discussion of $\mathrm{PD}^{2}$ can be found in Chapter 5.) Since tasks in $\mathrm{PD}^{2}$ are scheduled one quantum at a time, a task's execution time is assumed to be a multiple of the quantum size (and must be rounded up if this is not the case) and the scheduling of a task depends only on its weight. The main advantage of $\mathrm{PD}^{2}$ over all other aforementioned algorithms is that $\mathrm{PD}^{2}$ is the only algorithm that can guarantee that every job is scheduled before its deadline and no processor is over-allocated provided total utilization is at most the number of processors. However, since tasks are scheduled on a per-quantum basis, it is


Figure 1.6: A two-processor system scheduled by PD ${ }^{2}$.
possible that a task will be preempted and migrated every quantum, which can cause tasks to incur large migration/preemption costs. Another disadvantage of $\mathrm{PD}^{2}$ is that task execution times must be rounded up, which can cause the system to be underutilized. For example, if a task has an execution time of 3.1 and a period of 4 , then its weight would be $4 / 4=1$, since the execution time of 3.1 would be rounded up to one.

Example (Figure 1.6). Figure 1.6 shows a PD $^{2}$ schedule for the task system consider earlier in Figure 1.5.

Notice that, in this schedule, tasks are scheduled one quantum at a time. As a result, tasks may be preempted at the end of every quantum and may migrate nearly as frequently. Also note that a subtask (unlike a job) may be released before the deadline of its successor. Finally, notice that, in this schedule, every task is scheduled before its deadline and that tasks with the same weight (but different periods) receives allocations at the same rate, e.g., $T_{3}$ and $T_{4}$ receive approximately one allocation in any $7 / 3$-quantum interval.

### 1.2.5 Impact of Migrations and Preemptions

Given these five algorithms, it is obvious that in the absence of migration and preemption costs, $\mathrm{PD}^{2}$ should be the preferred algorithm since it is the only algorithm that can both guarantee that every job completes before its deadline and that the share of every task equals its weight. However, for many applications, migration and preemption costs may be sub-
stantial. Recently, our research group (Calandrino et al., 2006) constructed a multiprocessor testbed, called LITMUS ${ }^{\text {RT }}$, to compare different real-time scheduling algorithms. We then used this testbed to implement all of the aforementioned algorithms (except NP-PEDF) on a four-processor system (with 2.7 GHz processors) in an effort to assess the impact of migration and preemption costs. In our work, we varied the weight of tasks being scheduled and the amount of cache used by each task. (Migrations and preemptions cause a loss of cache affinity, and as a result, if a task utilizes a larger fraction of the cache, then that task has both higher migration and preemption costs.)

Our experiments assessed the performance of each algorithm as measured by the number of processors that would be required to schedule a number of randomly generated task sets with a maximal utilization of four. These experiments showed the following:

- The HRT performance of PEDF and GEDF improve (relative to the other algorithms) as migration and preemption costs increase and/or the weights of tasks decrease. However, PEDF always has better HRT performance than GEDF.
- The performance of $\mathrm{PD}^{2}$ improves (relative to the other algorithms) as migration and preemption costs decrease and/or the weights of the tasks increase.
- $\mathrm{PD}^{2}$ has virtually the same performance in both HRT and SRT systems.
- PEDF has virtually the same performance in both HRT and SRT systems.
- Both GEDF and NP-GEDF always perform better than any other algorithm for SRT systems. Furthermore, NP-GEDF has slightly better performance than GEDF. (The HRT performance of NP-GEDF was not considered since there does not exist a scheduability test for it that would return meaningful results.)

The reason why the performance of PEDF is adversely impacted by increasing task weights is that, if tasks have higher weights, then it is harder to produce a valid partitioning. Similarly, for GEDF, as task weights increase, it becomes more likely that a job will be tardy. As a result, additional processing capacity is needed to prevent such a scenario. The performance of $\mathrm{PD}^{2}$ is positively impacted by increasing task weights because $\mathrm{PD}^{2}$ is more likely to schedule
tasks with larger weights in consecutive quanta, thus reducing the number of migrations and preemptions. GEDF and NP-GEDF perform well for SRT systems because even though these algorithms can cause tasks to miss deadlines, they incur relatively little migration/preemption cost. Moreover, since tasks are not partitioned under GEDF and NP-GEDF, the performance (in terms of the required number of processors to guarantee bounded tardiness) does not substantially degrade as task weights increase. Given these results, it is easy to see that there there does not exist a single "best" multiprocessor scheduling algorithm, and the choice of which algorithm depends on the scenario in which it will be used. The theoretical and empirical results from this section are summarized in Tables. 1.1 and 1.2, respectively.

### 1.2.6 Research Needed

Under traditional task models (e.g., the sporadic model), the scheduability of a system is based on each task's WCET. The disadvantage of using WCETs is that a system may be deemed unschedulable even if it would function correctly most (or possibly all) of the time when deployed. Adaptive real-time scheduling algorithms allow per-task processor shares to be adjusted based upon run time conditions, instead of always using constant share allocations based upon WCETs. Prior to the research in the dissertation, for multiprocessor systems, one approach for reweighting a task had been proposed, as we will discuss in Section 2.1 (Srinivasan and Anderson, 2005); however, this approach only allows tasks to reweight at job boundaries. By delaying a task's reweighting request until its next job boundary, the system may "drift" from its "ideal" allocation by an arbitrarily large amount. As a result, for applications like Whisper and VEC, where timing constraints are continually changing, such delays may cause unacceptably poor performance. In this dissertation, we remedy this shortcoming by proposing a set of rules that allow tasks to reweight without causing unbounded drift, and an adaptive framework that determines at run time which adaptions are needed.

| Scheme | Desired $=$ Guaranteed Weight | Guaranteed Deadlines |
| :---: | :---: | :---: |
| PEDF | No | Yes |
| NP-PEDF | No | No |
| GEDF | Yes | No |
| NP-GEDF | Yes | No |
| PD $^{2}$ | Yes | Yes |


| Scheme | Migrations | Preemptions |
| :---: | :---: | :---: |
| PEDF | Never | At Job Completions and Releases |
| NP-PEDF | Never | Never |
| GEDF | At Job Completions and Releases | At Job Completions and Releases |
| NP-GEDF | In Between Jobs | Never |
| PD $^{2}$ | Every Quantum | Every Quantum |

Table 1.1: Summary of algorithms and their properties. Note that deadlines can be guaranteed under PEDF only at the expense of allowing guaranteed weights to be less than desired weights.

| Scheme <br> (Hard) | Light Tasks <br> (Hard) | Heavy Tasks <br> (Hard) |
| :---: | :---: | :---: |
| PEDF | Best | Poor |
| GEDF | Good | Worst |
| NP-GEDF | N/A | N/A |
| PD $^{2}$ | Worst | Best |


| Scheme <br> (Hard) | High Mig./Preemp. Costs <br> (Hard) | Low Mig./Preemp. Costs <br> (Hard) |
| :---: | :---: | :---: |
| PEDF | Best | Poor |
| GEDF | Good | Worst |
| NP-GEDF | N/A | N/A |
| PD $^{2}$ | Worst | Excellent |


| Scheme <br> (All cases) | Soft-Real Time |
| :---: | :---: |
| PEDF | Same as hard |
| GEDF | Excellent |
| NP-GEDF | Best |
| PD $^{2}$ | Same as hard |

Table 1.2: Empirical performance of the algorithms under different conditions. Heavy tasks have a weight of at least $1 / 2$, and light tasks have a weight less than $1 / 2$. (Results are relative to other algorithms.)

### 1.3 Adaptivity

This section provides a review of prior work on adaptive uniprocessor real-time schemes in which tasks are reweighted based on external and internal stimuli.

### 1.3.1 Leave/join Reweighting

Under leave/join reweighting (Srinivasan and Anderson, 2005), a task's weight is changed at job boundaries by forcing it to leave with its old weight and rejoin with its new weight.

Example (Figure 1.7). Consider the example Figure 1.7, which depicts a one-processor system with three tasks: $T_{1}:(1,2)$ that leaves at time $2 ; T_{2}$, which has an initial weight of $1 / 4$, an execution cost of 1 , and "initiates" a weight increase at time 2 to a weight of $3 / 4$, which is enacted by leave/join reweighting at $T_{2}^{1}$ 's deadline (i.e., time 4); and $T_{3}:(2,8)$. (a) illustrates the EDF schedule and (b) illustrates the ideal and actual allocations to task $T_{2}$.

Notice that, even though $T_{2}$ "initiates" its change at time 2 and capacity exists for $T_{2}$ to increase its weight, this weight change cannot be "enacted" until its deadline. This illustrates that the primary drawback to leave/join reweighting is that a task can only change its weight at job boundaries. As a result, over the time range $[2,4), T_{2}$ behaves as though it is a task of weight $1 / 4$, even though in the ideal system (which can instantly enact weight changes) $T_{2}$ would behave as a task of weight $3 / 4$, is depicted in Figure 1.7(b). As a result, $T_{2}$ 's allocation in the actual schedule "drifts" from its allocation in the ideal schedule by one time unit. (Briefly, drift is the difference between a task's ideal and actual allocations caused by a single reweighting event.) Moreover, since leave/join reweighting cannot enact a reweighting event until a job boundary, it is possible that a task can incur an arbitrarily large amount of drift for one reweighting event. Section 1.5.3 provides a more detailed discussion concerning drift.

### 1.3.2 Rate-Based Earliest Deadline

Under rate-based earliest-deadline (RBED) scheduling (Brandt et al., 2003), which was proposed for uniprocessor systems, tasks are scheduled on an EDF basis and can change their

(a)

(b)

Figure 1.7: The (a) EDF schedule and the (b) ideal and actual allocations to $T_{2}$ in a oneprocessor example of leave/join reweighting.
weights and periods via two different rules based on whether the execution time or period is changed. While these rules are more responsive than leave/join reweighting, it is still possible for a task to incur an arbitrarily large amount of drift under RBED. In Chapter 2, we will review this work in detail.

### 1.3.3 Proportional Share Scheduling

Under proportional share scheduling (Stoica et al., 1996), the guaranteed weight of each task is determined as a function of its desired weight and the desired weight of all other tasks. Specifically, the guaranteed weight of the task $T_{i}$ at time $t$ is defined as

$$
\begin{equation*}
\operatorname{Gwt}\left(T_{i}, t\right)=\frac{\mathrm{wt}\left(T_{i}\right)}{\mathrm{WT}(\tau, t)}, \tag{1.2}
\end{equation*}
$$

where $\mathrm{WT}(\tau, t)$ is the total desired weight of all active tasks in the system $\tau$ at time $t$. For example, consider a one-processor system that consists of four tasks: $T_{1}$ that has a desired weight of $1 / 2 ; T_{2}$ that has a desired weight of $1 / 4 ; T_{3}$ that has a desired weight of $1 / 4$; and $T_{4}$ that has a desired weight of $1 / 8$. If, at some time $t_{1}$ all four tasks are active, then the guaranteed weight for each task would be $1 /(1 / 2+1 / 4+1 / 4+1 / 8)=8 / 9$ of its desired weight. So, the desired weight of $T_{1}$ would be $4 / 9$. Alternatively, if at some time, $t_{2}$, only $T_{1}$ and $T_{2}$ were active, then the guaranteed weight of each would be $4 / 3$ of its desired weight for each task. So, the guaranteed weight of $T_{1}$ would be $2 / 3$. It is worthwhile to note that proportional share algorithms cannot change the weight of only one task in a system
without using leave/join techniques. In Section 2.3, we will review one of the more popular proportional share scheduling algorithms.

### 1.3.4 Adaptive Feedback-Based Frameworks

In adaptive feedback-based scheduling algorithms, the execution time of each job is unknown until it is complete. As a result, in these systems, each task's weight is defined as a function of its estimated execution time, which is calculated for a job by using the actual execution times of the task's prior jobs. Moreover, a user can fine-tune a feedback-based system to achieve desired behaviors.

Lu et al. ( Lu et al., 2000) were the first to propose such a scheduling algorithm, which was directed at uniprocessor systems. ${ }^{2}$ Under their algorithm, called FC-EDF ${ }^{2}$, each task has multiple versions (called service levels), each of which has a different level of $Q o S$ and a different nominal processor share, representing the fraction of a processor the task will require on average if it executes at that service level. A task can only execute at one service level at a time. In order to control the system, FC-EDF ${ }^{2}$ monitors the system's utilization and miss-ratio, i.e., the fraction of jobs with missed deadlines. In order to minimize the miss-ratio while maximizing utilization, FC-EDF $^{2}$ adjusts the set of scheduled tasks and their service levels. More recently, Lu et al. extended this work to create a comprehensive feedback scheduling framework (Lu et al., 2002) that more explicitly incorporates the value to the system associated with each service level. This framework is the basis for the approach we propose in this paper. One drawback of FC-EDF ${ }^{2}$ is that, because only the utilization and the system-wide miss-ratio are monitored, the system cannot identify whether an individual task has an actual execution time that deviates substantially from its estimated execution time. Thus, the system can only respond to differences between the actual and estimated execution times of tasks by changing the entire system instead of only a few tasks.

Alternatively, Abeni et al. have proposed a uniprocessor feedback algorithm in which each task has its own feedback-controller rather than one controller for the entire system (Abeni

[^1]et al., 2002). In order to attempt to maintain an accurate processor share for each task, their algorithm monitors, for each task, the difference between the estimated and actual execution times of each job. Once the system has calculated a new estimated execution time for a future job, it adjusts the task's weight. More recently, Cucinotta et al. extended this approach to provide stochastic guarantees concerning per-task processor shares (Cucinotta et al., 2004). One drawback of their approach is that it ignores the possibility that some tasks are more important than others.

In addition to general-purpose real-time scheduling algorithms, feedback-based scheduling has become increasingly important for managing control tasks, i.e., tasks that control external devices. In work by Martí et al. (Marti et al., 2004), an approach is proposed that is similar to that of Abeni et al., except that in (Marti et al., 2004), each period of each task has an associated "importance value" that denotes the task's value to the system in that period. By using importance values, Martí et al. determined the optimal period for each task via standard linear programming techniques. One limitation of this approach is that it cannot adjust the amount of time for which a task executes, like other approaches (e.g., like that of Lu et al. (Lu et al., 2002)).

### 1.4 Thesis Statement

There are two major limitations of prior work on adaptive uniprocessor systems. First, there does not exist a set of metrics for comparing different reweighting algorithms. Second, existing methods for changing the weights of tasks (i.e., leave/join reweighting and RBED) may give rise to an unacceptably long delay after initiating a weight change before it is enacted. For multiprocessor systems, the limitations are even more severe since the only method for changing the weight of a task is to use leave/join reweighting. Moreover, for multiprocessor systems, there is an inherent tradeoff, which must be explored, between the level of migration/preemption and the "accuracy" of the adaptive protocol. Also, there is no global feedback-based adaptive framework, which would be necessary for implementing applications like Whisper and VEC. The main thesis of this dissertation, which attempts to resolves these issues, is the following.

Multiprocessor real-time scheduling algorithms can be made more adaptive by allowing tasks to reweight between job releases. Feedback and optimization techniques can be used to determine at run time which reweighting events are needed. The accuracy of such an algorithm can be improved by allowing more frequent task migrations and preemptions; however, this accuracy comes at the expense of higher migration and preemption costs, which impacts average-case performance. Thus, there is a tradeoff between accuracy and average-case performance that will be dependent on the frequency of task migrations/preemptions and their cost.

### 1.5 Contributions

In this section, we briefly discuss the contributions of this dissertation.

### 1.5.1 Adaptable Task Model and Reweighting Algorithms

The first contribution we discuss in the dissertation is the adaptable task model (originally proposed in (Block et al., 2008b)). This model is an extension of the sporadic task model, where the weight of each task $T, \mathrm{wt}(T, t)$, is a function of time $t$, and a task's execution time can vary between job releases. A task $T$ changes weight or reweights at time $t+1$ if $\mathrm{wt}(T, t) \neq \mathrm{wt}(T, t+1)$. If a task $T$ changes weight at a time $t_{c}$ between the release and the deadline of some job $T_{i}^{j}$, then the following two actions may occur:
(i) If $T_{i}^{j}$ has not been scheduled by $t_{c}$, then $T_{i}^{j}$ may be "halted" at $t_{c}$.
(ii) $\mathrm{r}\left(T_{i}^{j+1}\right)$ may be redefined to be less than $\mathrm{d}\left(T_{i}^{j}\right)$.

In the sporadic model defined earlier, every job's deadline is at or before its successor's release. As we will discuss in Section 1.5.2, the reason why the above two actions may occur is that the value of $\mathrm{r}\left(T_{i}^{j+1}\right)$ may change as a result of a reweighting event. The reweighting rules we present later in this section state the conditions under which the above actions may occur and the number of time units before $\mathrm{d}\left(T_{i}^{j}\right)$ that job $T_{i}^{j+1}$ can be released.

As has already been discussed, when a task reweights, there can be a difference between when it initiates the change and when the change is enacted. The time at which the change
is initiated is a user-defined time; the time at which the change is enacted is dictated by a set of conditions that differ slightly for each type of multiprocessor reweighting algorithm. Furthermore, the release and deadline of a job of an adaptable task is defined based on the weight of the task when it was released.

### 1.5.2 Reweighting Rules.

The second major contribution of this dissertation is the construction of reweighting rules for PEDF, NP-PEDF, GEDF, NP-GEDF, and PD². (The reweighting rules for PEDF, and by extension for NP-PEDF, were proposed by Block and Anderson in (Block and Anderson, 2006); the reweighting rules for GEDF and NP-GEDF were proposed by Block et al. in (Block et al., 2008b); and the reweighting rules for PD ${ }^{2}$ were proposed by Block et al. in (Block et al., 2008a).) In all five reweighting algorithms, tasks change weight via one of two rules that are based on whether a task's active job is over- or under-allocated relative to an ideal schedule.

- If a task is under-allocated, then the change is enacted by immediately halting the active job and releasing a new job with the remaining execution time.
- If a task is over-allocated, then one of two actions occurs: (i) if the task increases its weight, then the change is enacted by immediately halting the active job and releasing the next job when the task's "ideal" allocation equals its "actual" allocation; (ii) if the task decreases its weight, then the active job immediately halts and when the task's "ideal" allocation equals its "actual" allocation, the change is enacted and the next job is released.

When a $T_{i}^{j}$ is halted at time $t, T_{i}^{j}$ is not scheduled after time $t$.
Example (Figure 1.8). Consider the examples depicted in Figure 1.8. Insets (a) and (b) pertain to a one-processor system with: $T_{1}:(1,2)$, which leaves the system at time $2 ; T_{2}:(1,6)$; and $T_{3}$, with an execution cost of 3 and an initial weight of $1 / 6$ that increases to $4 / 6$ at time 2 . Since the first jobs of $T_{2}$ and $T_{3}$ have the same deadline, we can arbitrarily choose which job has a higher scheduling priority. Inset (a) depicts the case where $T_{2}$ has higher priority, and as a result, $T_{3}$ is "under-allocated" relative to the ideal system when it reweights at time 2.

Thus, when $T_{3}$ reweights, its current job "halts" and its next job is immediately released. Inset (b) depicts the case where $T_{3}$ has higher priority, and as a result, $T_{3}$ is "over-allocated" relative to the ideal system, when it reweights at time 2 . Thus, when $T_{3}$ reweights, its current job "halts" and its next job is not released until the difference between $T_{3}$ 's ideal and actual allocations is zero (at time 3). The dotted lines denote the interval of the first job of $T_{3}$ that has been changed by the reweighting event. Inset (c) depicts the ideal allocation of $T_{3}$ as a function of time. Notice that, at time 2, when $T_{3}$ increases its weight from $1 / 6$ to $4 / 6$, the ideal allocation rate increases from $1 / 6$ to $4 / 6$, and that at time $3, T_{3}$ 's total ideal allocation equals 1. Inset ( d ) depicts a one-processor system with four tasks: $T_{1}:(1,2)$, which joins the system at time $1.5 ; T_{2}:(1,6) ; T_{3}:(1,6)$; and $T_{4}$, which has an initial weight of $4 / 6$ that decreases to $1 / 6$ at time 1. Since $T_{4}$ is over-allocated at time 1 and decreases its weight, its weight change is enacted when its ideal allocation equals its actual allocation (at time 1.5). Inset (e) depicts ideal allocation for $T_{4}$.

One important property of the above reweighting rules is that reweighting events are taskindependent. That is, if a task $T_{i}$ changes its weight, then only the releases and deadlines of jobs of $T_{i}$ change. As a result of this independence, a task can only increase its weight at time $t$ if enough capacity for the reweighting event is available. For example, if a two-processor task system has three tasks each of weight 0.6 , then none of these tasks can initiate a weight increase to 0.9 (since this would cause the system load to be 2.1 ), unless one of the other two tasks first decreases its weight or leaves the system. (If reweighting events were not task-independent, then it would be possible for a proportional-share scheduling algorithm to increase the weight of a task at the expense of the shares of other tasks in the system.)

### 1.5.3 Evaluating Algorithms

The next contribution of this dissertation is a comparison of the reweighting algorithm for different multiprocessor scheduling frameworks. In order to evaluate the reweighting algorithms, we use three different metrics: overload, tardiness, and drift. Overload is the maximal amount that a single processor is overutilized (assuming that the system is not overutilized). Tardiness is the maximal amount by which a task can miss a deadline. Drift is the maximal amount of


Figure 1.8: Several one-processor systems scheduled by EDF using our reweighting rules. (a) Increasing the weight of a task $\left(T_{3}\right)$ when it is under-allocated. (b) Increasing the weight of a task $\left(T_{3}\right)$ when it is over-allocated. (c) $T_{3}$ 's ideal allocation. (d) Decreasing the weight of a task $\left(T_{4}\right)$. (e) $T_{4}$ 's ideal allocation.

| Scheme | Tardiness | Drift | Overload |
| :---: | :---: | :---: | :---: |
| PD $^{2}$ | $0^{3}$ | 2 | 0 |
| PEDF | 1 | $\mathrm{e}_{\max }\left(T_{i}\right)$ | W |
| NP-PEDF | $\mathrm{e}\left(T_{i}^{j}\right)+\mathrm{e}_{\max }\left(T_{i}\right)+1$ | $\mathrm{e}_{\max }\left(T_{i}\right)$ | W |
| GEDF | $\kappa(m-1)+\mathrm{e}_{\max }\left(T_{i}\right)$ | $\mathrm{e}_{\max }\left(T_{i}\right)$ | 0 |
| NP-GEDF | $\kappa(m)+\mathrm{e}_{\max }\left(T_{i}\right)$ | $\mathrm{e}_{\max }\left(T_{i}\right)$ | 0 |

Table 1.3: Summary of worst-case results for reweighting systems.
computation time a task "loses" between the initiation and enaction of a reweighting event. A comparison of the adaptable scheduling algorithms considered in this dissertation is given in Table 1.3. In this table, $\mathrm{e}_{\max }\left(T_{i}\right)$ denotes the maximum execution time of any job of task $T_{i} ; \mathrm{wt}_{\max }\left(T_{i}\right)$ denotes the maximal weight of task $T_{i}$ at any time; and $W$ denotes the maximal weight of the $(m \cdot\lfloor 1 / X\rfloor+1)^{s t}$ heaviest task (by maximal weight) in $\tau$, where $m$ is the number of processors, $\tau$ is the set of all tasks, and $X$ is the maximal weight of the heaviest task. In addition,

$$
\kappa(\ell)=\frac{\sum_{T_{z} \in \mathcal{E}(\ell)} \mathrm{e}_{\max }\left(T_{z}\right)}{m-\sum_{T_{z} \in \mathcal{X}(T) \ell-1} \mathcal{W}\left(T_{z}\right)}+\mathrm{e}_{\max }\left(T_{i}\right)
$$

where $\mathcal{E}(\ell)$ is the set of $\ell$ tasks in $\tau$ with the highest maximal execution time and $\mathcal{X}(\ell-1)$ is the set of $\ell-1$ tasks in $\tau$ of largest maximal weight.

Overload. As mentioned above, overload is the maximal amount that a single processor is overutilized. For example, in Figure $1.4(\mathrm{~b})$, the system is not overutilized (it has two processors and its utilization is two), yet one processor is overloaded by $1 / 3$. Hence, overload is $1 / 3$. Assuming that any reasonable allocation decreasing partitioning algorithm (i.e., tasks are sorted decreasing by weight before being assigned to processors) is used, no processor can be overutilized by more than the weight of the lightest task assigned to it (Lopez et al., 2004). More specifically, for partitioned algorithms, no processor is overutilized by more than $W$, where $W$ is the weight of of the $\left(m \cdot\left\lfloor\frac{1}{X}\right\rfloor+1\right)^{s t}$ heaviest task and $X$ is the weight of the heaviest task (Block and Anderson, 2006). (In Figure 1.4, $T_{1}$ is the $\left(m \cdot\left\lfloor\frac{1}{X}\right\rfloor+1\right)^{\text {st }}$ heaviest task.) Thus, for partitioned algorithms, the overload error is $W$. As discussed in Section 1.4, under global scheduling algorithms, no single processor can be overutilized. Hence, for global scheduling algorithms, overload is zero.

Tardiness. Recall that one of the advantages for both PEDF and PD ${ }^{2}$ is that both of these algorithms have zero tardiness (but, for PEDF, this requires allowing guaranteed weights to be less than desired weights). GEDF has a maximal tardiness that is at most the value given in (1.1). For the adaptive variants of these algorithms, the tardiness for both PEDF and PD ${ }^{2}$ is still zero, whereas the tardiness bound for GEDF must be modified to include dynamic weights and execution times. Specifically, in Chapter 3, we establish that tardiness under adaptive GEDF is at most

$$
\begin{equation*}
\frac{\sum_{k=1}^{\Gamma(t)} \max _{\mathrm{e}}(k)}{m-\sum_{k=1}^{\Gamma(t)-1} \max _{\mathrm{wt}}(k)}+\mathrm{e}\left(T_{i}^{j}\right) \tag{1.3}
\end{equation*}
$$

provided that, for all $t, \mathrm{WT}(\tau, t) \leq m$, where $\mathrm{WT}(\tau, t)=\sum_{T_{i} \in \tau} \mathrm{wt}\left(T_{i}, t\right)$, and $\max _{\mathrm{e}}(k)$ and $\max _{\mathrm{wt}}(k)$ are, respectively, the $k^{\text {th }}$ largest maximal execution time and weight of any task, and

$$
\Gamma(t)=\left\{\begin{array}{lr}
\mathrm{WT}(\tau, t)-1, & \mathrm{WT}(\tau, t) \text { is integral } \\
\lfloor\mathrm{WT}(\tau, t)\rfloor, & \text { otherwise }
\end{array}\right.
$$

Note that, because NP-PEDF and NP-GEDF are non-preemptive, their tardiness bounds are slightly larger than those for PEDF (which is zero) and GEDF, respectively.

Drift. For most non-adaptive real-time scheduling algorithms, the difference between a task's actual and ideal allocation lies within some bounded range centered at zero. For example, under a uniprocessor EDF schedule of a sporadic system, this difference lies within $\left(-\mathrm{e}\left(T_{i}\right), \mathrm{e}\left(T_{i}\right)\right)$. When a weight change occurs, the same bounds are maintained except that they may be centered at a different value. For example, in Figure 1.8(b), the range is originally $(-1,1)$ for task $T_{3}$, but after the reweighting event, it is $(-4 / 6,8 / 6)$. This lost allocation is called drift. Given this loss (barring further reweighting events), $T_{i}$ 's drift will not change. In general, a task's drift per reweighting event will be non-negative if it increases its weight, or non-positive if it decreases its weight. For GEDF, NP-GEDF, PEDF, and NP-PEDF the maximal absolute value of the drift for $T_{i}^{j}$ is at most the maximal execution time of $T_{i}$ (Block and Anderson, 2006; Block et al., 2008b). (For the remainder of this dissertation, we will use

[^2]the term absolute drift to refer to the absolute value of the drift.) Under $\mathrm{PD}^{2}$ the maximal absolute drift is 2 (Block et al., 2008a).

There is one subtle issue with the reweighting rules given in Section 1.5.2 that is important to mention. When designing the above rules, one of the guiding principals was that the reweighting rules be task-independent. As a result, when a task's weight decreases, if the active job is over-allocated and decreases its weight, then the capacity gained by decreasing the weight cannot be "released" until at or after the time at which the ideal allocation for the task equals its actual allocation. For example, in Figure 1.8(d), the capacity created by decreasing $T_{4}$ 's weight from $4 / 6$ to $1 / 6$ is not "released" until the time at which $T_{4}$ 's actual allocation equals its ideal allocation at time 1.5. This occurs even though $T_{4}$ decrease its weight at time 1. If such a delay did not exist, then a task could artificially increase its weight by continually decreasing and increasing its weight.

Example (Figure 1.9). Consider the example Figure 1.9, which depicts a one-processor system with six tasks: $T_{1}:(8,10)$ and $T_{2}, \ldots, T_{6}$ each of which has an execution cost of one, an initial weight of $1 / 5$, initiates a weight decrease to a weight of 0 immediately after being scheduled, and joins the system as soon capacity is available. Since the tasks in the set $T_{2}, \ldots$, $T_{6}$ continually decrease and increase their weights, over the range $[0,10$ ), these tasks receive one-half of the available capacity, even though they should only receive one-fifth. As a result, $T_{1}^{1}$ misses its deadline by three time units. This example can be easily extended (by decreasing task weights) to construct scenarios with arbitrarily large tardiness. Since the capacity from a weight decrease cannot be released until ideal and actual allocations match, when a task decreases its weight there might exist an arbitrarily long length of time before the capacity is released. The problem with this delay is that, not only does it cause constant positive drift, but it may also cause a reweighting event initiation to be delayed for an arbitrarily long period of time. However, for any reweighting algorithm that both ensures bounded tardiness and uses task-independent reweighting, this delay is unavoidable.


Figure 1.9: A one-processor system scheduled by EDF, which illustrates why task leaves must be delayed.

Service Levels Changed


Figure 1.10: The AGEDF system.

### 1.5.4 AGEDF

The multiprocessor reweighting algorithms we have developed require that the desired weights be provided. Desired weights can be determined using feedback-based techniques that use run time conditions. It is desirable that weight adjustments be enacted in a way that attempts to maximize overall QoS. Moreover, the allocation scheme should not "over-react" in adjusting weights when transient overloads occur. While there has been extensive work on adaptive feedback-based frameworks for uniprocessor systems (Abeni et al., 2002; Cucinotta et al., 2004; Lu et al., 2000; Lu et al., 1999), there has been relatively little work on multiprocessor feedback-based adaptive frameworks and the work that has been done has focused on non-preemptive systems where worst-case execution times, best-case execution times, and deadlines are static (Al-Omari et al., 2003; Sahoo et al., 2002).

In this dissertation, we remedy the shortcomings of prior adaptive multiprocessor real-time systems by presenting an adaptive GEDF framework called AGEDF (depicted in Figure 1.10), which consists of three components: a predictor, which uses feedback techniques to estimate the processor shares of future jobs; an optimizer, which uses the estimated processor shares
of tasks to determine a new set of service levels; and several reweighting rules, which change the service levels to match those determined by the optimizer. To the best of our knowledge, this is the first such adaptive global framework for multiprocessor systems to be proposed. The reason why we constructed such a framework for GEDF and not for any of the other algorithms is because the feedback techniques used in this framework are best suited for SRT systems, and, as we discussed in Section 1.2.5, GEDF-based algorithms have superior SRT scheduability compared to PEDF- and PD²-based algorithms.

### 1.5.5 AGEDF Implementation

As mentioned above, our research group developed a testbed called LITMUS ${ }^{\text {RT }}$ that allows different multiprocessor scheduling algorithms to be linked as plug-in components (Calandrino et al., 2006). As a final contribution, we implemented AGEDF as a LITMUS ${ }^{\text {RT }}$ plug-in. Our implementation of AGEDF consists of both a user-space library and kernel support added to LITMUS ${ }^{\text {RT }}$. In this section, we briefly discuss both parts.

Because LITMUS ${ }^{\text {RT }}$ was designed for sporadic tasks provisioned using WCETs, several modifications were needed to support adaptable sporadic tasks. These included: adjusting the internal structure of a task to allow each task to have multiple service levels; disabling the enforcement of WCETs to allow tasks to overrun their expected allocation; and modifying LITMUS ${ }^{\text {RT }}$ to allow task statistics such as actual execution times to be gathered.

After making these changes, we implemented AGEDF by changing the GEDF scheduling algorithm (which had already been implemented in LITMUS ${ }^{\text {RT }}$ ) in two ways. First, we introduced a system call to query the kernel in order for a task to determine its current service level. Second, we implemented the feedback, optimization, and reweighting components in kernel space.

After modifying LITMUS ${ }^{\text {RT }}$, we then evaluated its performance by using the core operations of both Whisper and VEC (correlation computations and bilateral filters, respectively). In this evaluation, AGEDF proved to be an extensible scheduling framework that can be easily configured to support different optimization criteria. Moreover, it exhibited good performance in scenarios in which the use of a non-adaptive GEDF algorithm would result in significant
system over-utilization.

### 1.6 Organization

The organization of this dissertation is as follows. In Chapter 2, we review prior work on adaptive, real-time, and multimedia systems. In Chapters 3, 4, and 5, we define and prove the reweighting rules for GEDF, PEDF, and $\mathrm{PD}^{2}$, respectively. In Chapter 6, we define the adaptable framework AGEDF. In Chapter 7, we discuss the implementation of AGEDF under LITMUS ${ }^{\text {RT }}$ and present an experimental evaluation of AGEDF. Finally, we conclude in Chapter 8.

## CHAPTER 2

## PRIOR WORK

In this chapter, we review in detail prior work on adaptive real-time systems and feedbackcontrol theory. As mentioned in Chapter 1, the adaptive framework we propose consists of two components: one that changes the parameters of running tasks (e.g., periods and weights) and one that determines such parameters. In Sections 2.1-2.2, we review three approaches for changing the periods and weights of running tasks. In Section 2.4, we review feedback-control theory, which is used for determining task execution times in both Lu et al.'s Feedback-Control Real-Time Scheduling (FCS) framework (Lu et al., 2002), which we cover in Section 2.5, and Abeni et al.'s Adaptive Reservation-Based Scheduler (Abeni et al., 2002), which we cover in Section 2.6.

### 2.1 Leave/Join Reweighting

As was mentioned in Section 1.3.1, under leave/join reweighting (Srinivasan and Anderson, 2005), a task's weight is changed at job boundaries by forcing it to leave with its old weight and rejoin with its new weight.

Example (Figure 2.1). Consider the example in Figure 2.1, which depicts a one-processor system with three tasks: $T_{1}:(1,2)$, which leaves at time $2 ; T_{2}$, which has an initial weight of $1 / 4$, an execution cost of 1 , and "initiates" a weight increase at time 2 to a weight of $3 / 4$, which is enacted by leave/join reweighting at $T_{2}^{1}$ 's deadline (i.e., time 4 ); and $T_{3}:(2,8)$. (This is the same system that was depicted in Figure 2.1, but has been repeated here to improve readability.) Inset (a) illustrates the EDF schedule and inset (b) illustrates the ideal and actual allocations to task $T_{2}$. Notice that, even though $T_{2}$ initiates its change at time 2 and capacity exists for $T_{2}$ to increase its weight, this weight change cannot be enacted until


Figure 2.1: The (a) EDF schedule and (b) ideal and actual allocations to $T_{2}$ in a one-processor example of leave/join reweighting.
its deadline. This illustrates that the primary drawback of leave/join reweighting is that a task can only change its weight at job boundaries. As a result, $T_{2}$ 's allocation in the actual schedule drifts from its allocations in the ideal schedule by one quantum.

Since leave/join reweighting cannot enact a reweighting event until a job boundary, it is possible that a task can incur an arbitrarily large amount of drift for one reweighting event.

### 2.2 Rate-Based Earliest Deadline

Under rate-based earliest-deadline (RBED) scheduling (Brandt et al., 2003), tasks are scheduled on a uniprocessor on an EDF basis and can change their weights and periods via four different rules. Specifically, under RBED, if a task $T_{i}$ changes its weight or period at time $t$, then the active job $T_{i}^{j}$ of $T_{i}$ at $t$ is modified as follows, where $x$ denotes the amount of time for which $T_{i}^{j}$ has been scheduled before $t$ :

1. $T_{i}$ increases its period to $P . T_{i}^{j}$,s deadline and period are immediately increased in accordance with the new period. Moreover, if $T_{i}^{j}$,s deadline is increased to time $D$, then the execution time of $T_{i}^{j}$ is increased by $\operatorname{wt}\left(T_{i}\right) \cdot\left(D-\mathrm{d}\left(T_{i}^{j}\right)\right)$.
2. $T_{i}$ decreases its period to $P$. If $x \geq \mathrm{wt}\left(T_{i}\right) \cdot\left(t-\mathrm{r}\left(T_{i}^{j}\right)\right)$, then $T_{i}^{j}$,s deadline is changed to $\mathrm{r}\left(T_{i}^{j}\right)+\max \left(x / \mathrm{wt}\left(T_{i}\right), P\right)$; otherwise, $T_{i}^{j}$,s deadline is unchanged. Moreover, if $T_{i}^{j}$,s deadline is decreased to time $D$, then $T_{i}^{j}$,s execution time is reduced by $\mathrm{wt}\left(T_{i}\right) \cdot\left(\mathrm{d}\left(T_{i}^{j}\right)-\right.$
$D)$. Regardless of the value to which $T_{i}^{j}$,s deadline was changed, $T_{i}$ 's period is changed to $P$. (Thus, $\mathrm{d}\left(T_{i}^{j+1}\right)-\mathrm{r}\left(T_{i}^{j+1}\right)=P$.)
3. $T_{i}$ increases its weight from Ow to Nw . If $\mathrm{Nw}-\mathrm{Ow}$ is at most the amount of spare utilization (i.e., the total utilization of all other tasks at most $1+\mathrm{Ow}-\mathrm{Nw}$ ), then $T_{i}$ can increase its weight. If $T_{i}$ increases its weight, then $T_{i}^{j}$,s execution time increases by $\left(\mathrm{d}\left(T_{i}^{j}\right)-t\right) \cdot(\mathrm{Nw}-\mathrm{Ow})$.
4. $T_{i}$ decreases its weight from Ow to $\mathrm{Nw} . T_{i}$ can decrease its weight to Nw only if $\mathrm{Nw} \geq \mathrm{Ow}-\frac{x}{t-\mathrm{r}\left(T_{i}^{j}\right)}$. Moreover, if $T_{i}$ 's weight is decreased at time $t$, then $T_{i}^{j}$,s execution time is reduced by $(\mathrm{Ow}-\mathrm{Nw}) \cdot\left(\mathrm{d}\left(T_{i}^{j}\right)-t\right)$.

Intuitively, these rules seek to prevent a task from artificially increasing its allocations by repeatedly changing its weight and/or period. (Notice that in Rule $2, \mathrm{r}\left(T_{i}^{j}\right)+x / \mathrm{wt}\left(T_{i}\right) \leq \mathrm{d}\left(T_{i}^{j}\right)$ since $x \leq \mathrm{e}\left(T_{i}^{j}\right)$ and $\left.\mathrm{d}\left(T_{i}^{j}\right)=\mathrm{r}\left(T_{i}^{j}\right)+\mathrm{e}\left(T_{i}^{j}\right) / \mathrm{wt}\left(T_{i}\right).\right)$

Example (Figure 2.2). Rules 1 through 4 are illustrated in Figure 2.2 via the following examples (each involving three tasks executing on one processor):
(a) $T_{1}, T_{2}$, and $T_{3}$ all have an initial weight of $1 / 3$ and period of 3 . At time $1, T_{1}$ increases its period to 5 via Rule 1. As a result, $T_{1}^{1}$ 's execution time increases from 1 to $5 / 3$.
(b) $T_{1}, T_{2}$, and $T_{3}$ all have an initial weight of $1 / 3$ and period of 6 . At time $1, T_{1}$ decreases its period to 3 via Rule 2. As a result, $T_{1}^{1}$ 's execution time decreases from 2 to 1 .
(c) $T_{1}$ has an initial weight of $1 / 6$ and period of $3 . T_{2}$ and $T_{3}$ both have an initial weight of $1 / 3$ and period of 3 . At time $1, T_{1}$ increases its weight to $1 / 3$ via Rule 3 . As a result, $T_{1}^{1}$ 's execution time increases from $1 / 2$ to $5 / 6$.
(d) $T_{1}, T_{2}$, and $T_{3}$ all have an initial weight of $1 / 3$ and period of 6 . At time $3, T_{1}$ decreases its weight to 0 via Rule 4. As a result, $T_{1}^{1}$ 's execution time decreases from 2 to 0 .

The primary drawback to RBED is that a task cannot decrease its period if it is underallocated relative to the ideal schedule (i.e., $\left.x<\mathrm{wt}\left(T_{i}\right) \cdot\left(t-\mathrm{r}\left(T_{i}^{j}\right)\right)\right)$. Thus, if an under-allocated task wants to decrease its period, it may be forced to wait until its deadline. Moreover, if


Figure 2.2: Several one-processor examples of RBED. Insets (a)-(d) illustrate Rules 1 through 4, respectively.
a task decreases its period in order to increase its weight (one of only two ways for a task to change its weight), then the task can incur an arbitrarily large amount of drift for one reweighting event.

### 2.3 Earliest Eligible Virtual Deadline First

As mentioned in Section 1.3.3, under proportional share scheduling (Stoica et al., 1996), the guaranteed weight of a task $T_{i}$ at time $t$ is defined as

$$
\begin{equation*}
\operatorname{Gwt}\left(T_{i}, t\right)=\frac{\mathrm{wt}\left(T_{i}\right)}{\mathrm{WT}(\tau, t)}, \tag{2.1}
\end{equation*}
$$

where $\mathrm{wt}\left(T_{i}\right)$ is the desired weight of $T_{i}$ and $\mathrm{WT}(\tau, t)$ is the total desired weight of all tasks in the system $\tau$ at time $t .{ }^{1}$

One of the preferred algorithms for proportional share scheduling on uniprocessors is the earliest-eligible-virtual-deadline-first (EEVDF) algorithm, proposed by Stoica et al. (Stoica et al., 1996). Under EEVDF, task releases and deadlines are defined based on an additional notion of time called virtual time. Virtual time, unlike "real time," scales with the processor

[^3]load (i.e., as the processor load increases, virtual time slows down, and as the processor load decreases, virtual time speeds up). Intuitively, virtual time acts as a scaling factor for the system that allows task weights to be changed with minimal overhead. Specifically, the virtual time of the system $\tau$ at real time $t$ is
\[

$$
\begin{equation*}
\operatorname{vt}(\tau, t)=\int_{0}^{t} \frac{1}{\mathrm{WT}(\tau, u)} d u \tag{2.2}
\end{equation*}
$$

\]

EEVDF utilizes the notion of virtual time by assigning each job both a virtual deadline and virtual release, which are defined as

$$
\begin{aligned}
\operatorname{vr}\left(T_{i}^{j}\right) & =(j-1) \cdot \frac{\mathrm{e}\left(T_{i}\right)}{\mathrm{wt}\left(T_{i}\right)}+\Theta\left(T_{i}^{j}\right) \\
\operatorname{vd}\left(T_{i}^{j}\right) & =\operatorname{vr}\left(T_{i}^{j}\right)+\frac{\mathrm{e}\left(T_{i}\right)}{\mathrm{wt}\left(T_{i}\right)},
\end{aligned}
$$

respectively, where $j>0$ and $\Theta\left(T_{i}^{j+1}\right) \geq \Theta\left(T_{i}^{j}\right) \geq 0 . \Theta\left(T_{i}^{j}\right)$ is similar to the notion of a sporadic separation considered earlier except that it is measured in the virtual-time domain. Furthermore, under EEVDF, each task is scheduled on an EDF basis using virtual deadlines. Since virtual time scales with the system load and all deadlines are defined in terms of virtual time, the (real) time deadline of each job scales with the processor load so that no task misses its deadline.

Example (Figure 2.3). Consider the example in Figure 2.3, which depicts a one-processor system scheduled by EEVDF with six tasks, each of which has an execution time of one: $T_{1}$, which has a desired weight of $1 / 2$, is initially in the system, and leaves at (real) time 2 ; $T_{2}$ and $T_{3}$, both of which have a desired weight of $1 / 4$ and are initially in the system; $T_{4}$, which has a desired weight of $1 / 2$ and joins the system at (real) time 3 ; and $T_{5}$ and $T_{6}$, both of which have a desired weight of $1 / 2$ and join the system at (real) time 5 . Inset (a) shows the mapping of virtual time to (real) time, and inset (b) shows the EEVDF schedule. In this example, the system is fully utilized over the (real) time ranges $[0,2)$ and $[3,5$ ), and as a result, virtual time and (real) time progress at the same rate. Also, the system is underutilized over the (real) time range $[2,3$ ), and as a result, virtual time progresses faster than


Figure 2.3: A one-processor system scheduled by EEVDF. (a) The mapping of virtual time to real time. The total desired weight of active tasks for each time range $[t, t+1)$ is labeled across the top axis. (b) The EEVDF schedule.
(real) time. Finally, the system is over-utilized over the (real) time range [5, 9), and as a result, virtual time progresses slower than real time.

One final note: while EEVDF can change the guaranteed weights of all tasks proportionally with minimal overhead, it cannot change the guaranteed (or desired) weights of tasks independently of each other without using leave/join reweighting. For example, in Figure 2.3(b), the only way to increase $T_{2}$ 's share to $1 / 2$ is to decrease the processor load so that $\mathrm{wt}\left(T_{2}\right)$ accounts for half the processor load. Such a scenario occurs in the example over the (real) time range $[2,3)$.

### 2.4 Feedback-Control Theory

Feedback systems use the previous states of a system in order to predict and control the future behavior of the system. In this section, we describe the basics of feedback-control theory. The review presented in this section is taken from (Nise, 2004) and (Smith, 1997).

### 2.4.1 Basics of Feedback Theory

We begin by introducing some of the central definitions and terminology of feedback-control theory. Most feedback systems consist of the following components, as labeled in Figure 2.4: the reference input value, the output value, the actuator, the error, the plant, and the controller. The reference input value is the objective value for the system, while the output value is computed by the system. The actuator calculates the error by subtracting the output from the reference input value. The plant is the system we wish to control. The controller modifies the reference input value to change the behavior of the output. Depending on how frequently the system is sampled, i.e., the output of the system is fed back to the actuator, the system is classified as either an analog or a discrete system. Specifically, if the system is sampled continually, then it is an analog system; otherwise, it is a discrete system. Since we employ feedback techniques only at job completions (i.e., at a discrete set of times), we are only considered with discrete feedback-controlled systems, in this dissertation. In discrete feedback-controlled systems, the behavior of the plant and controller are specified as difference equations, i.e., for either a plant or a controller, if $x(k)$ is the output and $e(k)$ is the


Figure 2.4: A simple feedback-control loop.
input, then $x(k)$ is of the form

$$
x(k)=b_{n} e(k)+b_{n-1} e(k-1)+\ldots+b_{o} e(k-n)-a_{n-1} x(k-1)-\ldots-a_{o} x(k-n),
$$

for some value of $n$, where $a_{i}$ and $b_{i}$ are constants and $x(k)$ and $e(k)$ are real-valued discrete functions of time $k$. Since $x(k)$ and $e(k)$ are functions of time, we say that they are in the time domain. Notice that difference equations are linear. One of the requirements for using feedback techniques is that there is a linear relationship between the input and the output.

The performance of a feedback system is measured in terms of transient response, steadystate error, and stability. The transient response of a system is the initial response of the system to a change in reference input value, as depicted in Figure 2.5. The steady-state error denotes the difference between the output and the reference input value of the system as time increases (also depicted in Figure 2.5). A system is considered to be stable if every bounded reference input value causes the system's steady-state error to be bounded. For feedback systems, it is crucially important that the system be stable.

While the behavior of a system is often defined in the time domain, when analyzing a feedback system, it is often helpful to transform the formulas that define the plant to the frequency domain, i.e., as a function of frequencies. In order to make this transformation, we use the $z$-transform, which is defined as

$$
\begin{equation*}
F(z)=\sum_{k=0}^{\infty} f(k) z^{-k} \tag{2.3}
\end{equation*}
$$

for the function $f(k)$ in the time domain. For the function $f(k)$, the convention is to write the


Figure 2.5: An example of an over-damped, under-damped, and critically-damped feedback system responding to a step input.
$z$-transformed function as $F(z)$, where $z$ is a complex variable. A more detailed discussion of the $z$-transform can be found in (Nise, 2004).

By taking the $z$-transform of the plant (or controller), we get its transfer function, which relates the input of the plant (or controller) to its output. More specifically, if $i(k)$ is the input of a plant (or controller) and $p(k)$ is its output, then the transfer function for the plant (or controller) is given by $\frac{P(z)}{I(z)}$, where $I(z)$ is the $z$-transform of $i(k)$ and $P(z)$ is the $z$-transform of $p(k)$. For example, consider the system depicted in Figure 2.6 in which the difference equation for the controller is

$$
c(k+1)=a_{1} \cdot e(k)+a_{2} \sum_{j=1}^{j=k-1} e(j),
$$

and the plant is

$$
m(k)=b \cdot c(k) .
$$

In this example, the transfer function for the controller is given as

$$
\frac{C(z)}{E(z)}=\frac{a_{1}\left(z-\frac{a_{1}-a_{2}}{a_{1}}\right)}{z(z-1)},
$$



Figure 2.6: An example feedback-control loop, where the controller is defined as $c(k+1)=$ $a_{1} e(k)+a_{2} \sum_{j=0}^{j=k-1} e(j)$ and the plant is defined as $m(k)=b \cdot c(k)$.
and the transfer function for the plant is given as

$$
\frac{M(z)}{C(z)}=b .
$$

By combining the transfer function of the plant and the controller, we get the open-loop transfer function, which is so called because it ignores the feedback loop. In Figure 2.6, the open-loop transfer function is given by

$$
G(z)=\frac{M(z)}{E(z)}=\frac{M(z)}{C(z)} \frac{C(z)}{E(z)}=b \frac{a_{1}\left(z-\frac{a_{1}-a_{2}}{a_{1}}\right)}{z(z-1)} .
$$

The open-loop zeroes of a system are the values of $z$ such that the numerator of the open-loop transfer function equals zero. Similarly, the open-loop poles of a system are the values of $z$ such that the denominator of the open-loop transfer function equals zero. In the system depicted in Figure 2.6, the open-loop zero is $\frac{a_{1}-a_{2}}{a_{1}}$ and the open-loop poles are 0 and 1.

The closed-loop transfer function, which incorporates both the behavior of the controller and feedback loop, is given by

$$
\begin{equation*}
H(z)=\frac{G(z)}{1+G(z)} \tag{2.4}
\end{equation*}
$$

where $G(z)$ is the open-loop transfer function. In Figure 2.6, the closed-loop transfer function is

$$
H(z)=\frac{G(z)}{1+G(z)}=\frac{b a_{1}\left(z-\frac{a_{1}-a_{2}}{a_{1}}\right)}{z^{2}+\left(b a_{1}-1\right) z-b\left(a_{1}-a_{2}\right)}
$$

The closed-loop zeroes of the system are the values of $z$ such that the numerator of the closed-loop transfer function equals zero. Similarly, the closed-loop poles of the system are the values of $z$ such that the denominator of the closed-loop transfer function equals zero. In
our previous example, the closed-loop zero is $\frac{a_{1}-a_{2}}{a_{1}}$ and the closed-loop poles are

$$
\frac{\left(1-b a_{1}\right) \pm \sqrt{\left(b a_{1}-1\right)^{2}+4 b\left(a_{1}-a_{2}\right)}}{2}
$$

### 2.4.2 Feedback Characteristics

In this section, we describe how to determine the stability, transient response, and steadystate error for feedback systems. For the remainder of this section, we use $\mathcal{R}(\mathcal{P})$, and $\theta(\mathcal{P})$ to denote, respectively, the radius and angle (in radians) of the pole $\mathcal{P}$ in polar-complex form, and we use $\mathcal{P}_{m}$ to denote the closed-loop pole that is farthest from the origin.

Stability. Recall from Section 2.4.1 that a system is stable if every bounded reference input value causes the system's steady-state error to be bounded. The test for stability is based on the location of the closed-loop poles in the complex plane: a system is stable if all closed-loop poles are within the unit circle in the complex plane, i.e., $\mathcal{P}_{m}<1$. A system is unstable if any closed-loop pole is outside the unit circle, i.e., $\mathcal{P}_{m}>1$. A system is marginally stable, in which case the output neither converges nor diverges, if at least one pole is on the unit circle and no pole is outside of the unit circle, i.e., $\mathcal{P}_{m}=1$. Thus, in the system from Figure 2.6, if $a_{1}=2$ and $a_{2}=1$, then system is stable if $b \in(0,2 / 3)$, the system is marginally stable if $b=0$ or $b=2 / 3$, and the system is unstable if $b>2 / 3$ or $b<0$.

Transient response. Two of the most important types of feedback systems are first- and second-order systems. A feedback system is considered to be a first-order system if it has one closed-loop pole and a system is considered to be a second-order system if has two closed-loop poles. These two types of systems are important because both have a set of simple formulas for determining their transient response. (Higher-order systems often have a transient response that is too complex to determine without approximation.)

The transient response of a system is usually evaluated by examining the behavior of the output when the system incurs a step input, i.e., the reference input value suddenly increases to a given value. Since feedback systems use previous results to predict future results, a step reference input value represents the worst-case scenario - a sudden change from one value to
another. For a first-order system, the transient response is characterized by its settling time (i.e., the time it takes for the output to attain and stay within $2 \%$ of its steady-state value) and whether the output "overshoots" its final value. For example, such a scenario is depicted by the curve labeled under-damped in Figure 2.5. For a second-order system, the transient response is characterized by its settling time and whether it is under-damped, over-damped, or critically-damped (all three are depicted in Figure 2.5.) The settling time (where time is measured in terms of samples) of the system is given by the standard formula

$$
\left\lceil\frac{-4}{\ln \left(\mathcal{R}\left(\mathcal{P}_{m}\right)\right)}\right\rceil
$$

(This formula has a ceiling because time is discrete.)

For first-order systems, the output overshoots its final value if $\mathcal{P}_{m}<0$. For first-order systems, it is typically undesirable for the output to overshoot its final value. (This is not the case for second-order systems since for second-order systems overshooting may be the only way to achieve the specified settling time. For most first-order systems, it is possible to achieve a desired settling time without overshooting the output.)

If a second-order system is over-damped, then the output will never overshoot the reference input value for a step input. If a second-order system is under-damped, then the output will overshoot the reference input value for a step input. For under-damped systems, the percent overshoot is an additional characteristic of transient response. If a second-order system is critically-damped, then the settling time is as small as possible without causing the output to overshoot the reference input value. Whether a system is under-, over-, or criticallydamped depends on the location of the closed loop poles. If both poles are unique, real, and positive, then the system is over-damped. If both poles have the same radius, are real, and are positive, then the system is critically-damped. Otherwise, the system is under-damped. For under-damped systems, the percent overshoot is given by

$$
\begin{equation*}
e^{-\left(\zeta \pi / \sqrt{1-\zeta^{2}}\right)} \cdot 100 \tag{2.5}
\end{equation*}
$$

where $\zeta$ is a value called the damping ratio and is given by

$$
\zeta=\frac{-\ln \left(\mathcal{R}\left(\mathcal{P}_{m}\right)\right)}{\sqrt{\theta\left(\mathcal{P}_{m}\right)^{2}+\ln ^{2}\left(\mathcal{R}\left(\mathcal{P}_{m}\right)\right)}}
$$

For example, in the system from Figure 2.6 , if $a_{1}=2$ and $a_{2}=1$, then the system is underdamped for any value of $b \in(0,2 / 3)$. Alternatively, if $b=1, a_{1} \approx 0.102$, and $a_{2} \approx 0.3035$, then the system is critically-damped and has a settling time of 6 time units. Finally, if $b=1$, $a_{1}=1.4008$, and $a_{2}=1.60238$, then the system is under-damped, the settling time is 5 time units, and the percent overshoot is approximately $10.3 \%$.

Steady-state error. Finally, we turn our attention to steady-state error. The steady-state error of a system is measured based on the system's response to a step and/or a ramp input. The ramp input simulates a reference input value that constantly increases by a rate of $\mathcal{T}$ per time unit. The steady-state error for a system is determined by using the final value theorem, which states that if $E(z)$ is the $z$-transform of a system's error, then the steady state error is given by

$$
\begin{equation*}
\lim _{z \rightarrow 1} \frac{z-1}{z} E(z) \tag{2.6}
\end{equation*}
$$

Since $E(z)$ can be defined as

$$
\begin{equation*}
E(z)=R(z)-M(z) \tag{2.7}
\end{equation*}
$$

where $R(z)$ is the reference input value and $M(z)$ is the output, and $M(z)$ can be defined as

$$
M(z)=E(z) G(z)
$$

where $G(z)$ is the open-loop transfer function, we get

$$
\begin{equation*}
E(z)=\frac{R(z)}{1+G(z)} \tag{2.8}
\end{equation*}
$$

Since the $z$-transform of the step input is given by $R(z)=\frac{z}{z-1}$, from (2.6) and (2.8), we
can derive the steady-state error of a system in response to a step input as

$$
\begin{equation*}
\lim _{z \rightarrow 1} \frac{1}{1+G(z)} \tag{2.9}
\end{equation*}
$$

Since the $z$-transform of the ramp input is given by $R(z)=\frac{z \mathcal{T}}{(z-1)^{2}}$, from (2.6) and (2.8), we can derive the steady-state error of a system in response to a ramp input as

$$
\begin{equation*}
\frac{\mathcal{T}}{\lim _{z \rightarrow 1}(z-1) G(z)} \tag{2.10}
\end{equation*}
$$

The value to which the $(z-1)$ term is raised in the denominator of the open-loop transfer function, $G(z)$, is called the system type, and it is used to quickly determine if the steadystate error is zero, some non-zero constant, or $\infty$. Specifically, if $G(z)$ has a system type of zero (i.e., it has no $(z-1)$ terms in its denominator), then the system has a steady-state error of some constant value for the step input and $\infty$ for any ramp input. If $G(z)$ has a system type of one (i.e., it has one $(z-1)$ term in its denominator), then the system has a steady-state error of 0 for the step input and a constant for any ramp input. Finally, if $G(z)$ has a system type of two (i.e., it has two $(z-1)$ terms in its denominator), then the system has a steady-state error of 0 for any step or ramp input.

For example, in the system depicted in Figure 2.6, since there is one $(z-1)$ term in the denominator of $G(z)$, the system type is one. Thus, it has a steady-state error of zero for a step input and a constant steady-state error for any ramp input. Specifically, if $b=0.5$, $a_{1}=2$ and $a_{2}=1$, then the steady-state error for the ramp response is $2 \mathcal{T}$. Alternatively, if $b=1, a_{1} \approx 0.102$, and $a_{2} \approx 0.3035$, then the steady-state error for the ramp input is approximately $3.295 \mathcal{T}$. Finally, if $b=1, a_{1}=1.4008$, and $a_{2}=1.60238$, then the steady-state error is approximately $0.624 \mathcal{T}$.

### 2.4.3 Controllers

As mentioned above, the purpose of a controller is to improve the system response. There are three main types of of controllers: proportional-integral (PI) controllers; proportionalderivative (PD) controllers; and proportional-integral-derivative (PID) controllers. PI con-
trollers improve the steady-state error of the system by increasing the system-type. Timedomain definitions of such controllers are of the form

$$
c(k+1)=a \cdot e(k)+b \sum_{j=1}^{j=k-1} e(j),
$$

and the $z$-transform of such a controller is

$$
G(z)=\frac{C(z)}{E(z)}=\frac{a\left(z-\frac{a-b}{a}\right)}{z(z-1)} .
$$

PD Controllers improve the transient response of the system by an additional closed-loop zero. Time-domain definitions of such controllers are of the form

$$
c(k+1)=a \cdot e(k)+b(e(k)-e(k-1)),
$$

and the $z$-transform of such a controller is

$$
G(z)=\frac{C(z)}{E(z)}=\frac{(a+b) z-b}{z^{2}} .
$$

PID Controllers improve the transient response and the steady-state error of the system by combining both PI and PD techniques. Such controllers are of the form

$$
c(k+1)=a \cdot e(k)+b(e(k)-e(k-1))+d \sum_{j=1}^{j=k-1} e(k),
$$

and the $z$-transform of such a controller is

$$
G(z)=\frac{C(z)}{E(z)}=\frac{(a+b) z^{2}+(d-a-2 b) z+b}{z^{2}(z-1)} .
$$

The problem with PID controllers is that such controllers can easily increase the system beyond a second-order system. As a result, such systems are substantially more difficult to design for a specific transient response than PI or PD controllers.


Figure 2.7: A feedback-control loop with a disturbance.

### 2.4.4 Disturbance

In addition to controlling a system based on its reference input value, a feedback system can be designed to handle disturbances, which represent an additional (typically unwanted) source of input to the system. For example, if we were constructing a feedback-controlled cruise control system, then the reference input value would be the car's desired speed and the disturbance would be the slope of the road. The typical model for a feedback system with a disturbance is shown in Figure 2.7. For such a system, the output of the system, $M(z)$, is given by

$$
\begin{equation*}
M(z)=E(z) G_{1}(z) G_{2}(z)+D(z) G_{2}(z) \tag{2.11}
\end{equation*}
$$

where $E(z)$ is the $z$-transform of the error, $G_{1}(z)$ and $G_{2}(z)$ are the transfer functions for either a plant or a controller, and $D(z)$ is the $z$-transform of the disturbance.

When constructing a system that may have a disturbance, the primary design characteristic of interest is the steady-state error in response to a step input by the disturbance. Thus, just as in Section 2.4.2, to solve for the steady state error, we need to find a transfer function that relates $D(z)$ to $E(z)$. Recall from (2.7) that

$$
\begin{equation*}
E(z)=M(z)-R(z) . \tag{2.12}
\end{equation*}
$$

By substituting (2.12) into (2.11), we get

$$
\begin{equation*}
E(z)=\frac{1}{1+G_{1}(z) G_{2}(z)} R(z)-\frac{G_{2}(z)}{1+G_{1}(z) G_{2}(z)} D(z) \tag{2.13}
\end{equation*}
$$

which defines the relationship between $E(z)$ and both the reference input value, $R(z)$, and the disturbance, $D(z)$. By applying the final value theorem to (2.13), we can obtain the steady-state error of this system as

$$
\begin{equation*}
\lim _{z \rightarrow 1} \frac{z-1}{z}\left(\frac{1}{1+G_{1}(z) G_{2}(z)} R(z)-\frac{G_{2}(z)}{1+G_{1}(z) G_{2}(z)} D(z)\right) . \tag{2.14}
\end{equation*}
$$

By isolating the disturbance term from (2.14) and substituting the transfer function for a step input into $D(z)$, i.e., $D(z)=\frac{z}{z-1}$ we obtain that the steady-state error for a step input in the disturbance is

$$
\begin{equation*}
\lim _{z \rightarrow 1}\left(-\frac{1}{\frac{1}{G_{2}(z)}+G_{1}(z)}\right) . \tag{2.15}
\end{equation*}
$$

### 2.4.5 Feedback Theory For a Predictor

It is worthwhile to note that while feedback-based techniques are primarily used to control the behavior of a plant for which the (reference) input is known, another viable use for such techniques is to predict future values of a changing and unknown input. The design of such a system is exactly the same as the typical feedback system, except that the feedback loop does not directly impact the behavior of the system. (Thus, the plant and the controller can be one-in-the-same.) In such a system, the transient response describes the initial accuracy of predictions after there has been a change in the input, and the steady-state error describes the difference between the predicted and actual values as system time increases.

By using feedback-based techniques in the construction of the predictor, instead of using a simpler approach, such as setting the current value to equal the previous value, the predictor both produces values that are less susceptible to ephemeral fluctuations in the workload and is capable of closely tracking trends in the value (i.e., such systems have a bounded steady state error for the ramp input).

### 2.5 The FCS Framework

In this section, we review the uniprocessor Feedback-Control Real-Time Scheduling (FCS) framework proposed in (Lu et al., 2002). (The FC-EDF and FC-EDF ${ }^{2}$ algorithms proposed
in (Lu et al., 1999) and (Lu et al., 2000) and mentioned in Chapter 1 are subsets of this framework.) As was mentioned in Chapter 1, in the FCS framework, unlike in most work on real-time feedback scheduling algorithms, it is assumed that the execution time of a job is unknown until it is complete. Given this limitation, the FCS framework has two objectives. First, maintain the total system utilization at some user-defined value, $u_{s}$. Second, if the system has a utilization greater than one, then maintain the miss-ratio (i.e., the fraction of jobs with a missed deadline) at some user-defined value, $\mathrm{m}_{\mathrm{s}}$. In order to construct such a system, the FCS framework assumes a task model in which each task has multiple versions (called service levels), each of which has a period, an estimated execution time, which represents the amount of time the task will require on average if it executes at that service level, and a $Q o S$ value, which represents the value to system if the task finishes before its deadline while executing at that service level. ${ }^{2}$ Throughout this section, we will use $\mathrm{p}\left(T_{i}, k\right), \mathrm{e}\left(T_{i}, k\right)$, and $\mathrm{v}\left(T_{i}, k\right)$ to denote, respectively, the period, estimated execution time, and QoS value for the $k^{\text {th }}$ service level of $T_{i}$, where service levels are ordered increasingly by estimated weight, which is defined as $\mathrm{e}\left(T_{i}, k\right) / \mathrm{p}\left(T_{i}, k\right)$. A task can only execute at one service level at a time. Without loss of generality, we assume that for any task $T_{i}$, service level 0 is defined such that $\mathrm{p}\left(T_{i}, 0\right)=0, \mathrm{e}\left(T_{i}, 0\right)=0$, and $\mathrm{v}\left(T_{i}, 0\right)=0$.

### 2.5.1 The FCS Framework's Architecture

The major components of the FCS framework are depicted in Figure 2.8. At a high level, these components function as follows.

- At each instant, the pending job with the smallest deadline is scheduled. ${ }^{3}$
- After a sampling period of length $t$ time units, where $t$ is a user-defined value, several actions occur. First, the monitor calculates in the last $t$ time units both the fraction of jobs that missed their deadlines (i.e., the miss-ratio) and the fraction of time that the processor was scheduled (i.e., the utilization). Next, the controller calculates the change

[^4]Change in Estimated Utilization


Figure 2.8: The design of the FCS.
in the total estimated utilization that is required to maintain the actual utilization at $\mathrm{u}_{s}$ or the miss-ratio at $\mathrm{m}_{\mathrm{s}}$. Finally, the QoS actuator changes the set of running tasks and the service levels of running tasks to match the total estimated utilization as calculated by the controller.

It is important to note that in the FCS framework, the miss-ratio and the actual utilization are always calculated over the previous $t$ time units.

### 2.5.2 Feedback in the Controller and QoS Actuator

The heart of the FCS framework is the controller and the QoS actuator, which use feedback techniques in order to determine the required change to the estimated utilization. In order to calculate this change, two different feedback loops are used: one loop determines the change in the estimated utilization that would maintain the actual utilization at $\mathrm{u}_{s}$, and the other loop determines the change in the estimated utilization that would maintain the miss-ratio at $\mathrm{m}_{\mathrm{s}}$. The actual change in the estimated utilization is calculated by dynamically switching between the values produced by these two loops. In this section, we describe these two feedback loops.

Before continuing, there is one subtle issues that must be discussed. As was mentioned in Section 2.4.1, the relationship between the input and output for a plant or a controller must be linear. Unfortunately, for the loop that monitors the miss-ratio, the relationship between the estimated utilization (again, the input) and the miss-ratio (the output) is nonlinear. Moreover, for the loop that monitors the actual utilization, the relationship between
the estimated utilization (the input to the plant/controller) and the actual utilization (the output to the plant/controller) cannot be determined exactly. As a result, Lu et al. use a linear estimation in both the utilization and miss-ratio feedback loops.

Utilization feedback loop. Before presenting difference equations for the utilization feedback loop, depicted in Figure 2.9(a), we introduce a few definitions. Let the value a $(k)$ denote the fraction of time the processor was busy over the $k^{t h}$ sampling period, i.e., a(k) equals the fraction of time the processor was busy over the time range $[(k-1) \cdot q, k \cdot q)$. Let Ew $(k)$ denote the total estimated utilization of all tasks in the $k^{t h}$ sampling period. Let $\mathrm{G}_{\mathrm{u}}$ to denote the maximum value of $\mathrm{a}(k) / \mathrm{Ew}(k)$ for any value of $k$. As mentioned above, the relationship between the estimated utilization, $\mathrm{Ew}(k)$, and the actual utilization, $\mathrm{a}(k)$, cannot be determined exactly. As a result, the plant in the utilization feedback loop defines a relationship between the estimated utilization and the worst-case estimated utilization, which is denoted as $\mathbf{u}(k)$. Specifically, the difference equation for the plant is defined as

$$
\begin{align*}
\mathrm{u}(k)=\mathrm{G}_{\mathrm{u}} \mathrm{Ew}(k), & \text { if } \mathrm{G}_{\mathrm{u}} \mathrm{Ew}(k) \leq 1  \tag{2.16}\\
\mathrm{u}(k)=1, & \text { if } \mathrm{G}_{\mathrm{u}} \mathrm{Ew}(k)>1 \tag{2.17}
\end{align*}
$$

Notice that this system is linear so long as it is not saturated, i.e., as long as $\mathrm{G}_{\mathrm{u}} \mathrm{Ew}(k)<1$. Also note that in the absence of saturation, (2.16) can be rewritten as

$$
\begin{equation*}
\mathrm{u}(k)=\mathrm{u}(k-1)+\mathrm{G}_{\mathrm{u}} \mathrm{c}_{\mathrm{u}}(k-1) \tag{2.18}
\end{equation*}
$$

where $\mathrm{c}_{\mathrm{u}}(k)=\mathrm{Ew}(k+1)-\mathrm{Ew}(k)$. Thus, in the absence of saturation, the open-loop transfer function for the utilization of the system can be written as

$$
\begin{equation*}
\mathrm{P}_{\mathrm{u}}(z)=\frac{\mathrm{U}(z)}{\mathrm{C}_{\mathrm{u}}(z)}=\frac{\mathrm{G}_{\mathrm{u}}}{(z-1)} \tag{2.19}
\end{equation*}
$$

The controller defines a relationship between the error of the system, which is defined as $\epsilon_{u}(k)=\mathrm{u}_{s}-\mathrm{u}(k)$, and the change in estimated utilization, $\mathrm{c}_{\mathrm{u}}(k)$. Thus, the difference
equation for the controller is given as

$$
\begin{equation*}
\mathrm{c}_{\mathrm{u}}(k)=\mathrm{K}_{\mathrm{pu}} \epsilon_{u}(k), \tag{2.20}
\end{equation*}
$$

where $\mathrm{K}_{\mathrm{pu}}$ is a tunable parameter The open-loop transfer function for (2.20) is given by

$$
\begin{equation*}
\mathrm{H}_{u}(z)=\mathrm{K}_{\mathrm{pu}} . \tag{2.21}
\end{equation*}
$$

Notice that this is a proportional controller (i.e., a P Controller). Lu et al. do not include an integral controller in their design because the formulation of the plant already has a systemtype of one. Additionally, a derivative controller is not necessary since its primary purpose is to introudce an additional variable for controlling the transient response of the system, and for the types of inputs that Lu et al. are concerned with the system has a sufficient number of variables that control the system's behavior.

From (2.19) and (2.21), we can derive an open-loop transfer function for the plant and controller combined as

$$
\begin{equation*}
\mathrm{Q}_{u}(z)=\mathrm{H}_{u}(z) \mathrm{P}_{\mathrm{u}}(z)=\frac{\mathrm{K}_{\mathrm{pu}} \mathrm{G}_{\mathrm{u}}}{(z-1)}, \tag{2.22}
\end{equation*}
$$

and the closed-loop transfer function as

$$
\begin{equation*}
\mathrm{Y}_{u}(z)=\frac{\mathrm{Q}_{u}(z)}{1+\mathrm{Q}_{u}(z)}=\frac{\mathrm{K}_{\mathrm{pu}} \mathrm{G}_{\mathrm{u}}}{z-\left(1-\mathrm{K}_{\mathrm{pu}} \mathrm{G}_{u}\right)} \tag{2.23}
\end{equation*}
$$

Miss-ratio feedback loop. As mentioned above, the exact relationship between the estimated utilization, $\mathrm{Ew}(k)$, and the miss-ratio, $\mathrm{m}(k)$, in sampling period $k$ is nonlinear. As a result, Lu et al. approximate the relationship between the two by using the derivative of this relationship around the vicinity of $\mathrm{m}_{\mathrm{s}}$. (The notion of "around the vicinity of $\mathrm{m}_{\mathrm{s}}$ " is loosely defined by Lu et al., as we will discuss in Section 2.5.3.) Specifically, the plant is defined as

$$
\begin{align*}
\mathrm{m}(k)=0, & \text { if } \mathrm{G}_{\mathrm{u}} \mathrm{Ew}(k) \leq 1  \tag{2.24}\\
\mathrm{~m}(k)=\mathrm{m}(k-1)+\mathrm{G}_{\mathrm{m}} \mathrm{G}_{\mathrm{u}} \mathrm{c}_{\mathrm{m}}(k-1), & \text { if } \mathrm{G}_{\mathrm{u}} \mathrm{Ew}(k)>1, \tag{2.25}
\end{align*}
$$

where $\mathrm{c}_{\mathrm{m}}(k)=\mathrm{Ew}(k+1)-\mathrm{Ew}(k)$ and $\mathrm{G}_{\mathrm{m}}$ is the maximum value of

$$
\frac{d \mathrm{~m}(k)}{d\left(\mathrm{G}_{\mathrm{u}} \mathrm{Ew}(k)\right)}
$$

around the vicinity of $m_{s}$. Lu et al. suggest deriving both $\frac{d m(k)}{d\left(G_{u} E w(k)\right)}$ and $G_{m}$ experimentally. Notice that, in the absence of saturation, i.e., as long as $\mathrm{m}(k)>0$ holds, the plant is linear, and as a result, its open-loop transfer function is given by

$$
\begin{equation*}
\mathrm{P}_{\mathrm{m}}(z)=\frac{\mathrm{M}(z)}{\mathrm{C}_{\mathrm{m}}(z)}=\frac{\mathrm{G}_{\mathrm{m}} \mathrm{G}_{\mathrm{u}}}{(z-1)} \tag{2.26}
\end{equation*}
$$

Just as before, the controller defines a relationship between the error of the system, which is defined as $\epsilon_{m}(k)=\mathrm{m}_{\mathrm{s}}-\mathrm{m}(k)$, and the change in estimated utilization, $\mathrm{c}_{\mathrm{m}}(k)$. Thus, the difference equation for the controller is given as

$$
\begin{equation*}
\mathrm{c}_{\mathrm{m}}(k)=\mathrm{K}_{\mathrm{pm}} \epsilon_{u}(k), \tag{2.27}
\end{equation*}
$$

where $\mathrm{K}_{\mathrm{pm}}$ is a tunable parameter. The open-loop transfer function for (2.27) is given by

$$
\begin{equation*}
\mathrm{H}_{m}(z)=\mathrm{K}_{\mathrm{pm}} \tag{2.28}
\end{equation*}
$$

Just as before, this is also a proportional controller (i.e., a P controller).

From (2.26) and (2.28), we can derive an open-loop transfer function for the plant and controller combined as

$$
\begin{equation*}
\mathrm{Q}_{m}(z)=\mathrm{H}_{m}(z) \mathrm{P}_{\mathrm{m}}(z)=\frac{\mathrm{K}_{\mathrm{pm}} \mathrm{G}_{\mathrm{m}} \mathrm{G}_{\mathrm{u}}}{(z-1)}, \tag{2.29}
\end{equation*}
$$

and the closed-loop transfer function as

$$
\begin{equation*}
\mathrm{Y}_{m}(z)=\frac{\mathrm{Q}_{m}(z)}{1+\mathrm{Q}_{m}(z)}=\frac{\mathrm{K}_{\mathrm{pm}} \mathrm{G}_{\mathrm{m}} \mathrm{G}_{\mathrm{u}}}{z-\left(1-\mathrm{K}_{\mathrm{pm}} \mathrm{G}_{\mathrm{m}} \mathrm{G}_{\mathrm{u}}\right)} \tag{2.30}
\end{equation*}
$$


(a)

(b)

Figure 2.9: The feedback loops for the FCS framework. (a) The feedback loop for controlling the utilization. (b) The feedback loop for controlling the miss ratio.

Stability, steady-state error, and transient response. By using the analysis in Section 2.4.2, it is not difficult to show that the utilization feedback loop is stable iff

$$
0<\mathrm{K}_{\mathrm{pu}}<\frac{2}{\mathrm{G}_{\mathrm{u}}}
$$

and the miss-ratio feedback loop is stable iff

$$
0<\mathrm{K}_{\mathrm{pm}}<\frac{2}{\mathrm{G}_{\mathrm{m}} \mathrm{G}_{\mathrm{u}}}
$$

Additionally, since it can be shown that both the systems described by (2.22) and (2.29) have a system type of one, the steady-state error for a step input in the reference input value is zero. ( Lu et al. are not concerned with ramp inputs, although that could also be easily derived from the analysis presented in Section 2.4.2.)

Finally, since the systems described by (2.23) and (2.30) are first-order system, the utilization feedback loop does not overshoot its final value so long as

$$
0<\mathrm{K}_{\mathrm{pu}} \leq \frac{1}{\mathrm{G}_{\mathrm{u}}}
$$

and the miss-ratio feedback loop does not overshoot its final value so long as

$$
0<\mathrm{K}_{\mathrm{pm}} \leq \frac{1}{\mathrm{G}_{\mathrm{m}} \mathrm{G}_{\mathrm{u}}}
$$

Additionally, the settling time (where time is measured by the number of sampling periods) for the utilization feedback loop is

$$
\left\lceil\frac{-4}{\ln \left(1-\mathrm{K}_{\mathrm{pu}} \mathrm{G}_{\mathrm{u}}\right)}\right\rceil \text {, }
$$

and the settling time time for the miss-ratio feedback loop is

$$
\left\lceil\frac{-4}{\ln \left(1-\mathrm{K}_{\mathrm{pm}} \mathrm{G}_{\mathrm{m}} \mathrm{G}_{\mathrm{u}}\right)}\right\rceil .
$$

Loop switching. As we previously discussed, one of the complications with controlling the miss ratio and the utilization is that it is possible for both of these variables to saturate. (Recall that the utilization saturates at 1 and the miss ratio saturates at $0 \%$.) When one of these variable becomes saturated, it is no longer possible to use that variable to control the system in any meaningful way. To resolve this issue, Lu et al. switch between using the miss-ratio and utilization feedback loops. (Unfortunately, as we will discuss in Section 2.5.4, Lu et al. do not discuss how to handle the scenario when both the miss-ratio and utilization are saturated.) Specifically, the estimated utilization for sampling period $k+1$ is defined by

$$
\begin{equation*}
\left.\operatorname{Ew}(k+1)=\operatorname{Ew}(k)+\min \left(c_{\mathbf{u}}(k), \mathrm{c}_{\mathrm{m}}(k)\right)\right) . \tag{2.31}
\end{equation*}
$$

By combining the two controllers in such a fashion, it is possible to set a nominal desired utilization $u_{s}$ at which deadlines should not be missed, and set an acceptable value of $m_{s}$ that can be handled in overloaded scenarios.

Changing the system load. The QoS actuator changes the system load via a two-step process. First, the service levels are determined. Second, the service-level changes are enacted by leave/join reweighting. Since we have already discussed leave/join reweighting in Section 2.1,
for the remainder of this section, we discuss the highest-value-density first (HVDF) algorithm that is used to determine the service levels of tasks. The value-density of the service level $k>0$ of $T_{i}$ is defined as

$$
\frac{\mathrm{v}\left(T_{i}, k\right) \mathrm{p}\left(T_{i}, k\right)}{\mathrm{e}\left(T_{i}, k\right)}
$$

The HVDF algorithm determines the service levels for tasks as follows. First, the highest value-density is calculated for each task. Second, tasks are ordered by highest value-density service level from largest to smallest. Next, each task is assigned, in order, its highest valuedensity service level until the total estimated utilization of all tasks reaches a user-defined utilization threshold, which can be any value in the range $[0,1]$.

Example. As an example, suppose that the utilization threshold is 0.51 , and there are three tasks in the system each with three service levels and for any task $T_{i}$ and service level $k>0$, $\mathrm{p}\left(T_{i}, k\right)=100$ (recall that for service level $0, \mathrm{p}\left(T_{i}, 0\right)=0, \mathrm{e}\left(T_{i}, 0\right)=0$, and $\left.\mathrm{v}\left(T_{i}, 0\right)=0\right)$. For $T_{1}, \mathrm{e}\left(T_{1}, 1\right)=20, \mathrm{v}\left(T_{1}, 1\right)=0.5, \mathrm{e}\left(T_{1}, 2\right)=30$, and $\mathrm{v}\left(T_{1}, 2\right)=0.6$. For $T_{2}, \mathrm{e}\left(T_{2}, 1\right)=20$, $\mathrm{v}\left(T_{2}, 1\right)=0.2$, $\mathrm{e}\left(T_{2}, 2\right)=30$, and $\mathrm{v}\left(T_{2}, 2\right)=0.5$. For $T_{3}, \mathrm{e}\left(T_{3}, 1\right)=20, \mathrm{v}\left(T_{3}, 1\right)=0.2$, $\mathrm{e}\left(T_{3}, 2\right)=30$, and $\mathrm{v}\left(T_{3}, 2\right)=0.6$. The service levels with the highest value-densities in this example are service level 1 of $T_{1}$ (a value density of 2.5 ), service level 2 of $T_{2}$ (a value density of $1 . \overline{6}$ ), and service level 2 of $T_{3}$ (a value density of 2). Thus, according the HVDF, $T_{1}$ is first assigned service level 1. Next, $T_{3}$ is assigned service level 2. Finally, since no other service levels with a positive weight can be assigned without exceeding the desired estimated weight of $0.51, T_{2}$ is assigned service level 0 .

### 2.5.3 Assumptions of the FCS Framework

Feedback systems are typically designed for specific scenarios. For this reason, most feedback algorithms are based on assumptions about the behavior of the system. In this section, we discuss the major assumptions made in the FCS framework.

Assumption 1. Each service level has an expected execution time that represents the "average-case" case. In Whisper such an assumption would require that there be an average
signal-to-noise ratio for every microphone speaker pair. This is probably a valid assumption for a wide range of applications; however, it is not difficult to conceive of applications that are deployed in highly-variable environments for which this assumption does not hold.

Assumption 2. The value of $G_{u}$ for each task can be determined experimentally. This is probably a valid assumption, although depending on the application $G_{u}$ could be easily over-estimated.

Assumption 3. There exists some "vicinity" around the target miss-ratio for which the derivative of the relationship between the utilization and miss-ratio is a constant, i.e., the value $G_{m}$ exists and is a constant. This assumption is probably more questionable than the previous two. Since depending on the types of tasks in the system, the relationship between the miss-ratio and utilization could vary widely.

Example (Figure 2.10). Consider the example in Figure 2.10, which depicts a one-processor system with 4 tasks: $T_{1}:(7,8) ; T_{2}:(2,8) ; T_{3}:(2,8)$; and $T_{4}:(2,8)$. In inset (a), $T_{1}$ has the highest priority. In inset (b), $T_{1}$ has the lowest priority. In Figure 2.10(a), since $T_{1}^{1}$ is scheduled first, $T_{2}-T_{4}$ all miss deadlines. In Figure $2.10(\mathrm{~b})$, since $T_{1}^{1}$ is scheduled last, $T_{2}-T_{4}$ all make their deadlines.

Notice that, even though these two schedules are both valid under EDF, the miss-ratio differs dramatically. It is worth mentioning that Lu et al. suggest that the value of $\mathrm{G}_{\mathrm{m}}$ should be determined experimentally. Thus, while it is conceivable that for some applications it may be feasible to empirically determine the value of $G_{m}$, such a value may not accurately reflect the typical behavior of the system.

### 2.5.4 Limitations of the FCS Framework

We conclude our discussion of the FCS Framework with a discussion of its limitations.

Limitation 1. The FCS framework adjusts tasks based on only the system-wide miss-ratio and utilization. As a result, if only a few tasks have actual utilizations that differ substantially


Figure 2.10: An example one of the FCS framework's assumptions. (a) $T_{1}$ has the highest scheduling priority. (b) $T_{1}$ has the lowest scheduling priority.
from their estimated utilizations, then the system is incapable of adjusting only those few tasks.

Limitation 2. Lu et al. state, without proof, that both the miss-ratio and utilization cannot be saturated at the same time in a given sampling period. While it is true that if the system is over-utilized, then jobs will eventually start missing their deadlines, they claim that in each sampling period the utilization and the miss-ratio are not both saturated. This stronger claim is not always true.

Example (Figure 2.11). Consider the example in Figure 2.11, which depicts a one-processor system with four tasks: $T_{1}:(3,8) ; T_{2}:(3,8) ; T_{3}:(3,8)$; and $T_{4}:(3,8)$. In this system, the missratio and the utilization are monitored every three time units (as denoted by the dashed line). Thus, for the first two sampling periods, the miss-ratio is $0 \%$ and the utilization is one. As a


Figure 2.11: An example one of the FCS framework's limitations.
result, both variables are saturated.

It is worthwhile to note that Lu et al. do not offer any guidelines for choosing the sampling period. However, even if we assume that the sampling periods are substantially larger than task periods (which would mitigate this limitation), it is not hard to construct example systems for which there exists at least one sampling period where both the miss-ratio and the utilization are saturated.

### 2.6 The Constant Bandwidth Server Feedback Scheduler

As mentioned in Section 2.5.4, one of the main limitations of the approach in (Lu et al., 2002) is that the system cannot identify whether an individual task has an actual execution time that deviates substantially from its estimated execution time. A uniprocessor feedbackcontrolled real-time scheduling algorithm that does not have this limitation was proposed in (Abeni et al., 2002). To accurately assign each task a weight, their algorithm monitors, for each task, the difference between the estimated and actual execution times of each job. Once the system has calculated a new estimated execution time for a future job, it adjusts the maximum fraction of the processor allocated to the task in order to reduce the number of deadline misses while ensuring that the task receives an accurate fraction of the processor.

### 2.6.1 Constant Bandwidth Server

Before describing Abeni et al.'s work in detail, we first review the Constant Bandwidth Server (CBS), first proposed by (Abeni and Buttazzo, 1998), which is the scheduling framework their system is built upon. Under CBS, each task is defined by a triple $\left(\mathrm{p}\left(T_{i}\right)\right.$, band $\left.\left(T_{i}\right), \operatorname{PR}\left(T_{i}\right)\right)$, where $\mathrm{p}\left(T_{i}\right)$ denotes the period of $T_{i}$, band $\left(T_{i}\right)$ is the bandwidth of $T_{i}$-which is the fraction of a processor allocated to $T_{i}$ (this is similar to notion of a task weight)—and $\operatorname{PR}\left(T_{i}\right)$ is the period of reservation - in every $\operatorname{PR}\left(T_{i}\right)$ time units, $T_{i}$ can be scheduled for up to band $\left(T_{i}\right) \cdot \operatorname{PR}\left(T_{i}\right)$ time units. The value band $\left(T_{i}\right) \cdot \operatorname{PR}\left(T_{i}\right)$ is called the budget of $T_{i}$ and is denoted $\operatorname{budg}\left(T_{i}\right)$. At time $q \cdot \operatorname{PR}\left(T_{i}\right)$ (where $q \geq 1$ is an integer), $T_{i}$ experiences a budget renewal. (Notice that a task's period defines when its jobs maybe released, but the period of reservation defines when the task's budget is renewed.) Tasks are scheduled on an earliest-budget-renewal-time-first basis. It is important to note that $\operatorname{PR}\left(T_{i}\right)$ can be defined as any value so long as $\mathrm{p}\left(T_{i}\right)=k \cdot \operatorname{PR}\left(T_{i}\right)$ for some integer value $k \geq 1$. One final note: if a job must execute for more than its alloted budget, then it will continue executing by using the next job's budget.

Example (Figure 2.12). Consider the example in Figure 2.12, which depicts a one-processor system scheduled by CBS with three tasks: $T_{1}$, which is defined by the triple $(4,1 / 4,4), T_{2}$, which is defined by the triple $(8,1 / 4,4)$, and $T_{3}$, which is defined by the triple $(4,1 / 2,2)$. Also, $T_{3}^{1}$ requires one additional unit of execution beyond its budget. A budget renewal is denoted by a large down-arrow. Notice that $T_{3}^{1}$ uses one time unit of $T_{3}^{2}$ 's budget. As such, $T_{3}^{1}$ executes for three time units, even though $T_{3}^{1}$ is allocated only two time units over the range $[0,4)$. As a result, $T_{3}^{1}$ 's remaining execution time is scheduled using $T_{3}^{2}$ 's budget over the time range $[4,5)$.

### 2.6.2 Feedback Framework

The objective of the CBS feedback scheduling algorithm is to maintain the smallest possible value of band $\left(T_{i}\right)$ such that each job of $T_{i}$ completes by its deadline. In order to satisfy this design object, each task has a feedback-control loop, as depicted in Figure 2.13, that monitors the "scheduling error" for each job. The scheduling error for the job $T_{i}^{j}$, denoted $\epsilon\left(T_{i}^{j}\right)$, is the
$\mathrm{T}_{\mathbf{i}}^{\mathrm{j}} \quad \mathbf{T}_{\mathbf{i}}^{\mathbf{j}}$ Scheduled $\downarrow$ Job release $\uparrow$ Job deadline $\boldsymbol{\&}$ budget renewal $\nabla^{\boldsymbol{j}}$ Budget renewal


Figure 2.12: A one-processor example of the CBS with three tasks.


Figure 2.13: The adaptive reservation-based feedback design.
difference between its period, $\mathbf{p}\left(T_{i}\right)$, and the time that the job would finish if it were assigned to a processor with a speed of band $\left(T_{i}\right)$. Specifically,

$$
\epsilon\left(T_{i}^{j+1}\right)=\left\{\begin{array}{ll}
\epsilon\left(T_{i}^{j}\right)+\operatorname{Ae}\left(T_{i}^{j}\right) \cdot \operatorname{band}\left(T_{i}^{j}\right)-\mathrm{p}\left(T_{i}\right) & \epsilon\left(T_{i}^{j}\right) \geq 0  \tag{2.32}\\
\operatorname{Ae}\left(T_{i}^{j}\right) \cdot \operatorname{band}\left(T_{i}^{j}\right)-\mathrm{p}\left(T_{i}\right) & \epsilon\left(T_{i}^{j}\right)<0
\end{array},\right.
$$

where $\operatorname{band}\left(T_{i}^{j}\right)$ is the bandwidth assigned to $T_{i}$ when $T_{i}^{j}$ is released. Recall that $\operatorname{Ae}\left(T_{i}^{j}\right)$ is the actual execution time of $T_{i}^{j}$. The reason why the term $\epsilon\left(T_{i}^{j}\right)$ is included in the calculation of $\epsilon\left(T_{i}^{j+1}\right)$ when $\epsilon\left(T_{i}^{j}\right) \geq 0$ is because, in this case, $T_{i}^{j}$ overran its budget. As a result, the first part of $T_{i}^{j+1}$,s budget is dedicated to finishing $T_{i}^{j}$. For example, in Figure 2.12, $\epsilon\left(T_{3}^{1}\right)=1$, and as a result, one quantum of $T_{3}^{2}$ 's budget is spent completing $T_{3}^{1}$. Upon the completion of a job, $T_{i}^{j}$, the scheduling error is fed into a PI controller, which we will describe in Section 2.6.3. Notice that, in this system, unlike Lu et al.'s FCS framework, there is an exact linear relationship between the input (i.e., the actual execution time) and the output (i.e., the scheduling error). (This linear relationship could not be enacted under the FCS framework because of its use of the miss-ratio and system utilization.)

### 2.6.3 Stability

One complicating factor in this system is that the behavior of the plant, i.e., the scheduling error, changes dramatically depending on whether $\epsilon\left(T_{i}^{j}\right)<0$. To resolve this issue, Abeni et al. assume that the equilibrium value of $\epsilon\left(T_{i}^{j}\right)$, denoted $\bar{\epsilon}\left(T_{i}\right)$, is far enough away from 0 that the system does not frequently switch between the two modes. By making such an assumption, the objective of the system is for $\epsilon\left(T_{i}^{j}\right)=\bar{\epsilon}\left(T_{i}\right)$ to hold. (In Section 2.6.4, we disucss the validity of this assumption.) After making such an assumption, Abeni et al. construct two different system equations based on whether $\bar{\epsilon}\left(T_{i}\right)<0$. In the remainder of this section, we explain these equations.

Case 1: $\bar{\epsilon}\left(T_{i}\right) \geq 0$. We first consider the case where $\bar{\epsilon}\left(T_{i}\right) \geq 0$. In this case, assuming that the variation around the equilibrium quantities for scheduling error, execution time, and bandwidth are small, the formula for the plant can be written as

$$
\begin{equation*}
\Delta \epsilon\left(T_{i}^{j+1}\right)=\Delta \epsilon\left(T_{i}^{j}\right)+\frac{\mathrm{p}\left(T_{i}\right)}{\overline{\mathrm{u}}\left(T_{i}\right)} \Delta \mathrm{u}\left(T_{i}^{j}\right)+\overline{\mathrm{u}}\left(T_{i}\right) \Delta \operatorname{Ae}\left(T_{i}^{j}\right) \tag{2.33}
\end{equation*}
$$

where $\Delta \epsilon\left(T_{i}^{j}\right)=\epsilon\left(T_{i}^{j}\right)-\bar{\epsilon}\left(T_{i}\right), \mathrm{u}\left(T_{i}^{j}\right)=\frac{1}{\operatorname{band}\left(T_{i}^{j}\right)}, \overline{\mathrm{u}}\left(T_{i}\right)$ is the equilibrium value of $\mathrm{u}\left(T_{i}^{j}\right)$, $\Delta \mathrm{u}\left(T_{i}^{j}\right)=\mathrm{u}\left(T_{i}^{j}\right)-\overline{\mathrm{u}}\left(T_{i}\right), \Delta \mathrm{Ae}\left(T_{i}^{j}\right)=\mathrm{Ae}\left(T_{i}^{j}\right)-\overline{\mathrm{Ae}}\left(T_{i}\right)$, and $\overline{\mathrm{Ae}}\left(T_{i}\right)$ is the equilibrium value of $\mathrm{Ae}\left(T_{i}^{j}\right)$. By taking the $z$-transform of (2.33), we get

$$
\begin{equation*}
\mathrm{E}(z)=\mathrm{F}_{c}(z) \mathrm{AE}(z)+\mathrm{F}_{u}(z) \mathrm{U}(z) \tag{2.34}
\end{equation*}
$$

where $\mathrm{E}(z)$ is the $z$-transform of $\epsilon\left(T_{i}^{j}\right), \mathrm{AE}(z)$ is the $z$-transform of $\Delta \mathrm{Ae}\left(T_{i}^{j}\right), \mathrm{U}(z)$ is the $z$-transform of $\Delta \mathrm{u}\left(T_{i}^{j}\right), \mathrm{F}_{c}(z)=\frac{\overline{\mathrm{u}}\left(T_{i}\right)}{z-1}$, and $\mathrm{F}_{u}(z)=\frac{\mathrm{p}\left(T_{i}\right)}{\overline{\mathrm{u}}\left(T_{i}\right)(z-1)}$.

The relationship between the the value $\Delta \epsilon\left(T_{i}^{j}\right)$ and $\Delta \mathrm{u}\left(T_{i}^{j}\right)$ is specified as a PI controller, which is given as

$$
\begin{equation*}
\Delta \mathrm{u}\left(T_{i}^{j}\right)=-a \cdot \Delta \epsilon\left(T_{i}^{j}\right)+b \sum_{q=1}^{j-1}\left(-\Delta \epsilon\left(T_{i}^{q}\right)\right) \tag{2.35}
\end{equation*}
$$

where both $a$ and $b$ are constants determined by the system designer. The $z$-transform of
(2.35) is given as

$$
\begin{equation*}
\frac{\mathrm{U}(z)}{-\mathrm{E}(z)}=\mathrm{G}(z)=\frac{\alpha z+\beta}{z-1} \tag{2.36}
\end{equation*}
$$

where $\alpha=a$ and $\beta=b-a$. By substituting (2.36) into (2.34), we get the closed-loop transfer function, which is defined as

$$
\begin{equation*}
\mathrm{E}(z)=\frac{\mathrm{F}_{c}(z)}{1+\mathrm{G}(z) \mathrm{F}_{u}(z)} \mathrm{AE}(z) \tag{2.37}
\end{equation*}
$$

By expanding (2.37), we get

$$
\begin{equation*}
\mathrm{E}(z)=\frac{\overline{\mathrm{u}}\left(T_{i}\right)(z-1)}{z^{2}+\left(\frac{\mathrm{p}\left(T_{i}\right)}{\overline{\mathrm{u}}\left(T_{i}\right)} \alpha-2\right) z+\beta \frac{\mathrm{p}\left(T_{i}\right)}{\overline{\mathrm{u}}\left(T_{i}\right)}+1} \mathrm{AE}(z) . \tag{2.38}
\end{equation*}
$$

Notice that the closed-loop poles of (2.38) are the values of $\mathcal{P}_{1}$ and $\mathcal{P}_{2}$ such that

$$
z^{2}+\left(\frac{\mathrm{p}\left(T_{i}\right)}{\overline{\mathrm{u}}\left(T_{i}\right)} \alpha-2\right) z+\beta \frac{\mathrm{p}\left(T_{i}\right)}{\overline{\mathrm{u}}\left(T_{i}\right)}+1=z^{2}-\left(\mathcal{P}_{1}+\mathcal{P}_{2}\right) z+\mathcal{P}_{1} \mathcal{P}_{2} .
$$

Thus, in terms of $\alpha$ and $\beta$, the closed-loop poles for (2.37) are the values of $\mathcal{P}_{1}$ and $\mathcal{P}_{2}$ such that

$$
\begin{equation*}
\alpha=\frac{\overline{\mathrm{u}}\left(T_{i}\right)\left(2-\mathcal{P}_{1}-\mathcal{P}_{2}\right)}{\mathrm{p}\left(T_{i}\right)}, \quad \beta=\frac{\overline{\mathrm{u}}\left(T_{i}\right)\left(\mathcal{P}_{1} \mathcal{P}_{2}\right)}{\mathrm{p}\left(T_{i}\right)} . \tag{2.39}
\end{equation*}
$$

Since (2.37) is a second-order system, the system is stable if the distance of both poles from the origin is less than one. Thus, for a given $\alpha$ and $\beta$, the system is stable if the values of $\mathcal{P}_{1}$ and $\mathcal{P}_{2}$ that satisfy (2.39) also satisfy.

$$
\left\|\mathcal{P}_{1}\right\|<1, \quad\left\|\mathcal{P}_{2}\right\|<1
$$

Case 2: $\bar{\epsilon}\left(T_{i}\right)<0$. From the above result, it is easy to show (by redoing the calculations) that if $\bar{\epsilon}\left(T_{i}\right)<0$, then for a given value of $\alpha$ and $\beta$, the system is stable if the values of $\mathcal{P}_{1}$ and $\mathcal{P}_{2}$ that satisfy

$$
\alpha=\frac{\overline{\mathrm{u}}\left(T_{i}\right)\left(1-\mathcal{P}_{1}-\mathcal{P}_{2}\right)}{\mathrm{p}\left(T_{i}\right)}, \quad \beta=\frac{\overline{\mathrm{u}}\left(T_{i}\right)\left(\mathcal{P}_{1} \mathcal{P}_{2}\right)}{\mathrm{p}\left(T_{i}\right)}
$$

also satisfy

$$
\left\|\mathcal{P}_{1}\right\|<1, \quad\left\|\mathcal{P}_{2}\right\|<1 .
$$

### 2.6.4 Scheduling Error Assumption

The most important assumption made by Abeni et al. is that the the scheduling error has an average value and there is little variance around this value. For applications that are not deployed in highly variable environments, this is probably a reasonable assumption. For highly variable scenarios, this assumption could be problematic because if the scheduling error rapidly switches between positive and negative values, then their analysis is invalidated.

### 2.6.5 Limitations

The major limitations of this work stem from the fact that its only objective is to determine an accurate bandwidth value for each task. As a result, the system is not capable of minimizing deadline misses when the system is overloaded. Specifically, if the system is overloaded, i.e., $\sum_{T_{i} \in \tau}$ band $\left(T_{i}\right)>1$, then the system will scale down the bandwidth of all tasks such that $\sum_{T_{i} \in \tau} \operatorname{band}\left(T_{i}\right) \leq 1$. As a result, in an overloaded scenario, it is possible that every task will miss its deadlines regardless of relative importance. Moreover, because the only manipulated variable is the bandwidth, it is not possible to mitigate an overloaded scenario by increasing either the period or reservation period of a task (which would have the impact of reducing the task's QoS to prevent deadline misses).

### 2.7 Conclusion

In this chapter, we reviewed several different techniques for changing the weight of a task (i.e., leave/join reweighting, RBED scheduling, and EEVDF scheduling). In addition, we discussed the basics of feedback control theory that will be used in this dissertation. Finally, we concluded this chapter with a discussion of two uniprocessor feedback-based approaches for adapting to external stimuli (i.e., the FCS framework and the feedback-controlled reservation based scheduling algorithm). In the remainder of this dissertation, we will extend the work reviewed in this chapter to function under a multiprocessor environment.

## CHAPTER 3

## GEDF and NP-GEDF*

In this chapter, we present the rules for reweighting tasks for both the GEDF and NP-GEDF scheduling algorithms. Before doing so so, we first define the "adaptable sporadic task model" as well as three theoretical scheduling algorithms that will be useful for describing these reweighting rules. (To improve readability, all of the terms in this chapter are summarized in Table 3.1.)

### 3.1 Adaptable Sporadic Task System

An adaptable sporadic task system is an extension of a sporadic task system, where the weight of each task $T_{i}$ is a function of time $t$, denoted $w t\left(T_{i}, t\right)$, and its execution time can vary with each job $T_{i}^{j}$, denoted e $\left(T_{i}^{j}\right)$. (The behavior of an adaptable sporadic task is defined by an execution time and weight instead of an execution time and period pair-the two approaches are equivalent but the former results in less complex reweighting rules.) For simplicity, if every job of a task $T_{i}$ has the same execution time, then we will denote this time by e $\left(T_{i}^{j}\right)$, and if the weight of task $T_{i}$ is constant, then we denote its weight as $w t\left(T_{i}\right)$.

For adaptable sporadic tasks, the absolute deadline of a job $T_{i}^{j}$, denoted $\mathrm{d}\left(T_{i}^{j}\right)$, is defined as

$$
\mathrm{d}\left(T_{i}^{j}\right)=\mathrm{r}\left(T_{i}^{j}\right)+\mathrm{e}\left(T_{i}^{j}\right) / \mathrm{wt}\left(T_{i}, \mathrm{r}\left(T_{i}^{j}\right)\right)
$$

(Recall that $\mathrm{r}\left(T_{i}^{j}\right)$ is the release time of $T_{i}^{j}$.) In the absence of reweighting, consecutive job releases $\left(\mathrm{r}\left(T_{i}^{j}\right)\right.$ and $\left.\mathrm{r}\left(T_{i}^{j+1}\right)\right)$ of a task $T_{i}$ must be separated by at least $\mathrm{e}\left(T_{i}^{j}\right) / \mathrm{wt}\left(T_{i}, \mathrm{r}\left(T_{i}^{j}\right)\right)$

[^5]| Notation | Definition |
| :---: | :---: |
| $\mathrm{wt}\left(T_{i}, t\right)$ | Weight of $T_{i}$ at time $t$. |
| wt $\left(T_{i}\right)$ | Weight of the $T_{i}$ that does not change its weight. |
| Swt( $T_{i}, t$ ) | Scheduling weight of $T_{i}$ at time $t$. |
| $\mathrm{e}\left(T_{i}^{j}\right)$ | Worst-case execution time of job $T_{i}^{j}$. |
| $\mathrm{e}\left(T_{i}\right)$ | Worst-case execution time for all jobs of $T_{i}$. |
| $\mathrm{e}_{\text {max }}\left(T_{i}\right)$ | Maximal worst-case execution time of all jobs of $T_{i}$. |
| $\mathrm{Ae}\left(T_{i}^{j}\right)$ | Actual execution time of job $T_{i}^{j}$. |
| $\mathrm{r}\left(T_{i}^{j}\right)$ | Release time of $T_{i}^{j}$. |
| $\mathrm{d}\left(T_{i}^{j}\right)$ | Deadline of $T_{i}^{j}$. |
| $\theta\left(T_{i}^{j}\right)$ | IS separation between $T_{i}^{j-1}$ and $T_{i}^{j}$. |
| SW | Non-clairvoyant scheduling-weight scheduling algorithm. While a task is active, this algorithm allocates the task its scheduling weight at each instant. |
| $\mathcal{S W}$ | SW schedule of a task system $\tau$. |
| CSW | Clairvoyant scheduling-weight scheduling algorithm. While a task is active and the allocation to its active job is less than its actual execution time, this algorithm allocates the task its scheduling weight at each instant. |
| CSW | CSW schedule of task system $\tau$. |
| IDEAL | Ideal scheduling algorithm. While a task is active, this algorithm allocates each task its weight at each instant. |
| I | IDEAL schedule of task system $\tau$. |
| $\mathcal{S}$ | Actual (i.e., GEDF or NP-GEDF) schedule of task system $\tau$. |
| $\mathrm{A}\left(\mathcal{B}, T_{i}^{j}, t_{1}, t_{2}\right)$ | Allocation to $T_{i}^{j}$ in the schedule $\mathcal{B}$ over $\left[t_{1}, t_{2}\right)$. |
| $\mathrm{A}\left(\mathcal{B}, T_{i}, t_{1}, t_{2}\right)$ | Allocation to $T_{i}$ in the schedule $\mathcal{B}$ over $\left[t_{1}, t_{2}\right)$. |
| $\operatorname{dev}\left(T_{i}^{j}, t\right)$ | Deviance of $T_{i}^{j}: \mathrm{A}\left(\mathcal{S} W, T_{i}^{j}, 0, t\right)-\mathrm{A}\left(\mathcal{S}, T_{i}^{j}, 0, t\right)$. |
| $\operatorname{drift}\left(T_{i}, t\right)$ | Drift of $T_{i}$ : $\mathrm{A}\left(\mathcal{I}, T_{i}, 0, t\right)-\mathrm{A}\left(\mathcal{C} S W, T_{i}, 0, t\right)$. |
| Ow | Scheduling weight before a reweighting event. |
| Nw | New weight after a reweighting event. |
| $\operatorname{REM}\left(T_{i}^{j}, t\right)$ | Remaining execution time of $T_{i}^{j}$ at $t$. $\mathrm{e}\left(T_{i}^{j}\right)-\mathrm{A}\left(\mathcal{S}, T_{i}^{j}, 0, t\right)$. |
| nextE ( $\left.T_{i}^{j}, t\right)$ | If REM $\left(T_{i}^{j}, t\right)>0$, then $\operatorname{nextE}\left(T_{i}^{j}, t\right)=\operatorname{REM}\left(T_{i}^{j}, t\right) ; \operatorname{else}, \operatorname{nextE}\left(T_{i}^{j}, t\right)=\mathrm{e}\left(T_{i}^{j+1}\right)$. |

Table 3.1: Summary of notation used in this chapter.


Figure 3.1: A one-processor example of an adaptable sporadic task system.
time units.

Example (Figure 3.1). Consider the example in Figure 3.1, which depicts a one-processor system with four tasks: $T_{1}$, which joins the system at time 3 with $\mathrm{wt}\left(T_{1}\right)=2 / 5$ and $\mathrm{e}\left(T_{1}\right)=2$; $T_{2}$ and $T_{3}$, both of which have a weight of $1 / 6$ and an execution time of 1 ; and $T_{4}$, which has, $\mathrm{e}\left(T_{4}^{1}\right)=2, \mathrm{e}\left(T_{4}^{2}\right)=1$, and an initial weight of $2 / 3$ that decreases to $1 / 5$ at time 3 . Notice that since $\mathrm{e}\left(T_{4}^{1}\right)=2$ and $\operatorname{wt}\left(T_{4}, \mathrm{r}\left(T_{4}^{1}\right)\right)=2 / 3$, both $\mathrm{d}\left(T_{4}^{1}\right)=3$ and $\mathrm{r}\left(T_{4}^{2}\right)=3$. Also, since $\mathrm{e}\left(T_{4}^{2}\right)=1$ and $\operatorname{wt}\left(T_{4}, \mathrm{r}\left(T_{4}^{2}\right)\right)=1 / 5, \mathrm{~d}\left(T_{4}^{2}\right)=8$.

A task $T_{i}$ changes weight or reweights at time $t$ if $\operatorname{wt}\left(T_{i}, t-\epsilon\right) \neq \mathrm{wt}\left(T_{i}, t\right)$ where $\epsilon \rightarrow 0^{+}$. For example, in the system depicted in Figure 3.1, $T_{4}$ reweights at time 3.

We now explain some of the issues involved in processing such a reweighting event. If a task $T_{i}$ changes weight at a time $t_{c}$ between the release and the deadline of some job $T_{i}^{j}$, then the following two actions may occur:

- The execution time of $T_{i}^{j}$ may be reduced to the amount of time for which $T_{i}^{j}$ has executed prior to $t_{c}$, and the execution time of $T_{i}^{j+1}$ may be redefined to be the amount of time "lost" by reducing the execution time of $T_{i}^{j}$.
- $\mathrm{r}\left(T_{i}^{j+1}\right)$ may be redefined to be less than $\mathrm{r}\left(T_{i}^{j}\right)+\mathrm{e}\left(T_{i}^{j}\right) / \mathrm{wt}\left(T_{i}, \mathrm{r}\left(T_{i}^{j}\right)\right)$. In this case, since $\mathrm{d}\left(T_{i}^{j}\right)=\mathrm{r}\left(T_{i}^{j}\right)+\mathrm{e}\left(T_{i}^{j}\right) / \mathrm{wt}\left(T_{i}, \mathrm{r}\left(T_{i}^{j}\right)\right)$, jobs $T_{i}^{j}$ and $T_{i}^{j+1}$ will "overlap." (For a standard sporadic task, every job's deadline is at or before its successor's release.)

The reweighting rules we present in Section 3.4 state under what conditions the above actions occur and by how much before $\mathrm{r}\left(T_{i}^{j}\right)+\mathrm{e}\left(T_{i}^{j}\right) / \mathrm{wt}\left(T_{i}, \mathrm{r}\left(T_{i}^{j}\right)\right)$ the job $T_{i}^{j+1}$ can be released. When the execution time of a job is reduced at a time $t$, then we say that it is "halted." Specifically, if a job $T_{i}^{j}$ is halted at time $t$, then $\operatorname{Ae}\left(T_{i}^{j}\right)$ is set to $\mathrm{A}\left(\mathcal{S}, T_{i}^{j}, 0, t\right)$.

Initiate and enact. As mentioned in Section 1.3.1, when a task reweights, there can be a difference between when it "initiates" the change and when the change is "enacted." The time at which the change is initiated is a user-defined time; the time at which the change is enacted is dictated by a set of conditions described in Section 3.4. We use the scheduling weight of a task $T_{i}$ at time $t$, denoted $\operatorname{Swt}\left(T_{i}, t\right)$, to represent the "last enacted weight of $T_{i}$." Formally, $\operatorname{Swt}\left(T_{i}, t\right)$ equals $\operatorname{wt}\left(T_{i}, u\right)$, where $u$ is the last time at or before $t$ that a weight change was enacted for $T_{i}$ (assuming an initial weight change occurred when $T_{i}$ joined the system). It is important to note that for adaptable sporadic tasks, we compute task deadlines and releases using scheduling weights. Hence, we have the following formulas:

$$
\begin{aligned}
\mathrm{r}\left(T_{i}^{1}\right) & =\theta\left(T_{i}^{1}\right) \\
\mathrm{d}\left(T_{i}^{j}\right) & =\mathrm{r}\left(T_{i}^{j}\right)+\mathrm{e}\left(T_{i}^{j}\right) / \operatorname{Swt}\left(T_{i}, \mathrm{r}\left(T_{i}^{j}\right)\right) \\
\mathrm{r}\left(T_{i}^{j+1}\right) & =\mathrm{d}\left(T_{i}^{j}\right)+\theta\left(T_{i}^{j+1}\right),
\end{aligned}
$$

where $\theta\left(T_{i}^{j}\right) \geq 0$. The third equation only applies in the absence of reweighting events, which may cause release times to be redefined.

Because the reweighting rules may cause $\mathrm{r}\left(T_{i}^{j+1}\right)<\mathrm{d}\left(T_{i}^{j}\right)$, we must slightly modify the definition of "window," "active," and "inactive" presented in Section 1.2.

Definition 3.1 (Window, Active, and Inactive). If $T_{i}^{j}$ is a job in the adaptable sporadic task system, $T$, then the window of $T_{i}^{j}$ defined as the range $\left[\mathrm{r}\left(T_{i}^{j}\right), \min \left(\mathrm{d}\left(T_{i}^{j}\right), \mathrm{r}\left(T_{i}^{j+1}\right)\right)\right)$. Furthermore, job $T_{i}^{j}$ is active at time $t \mathrm{iff} t$ is in $T_{i}^{j}$,s window (i.e., $t \in\left[\mathrm{r}\left(T_{i}^{j}\right), \min \left(\mathrm{d}\left(T_{i}^{j}\right), \mathrm{r}\left(T_{i}^{j+1}\right)\right)\right)$ ), and inactive otherwise.

### 3.2 The SW Scheduling Algorithm and Deviance

The scheduling weight (SW) scheduling algorithm is a theoretical scheduling algorithm that is used to determine which reweighting rule to apply as well to prove tardiness bounds. Under the SW scheduling algorithm, at each instant $t$, each active job $T_{i}^{j}$ in $\tau$ is allocated a share equal to its scheduling weight $\operatorname{Swt}\left(T_{i}, t\right)$. Hence, if a job $T_{i}^{j}$ is active over the range $\left[t_{1}, t_{2}\right)$, then over this range, $T_{i}^{j}$ is allocated $\int_{t_{1}}^{t_{2}} \operatorname{Swt}\left(T_{i}, u\right) d u$ time. (If a job is inactive, then it receives no allocations in SW.) Throughout this dissertation we use $\mathcal{S} W$ to denote the SW schedule of a task system $\tau$.

Example (Figure 3.2). Consider the example in Figure 3.2, which depict a one-processor system with four tasks: $T_{1}$, which has $\mathrm{e}\left(T_{1}\right)=1$ and $\operatorname{wt}\left(T_{1}\right)=1 / 3 ; T_{2}$, which has $\mathrm{e}\left(T_{2}\right)=1$ and $\mathrm{wt}\left(T_{2}\right)=1 / 6 ; T_{3}$, which has $\mathrm{e}\left(T_{3}\right)=2$ and $\mathrm{wt}\left(T_{3}\right)=1 / 4$ and leaves at time 8 ; and $T_{4}$, which has $\mathrm{e}\left(T_{4}\right)=4$ and an initial weight of $1 / 4$ and initiates and enacts a weight increase to $1 / 2$ at time 8. (As we will discuss in Section 3.4, the reweighting rules stop $T_{4}^{1}$ from receiving more than one unit of allocation and cause $T_{4}^{2}$ to be released with the remaining three units of execution at time 8.) Inset (a) depicts the GEDF schedule. Inset (b) depicts the SW schedule. Notice that, $T_{4}$ initiates and enacts a weight increase from $1 / 4$ to $1 / 2$ at time 8. Inset (c) depicts the allocations to $T_{4}$ in the GEDF and SW. Hence, before time 8, in the SW schedule, $T_{4}$ receives $1 / 4$ of the processor at each instant, and after time $8, T_{4}$ receives $1 / 2$ of the processor at each instant.

Example (Figure 3.3). Consider the example in Figure 3.3, which depicts a one-processor system with four tasks (where the execution time of each job is one): $T_{1}$, which joins the system at time 2 with $\mathrm{wt}\left(T_{1}\right)=1 / 2 ; T_{2}$ and $T_{3}$, both of which have a weight of $1 / 6$; and $T_{4}$, which has an initial weight of $1 / 2$ and initiates a weight decrease to $1 / 6$ at time 1 that is enacted at time 2. (As we will discuss in Section 3.4, the reweighting rules do not allow $T_{2}$ to decrease its weight immediately.) Inset (a) depicts the GEDF schedule. Inset (b) depicts the SW schedule. Inset (c) depicts the allocations to $T_{4}$ by the GEDF and SW scheduling algorithms. Notice that $T_{4}$ initiates a weight decrease at time 1 from $1 / 2$ to $1 / 6$ that is enacted at time 2 . As a result, in $\mathcal{S} W$ over the range $[1,2), T_{4}$ receives $1 / 2$ of the processor


Figure 3.2: A one-processor example of a task that increases its weight. (a) The GEDF schedule. (b) The SW schedule. (c) The allocations to $T_{4}$ in the GEDF and schedules.

(a)

(b)

$$
-\cdots={ }^{-\cdots} \text { SW allocation for } \mathrm{T}_{4} \quad \text { Actual allocation for } \mathrm{T}_{4}
$$


(c)

Figure 3.3: A one-processor example of a task that decreases its weight. (a) The GEDF schedule. (b) The SW schedule. (c) The allocations to $T_{4}$ in the GEDF and SW schedules.
at each instant, even though over $t \in[1,2), \operatorname{wt}\left(T_{4}, t\right)=1 / 6$.

The deviance of job $T_{i}^{j}$ of task $T_{i}$ at time $t$ is defined as $\operatorname{dev}\left(T_{i}^{j}, t\right)=\mathrm{A}\left(\mathcal{S} W, T_{i}^{j}, 0, t\right)-$ $\mathrm{A}\left(\mathcal{S}, T_{i}^{j}, 0, t\right)$, where $\mathcal{S}$ is the actual schedule. The deviance of a job $T_{i}^{j}$ represents the difference between $T_{i}^{j}$, s actual and SW allocations up to time $t$. If the deviance is negative, then the job has received a greater allocation in the actual schedule, and if the deviance is positive, then the job has received a greater allocation in the SW schedule. For example, in the system depicted in Figure 3.2, at time $8, T_{4}^{1}$ 's deviance is positive since it has been allocated more capacity in the SW schedule than in the actual schedule. On the other hand, in the system depicted in Figure 3.3, $T_{4}^{1}$ has negative deviance at time 1 because it has been allocated more capacity in the actual schedule than in the SW schedule. Whether the deviance of a task is positive or negative will determine which reweighting rule can be applied.

### 3.3 Modifications

In the adaptable sporadic task model, as presented in this chapter, the desired and guaranteed weight of each task is the same. Additionally, the fundamental unit for scheduling is a job. Both of these assumptions do not hold when scheduling tasks under PEDF or Pfair-based algorithms. Specifically, as we will discuss in Chapter 4, under our adaptive PEDF scheduling algorithm, a task's guaranteed and desired weight may differ. Also, as we will discuss in Chapter 5, under Pfair-based scheduling algorithms, the fundamental unit of scheduling is a subtask, not a job. As a result, the adaptable sporadic task model and theoretical scheduling algorithms presented here will be slightly modified in subsequent chapters to accommodate these differences.

### 3.4 Task Reweighting

Having defined the adaptable sporadic task model and its associated theoretical scheduling algorithms, we now define the reweighting rules for the GEDF and NP-GEDF scheduling algorithms and prove their tardiness and drift bounds.

We begin our discussion of reweighting under GEDF and NP-GEDF by defining the reweighting rules for GEDF. Then, in Section 3.4.2, we explain how these rules can be modified for NP-GEDF. Before continuing, we introduce one assumption that we make for simplicity: we assume that the actual execution time for any job is equal to its specified execution time, unless a task reweights when it has an active job. Then and only then can the actual execution time of a job be less than its execution time. ${ }^{1}$ In this scenario, the actual execution time of the job is determined by the rules we present shortly.

### 3.4.1 Reweighting Under GEDF

Let $\tau$ be a task system in which some task $T_{i}$ initiates a weight change to weight Nw at time $t_{c}$. Let Ow be the last scheduling weight of $T_{i}$ before the change is initiated at $t_{c}$. Let $\mathcal{S}$ be the $m$-processor GEDF schedule of $\tau$. Let $T_{i}^{j}$ be last-released job of $T_{i}$ before $t_{c}$. If $T_{i}^{j}$ does not exist or $T_{i}^{j}$ is inactive at $t_{c}$ before the reweighting event is initiated (i.e., $t_{c} \geq \mathrm{d}\left(T_{i}^{j}\right)$ ), then the weight change is immediately enacted, and future jobs of $T_{i}$ are released with the new weight. In the following rules, we consider the remaining possibility, i.e., $T_{i}^{j}$ exists and is active at $t_{c}$. (Notice that if $t_{c}=\mathrm{d}\left(T_{i}^{k}\right)=\mathrm{r}\left(T_{i}^{k+1}\right)$, then $T_{i}^{k}$ is the last-released job of $T_{i}$ before $t_{c}$, and it is not active at $t_{c}$. Therefore, the change is immediately enacted and $T_{i}^{k+1}$ is released with the new weight.)

Let $\operatorname{REM}\left(T_{i}^{j}, t_{c}\right)=\mathrm{e}\left(T_{i}^{j}\right)-\mathrm{A}\left(\mathcal{S}, T_{i}^{j}, 0, t_{c}\right)$. Note that $\operatorname{REM}\left(T_{i}^{j}, t_{c}\right)$ denotes the actual remaining computation in $T_{i}$ 's current job. Let nextE $\left(T_{i}^{j}, t_{c}\right)$ equal $\operatorname{REM}\left(T_{i}^{j}, t_{c}\right)$, if $\operatorname{REM}\left(T_{i}^{j}, t_{c}\right)>$ 0 ; otherwise, if $\operatorname{REM}\left(T_{i}^{j}, t_{c}\right)=0$, then let $\operatorname{nextE}\left(T_{i}^{j}, t_{c}\right)$ equal the value of $\mathrm{e}\left(T_{i}^{j+1}\right)$ had the weight-change event not occurred. Since $\operatorname{nextE}\left(T_{i}^{j}, t_{c}\right)$ is only used to determine the execution time of the next job released, it can be calculated at time $\mathrm{r}\left(T_{i}^{j+1}\right)$. (Notice that, if $T_{i}$ has no next job $T_{i}^{j+1}$ to release at the time specified in the rules below, then $\operatorname{next} \mathrm{E}\left(T_{i}^{j}, t_{c}\right)=0$. In this case, the rules are applied as stated, except that $T_{i}^{j+1}$ is not released.)

As was mentioned earlier, the choice of which rule to apply depends on whether deviance is positive or negative. If positive, then we say that $T_{i}$ is positive-changeable at time $t_{c}$

[^6]from weight Ow to Nw ; otherwise $T_{i}$ is negative-changeable at time $t_{c}$ from weight Ow to Nw . Because $T_{i}$ initiates its weight change at time $t_{c}, \mathrm{wt}\left(T_{i}, t_{c}\right)=\mathrm{Nw}$ holds; however, $T_{i}$ 's scheduling weight does not change until the weight change has been enacted, as specified in the rules below. Note that, if $t_{c}$ occurs between the initiation and enaction of a previous reweighting event of $T_{i}$, then the previous event is canceled, i.e., treated as if it had not occurred. As discussed later, any "error" associated with canceling a reweighting event like this is accounted for when determining drift (formally defined in Section 3.5.3).

Rule P: If $T_{i}$ is positive-changeable at time $t_{c}$ from weight Ow to Nw , then one of two actions is taken: (i) if $\mathrm{d}\left(T_{i}^{j}\right)-t_{c}>\operatorname{REM}\left(T_{i}^{j}, t_{c}\right) / \mathrm{Nw}$, then immediately, $T_{i}^{j}$ is halted, the weight change is enacted, a new job with an execution time of $\operatorname{next} \mathrm{E}\left(T_{i}^{j}, t_{c}\right)$ is released (if $\left.\operatorname{nextE}\left(T_{i}^{j}, t_{c}\right)>0\right)$, and $T_{i}^{j}$ becomes inactive; (ii) otherwise, at time $\mathrm{d}\left(T_{i}^{j}\right)$, the weight change is enacted, i.e., the scheduling weight of $T_{i}$ does not change until the end of its current job.

Rule N: If $T_{i}$ is negative-changeable at time $t_{c}$ from weight Ow to Nw , then one of two actions is taken: (i) if $\mathrm{Nw}>\mathrm{Ow}$, then immediately, $T_{i}^{j}$ is halted and its weight change is enacted, and at time $t_{r}$, a new job with an execution time of nextE $\left(T_{i}^{j}, t_{c}\right)$ is released (if nextE $\left(T_{i}^{j}, t_{c}\right)>0$ ) and $T_{i}^{j}$ becomes inactive, where $t_{r}$ is the smallest time at or after $t_{c}$ such that $\operatorname{dev}\left(T_{i}^{j}, t_{r}\right)=0$ holds; (ii) otherwise, at time $t_{e}$, the weight change is enacted, a new job with an execution time of $\operatorname{next} \mathrm{E}\left(T_{i}^{j}, t_{c}\right)$ is released (if $\operatorname{next} \mathrm{E}\left(T_{i}^{j}, t_{c}\right)>0$ ), and $T_{i}^{j}$ becomes inactive, where $t_{e}=\min \left(t_{r}, \mathrm{~d}\left(T_{i}^{j}\right)\right)$, and $t_{r}$ is smallest time at or after $t_{c}$ such that $\operatorname{dev}\left(T_{i}^{j}, t_{r}\right)=0$ holds.

Intuitively, Rule P changes a task's weight by halting its current job and issuing a new job with an execution time of $\operatorname{nextE}\left(T_{i}^{j}, t_{c}\right)$ with the new weight if doing so would improve its deadline.

Example (Figure 3.4). Consider the example in Figure 3.4, which depicts a one-processor system with four tasks (where the execution time of each job is one): $T_{1}$, which has $\mathrm{wt}\left(T_{1}\right)=$ $1 / 2$ and leaves at time at time $2 ; T_{2}$ and $T_{3}$, both of which have a weight of $1 / 6$; and $T_{4}$, which has an initial weight of $1 / 6$ that increases to $4 / 6$ at time 2 . In this system, $T_{4}$ initially has the


Figure 3.4: A one-processor example of reweighting via Case (i) of Rule P under GEDF. (a) The GEDF schedule. (b) $T_{4}$ 's allocations in the IDEAL, CSW, and SW schedules.


Figure 3.5: A one-processor example of reweighting via Case (ii) of Rule P under GEDF. (a) The GEDF schedule. (b) $T_{3}$ 's allocations in the IDEAL, CSW, and SW schedules.
lowest scheduling priority (there is a deadline tie). Inset (a) depicts the GEDF schedule. Inset (b) depicts $T_{4}$ 's allocations in the SW schedule and also in the other schedules, the IDEAL and CSW schedules, which are formally defined later in Section 3.5.2. Since $T_{4}$ is not scheduled by time 2 and because $\mathrm{d}\left(T_{4}^{1}\right)-t_{c}>\operatorname{REM}\left(T_{4}^{1}, t_{c}\right) / \mathrm{Nw}$, i.e., $\mathrm{d}\left(T_{4}^{1}\right)-2>0 /(4 / 6)$, it has positive deviance and changes its weight via Case (i) of Rule P. This, in turn, causes $T_{4}^{1}$ to be halted, $T_{4}^{2}$ to be released at time 2 with a deadline of $7 / 2$, and $T_{4}$ 's drift to become $2 / 6$. Note that halting $T_{4}$ 's current job and issuing a new job with an execution time of one improves $T_{4}$ 's scheduling priority, i.e., $\mathrm{d}\left(T_{4}^{1}\right)=6>\frac{7}{2}=\mathrm{d}\left(T_{4}^{2}\right)$.

Example (Figure 3.5). Consider the example in Figure 3.5, which depicts a one-processor system with three tasks (where the execution time of each job is one): $T_{1}$, which has wt $\left(T_{1}\right)=$ $1 / 3 ; T_{2}$, which has $w t\left(T_{2}\right)=1 / 4$; and $T_{3}$, which has an initial weight of $1 / 4$ that increases to $1 / 3$ at time 2. Inset (a) depicts the GEDF schedule. Inset (b) depicts $T_{3}$ 's allocations in the IDEAL, CSW, and SW schedules. (Again, the IDEAL and CSW schedules are formally defined later in Section 3.5.2.) Since $T_{3}^{1}$ has not been scheduled by time 2 its deviance is positive; furthermore, since $\mathrm{d}\left(T_{3}^{1}\right)-2<\operatorname{REM}\left(T_{3}^{1}, 2\right) /(1 / 3), T_{1}$ enacts its weight change via Case (ii) of Rule P. Notice that if $T_{3}^{1}$ had been halted at time 2 and released a new job of weight $1 / 3$, the deadline of this new job would equal time 5 (since $5=2+1 /(1 / 3)$ ). Thus, if we were to enact the change via Case (i) of Rule P , then we would increase the deadline of the first scheduled job of $T_{3}$, even though the weight of the task increased (i.e., such a change would decrease the scheduling priority of $T_{3}$ ). Therefore, we enact the weight change via Case (ii) of Rule P , which delays enacting the weight change until the deadline of $T_{3}^{1}$.

Rule N changes the weight of a task by one of two approaches: (i) if a task increases its weight, then Rule N causes the release time of its next job to be adjusted so that it is commensurate with the new weight; (ii) if a task decreases its weight, then Rule N causes the next job to be issued with a deadline that is commensurate with the new weight at the end of the current job.

Example (Figure 3.6). Consider the example in Figure 3.6, which depicts a one-processor system with four tasks (where the execution time of each job is one): $T_{1}$, which has wt $\left(T_{1}\right)=$ $1 / 2$ and leaves at time at time $2 ; T_{2}$ and $T_{3}$, both of which have a weight of $1 / 6 ;$ and $T_{4}$,


Figure 3.6: A one-processor example of reweighting via Case (i) of Rule N under GEDF. (a) The GEDF schedule. (b) $T_{4}$ 's allocations in the IDEAL, CSW, and SW schedules.


Figure 3.7: A one-processor example of reweighting via Case (ii) of Rule N under GEDF. (a) The GEDF schedule. (b) $T_{4}$ 's allocations in the IDEAL, CSW, and SW schedules.
which has an initial weight of $1 / 6$ that increases to $4 / 6$ at time 2 . This is the same system as in Figure 3.4 except that $T_{4}$ has a higher priority than both $T_{2}$ and $T_{3}$. Inset (a) depicts the GEDF schedule. Inset (b) depicts $T_{4}$ 's allocations in the IDEAL, CSW, and SW schedules. (Again, the IDEAL and CSW schedules are formally defined later in Section 3.5.2.) Since $T_{4}$ has been scheduled by time 2 , it has negative deviance and thus because it increases its weight, the change is enacted via Case (i) of Rule N. Thus, its next job is released time of 3 , which is such that $\operatorname{dev}\left(T_{4}, 3\right)=\int_{0}^{3} \operatorname{Swt}\left(T_{i}, u\right) d u-\mathrm{A}\left(\mathcal{S}, T_{4}, 0,3\right)=1-1=0$. By releasing the next job of $T_{4}$ at time 3 , the drift incurred is zero.

Example (Figure 3.7). Consider the example in Figure 3.7, which depicts a one-processor system with four tasks (where the execution time of each job is one): $T_{1}$, which joins the system at time 2 and has $\operatorname{wt}\left(T_{1}\right)=1 / 2 ; T_{2}$ and $T_{3}$, both of which have a weight of $1 / 6$; and $T_{4}$, which has an initial weight of $1 / 2$ that initiates a weight decrease to $1 / 6$ at time 1 that is enacted at time 2. Inset (a) depicts the GEDF schedule. Inset (b) depicts $T_{4}$ 's allocations in the IDEAL, CSW, and SW schedules. (Again, the IDEAL and CSW schedules are formally defined later in Section 3.5.2.) Since $T_{4}$ has negative deviance at time 1 and it decreases its weight, this weight change is enacted via Case (ii) of Rule N , causing $T_{4}$ 's next job to have a deadline of 8 and $T_{4}$ to have a drift of $-1 / 3$.

Notice that if $T_{i}$ initiates a weight change at time $t_{c}$ while some job $T_{i}^{k}$ of $T_{i}$ (not necessarily its last-released job) has missed its deadline, then the Rules P and N specify that one of two actions is taken. If no job of $T_{i}$ is active at $t_{c}$, then the weight change is enacted immediately. If there is a job $T_{i}^{j}$ that is active at $t_{c}$, then since $T_{i}^{j}$ has not been scheduled (because the earlier job $T_{i}^{k}$ has missed its deadline), it follows that $T_{i}$ is positive-changeable, and thus the weight change is enacted via Rule P (which may cause $T_{i}^{j}$ but not $T_{i}^{k}$ to halt). Notice that $T_{i}^{k}$ is unaffected in both cases.

It is important to remember that when the Rules P and N halt a job, they do not abandon the computation that the job was performing. Rather, these rules split that computation across two jobs. Since these rules change the ordering of a task in the priority queues that determine scheduling, the time complexity for reweighting one task is $O(\log N)$, where $N$ is the number of tasks in the system (assuming priority queues are implement using binomial


Figure 3.8: A one-processor example of canceling a reweighting event.
heaps).

Canceled reweighting events. We now introduce a property about the relationship between the initiation and enactment of a reweighting event in the case that some such events are canceled due to later reweighting events. Notice that, once a task $T_{i}$ initiates a weight change at $t_{c}$, this weight change is eventually either canceled by another weight change or enacted. Further, Rules P and N enact any non-canceled reweighting event no later than the deadline of the last-released job $T_{i}^{j}$ of $T_{i}$ at $t_{c}$ (if it exists and if $t_{c} \leq \mathrm{d}\left(T_{i}^{j}\right)$ ).

Example (Figure 3.8). Consider the example in Figure 3.8, which depicts a one-processor system with three tasks: $T_{1}$ and $T_{2}$, both of which have an execution time of 2 and a weight of $1 / 3$; and $T_{3}$, which has $\mathrm{e}\left(T_{3}\right)=2$ and an initial weight of $1 / 3$ that changes to $1 / 10$ at time 3 via Case (ii) of Rule N and then to $1 / 4$ at time 5 via Case (ii) of Rule N. Notice that, because the change initiated at time 3 is via Case (ii) of Rule N , the change is not enacted until time 6 . As a result, when a change is initiated at time 5 , this new change cancels the previous change. Even though the change initiated at time 3 is canceled, the time of the next weight enactment is still at time 6 .

From Figure 3.8, we can see that, once a reweighting event has been initiated during an active job, some weight change will be enacted by the earlier of the deadline of that job


Figure 3.9: A one-processor example of NP-GEDF. (a) $T_{3}$ has a lower scheduling-priority than $T_{2}$. (a) $T_{2}$ has a lower scheduling-priority than $T_{3}$.
or when the job becomes inactive (which may be earlier, by Rules P and N). Property (X) formalizes this idea.
(X) If a task $T_{i}$ initiates a weight change at time $t_{c}$ and the job $T_{i}^{j}$ is active at $t_{c}$, then some weight change is enacted according to Rule P or N by either $\mathrm{d}\left(T_{i}^{j}\right)$ or when $T_{i}^{j}$ becomes inactive, whichever is first.

### 3.4.2 Modifications for NP-GEDF

In order to adapt Rules P and N to work for NP -GEDF, the only modification we need to make is when these rules are initiated. If a task with an active job reweights before or after that job has been scheduled, then Rules P and N are initiated as before. (Note that after the active job $T_{i}^{j}$ has been released, if $T_{i}^{j}$ has not been scheduled, then $T_{i}$ is positive changeable, and if $T_{i}^{j}$ has been scheduled, then $T_{i}$ is negative changeable.) However, if a task changes its weight while the active job $T_{i}^{j}$ is executing, then the initiation of the weight change is delayed until $T_{i}^{j}$ has completed or $T_{i}^{j}$ is no longer active, whichever is first. Note that, if a task $T_{i}$ changes its weight from Ow to Nw at time $t_{c}$ in NP-GEDF, then $\mathrm{wt}\left(T_{i}, t_{c}\right)=\mathrm{Nw}$ holds, regardless of whether the initiation of Rule P or N must be delayed.

Example (Figure 3.9). Consider the example in Figure 3.9, which depicts the NP-GEDF schedule of a one-processor system with three tasks: $T_{1}$, which has $\mathrm{e}\left(T_{1}\right)=1$ and $\operatorname{wt}\left(T_{1}\right)=1 / 2$
that leaves at time $2 ; T_{2}$, which has $\mathrm{e}\left(T_{2}\right)=1$ and $\operatorname{wt}\left(T_{2}\right)=1 / 6$; and $T_{3}$, which has $\mathrm{e}\left(T_{3}\right)=2$ and an initial weight of $1 / 3$ that increases to $4 / 6$ at time 2 . In inset (a), $T_{3}$ has the lowest scheduling priority (there is a deadline tie). Since $T_{3}$ is not scheduled by time 2 , it has positive deviance and changes its weight via Rule P , causing $T_{3}^{1}$ to be halted, $T_{3}^{2}$ to be released at time 2 with a deadline of 5 . Inset (b) depicts the same scenario as in (a) except that $T_{3}$ has higher priority than $T_{2}$. Since $T_{3}$ is scheduled at time 2 , and the system is schedule by NP-GEDF, the initiation of the reweighting event is delayed until $T_{3}$ stops executing at time 3 . Since $T_{3}^{1}$ is complete by time 2, it has negative deviance and changes its weight via Rule N , causing its next job to have a release time of $9 / 2$.

### 3.5 Tardiness and Drift Bounds

In this section, we formally present and prove tardiness and drift bounds for the GEDF and NP-GEDF reweighting algorithms.

### 3.5.1 Tardiness Bounds

Instead of deriving tardiness bounds for GEDF or NP-GEDF when scheduling adaptable sporadic tasks from scratch (which would be quite tedious), we instead leverage the results reported by Devi and Anderson in (Devi and Anderson, 2008) concerning tardiness bounds that can be guaranteed under GEDF and NP-GEDF when scheduling sporadic tasks. In addition to deriving tardiness bounds under GEDF and NP-GEDF for sporadic task systems, Devi and Anderson also proposed an extension to the sporadic task model, referred to as the extended sporadic task model, and determined tardiness bounds that can be guaranteed to task systems that conform to the extended sporadic task model. We will show that any adaptable sporadic task system can be modeled as an extended sporadic task system; hence, tardiness bounds derived for extended sporadic task systems can be applied to adaptable sporadic task systems as well. We begin by describing the extended sporadic task model.

The extended sporadic task model. In the conventional sporadic task model, the number of tasks in a task system is fixed, and the sum of the weights of all its tasks is assumed
to be at most $m$ (the number of processors). On the other hand, in the extended model, the number of tasks associated with a task system is allowed to vary and the total weight of all tasks is allowed to exceed $m$. (The number of tasks could potentially be infinite.) Further, each task is assigned a static weight, and all jobs (except possibly the final job) of a task have equal execution times. However, to prevent overload, at any given time, only a subset of tasks whose total weight is at most $m$ is allowed to be effective, ${ }^{2}$ i.e., is allowed to release jobs. Additionally, the final job of a task can stop ${ }^{3}$ at some time $t_{s}$ before its deadline, provided the allocation that the job receives in the actual schedule is at most the allocation it receives up to $t_{s}$ in the SW schedule, ${ }^{4}$ i.e., the last job has non-negative deviance, and the job is not executing in a non-preemptive segment at $t_{s}$. When a job stops, its execution time is altered to equal the amount of time that the job actually executed for in the actual schedule up to time $t_{s}$. Thus, at any time $t$, each task $T_{i}$ can be in one of the following states.

- Effective, if the first job of $T_{i}$ is released at or before $t$, the deadline of its final job is after $t$, and its final job has not stopped at or before $t$. A task whose final job has its deadline at or before $t$ is not considered effective at $t$ even if the final job is pending at $t$.
- Ineffective, if the release time of the first job of $T_{i}$ is after $t$.
- Terminated, if the deadline of $T_{i}$ 's final job is before $t$.
- Stopped, if the final job has stopped but its deadline has not elapsed.

As can be easily seen, a task that is either ineffective or terminated at time $t$ cannot have active or effective jobs at $t$.

In order to provide tardiness bounds for extended sporadic task systems, Devi and Anderson proposed partitioning the set of all tasks associated with a task system into $N$ task classes such that the following hold: (i) effective intervals are disjoint for every two tasks in

[^7]| $\downarrow$ | Job release |
| :---: | :---: |
|  | Job deadline |
|  | Job release/deadline |
|  | Reweighting event enactment |
| $\square$ | Scheduled |
|  | Job layout without stop |
|  | Task is efective/ Subtask range |



Figure 3.10: A one-processor example of task classes. (a) An extended sporadic task system. The effective range for each task is denoted by a dashed rectangle with rounded corners. (b) The same system as in Figure 3.4. Subtasks are denoted by a dashed rectangle with rounded corners.
each class and (ii) tasks within a class are governed by precedence constraints, i.e., the first job of a task cannot begin execution until all jobs of all tasks with earlier effective intervals in its class have completed execution. The second requirement implies that tasks that are not bound by precedence constraints should belong to different classes even if their effective intervals are disjoint.

Example (Figure 3.10). Consider the example in Figure 3.10(a), which depicts five tasks, each with an execution time of one: $T_{1}$, which has $\mathrm{wt}\left(T_{1}\right)=1 / 2$ that leaves at time $2 ; T_{2}$ and $T_{3}$, each of which have a weight of $1 / 6 ; T_{4}$, which has $w t\left(T_{4}\right)=1 / 6$ and stops at time 2 ; and $T_{5}$, which has $\mathrm{wt}\left(T_{5}\right)=4 / 6$, is in the same task class as $T_{4}$, and becomes effective as soon as $T_{4}$ stops. $T_{4}^{1}$ has a lower scheduling-priority than $T_{2}^{1}$ and $T_{3}^{1}$. Inset (b) depicts the same system as in Figure 3.4 (with the tasks renumbered for clarity). Notice that $T_{4}$ can stop at time 2 because its deviance is zero. Also note that, since $T_{4}$ and $T_{5}$ are in the same task class, only one of them can be effective at the same time.

Let $T^{[\ell]}$ denote task class $\ell$, and let $\mathrm{e}_{\max }\left(T^{[\ell]}\right)$ and $\mathcal{W}\left(T^{[\ell]}\right)$ denote the maximum execution time and weight, respectively, of any task in $T^{[\ell]}$. In (Devi and Anderson, 2008), it is shown that the tardiness for any task of any task class $T^{[i]}$ of an extended sporadic task system $T$
under global EDF is at most

$$
\begin{equation*}
\frac{\sum_{T^{[\ell]} \in \mathcal{E}(T, m-1)} \mathrm{e}_{\max }\left(T^{[\ell]}\right)}{m-\sum_{T^{[\ell]} \in \mathcal{X}(T, m-2)} \mathcal{W}\left(T^{[\ell]}\right)}+\mathrm{e}_{\max }\left(T^{[i]}\right), \tag{3.1}
\end{equation*}
$$

and that under global NP-EDF is at most

$$
\begin{equation*}
\frac{\sum_{T^{[\ell]} \in \mathcal{E}(T, m)} \mathrm{e}_{\max }\left(T^{[\ell]}\right)}{m-\sum_{T^{[\ell]} \in \mathcal{X}(T, m-1)} \mathcal{W}\left(T^{[\ell]}\right)}+\mathrm{e}_{\max }\left(T^{[i]}\right), \tag{3.2}
\end{equation*}
$$

where $\mathcal{E}(T, k)$ and $\mathcal{X}(T, k)$ are subsets of $k$ task classes of $T$ with the highest execution times and weights, respectively, for any of their tasks (i.e., with the highest values for $\mathrm{e}_{\max }\left(T^{[\ell]}\right)$ and $\mathcal{W}\left(T^{[\ell]}\right)$, respectively $)$.

Extended and adaptable sporadic task systems. We now show how an adaptable sporadic task system can be modeled as an extended sporadic task system. We initially assume that no task changes its weight by Case (i) of Rule N, i.e., no negative-changeable task halts. Such weight changes are considered afterwards. We first show that each task of an adaptable sporadic task system can be modeled as a task class of an extended sporadic task system. For this, we decompose each adaptable sporadic task into disjoint "subtasks ${ }^{5}$," where a subtask $\operatorname{sub}\left(T_{i}, j\right)$ of a adaptable sporadic task $T_{i}$ is a "maximal" set of jobs with the following properties: (i) the jobs in $\operatorname{sub}\left(T_{i}, j\right)$ are consecutive jobs of $T_{i}$; (ii) each job is released between the same pair of two consecutive weight-change enactments for $T_{i}$; (iii) each job has the same execution time; and (iv) no new job can be added to $\operatorname{sub}\left(T_{i}, j\right)$ without violating one or more of properties (i), (ii), and (iii), and in that sense, $\operatorname{sub}\left(T_{i}, j\right)$ is maximal. As an example, consider Figure 3.10(b), which depicts the same system as in Figure 3.4 with the subtasks marked.

Recall that in an extended sporadic task system, if a job $T_{i}^{j}$ "stops" at time $t_{s}$, then $t_{s}<\mathrm{d}\left(T_{i}^{j}\right), T_{i}^{j}$,s deviance is non-negative, and when a job stops, its actual execution time is set to the value that the job had executed for in the actual schedule up to time $t_{s}$. Since we are assuming that only positive-changeable tasks may halt, which by definition have an

[^8]active job that has non-negative deviance, it is easy to see that for positive-changeable tasks "stopping" in an extended sporadic task system has the same effect as "halting" an adaptable sporadic task system. For example, the impact on the system when $T_{4}^{1}$ stops in Figure 3.10(a) is the same as when $T_{6}^{1}$ halts in Figure 3.10(b).

By definition, all jobs of a subtask have equal execution times, and because all such jobs are released between two consecutive weight-change enactments, each subtask has a static scheduling weight. Also, as explained above, halting is the same as stopping for a positivechangeable task. Hence, if no task changes its weight via Case (i) of Rule N, then it follows that each subtask in the adaptable sporadic task model corresponds to a task, with a static weight and execution time, in the extended sporadic task model, and each task in a adaptable sporadic task model that consists of subtasks corresponds to a task class, composed of tasks with different weights or execution times or both, of the extended sporadic task model. Also note that intervals within which subtasks of an adaptable sporadic task are effective are disjoint. For example, consider insets (a) and (b) of Figure 3.10, which despite the notation change, have the identical schedules.

We now explain that an adaptable sporadic task can be modeled as an extended sporadic task even if tasks change their weight via Case (i) of Rule N. Before we continue, notice that the one difference between positive and negative-changeable tasks with respect to halting at time $t_{h}$ is as follows: if $T_{i}$ is positive-changeable at $t_{h}$ and its job $T_{i}^{j}$ is active at that time, then $T_{i}^{j+1}$, i.e., the next job of $T_{i}$, may be released at $t_{h}$, whereas if $T_{i}$ is negative-changeable, then $T_{i}^{j+1}$ may not be released until time $t_{r}>t_{h}$, where $t_{r}$ is the earliest time at which the allocations to $T_{i}^{j}$ are equal in the actual schedule and under SW , i.e., the next time $T_{i}^{j}$, s deviance is zero. Thus, if $T_{i}$ is negative-changeable and halts at $t_{h}$, then $T_{i}^{j}$ may be thought of as stopping at time $t_{r}$, where $t_{r}$ is as defined above. However, notice that by Case (i) of Rule N , if $T_{i}^{j}$ halts at time $t_{h}$, then $T_{i}$ enacts a weight increase at time $t_{h}$. Thus, the time $t_{r}$ is calculated using dynamic weights, which are not explicitly included in Devi and Anderson's extended sporadic model.

Example (Figure 3.11). Consider the example in Figure 3.11(a), which depicts five tasks, each with an execution time of one: $T_{1}$, which has $\operatorname{wt}\left(T_{1}\right)=1 / 2$ and leaves at time $2 ; T_{2}$


Figure 3.11: A one-processor example of task classes. (a) An extended sporadic task system. The effective range for each task is denoted by a dashed rectangle with rounded corners. (b) The same system as in Figure 3.6. Subtasks are denoted by a dashed rectangle with rounded corners.
and $T_{3}$, each of which has a weight of $1 / 6 ; T_{4}$, which has an initial weight of $1 / 6$, at time 2 increases its weight to $4 / 6$, and at time 3 stops; and $T_{5}$, which has $\operatorname{wt}\left(T_{5}\right)=4 / 6$, is in the same task class as $T_{4}$, and becomes effective as soon as $T_{4}$ stops. $T_{4}^{1}$ has a higher schedulingpriority than $T_{2}^{1}$ and $T_{3}^{1}$. Inset (b) depicts the same system as in Figure 3.6 (with the tasks renumbered for clarity). Since $T_{4}$ and $T_{5}$ are in the same task class, only one of them can be effective at the same time. Notice that $T_{4}$ can stop at time 3 because its actual allocation until then is no greater than its SW allocation; however, if we were using static weights than $T_{4}^{1}$ would not be able to stop until time 6 .

Even though dynamic weights are not explicitly included in the extended sporadic model, the bounds in (3.1) and (3.2) still hold if the only time a task is allowed to change its weight is when its final job has finished executing and the weight change is an increase, which is exactly the scenario that arises when a task changes its weight via Case (i) of Rule N. The reason why (3.1) and (3.2) still hold in the presence of such weight changes is because the extended sporadic task model only requires that the total allocation in the SW schedule to a stopping job $T_{i}^{j}$ at the time it stops be at least the allocation $T_{i}^{j}$ received in the actual schedule; increasing the rate at which the $T_{i}^{j}$ is allocated time in the SW schedule is not an issue as long as the total SW allocation to all tasks that are effective is at most $m$ at each
instant.

Informally, this holds by the following reasoning. Increasing a task $T_{i}$ 's allocation rate in the SW schedule would cause $T_{i}$ 's period to decrease. As a result, jobs of $T_{i}$ would have a higher relative scheduling priority after the change. However, since the final job of $T_{i}$ has already completed execution in the actual schedule when it halts, this increase in priority does not impact any job. In particular, it is not possible for another job to have a lower priority than the halting job $T_{i}^{j}$ before the weight change and a higher priority after the weight change. Therefore, scheduling $T_{i}^{j}$ (the halting job) in the past with a lower priority does not adversely impact how any other job was scheduled in the past. Since $T_{i}^{j}$ has already completed execution before its deadline, $T_{i}^{j}$ 's tardiness is not impacted either.

Hence, if the task $T_{i}$ enacts a weight change via Case (i) of Rule N at $t_{h}$ and the system is not over-utilized after the change, no other job will be impacted. Thus, since a weight change is enacted at $t_{h}$ regardless of whether $T_{i}$ is positive- or negative-changeable, and $T_{i}^{j+1}$ is released at or after $t_{h}$, in both the cases, $T_{i}^{j}$ and $T_{i}^{j+1}$ belong to different subtasks of $T_{i}$. Hence, the definition of a subtask is unaltered even in the presence of negative-changeable jobs by Case (i) of Rule N , and the correspondence described earlier between a adaptable sporadic task and an extended sporadic task holds.

Thus, the tardiness bounds specified in (3.1) and (3.2) and that can be guaranteed to extended sporadic task systems are also applicable to adaptable sporadic task systems if task class $T^{[\ell]}$ is replaced by adaptable sporadic task $T_{z}$, and $\mathrm{e}_{\max }\left(T_{z}\right)$ and $\mathcal{W}\left(T_{z}\right)$ are taken as the maximum execution time of any job of $T_{z}$ and the maximum weight assigned to $T_{z}$ at any time. (It should be noted that the tardiness bounds hold only if the sum of the weights of all tasks that are active at any instant is at most $m$.) Thus, we have the following theorem.

Theorem 3.1. Let $\tau$ be an adaptable sporadic task system, where for any $t \geq 0$, $\sum_{T_{i} \in \tau} \operatorname{Swt}\left(T_{i}, t\right) \leq m$. Then, for any task $T_{i}$, GEDF on $m$ processors ensures a tardiness of at most

$$
\frac{\sum_{T_{z} \in \mathcal{E}(m-1)} \mathrm{e}_{\max }\left(T_{z}\right)}{m-\sum_{T_{z} \in \mathcal{X}(m-2)} \mathrm{wt}_{\max }\left(T_{z}\right)}+\mathrm{e}_{\max }\left(T_{i}\right)
$$

and NP-GEDF on $m$ processors ensures a tardiness of at most

$$
\frac{\sum_{T_{z} \in \mathcal{E}(T, m)} \mathrm{e}_{\max }\left(T_{z}\right)}{m-\sum_{T_{z} \in \mathcal{X}(T, m-1)} \mathcal{W}\left(T_{z}\right)}+\mathrm{e}_{\max }\left(T_{i}\right) .
$$

### 3.5.2 Additional Theoretical Algorithms

"Drift bounds" (formally defined in Section 3.5.3) reflect a reweighting algorithm's accuracy at creating a job set that mimics an "ideal" task system, in which weight changes can always be initiated and enacted instantaneously. In order to define drift and prove drift bounds for the reweighting rules proposed in Section 3.4, we introduce two additional theoretical scheduling algorithms that are able to preempt and swap tasks at arbitrarily small intervals: the clairvoyant scheduling-weight (CSW) scheduling algorithm and the ideal (IDEAL) scheduling algorithm. The CSW scheduling algorithm allocates each task a fraction of the system equal to its scheduling weight, and will not allocate capacity to a task if its active job has received an allocation equal to its actual execution time. The IDEAL scheduling algorithm allocates each task a fraction of the system equal to its weight (i.e., not its scheduling weight) at each instant. Further, the IDEAL scheduling algorithm continually allocates capacity to a task as long as it has an active job.

We now discuss these two algorithms in more detail, by exploring their differences when scheduling the two example systems presented in Figures 3.12, 3.13, and 3.14.

Example (Figures 3.12 and 3.13). Consider the example in Figures 3.12 and 3.13, which depict a one-processor system with four tasks: $T_{1}$, which has $\mathrm{e}\left(T_{1}\right)=1$ and $\operatorname{wt}\left(T_{1}\right)=1 / 3 ; T_{2}$, which has $\mathrm{e}\left(T_{2}\right)=1$ and $\mathrm{wt}\left(T_{2}\right)=1 / 6 ; T_{3}$, which has $\mathrm{e}\left(T_{3}\right)=2$ and $\operatorname{wt}\left(T_{3}\right)=1 / 4$ and leaves at time 8 ; and $T_{4}$, which has $\mathrm{e}\left(T_{4}\right)=4$ and an initial weight of $1 / 4$ and initiates and enacts a weight increase to $1 / 2$ at time 8 (the same system as in Figure 3.2). Figure 3.12(a) depicts the GEDF schedule. Figure 3.12 (b) depicts the allocations to $T_{4}$ in the GEDF, IDEAL, SW, and CSW scheduling algorithms. Figure 3.13(a) depicts the SW schedule. Figure 3.13(b) depicts the CSW schedule. Figure 3.13(c) depicts the IDEAL schedule. Notice that $T_{4}^{1}$ receives no


Figure 3.12: A one-processor example of a task that increases its weight. (a) The GEDF schedule. (b) The allocations to $T_{4}$ in the GEDF, SW, CSW, and IDEAL schedules.


Figure 3.13: A continuation of Figure 3.12 that depicts (a) SW, (b) CSW, and (c) IDEAL schedules.
allocations in the CSW schedule once it has received one unit of execution (the amount $T_{4}^{1}$ is allocated in the GEDF schedule).

Example (Figure 3.14). Consider the example in Figure 3.14, which depicts a one-processor system with four tasks (where the execution time of each job is one): $T_{1}$, which joins the system at time 2 with a $\operatorname{wt}\left(T_{1}\right)=1 / 2 ; T_{2}$ and $T_{3}$, both of which have a weight of $1 / 6 ;$ and $T_{4}$, which has an initial weight of $1 / 2$ and initiates a weight decrease to $1 / 6$ at time 1 that is enacted at time 2 (the same system as in Figure 3.3). Inset (a) depicts the GEDF schedule. Inset (b) depicts the allocations to $T_{4}$ in the GEDF, IDEAL, SW, and CSW scheduling algorithms. Inset (c) depicts the SW schedule. Inset (d) depicts the CSW schedule. Inset (e) depicts the IDEAL schedule.

The CSW scheduling algorithm. CSW is a theoretical scheduling algorithm that is used as a reference for calculating drift. Under the CSW scheduling algorithm, at each instant $t$, each job of each task $T_{i}$ that is both active and incomplete (in the CSW schedule) is allocated a fraction of a processor equal to $\operatorname{Swt}\left(T_{i}, t\right)$. Furthermore, we consider CSW to be "clairvoyant" in the sense that CSW uses the actual execution time of $T_{i}^{j}$ to determine if $T_{i}^{j}$ has completed before it halts. More specifically, for any schedule $\mathcal{C} S W$ under CSW of any task system $\tau$, we say that $T_{i}^{j}$ has completed by time $t$ in $\mathcal{C} S W$ iff $T_{i}^{j}$ has executed for $\operatorname{Ae}\left(T_{i}^{j}\right)$ by $t$. Thus, the difference between SW and CSW is that a job in SW will not stop receiving allocations as long as its active, whereas a job in CSW will stop receiving allocations as soon as it receives its actual execution time. Throughout this dissertation we use $\mathcal{C} S W$ to denote the CSW schedule of a task system $\tau$.

Example (Figures 3.12 and 3.13). Consider the system depicted in Figures 3.12 and 3.13 (described above). In this system, the reweighting rules Rule P stops $T_{4}^{1}$ from scheduling its second unit of execution. As a result, $\operatorname{Ae}\left(T_{4}^{1}\right)=1$. So, as illustrated in Figure $3.13(\mathrm{~b}), T_{4}^{1}$ stops receiving allocation in the CSW schedule once it has been allocated one unit of execution (at time 4). $T_{4}$ resumes execution once its next job has been released at time 8. This differs from the SW schedule, where $T_{4}$ receives allocations over the range $[4,8)$. Notice that, in the example depicted in Figure $3.14, \mathrm{Ae}\left(T_{4}^{1}\right)=\mathrm{e}\left(T_{4}^{1}\right)$, so $T_{4}^{1}$ 's allocations are identical in both


Figure 3.14: A one-processor example of a task that decreases its weight. (a) The GEDF schedule. (b) The allocations to $T_{4}$ in the GEDF, IDEAL, SW, and CSW scheduling algorithms. (c) The SW schedule. (c) The CSW schedule. (d) The IDEAL schedule.
the $\mathcal{S} W$ and $\mathcal{C} S W$ schedules. Notice that, in both the CSW and actual schedule, the total time allocated to a job is the same. Thus, the CSW algorithm more accurately represents the behavior of a task than the SW algorithm.

The IDEAL scheduling algorithm. Under the ideal (IDEAL) scheduling algorithm, at each instant $t$, each task $T_{i}$ in $\tau$ with an active job at $t$ is allocated a fraction of the system equal to its weight (i.e., $\operatorname{wt}\left(T_{i}, t\right)$ ). Hence, if $\mathcal{I}$ is the IDEAL schedule of $\tau$ and $T_{i}$ is active over the interval $\left[t_{1}, t_{2}\right)$, then over $\left[t_{1}, t_{2}\right)$, the task $T_{i}$ is allocated $\mathrm{A}\left(\mathcal{I}, T_{i}^{j}, t_{1}, t_{2}\right)=\int_{t_{1}}^{t_{2}} \mathrm{wt}\left(T_{i}, u\right) d u$ time. As mentioned earlier, the IDEAL algorithm is similar to SW, with one major exception: each task receives an allocation equal to its weight, whereas under SW, each task receives an allocation equal to its scheduling weight. Throughout this dissertation we use $\mathcal{I}$ to denote the IDEAL schedule of a task system $\tau$.

Example (Figures 3.12-3.14). Notice that in the system depicted in Figure 3.14, the reweighting event initiated at time 1 by $T_{4}$ is not enacted until time 2 . As a result, over the range $[1,2)$, in the IDEAL schedule, $T_{4}$ receives $\mathrm{wt}\left(T_{4}, t\right)=1 / 6$ at each instant, whereas in the SW schedule, $T_{4}$ receives $\operatorname{Swt}\left(T_{4}, t\right)=1 / 2$ at each instant. On the other hand, in the systems depicted in Figures 3.12 and 3.13, $T_{4}$ 's reweighting event is enacted as soon as it is initiated. As a result, the IDEAL and SW schedules are the same.

### 3.5.3 Drift

For most real-time scheduling algorithms, the difference between the IDEAL and actual allocations a task receives lies within some bounded range centered at zero. For example, under a uniprocessor EDF schedule, the difference between the ideal and actual allocations for a task lies within $\left(-\mathrm{e}_{\max }\left(T_{i}\right), \mathrm{e}_{\max }\left(T_{i}\right)\right)$ (assuming the processor is not over-utilized). When a weight change occurs, the same bounds are maintained except that they may be centered at a different value. For example, in Figure 3.12, the range for $T_{4}$ is originally $(-4,4)$, but after the reweighting event, it is $(-3,5)$. This lost allocation is called drift. Given this loss (barring further reweighting events) $T_{i}$ 's drift will not change. In general, a task's drift per reweighting event will be non-negative if it increases its weight, and a task's drift per reweighting event
will be non-positive if it decreases its weight. The drift of a task $T_{i}$ at time $t$ is defined as

$$
\operatorname{drift}\left(T_{i}, t\right)=\mathrm{A}\left(\mathcal{I}, T_{i}, 0, t\right)-\mathrm{A}\left(\mathcal{C} S W, T_{i}, 0, t\right) .
$$

(Notice that drift is defined in terms of CSW instead of SW. This is because, as we discussed earlier, the CSW scheduling algorithm is a more accurate representation of the actual schedule than the SW scheduling algorithm.) The per-event absolute drift is the absolute value of the amount of drift that is incurred as a result of a reweighting event. For example, if the perevent absolute drift is $\mathrm{e}_{\max }\left(T_{i}\right)$, then after $n$ reweighting events, the maximal absolute drift is $n \cdot \mathrm{e}_{\max }\left(T_{i}\right)$.

Under GEDF, the drift per reweighting event is bounded as follows.
Theorem 3.2. The per-event absolute drift under GEDF for each task $T_{i}$ is at most $\mathrm{e}_{\max }\left(T_{i}\right)$.
Proof. We first show that for any job $T_{i}^{j}$ of any task $T_{i}, \mathrm{~A}\left(\mathcal{C} S W, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t_{A}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right)$, where $t_{A}$ is the time that $T_{i}^{j}$ becomes inactive. Notice that, since a job is inactive by its deadline (i.e., $t_{A} \leq \mathrm{d}\left(T_{i}^{j}\right)$ ), and the deadline of $T_{i}^{j}$ is defined as $\mathrm{r}\left(T_{i}^{j}\right)+\mathrm{e}\left(T_{i}^{j}\right) / \operatorname{Swt}\left(T_{i}, \mathrm{r}\left(T_{i}^{j}\right)\right)$, it follows that if $T_{i}$ does not enact a weight change over the range $\left[\mathrm{r}\left(T_{i}^{j}\right), t_{A}\right)$, i.e., $T_{i}$ 's scheduling weight is static over the range $\left[\mathrm{r}\left(T_{i}^{j}\right), t_{A}\right)$, then $\mathrm{A}\left(\mathcal{C} S W, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t_{A}\right) \leq \mathrm{e}\left(T_{i}^{j}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right)$.

Thus, in order for, $\mathrm{A}\left(\mathcal{C} S W, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t_{A}\right)>\mathrm{e}_{\max }\left(T_{i}\right)$ to hold, it must be that $T_{i}$ enacted a weight change over the range $\left[\mathrm{r}\left(T_{i}^{j}\right), t_{A}\right)$. Thus, we assume that $T_{i}$ enacts a change over the range $\left[\mathrm{r}\left(T_{i}^{j}\right), t_{A}\right)$, and that $t_{e}$ is the last such time. Notice that if the change enacted at $t_{e}$ is by Case (i) or (ii) of Rule P or Case (ii) of Rule N , then $T_{i}^{j}$ becomes inactive at $t_{e}$, which contradicts our assumption that $t_{e}<t_{A}$. Thus, the change enacted at $t_{e}$ must be by Case (i) of Rule N. By Case (i) of Rule $\mathrm{N}, T_{i}^{j}$ becomes inactive at the first time at or after $t_{e}$ such that the deviance of $T_{i}^{j}$ equals zero, i.e., $t_{A}$ is the smallest time such that $\int_{\mathrm{r}\left(T_{i}^{j}\right)}^{t_{A}} \operatorname{Swt}\left(T_{i}, u\right) d u=\operatorname{Ae}\left(T_{i}^{j}\right)$. Since $\mathrm{A}\left(\mathcal{C} S W, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t_{A}\right) \leq \int_{\mathrm{r}\left(T_{i}^{j}\right)}^{t_{A}} \operatorname{Swt}\left(T_{i}, u\right) d u$ and since $\operatorname{Ae}\left(T_{i}^{j}\right) \leq \mathrm{e}\left(T_{i}^{j}\right) \leq \mathrm{e}_{\max }\left(T_{i}^{j}\right)$, it follows that

$$
\begin{equation*}
\mathrm{A}\left(\mathcal{C} S W, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t_{A}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right) \tag{3.3}
\end{equation*}
$$

Let $t_{c}$ be a time such that some task $T_{i}$ initiates a weight change. Let $t_{e}$ denote the next
time that the change initiated at $t_{c}$ is either canceled or enacted, whichever is first. Thus, by the definition of canceled, no weight change is initiated over the range $\left(t_{c}, t_{e}\right)$. We show that regardless of which rule $T_{i}$ uses to change its weight, the drift incurred by this initiation is at $\operatorname{most} \mathrm{e}_{\max }\left(T_{i}\right)$. Let $T_{i}^{j}$ be the last-released job of $T_{i}$ before $t_{c}$. Notice that if $T_{i}^{j}$ is not active at $t_{c}$ or $T_{i}^{j}$ does not exist, then the change is immediately enacted and no job is halted. Since the only two potential sources of drift are delays in enacting a weight change and halting a job (which causes the actual execution time to be lower than the execution time), it follows that if $T_{i}^{j}$ does not exist or $T_{i}^{j}$ is inactive at $t_{c}$ then no drift is incurred. Thus, for the rest of this proof, we assume that $T_{i}^{j}$ is active at $t_{c}$, and we let $t_{A}$ denote the time that $T_{i}^{j}$ becomes inactive. By Property ( X ) and the fact that $t_{A} \leq \mathrm{d}\left(T_{i}^{j}\right)$ holds, we have

$$
\begin{equation*}
t_{c} \leq t_{e} \leq t_{A} \leq \mathrm{d}\left(T_{i}^{j}\right) \tag{3.4}
\end{equation*}
$$

If $T_{i}$ changes its weight at time $t_{c}$ via Case (i) of Rule P , then since this weight change is immediately enacted (i.e., $t_{c}=t_{e}$ ), it is as though allocation equal to $\mathrm{A}\left(\mathcal{I}, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t_{c}\right)$ $\mathrm{A}\left(\mathcal{C} S W, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t_{c}\right)$ is "lost." For example in Figure 3.4, the task $T_{4}$ "loses" an allocation of 2/6. Notice that, per reweighting event, $\mathrm{A}\left(\mathcal{I}, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t_{c}\right)-\mathrm{A}\left(\mathcal{C} S W, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t_{c}\right) \leq$ $\mathrm{e}_{\max }\left(T_{i}\right)$. Also note that since, by (3.3) and (3.4), $\mathrm{A}\left(\mathcal{C} S W, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t_{c}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right)$, it follows that $-\mathrm{e}_{\max }\left(T_{i}\right) \leq \mathrm{A}\left(\mathcal{I}, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t_{c}\right)-\mathrm{A}\left(\mathcal{C} S W, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t_{c}\right)$. Thus, since $-\mathrm{e}_{\max }\left(T_{i}\right) \leq$ $\mathrm{A}\left(\mathcal{I}, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t_{c}\right)-\mathrm{A}\left(\mathcal{C} S W, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t_{c}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right)$, it follows that the absolute drift is at $\operatorname{most} \mathrm{e}_{\max }\left(T_{i}\right)$.

Suppose that $T_{i}$ initiates a change to weight Nw via Case (ii) of Rule P at $t_{c}$. Since this change is enacted or canceled at time $t_{e}$, it is as though allocation equal to $\mathrm{A}\left(\mathcal{I}, T_{i}, t_{c}, t_{e}\right)-$ $\mathrm{A}\left(\mathcal{C} S W, T_{i}, t_{c}, t_{e}\right)$ is "lost." Recall that a task only changes its weight via Case (ii) of Rule P if the deviance of $T_{i}$ is positive and if $\mathrm{d}\left(T_{i}^{j}\right)-t_{c} \leq \operatorname{REM}\left(T_{i}^{j}, t_{c}\right) / \mathrm{Nw}$. Since by (3.4), $t_{e} \leq \mathrm{d}\left(T_{i}^{j}\right)$ and $\operatorname{REM}\left(T_{i}^{j}, t_{c}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right)$, the previous inequality can be rewritten as

$$
\begin{equation*}
\mathrm{Nw} \cdot\left(t_{e}-t_{c}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right) . \tag{3.5}
\end{equation*}
$$

Recall that no change is initiated over the range $\left(t_{c}, t_{e}\right)$. Thus, by definition, $\mathrm{A}\left(\mathcal{I}, T_{i}, t_{c}, t_{e}\right)=$
$\int_{t_{c}}^{t_{e}}(\mathrm{Nw}) d u=\mathrm{Nw} \cdot\left(t_{e}-t_{c}\right)$. Hence, by (3.5), $\mathrm{A}\left(\mathcal{I}, T_{i}, t_{c}, t_{e}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right)$. Furthermore, by (3.3) and (3.4), $\mathrm{A}\left(\mathcal{C} S W, T_{i}, t_{c}, t_{e}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right)$. Thus, $-\mathrm{e}_{\max }\left(T_{i}\right) \leq \mathrm{A}\left(\mathcal{I}, T_{i}, t_{c}, t_{e}\right)-$ $\mathrm{A}\left(\mathcal{C} S W, T_{i}, t_{c}, t_{e}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right)$. Since $\mathrm{A}\left(\mathcal{I}, T_{i}, t_{c}, t_{e}\right)-\mathrm{A}\left(\mathcal{C} S W, T_{i}, t_{c}, t_{e}\right)$ denotes the lost allocation for this reweighting event, it follows that the absolute drift is at most $\mathrm{e}_{\max }\left(T_{i}\right)$. For example, in Figure 3.5, over the range $[2,4)$ for task $T_{3}$, an allocation of $\mathrm{A}\left(\mathcal{I}, T_{3}, 2,4\right)$ $\mathrm{A}\left(\mathcal{C S W}, T_{3}, 2,4\right)=2 / 3-2 / 4=1 / 6$ is lost.

Suppose that $T_{i}$ changes its weight to Nw at time $t_{c}$ via Rule N. If $T_{i}$ decreases its weight, then it is as though the allocation equal to $\mathrm{A}\left(\mathcal{I}, T_{i}, t_{c}, t_{e}\right)-\mathrm{A}\left(\mathcal{C} S W, T_{i}, t_{c}, t_{e}\right)$ is lost. Furthermore, since it was a weight decrease initiated at $t_{c}$, it follows that $\mathrm{A}\left(\mathcal{I}, T_{i}, t_{c}, t_{e}\right)<$ $\mathrm{A}\left(\mathcal{C} S W, T_{i}, t_{c}, t_{e}\right)$. Thus, since by (3.3) and (3.4), $\mathrm{A}\left(\mathcal{C} S W, T_{i}, t_{c}, t_{e}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right)$, it follows that $\mathrm{A}\left(\mathcal{I}, T_{i}, t_{c}, t_{e}\right)<\mathrm{A}\left(\mathcal{C} S W, T_{i}, t_{c}, t_{e}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right)$. Thus, $-\mathrm{e}_{\max }\left(T_{i}\right) \leq \mathrm{A}\left(\mathcal{I}, T_{i}, t_{c}, t_{e}\right)-$ $\mathrm{A}\left(\mathcal{C} S W, T_{i}, t_{c}, t_{e}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right)$. Since $\mathrm{A}\left(\mathcal{I}, T_{i}, t_{c}, t_{e}\right)-\mathrm{A}\left(\mathcal{C} S W, T_{i}, t_{c}, t_{e}\right)$ denotes the lost allocation for this reweighting event, it follows that the absolute drift incurred is at most $\mathrm{e}_{\max }\left(T_{i}\right)$. For example, in Figure 3.7, the drift incurred by $T_{4}$ is $-1 / 3$, i.e., $\operatorname{drift}\left(T_{4}, t\right)=$ $-1 / 3$, where $t \geq 2$. If $T_{i}$ increases its weight (Case (i)), then it incurs zero drift, since it immediately enacts the weight change (i.e., the scheduling weight changes immediately). Hence, the absolute drift incurred by this reweighting event is less than $\mathrm{e}_{\max }\left(T_{i}\right)$. For example, in Figure 3.6, the drift incurred by $T_{4}$ is 0 , i.e., $\operatorname{drift}\left(T_{4}, t\right)=0$, where $t \geq 2$.

Notice that the presence of jobs that miss their deadlines does not affect the drift bounds. The reason for this is that the reweighting rules are only based on the state of the active job at the time the reweighting event is initiated. Thus, if a job has not been scheduled by the time it reweights, then it does not matter whether a predecessor prevented the job from being scheduled or the job had the lowest scheduling priority, only that the job has not been scheduled, and therefore is positive-changeable.

Example (Figure 3.15). Consider the example in Figure 3.15, which depicts the partial schedule for a task $T_{i}$ that has all of the following characteristics: an initial weight of $1 / 10$ that increases to $1 / 2$ at time $t_{c}$; a job $T_{i}^{j-1}$ that has a deadline at $t_{r}=t_{c}-5$, an execution time of 14 , and misses its deadline by 11 time units; and all jobs released after $T_{i}^{j-1}$ have an execution time of 1 . Inset (a) depicts the GEDF schedule. Inset (b) depicts the CSW

(a)


(b)

Figure 3.15: A partial schedule that illustrates drift when tasks miss deadlines. The partial (a) GEDF and (b) CSW and IDEAL schedules for the task $T_{i}$. The difference between $T_{i}$ 's allocation in CSW and IDEAL are labeled above inset (b).
and IDEAL schedules. Notice that, even though $T_{i}^{j-1}$ misses its deadline at time $t_{r}$, when $T_{i}$ initiates the change at time $t_{c}, T_{i}^{j}$ is the active job, and since it has not been scheduled, $T_{i}$ is positive-changeable at $t_{c}$. Therefore, by Rule $\mathrm{P}, T_{i}^{j}$ (but not $T_{i}^{j-1}$ ) is halted and $T_{i}^{j+1}$ is immediately released with the new weight, which incurs a drift of $1 / 2$.

Modifications for NP-GEDF. Note that delaying the initiation of a reweighting event due to non-preemptivity does not substantially increase the drift incurred per reweighting event, since the longest a reweighting event can be delayed is the execution time of the active job of the task being reweighted.

Suppose that the task $T_{i}$ initiates a weight change at time $t_{c}$. If $T_{i}^{j}$ is active at $t_{c}$, and if $T_{i}$ 's reweighting event is delayed until some time $t$ (by a non-preemptive section), then at $t$ either (a) $T_{i}^{j}$ has a non-positive deviance (i.e., $T_{i}^{j}$ completes before its deadline), or (b) $t$ is
the first time that $T_{i}^{j}$ becomes inactive (i.e., $t=\min \left(r\left(T_{i}^{j+1}\right), \mathrm{d}\left(T_{i}^{j}\right)\right)$.)
If Case (a) occurs, then $T_{i}$ is negative-changeable at $t$, and $T_{i}^{j}$ is active at $t$. Hence, if $T_{i}$ increases its weight, then the only drift $T_{i}$ will incur for this reweighting event results from delaying the initiation of the event, i.e., at most $\mathrm{e}_{\max }\left(T_{i}\right)$. If $T_{i}$ decreases its weight, then delaying the reweighting event will not affect drift, since the enactment of the reweighting event would occur when $T_{i}^{j}$ becomes inactive, regardless of whether the initiation of the reweighting event was delayed or not.

Example (Figure 3.16). Consider the example in Figure 3.16, which depicts a one-processor system scheduled by NP-GEDF with two tasks: $T_{1}$, which has $\mathrm{wt}\left(T_{1}\right)=3 / 10$ and $\mathrm{e}\left(T_{1}\right)=3$; and $T_{2}$, which has e $\left(T_{2}\right)=2$ and an initial weight of $1 / 5$ and initiates a weight increase to $1 / 2$ at time 4. Inset (a) depicts the NP-GEDF schedule. Inset (b) depicts the CSW schedule. Inset (c) depicts the IDEAL schedule. Inset (d) depicts $T_{2}$ 's allocations in the CSW and IDEAL schedules. Notice that $T_{2}$ 's weight change is delayed from time 4 to time 5 because $T_{2}$ is non-preemptively executing at time 4 . As a result, $T_{2}$ is negative-changeable at time 5 . Also note that, $T_{2}^{2}$ is released when $T_{2}$ 's actual allocation equals its allocation in the CSW schedule at time 7, i.e., when $T_{2}^{2}$ 's deviance equals zero.

If Case (b), mentioned earlier, occurs, then either no job of $T_{i}$ is active at $t$ or $T_{i}^{j+1}$ is active at $t$. If no job of $T_{i}$ is active at $t$, then the change is enacted immediately, and the drift that the task incurs from the reweighting event is a result of delaying the initiation of the event, i.e., $\mathrm{e}_{\max }\left(T_{i}\right)$. If $T_{i}^{j+1}$ is active at $t$, then since $t=\min \left(\mathrm{r}\left(T_{i}^{j+1}\right), \mathrm{d}\left(T_{i}^{j}\right)\right)$, it must be the case that $\mathrm{r}\left(T_{i}^{j+1}\right)=t$. As a result, the weight change is enacted immediately and $T_{i}^{j+1}$ is released with the new weight. Hence, the only drift that is incurred is as a result of delaying the initiation of the reweighting event, i.e., at most $\mathrm{e}_{\max }\left(T_{i}\right)$.

Example (Figure 3.17). Consider the example in Figure 3.17, which depicts a partial NP-GEDF schedule for a task $T_{i}$, which has an initial weight of $1 / 10$ that increases to $1 / 2$ at time $t_{c}$ while the last-released job of $T_{i}$ before $t_{c}, T_{i}^{j}$, is both active and being scheduled. Note that $T_{i}^{j}$ has an execution time of four, and all jobs released after $T_{i}^{j}$ have an execution time of one. Moreover, $T_{i}^{j}$ does not complete execution until after its deadline. Inset (a) depicts the NP-GEDF schedule. Inset (b) depicts the CSW schedule. Inset (c) depicts


Figure 3.16: A one-processor example of drift in NP-GEDF, where $T_{2}^{1}$ completes before its deadline. (a) The NP-GEDF schedule. (b) The CSW schedule. (c) The IDEAL schedule. (d) $T_{2}$ 's allocations in the CSW and IDEAL schedules.
the IDEAL schedule. Inset (d) depicts $T_{i}$ 's allocation in the CSW and IDEAL schedules. Because $T_{i}^{j}$ is not complete by its deadline, the initiation of the weight change is delayed until $t=\mathrm{d}\left(T_{i}^{j}\right)=\mathrm{r}\left(T_{i}^{j+1}\right)$. Recall that, if a weight change is initiated when $\mathrm{d}\left(T_{i}^{j}\right)=\mathrm{r}\left(T_{i}^{j+1}\right)$, then the weight change is immediately enacted and $T_{i}^{j+1}$ is released with the new weight (even though $T_{i}^{j}$ has not yet completed execution). Thus, the only source of drift is because the initiation of the reweighting event is delayed.

From the reasoning presented in these examples, we can see that the following theorem holds.

Theorem 3.3. The per-event absolute drift under NP-GEDF for each task $T_{i}$ is at most $\mathrm{e}_{\max }\left(T_{i}\right)$.

### 3.6 Conclusion

In this chapter, we presented the adaptable sporadic task model as well as the rules for reweighting a task under the GEDF and NP-GEDF scheduling algorithms. In addition, we

| X | Fraction X of the processor scheduling the task |
| :---: | :---: |
| . Job lay | ut without reweighting $\downarrow$ Job released $\uparrow$ Job deadline |
| Reweig | ting event initiated $\quad \downarrow$ Reweighting event enacted |




(d)

Figure 3.17: A partial schedule of a one-processor example of drift in NP-GEDF. $T_{i}^{j}$ completes after its deadline. (a) The NP-GEDF schedule. (b) The CSW schedule. (c) The IDEAL schedule. (d) $T_{i}$ 's allocations in the CSW and IDEAL schedules.
proved tardiness bounds our reweighting rules by leveraging prior work by Devi and Anderson. In addition, we proved that the absolute value of the drift that can be incurred per reweighting event is at most the maximal execution time of a task.

## CHAPTER 4

## PEDF and NP-PEDF*

In this chapter, we examine the issue of reweighting in the context of partitioned algorithms. Because there cannot exist an optimal ${ }^{1}$ partitioned scheduling algorithm, we focus our attention on different heuristic tradeoffs that can minimize different sources of error. Before discussing these tradeoffs in detail, we first define some necessary notation and consider a fundamental limitation of all partitioning algorithms.

### 4.1 Preliminaries

In this section, we introduce a few terms that will facilitate our discussion of partitioned systems. We denote the $q^{\text {th }}$ processor in the system, where processors are ordered by some arbitrary method, as $P_{[q]}$. As a shorthand, we use $T_{i} \in P_{[q]}$ to denote that $T_{i}$ is assigned to $P_{[q]}$. We denote the set of tasks that are assigned to $P_{[q]}$ at time $t$ as $\operatorname{ASSN}\left(P_{[q]}, t\right)$. We denote the set of tasks that are assigned to $P_{[q]}$ and active at time $t$ as $\operatorname{ACT}\left(P_{[q]}, t\right)$. (Recall that a task $T_{i}$ is active at time if it has an active job at time $t$, and a job $T_{i}^{j}$ is active at time $t$ if $t \in\left[\mathrm{r}\left(T_{i}^{j}\right), \min \left(\mathrm{d}\left(T_{i}^{j}\right), \mathrm{r}\left(T_{i}^{j+1}\right)\right)\right)$. .) We denote the desired and guaranteed weight of $T_{i}$ at time $t$ as $\operatorname{Dwt}\left(T_{i}, t\right)$ and $\operatorname{Gwt}\left(T_{i}, t\right)$, respectively. (As we discuss in Section 4.4.3, when a task's guaranteed and desired weight differ, the releases and deadlines of its jobs will be based on its guaranteed weight.) If a task's desired weight does not change with time, then we denote

[^9]this value as $\operatorname{Dwt}\left(T_{i}\right)$.
We say that $T_{i} \in P_{[q]}$ is the heaviest task assigned to $P_{[q]}$ iff $T_{i}$ has the largest desired weight of any task assigned to $P_{[q]}$. Similarly, we say that $T_{i} \in P_{[q]}$ is the lightest task assigned to $P_{[q]}$ iff $T_{i}$ has the smallest desired weight of any task assigned to $P_{[q]}$. We say that a processor $P_{[q]}$ is over-utilized by $x$ iff it has been assigned tasks with a total desired weight of $1+x$. Similarly, we say that $P_{[q]}$ is under-utilized by $x$ iff it has been assigned tasks with a total desired weight of $1-x$. Additionally, we say that $P_{[q]}$ is fully-utilized iff it has been assigned tasks with a total desired weight of 1 . If $P_{[q]}$ is over-utilized by $x$ at time $t$, then we denote this value as $\omega\left(P_{[q]}, t\right)$; if $P_{[q]}$ is not over-utilized at $t$, then $\omega\left(P_{[q]}, t\right)=0$. These terms (as well as other terms used throughout this chapter) are summarized in Table 4.1.

### 4.2 A Limitation of Partitioning Schemes

As was mentioned in Section 1.2.2, under any partitioning scheme, there exist task systems where only a subset of tasks can receive their desired allocation even though the total weight of all tasks is at most the number of processors. For example, consider a two-processor system with three identical periodic tasks with an execution cost of 2.0 and a period of 3.0. Because tasks are partitioned, one processor will be assigned two of these tasks, thus over-utilizing it. There are two approaches for handling this problem. First, we could cap the total utilization of all tasks in the system. Unfortunately, under any $M$-processor partitioning scheme, a cap of approximately $M / 2$ is required in the worst case (Carpenter et al., 2004), which implies that as much as half the system's processing capacity could be lost. Such caps are due to connections to bin-packing.

An alternative approach is to assign a subset of tasks in the system guaranteed weights that are less than their desired weights. Although allocating a task a weight less than its desired weight is obviously undesirable, such an approach can guarantee that the system's overall capacity does not have to be restricted, which is a significant advantage in computationallyintensive systems like Whisper and VEC. Moreover, allowing the guaranteed weights of tasks to be somewhat malleable circumvents any bin-packing-like intractabilities that might otherwise arise - with frequent weight changes, such intractabilities would have to be dealt with

| Notation | Definition |
| :---: | :---: |
| $P_{[q]}$ | The $q^{t h}$ processor. |
| $T_{i} \in P_{[q]}$ | $T_{i}$ is assigned to $P_{[q]}$. |
| $\operatorname{ASSN}\left(P_{\text {[q] }}, t\right)$ | Set of tasks assigned to $P_{[q]}$ at time $t$. |
| ACT ( $\left.P_{[q]}, t\right)$ | Set of tasks assigned to $P_{[q]}$ that are active at time $t$. |
| $\mathrm{e}\left(T_{i}^{j}\right)$ | WECT of job $T_{i}^{j}$. |
| $\mathrm{e}_{\max }\left(T_{i}\right)$ | Maximal worst-case execution time of all jobs of $T_{i}$. |
| $\mathrm{Ae}\left(T_{i}^{j}\right)$ | Actual execution time of job $T_{i}^{j}$. |
| $\mathrm{r}\left(T_{i}^{j}\right)$ | Release time of $T_{i}^{j}$. |
| $\mathrm{d}\left(T_{i}^{j}, t\right)$ | Perceived deadline of $T_{i}^{j}$ at time $t$. |
| $\mathrm{d}\left(T_{i}^{j}\right)$ | Deadline of $T_{i}^{j}$. |
| $\theta\left(T_{i}^{j}\right)$ | IS separation between $T_{i}^{j-1}$ and $T_{i}^{j}$. |
| Dwt $\left(T_{i}, t\right)$ | $T_{i}$ 's desired weight at time $t$. |
| $\operatorname{Gwt}\left(T_{i}, t\right)$ | $T_{i}$ 's guaranteed weight at time $t$. |
| $\operatorname{SDwt}\left(T_{i}, t\right)$ | $T_{i}$ 's desired scheduling weight at time $t$. |
| SGwt ( $T_{i}, t$ ) | $T_{i}$ 's guaranteed scheduling weight at time $t$. |
| $\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)$ | $P_{[q]}$ 's desired weight scaling factor: $\max \left(1, \sum_{T_{i} \in \operatorname{ACT}\left(P_{[q]}, t\right)} \operatorname{Dwt}\left(T_{i}, t\right)\right)$. |
| $\mathcal{T S}\left(P_{[q]}, t\right)$ | $P_{[q]}$ 's desired scheduling weight scaling factor: $\max \left(1, \sum_{T_{i} \in \operatorname{ACT}\left(P_{[q]}, t\right)} \operatorname{SDwt}\left(T_{i}, t\right)\right)$. |
| $\omega\left(P_{[q]}, t\right)$ | $\max \left(0,1-\sum_{T_{i} \in \operatorname{ACT}\left(P_{\text {Pqd }}, t\right)} \operatorname{Dwt}\left(T_{i}, t\right)\right)$. |
| $\operatorname{Irem}\left(T_{i}^{j}, t\right)$ | $\mathrm{e}\left(T_{i}^{j}\right)-\int_{\mathrm{r}\left(T_{i}^{j}\right)}^{t} \operatorname{SGwt}\left(T_{i}, u\right) d u$. |
| SW | Scheduling-weight scheduling algorithm. |
| $\mathcal{S W}$ | SW schedule of a task system $\tau$. |
| CSW | Clairvoyant scheduling-weight scheduling algorithm. |
| CSW | CSW schedule of a task system $\tau$. |
| IDEAL | Ideal scheduling algorithm. |
| I | IDEAL schedule of a task system $\tau$. |
| PT | Partial ideal scheduling algorithm. |
| $\mathcal{P} T$ | PT schedule of a task system $\tau$. |
| $\mathcal{S}$ | Actual schedule (i.e., GEDF or NP-GEDF) of task system $\tau$. |
| $\mathrm{A}\left(\mathcal{B}, T_{i}^{j}, t_{1}, t_{2}\right)$ | Allocation to $T_{i}^{j}$ in the schedule $\mathcal{B}$ over $\left[t_{1}, t_{2}\right)$. |
| $\mathrm{A}\left(\mathcal{B}, T_{i}, t_{1}, t_{2}\right)$ | Allocation to $T_{i}$ in the schedule $\mathcal{B}$ over [ $\left.t_{1}, t_{2}\right)$. |
| $\operatorname{dev}\left(T_{i}^{j}, t\right)$ | Deviance of $T_{i}^{j}: \mathrm{A}\left(\mathcal{S} W, T_{i}^{j}, 0, t\right)-\mathrm{A}\left(\mathcal{S}, T_{i}^{j}, 0, t\right)$. |
| $\operatorname{drift}\left(T_{i}, t\right)$ | Drift of $T_{i}: \mathrm{A}\left(\mathcal{I}, T_{i}, 0, t\right)-\mathrm{A}\left(\mathcal{C} S W, T_{i}, 0, t\right)$. |
| Pdrift $\left(T_{i}, t\right)$ | Partial drift of $T_{i}: \mathrm{A}\left(\mathcal{P} T, T_{i}, 0, t\right)-\mathrm{A}\left(\mathcal{C} S W, T_{i}, 0, t\right)$. |
| Ow | Desired scheduling weight before a reweighting event. |
| Nw | New desired weight after a reweighting event. |
| $\operatorname{REM}\left(T_{i}^{j}, t\right)$ | Remaining execution time of $T_{i}^{j}$ at $t: \mathrm{e}\left(T_{i}^{j}\right)-\mathrm{A}\left(\mathcal{S}, T_{i}^{j}, 0, t\right)$. |
| $\mathrm{nextE}\left(T_{i}^{j}, t\right)$ | If $\operatorname{REM}\left(T_{i}^{j}, t\right)>0$, then $\operatorname{nextE}\left(T_{i}^{j}, t\right)=\operatorname{REM}\left(T_{i}^{j}, t\right)$; else, $\operatorname{nextE}\left(T_{i}^{j}, t\right)=\mathrm{e}\left(T_{i}^{j+1}\right)$. |
| $\mathbf{H}\left(T_{i}^{j}, t\right)$ | Jobs with a scheduling priority higher than or equal to $T_{i}^{j}$,s that are assigned to the same processor as $T_{i}^{j}$ and are both active and pending at time $t$. |
| $\operatorname{lag}\left(T_{i}^{j}, t\right)$ | Lag of $T_{i}^{j}$ at $t: \mathrm{A}\left(\mathcal{C} S W, T_{i}^{j}, \mathrm{r}\left(T_{i}^{j}\right), t\right)-\mathrm{A}\left(\mathcal{S}, T_{i}^{j}, \mathrm{r}\left(T_{i}^{j}\right), t\right)$. |
| $\operatorname{LAG}(G, t)$ | Lag of the job group $G$ at time $t$ : $\sum_{T_{i}^{j} \in G} \operatorname{lag}\left(T_{i}^{j}, t\right)$. |

Table 4.1: Summary of notation used in this chapter.
frequently at run-time. Note that we are still able to offer some service guarantees with this approach, as discussed in Sections 4.7 and 4.9. (In particular, if the maximum amount by which a processor is over-utilized is relatively small, then the resulting guaranteed weight may be acceptable.) For these reasons, we use this approach in the schemes we propose. To the best of our knowledge, we are the first to suggest using such an approach to schedule dynamically-changing multiprocessor workloads. The fundamental limitation of partitioned schemes noted at the beginning of this chapter is formalized by the following theorem.

Theorem 4.1. For any partitioned scheduling algorithm and any integers $M$ and $k$ such that $M \geq 2$ and $k \geq M+1$, there exists an $M$-processor system $\tau$ with $k$ tasks that have a total desired weight at most $M$ where at least one processor is assigned a set of tasks that have a total desired weight greater than one.

Proof. In order to prove Theorem 4.1, we construct a system that satisfies the theorem for any value of $M$ and $k$ such that $M \geq 2$ and $k \geq M+1$. Let the first $M$ tasks of $\tau$ have a desired weight $X=1-\epsilon$, where $0<\epsilon<0.5$. Let the $(M+1)^{s t}$ task of $\tau$ have a desired weight of $W=\min (M \cdot \epsilon-\delta, 1-\epsilon)$, where $0<\delta<\epsilon$, and let the total desired weight of the remaining $k-(M+1)$ tasks of $\tau$ be $\delta$. (For example, if $\epsilon=1 / 3, k=3$, and $M=2$, then the system consists of three tasks of weight $2 / 3$.) By definition, the desired weight of the first $(M+1)^{s t}$ tasks have a total desired weight of at most $M \cdot(1-\epsilon)+M \cdot \epsilon-\delta=M-\delta$. Since the total desired weight of the remaining $k-(M+1)$ tasks is $\delta$, the total desired weight of all the tasks is at most $M$.

Thus, it remains to be shown that one processor is assigned tasks with a total desired weight greater than one. Notice that no matter how the first $M+1$ tasks are partitioned, at least one processor will been assigned two of these tasks, i.e., at least one processor will be assigned either two of the first $M$ tasks or one of the first $M$ tasks and the $(M+1)^{s t}$ task. If two of the first $M$ tasks are assigned to the same processor, then the total desired weight on that processor is at least $2 \cdot(1-\epsilon)$. In this case, since $\epsilon<0.5$, the processor will have been assigned tasks with total desired weight greater than one. If one processor is assigned one of the first $M$ tasks and the $(M+1)^{s t}$ task, then the total desired weight of the tasks assigned to this processor would be $1-\epsilon+W$. Thus, it remains to be shown that $W>\epsilon$. Thus, we
consider two cases depending on whether $W=1-\epsilon$ or $W=M \cdot \epsilon-\delta$. If $W=1-\epsilon$, then since $\epsilon<0.5, W>\epsilon$. If $W=M \cdot \epsilon-\delta$, then since $M \geq 2$ and $\delta<\epsilon$, it follows that $W>\epsilon$. This completes the proof.

### 4.3 Partitioning and Repartitioning

The problem of assigning tasks to processors is equivalent to the NP-hard bin-packing problem. Given that reweighting events may be frequent, an optimal assignment of tasks to processors is not realistic to maintain. In our approach, we partition $N$ tasks onto $M$ processors in $O(M+N \log N)$ time by first sorting them by desired weight from heaviest to lightest, and by then placing each on the processor that is the "best fit" (this partitioning method is called descending best-fit). We chose this method because it falls within a class of bin-packing heuristics called reasonable allocation decreasing (RAD), which has been shown by Lopez et al. to produce better packings than other types of heuristics (Lopez et al., 2004). Most importantly, the "descending best-fit" strategy can guarantee that no processor is overutilized by more than $W$, where $W$ is the desired weight of the $\left(M \cdot\left\lfloor\frac{1}{X}\right\rfloor+1\right)^{\text {st }}$ heaviest task in the system and $X$ is the desired weight of the heaviest task in the system. Also, under this strategy, no processor is over-utilized by more than the desired weight of the lightest task assigned to it.

As tasks are reweighted, the likelihood of a processor becoming substantially over-utilized increases dramatically, creating significant overall error (however assessed) on these processors. The extent of overall error can be controlled by repartitioning the system. In order to give the user control over migration overhead, we introduce $\alpha$-partitioning: if a reweighting event causes the total desired weight of all tasks assigned to any one processor to be at least $1+\alpha$, the system is reset, where $\alpha$ is a user-defined value. A reset causes the set of tasks to be repartitioned (using the descending best-fit method described earlier) and each active task to issue a new job with the remaining execution time of its pending job ("pending" is formally defined in Section 1.2). In Section 4.6, we formally define the rule for resetting a system.

| Metric Name | Metric Formula |
| :---: | :---: |
| MAOE | $\max _{T_{i} \in \operatorname{ASSN}\left(P_{[q]}, t\right)}\left\{\operatorname{Dwt}\left(T_{i}, t\right)-\operatorname{Gwt}\left(T_{i}, t\right)\right\}$ |
| AAOE | $\frac{1}{n} \cdot \sum_{T_{i} \in \operatorname{ASSN}\left(P_{[q]}, t\right)}\left(\operatorname{Dwt}\left(T_{i}, t\right)-\operatorname{Gwt}\left(T_{i}, t\right)\right)$ |
| MROE | $\max _{T_{i} \in \operatorname{ASSN}\left(P_{[q]}, t\right)}\left\{\frac{\operatorname{Dwt}\left(T_{i}, t\right)-\operatorname{Gwt}\left(T_{i}, t\right)}{\operatorname{Dt}\left(T_{i}, t\right)}\right\}$ |
| AROE | $\frac{1}{n} \cdot \sum_{T_{i} \in \operatorname{ASSN}\left(P_{[q]}, t\right)}\left(\frac{\operatorname{Dwt}\left(T_{i}, t\right)-\operatorname{cwt}\left(T_{i}, t\right)}{\operatorname{Dwt}(T, t)}\right)$ |

Table 4.2: The MAOE, AAOE, MROE, and AROE metrics.

### 4.4 Allowing Guaranteed and Desired Weights to Differ

Given that we allow a task's guaranteed and desired weight to differ, two questions arise. First, how do we determine each task's guaranteed weight? Second, how do we modify the adaptable sporadic task model (presented in Section 3.1) to accommodate a task with a different guaranteed and desired weight? In this section, we discuss several possible answers to these questions.

### 4.4.1 Determining Guaranteed Weights

When a processor $P_{[q]}$ is over-utilized, there is a degradation in the system's performance because at least one task assigned to $P_{[q]}$ will have a guaranteed weight that is less than its desired weight. Unfortunately, it is not immediately clear how this degradation should be measured. For example, should we measure the absolute or relative difference between each task's desired and guaranteed weight? In this section, we propose four different metrics for measuring the system's degradation, summarized in Tables 4.2 and 4.3, and explain how to determine the guaranteed weight of tasks in order to minimizing each of these metrics. ${ }^{2}$ We illustrate these metrics via the following example.

Example (Figure 4.1). Consider the example in Figure 4.1, which depicts one processor that is assigned four tasks: $T_{1}$, which has a desired weight of $0.36 ; T_{2}$ and $T_{3}$, both of which have a have a desired weight of 0.30 ; and $T_{4}$, which has a desired weight of 0.24 . Inset (a) depicts the desired weight for each task. Insets (b), (c), and (d) depict, respectively, the

[^10]| Metric Name | Guaranteed Weight Assignment |
| :---: | :---: |
| MAOE | $\operatorname{Gwt}\left(T_{i}, t\right)=\operatorname{Dwt}\left(T_{i}, t\right)-\frac{\omega\left(P_{[q,}, t\right)}{\left.\operatorname{ASSN}\left(P_{[q]}, t\right)\right]}$ |
| AAOE | $1=\sum_{T_{z} \in \operatorname{ASSN}\left(P_{[q]}, t\right)} \operatorname{Gwt}\left(T_{z}, t\right)$ |
| MROE | $\operatorname{Gwt}\left(T_{i}, t\right)=\frac{\sum_{T_{z} \in \operatorname{ASSN}\left(P_{i, t}, t\right)} \operatorname{Dwt}\left(T_{z}, t\right)}{}$ |
| AROE | $\operatorname{Gwt}\left(T_{i}, t\right)= \begin{cases}\operatorname{Dwt}\left(T_{i}, t\right)-\omega\left(P_{[q]}, t\right) & \text { if } T_{i}=T_{H} \\ \operatorname{Dwt}\left(T_{i}, t\right) & \text { otherwise }\end{cases}$ |

Table 4.3: The guaranteed weight assignments for $T_{i} \in \operatorname{ASSN}\left(P_{[q]}, t\right)$ that minimize the MAOE, AAOE, MROE, and AROE metrics. $T_{H}$ is the heaviest task in $\operatorname{ASSN}\left(P_{[q]}, t\right)$.
guaranteed weights for the four tasks when such weights are chosen to minimize the MAOE, MROE, and AAOE metrics (discussed shortly).

Absolute Error. The first two metrics we consider are based on the absolute difference between the guaranteed and desired weight of a task. Specifically, the maximal absolute overall error (MAOE) on an over-utilized processor, $P_{[q]}$, is given by

$$
\max _{T_{i} \in \operatorname{ASSN}\left(P_{[q]}, t\right)}\left\{\operatorname{Dwt}\left(T_{i}, t\right)-\operatorname{Gwt}\left(T_{i}, t\right)\right\} .
$$

To minimize this metric, the difference between the guaranteed and the desired weight for each task should be the same. Specifically, to minimize the MAOE metric, the guaranteed weight for $T_{i} \in P_{[q]}$ at time $t$ is specified as

$$
\begin{equation*}
\operatorname{Gwt}\left(T_{i}, t\right)=\operatorname{Dwt}\left(T_{i}, t\right)-\frac{\omega\left(P_{[q]}, t\right)}{\left|\operatorname{ASSN}\left(P_{[q]}, t\right)\right|} . \tag{4.1}
\end{equation*}
$$

Notice that, in Figure 4.1, since the processor is over-utilized by 0.20 and there are four tasks assigned to it, MAOE is minimized by setting each task's guaranteed weight to be $\frac{0.2}{4}$ less than its desired weight.

It is worthwhile to note that (4.1) only produces non-negative guaranteed weights if the desired weight for each task is at least $\frac{\omega\left(P_{[q]}, t\right)}{\left|\operatorname{ASSN}\left(P_{[q]}, t\right)\right|}$. This condition is satisfied when $P_{[q]}$ is over-utilized by less than smallest desired weight of any task assigned to $P_{[q]}$. Since, as we discussed in Section 4.3, such a property can be guaranteed by any RAD partitioning


Figure 4.1: (a) The desired weights for four tasks assigned to one processor. Guaranteed weights for the four tasks when the guaranteed weights are chosen to minimize the (b) MAOE, (c) MROE and (d) AAOE metrics.
algorithm, it is possible to repartition the system in order to guarantee that (4.1) returns valid results. If repartitioning cannot be performed (e.g., for application-oriented reasons), then it is possible to use an iterative approach for determining the guaranteed weight of tasks.

The average absolute overall error (AAOE) on an over-utilized processor, $P_{[q]}$, is given by

$$
\frac{1}{n} \cdot \sum_{T_{i} \in \operatorname{ASSN}\left(P_{[q]}, t\right)}\left(\operatorname{Dwt}\left(T_{i}, t\right)-\operatorname{Gwt}\left(T_{i}, t\right)\right) .
$$

It is easy to show that this metric is minimized whenever the guaranteed weight of all tasks assigned to a processor sum to 1 . Since any reasonable method for determining the guaranteed weights of tasks will minimize the AAOE metric, this metric is of little value.

Relative error. The next two metrics we consider are based on the relative difference between the guaranteed and desired weights of a task. The maximal relative overall error (MROE) on an over-utilized processor, $P_{[q]}$, is given by

$$
\max _{T_{i} \in \operatorname{ASSN}\left(P_{[q]}, t\right)}\left\{\frac{\operatorname{Dwt}\left(T_{i}, t\right)-\operatorname{Gwt}\left(T_{i}, t\right)}{\operatorname{Dwt}\left(T_{i}, t\right)}\right\} .
$$

This metric is minimized when all task shares are scaled by the same value. Specifically, we define the guaranteed weight of $T_{i}$ assigned to the processor $P_{[q]}$ at time $t$ as

$$
\begin{equation*}
\operatorname{Gwt}\left(T_{i}, t\right)=\frac{\operatorname{Dwt}\left(T_{i}, t\right)}{\mathcal{T D}\left(P_{[q]}, t\right)}, \tag{4.2}
\end{equation*}
$$

where

$$
\begin{equation*}
\mathcal{T D}\left(P_{[q]}, t\right)=\max \left(1, \sum_{T_{i} \in \operatorname{ASSN}\left(P_{[q]}, t\right)} \operatorname{Dwt}\left(T_{i}, t\right)\right) \tag{4.3}
\end{equation*}
$$

For example, consider the system in Figure 4.1(c). The depicted set of tasks over-utilizes the processor by 0.2 , so each task's guaranteed weight should be $\frac{1}{1.2}$ times its desired weight. This scaling is the same as the proportional-share scaling used in EEVDF (Stoica et al., 1996).

The average relative overall error (AROE) on an over-utilized processor $P_{[q]}$ is given by

$$
\frac{1}{n} \cdot \sum_{T_{i} \in P_{[q]}} \frac{\operatorname{Dwt}\left(T_{i}, t\right)-\operatorname{Gwt}\left(T_{i}, t\right)}{\operatorname{Dwt}(T, t)}
$$

This metric is minimized when the guaranteed weight for the heaviest task on an over-utilized processor $P_{[q]}$ is less than its desired weight by $\omega\left(P_{[q]}, t\right)$, and the guaranteed weight of every other task equals its desired weight. Specifically, the guaranteed weight of $T_{i} \in \operatorname{ASSN}\left(P_{[q]}, t\right)$ is defined as

$$
\operatorname{Gwt}\left(T_{i}, t\right)= \begin{cases}\operatorname{Dwt}\left(T_{i}, t\right)-\omega\left(P_{[q]}, t\right), & \text { if } T_{i} \text { is the heaviest task in } \operatorname{ASSN}\left(P_{[q]}, t\right)  \tag{4.4}\\ \operatorname{Dwt}\left(T_{i}, t\right), & \text { otherwise. }\end{cases}
$$

For example, consider the system depicted in Figure 4.1(d). In this system, since $T_{1}$ has the largest desired weight, its guaranteed weight is its desired weight minus the amount the processor is over-utilized, i.e., $0.36-0.20=0.16$.

Notice that (4.4) is valid only if $\omega\left(P_{[q]}, t\right) \leq \operatorname{Dwt}\left(T_{H}, t\right)$, where $T_{H}$ is the heaviest task in $\operatorname{ASSN}\left(P_{[q]}, t\right)$. As with the MAOE metric, this condition is satisfied by repartitioning the system by using any RAD partitioning algorithm. If repartitioning cannot be performed (e.g., for application-oriented reasons), then it is possible to reduce the AROE by itierativly choosing the smallest possible guaranteed weight for tasks from the heaviest to lightest task
in $\operatorname{ASSN}\left(P_{[q]}, t\right)$ until the total guaranteed weight equals one.
To simplify our discussion, we focus on minimizing MROE. It is worthwhile to note that under our adaptable framework, any method for determining the guaranteed weights of tasks on an over-utilized processor can be used so long as the sum of these weights is at most one. In Section 4.10, we explain how our system can be extended to accommodate any such method, including methods based the AROE and MAOE metrics.

### 4.4.2 The Adaptable Sporadic Task Model, Revisted

In this section, we incorporate the notions of guaranteed and desired weights into the adaptable sporadic task model that was presented in Section 3.1. It is important to note that, unless otherwise specified, all references to the adaptable sporadic task model in this chapter are to the task model presented in this section. We begin our discussion by formally defining when a task changes its weight. A task $T_{i}$ changes its desired weight at time $t$ if $\operatorname{Dwt}\left(T_{i}, t-\epsilon\right) \neq \operatorname{Dwt}\left(T_{i}, t\right)$ where $\epsilon \rightarrow 0^{+}$. Similarly, a task, $T_{i}$, changes its guaranteed weight at time $t$ if $\operatorname{Gwt}\left(T_{i}, t-\epsilon\right) \neq \operatorname{Gwt}\left(T_{i}, t\right)$ where $\epsilon \rightarrow 0^{+}$.

Notice that, since we are attempting to minimize MROE, by (4.2), the guaranteed weight for every task on an over-utilized processor is a function of its desired weight and the desired weight of every other task assigned to the same processor. As a result, when a task, $T_{i}$, changes its desired weight at time $t$, the guaranteed weight for $T_{i}$ and for every other task assigned to same processor may change.

Additionally, different actions may occur depending on whether the desired or guaranteed weight of a task changes. Specifically, if a task, $T_{i}$, changes its desired weight at time $t_{c}$ when a job $T_{i}^{j}$ is active, then the following two actions may occur.

- The execution time of $T_{i}^{j}$ may be reduced to the amount of time for which $T_{i}^{j}$ has executed prior to $t_{c}$, and the execution time of $T_{i}^{j+1}$ may be redefined to be the amount of time "lost" by reducing the execution time of $T_{i}^{j}$.
- $\mathrm{r}\left(T_{i}^{j+1}\right)$ may be redefined to be less than $\mathrm{d}\left(T_{i}^{j}\right)$, which would cause jobs $T_{i}^{j}$ and $T_{i}^{j+1}$ to "overlap."

If the guaranteed weight but not the desired weight of a task changes, then the following two actions may occur

- The deadline of $T_{i}^{j}$ may be changed.
- The release time of $T_{i}^{j+1}$ (if it exists) may be changed.

Scheduling weights. Just as with the adaptable sporadic task model presented in Section 3.1, there can be a difference between when a desired weight change is initiated and when it is enacted. We use the desired scheduling weight of a task $T_{i}$ at time $t$, denoted $\operatorname{SDwt}\left(T_{i}, t\right)$, to represent the "last enacted desired weight of $T_{i}$." Formally, $\operatorname{SDwt}\left(T_{i}, t\right)$ equals $\operatorname{Dwt}\left(T_{i}, u\right)$, where $u$ is the last time at or before $t$ that a weight change was enacted for $T_{i}$. It is important to note that we use the desired scheduling weight of a task to compute the guaranteed scheduling weight of a task, denoted $\operatorname{SGwt}\left(T_{i}, t\right)$. In turn, the guaranteed scheduling weight is used to compute the deadlines and releases of tasks.

Formally, the guaranteed scheduling weight of $T_{i}$ that is assigned to the processor $P_{[q]}$ at time $t$ is

$$
\begin{equation*}
\operatorname{SGwt}\left(T_{i}, t\right)=\frac{\operatorname{SDwt}\left(T_{i}, t\right)}{\mathcal{T S}\left(P_{[q]}, t\right)} \tag{4.5}
\end{equation*}
$$

where

$$
\begin{equation*}
\mathcal{T S}\left(P_{[q]}, t\right)=\max \left(1, \sum_{T_{i} \in \operatorname{ASSN}\left(P_{[q]}, t\right)} \operatorname{SDwt}\left(T_{i}, t\right)\right) \tag{4.6}
\end{equation*}
$$

Notice that, by (4.5), the sum of the guaranteed scheduling weights of all active tasks assigned to a processor is at most one, which is formalized by the following property
(W) For any processor $P_{[q]}$ and any time $t, \sum_{T_{i}^{j} \in \operatorname{ASSN}\left(P_{[q]}, t\right)} \operatorname{SGwt}\left(T_{i}, t\right) \leq 1$.

Because the rules for changing the guaranteed weight of a task are simpler than changing its desired weight, it is possible to integrate them directly into the definition of a job's release and deadline. To do so, we introduce the notion of a perceived deadline of $T_{i}^{j}$ at time $t$, denoted $\mathrm{d}\left(T_{i}^{j}, t\right)$, which represents what the deadline of $T_{i}^{j}$ would be if its guaranteed weight did not change. As a shorthand, we use $\mathrm{d}\left(T_{i}^{j}\right)$ to denote the time $u$ such that $u=\mathrm{d}\left(T_{i}^{j}, u\right)$. ( $\mathrm{d}\left(T_{i}^{j}\right)$ represents the actual deadline of the job, but it cannot be determined until it is reached
(see the example below). Formally, we define $\mathrm{d}\left(T_{i}^{j}, t\right)$ as,

$$
\mathrm{d}\left(T_{i}^{j}, t\right)= \begin{cases}t+\frac{\operatorname{Irem}\left(T_{i}^{j}, t\right)}{\operatorname{SGwt}\left(T_{i}, t\right)}, & \text { if } \operatorname{SGwt}\left(T_{i}, t\right)>0  \tag{4.7}\\ \infty, & \text { otherwise }\end{cases}
$$

where

$$
\operatorname{Irem}\left(T_{i}^{j}, t\right)= \begin{cases}\mathrm{e}\left(T_{i}^{j}\right), & \text { if } t<\mathrm{r}\left(T_{i}^{j}\right)  \tag{4.8}\\ \mathrm{e}\left(T_{i}^{j}\right)-\int_{\mathrm{r}\left(T_{i}^{j}\right)}^{t} \operatorname{SGwt}\left(T_{i}, u\right) d u, & \text { otherwise. }\end{cases}
$$

Additionally, we define $\mathrm{r}\left(T_{i}^{j}\right)$ as,

$$
\begin{array}{rr}
\mathrm{r}\left(T_{i}^{j}\right)=\theta\left(T_{i}^{j}\right), & j=1 \\
\mathrm{r}\left(T_{i}^{j}\right)=\mathrm{d}\left(T_{i}^{j-1}\right)+\theta\left(T_{i}^{j}\right), & j>1, \tag{4.10}
\end{array}
$$

where $\theta\left(T_{i}^{j}\right) \geq 0$. Notice that, if for some job $T_{i}^{j}$ and $t>\mathrm{r}\left(T_{i}^{j}\right), \operatorname{SGwt}\left(T_{i}, t\right)=0$, then it is possible that $\mathrm{d}\left(T_{i}^{j}\right)$ cannot be reached since $\mathrm{d}\left(T_{i}^{j}, t\right)=\infty$. The scenario where $\operatorname{SGwt}\left(T_{i}, t\right)=0$ is a special case that is used to represent $T_{i}$ leaving the system. Thus, once $\operatorname{SGwt}\left(T_{i}, t\right)=0$ holds, $T_{i}$ cannot release any more jobs and is no longer allocated any capacity in any schedule. As a result, if $\mathrm{d}\left(T_{i}^{j}, t\right)=\infty$, then we set $\mathrm{d}\left(T_{i}^{j}\right)$ as $\infty$. Notice that all other terms of $T_{i}$ are still well-defined since the only other term that is defined using $\mathrm{d}\left(T_{i}^{j}\right)$ is $\mathrm{r}\left(T_{i}^{j+1}\right)$; however, if $\mathrm{d}\left(T_{i}^{j}\right)=\infty$, then $T_{i}^{j+1}$ does not exist.

From the definition of $\mathrm{d}\left(T_{i}^{j}, t\right)$ in (4.7) it is not hard to see that the following property holds.
(D) For any two times $u_{1}$ and $u_{2}$ such that $\mathrm{r}\left(T_{i}^{j}\right) \leq u_{1} \leq u_{2} \leq \mathrm{d}\left(T_{i}^{j}\right), \int_{u_{1}}^{u_{2}} \operatorname{SGwt}\left(T_{i}, t\right) d t \leq$ $\mathrm{e}\left(T_{i}^{j}\right)$.

Notice that, since for $t \in\left[\mathrm{r}\left(T_{i}^{j}\right), \mathrm{d}\left(T_{i}^{j}\right)\right), \int_{\mathrm{r}\left(T_{i}^{j}\right)}^{t} \mathrm{SGwt}\left(T_{i}, u\right) \leq \mathrm{e}\left(T_{i}^{j}\right)$, by (4.8), it follows that for any time $t$

$$
\begin{equation*}
\operatorname{Irem}\left(T_{i}^{j}, t\right) \leq \mathrm{e}\left(T_{i}^{j}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right) \tag{4.11}
\end{equation*}
$$

It is important to note that the PEDF and NP-PEDF scheduling algorithms discussed in this chapter prioritize jobs based on their perceived deadlines; however, even though jobs are
prioritized based on their perceived deadlines, the relative scheduling priority between any two jobs is time invariant. Specifically, it two jobs, $T_{i}^{j}$ and $T_{a}^{b}$, are assigned to the same processor and at some time $t_{1} \geq \max \left(\mathrm{r}\left(T_{i}^{j}\right), \mathrm{r}\left(T_{a}^{b}\right)\right), \mathrm{d}\left(T_{i}^{j}, t_{1}\right)<\mathrm{d}\left(T_{a}^{b}, t_{1}\right)$ holds, then at any time $t \geq \max \left(\mathrm{r}\left(T_{i}^{j}\right), \mathrm{r}\left(T_{a}^{b}\right)\right), \mathrm{d}\left(T_{i}^{j}, t\right)<\mathrm{d}\left(T_{a}^{b}, t\right)$ holds. Intuitively, the reason for this behavior is that the deadlines for both jobs always scale by the same factor, i.e., $\frac{1}{\operatorname{TS}\left(P_{[q]}, t\right)}$. Also, when the desired weight of a task changes, the relative scheduling priority between any two released jobs of any two tasks is unchanged. (When task's desired weight is changed, it is possible that the current job may halt and release a new job with a higher or lower scheduling priority. Such an action would change the relative scheduling priority between two tasks but not between the jobs themselves because the original job was halted.)

Example (Figure 4.2). Consider the example in Figure 4.2, which depicts a processor that is assigned four tasks: $T_{1}$, which has $\mathrm{e}\left(T_{1}\right)=3$ and $\operatorname{Dwt}\left(T_{1}\right)=2 / 5 ; T_{2}$, which has $\mathrm{e}\left(T_{2}\right)=2$ and $\operatorname{Dwt}\left(T_{2}\right)=1 / 3 ; T_{3}$, which has $\mathrm{e}\left(T_{3}\right)=1$ and $\operatorname{Dwt}\left(T_{3}\right)=1 / 3$; and $T_{4}$, which has $\mathrm{e}\left(T_{4}\right)=4$ and $\operatorname{Dwt}\left(T_{4}\right)=4 / 15$. The total desired weight is $\frac{4}{3}$. $T_{4}^{1}$ 's perceived deadline is shown above each inset. Inset (a) depicts the scenario where $T_{1}$ never leaves. Inset (b) depicts the scenario where $T_{1}$ leaves at time 10 causing the processor to be under-utilized. Notice that, in inset (b), $T_{4}^{1}$ 's perceived deadline changes when $T_{1}$ leaves (at time 10) from 20 to $18 . \overline{3}$ because $T_{4}^{1}$ 's guaranteed scheduling weight changes. This differs from inset (a), in which $T_{4}^{1}$ 's guaranteed scheduling weight does not change, and as a result, $T_{4}^{1}$ 's perceived deadline remains constant. One final note: while, in inset (a), $\mathrm{d}\left(T_{4}^{1}\right)=20$, and in inset (b), $\mathrm{d}\left(T_{4}^{1}\right)=18 . \overline{3}$, neither of these values are known until these corresponding points in time are reached.

Because the reweighting rules may cause $\mathrm{r}\left(T_{i}^{j+1}\right)<\mathrm{d}\left(T_{i}^{j}\right)$, we must slightly modify the definition of "window," "active," and "inactive" presented in Section 1.2.

Definition 4.1 (Window, Active, and Inactive). If $T_{i}^{j}$ is a job in the adaptable sporadic task system, $T$, then the window of $T_{i}^{j}$ is defined as the range $\left[\mathrm{r}\left(T_{i}^{j}\right), \min \left(\mathrm{d}\left(T_{i}^{j}\right), \mathrm{r}\left(T_{i}^{j+1}\right)\right)\right)$. Furthermore, the job $T_{i}^{j}$ is active at time $t$ iff $t$ is in $T_{i}^{j}$,s window (i.e., $t \in$ $\left[\mathrm{r}\left(T_{i}^{j}\right), \min \left(\mathrm{d}\left(T_{i}^{j}\right), \mathrm{r}\left(T_{i}^{j+1}\right)\right)\right)$, and is inactive otherwise.


Figure 4.2: A one-processor system with four tasks. (a) $T_{1}$ never leaves. (b) $T_{1}$ leaves at time 10. The perceived deadline for $T_{4}^{1}$ is shown above each figure.

### 4.4.3 Modifying the SW Algorithm

Having extended the adaptable sporadic task model presented in Section 3.1, we now extend the SW theoretical algorithm, presented in Section 3.2, to incorporate desired and guaranteed weights. Under the SW scheduling algorithm, at each instant $t$, each active job $T_{i}^{j}$ in $\tau$ is allocated a fraction of the system equal to its guaranteed scheduling weight $\operatorname{SGwt}\left(T_{i}, t\right)$. Hence, if a job $T_{i}^{j}$ is active over the range $\left[t_{1}, t_{2}\right)$, then over this range, $T_{i}^{j}$ is allocated $\int_{t_{1}}^{t_{2}} \operatorname{SGwt}\left(T_{i}, u\right) d u$ time. Throughout this chapter, we use $\mathcal{S} W$ to denote the SW schedule of a task system $\tau$.

Deviance. The deviance of the job $T_{i}^{j}$ of the task $T_{i}$ at time $t$ is defined as

$$
\begin{equation*}
\operatorname{dev}\left(T_{i}^{j}, t\right)=\mathrm{A}\left(\mathcal{S} W, T_{i}^{j}, 0, t\right)-\mathrm{A}\left(\mathcal{S}, T_{i}^{j}, 0, t\right) \tag{4.12}
\end{equation*}
$$

where $\mathcal{S}$ is the actual schedule.

Example (Figure 4.3). Consider the example in Figure 4.3, which depicts one processor that is assigned four tasks: $T_{1}$, which has $\mathrm{e}\left(T_{1}\right)=1$ and $\operatorname{Dwt}\left(T_{1}\right)=1 / 3 ; T_{2}$, which has $\mathrm{e}\left(T_{2}\right)=1$ and $\operatorname{Dwt}\left(T_{2}\right)=1 / 6 ; T_{3}$, which has $\mathrm{e}\left(T_{3}\right)=2, \operatorname{Dwt}\left(T_{3}\right)=1 / 4$, and leaves at time $8 ;$ and $T_{4}$, which has $\mathrm{e}\left(T_{4}\right)=4$ and an initial desired weight of $1 / 4$ that increases to $1 / 2$ at time 8. (The reweighting rules presented in Section 4.5 stop $T_{4}^{1}$ from receiving more than one unit of allocation and cause $T_{4}^{2}$ to be released with the remaining three units of execution at time 8.) Inset (a) depicts the PEDF schedule. Inset (b) depicts the SW schedule. Inset (c) depicts the allocations to $T_{4}$ in the PEDF and SW schedules. Notice that $T_{4}^{1}$ has positive deviance at time 8.

Example (Figure 4.4). Consider the example in Figure 4.4, which depicts one processor that is assigned four tasks (where the execution time of each job is one): $T_{1}$, which has $\operatorname{Dwt}\left(T_{1}\right)=1 / 2 ; T_{2}$ and $T_{3}$, both of which have a desired weight of $1 / 6$; and $T_{4}$, which has an initial desired weight of $1 / 2$ that initiates a desired weight decrease to $1 / 6$ at time 1 that is enacted at time $2 . \overline{6}$. (The reweighting rules presented in Section 4.5 do not allow $T_{2}$ to decrease its desired weight immediately.) Notice that $T_{4}^{1}$ has negative deviance at time 1.

### 4.5 Changing Desired Weights

We now introduce two rules for changing the desired weight of a task. (These rules are similar to reweighting rules for changing the weight of a task under GEDF that were presented in Section 3.4.) These rules work by modifying future release times and deadlines. (The rules below are applied on a single processor; reweighting events that trigger a repartitioning are discussed in Section 4.6.) We first describe how to change the desired weight of a task under PEDF in Section 4.5.1 and then explain how to change the desired weight of a task under NP-PEDF in Section 4.5.2.

### 4.5.1 Changing Desired Weights in PEDF

Let $\tau$ be a task system in which some task $T_{i}$ initiates a desired weight change to Nw at time $t_{c}$. Let Ow be the last desired scheduling weight of $T_{i}$ before the change is initiated at $t_{c}$. Let


Figure 4.3: A one-processor system with four tasks. (a) The PEDF schedule. (b) The SW schedule. (c) The allocations to $T_{4}$ in the PEDF and SW schedules.


Figure 4.4: A one-processor system with four tasks. (a) The PEDF schedule. (b) The SW schedule. (c) The allocations to $T_{4}$ in the PEDF and SW scheduling algorithms.
$\mathcal{S}$ be the $M$-processor PEDF schedule of $\tau$. Let $T_{i}^{j}$ be last-released job of $T_{i}$ before $t_{c}$. If $T_{i}^{j}$ does not exist or $T_{i}^{j}$ is inactive at $t_{c}$ before the reweighting event is initiated (i.e., $t_{c} \geq \mathrm{d}\left(T_{i}^{j}\right)$ ), then the desired weight change is immediately enacted, and future jobs of $T_{i}$ are released with the new desired weight. In the following rules, we consider the remaining possibility, i.e., $T_{i}^{j}$ exists and is active at $t_{c}$. (Notice that, if $t_{c}=\mathrm{d}\left(T_{i}^{k}\right)=\mathrm{r}\left(T_{i}^{k+1}\right)$, then $T_{i}^{k}$ is the last-released job of $T_{i}$ before $t_{c}$, and it is not active at $t_{c}$. Therefore, the change is immediately enacted and $T_{i}^{k+1}$ is released with the new desired weight.)

Let $\operatorname{REM}\left(T_{i}^{j}, t\right)=\mathrm{e}\left(T_{i}^{j}\right)-\mathrm{A}\left(\mathcal{S}, T_{i}^{j}, 0, t\right)$. Note that $\operatorname{REM}\left(T_{i}^{j}, t\right)$ denotes the actual remaining computation in $T_{i}$ 's current job. Since for any time $t, \mathrm{~A}\left(\mathcal{S}, T_{i}^{j}, 0, t\right) \geq 0$,

$$
\begin{equation*}
\operatorname{REM}\left(T_{i}^{j}, t\right) \leq \mathrm{e}\left(T_{i}^{j}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right) . \tag{4.13}
\end{equation*}
$$

Let $\operatorname{nextE}\left(T_{i}^{j}, t_{c}\right)$ equal $\operatorname{REM}\left(T_{i}^{j}, t_{c}\right)$, if $\operatorname{REM}\left(T_{i}^{j}, t_{c}\right)>0$; otherwise, if $\operatorname{REM}\left(T_{i}^{j}, t_{c}\right)=0$, then let $\operatorname{next} \mathrm{E}\left(T_{i}^{j}, t_{c}\right)$ equal the value of $\mathrm{e}\left(T_{i}^{j+1}\right)$ had the desired weight change event not occurred. Since $\operatorname{next} \mathrm{E}\left(T_{i}^{j}, t_{c}\right)$ is only used to determine the execution time of the next job released, it can be calculated at time $\mathrm{r}\left(T_{i}^{j+1}\right)$. (Notice that, if $T_{i}$ has no next job $T_{i}^{j+1}$ to release at the time specified in the rules below, then $\operatorname{nextE}\left(T_{i}^{j}, t_{c}\right)=0$. In this case, the rules are applied as stated, except that $T_{i}^{j+1}$ is not released.)

As was the case for reweighting under GEDF, the choice of which rule to apply depends on whether deviance is positive or negative. If positive, then we say that $T_{i}$ is positive-changeable at time $t_{c}$ from a desired weight of Ow to Nw ; otherwise $T_{i}$ is negative-changeable at time $t_{c}$ from a desired weight of Ow to Nw . Because $T_{i}$ initiates its desired weight change at time $t_{c}, \operatorname{Dwt}\left(T_{i}, t_{c}\right)=\mathrm{Nw}$ holds; however, $T_{i}$ 's desired scheduling weight does not change until the desired weight change has been enacted, as specified in the rules below. Note that, if $t_{c}$ occurs between the initiation and enaction of a previous reweighting event of $T_{i}$, then the previous event is canceled, i.e., treated as if it had not occurred. As discussed later, any "error" associated with canceling a reweighting event like this is accounted for when determining drift (formally defined in Section 4.9).

Rule P: If $T_{i}$ is positive-changeable at time $t_{c}$ from a desired weight of Ow to Nw , then one
of two actions is taken: (i) if $\frac{\operatorname{Irem}\left(T_{i}^{j}, t_{c}\right)}{\mathrm{Ow}}>\frac{\mathrm{REM}\left(T_{i}^{j}, t_{c}\right)}{\mathrm{Nw}}$, then immediately, $T_{i}^{j}$ is halted, the desired weight change is enacted, a new job $T_{i}^{j+1}$ with an execution time of next $\mathrm{E}\left(T_{i}^{j}, t_{c}\right)$ is released (if next $\mathrm{E}\left(T_{i}^{j}, t_{c}\right)>0$ ), and $T_{i}^{j}$ becomes inactive; (ii) otherwise, at time $\mathrm{d}\left(T_{i}^{j}\right)$, the desired weight change is enacted, i.e., the desired scheduling weight of $T_{i}$ does not change until the end of its current job.

Rule $\mathbf{N}$ : If $T_{i}$ is negative-changeable at time $t_{c}$ from a desired weight of Ow to Nw , then one of two actions is taken: (i) if $\mathrm{Nw}>\mathrm{Ow}$, then immediately, $T_{i}^{j}$ is halted and its desired weight change is enacted, and at time $t_{r}$, a new job $T_{i}^{j+1}$ with an execution time of $\operatorname{next} \mathrm{E}\left(T_{i}^{j}, t_{c}\right)$ is released (if next $\mathrm{E}\left(T_{i}^{j}, t_{c}\right)>0$ ) and $T_{i}^{j}$ becomes inactive, where $t_{r}$ is the smallest time at or after $t_{c}$ such that $\operatorname{dev}\left(T_{i}^{j}, t_{r}\right)=0$ holds; (ii) otherwise, at time $t_{e}$, the desired weight change is enacted, a new job with an execution time of nextE $\left(T_{i}^{j}, t_{c}\right)$ is released (if next $\mathrm{E}\left(T_{i}^{j}, t_{c}\right)>0$ ), and $T_{i}^{j}$ becomes inactive, where $t_{e}=\min \left(t_{r}, \mathrm{~d}\left(T_{i}^{j}\right)\right)$, and $t_{r}$ is smallest time at or after $t_{c}$ such that $\operatorname{dev}\left(T_{i}^{j}, t_{r}\right)=0$ holds.

Intuitively, Rule P changes a task's desired weight by halting its current job and issuing a new job with an execution time of next $\mathrm{E}\left(T_{i}^{j}, t_{c}\right)$ with the new desired weight if doing so would improve its scheduling priority. Notice that, at time $t$, job $T_{i}^{j}$ has a higher scheduling priority than job $T_{\ell}^{w}$ if

$$
\frac{\operatorname{Irem}\left(T_{i}^{j}, t_{c}\right)}{\operatorname{SGwt}\left(T_{i}, t_{c}\right)}<\frac{\operatorname{Irem}\left(T_{\ell}^{w}, t_{c}\right)}{\operatorname{SGwt}\left(T_{\ell}, t_{c}\right)}
$$

Hence, it is not difficult to show that if $\frac{\operatorname{Irem}\left(T_{i}^{j}, t_{c}\right)}{\mathrm{Ow}}>\frac{\mathrm{REM}\left(T_{i}^{j}, t_{c}\right)}{\mathrm{Nw}}$ holds, then halting $T_{i}^{j}$ and issuing a new job with an execution time of next $\mathrm{E}\left(T_{i}^{j}, t_{c}\right)$ would improve $T_{i}$ 's scheduling priority.

Example (Figure 4.5). Consider the example of Case (i) of Rule P illustrated in Figure 4.5, which depicts one processor that is assigned four tasks (where the execution cost of each job is one): $T_{1}$, which has $\operatorname{Dwt}\left(T_{1}\right)=1 / 2$ and leaves at time at time $2 ; T_{2}$ and $T_{3}$, both of which have a desired weight of $1 / 6$; and $T_{4}$, which has an initial desired weight of $1 / 6$ that increases to $4 / 6$ at time 2 . In this system, $T_{4}$ initially has the lowest scheduling priority (there is a deadline tie). Inset (a) depicts the PEDF schedule. Inset (b) depicts $T_{4}$ 's allocations in the SW schedule as well three other schedules, namely, the CSW, IDEAL, and PT schedules, which
are defined later in Sections 4.7.1 and 4.8. Since $T_{4}$ is not scheduled by time 2, it has positive deviance and changes its weight via Rule P, causing $T_{4}^{1}$ to be halted, $T_{4}^{2}$ to be released at time 2 with a deadline of $7 / 2$, and $T_{4}$ 's drift to become $2 / 6$. Note that halting $T_{4}$ 's current job and issuing a new job with an execution time of one improves $T_{4}$ 's scheduling priority, i.e., $\frac{\operatorname{Irem}\left(T_{4}^{1}, 2\right)}{\mathrm{Ow}}=\frac{4 / 6}{1 / 6}=4>6 / 4=\frac{1}{4 / 6}=\frac{\operatorname{REM}\left(T_{4}^{1}, 2\right)}{\mathrm{Nw}}$. Notice also that the third job of $T_{4}$ is issued $6 / 4$ time units after time 2 . This spacing is in keeping with a new job of desired weight $4 / 6$ issued at time 2.

Example (Figure 4.6). Consider the example of Case (ii) of Rule P illustrated in Figure 4.6, which depicts one processor that is assigned three tasks (where the execution cost of each job is one): $T_{1}$, which has $\operatorname{Dwt}\left(T_{1}\right)=1 / 3 ; T_{2}$, which has $\operatorname{Dwt}\left(T_{2}\right)=1 / 4$; and $T_{3}$, which has an initial desired weight of $1 / 4$ that initiates an increase to $1 / 3$ at time 2 . Inset (a) depicts the PEDF schedule. Inset (b) depicts $T_{3}$ 's allocations in the SW schedule. (Again, the CSW, IDEAL, and PT scheduling algorithms are defined later in Sections 4.7.1 and 4.8.) Since $T_{1}^{1}$ has not been scheduled by time 2, its deviance is positive; furthermore, since $\frac{\operatorname{Irem}\left(T_{3}^{1}, 2\right)}{\mathrm{Ow}}=\frac{.50}{.25}=$ $2<3=\frac{1}{1 / 3}=\frac{\operatorname{REm}\left(T_{3}^{1}, 2\right)}{N_{w}}, T_{1}$ enacts its weight change via Case (ii) of Rule P. Notice that, if $T_{3}^{1}$ had been halted at time 2 and released a new job of desired weight $1 / 3$, its scheduling priority would be decreased. Thus, if we were to enact the change via Case (i) of Rule P, then we would in effect be increasing the deadline of the first scheduled job of $T_{3}$, even though the desired weight of the task increased.

Rule N changes the desired weight of a task by one of two approaches: (i) if a task increases its desired weight, then Rule N causes the release time of its next job to be adjusted so that it is commensurate with the new desired weight; (ii) if a task decreases its desired weight, then Rule N causes the next job to be issued with a deadline that is commensurate with the new desired weight at the end of the current job.

Example (Figure 4.7). Consider the example of Case (i) of Rule N illustrated in Figure 4.7, which depicts the same system as in Figure 4.5 , except that $T_{4}$ has the highest priority. Inset (a) depicts the PEDF schedule. Inset (b) depicts $T_{3}$ 's allocations in the SW schedule. (Again, the CSW, IDEAL, and PT scheduling algorithms are defined later in Sections 4.7.1


Figure 4.5: An illustration of reweighting via Case (i) of Rule P under PEDF. (a) The PEDF schedule for a one-processor systems with four tasks. ( $T_{4}$ has the lowest scheduling priority.) (b) $T_{4}$ 's allocations in the IDEAL, CSW, PT, and SW schedules.


Figure 4.6: An illustration of reweighting via Case (ii) of Rule P under PEDF. (a) The PEDF schedule for a one-processor systems with four tasks. (b) $T_{3}$ 's allocations in the IDEAL, CSW, PT, and SW schedules.
and 4.8.) Notice that the second job of $T_{4}$ is released at time 3 , which is the time such that $\operatorname{dev}\left(T_{4}, 3\right)=\int_{0}^{3} \operatorname{SGwt}\left(T_{i}, u\right) d u-\mathrm{A}\left(\mathcal{S}, T_{4}, 0,3\right)=1-1=0$.

Example (Figure 4.8). Consider the example of Case (ii) of Rule N illustrated in Figure 4.8, which depicts a one-processor systems with four tasks (where the execution cost of each job is one): $T_{1}$, which has $\operatorname{Dwt}\left(T_{1}\right)=1 / 2$ and joins the system at time $2 ; T_{2}$ and $T_{3}$, both of which have a desired weight of $1 / 6$; and $T_{4}$, which has an initial desired weight of $1 / 2$ that initiates a weight decrease to $1 / 6$ at time 1 that is enacted at time 2 . Inset (a) depicts the PEDF schedule. Inset (b) depicts $T_{4}$ 's allocations in the SW schedule. (Again, the CSW, IDEAL, and PT scheduling algorithms are defined later in Sections 4.7.1 and 4.8.) Since $T_{4}$ has negative deviance at time 1 , and it decreases its desired weight, this change is enacted via Case (ii) of Rule N , causing $T_{4}$ 's next job to have a deadline of 8 and $T_{4}$ to have a drift of $-1 / 3$.

It is important to remember that when the Rules P and N halt a job, they do not abandon the computation that the job was performing. Rather, these rules split that computation across two jobs. Since these rules change the ordering of a task in the priority queues that determine scheduling, the time complexity for reweighting one task is $O(\log N)$, where $N$ is the number of tasks in the system (assuming that binomial heaps are used to implement the priority queues).

Canceled reweighting events. We now introduce a property about the relationship between the initiation and enactment of a desired weight change in the case that some such changes are canceled due to later desired weight changes. Notice that, once a task $T_{i}$ initiates a desired weight change at $t_{c}$, this desired weight change is eventually either canceled by another desired weight change or enacted. Further, Rules P and N enact any non-canceled desired weight change no later than the deadline of the last-released job $T_{i}^{j}$ of $T_{i}$ at $t_{c}$ (if it exists and if $\left.t_{c} \leq \mathrm{d}\left(T_{i}^{j}\right)\right)$.

Example (Figure 4.9). Consider the example in Figure 4.9, which depicts one processor that is assigned three tasks: $T_{1}$ and $T_{2}$, both of which have an execution time of 2 and a weight of $1 / 3$; and $T_{3}$, which has $\mathrm{e}\left(T_{3}\right)=2$ and an initial weight of $1 / 3$ that changes to $1 / 10$


Figure 4.7: An illustration of reweighting via Case (i) of Rule N under PEDF. (a) The PEDF schedule for a one-processor systems with four tasks. ( $T_{4}$ has the highest scheduling priority.) (b) $T_{4}$ 's allocations in the IDEAL, CSW, PT, and SW schedules.


Figure 4.8: An illustration of reweighting via Case (ii) of Rule N under PEDF. (a) The PEDF schedule for a one-processor systems with four tasks. (b) $T_{4}$ 's allocations in the IDEAL, CSW, PT, and SW schedules.


Figure 4.9: A one-processor example of canceling a reweighting event.
at time 3 via Case (ii) of Rule $N$ and then to $1 / 4$ at time 5 via Case (ii) of Rule N. Notice that, because the change initiated at time 3 is via Case (ii) of Rule N , the change is not enacted until time 6. As a result, when a change is initiated at time 5 , this new change cancels the previous change. Even though the change initiated at time 3 is canceled, the time of the next weight enactment is still at time 6 .

From this example, we can see that, once a desired weight change has been initiated during an active job, some desired weight change will be enacted by the earlier of the deadline of that job or when the job becomes inactive (which may be earlier, by Rules P and N). Property (X) formalizes this idea.
(X) If a task $T_{i}$ initiates a desired weight change at time $t_{c}$ and the job $T_{i}^{j}$ is active at $t_{c}$, then some desired weight change is enacted according to Rule P or N by either $\mathrm{d}\left(T_{i}^{j}\right)$ or when $T_{i}^{j}$ becomes inactive, whichever is first.

### 4.5.2 Modifications for NP-PEDF

In order to adapt Rules P and N to work for NP-PEDF, the only modification we need to make is when these rules are initiated. If a task with an active job changes its desired weight before or after that job has been scheduled, then Rules P and N are initiated as before. (Note that, if the active job has not been scheduled, then its deviance is positive, and if the active


Figure 4.10: A one-processor example of reweighting under NP-PEDF. (a) A reweighting event is initiated before $T_{3}^{1}$ is scheduled. ( $T_{3}^{1}$ has a lower scheduling priority than $T_{2}^{1}$.) (b) A reweighting event is initiated while $T_{3}^{1}$ is being scheduled. ( $T_{3}^{1}$ has a lower scheduling priority than $T_{2}^{1}$.)
job has been scheduled, then its deviance is negative.) However, if a task changes its desired weight while the active job $T_{i}^{j}$ is executing, then the initiation of the desired weight change is delayed until $T_{i}^{j}$ has completed or $T_{i}^{j}$ is no longer active, whichever is first. Note that, if a task $T_{i}$ changes its desired weight from Ow to Nw at time $t_{c}$ in NP-PEDF, then $\operatorname{Dwt}\left(T_{i}, t_{c}\right)=\mathrm{Nw}$ holds, regardless of whether the initiation of Rule P or N must be delayed.

Example (Figure 4.10). Consider the example in Figure 4.10, which depicts one processor that is assigned three tasks: $T_{1}$, which has $\mathrm{e}\left(T_{1}\right)=1, \operatorname{Dwt}\left(T_{1}\right)=1 / 2$, and leaves at time $2 ; T_{2}$, which has $\mathrm{e}\left(T_{2}\right)=1$ and $\operatorname{Dwt}\left(T_{2}\right)=1 / 6$; and $T_{3}$, which has $\mathrm{e}\left(T_{3}\right)=2$ and an initial desired weight of $1 / 3$ and initiates an increases to $4 / 6$ at time 2 . Inset (a) depicts the scenario where $T_{3}^{1}$ has the lowest scheduling priority (there is a deadline tie). Since $T_{3}$ is not scheduled by time 2, it has positive deviance and changes its weight via Rule P , causing $T_{3}^{1}$ to be halted, and $T_{3}^{2}$ to be released at time 2 with a deadline of 5 . Inset (b) depicts the same scenario as in (a) except that $T_{3}$ has higher priority than $T_{2}$. Since $T_{3}$ is scheduled at time 2 , and the system is schedule by NP-PEDF, the initiation of the reweighting event is delayed until $T_{3}$ stops executing at time 3 . Since $T_{3}^{1}$ is complete by time 2 , it has negative deviance and changes its weight via Rule N , causing its next job to have a release time of $9 / 2$.

### 4.6 Resetting Rules

In this section, we formally define the rules for resetting a system. Let $\tau$ be a task system in which the system is reset at time $t_{c}$. Let $\mathcal{S}$ be an $M$-processor PEDF schedule of $\tau$. Let $T_{i}^{j}$ be the last-released job (if any) of some task $T_{i}$ before $t_{c}$. Let $P_{[q]}$ and $P_{[\ell]}$ be, respectively, the processor that $T_{i}$ is assigned to before and after the system is reset at time $t_{c}$. Then, there are two possibilities. First, $T_{i}^{j}$ exists and is active at $t_{c}$, (e.g., $\mathrm{r}\left(T_{i}^{j}\right) \leq t_{c}<\mathrm{d}\left(T_{i}^{j}\right)$ ). Second, $T_{i}^{j}$ either is inactive at $t_{c}$ or does not exist (because $T_{i}^{1}$ is not released until after $t_{c}$ ). Thus, we have the following rule, with two cases, for resetting a task

Rule R: (i) If $T_{i}^{j}$ exists and is active at $t_{c}$, then $T_{i}^{j}$ is halted, and $T_{i}^{j+1}$ is released immediately on $P_{[\ell]}$, with an execution time of $\operatorname{nextE}\left(T_{i}^{j}, t_{c}\right)$. (ii) If $T_{i}^{j}$ either does not exist or is inactive at $t_{c}$, then the next job of $T_{i}$ is released at $t_{c}+\theta\left(T_{i}^{j+1}\right)$ on $P_{[\ell]}$ (or at $t_{c}+\theta\left(T_{i}^{1}\right)$, if $T_{i}^{j}$ does not exist).

Notice that, by this rule, whenever a system is reset, all active jobs are halted, and new jobs are released for the newly repartitioned system. Thus, if some task $T_{a}$ is assigned to $P_{[z]}$ when $T_{a}^{b}$ is released, then $T_{a}^{b}$ will only be scheduled on $P_{[z]}$.

Example (Figure 4.11). Consider the example in Figure 4.11, which depicts a two-processor system that is assigned four tasks: $T_{1}$, which has $\mathrm{e}\left(T_{1}\right)=1$ and an initial desired weight of $1 / 2$ that decreases to $1 / 4$ at time $2 ; T_{2}$, which has $e\left(T_{2}\right)=2$ and $\operatorname{Dwt}\left(T_{2}\right)=1 / 2 ; T_{3}$, which has $\mathrm{e}\left(T_{3}\right)=3$ and an initial desired weight of $1 / 2$ that increases to $3 / 4$ at time 2 ; and $T_{4}$, which has $\mathrm{e}\left(T_{4}\right)=4$ and $\operatorname{Dwt}\left(T_{4}\right)=1 / 2$. Initially, $T_{1}$ and $T_{2}$ are assigned to the first processor and $T_{3}$ and $T_{4}$ are assigned to the second. Inset (a) depicts the PEDF schedule (the other insets are considered later). At time 2, the system is repartitioned, i.e., is reset, and $T_{1}$ and $T_{3}$ are assigned to one processor and $T_{2}$ and $T_{4}$ are assigned to the other processor.

Complications with NP-PEDF. Under NP-PEDF, repartitioning is slightly more complex since a job cannot be preempted if it is currently being scheduled. Thus, the entire system cannot be reset at the same time. Thus, there are two viable options for reseting an NP-PEDFscheduled system. First, whenever the system is reset, ignore non-preemptive behavior, and


Figure 4.11: (a) An illustration of resetting in PEDF. Insets (b)-(e) depict, respectively, the allocations for $T_{1}, \ldots, T_{4}$.

| Property | Definition |
| :---: | :---: |
| $\begin{gathered} \mathbf{( S W - 1 )} \\ {[\text { Page 127] }} \end{gathered}$ | If the system is not reset over the range $\left(t_{1}, t_{2}\right]$, then the function $\mathrm{A}\left(\mathcal{C} S W, T_{i}, 0, t\right)$ is continuous for any $t \in\left[t_{1}, t_{2}\right]$. |
| $\begin{gathered} (\text { SW-2) } \\ {[\text { Page 127] }} \end{gathered}$ | For any job $T_{i}^{j}$ and any $t \geq \mathrm{r}\left(T_{i}^{j}\right), \mathrm{A}\left(\mathcal{C} S W, T_{i}^{j}, \mathrm{r}\left(T_{i}^{j}\right), t\right) \leq \mathrm{Ae}\left(T_{i}^{j}\right) \leq \mathrm{e}\left(T_{i}^{j}\right)$. |
| $\begin{gathered} \text { (SW-3) } \\ {[\text { Page 127] }} \end{gathered}$ | Any job $T_{i}^{j}$ is complete in the CSW schedule by time $\mathrm{d}\left(T_{i}^{j}\right)$. |
| $\begin{gathered} (\mathrm{V}) \\ {[\text { Page } 128]} \end{gathered}$ | For the jobs $T_{i}^{j}$ and $T_{i}^{j+1}$, if $\mathrm{d}\left(T_{i}^{j}\right)>\mathrm{r}\left(T_{i}^{j+1}\right)$, then $T_{i}^{j}$ is complete in both the CSW and actual schedules by time $\mathrm{r}\left(T_{i}^{j+1}\right)$. |
| $\begin{gathered} (\mathbf{L - 1}) \\ {[\text { Page 129] }} \end{gathered}$ | For any time $t \leq \mathrm{r}\left(T_{i}^{j}\right), \operatorname{lag}\left(T_{i}^{j}, t\right)=0$. |
| $\begin{gathered} (\mathbf{L - 2 )} \\ {[\text { Page 129] }} \\ \hline \end{gathered}$ | If $\operatorname{lag}\left(T_{i}^{j}, t\right)>0$ for some time $t \geq \mathrm{d}\left(T_{i}^{j}\right)$, then the value of $\operatorname{lag}\left(T_{i}^{j}, t\right)$ denotes the amount of time remaining to be scheduled for $T_{i}^{j}$ after time $t$. |
| $\begin{gathered} (\mathbf{L}-3) \\ {[\text { Page 129] }} \end{gathered}$ | If the system is not reset over the time range $\left(t_{1}, t_{2}\right]$, then the function $\operatorname{lag}\left(T_{i}^{j}, t\right)$ is continuous over $\left[t_{1}, t_{2}\right]$. |
| $\begin{gathered} (\mathbf{H}-\mathbf{1}) \\ {[\text { Page 130] }} \end{gathered}$ | If a job $T_{a}^{b}$ has a scheduling priority at least $T_{i}^{j}$ 's and is assigned to the same processor as $T_{i}^{j}$, then $T_{a}^{b} \in \mathrm{H}\left(T_{i}^{j}, t\right)$ for every value of $t \in\left[\mathrm{r}\left(T_{a}^{b}\right), t_{e}\right)$, where $t_{e}$ is the first time at which $T_{a}^{b}$ is both inactive and not pending. |
| $\begin{gathered} (\mathbf{H}-2) \\ {[\text { Page 130 }]} \\ \hline \end{gathered}$ | If the system is not reset over the range $\left(t_{1}, t_{2}\right]$, then for any job, $T_{i}^{j}$ the function $\operatorname{LAG}\left(\mathrm{H}\left(T_{i}^{j}, t\right), t\right)$ is continuous over the time range $\left[t_{1}, t_{2}\right]$. |

Table 4.4: Summary of properties used in Section 4.7.
immediately halt all active tasks. Second, apply Rule R to all tasks that are not currently running when the system reset at $t_{c}$, and migrate the tasks that are scheduled at $t_{c}$ when they complete. Notice that the first option can only be used in systems where non-preemptive behavior is desirable but not required. The second option is more complex to implement than the first option, but is the only viable option when non-preemptive behavior is required.

### 4.7 Scheduling Correctness

In this section, we prove that no job misses a deadline under our adaptive PEDF scheduling algorithm and that jobs have bounded tardiness under our adaptive NP-PEDF scheduling algorithm. The properties used throughout this section are summarized in Table 4.4.

### 4.7.1 The CSW Algorithm

Scheduling correctness is established by considering the clairvoyant scheduling-weight (CSW) scheduling algorithm. Under it, at each instant $t$, each job of each task $T_{i}$ that is both active and incomplete (in the CSW schedule) is allocated a fraction of a processor equal to

SGwt $\left(T_{i}, t\right)$. We consider CSW to be "clairvoyant" in the sense that CSW does not allocate a job $T_{i}^{j}$ more than $\operatorname{Ae}\left(T_{i}^{j}\right)$ time. More specifically, for any schedule $\mathcal{C} S W$ under CSW of any task system $\tau$, we say that $T_{i}^{j}$ has completed by time $t$ in $\mathcal{C} S W$ iff $T_{i}^{j}$ has executed for $\operatorname{Ae}\left(T_{i}^{j}\right)$ by $t$. For example, in Figure 4.5, $T_{4}^{1}$ receives no allocations in the CSW scheduling algorithm since $\operatorname{Ae}\left(T_{4}^{1}\right)=0$. In addition, if $\mathcal{S}$ is the actual schedule for a task system $\tau, \mathcal{C} S W$ is the CSW schedule of $\tau$, the system is reset (re-partitioned) at some time $t_{r}$, and some task $T_{i}$ has received $X$ more time units of allocation in $\mathcal{S}$ than in $\mathcal{C} S W$, then for any $t \geq t_{r}$, the value $X$ is added to $\mathrm{A}\left(\mathcal{C} S W, T_{i}, 0, t\right)$. For example, in Figure 4.11, when the system is reset at time $2, T_{3}$ is allocated more capacity in the actual schedule than in the CSW schedule. As a result, $T_{3}$ 's allocations in the CSW schedule jumps at time 2 from 1 to 2 .

From Figure 4.11, we can see that the following property holds:
(SW-1) If the system is not reset over the range $\left(t_{1}, t_{2}\right]$, then the function $\mathrm{A}\left(\mathcal{C} S W, T_{i}, 0, t\right)$ is continuous for any $t \in\left[t_{1}, t_{2}\right]$.

By the definition of the CSW scheduling algorithm, we have the following property:
(SW-2) For any job $T_{i}^{j}$ and any two times $t_{a}$ and $t_{b}$, where $r\left(T_{i}^{j}\right) \leq t_{a} \leq t_{b}$,

$$
\mathrm{A}\left(\mathcal{C} S W, T_{i}^{j}, t_{a}, t_{b}\right) \leq \mathrm{Ae}\left(T_{i}^{j}\right) \leq \mathrm{e}\left(T_{i}^{j}\right)
$$

In addition, by the definition of $\mathrm{d}\left(T_{i}^{j}\right)$, it is not difficult to see that the following property holds:
(SW-3) Any job $T_{i}^{j}$ is complete in the CSW schedule by time $\mathrm{d}\left(T_{i}^{j}\right)$.
Throughout this chapter, we use $\mathcal{C} S W$ to denote the CSW schedule of a task system $\tau$.
Notice that, since the CSW schedule is identical to the SW schedule, except for jobs that halt, which receive a smaller allocation in the CSW schedule, it follows that for any job $T_{i}^{j}$ and times $t_{1}$ and $t_{2}$, where $t_{1} \leq t_{2}$, the following property holds:

$$
\begin{equation*}
\mathrm{A}\left(\mathcal{C} S W, T_{i}^{j}, t_{1}, t_{2}\right) \leq \mathrm{A}\left(\mathcal{S} W, T_{i}^{j}, t_{1}, t_{2}\right) \tag{4.14}
\end{equation*}
$$

Moreover, since for the same times $t_{1}$ and $t_{2}, \mathrm{~A}\left(\mathcal{S} W, T_{i}^{j}, t_{1}, t_{2}\right)=\int_{t_{1}}^{t_{2}} \operatorname{SGwt}\left(T_{i}, u\right) d u$, the reweighting rules allow at most one job of a task to be active at any given point in time $t$,
and by the Property $(\mathrm{W}), \sum_{T_{i}^{j} \in \operatorname{ASSN}\left(P_{[q]}, t\right)} \operatorname{SGwt}\left(T_{i}, t\right) \leq 1$, holds for any job, at any time $t$, it follows that is not difficult to prove the following lemma.

Lemma 4.1. For any time interval $\left[t_{1}, t_{2}\right]$ on any processor $P_{[q]}$,

$$
\sum_{T_{i}^{j} \in \operatorname{ASSN}\left(P_{[q]}, t\right)} \mathrm{A}\left(\mathcal{C S W}, T_{i}^{j}, t_{1}, t_{2}\right) \leq \sum_{T_{i}^{j} \in \operatorname{ASSN}\left(P_{[q]}, t\right)} \mathrm{A}\left(\mathcal{S} W, T_{i}^{j}, t_{1}, t_{2}\right) \leq t_{2}-t_{1}
$$

### 4.7.2 Overlap

One consequence of the reweighting rules is that for some job, $T_{i}^{j}$, it is possible that $\mathrm{r}\left(T_{i}^{j+1}\right)<$ $\mathrm{d}\left(T_{i}^{j}\right)$. Such a scenario can arise through one of four possibilities: at some time $t, T_{i}$ changes its desired weight via Case (i) of Rule P; $T_{i}$ changes its desired weight via Case (i) of Rule N; $T_{i}$ changes its desired weight via Case (ii) of Rule N ; or $T_{i}$ is reset via Case (i) of Rule R. In all four of these scenarios, $T_{i}^{j}$ is halted (and thus completes in the actual schedule) at time $t$. Moreover, in all three cases, $T_{i}^{j}$ is complete in the CSW schedule by time $\mathrm{r}\left(T_{i}^{j+1}\right) \geq t$. (In Case (i) of Rule P, $T_{i}^{j}$ has positive deviance, which guarantees that in the CSW schedule $T_{i}^{j}$ has received an allocation of $\operatorname{Ae}\left(T_{i}^{j}\right)$ by time $t$. In Case (i) of Rule $\mathrm{N}, T_{i}^{j+1}$ is released when its deviance is zero, which guarantees that in the CSW schedule $T_{i}^{j}$ has received an allocation of $\operatorname{Ae}\left(T_{i}^{j}\right)$ by time $r\left(T_{i}^{j+1}\right)$. In Case (i) of Rule R, if $T_{i}^{j}$ has a smaller allocation in the CSW schedule than the actual schedule, then this difference is added to $T_{i}^{j}$,s allocation at time $t$, which again guarantees that in the CSW schedule $T_{i}^{j}$ has received an allocation of $\mathrm{Ae}\left(T_{i}^{j}\right)$ by time $\mathrm{r}\left(T_{i}^{j+1}\right)$.) Thus, we have the following property:
(V) For the jobs $T_{i}^{j}$ and $T_{i}^{j+1}$, if $\mathrm{d}\left(T_{i}^{j}\right)>\mathrm{r}\left(T_{i}^{j+1}\right)$, then $T_{i}^{j}$ is complete in both the CSW and actual schedules by time $\mathrm{r}\left(T_{i}^{j+1}\right)$.

### 4.7.3 Lag

Before continuing, we introduce the notion of lag, which represents the difference between the actual allocations to a task and its CSW allocations. Formally, the lag of a job at time $t$ is
given by

$$
\begin{equation*}
\operatorname{lag}\left(T_{i}^{j}, t\right)=\mathrm{A}\left(\mathcal{C} S W, T_{i}^{j}, \mathrm{r}\left(T_{i}^{j}\right), t\right)-\mathrm{A}\left(\mathcal{S}, T_{i}^{j}, \mathrm{r}\left(T_{i}^{j}\right), t\right) \tag{4.15}
\end{equation*}
$$

where $\mathcal{C} S W$ is the CSW schedule for the system and $\mathcal{S}$ is the actual schedule for the system. Notice that lag is similar to deviance, except that while deviance relates the actual schedule to the SW schedule, lag relates the actual schedule to the CSW schedule. For example, in Figure 4.5, $\operatorname{lag}\left(T_{4}^{1}, 2\right)=1 / 3$ and $\operatorname{lag}\left(T_{4}^{1}, 3\right)=-1 / 6$.

We now present a few simple properties about lag will be used in the following proofs.
(L-1) For any time $t \leq \mathrm{r}\left(T_{i}^{j}\right), \operatorname{lag}\left(T_{i}^{j}, t\right)=0$.
(L-2) If $\operatorname{lag}\left(T_{i}^{j}, t\right)>0$ for some time $t \geq \mathrm{d}\left(T_{i}^{j}\right)$, then the value of $\operatorname{lag}\left(T_{i}^{j}, t\right)$ denotes the amount of time remaining to be scheduled for $T_{i}^{j}$ after time $t$.
(L-3) If the system is not reset over the time range $\left(t_{1}, t_{2}\right]$, then the function $\operatorname{lag}\left(T_{i}^{j}, t\right)$ is continuous over $\left[t_{1}, t_{2}\right]$.

Property (L-1) is trivially true, since for any $t \leq \mathrm{r}\left(T_{i}^{j}\right), \mathrm{A}\left(\mathcal{C} S W, T_{i}^{j}, \mathrm{r}\left(T_{i}^{j}\right), t\right)=$ $\mathrm{A}\left(\mathcal{S}, T_{i}^{j}, \mathrm{r}\left(T_{i}^{j}\right), t\right)=0$. Property (L-2) holds since by time $t \geq \mathrm{d}\left(T_{i}^{j}\right), \mathrm{A}\left(\mathcal{C} S W, T_{i}^{j}, \mathrm{r}\left(T_{i}^{j}\right), t\right)=$ $\mathrm{Ae}\left(T_{i}^{j}\right)$ and $\mathrm{Ae}\left(T_{i}^{j}\right)-\mathrm{A}\left(\mathcal{S}, T_{i}^{j}, \mathrm{r}\left(T_{i}^{j}\right), t\right)$ denotes the amount of time remaining to be scheduled for $T_{i}^{j}$ at and after $t$. Property (L-3) holds because as long as the system is not reset over the range $\left(t_{1}, t_{2}\right]$, both $\mathrm{A}\left(\mathcal{S}, T_{i}^{j}, \mathrm{r}\left(T_{i}^{j}\right), t\right)$ and $\mathrm{A}\left(\mathcal{C} S W, T_{i}^{j}, \mathrm{r}\left(T_{i}^{j}\right), t\right)$ are continuous for any $t \in\left[t_{1}, t_{2}\right]$. (Notice that the function $\mathrm{A}\left(\mathcal{S}, T_{i}^{j}, \mathrm{r}\left(T_{i}^{j}\right), t\right)$ is continuous by definition and the function $\mathrm{A}\left(\mathcal{C} S W, T_{i}^{j}, \mathrm{r}\left(T_{i}^{j}\right), t\right)$ is continuous by Property (SW-1).)

For brevity, we use $\operatorname{LAG}(G, t)$ to denote

$$
\operatorname{LAG}(G, t)=\sum_{T_{i}^{j} \in G} \operatorname{lag}\left(T_{i}^{j}, t\right),
$$

where $G$ is some set of jobs.
We now prove a simple property about lag in PEDF schedules.

Lemma 4.2. If no deadline of a job of task $T_{i}$ has been missed by time $t$ and the job $T_{i}^{j}$ is not pending at time $t$, then $\operatorname{lag}\left(T_{i}^{j}, t\right) \leq 0$.

Proof. Let $T_{i}^{j}$ be as defined in the lemma, i.e., $T_{i}^{j}$ is not pending at time $t$. Since all previous jobs of $T_{i}$ have completed before their deadlines, by the definition of pending, either $t<\mathrm{r}\left(T_{i}^{j}\right)$ or $t \geq \mathrm{r}\left(T_{i}^{j}\right) \wedge \mathrm{A}\left(\mathcal{S}, T_{i}^{j}, \mathrm{r}\left(T_{i}^{j}\right), t\right)=\mathrm{Ae}\left(T_{i}^{j}\right)$ holds. By Property $(\mathrm{L}-1)$, if $t<\mathrm{r}\left(T_{i}^{j}\right)$, then $\operatorname{lag}\left(T_{i}^{j}, t\right)=0$. If $t \geq \mathrm{r}\left(T_{i}^{j}\right) \wedge \mathrm{A}\left(\mathcal{S}, T_{i}^{j}, \mathrm{r}\left(T_{i}^{j}\right), t\right)=\mathrm{Ae}\left(T_{i}^{j}\right)$ holds, then by Property (SW2), $\mathrm{A}\left(\mathcal{S}, T_{i}^{j}, \mathrm{r}\left(T_{i}^{j}\right), t\right)=\mathrm{Ae}\left(T_{i}^{j}\right) \geq \mathrm{A}\left(\mathcal{C} S W, T_{i}^{j}, \mathrm{r}\left(T_{i}^{j}\right), t\right)$ holds. This, in turn, implies that $\operatorname{lag}\left(T_{i}^{j}, t\right) \leq 0$.

### 4.7.4 Higher-Priority Jobs

In the following proofs, we use the term $\mathrm{H}\left(T_{i}^{j}, t\right)$ to denote the set of jobs that are assigned to the same processor as $T_{i}^{j}$, have a scheduling priority higher than or equal to $T_{i}^{j}$, and are either pending or active at time $t$. For example, in Figure 4.3, $\mathrm{H}\left(T_{3}^{1}, 0\right)=\left\{T_{1}^{1}, T_{2}^{1}, T_{3}^{1}\right\}$, $\mathrm{H}\left(T_{3}^{1}, 3\right)=\left\{T_{1}^{2}, T_{2}^{1}, T_{3}^{1}\right\}$ and $\mathrm{H}\left(T_{3}^{1}, 7\right)=\left\{T_{3}^{1}\right\}$. Since, as we discussed in Section 4.4.2, the relative scheduling priority between any two jobs is time invariant, we have the following property:
(H-1) If a job $T_{a}^{b}$ has a scheduling priority at least $T_{i}^{j}$ 's and is assigned to the same processor as $T_{i}^{j}$, then $T_{a}^{b} \in \mathrm{H}\left(T_{i}^{j}, t\right)$ for every value of $t \in\left[\mathrm{r}\left(T_{a}^{b}\right), t_{e}\right)$, where $t_{e}$ is the first time at which $T_{a}^{b}$ is both inactive and not pending.

Notice that, if a job $T_{a}^{b}$ becomes an element of $\mathrm{H}\left(T_{i}^{j}, t\right)$ at some time $t$, then by Property (H-1) it must be that $T_{a}^{b}$ was released at time $t$, and thus, by Property $(\mathrm{L}-1), \operatorname{lag}\left(T_{a}^{b}, t\right)=0$ holds. Moreover, by Property (H-1), the only way a job $T_{a}^{b}$ can leave the set $\mathrm{H}\left(T_{i}^{j}, t\right)$ is if by time $t, T_{a}^{b}$ is both inactive and not pending. If $T_{a}^{b}$ is inactive at $t$, then $\mathrm{r}\left(T_{a}^{b+1}\right) \leq t$ or $\mathrm{d}\left(T_{a}^{b}\right) \leq t$ (or both). If $\mathrm{r}\left(T_{a}^{b+1}\right) \leq t$ holds, then by Property (V), $T_{a}^{b}$ is complete in both the actual and CSW schedules, in which case $\operatorname{lag}\left(T_{a}^{b}, t\right)=0$. If $\mathrm{d}\left(T_{a}^{b}\right) \leq t$, then since $T_{a}^{b}$ is not pending at time $t$, i.e., $\operatorname{Ae}\left(T_{a}^{b}\right)=\mathrm{A}\left(\mathcal{S}, T_{a}^{b}, \mathrm{r}\left(T_{a}^{b}\right), t\right)$, by Property $(\mathrm{L}-2), \operatorname{lag}\left(T_{a}^{b}, t\right)=0$. Thus, given that the lag of a job is zero when it either joins or leaves $\mathrm{H}\left(T_{i}^{j}, t\right)$ and that (by Property (L-3)) $\operatorname{lag}\left(T_{a}^{b}, t\right)$ is continuous so long as the system is not reset, we have the following property:
(H-2) If the system is not reset over the range $\left(t_{1}, t_{2}\right]$, then for any job $T_{i}^{j}$, the function $\operatorname{LAG}\left(\mathrm{H}\left(T_{i}^{j}, t\right), t\right)$ is continuous over the time range $\left[t_{1}, t_{2}\right]$.

### 4.7.5 PEDF Correctness

We now show that in a PEDF scheduled system, no job misses a deadline.

Theorem 4.2. Let $\tau$ be an adaptable sporadic task system. Then, for any job $T_{i}^{j}$ that is in $\tau, T_{i}^{j}$ completes by $\mathrm{d}\left(T_{i}^{j}\right)$ under PEDF.

Proof. Suppose, to derive a contradiction, that there exists some job $T_{i}^{j}$ in $\tau$ such that $T_{i}^{j}$ misses its deadline at time $t_{d}$. Without loss of generality, let $t_{d}$ be the first time that a job deadline is missed in $\tau$. Notice that, by Property (V), if $\mathrm{d}\left(T_{i}^{j-1}\right)>\mathrm{r}\left(T_{i}^{j}\right)$, then $T_{i}^{j-1}$ must be complete by $\mathrm{r}\left(T_{i}^{j}\right)$. Moreover, if $\mathrm{d}\left(T_{i}^{j-1}\right) \leq \mathrm{r}\left(T_{i}^{j}\right)<\mathrm{d}\left(T_{i}^{j}\right)=t_{d}$, then by the definition of $t_{d}$, $T_{i}^{j-1}$ is complete by $\mathrm{r}\left(T_{i}^{j}\right)$. Thus, in either case, $T_{i}^{j-1}$ is complete by $\mathrm{r}\left(T_{i}^{j}\right)$.

By Property (L-2), the value of $\operatorname{lag}\left(T_{a}^{b}, t\right)$ at any time $t \geq \mathrm{d}\left(T_{a}^{b}\right)$ represents $T_{a}^{b}$ 's remaining computation time at time $t$. Thus, if $T_{i}^{j}$ misses a deadline at time $t_{d}$, then $\operatorname{lag}\left(T_{i}^{j}, t_{d}\right)>0$ must hold. Moreover, since every job in $\operatorname{LAG}\left(\mathrm{H}\left(T_{i}^{j}, t_{d}\right), t_{d}\right)$ has a deadline at or before $t_{d}$, by Property (L-2), $\operatorname{LAG}\left(\mathrm{H}\left(T_{i}^{j}, t_{d}\right), t_{d}\right)$ represents the amount of computation remaining for all jobs in $\operatorname{LAG}\left(\mathrm{H}\left(T_{i}^{j}, t_{d}\right), t_{d}\right)$, including $T_{i}^{j}$. Thus, if $T_{i}^{j}$ misses a deadline at $t_{d}$, then $\operatorname{LAG}\left(\mathrm{H}\left(T_{i}^{j}, t_{d}\right), t_{d}\right)>0$ must hold. The objective of this proof is to show that $\operatorname{LAG}\left(\mathrm{H}\left(T_{i}^{j}, t_{d}\right), t_{d}\right) \leq 0$, which contradicts our assumption and completes the proof.

Before continuing, we introduce a few terms. Let $P_{[q]}$ denote the processor that $T_{i}^{j}$ is assigned to at time $t_{d}$. Let $\mathcal{S}$ denote the PEDF schedule of $\tau$. Let $t_{1}$ be the first time before $t_{d}$ such that over the range $\left(t_{1}, t_{d}\right]$ the system is not reset, and over the range $\left[t_{1}, t_{d}\right], P_{[q]}$ is continually scheduling jobs from $\mathrm{H}\left(T_{i}^{j}, t\right)$. Notice that $t_{1}$ must exist because $T_{i}^{j}$ is pending at time $t_{d}$, and since (as we established above) $T_{i}^{j-1}$ is complete by $\mathrm{r}\left(T_{i}^{j}\right)$, either $T_{i}^{j}$ or some higher-priority job is scheduled immediately before $t_{d}$. In addition, the system is not reset at $t_{d}$ because if it were, then $T_{i}^{j}$ would be halted at $t_{d}$ (and thus complete). Figure 4.12 illustrates both $t_{1}$ and $t_{d}$.

There are two remaining components of this proof. First, we show that $\operatorname{LAG}\left(\mathrm{H}\left(T_{i}^{j}, t_{1}\right), t_{1}\right) \leq 0$ holds. Second, we show that for any $t \in\left[t_{1}, t_{d}\right], \operatorname{LAG}\left(\mathrm{H}\left(T_{i}^{j}, t\right), t\right) \leq 0$ holds, which implies that $\operatorname{lag}\left(T_{i}^{j}, t_{d}\right) \leq 0$ holds.

Claim 4.1. $\operatorname{LAG}\left(\mathrm{H}\left(T_{i}^{j}, t_{1}\right), t_{1}\right) \leq 0$.


Figure 4.12: An illustration of time $t_{1}$ and $t_{d}$.

Proof. By Rule R, if the system is reset at time $t_{1}$ or $t_{1}=0$, then every job that is active at $t_{1}$ must also be released at $t_{1}$. Thus, by Property (L-1) the lag of every active job at $t_{1}$ equals 0 . This, in turn, implies that $\operatorname{LAG}\left(\mathrm{H}\left(T_{i}^{j}, t_{1}\right), t_{1}\right)=0$. Thus, for the remainder of the proof of Claim 4.1, we assume that $t_{1}>0$ and the system is not reset at time $t_{1}$.

In this case, by the definition of $t_{1}$, no job with a priority higher than or equal to $T_{i}^{j}$,s is scheduled $t_{1}$. Thus, there must exist some value $\epsilon_{1}>0$ such that for any time $t \in\left[t_{1}-\epsilon_{1}, t_{1}\right)$, either $P_{[q]}$ is idle or some job $T_{a}^{b}$ with a scheduling priority lower than any job in $\mathbf{H}\left(T_{i}^{j}, t\right)$ is scheduled on $P_{[q]}$. For example, in Figure 4.12, $T_{3}$ is scheduled over $\left[t_{1}-1, t_{1}\right)$. Regardless of whether $P_{[q]}$ is idle or scheduling a lower-priority job over the range $\left[t_{1}-\epsilon_{1}, t_{1}\right)$, no job in $\mathrm{H}\left(T_{i}^{j}, t\right)$ is pending over $t \in\left[t_{1}-\epsilon_{1}, t_{1}\right)$. For example, in Figure 4.12, neither $T_{1}$ or $T_{2}$ is pending over the range $\left[t_{1}-1, t_{1}\right)$. As a result, at time $t_{1}$ every job in $\mathrm{H}\left(T_{i}^{j}, t_{1}\right)$ is either released, i.e., becomes a pending job, or is active but not pending. Thus, by Property (L-1), Lemma 4.2, and because no deadline is missed by $t_{1}\left(\right.$ since $\left.t_{1}<t_{d}\right), \operatorname{LAG}\left(\mathrm{H}\left(T_{i}^{j}, t_{1}\right), t_{1}\right) \leq 0$.

Claim 4.2. For any $t \in\left[t_{1}, t_{d}\right], \operatorname{LAG}\left(\mathrm{H}\left(T_{i}^{j}, t\right), t\right) \leq 0$.

Proof. To derive a contradiction, we assume that there exists some $t \in\left[t_{1}, t_{d}\right]$ such that $\operatorname{LAG}\left(\mathrm{H}\left(T_{i}^{j}, t\right), t\right)>0$. Since the system is not reset over the time range $\left(t_{1}, t_{d}\right]$, by Property (H-2), it follows that $\operatorname{LAG}\left(\mathrm{H}\left(T_{i}^{j}, t\right), t\right)$ is continuous over the range $\left[t_{1}, t_{d}\right]$. As a result, by Claim 4.1 $\left(\operatorname{LAG}\left(\mathrm{H}\left(T_{i}^{j}, t_{1}\right), t_{1}\right) \leq 0\right)$ and our assumption that $\operatorname{LAG}\left(\mathrm{H}\left(T_{i}^{j}, t\right), t\right)>0$ holds, it follows, by the definition of continuity, that there exists a time $t_{z} \in\left[t_{1}, t_{d}\right)$ and a value


Figure 4.13: An illustration of time $t_{z}$ and $\epsilon_{z}$.
$\epsilon_{z} \in\left(0, t_{d}-t_{z}\right)$ such that the following conditions hold
(i) $\operatorname{LAG}\left(\mathrm{H}\left(T_{i}^{j}, t_{z}\right), t_{z}\right)=0$,
(ii) $\operatorname{LAG}\left(\mathrm{H}\left(T_{i}^{j}, t\right), t\right)$ is strictly monotonically increasing over the range $t \in\left[t_{z}, t_{z}+\epsilon_{z}\right)$,

Such a scenario is illustrated in Figure 4.13.
We now show that $t_{z}$ and $\epsilon_{z}$ cannot exist, which completes the proof of Claim 4.2. Since, by the definition of $t_{1}$, over the range $\left[t_{1}, t_{d}\right], P_{[q]}$ is continually scheduling jobs from $\mathrm{H}\left(T_{i}^{j}, t\right)$, it follows that some job from $\mathrm{H}\left(T_{i}^{j}, t\right)$ is scheduled for some range $\left(t_{3}, t_{4}\right)$ that is contained within $\left(t_{z}, t_{z}+\epsilon_{z}\right)$. Thus, by Lemma 4.1,

$$
\mathrm{A}\left(\mathcal{C} S W, \operatorname{ASSN}\left(P_{[q]}, t\right), t_{3}, t_{4}\right) \leq t_{4}-t_{3}=\mathrm{A}\left(\mathcal{S}, \mathrm{H}\left(T_{i}^{j}, t\right), t_{4}, t_{3}\right)
$$

Moreover, since $\mathrm{A}\left(\mathcal{C S W}, \mathrm{H}\left(T_{i}^{j}, t\right), t_{3}, t_{4}\right) \leq \mathrm{A}\left(\mathcal{C} S W, \operatorname{ASSN}\left(P_{[q]}, t\right), t_{3}, t_{4}\right)$ it follows that

$$
\mathrm{A}\left(\mathcal{C} S W, \mathrm{H}\left(T_{i}^{j}, t\right), t_{3}, t_{4}\right) \leq \mathrm{A}\left(\mathcal{C} S W, \operatorname{ASSN}\left(P_{[q]}, t\right), t_{3}, t_{4}\right) \leq \mathrm{A}\left(\mathcal{S}, \mathrm{H}\left(T_{i}^{j}, t\right), t_{4}, t_{3}\right)
$$

Thus,

$$
\operatorname{LAG}\left(\mathrm{H}\left(T_{i}^{j}, t\right), t_{4}\right)-\operatorname{LAG}\left(\mathrm{H}\left(T_{i}^{j}, t\right), t_{3}\right) \leq 0 .
$$

This, in turn, implies that over the range $\left(t_{3}, t_{4}\right) \operatorname{LAG}\left(\mathrm{H}\left(T_{i}^{j}, t\right), t\right)$ is not strictly monotonically increasing, which contradicts Property (ii) of $t_{z}$ and $\epsilon_{z}$. Thus, $t_{z}$ and $\epsilon_{z}$ cannot exist, which completes the proof of the claim.

Since Claim 4.2 implies that $\operatorname{LAG}\left(\mathrm{H}\left(T_{i}^{j}, t_{d}\right), t_{d}\right) \leq 0$ holds, it follows that at $t_{d}$ no computation time remains for any job in $\mathrm{H}\left(T_{i}^{j}, t_{d}\right)$, which includes $T_{i}^{j}$. Thus, $T_{i}^{j}$ is complete by its deadline.

### 4.7.6 NP-PEDF Correctness

In this section, we show that the tardiness of a task $T_{i}$ under NP-PEDF is at most $\mathcal{X}$, where $\mathcal{X}$ is the largest execution time of any task. Since the proof for tardiness bounds under NP-PEDF is similar to the scheduling correctness proof for PEDF, rather than repeat the entire proof, we state where they differ.

The primary difference between the proofs for NP-PEDF and PEDF is that in the proof of Theorem 4.2, if the system is not reset at $t_{1}$ and $t_{1}>0$, then we can guarantee that no job in $\mathrm{H}\left(T_{i}^{j}, t\right)$ is pending immediately before $t_{1}$. However, in proving the tardiness bounds for NP-PEDF this is not the case. For example, consider the scenario, depicted in Figure 4.14, where a job not in $\mathrm{H}\left(T_{i}^{j}, t\right)$ becomes pending immediately after a lower-priority job begins being scheduled. In this scenario, the job in $\mathrm{H}\left(T_{i}^{j}, t\right)$ must wait until the lower-priority job completes before it begins being scheduled. Since this delay can be up to $\mathcal{X}$ time units long, and since it is possible that the sum of the guaranteed weights of all tasks in $\mathrm{H}\left(T_{i}^{j}, t\right)$ may be close to one, it is possible that $\operatorname{LAG}\left(\mathrm{H}\left(T_{i}^{j}, t_{1}\right), t_{1}\right)$ may be close to $\mathcal{X}$. As a result, $\operatorname{LAG}\left(\mathrm{H}\left(T_{i}^{j}, t_{d}\right), t_{d}\right)$ may be close to $\mathcal{X}$, which implies that the amount of work remaining for all tasks in $\operatorname{LAG}\left(\mathrm{H}\left(T_{i}^{j}, t_{d}\right), t_{d}\right)$ is at most $\mathcal{X}$, which means that $T_{i}^{j}$ may miss its deadline by at most $\mathcal{X}$.

One final note: notice that if a task cannot be migrated immediately when the system is reset (because it is non-preemptable and is being scheduled), then this does not impact the correctness proof. The reason why is because delaying a task's migration does not cause the guaranteed weight of all tasks assigned to a processor to be larger than one.

From this discussion, we have the following theorem.


Figure 4.14: A one-processor example of an NP-PEDF system where two tasks are released after a lower-priority job begins executing.

Theorem 4.3. Let $\tau$ be an adaptable sporadic task system. Then, for any job $T_{i}^{j}$ of a task in $\tau, T_{i}^{j}$ has tardiness at most $\mathcal{X}$ under NP-PEDF, where $\mathcal{X}$ is the largest execution time of any job in $\tau$.

### 4.8 The IDEAL and PT Algorithms

In section, Section 4.9, we turn our attention to proving "drift" bounds. For this purpose, we introduce two new theoretical scheduling algorithms, namely the IDEAL and PT algorithms.

Under the IDEAL scheduling algorithm, at each instant $t$, each $\operatorname{task} T_{i}$ in $\tau$ with an active job at $t$ is allocated a fraction of the system equal to its guaranteed weight, $\operatorname{Gwt}\left(T_{i}, t\right)$. Hence, if $\mathcal{I}$ is the IDEAL schedule of $\tau$ and $T_{i}$ is active over the interval $\left[t_{1}, t_{2}\right)$, then over $\left[t_{1}, t_{2}\right)$, the task $T_{i}$, assigned to $P_{[q]}$, is allocated

$$
\begin{equation*}
\mathrm{A}\left(\mathcal{I}, T_{i}, t_{1}, t_{2}\right)=\int_{t_{1}}^{t_{2}} \operatorname{Gwt}\left(T_{i}, u\right) d u=\int_{t_{1}}^{t_{2}} \frac{\operatorname{Dwt}\left(T_{i}, u\right)}{\mathcal{T D}\left(P_{[q]}, u\right)} d u \tag{4.16}
\end{equation*}
$$

time. Throughout this chapter, we use $\mathcal{I}$ to denote an IDEAL schedule of the task system $\tau$.
In the partial ideal (PT) scheduling algorithm, each task $T_{i}$ with an active job at each instant is allocated a fraction of the system equal to

$$
\frac{\operatorname{Dwt}\left(T_{i}, t\right)}{\mathcal{T S}\left(P_{[q]}, t\right)}
$$

Hence, if $\mathcal{P} T$ is a PT schedule of $\tau$ and the task $T_{i}$ is active over the interval $\left[t_{1}, t_{2}\right)$, then
over $\left[t_{1}, t_{2}\right), T_{i}$ is allocated

$$
\begin{equation*}
\mathrm{A}\left(\mathcal{P} T, T_{i}, t_{1}, t_{2}\right)=\int_{t_{1}}^{t_{2}} \frac{\operatorname{Dwt}\left(T_{i}, u\right)}{\mathcal{T} \mathrm{S}\left(P_{[q]}, u\right)} d u \tag{4.17}
\end{equation*}
$$

time. The PT scheduling algorithm will be used to help calculate the drift incurred by changing the desired weight of a task. Throughout this chapter, we use $\mathcal{P} T$ to denote a PT schedule of the task system $\tau$.

Example (Figures 4.15 and 4.16). Consider the example in Figures 4.15 and 4.16, which pertain to a one-processor system that is assigned four tasks: $T_{1}$, which has $\mathrm{e}\left(T_{1}\right)=1$ and $\operatorname{Dwt}\left(T_{1}\right)=1 / 3 ; T_{2}$, which has $\mathrm{e}\left(T_{2}\right)=1$ and $\operatorname{Dwt}\left(T_{2}\right)=1 / 6 ; T_{3}$, which has $\mathrm{e}\left(T_{3}\right)=2$, $\operatorname{Dwt}\left(T_{3}\right)=1 / 4$, and leaves at time 8 ; and $T_{4}$, which has $\mathrm{e}\left(T_{4}\right)=4$ and an initial desired weight of $1 / 4$ that increases to $1 / 2$ at time 8 via Case (i) of Rule P . (This is the same system as in Figure 4.3.) Figure 4.15(a) depicts the PEDF schedule. Figure 4.15(b) depicts the allocations to $T_{4}$ in the PEDF, IDEAL, SW, CSW, and PT scheduling algorithms. Figure 4.16(a) depicts the SW, IDEAL, and PT schedules. Figure 4.16(b) depicts the CSW schedule. Notice that $T_{4}^{1}$ receives no allocations in CSW once it has received one unit of execution (the amount that $T_{4}^{1}$ is allocated in the PEDF schedule).

Example (Figure 4.17). Consider the example in Figure 4.17, which depicts one processor that is assigned four tasks (where the execution time of each job is one): $T_{1}$, which has $\operatorname{Dwt}\left(T_{1}\right)=1 / 2 ; T_{2}$ and $T_{3}$, both of which have a desired weight of $1 / 6$; and $T_{4}$, which has an initial desired weight of $1 / 2$ that initiates a desired weight decrease to $1 / 6$ at time 1 that is enacted at time $2 . \overline{6}$ via Case (ii) of Rule N. (This is the same system as in Figure 4.4.)

Notice that, since the IDEAL scheduling algorithm allocates capacity to each task based on its guaranteed weight (rather than based on its guaranteed scheduling weight), when one task initiates a decrease in its desired weight on an over-utilized processor, the allocation to all other tasks in the IDEAL schedule immediately increases. For example, in Figure 4.17(e), when $T_{4}$ initiates a desired weight decrease at time 1 , the rate of allocation to all other tasks immediately increases even though the desired weight change is not enacted until time $2 \cdot \overline{6}$. (In the IDEAL schedule, before time $1, T_{1}$ is allocated $3 / 8$ of the processor at each instant,


Figure 4.15: A one-processor system with four tasks. (a) The PEDF schedule. (b) The allocations to $T_{4}$ in the PEDF, IDEAL, SW, CSW, and PT scheduling algorithms. (Figure 4.16 depicts the IDEAL, SW, CSW, and PT schedules.)


Figure 4.16: The IDEAL, CSW, SW, PT schedules for the same system as in Figure 4.15. (a) The SW, IDEAL, and PT schedules. (b) The CSW schedule.
and after time 1 it is allocated $1 / 2$ of the processor at each instant.)

It is important to note that since the scaling factor for each task in the PT scheduling algorithm is based on the the total guaranteed scheduling weight (rather than the total guaranteed weight), when one task initiates a decrease in its desired weight on an over-utilized processor, the allocation to all other tasks in the PT schedule remains the same until the desired weight change is enacted. For example, in Figure 4.17(f), when $T_{4}$ initiates a desired weight decrease at time 1, the rate of allocation to all other tasks remains the same until the change is enacted at time $\frac{8}{3}$. This example illustrates the difference between the IDEAL and PT scheduling algorithms.


Figure 4.17: A one-processor system with four tasks. (a) The PEDF schedule. (b) The allocations to $T_{4}$ in the PEDF, IDEAL, SW, CSW, and PT scheduling algorithms. (c) The SW schedule. (d) The CSW schedule. (e) The IDEAL schedule. (f) The PT schedule.

| Property | Definition |
| :---: | :--- |
| (D) <br> [Page 110] | For any two times $u_{1}$ and $u_{2}$ such that $\mathrm{r}\left(T_{i}^{j}\right) \leq u_{1} \leq u_{2} \leq \mathrm{d}\left(T_{i}^{j}\right), \int_{u_{1}}^{u_{2}} \operatorname{SGwt}\left(T_{i}, t\right) d t \leq$ <br> (SW-2) <br> $[$ Page 127] For any job $T_{i}^{j}$ and any $t \geq \mathrm{r}\left(T_{i}^{j}\right), \mathrm{A}\left(\mathcal{C} S W, T_{i}^{j}, \mathrm{r}\left(T_{i}^{j}\right), t\right) \leq \mathrm{Ae}\left(T_{i}^{j}\right) \leq \mathrm{e}\left(T_{i}^{j}\right)$. |
| (SW-2) <br> $[$ Page 127] | Any job $T_{i}^{j}$ is complete in the CSW schedule by time $\mathrm{d}\left(T_{i}^{j}\right)$. |
| $\mathbf{( X )}$ <br> $[$ Page 122] | If a task $T_{i}$ initiates a desired weight change at time $t_{c}$ and the job $T_{i}^{j}$ is active at $t_{c}$, <br> then some desired weight change is enacted according to Rule P or N by either $\mathrm{d}\left(T_{i}^{j}\right)$ <br> or when $T_{i}^{j}$ becomes inactive, whichever is first. |
| $\mathbf{( V )}$ <br> $[$ Page 128] | For the jobs $T_{i}^{j}$ and $T_{i}^{j+1}$, if $\mathrm{d}\left(T_{i}^{j}\right)>\mathrm{r}\left(T_{i}^{j+1}\right)$, then $T_{i}^{j}$ is complete in both the CSW <br> and actual schedules by time $\mathrm{r}\left(T_{i}^{j+1}\right)$. |

Table 4.5: Summary of properties used in Section 4.9.

### 4.9 Drift

We now turn our attention to the issue of measuring drift under PEDF. The properties used throughout this section are summarized in Table 4.5. We begin this section, by formally defining drift. The drift of a task $T_{i}$ at time $t$ is defined as

$$
\begin{equation*}
\operatorname{drift}\left(T_{i}, t\right)=\mathrm{A}\left(\mathcal{I}, T_{i}, 0, t\right)-\mathrm{A}\left(\mathcal{C} S W, T_{i}, 0, t\right) \tag{4.18}
\end{equation*}
$$

In this section, we show that at time $t$ for $T_{i}$, if $t$ satisfies one of conditions (T-1), .., (T-3) (defined below), then the value of $\operatorname{drift}\left(T_{i}, t\right)$ is bounded by $[-\mathcal{X} \cdot \mathcal{Q}, \mathcal{X} \cdot \mathcal{Q}]$, where $\mathcal{X}$ is the maximal execution time of any task in the system and $\mathcal{Q}$ denotes the number of system resets plus the number of desired weight changes for any task assigned to the same processor as $T_{i}$ that are initiated before $t$. In the following conditions, $T_{i}^{j}$ denotes the last-released job (if any) of $T_{i}$ before $t$.
(T-1) The first job of $T_{i}$ (if it exists) is released at or after $t$.
$\mathbf{( T - 2 )} \mathrm{d}\left(T_{i}^{j}\right) \leq t$.
(T-3) $r\left(T_{i}^{j}\right)<t_{e} \leq t$, where $t_{e}$ is the time that the last-initiated change (at or before $t$ ) by $T_{i}$ was enacted.

The reason why we must constrain the times at which the drift is measured is because it is possible for $T_{i}$ to incur drift for a reweighting event of $T_{i}$ that has yet to be initiated. For example, in Figure $4.5, T_{4}$ incurs drift over the range $[0,2)$ even though it does not initiate a desired weight change until time 2. This complication arises because the reweighting rules may halt the last-released job of a task. Since the CSW scheduling algorithm is clairvoyant, it accounts for this drift before the reweighting event occurs. Thus, to accurately assess the drift $T_{i}$ incurs per reweighting event, we can only measure the drift at the times described above.

We begin our discussion by first calculating the drift that is incurred in a system with no system resets. We then factor in resets in Section 4.9.4.

Lemma 4.3. For any job $T_{i}^{j}$ of any task $T_{i}$ in a adaptable sporadic task system scheduled by PEDF,

$$
\mathrm{A}\left(\mathcal{C} S W, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t_{I}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right)
$$

where $t_{I}$ is the time that $T_{i}^{j}$ becomes inactive.
Proof. By Property (SW-2), $\mathrm{A}\left(\mathcal{C} S W, T_{i}^{j}, \mathrm{r}\left(T_{i}^{j}\right), t\right) \leq \mathrm{e}_{\max }\left(T_{i}\right)$ holds for any $t \geq \mathrm{r}\left(T_{i}^{j}\right)$, including $t_{I}$. Thus, if it can be shown that no other job of $T_{i}$ receives allocations in $\mathcal{C} S W$ over the range $\left[\mathrm{r}\left(T_{i}^{j}\right), t_{I}\right)$, then the proof is complete. Since $T_{i}^{j}$ becomes inactive at $\min \left(\mathrm{d}\left(T_{i}^{j}\right), \mathrm{r}\left(T_{i}^{j+1}\right)\right)$, it follows that no job of $T_{i}$ that is released after $\mathrm{r}\left(T_{i}^{j}\right)$ recives any allocations in $\mathcal{C} S W$ over the range $\left[\mathrm{r}\left(T_{i}^{j}\right), t_{I}\right)$.

Thus, it remains to be shown that no job $T_{i}^{a}$ that is released before $\mathrm{r}\left(T_{i}^{j}\right)$ receives any allocations in $\mathcal{C} S W$ over the range $\left[\mathrm{r}\left(T_{i}^{j}\right), t_{I}\right)$. By Property $(\mathrm{V})$, if $\mathrm{d}\left(T_{i}^{a}\right)>\mathrm{r}\left(T_{i}^{j}\right)$, then $T_{i}^{a}$ is complete by $\mathrm{r}\left(T_{i}^{j}\right)$. By Property (SW-3), if $\mathrm{d}\left(T_{i}^{a}\right) \leq \mathrm{r}\left(T_{i}^{j}\right)$, then $T_{i}^{a}$ is complete by $\mathrm{r}\left(T_{i}^{j}\right)$. Thus, in either case, $T_{i}^{a}$, receives no allocations in CSW over the range $\left[\mathrm{r}\left(T_{i}^{j}\right), t_{I}\right)$. This completes the proof.

Two types of drift. One complication with calculating the drift associated with a nonresetting reweighting event is that on an over-utilized processor a task may incur drift even when its desired weight does not change. For example, in Figure $4.17, T_{1}, T_{2}$, and $T_{3}$, incur drift over the range $[1,8 / 3)$, during which time $T_{4}$ has initiated a reweighting event that has
not yet been enacted. This behavior is the result of defining each task's guaranteed weight as a function of its desired weight and the desired weight of every task in the system. Thus, when $T_{4}$ initiates its desired weight decrease at time 1, in the IDEAL schedule, in Figure 4.17(e), the guaranteed weight of every other task immediately increases. However, in the CSW schedule, in Figure $4.17(\mathrm{~d})$, the guaranteed weight of $T_{1}, T_{2}$, and $T_{3}$ cannot change until $T_{4}$ 's decrease is enacted at time $8 / 3$.

As a result, a task $T_{i}$ can incur two types of drift: (i) $T_{i}$ can incur drift because it changes its desired weight and (ii) $T_{i}$ can incur drift because another task on the same processor initiated a desired weight change that is not immediately enacted. (Recall that, for now, we are ignoring the drift incurred by resetting the system.) In order to determine the total amount of drift caused by a reweighting event, we consider these two types of drift separately. In order to differentiate between type-(i) and -(ii) drift, we use the PT scheduling algorithm (which was defined in Section 4.8).

Recall that under the PT scheduling algorithm at each instant each active task $T_{i}$ (assigned to $\left.P_{[q]}\right)$ receives an allocation equal to

$$
\frac{\operatorname{Dwt}\left(T_{i}, t\right)}{\mathcal{T S}\left(P_{[q]}, t\right)}
$$

Since the scaling factor in the PT scheduling algorithm, i.e., $\frac{1}{\operatorname{TS}\left(P_{[q]}, t\right)}$, is defined in terms of the total desired scheduling weight instead of the total desired weight (i.e., in terms of $\mathcal{T S}\left(P_{[q]}, t\right)$ instead of $\left.\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)\right)$, the only time there is a difference between $T_{i}$ 's allocations in the PT and CSW schedules, is when $T_{i}$ 's desired weight is changed. We describe this difference as the partial drift of a task $T_{i}$ at time $t$, which is formally defined as

$$
\begin{equation*}
\operatorname{Pdrift}\left(T_{i}, t\right)=\mathrm{A}\left(\mathcal{P} T, T_{i}, 0, t\right)-\mathrm{A}\left(\mathcal{C} S W, T_{i}, 0, t\right) \tag{4.19}
\end{equation*}
$$

For example, in Figures 4.15 and 4.16, only $T_{4}$ incurs partial drift because it is the only task that changes its desired weight. Also, in Figure 4.17, $T_{4}$ is the only task to incur partial drift because it is the only task that changes its desired weight. Thus, the partial drift incurred by $T_{i}$ changing its desired weight via Rules P and N is equal to drift of type (i).

In addition, since the only difference between a task's allocations in the IDEAL and PT schedules is the scaling factor (i.e., in the IDEAL schedule $T_{i}$ receives $\frac{\operatorname{Dwt}\left(T_{i}, t\right)}{\mathcal{T D}\left(P_{[q]}, t\right)}$ allocations at each instant, while in the PT, schedule $T_{i}$ receives $\frac{\operatorname{Dwt}\left(T_{i}, t\right)}{\operatorname{TS}\left(P_{[q]}, t\right)}$ allocations at each instant), the amount of type (ii) drift a task incurs is equal to the difference in its allocations between the PT and IDEAL schedules.

### 4.9.1 Partial Drift

In this section, we establish bounds on the partial drift incurred by a reweighting event.

Lemma 4.4. Under PEDF, let $T_{i}$ be a task that initiates a desired weight change at $t_{c}$. Assume that $T_{i}$ has released a job before time $t_{c}$ and let $T_{i}^{j}$ be its last-released job before $t_{c}$. Let $t_{e}$ denote the first time at or after $t_{c}$ at which $T_{i}$ enacts a desired weight change. Let $t_{I}$ denote the time that $T_{i}^{j}$ becomes inactive. If $t_{c}$ is the first change initiated by $T_{i}$ in the range $\left(\mathrm{r}\left(T_{i}^{j}\right), t_{I}\right]$, then for any $t_{b} \in\left[t_{e}, t_{I}\right]$ and any time $t_{a} \in\left[r\left(T_{i}^{j}\right), t_{b}\right], \operatorname{Pdrift}\left(T_{i}, t_{b}\right)-\operatorname{Pdrift}\left(T_{i}, t_{a}\right)$ is bounded by $\left[-\mathrm{e}_{\max }\left(T_{i}\right) \cdot\left(\mathcal{G}+\mathcal{P}_{1}\right), \mathrm{e}_{\max }\left(T_{i}\right) \cdot\left(\mathcal{P}_{1}+\mathcal{P}_{2}\right)\right]$, where $\mathcal{G}, \mathcal{P}_{1}$, and $\mathcal{P}_{2}$ denote, respectively, the number of weight decreases, the number of weight increases via Case (i) of Rule P, and the number of weight increases via Case (ii) of Rule P by $T_{i}$ that were initiated at or before $t_{b}$ and enacted or canceled after $t_{a}$.

Proof. Let $t_{c}, t_{e}, t_{b}, t_{a}, t_{I}$, and $T_{i}^{j}$ be as defined in the statement of the lemma. If $t_{I}<t_{e}$, then the lemma is vacuously true, so assume that $t_{e} \leq t_{I}$. Notice that, by the definition of inactive (Definition 4.1), $t_{I}=\min \left(\mathrm{d}\left(T_{i}^{j}\right), \mathrm{r}\left(T_{i}^{j+1}\right)\right)$. Thus, by the statement of the lemma, we have

$$
\begin{equation*}
t_{c} \leq t_{e} \leq t_{b} \leq t_{I} \leq \mathrm{d}\left(T_{i}^{j}\right) \tag{4.20}
\end{equation*}
$$

Suppose for the moment that $T_{i}$ initiates a desired weight change at time $t_{c}^{\prime}$ such that $t_{e}<t_{c}^{\prime} \leq t_{I}$. Notice that $t_{c}^{\prime}$ 's existence implies that the change enacted at $t_{e}$ must have been by Case (i) of Rule N , since for all other rules, $t_{e}=t_{I}$ (specifically, all other scenarios release $T_{i}^{j+1}$ at $t_{e}$ ). Thus, by Case (i) of Rule $\mathrm{N}, T_{i}^{j}$ is halted before $t_{c}^{\prime}$. Thus, we having the following property:
(HA) No change initiated in the range $\left(t_{e}, t_{I}\right]$ can halt $T_{i}^{j}$.

Having established (4.20) and Property (HA), we now show that the value of $\operatorname{Pdrift}\left(T_{i}, t_{b}\right)-$ $\operatorname{Pdrift}\left(T_{i}, t_{a}\right)$ is bounded. Notice that, by (4.19),

$$
\begin{equation*}
\operatorname{Pdrift}\left(T_{i}, t_{b}\right)-\operatorname{Pdrift}\left(T_{i}, t_{a}\right)=\mathrm{A}\left(\mathcal{P} T, T_{i}, t_{a}, t_{b}\right)-\mathrm{A}\left(\mathcal{C} S W, T_{i}, t_{a}, t_{b}\right) \tag{4.21}
\end{equation*}
$$

By adding and subtracting $\int_{t_{a}}^{t_{b}} \mathrm{SGwt}\left(T_{i}, u\right) d u$, the right-hand side of the above formula can be rewritten as

$$
\mathrm{A}\left(\mathcal{P} T, T_{i}, t_{a}, t_{b}\right)-\int_{t_{a}}^{t_{b}} \operatorname{SGwt}\left(T_{i}, u\right) d u+\int_{t_{a}}^{t_{b}} \operatorname{SGwt}\left(T_{i}, u\right) d u-\mathrm{A}\left(\mathcal{C} S W, T_{i}, t_{a}, t_{b}\right) .
$$

We now bound the terms $\mathrm{A}\left(\mathcal{P} T, T_{i}, t_{a}, t_{b}\right)-\int_{t_{a}}^{t_{b}} \operatorname{SGwt}\left(T_{i}, u\right) d u$ and $\int_{t_{a}}^{t_{b}} \operatorname{SGwt}\left(T_{i}, u\right) d u-$ $\mathrm{A}\left(\mathcal{C} S W, T_{i}, t_{a}, t_{b}\right)$.

Bounding $\mathrm{A}\left(\mathcal{P} T, T_{i}, t_{a}, t_{b}\right)-\int_{t_{a}}^{t_{b}} \mathrm{SGwt}\left(T_{i}, u\right) d u$. Notice that, since $t_{c}$ is the first change initiated by $T_{i}$ over the range $\left(\mathrm{r}\left(T_{i}^{j}\right), t_{I}\right]$, it follows that for any $t \in\left[\mathrm{r}\left(T_{i}^{j}\right), t_{c}\right)$, $\operatorname{Dwt}\left(T_{i}, t\right)=$ $\operatorname{SDwt}\left(T_{i}, t\right)$. Thus, since, by (4.5), $\mathrm{SGwt}\left(T_{i}, t\right)=\frac{\mathrm{SDwt}\left(T_{i}, t\right)}{\mathrm{TS}\left(P_{[q]}, t\right)}$ and $T_{i}$ receives $\frac{\operatorname{Dwt}\left(T_{i}, t\right)}{\mathcal{T S}\left(P_{[q]}, t\right)}$ allocations at each instant in the PT schedule, it follows that $\mathrm{A}\left(\mathcal{P} T, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t_{c}\right)-\int_{\mathrm{r}\left(T_{i}^{j}\right)}^{t_{c}} \mathrm{SGwt}\left(T_{i}, u\right) d u=$ 0 , and, by extension, if $t_{a} \in\left[\mathrm{r}\left(T_{i}^{j}\right), t_{b}\right]$, then $\mathrm{A}\left(\mathcal{P} T, T_{i}, t_{a}, t_{c}\right)-\int_{t_{a}}^{t_{c}} \mathrm{SGwt}\left(T_{i}, u\right) d u=0$. Thus, for the remainder of this case we bound the value of $\mathrm{A}\left(\mathcal{P} T, T_{i}, \max \left(t_{c}, t_{a}\right), t_{b}\right)-$ $\int_{\max \left(t_{c}, t_{a}\right)}^{t_{b}} \operatorname{SGwt}\left(T_{i}, u\right) d u$.

Notice that, if $t_{c}=t_{b}$, then $t_{b}=t_{e}=\max \left(t_{c}, t_{a}\right)$ (since, by the definition of $t_{b}$ and $t_{a}$ given in the statement of the lemma, both $t_{a} \leq t_{b}$ and $t_{e} \leq t_{b}$ hold, and by (4.20), $t_{c} \leq t_{e}$ ) and $\mathrm{A}\left(\mathcal{P} T, T_{i}, \max \left(t_{c}, t_{a}\right), t_{b}\right)-\int_{\max \left(t_{c}, t_{a}\right)}^{t_{b}} \mathrm{SGwt}\left(T_{i}, u\right) d u=0$. Thus, for the remainder of this case, we assume that $t_{c}<t_{b}$.

Let $z$ denote the number of times $T_{i}$ initiates or enacts a desired weight change over the range $\left[\max \left(t_{c}, t_{a}\right), t_{b}\right)$. We begin by decomposing the interval $\left[\max \left(t_{c}, t_{a}\right), t_{b}\right)$ into several subregions. Let $t_{1}=\max \left(t_{c}, t_{a}\right)$ and let $t_{z+1}=t_{b}$. For $k \in\{2,3, \ldots, z\}$, let $t_{k}$ be the first time after $t_{k-1}$ such that $T_{i}$ either initiates or enacts a desired weight change. Notice that, by statement of the lemma, $\mathrm{r}\left(T_{i}^{j}\right)<\max \left(t_{c}, t_{a}\right)$, and by the definition of our decomposition,
$\max \left(t_{c}, t_{a}\right) \leq t_{k}$ for any $k \in\{1, \ldots, z+1\}$. Thus,

$$
\begin{equation*}
\text { for any } k \in\{1, \ldots, z+1\}, \mathrm{r}\left(T_{i}^{j}\right)<\max \left(t_{c}, t_{a}\right) \leq t_{k} . \tag{4.22}
\end{equation*}
$$

Notice that, if we can bound the value of $\mathrm{A}\left(\mathcal{P} T, T_{i}, t_{k}, t_{k+1}\right)-\int_{t_{k}}^{t_{k+1}} \operatorname{SGwt}\left(T_{i}, u\right) d u$ for every value of $k \in\{2,3, \ldots, z\}$, then we can compute a bound for $\mathrm{A}\left(\mathcal{P} T, T_{i}, t_{a}, t_{b}\right)$ $\int_{t_{a}}^{t_{b}} \operatorname{SGwt}\left(T_{i}, u\right) d u$.

Example (Figure 4.18). Consider the example in Figure 4.18, which depicts one processor that is assigned three tasks: $T_{1}$ and $T_{2}$, both of which have an execution time of 3 and a desired weight of $1 / 4$; and $T_{3}$, which has $\mathrm{e}\left(T_{3}\right)=2$ and an initial desired weight of $1 / 6$ that changes to $1 / 10$ at time 3 via Case (ii) of Rule N , then to $1 / 2$ at time 5 via Case (i) of Rule N , and then again to $1 / 3$ at time 7 via Case (ii) of Rule N. Notice that, because the change initiated at time 3 is via Case (ii) of Rule N , the change is not enacted until $T_{3}^{1}$ 's deadline. As a result, when a change is initiated at time 6 , this new change cancels the previous change. Moreover, since the change enacted at time 6 is a weight increase, it is immediately enacted via Case (i) of Rule N. Finally, before the deviance of $T_{3}^{1}$ becomes zero, $T_{3}$ initiates a weight decrease at time 7, which is enacted once the $T_{3}^{1}$ 's deviance becomes zero at time 8. Notice that, if $t_{a}=\mathrm{r}\left(T_{3}^{1}\right)$ and $t_{b}=8$, then there are three weight initiations and two enactments over the range $\left[\mathrm{r}\left(T_{3}^{1}\right), t_{b}\right)$. Thus, $t_{1}=3, t_{2}=6, t_{3}=7$, and $t_{4}=t_{b}=8$.

To continue the proof, let $k$ denote some value in $\{1,2, \ldots, z\}$. If $T_{i}$ enacts a change at $t_{k}$, then since $T_{i}$ does not initiate or enact a desired weight change until $t_{k+1}, \operatorname{SGwt}\left(T_{i}, u\right)=$ $\operatorname{Gwt}\left(T_{i}, u\right)$. Thus, $\mathrm{A}\left(\mathcal{P} T, T_{i}, t_{k}, t_{k+1}\right)=\int_{t_{k}}^{t_{k+1}} \mathrm{SGwt}\left(T_{i}, u\right) d u$ holds. Thus, we henceforth assume that $T_{i}$ initiates a change at $t_{k}$. Since $t_{k}<t_{k+1}$, the change initiated at $t_{k}$ must have been by either Case (ii) of Rule P or Case (ii) of Rule N. We now consider these two possibilities.

Sub-Case 1: Case (ii) of Rule P. By definition, $T_{i}$ does not enact any desired weight changes over the range $\left(t_{k}, t_{k+1}\right)$. Thus, we have the following property
(SC) The value $\operatorname{SDwt}\left(T_{i}, t\right)$ is constant within $\left[t_{k}, t_{k+1}\right)$.


Figure 4.18: A one-processor example of the task decomposition in Lemma 4.4.

Thus, for brevity, we denote $\operatorname{SDwt}\left(T_{i}, t_{k}\right)$ as Ow .
Suppose that $T_{i}$ initiates a change in its desired weight from Ow to Nw via Case (ii) of Rule P at $t_{k}$. Thus, for $t \in\left[t_{k}, t_{k+1}\right)$, $\operatorname{Dwt}\left(T_{i}, t\right)=\mathrm{Nw}$. Since, by Property (SC), for $t \in\left[t_{k}, t_{k+1}\right), \operatorname{SDwt}\left(T_{i}, t\right)=\mathrm{Ow}, \mathrm{A}\left(\mathcal{P} T, T_{i}, t_{k}, t_{k+1}\right)=\int_{t_{k}}^{t_{k+1}} \frac{\operatorname{Dwt}\left(T_{i}, t\right)}{\mathcal{T S}(P[q], t)} d t$, and, by (4.5), $\operatorname{SGwt}\left(T_{i}, t\right)=\frac{\operatorname{SDwt}\left(T_{i}, t\right)}{\operatorname{TS}\left(P_{[q]}, t\right)}$, it follows that

$$
\begin{equation*}
\mathrm{A}\left(\mathcal{P} T, T_{i}, t_{k}, t_{k+1}\right)-\int_{t_{k}}^{t_{k+1}} \operatorname{SGwt}\left(T_{i}, t\right) d t=\int_{t_{k}}^{t_{k+1}} \frac{\mathrm{Nw}-\mathrm{Ow}}{\mathcal{T} \mathrm{~S}\left(P_{[q]}, t\right)} d t . \tag{4.23}
\end{equation*}
$$

Thus, for the remainder of this subcase, we bound the value of $\int_{t_{k}}^{t_{k+1}} \frac{\mathrm{Nw}-\mathrm{Ow}}{\mathrm{TS}\left(P_{[q]}, t\right)} d t$.
Recall that a task only changes its weight via Case (ii) of Rule P if the deviance of $T_{i}$ is positive and the following equation holds:

$$
\begin{equation*}
\frac{\operatorname{Irem}\left(T_{i}^{j}, t_{k}\right)}{\operatorname{Ow}} \leq \frac{\operatorname{REM}\left(T_{i}^{j}, t_{k}\right)}{\mathrm{Nw}} \tag{4.24}
\end{equation*}
$$

By our decomposition and (4.20), $t_{k+1} \leq t_{e} \leq \mathrm{d}\left(T_{i}^{j}\right)$. Thus, by Property (D), $\int_{\mathrm{r}\left(T_{i}^{j}\right)}^{t_{k+1}} \operatorname{SGwt}\left(T_{i}, u\right) d u \leq \mathrm{e}\left(T_{i}\right)$. Notice that, by (4.8), $\operatorname{Irem}\left(T_{i}^{j}, t_{k}\right)=\mathrm{e}\left(T_{i}^{j}\right)-\int_{\mathrm{r}\left(T_{i}^{j}\right)}^{t_{k}} \operatorname{SGwt}\left(T_{i}, u\right) d u$. By substituting $\int_{\mathrm{r}\left(T_{i}^{j}\right)}^{t_{k+1}} \operatorname{SGwt}\left(T_{i}, u\right) d u \leq \mathrm{e}\left(T_{i}\right)$ into $\operatorname{Irem}\left(T_{i}^{j}, t_{k}\right)=\mathrm{e}\left(T_{i}^{j}\right)-\int_{\mathrm{r}\left(T_{i}^{j}\right)}^{t_{k}} \operatorname{SGwt}\left(T_{i}, u\right) d u$, it follows that $\int_{t_{k}}^{t_{k+1}} \operatorname{SGwt}\left(T_{i}, u\right) d u \leq \operatorname{Irem}\left(T_{i}^{j}, t_{k}\right)$.

Since, by (4.5), $\operatorname{SGwt}\left(T_{i}, t\right)=\frac{\operatorname{SDwt}\left(T_{i}, t\right)}{\operatorname{TS}\left(P_{[q]}, t\right)}$, and for any $t \in\left[t_{k}, t_{k+1}\right), \operatorname{SDwt}\left(T_{i}, t\right)=$ $\operatorname{SDwt}\left(T_{i}, t_{k}\right)=\mathrm{Ow}$, by (4.11), we have

$$
\begin{equation*}
\int_{t_{k}}^{t_{k+1}} \frac{\mathrm{Ow}}{\mathcal{T S}\left(P_{[q]}, t\right)} d t=\int_{t_{k}}^{t_{k+1}} \operatorname{SGwt}\left(T_{i}, u\right) d u \leq \operatorname{Irem}\left(T_{i}^{j}, t_{k}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right) . \tag{4.25}
\end{equation*}
$$

We now consider two scenarios depending on whether $\mathrm{Ow}>\mathrm{Nw}$ or $\mathrm{Ow}<\mathrm{Nw}$. We first consider the scenario where $\mathrm{Ow}>\mathrm{Nw}$. In this case, since, by (4.6), $\mathcal{T S}\left(P_{[q]}, t\right) \geq 1$, it is trivial to show that

$$
0 \geq \int_{t_{k}}^{t_{k+1}} \frac{\mathrm{Nw}-\mathrm{Ow}}{\mathcal{T} \mathrm{~S}\left(P_{[q]}, t\right)} d t
$$

holds. By (4.25) and because $\mathrm{Nw}-\mathrm{Ow} \geq-\mathrm{Ow}$ and $\mathcal{T S}\left(P_{[q]}, t\right) \geq 1$, it follows that

$$
0 \geq \int_{t_{k}}^{t_{k+1}} \frac{\mathrm{Nw}-\mathrm{Ow}}{\mathcal{T} \mathrm{~S}\left(P_{[q]}, t\right)} d t \geq \int_{t_{k}}^{t_{k+1}} \frac{-\mathrm{Ow}}{\mathcal{T} \mathrm{~S}\left(P_{[q]}, t\right)} d t \geq-\mathrm{e}_{\max }\left(T_{i}\right)
$$

Thus, by (4.23), if $T_{i}$ initiates a weight decrease via Case (ii) of Rule P at time $t_{k}$, then

$$
\begin{equation*}
0 \geq \mathrm{A}\left(\mathcal{P} T, T_{i}, t_{k}, t_{k+1}\right)-\int_{t_{k}}^{t_{k+1}} \operatorname{SGwt}\left(T_{i}, u\right) d u \geq-\mathrm{e}_{\max }\left(T_{i}\right) . \tag{4.26}
\end{equation*}
$$

We now consider the scenario where $\mathrm{Ow}<\mathrm{Nw}$. In this case, since by (4.6), $\mathcal{T} \mathrm{S}\left(P_{[q]}, t\right) \geq 1$, it is trivial to show that

$$
0 \leq \int_{t_{k}}^{t_{k+1}} \frac{\mathrm{Nw}-\mathrm{Ow}}{\mathcal{T} \mathrm{~S}\left(P_{[q]}, t\right)} d t
$$

holds. By (4.24) and (4.25), it follows that

$$
\int_{t_{k}}^{t_{k+1}} \frac{1}{\mathcal{T S}\left(P_{[q]}, t\right)} d t \leq \frac{\operatorname{Irem}\left(T_{i}^{j}, t_{k}\right)}{\mathrm{Ow}} \leq \frac{\operatorname{REM}\left(T_{i}^{j}, t_{k}\right)}{\mathrm{Nw}}
$$

Thus, by (4.13),

$$
\int_{t_{k}}^{t_{k+1}} \frac{\mathrm{Nw}}{\mathcal{T S}\left(P_{[q]}, t\right)} d t \leq \operatorname{REM}\left(T_{i}^{j}, t_{c}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right)
$$

Since, $N w>O w \geq 0$,

$$
0 \leq \int_{t_{k}}^{t_{k+1}} \frac{\mathrm{Nw}-\mathrm{Ow}}{\mathcal{T} \mathrm{~S}\left(P_{[q]}, t\right)} d t \leq \operatorname{REM}\left(T_{i}^{j}, t_{c}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right)
$$

Thus, by (4.23), if $T_{i}$ initiates a weight increase via Case (ii) of Rule P at time $t_{k}$, then

$$
\begin{equation*}
0 \leq \mathrm{A}\left(\mathcal{P} T, T_{i}, t_{k}, t_{k+1}\right)-\int_{t_{k}}^{t_{k+1}} \operatorname{SGwt}\left(T_{i}, t\right) d t \leq \mathrm{e}_{\max }\left(T_{i}\right) . \tag{4.27}
\end{equation*}
$$

Case (ii) of Rule N. Suppose that $T_{i}$ changes its desired weight to Nw at time $t_{k}$ via Rule N. If $T_{i}$ initiates a desired weight decrease at time $t_{k}$ (i.e., the change initiated is via Case (ii) of Rule N ), it follows that for any $t \in\left[t_{k}, t_{k+1}\right)$, $\operatorname{Dwt}\left(T_{i}, t\right) \leq \operatorname{SDwt}\left(T_{i}, t\right)$. Thus, since $\mathrm{A}\left(\mathcal{P} T, T_{i}, t_{k}, t_{k+1}\right)=\int_{t_{k}}^{t_{k+1}} \frac{\operatorname{SDwt}\left(T_{i}, u\right)}{\mathcal{T S}\left(P_{[q]}, t\right)} d u($ by $(4.17))$, and $\operatorname{SGwt}\left(T_{i}, t\right)=\frac{\operatorname{SDwt}\left(T_{i}, t\right)}{\operatorname{TS}\left(P_{[q]}, t\right)}($ by $(4.5))$, it follows that $0 \leq \mathrm{A}\left(\mathcal{P} T, T_{i}, t_{k}, t_{k+1}\right) \leq \int_{t_{k}}^{t_{k+1}} \operatorname{SGwt}\left(T_{i}, u\right) d u$ holds.

We now show that $\int_{t_{k}}^{t_{k+1}} \operatorname{SGwt}\left(T_{i}, u\right) d u \leq \mathrm{e}_{\max }\left(T_{i}\right)$. By the definition of our decomposition and (4.20), $t_{k+1} \leq t_{b} \leq \mathrm{d}\left(T_{i}^{j}\right)$. Thus, by Property (D), $\int_{\mathrm{r}\left(T_{i}^{j}\right)}^{t_{k+1}} \mathrm{SGwt}\left(T_{i}, u\right) d u \leq \mathrm{e}\left(T_{i}\right)$. Since, by (4.22), $\mathrm{r}\left(T_{i}^{j}\right)<t_{k}$, and $\operatorname{SGwt}\left(T_{i}, t\right) \geq 0$ for all $t$, it follows that $\int_{t_{k}}^{t_{k+1}} \operatorname{SGwt}\left(T_{i}, u\right) d u \leq$ $\int_{\mathrm{r}\left(T_{i}^{j}\right)}^{t_{k+1}} \operatorname{SGwt}\left(T_{i}, u\right) d u \leq \mathrm{e}\left(T_{i}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right)$. Thus, if $T_{i}$ initiates a weight decrease via Case (ii) of Rule N at time $t_{k}$, then

$$
\begin{equation*}
0 \geq \mathrm{A}\left(\mathcal{P} T, T_{i}, t_{k}, t_{k+1}\right)-\int_{t_{k}}^{t_{k+1}} \operatorname{SGwt}\left(T_{i}, u\right) d u \geq-\mathrm{e}_{\max }\left(T_{i}\right) \tag{4.28}
\end{equation*}
$$

By combining (4.26), (4.27), and (4.28), it follows that

$$
\begin{equation*}
-\mathrm{e}_{\max }\left(T_{i}\right) \cdot \mathcal{G} \leq \mathrm{A}\left(\mathcal{P} T, T_{i}, t_{a}, t_{b}\right)-\int_{t_{a}}^{t_{b}} \mathrm{SGwt}\left(T_{i}, u\right) d u \leq \mathrm{e}_{\max }\left(T_{i}\right) \cdot \mathcal{P}_{2} \tag{4.29}
\end{equation*}
$$

where $\mathcal{G}$ denotes the number of weight decreases by $T_{i}$ that were initiated at or before $t_{b}$ and enacted or canceled after $t_{a}$, over the range and $\mathcal{P}_{2}$ denotes the number of desired weight increases via Case (ii) of Rule P that were initiated at or before $t_{b}$ and enacted or canceled
after $t_{a}$.

Bounding $\int_{t_{a}}^{t_{b}} \operatorname{SGwt}\left(T_{i}, u\right) d u-\mathrm{A}\left(\mathcal{C} S W, T_{i}, t_{a}, t_{b}\right)$. Notice that $\int_{t_{a}}^{t_{b}} \operatorname{SGwt}\left(T_{i}, u\right) d u=$ $\mathrm{A}\left(\mathcal{C} S W, T_{i}, t_{a}, t_{b}\right)$ unless $T_{i}^{j}$ is halted by a desired weight change. By Property (HA), no change initiated over the range $\left(t_{e}, t_{I}\right]$ can halt $T_{i}^{j}$. Thus, if $\int_{t_{a}}^{t_{b}} \operatorname{SGwt}\left(T_{i}, u\right) d u \neq$ $\mathrm{A}\left(\mathcal{C} S W, T_{i}, t_{a}, t_{b}\right)$, then $T_{i}^{j}$ enacted a desired weight change at $t_{e}$ that halts $T_{i}^{j}$. Moreover, by the reweighting rules, if the change enacted at $t_{e}$ halts $T_{i}^{j}$, then that change must also have been initiated at $t_{e}$ and must have been by either Case (i) of Rule P or Case (i) of Rule N .

Suppose that the change initiated at $t_{e}$ is via Case (i) of Rule N. Notice that, by Property $(\mathrm{SW}-2), \mathrm{A}\left(\mathcal{C} S W, T_{i}, t_{a}, t_{e}\right) \leq \mathrm{Ae}\left(T_{i}^{j}\right)$. Moreover, at the smallest value of $t_{r}$ such that $\mathrm{A}\left(\mathcal{S} W, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t_{r}\right)=\mathrm{Ae}\left(T_{i}^{j}\right)$ holds, $T_{i}^{j}$ becomes inactive, i.e., $t_{I}=t_{r}$. Recall, from the definition of $\mathcal{C} S W$, that $\mathrm{A}\left(\mathcal{C} S W, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t\right)=\mathrm{A}\left(\mathcal{S} W, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t\right)$ up to the first $t$ such that $\mathrm{A}\left(\mathcal{C} S W, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t\right)=\operatorname{Ae}\left(T_{i}^{j}\right)$. Thus, since, by $(4.20), t_{e} \leq t_{b} \leq t_{I}, T_{i}^{j}$ receives an allocation of $\operatorname{SGwt}\left(T_{i}, t\right)$ at each instant $t \in\left[\mathrm{r}\left(T_{i}^{u}\right), t_{b}\right)$ in $\mathcal{C} S W$. Therefore, since $t_{a} \in\left[\mathrm{r}\left(T_{i}^{j}\right), t_{b}\right]$ (by the statement of the lemma), we have $\int_{t_{a}}^{t_{b}} \operatorname{SGwt}\left(T_{i}, u\right) d u=\mathrm{A}\left(\mathcal{C} S W, T_{i}, t_{a}, t_{b}\right)$. Thus, if $\int_{t_{a}}^{t_{b}} \operatorname{SGwt}\left(T_{i}, u\right) d u \neq \mathrm{A}\left(\mathcal{C} S W, T_{i}, t_{a}, t_{b}\right)$, then the change must have been via Case (i) of Rule P.

If the change initiated at $t_{e}$ is via Case (i) of Rule P , then $T_{i}^{j}$ has positive deviance at $t_{e}$ and $t_{e}=t_{I}$. Thus, by (4.20), $t_{e}=t_{b}$. Since $t_{e}=t_{b}=t_{I} \leq \mathrm{d}\left(T_{i}^{j}\right)$ (by (4.20)), by Property (D), $0 \leq \int_{t_{a}}^{t_{b}} \mathrm{SGwt}\left(T_{i}, u\right) d u \leq \mathrm{e}_{\max }\left(T_{i}\right)$. Also, by Lemma 4.3, since $t_{b}=t_{I}$, $0 \leq \mathrm{A}\left(\mathcal{C} S W, T_{i}, t_{a}, t_{b}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right)$. Thus,

$$
\begin{equation*}
-\mathrm{e}_{\max }\left(T_{i}\right) \leq \int_{t_{a}}^{t_{b}} \operatorname{SGwt}\left(T_{i}, u\right) d u-\mathrm{A}\left(\mathcal{C} S W, T_{i}, t_{a}, t_{b}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right) \tag{4.30}
\end{equation*}
$$

From (4.29) and (4.30), we have $-\mathrm{e}_{\max }\left(T_{i}\right) \cdot\left(\mathcal{G}+\mathcal{P}_{1}\right) \leq \mathrm{A}\left(\mathcal{P} T, T_{i}, t_{a}, t_{b}\right)-$ $\mathrm{A}\left(\mathcal{C} S W, T_{i}, t_{a}, t_{b}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right) \cdot\left(\mathcal{P}_{1}+\mathcal{P}_{2}\right)$, where $\mathcal{G}, \mathcal{P}_{1}$, and $\mathcal{P}_{2}$ are as defined in the statement of the lemma. Thus, by (4.19)

$$
-\mathrm{e}_{\max }\left(T_{i}\right) \cdot\left(\mathcal{G}+\mathcal{P}_{1}\right) \leq \operatorname{Pdrift}\left(T_{i}, t_{b}\right)-\operatorname{Pdrift}\left(T_{i}, t_{a}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right) \cdot\left(\mathcal{P}_{1}+\mathcal{P}_{2}\right)
$$

which completes the proof of Lemma 4.4.

Notice that, if a task $T_{i}$ is inactive over the range $\left[t_{a}, t_{b}\right)$, then $\operatorname{Pdrift}\left(T_{i}, t_{b}\right)-\operatorname{Pdrift}\left(T_{i}, t_{a}\right)=$ 0 . Also notice that if $T_{i}$ never changes its desired weight while the job $T_{i}^{j}$ is active, then for any two times $t_{a}$ and $t_{b}$ such that $\mathrm{r}\left(T_{i}^{j}\right) \leq t_{a} \leq t_{b} \leq t_{I}$, where $t_{I}$ is the time that $T_{i}^{j}$ becomes inactive, $\operatorname{Pdrift}\left(T_{i}, t_{b}\right)-\operatorname{Pdrift}\left(T_{i}, t_{a}\right)=0$. Thus, a task $T_{i}$ can only incur partial drift over a range in which a job of $T_{i}$ is active and $T_{i}$ initiates a desired weight change. Such a scenario is addressed in Lemma 4.4. Thus, by iteratively applying Lemma 4.4, it is possible to show that the partial drift incurred over any range $\left[t_{a}, t_{b}\right]$ is bounded by $\left[-\mathrm{e}_{\max }\left(T_{i}\right) \cdot\left(\mathcal{G}+\mathcal{P}_{1}\right), \mathrm{e}_{\max }\left(T_{i}\right) \cdot\right.$ $\left.\left(\mathcal{P}_{1}+\mathcal{P}_{2}\right)\right]$, where $\mathcal{G}, \mathcal{P}_{1}$, and $\mathcal{P}_{2}$ denote, respectively, the number of weight decreases by $T_{i}$, the number of weight increases by $T_{i}$ via Case (i) of Rule P , and the number of weight increases by $T_{i}$ via Case (ii) of Rule P that were initiated at or before $t_{b}$ and enacted or canceled after $t_{a}$. Assuming that $t_{b}$ and $T_{i}$ satisfy one of conditions (T-1), $\ldots$, (T-3).

This bound is summarized in the following corollary

Corollary 4.1. Let $T_{i}$ denote $a$ task that is assigned to $P_{[q]}$ over some range $\left[t_{a}, t_{b}\right)$. If $T_{i}$ and $t_{b}$ satisfy one of Conditions (T-1), ..., (T-3). Then, $\operatorname{Pdrift}\left(T_{i}, t_{b}\right)-\operatorname{Pdrift}\left(T_{i}, t_{a}\right)$ is bounded by $\left[-\mathrm{e}_{\max }\left(T_{i}\right) \cdot\left(\mathcal{G}+\mathcal{P}_{1}\right), \mathrm{e}_{\max }\left(T_{i}\right) \cdot\left(\mathcal{P}_{1}+\mathcal{P}_{2}\right)\right]$, where $\mathcal{G}, \mathcal{P}_{1}$, and $\mathcal{P}_{2}$ denote, respectively, the number of weight decreases, the number of weight increases via Case (i) of Rule $P$, and the number of weight increase via Case (ii) of Rule $P$ by $T_{i}$ that were initiated at or before $t_{b}$ and enacted or canceled after $t_{a}$.

Example (Figure 4.19). Consider the example in Figure 4.9, which depicts one processor that is assigned three tasks: $T_{1}$ and $T_{2}$, both of which have an execution time of 2 and a weight of $1 / 3$; and $T_{3}$, which has $\mathrm{e}\left(T_{3}\right)=1$ and an initial weight of $1 / 10$ that changes to $1 / 7$ at time 1 via Case (i) of Rule P and then to $1 / 3$ at time 3 via Case (i) of Rule P. Inset (a) depicts the PEDF schedule. Inset (b) depicts $T_{3}$ 's partial drift. By Lemma 4.4, the maximal partial drift incurred by $T_{3}$ is one (its execution time) over both ranges $[0,1$ ) and $[1,3)$ (even though the actual partial drift incurred is 0.1 over the range $[0,1)$ and $2 / 7$ over the range $[1,3))$. Notice that, Lemma 4.4 cannot be used to determine the maximal partial drift over the range $[1,2.5)$ because $T_{3}^{2}$ has not enacted a weight change. Moreover, notice that the
last-initiated change at or before time 2.5 was enacted at time 1, i.e., $t_{e}=1$, and $r\left(T_{3}^{2}\right)=1$. Hence, $t_{e}=\mathrm{r}\left(T_{3}^{2}\right)$. Thus, Condition (T-3) does not hold at time 2.5 (even though $1=t_{e} \leq 2.5$ holds). On the other hand, the last-initiated change at or before time 3 was enacted at time 3, i.e., $t_{e}=3$. Thus, in this case, $\mathrm{r}\left(T_{3}^{2}\right)<t_{e}$. Hence, Condition (T-3) holds at time 3 because $\mathrm{r}\left(T_{3}^{2}\right)<t_{e} \leq 3$. Corollary 4.1 is used to sum the partial drift over multiple jobs, i.e., by Corollary 4.1, the maximal partial drift over the range $[0,3)$ is two.

### 4.9.2 Relationship Between PT and IDEAL

Now that we have established the partial drift incurred by a change in the desired weight of a task, we can determine the difference between a task's allocation in the IDEAL and PT schedules. In this section, we show that this difference up to time $t$, which can be easily determined from (4.16) and (4.17) as $\int_{0}^{t} \operatorname{Dwt}\left(T_{i}, u\right) \cdot\left(\frac{1}{\mathcal{T D}\left(P_{[q]}, u\right)}-\frac{1}{\mathcal{T S}\left(P_{[q]}, u\right)}\right) d u$, is bounded by the range $[-\mathcal{C} \cdot \mathcal{X}, \mathcal{D} \cdot \mathcal{X}]$, where $\mathcal{C}$ denotes the number of weight increases via Case (ii) of Rule $\mathrm{P}, \mathcal{D}$ denotes the number of weight decreases via Case (ii) of Rule P or Case (ii) of Rule N , and $\mathcal{X}$ denotes the maximal execution time of any task assigned to the same processor as $T_{i}$. Before we can establish this bound, we must establish some preliminarily lemmas.

Lemma 4.5. If a task $T_{i}$ initiates a desired weight decrease at time $t_{c}$ from Ow to Nw while $T_{i}^{j}$ is active, $T_{i}$ is assigned to $P_{[q]}$, and the decrease is enacted or canceled at time $t_{e}$, then $0 \leq \int_{t_{c}}^{t_{e}} \frac{\operatorname{SDwt}\left(T_{i}, t\right)-\operatorname{Dwt}\left(T_{i}, t\right)}{\operatorname{TS}\left(P_{[q]}, t\right) T \mathrm{D}\left(P_{[q]}, t\right)} d t \leq \mathrm{e}_{\max }\left(T_{i}\right)$.

Proof. Notice that, if $t_{c}=t_{e}$, then lemma is trivially true. Thus, for the remainder of this proof we assume that $t_{c}<t_{e}$. Since $t_{c}<t_{e}$, the change initiated at $t_{c}$ is not immediately enacted. Moreover, by the definition of $t_{c}$ and $t_{e}, T_{i}$ does not initiate a desired weight change over the range $\left(t_{c}, t_{e}\right)$. Thus, both the desired weight and the desired scheduling weight of $T_{i}$ are constant over the range $\left[t_{c}, t_{e}\right)$. Specifically, for any $t \in\left[t_{c}, t_{e}\right)$, both

$$
\begin{equation*}
\operatorname{SDwt}\left(T_{i}, t\right)=\mathrm{Ow} \tag{4.31}
\end{equation*}
$$

| $\square$ | Scheduled $\downarrow$ Job released $\uparrow$ Job deadline -- Job Layout without reweighting |
| :--- | :--- |
|  |  |
| Reweighting event initiated $\forall$ Reweighting event enacted |  |


(a)

(b)

Figure 4.19: A one-processor example of multiple reweighting events. (a) The PEDF schedule. (b) $T_{3}$ 's partial drift.
and $\operatorname{Dwt}\left(T_{i}, t\right)=$ Nw hold. Thus, if

$$
\begin{equation*}
0 \leq \int_{t_{c}}^{t_{e}} \frac{\mathrm{Ow}-\mathrm{Nw}}{\mathcal{T S}\left(P_{[q]}, t\right) \mathcal{T D}\left(P_{[q]}, t\right)} d t \leq \mathrm{e}_{\max }\left(T_{i}\right) \tag{4.32}
\end{equation*}
$$

holds, then the proof is complete.
Since $\mathrm{Ow}>\mathrm{Nw}, t_{c}<t_{e}, \mathcal{T S}\left(P_{[q]}, t\right) \geq 1$ (by (4.6)), and $\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right) \geq 1$ (by (4.3)), it follows that

$$
0 \leq \int_{t_{c}}^{t_{e}} \frac{\mathrm{Ow}-\mathrm{Nw}}{\mathcal{T S}\left(P_{[q]}, t\right) \mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)} d t
$$

holds. In, the remainder of this proof, we establish the upper bound of (4.32).
Because $t_{e}>t_{c}$, the change initiated at time $t_{c}$ was either by Case (ii) of Rule P or Case (ii) of Rule N. In either case, by Property (X), the weight change is either enacted or canceled by $\mathrm{d}\left(T_{i}^{j}\right)$, i.e., $t_{e} \leq \mathrm{d}\left(T_{i}^{j}\right)$. Since $\mathrm{r}\left(T_{i}^{j}\right) \leq t_{c}<t_{e} \leq \mathrm{d}\left(T_{i}^{j}\right)$, by Property (D),

$$
\int_{t_{c}}^{t_{e}} \operatorname{SGwt}\left(T_{i}, t\right) d t \leq \mathrm{e}\left(T_{i}^{j}\right)
$$

Moreover, since, by (4.31), for any $t \in\left[t_{c}, t_{e}\right), \operatorname{SDwt}\left(T_{i}, t\right)=\mathrm{Ow}$, and by (4.5), $\operatorname{SGwt}\left(T_{i}, t\right)=$ $\frac{\operatorname{SDwt}\left(T_{i}, t\right)}{\operatorname{TS}\left(P_{[q]}, t\right)}$, we have

$$
\begin{equation*}
\int_{t_{c}}^{t_{e}} \frac{\mathrm{Ow}}{\mathcal{T} \mathrm{~S}\left(P_{[q]}, t\right)} d t \leq \mathrm{e}\left(T_{i}^{j}\right) \tag{4.33}
\end{equation*}
$$

Since, by the statement of the lemma, $\mathrm{Ow}>\mathrm{Nw} \geq 0$, and by (4.3), $\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right) \geq 1$, we have $\mathrm{Ow}-\mathrm{Nw} \leq \mathrm{Ow}$ and $\mathcal{T S}\left(P_{[q]}, t\right) \mathcal{T} \mathrm{D}\left(P_{[q]}, t\right) \geq \mathcal{T S}\left(P_{[q]}, t\right)$. Hence, by (4.33),

$$
\int_{t_{c}}^{t_{e}} \frac{\mathrm{Ow}-\mathrm{Nw}}{\mathcal{T S}\left(P_{[q]}, t\right) \mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)} d t \leq \int_{t_{c}}^{t_{e}} \frac{\mathrm{Ow}}{\mathcal{T} \mathrm{~S}\left(P_{[q]}, t\right)} d t \leq \mathrm{e}\left(T_{i}^{j}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right)
$$

Thus, it follows that the upper bound of (4.32) holds, which completes the proof.
Notice that, if a task initiates a desired weight increase at $t_{c}$ via Case (i) of Rule P or Case (i) of Rule N , then it is enacted at time $t_{c}$. Similarly, if the last-released job of $T_{i}$ is not active or $T_{i}$ has not release such a job, then the change is enacted at $t_{c}$. From these observations, we have the following lemma.

Lemma 4.6. Let $T_{i} \in P_{[q]}$ be a task that initiates a desired weight increase from Ow to Nw
at time $t_{c}$, let $T_{i}^{j}$ be the last-released job (if any) of $T_{i}$, and let $t_{e}$ denote the time the change is enacted or canceled. If $T_{i}^{j}$ is not active at $t_{c}$ or no such job exists, or the change initiated at $t_{c}$ is via either Case (i) of Rule P or Case (i) of Rule $N$, then $\int_{t_{c}}^{t_{e}} \frac{\operatorname{SDwt}\left(T_{i}, t\right)-\operatorname{Dwt}\left(T_{i}, t\right)}{\operatorname{TS}\left(P_{[q]}, t\right) T \mathrm{D}\left(P_{[q]}, t\right)} d t=0$.

Lemma 4.7. If a task $T_{i}$ initiates a desired weight increase at time $t_{c}$ via Case (ii) of Rule P from Ow to Nw while $T_{i}^{j}$ is active, $T_{i}$ is assigned to $P_{[q]}$, and the increase is enacted or canceled at time $t_{e}$, then $0 \leq \int_{t_{c}}^{t_{e}} \frac{\operatorname{Dwt}\left(T_{i}, t\right)-\operatorname{SDwt}\left(T_{i}, t\right)}{\operatorname{TS}\left(P_{[q]}, t\right) T \mathrm{D}\left(P_{[q]}, t\right)} d t \leq \mathrm{e}_{\max }\left(T_{i}\right)$.

Proof. Notice that, if $t_{c}=t_{e}$, then lemma is trivially true. Thus, for the remainder of the proof, we assume that $t_{c}<t_{e}$. Since $t_{c}<t_{e}$, the change initiated at $t_{c}$ is not immediately enacted. Moreover, by the definition of $t_{c}$ and $t_{e}, T_{i}$ does not initiate a desired weight change over the range $\left(t_{c}, t_{e}\right)$. Thus, both the desired weight and the desired scheduling weight of $T_{i}$ are constant over the range $\left[t_{c}, t_{e}\right)$. Specifically, for any $t \in\left[t_{c}, t_{e}\right)$, both $\operatorname{SDwt}\left(T_{i}, t\right)=\mathrm{Ow}$ and $\operatorname{Dwt}\left(T_{i}, t\right)=\mathrm{Nw}$ hold. Thus, if

$$
\begin{equation*}
0 \leq \int_{t_{c}}^{t_{e}} \frac{\mathrm{Nw}-\mathrm{Ow}}{\mathcal{T} \mathrm{~S}\left(P_{[q]}, t\right) \mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)} d t \leq \mathrm{e}_{\max }\left(T_{i}\right) . \tag{4.34}
\end{equation*}
$$

holds, then the proof is complete.
Since $\mathrm{Nw}>\mathrm{Ow}, t_{c}<t_{e}, \mathcal{T S}\left(P_{[q]}, t\right) \geq 1$ (by (4.6)), and $\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right) \geq 1$ (by (4.3)), it follows that

$$
0 \leq \int_{t_{c}}^{t_{e}} \frac{\mathrm{Nw}-\mathrm{Ow}}{\mathcal{T} \mathrm{~S}\left(P_{[q]}, t\right) \mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)} d t
$$

holds. In the remainder of this proof, we establish the upper bound of (4.34). By Property (X), the weight change initiated at $t_{c}$ is either enacted or canceled by $\mathrm{d}\left(T_{i}^{j}\right)$, i.e., $t_{e} \leq \mathrm{d}\left(T_{i}^{j}\right)$. Thus, by Property (D), $\int_{\mathrm{r}\left(T_{i}^{j}\right)}^{t_{e}} \mathrm{SGwt}\left(T_{i}, u\right) d u \leq \mathrm{e}\left(T_{i}\right)$. By (4.8), $\operatorname{Irem}\left(T_{i}^{j}, t_{c}\right)=$ $\mathrm{e}\left(T_{i}^{j}\right)-\int_{\mathrm{r}\left(T_{i}^{j}\right)}^{t_{c}} \mathrm{SGwt}\left(T_{i}, u\right) d u$. By substituting $\int_{\mathrm{r}\left(T_{i}^{j}\right)}^{t_{e}} \mathrm{SGwt}\left(T_{i}, u\right) d u \leq \mathrm{e}\left(T_{i}\right)$ into $\operatorname{Irem}\left(T_{i}^{j}, t_{c}\right)=$ $\mathrm{e}\left(T_{i}^{j}\right)-\int_{\mathbf{r}\left(T_{i}^{j}\right)}^{t_{c}} \mathrm{SGwt}\left(T_{i}, u\right) d u$, it follows that $\int_{t_{c}}^{t_{e}} \operatorname{SGwt}\left(T_{i}, u\right) d u \leq \operatorname{Irem}\left(T_{i}^{j}, t_{c}\right)$.

Since, by (4.5), $\operatorname{SGwt}\left(T_{i}, t\right)=\frac{\operatorname{SDwt}\left(T_{i}, t\right)}{\operatorname{TS}\left(P_{[q]}, t\right)}$, for any $t \in\left[t_{c}, t_{e}\right), \operatorname{SDwt}\left(T_{i}, t\right)=\operatorname{SDwt}\left(T_{i}, t_{c}\right)=$ Ow, and by (4.11), $\operatorname{Irem}\left(T_{i}^{j}, t_{v}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right)$, we have

$$
\begin{equation*}
\int_{t_{c}}^{t_{e}} \frac{\mathrm{Ow}}{\mathcal{T S}\left(P_{[q]}, t\right)} d t=\int_{t_{c}}^{t_{e}} \operatorname{SGwt}\left(T_{i}, u\right) d u \leq \operatorname{Irem}\left(T_{i}^{j}, t_{v}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right) . \tag{4.35}
\end{equation*}
$$



Figure 4.20: Decomposition of a task for Lemma 4.7.

We now decompose the interval $\left[t_{c}, t_{e}\right)$ into the set of regions $\left\{\left[t_{1}, t_{2}\right),\left[t_{2}, t_{3}\right), \ldots,\left[t_{z}, t_{z+1}\right)\right\} . \quad$ Let $t_{1}=t_{c}$ and $t_{z+1}=t_{e} . \quad$ For, $k \in\{1,2, \ldots, z-1\}$, let $t_{k+1}$ equal the first time after $t_{k}$ such that the value of either $\mathcal{T S}\left(P_{[q]}, t\right)$ or $\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)$ changes. Note that the value of either $\mathcal{T S}\left(P_{[q]}, t\right)$ or $\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)$ can only change as a result of a task enacting or initiating a desired weight change. Thus, there exist a discrete number of times over the region $\left[t_{c}, t_{e}\right)$ that the value of either $\mathcal{T S}\left(P_{[q]}, t\right)$ or $\mathcal{T D}\left(P_{[q]}, t\right)$ can change, which makes this decomposition possible. An example decomposition is given in Figure 4.20.

Since, by the definition of our decomposition, the value of $\mathcal{T S}\left(P_{[q]}, t\right)$ does not change within $\left[t_{k}, t_{k+1}\right)$, where $k=\{1,2, \ldots, z\}$, (4.35) implies

$$
\begin{equation*}
\mathrm{Ow} \cdot\left(\sum_{k=1, \ldots, z} \frac{t_{k+1}-t_{k}}{\mathcal{T S}\left(P_{[q]}, t_{k}\right)}\right) \leq \operatorname{Irem}\left(T_{i}^{j}, t_{c}\right) . \tag{4.36}
\end{equation*}
$$

Additionally, since the value of $\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)$ also does not change within $\left[t_{k}, t_{k+1}\right)$, where $k=\{1,2, \ldots, z\}$, we have

$$
\begin{equation*}
\int_{t_{c}}^{t_{e}} \frac{\mathrm{Nw}-\mathrm{Ow}}{\mathcal{T S}\left(P_{[q]}, t\right) \mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)} d t=(\mathrm{Nw}-\mathrm{Ow}) \sum_{k=1, \ldots, z} \frac{t_{k+1}-t_{k}}{\mathcal{T} \mathrm{D}\left(P_{[q]}, t_{k}\right) \mathcal{T S}\left(P_{[q]}, t_{k}\right)} . \tag{4.37}
\end{equation*}
$$

Notice that, since $\mathrm{Nw}>\mathrm{Ow}$ and $\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right) \geq 1$ holds (by (4.3)), it follows that

$$
(\mathrm{Nw}-\mathrm{Ow}) \sum_{k=1, \ldots, z} \frac{t_{k+1}-t_{k}}{\mathcal{T} \mathrm{D}\left(P_{[q]}, t_{k}\right) \mathcal{T} \mathrm{S}\left(P_{[q]}, t_{k}\right)} \leq(\mathrm{Nw}-\mathrm{Ow}) \sum_{k=1, \ldots, z} \frac{t_{k+1}-t_{k}}{\mathcal{T S}\left(P_{[q]}, t_{k}\right)}
$$

Thus, by (4.36) and the fact that Ow $>0$ holds, it follows that

$$
(\mathrm{Nw}-\mathrm{Ow}) \sum_{k=1, \ldots, z} \frac{t_{k+1}-t_{k}}{\mathcal{T} \mathrm{D}\left(P_{[q]}, t_{k}\right) \mathcal{T S}\left(P_{[q]}, t_{k}\right)} \leq(\mathrm{Nw}-\mathrm{Ow}) \frac{\operatorname{Irem}\left(T_{i}^{j}, t_{c}\right)}{\mathrm{Ow}}
$$

Thus, by (4.37),

$$
\int_{t_{c}}^{t_{e}} \frac{\mathrm{Nw}-\mathrm{Ow}}{\mathcal{T} \mathrm{~S}\left(P_{[q]}, t\right) \mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)} d t \leq(\mathrm{Nw}-\mathrm{Ow}) \frac{\operatorname{Irem}\left(T_{i}^{j}, t_{c}\right)}{\mathrm{Ow}}
$$

Thus, if we can show that

$$
(\mathrm{Nw}-\mathrm{Ow}) \frac{\operatorname{Irem}\left(T_{i}^{j}, t_{c}\right)}{\mathrm{Ow}} \leq \mathrm{e}_{\max }\left(T_{i}\right)
$$

holds, then the proof is complete.
Recall that the change initiated at $t_{c}$ was via Case (ii) of Rule P. Thus,

$$
\frac{\operatorname{Irem}\left(T_{i}^{j}, t_{c}\right)}{\mathrm{Ow}} \leq \frac{\operatorname{REM}\left(T_{i}^{j}, t_{c}\right)}{\mathrm{Nw}}
$$

Therefore, since Nw - Ow $>0$,

$$
(\mathrm{Nw}-\mathrm{Ow}) \frac{\operatorname{Irem}\left(T_{i}^{j}, t_{c}\right)}{\mathrm{Ow}} \leq(\mathrm{Nw}-\mathrm{Ow}) \frac{\operatorname{REM}\left(T_{i}^{j}, t_{c}\right)}{\mathrm{Nw}}
$$

By rearranging terms, we get

$$
(\mathrm{Nw}-\mathrm{Ow}) \frac{\operatorname{Irem}\left(T_{i}^{j}, t_{c}\right)}{\mathrm{Ow}} \leq\left(1-\frac{\mathrm{Ow}}{\mathrm{Nw}}\right) \operatorname{REM}\left(T_{i}^{j}, t_{c}\right)
$$

Since $\mathrm{Ow}<\mathrm{Nw}$ and $\operatorname{REM}\left(T_{i}^{j}, t_{c}\right) \leq \mathrm{e}\left(T_{i}^{j}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right)$ (by (4.13)), it follows that

$$
(\mathrm{Nw}-\mathrm{Ow}) \frac{\operatorname{Irem}\left(T_{i}^{j}, t_{c}\right)}{\mathrm{Ow}} \leq\left(1-\frac{\mathrm{Ow}}{\mathrm{Nw}}\right) \mathrm{REM}\left(T_{i}^{j}, t_{c}\right) \leq \mathrm{REM}\left(T_{i}^{j}, t_{c}\right) \leq \mathrm{e}\left(T_{i}^{j}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right)
$$

This completes the proof.

Before continuing, we introduce an additional function that will facilitate our discussion. Let $\gamma\left(T_{i}, t\right)$ be defined as

$$
\gamma\left(T_{i}, t\right)= \begin{cases}0, & \text { if } T_{i} \in \operatorname{ACT}\left(P_{[q]}, t\right)  \tag{4.38}\\ \frac{\operatorname{SDwt}\left(T_{i}, t\right)-\operatorname{Dwt}\left(T_{i}, t\right)}{\operatorname{TS}\left(P_{[q]}, t\right) \mathcal{T D}\left(P_{[q]}, t\right)}, & \text { otherwise },\end{cases}
$$

where $P_{[q]}$ is the processor that $T_{i}$ is assigned to at time $t$.
Lemma 4.8. Let $t_{a}$ and $t_{b}$ be any two times such that $t_{a}<t_{b}$, and the system is not reset over the interval $\left(t_{a}, t_{b}\right)$. Let $\mathcal{D}$ denote the number of desired weight decreases for $T_{i}$ that are both initiated at or before $t_{b}$ and enacted or canceled after $t_{a}$, and let $\mathcal{C}$ denote the number of desired weight increases via Case (ii) of Rule $P$ for $T_{i}$ that are both initiated at or before $t_{b}$ and enacted or canceled after $t_{a}$. Then,

$$
\begin{aligned}
-\mathcal{C} \cdot \mathrm{e}_{\max }\left(T_{i}\right) & \leq \int_{t_{a}}^{t_{b}} \min \left(0, \gamma\left(T_{i}, t\right)\right) d t \\
\mathcal{D} \cdot \mathrm{e}_{\max }\left(T_{i}\right) & \geq \int_{t_{a}}^{t_{b}} \max \left(0, \gamma\left(T_{i}, t\right)\right) d t .
\end{aligned}
$$

Proof. Notice that, if $t_{a}=t_{b}$, then the lemma is trivially true. Thus, for the remainder of this proof we assume that $t_{a}<t_{b}$. If $T_{i}$ initiates a change before $t_{a}$ that is not enacted or canceled until after $t_{a}$, then let $t_{a}^{\prime}$ denote the time that change was initiated; otherwise, let $t_{a}^{\prime}=t_{a}$. If $T_{i}$ initiates a change before $t_{b}$ that is not enacted or canceled until after $t_{b}$, then let $t_{b}^{\prime}$ denote the time that change is enacted or canceled; otherwise, let $t_{b}^{\prime}=t_{b}$. Since $t_{a}^{\prime} \leq t_{a}$ and $t_{b} \leq t_{b}^{\prime}$, it suffices to bound $\int_{t_{a}^{\prime}}^{t_{b}^{\prime}} \min \left(0, \gamma\left(T_{i}, t\right)\right) d t$ and $\int_{t_{a}^{\prime}}^{t_{b}^{\prime}} \max \left(0, \gamma\left(T_{i}, t\right)\right) d t$.

We now decompose the internal $\left[t_{a}^{\prime}, t_{b}^{\prime}\right)$ into the set of regions $\left\{\left[t_{1}, t_{2}\right),\left[t_{2}, t_{3}\right), \ldots,\left[t_{z}, t_{z+1}\right)\right\}$. Let $t_{1}=t_{a}^{\prime}$, let $t_{z+1}=t_{b}^{\prime}$, and for $k \in\{1,2, \ldots, z-1\}$, let $t_{k+1}$ the next time after $t_{k}$ that $T_{i}$ initiated or enacted a desired weight change. Notice that, by the definitions of $\mathcal{C}$ and $\mathcal{D}$ given in the statement of the lemma, there are $\mathcal{C}$ values of $k \in\{1,2, \ldots, z\}$ such that at $t_{k}, T_{i}$ initiated a desired weight increase via Case (ii) of Rule P , and $\mathcal{D}$ values of $k \in\{1,2, \ldots, z\}$ such that at $t_{k}, T_{i}$ initiated a desired weight decrease. Thus, we can complete the proof by showing that, for any value of $k \in\{1,2, \ldots z\}$,

- if $T_{i}$ initiates a desired weight increase via Case (ii) of Rule P at time $t_{k}$, then $-\mathrm{e}_{\max }\left(T_{i}\right) \leq \int_{t_{k}}^{t_{k+1}} \min \left(0, \gamma\left(T_{i}, t\right)\right) d t ;$
- if $T_{i}$ initiates a desired weight decrease at time $t_{k}$, then $\mathrm{e}_{\max }\left(T_{i}\right) \geq$ $\int_{t_{k}}^{t_{k+1}} \max \left(0, \gamma\left(T_{i}, t\right)\right) d t ;$
- if $T_{i}$ initiates any other type of desired weight change at $t_{k}$ or does not initiate a desired weight change at $t_{k}$, then $0=\int_{t_{k}}^{t_{k+1}} \max \left(0, \gamma\left(T_{i}, t\right)\right) d t=\int_{t_{k}}^{t_{k+1}} \min \left(0, \gamma\left(T_{i}, t\right)\right) d t$.

Before we discuss these three case, notice that, for all values of $k \in\{1,2, \ldots, z\}$, for any $t \in\left[t_{k}, t_{k+1}\right), \operatorname{SDwt}\left(T_{i}, t\right)=\operatorname{SDwt}\left(T_{i}, t_{k}\right)$ and $\operatorname{Dwt}\left(T_{i}, t\right)=\operatorname{Dwt}\left(T_{i}, t_{k}\right)$ both hold since $T_{i}$ does not initiate or enact a change within $\left(t_{k}, t_{k+1}\right)$, Thus, for any $k \in\{1,2, \ldots, z\}$, if $\operatorname{Dwt}\left(T_{i}, t_{k}\right)=\operatorname{SDwt}\left(T_{i}, t_{k}\right)$ holds, then $\int_{t_{k}}^{t_{k+1}} \min \left(0, \gamma\left(T_{i}, t\right)\right) d t=\int_{t_{k}}^{t_{k+1}} \max \left(0, \gamma\left(T_{i}, t\right)\right) d t=0$ holds as well.
$T_{i}$ initiates a desired weight increase via Case (ii) of Rule P at time $t_{k}$. If the change initiated at $t_{k}$ is enacted immediately, then $\operatorname{Dwt}\left(T_{i}, t_{k}\right)=\operatorname{SDwt}\left(T_{i}, t_{k}\right)$, which as we already discussed implies $\int_{t_{k}}^{t_{k+1}} \min \left(0, \gamma\left(T_{i}, t\right)\right) d t=\int_{t_{k}}^{t_{k+1}} \max \left(0, \gamma\left(T_{i}, t\right)\right) d t=0$. Thus, we assume that the change initiated at $t_{k}$ is not immediately enacted. Thus, by the definition of our decomposition, $t_{k+1}$ represents the time that $T_{i}$ enacts or cancels the change initiated at $t_{k}$. Thus, by Lemma 4.7, $0 \leq \int_{t_{k}}^{t_{k+1}} \frac{\operatorname{Dwt}\left(T_{i}, t\right)-\mathrm{SDwt}\left(T_{i}, t\right)}{\operatorname{TS}\left(P_{[q]}, t\right) T \mathrm{D}\left(P_{[q]}, t\right)} d t \leq \mathrm{e}_{\max }\left(T_{i}\right)$, which implies

$$
\begin{equation*}
-\mathrm{e}_{\max }\left(T_{i}\right) \leq \int_{t_{k}}^{t_{k+1}} \min \left(0, \gamma\left(T_{i}, t\right)\right) d t \tag{4.39}
\end{equation*}
$$

$T_{i}$ initiates a desired weight decrease at time $t_{k}$. If the change initiated at $t_{k}$ is enacted immediately, then $\operatorname{Dwt}\left(T_{i}, t_{k}\right)=\operatorname{SDwt}\left(T_{i}, t_{k}\right)$, which as we already discussed implies $\int_{t_{k}}^{t_{k+1}} \min \left(0, \gamma\left(T_{i}, t\right)\right) d t=\int_{t_{k}}^{t_{k+1}} \max \left(0, \gamma\left(T_{i}, t\right)\right) d t=0$. Thus, we assume that the change initiated at $t_{k}$ is not immediately enacted. In this case, by the definition of our decomposition, $t_{k+1}$ represents the time that $T_{i}$ enacts or cancels the change initiated at $t_{k}$. Thus, by Lemma 4.5, $0 \leq \int_{t_{k}}^{t_{k+1}} \frac{\operatorname{SDwt}\left(T_{i}, t\right)-\operatorname{Dwt}\left(T_{i}, t\right)}{\mathcal{T S}\left(P_{[q]}, t\right) T \mathrm{D}\left(P_{[q]}, t\right)} d t \leq \mathrm{e}_{\max }\left(T_{i}\right)$, which implies

$$
\begin{equation*}
\mathrm{e}_{\max }\left(T_{i}\right) \geq \int_{t_{k}}^{t_{k+1}} \max \left(0, \gamma\left(T_{i}, t\right)\right) d t \tag{4.40}
\end{equation*}
$$

$T_{i}$ initiates alternative change at $t_{k}$ or does not initiate a change at $t_{k}$. If $T_{i}$ initiates an alternative type of desired weight change or does not initaite a desired weight change at $t_{k}$, then $\operatorname{Dwt}\left(T_{i}, t_{k}\right)=\operatorname{SDwt}\left(T_{i}, t_{k}\right)$, which as we already discussed implies

$$
\begin{equation*}
\int_{t_{k}}^{t_{k+1}} \min \left(0, \gamma\left(T_{i}, t\right)\right) d t=\int_{t_{k}}^{t_{k+1}} \max \left(0, \gamma\left(T_{i}, t\right)\right) d t=0 . \tag{4.41}
\end{equation*}
$$

Since, there are at most $\mathcal{C}$ values of $k \in\{1,2, \ldots, z\}$ such that (4.39) holds, $\mathcal{D}$ values of $k \in\{1,2, \ldots, z\}$ such that (4.40) holds, and for every other value of $k \in\{1,2, \ldots, z\}$, (4.41) holds, we have

$$
-\mathcal{C} \cdot \mathrm{e}_{\max }\left(T_{i}\right) \leq \int_{t_{a}}^{t_{b}} \min \left(0, \gamma\left(T_{i}, t\right)\right) d t \leq \int_{t_{a}^{\prime}}^{t_{b}^{\prime}} \min \left(0, \gamma\left(T_{i}, t\right)\right) d t
$$

and

$$
\mathcal{D} \cdot \mathrm{e}_{\max }\left(T_{i}\right) \geq \int_{t_{a}}^{t_{b}} \max \left(0, \gamma\left(T_{i}, t\right)\right) d t \geq \int_{t_{a}^{\prime}}^{t_{b}^{\prime}} \max \left(0, \gamma\left(T_{i}, t\right)\right) d t .
$$

This completes the proof.

Having proven the prerequisite lemmas, we can now show that the difference between $\frac{1}{\mathcal{T D}\left(P_{[q]}, t\right)}$ and $\frac{1}{\mathcal{T S}\left(P_{[q]}, t\right)}$ is a function of the number of reweighting events. Bounding this difference will enable us to bound the difference between $T_{i}$ 's allocations in the IDEAL and PT schedules, which is given by $\int_{0}^{t} \operatorname{Dwt}\left(T_{i}, u\right) \cdot\left(\frac{1}{\mathcal{T D}\left(P_{[q]}, u\right)}-\frac{1}{\mathcal{T S}\left(P_{[q]}, u\right)}\right) d u$.

Lemma 4.9. Let $t_{a}$ and $t_{b}$ be any two times such that $t_{a}<t_{b}$, and the system is not reset over the interval $\left(t_{a}, t_{b}\right)$. If $\mathcal{D}$ denotes the number of desired weight decreases by tasks assigned to $P_{[q]}$ that are both initiated at or before $t_{b}$ and enacted or canceled after $t_{a}$ and $\mathcal{C}$ denotes the number of desired weight increases via Case (ii) of Rule $P$ by tasks assigned to $P_{[q]}$ that are both initiated at or before $t_{b}$ and enacted or canceled after $t_{a}$, then

$$
-\mathcal{C} \cdot \mathcal{X} \leq \int_{t_{a}}^{t_{b}} \min \left(0, \frac{1}{\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)}-\frac{1}{\mathcal{T} \mathrm{~S}\left(P_{[q]}, t\right)}\right) d t
$$



Figure 4.21: Decomposition of a task for Lemma 4.9.

$$
\mathcal{D} \cdot \mathcal{X} \geq \int_{t_{a}}^{t_{b}} \max \left(0, \frac{1}{\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)}-\frac{1}{\mathcal{T} \mathrm{~S}\left(P_{[q]}, t\right)}\right) d t
$$

where $\mathcal{X}$ is the largest execution time of any task assigned to $P_{[q]}$ over the range $\left[t_{a}, t_{b}\right)$.

Proof. We begin by decomposing the interval $\left[t_{a}, t_{b}\right)$ into several subregions. Let $t_{1}=t_{a}$ and let $t_{z+1}=t_{b}$. For $k \in\{1,2, \ldots, z-1\}$, let $t_{k+1}$ be the first time after $t_{k}$ such that at least one of the following conditions holds:

1. If $\mathcal{T} \mathrm{S}\left(P_{[q]}, t_{k}\right)=1$, then $\mathcal{T} \mathrm{S}\left(P_{[q]}, t_{k+1}\right)>1$.
2. If $\mathcal{T} \mathrm{D}\left(P_{[q]}, t_{k}\right)=1$, then $\mathcal{T} \mathrm{D}\left(P_{[q]}, t_{k+1}\right)>1$.
3. If $\mathcal{T S}\left(P_{[q]}, t_{k}\right)>1$, then $\mathcal{T S}\left(P_{[q]}, t_{k+1}\right)=1$.
4. If $\mathcal{T} \mathrm{D}\left(P_{[q]}, t_{k}\right)>1$, then $\mathcal{T} \mathrm{D}\left(P_{[q]}, t_{k+1}\right)=1$.

In other words, $t_{k+1}$ denotes the time at which the total desired weight or total desired scheduling weight changes from over-utilizing $P_{[q]}$ to either under- or fully-utilizing $P_{[q]}$, or vice versa. An example of this decomposition is given in Figure 4.21. Notice that it is possible that neither the total desired scheduling weight nor the total desired weight will change from over-utilizing to under/fully-utilizing a processor or vice versa. In this case, $t_{1}=t_{a}$ and $t_{2}=t_{b}$.

Before continuing, notice that $\frac{1}{\mathcal{T D}\left(P_{[q]}, t\right)}-\frac{1}{\mathcal{T S}\left(P_{[q]}, t\right)}$ can be rearranged to equal $\frac{\mathcal{T S}\left(P_{[q]}, t\right)-\mathcal{T} \mathbf{D}\left(P_{[q]}, t\right)}{\mathcal{T D}\left(P_{[q]}, t\right) \mathcal{T}\left(P_{[q]}, t\right)}$. Thus, since our decomposition completely covers the range $\left[t_{a}, t_{b}\right)$, we
have

$$
\begin{align*}
\int_{t_{a}}^{t_{b}} \min \left(0, \frac{1}{\mathcal{T D}\left(P_{[q]}, t\right)}-\frac{1}{\mathcal{T S}\left(P_{[q]}, t\right)}\right) d t & =\quad \int_{t_{a}}^{t_{b}} \min \left(0, \frac{\mathcal{T S}\left(P_{[q], t)}-\mathcal{T} \mathbf{D}\left(P_{[q]}, t\right)\right.}{\mathcal{T D}\left(P_{[q]}, t\right) \mathcal{T S}\left(P_{[q}, t\right)}\right) d t  \tag{4.42}\\
& =\sum \int_{t_{k}}^{t_{k+1}} \min \left(0, \frac{\mathcal{T S}\left(P_{[q]}, t\right)-\mathcal{T D}\left(P_{[q]}, t\right)}{\mathcal{T D}\left(P_{[q]}, t\right) \mathcal{T S}\left(P_{[q]}, t\right)}\right) d t \\
\int_{t_{a}}^{t_{b}} \max \left(0, \frac{1}{\mathcal{T D}\left(P_{[q]}, t\right)}-\frac{1}{\mathcal{T S}\left(P_{[q]}, t\right)}\right) d t & =\quad \int_{t_{a}}^{t_{b}} \max \left(0, \frac{\mathcal{T S}\left(P_{[q], t)-\mathcal{T}\left(P_{[q]}, t\right)}^{\mathcal{T D}\left(P_{[q]}, t\right) \mathcal{T S}\left(P_{[q]}, t\right)}\right) d t}{}\right.  \tag{4.43}\\
& =\sum \int_{t_{k}}^{t_{k+1}} \max \left(0, \frac{\mathcal{T S}\left(P_{[q]}, t\right)-\mathcal{T D}\left(P_{[q]}, t\right)}{\mathcal{T D}\left(P_{[q]}, t\right) \mathcal{T S}\left(P_{[q]}, t\right)}\right) d t
\end{align*}
$$

Thus, it suffices to bound the values of

$$
\int_{t_{k}}^{t_{k+1}} \min \left(0, \frac{\mathcal{T} \mathrm{~S}\left(P_{[q]}, t\right)-\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)}{\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right) \mathcal{T} \mathrm{S}\left(P_{[q]}, t\right)}\right) d t
$$

and

$$
\int_{t_{k}}^{t_{k+1}} \max \left(0, \frac{\mathcal{T} \mathrm{~S}\left(P_{[q]}, t\right)-\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)}{\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right) \mathcal{T} \mathrm{S}\left(P_{[q]}, t\right)}\right) d t
$$

for each interval $\left[t_{k}, t_{k+1}\right)$. We now consider the different possible values for $\mathcal{T S}\left(P_{[q]}, t\right)$ and $\mathcal{T S}\left(P_{[q]}, t\right)$ for any $t \in\left[t_{k}, t_{k+1}\right)$.

Interval type 1: $\mathcal{T S}\left(P_{[q]}, t\right)=1$ and $\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)=1$. If $\mathcal{T S}\left(P_{[q]}, t\right)=1$ and $\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)=1$ both hold, then

$$
\begin{equation*}
\frac{\mathcal{T S}\left(P_{[q]}, t\right)-\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)}{\mathcal{T D}\left(P_{[q]}, t\right) \mathcal{T} \mathrm{S}\left(P_{[q]}, t\right)}=0 \tag{4.44}
\end{equation*}
$$

Interval type 2: $\mathcal{T S}\left(P_{[q]}, t\right)>1$ and $\mathcal{T D}\left(P_{[q]}, t\right)>1$. By (4.3) and (4.6), if $\mathcal{T S}\left(P_{[q]}, t\right)>1$ and $\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)>1$ both hold, then $\mathcal{T S}\left(P_{[q]}, t\right)=\sum_{T_{i} \in \operatorname{ASSN}\left(P_{[q]}, t\right)} \operatorname{SDwt}\left(T_{i}, t\right)$ and $\mathcal{T D}\left(P_{[q]}, t\right)=$ $\sum_{T_{i} \in \operatorname{ASSN}\left(P_{[q]}, t\right)} \operatorname{Dwt}\left(T_{i}, t\right)$ hold. Thus,

$$
\frac{\mathcal{T S}\left(P_{[q]}, t\right)-\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)}{\mathcal{T D}\left(P_{[q]}, t\right) \mathcal{T S}\left(P_{[q]}, t\right)}
$$

can be rewritten as

$$
\frac{\sum_{T_{i} \in \operatorname{ASSN}\left(P_{[q]}, t\right)} \operatorname{SDwt}\left(T_{i}, t\right)-\sum_{T_{i} \in \operatorname{ASSN}\left(P_{[q]}, t\right)} \operatorname{Dwt}\left(T_{i}, t\right)}{\mathcal{T D}\left(P_{[q]}, t\right) \mathcal{T} \mathrm{S}\left(P_{[q]}, t\right)}
$$

which equals

$$
\sum_{T_{i} \in \operatorname{ASSN}\left(P_{[q]}, t\right)} \frac{\operatorname{SDwt}\left(T_{i}, t\right)-\operatorname{Dwt}\left(T_{i}, t\right)}{\mathcal{T D}\left(P_{[q]}, t\right) \mathcal{T} \mathrm{S}\left(P_{[q]}, t\right)}=\sum_{T_{i} \in P_{[q]}} \gamma\left(T_{i}, t\right)
$$

Thus,

$$
\begin{equation*}
\frac{\mathcal{T} \mathrm{S}\left(P_{[q]}, t\right)-\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)}{\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right) \mathcal{T} \mathrm{S}\left(P_{[q]}, t\right)}=\sum_{T_{i} \in P_{[q]}} \gamma\left(T_{i}, t\right) \tag{4.45}
\end{equation*}
$$

Interval type 3: $\mathcal{T S}\left(P_{[q]}, t\right)>1$ and $\mathcal{T D}\left(P_{[q]}, t\right)=1$. By (4.6), if $\mathcal{T S}\left(P_{[q]}, t\right)>1$ holds, then $\mathcal{T} \mathrm{S}\left(P_{[q]}, t\right)=\sum_{T_{i} \in \operatorname{ASSN}\left(P_{[q]}, t\right)} \operatorname{SDwt}\left(T_{i}, t\right)$ holds. Thus,

$$
\frac{\mathcal{T S}\left(P_{[q]}, t\right)-\mathcal{T D}\left(P_{[q]}, t\right)}{\mathcal{T D}\left(P_{[q]}, t\right) \mathcal{T S}\left(P_{[q]}, t\right)}=\frac{\sum_{T_{i} \in \operatorname{ASSN}\left(P_{[q]}, t\right)} \operatorname{SDwt}\left(T_{i}, t\right)-1}{\mathcal{T D}\left(P_{[q]}, t\right) \mathcal{T S}\left(P_{[q]}, t\right)} .
$$

By (4.3), $\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)=1$, implies $\sum_{T_{i} \in \operatorname{ASSN}\left(P_{[q]}, t\right)} \operatorname{Dwt}\left(T_{i}, t\right) \leq 1$. Thus, since $1<\mathcal{T S}\left(P_{[q]}, t\right)=$ $\sum_{T_{i} \in \operatorname{ASSN}\left(P_{[q]}, t\right)} \operatorname{SDwt}\left(T_{i}, t\right)$, we have

$$
0 \leq \frac{\sum_{T_{i} \in \operatorname{ASSN}\left(P_{[q]}, t\right)} \operatorname{SDwt}\left(T_{i}, t\right)-1}{\mathcal{T D}\left(P_{[q]}, t\right) \mathcal{T S}\left(P_{[q]}, t\right)} \leq \sum_{T_{i} \in \operatorname{ASSN}\left(P_{[q]}, t\right)} \frac{\operatorname{SDwt}\left(T_{i}, t\right)-\operatorname{Dwt}\left(T_{i}, t\right)}{\mathcal{T D}\left(P_{[q]}, t\right) \mathcal{T S}\left(P_{[q]}, t\right)}=\sum_{T_{i} \in P_{[q]}} \gamma\left(T_{i}, t\right) .
$$

Thus,

$$
\begin{equation*}
0 \leq \frac{\mathcal{T S}\left(P_{[q]}, t\right)-\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)}{\mathcal{T D}\left(P_{[q]}, t\right) \mathcal{T S}\left(P_{[q]}, t\right)}=\sum_{T_{i} \in P_{[q]}} \gamma\left(T_{i}, t\right) \tag{4.46}
\end{equation*}
$$

Interval type 4: $\mathcal{T S}\left(P_{[q]}, t\right)=1$ and $\mathcal{T D}\left(P_{[q]}, t\right)>1$. By (4.3), if $\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)>1$ holds, then $\mathcal{T} D\left(P_{[q]}, t\right)=\sum_{T_{i} \in \operatorname{ASSN}\left(P_{[q]}, t\right)} \operatorname{Dwt}\left(T_{i}, t\right)$ holds. Thus,

$$
\frac{\mathcal{T S}\left(P_{[q,}, t\right)-\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)}{\mathcal{T D}\left(P_{[q]}, t\right) \mathcal{T S}\left(P_{[q]}, t\right)}=\frac{1-\sum_{T_{i} \in \operatorname{ASSN}\left(P_{[q]}, t\right)} \operatorname{Dwt}\left(T_{i}, t\right)}{\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right) \mathcal{T} \mathrm{S}\left(P_{[q]}, t\right)}
$$

By (4.6), $\mathcal{T S}\left(P_{[q]}, t\right)=1$ implies $\sum_{T_{i} \in \operatorname{ASSN}\left(P_{[q]}, t\right)} \operatorname{SDwt}\left(T_{i}, t\right) \leq 1$. Thus, since $1<\mathcal{T S}\left(P_{[q]}, t\right)=$ $\sum_{T_{i} \in \operatorname{ASSN}\left(P_{[q]}, t\right)} \operatorname{SDwt}\left(T_{i}, t\right)$, we have

$$
0 \geq \frac{1-\sum_{T_{i} \in \operatorname{ASSN}\left(P_{[q]}, t\right)} \operatorname{Dwt}\left(T_{i}, t\right)}{\mathcal{T D}\left(P_{[q]}, t\right) \mathcal{T S}\left(P_{[q]}, t\right)} \geq \sum_{T_{i} \in \operatorname{ASSN}\left(P_{[q]}, t\right)} \frac{\operatorname{SDwt}\left(T_{i}, t\right)-\operatorname{Dwt}\left(T_{i}, t\right)}{\mathcal{T D}\left(P_{[q]}, t\right) \mathcal{T S}\left(P_{[q]}, t\right)}=\sum_{T_{i} \in P_{[q]}} \gamma\left(T_{i}, t\right)
$$

Thus,

$$
\begin{equation*}
0 \geq \frac{\mathcal{T} \mathrm{S}\left(P_{[q]}, t\right)-\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)}{\mathcal{T D}\left(P_{[q]}, t\right) \mathcal{T} \mathrm{S}\left(P_{[q]}, t\right)}=\sum_{T_{i} \in P_{[q]}} \gamma\left(T_{i}, t\right) \tag{4.47}
\end{equation*}
$$

Putting it together. From (4.44)-(4.47), it follows that for any interval $\left[t_{k}, t_{k+1}\right)$,

$$
\int_{t_{k}}^{t_{k+1}} \min \left(0, \frac{\mathcal{T S}\left(P_{[q]}, t\right)-\mathcal{T D}\left(P_{[q]}, t\right)}{\mathcal{T D}\left(P_{[q]}, t\right) \mathcal{T S}\left(P_{[q]}, t\right)}\right) d t \geq \sum_{T_{i} \in P_{[q]}} \int_{t_{k}}^{t_{k+1}} \min \left(0, \gamma\left(T_{i}, t\right)\right) d t
$$

and

$$
\int_{t_{k}}^{t_{k+1}} \max \left(0, \frac{\mathcal{T} \mathrm{~S}\left(P_{[q]}, t\right)-\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)}{\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right) \mathcal{T} \mathrm{S}\left(P_{[q]}, t\right)}\right) d t \leq \sum_{T_{i} \in P_{[q]}} \int_{t_{k}}^{t_{k+1}} \max \left(0, \gamma\left(T_{i}, t\right)\right) d t
$$

Thus, from (4.42),

$$
\begin{align*}
\int_{t_{a}}^{t_{b}} \min \left(0, \frac{\mathcal{T} \mathrm{~S}\left(P_{[q]}, t\right)-\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)}{\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right) \mathcal{T S}\left(P_{[q]}, t\right)}\right) d t & \geq \sum_{k=\{1, \ldots, z\}} \sum_{T_{i} \in P_{[q]}} \int_{t_{k}}^{t_{k+1}} \min \left(0, \gamma\left(T_{i}, t\right)\right) d t \\
& \geq \sum_{T_{i} \in P_{[q]}} \int_{t_{a}}^{t_{b}} \min \left(0, \gamma\left(T_{i}, t\right)\right) d t \tag{4.48}
\end{align*}
$$

and from (4.43),

$$
\begin{align*}
\int_{t_{a}}^{t_{b}} \max \left(0, \frac{\mathcal{T} \mathrm{~S}\left(P_{[q]}, t\right)-\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)}{\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right) \mathcal{T S}\left(P_{[q]}, t\right)}\right) d t & \leq \sum_{k=\{1, \ldots, z\}} \sum_{T_{i} \in P_{[q]}} \int_{t_{k}}^{t_{k+1}} \max \left(0, \gamma\left(T_{i}, t\right)\right) d t \\
& \leq \sum_{T_{i} \in P_{[q]}} \int_{t_{a}}^{t_{b}} \max \left(0, \gamma\left(T_{i}, t\right)\right) d t \tag{4.49}
\end{align*}
$$

By Lemma 4.8, we have

$$
-\sum_{T_{i} \in P_{[q]}} \mathcal{C}_{i} \cdot \mathrm{e}_{\max }\left(T_{i}\right) \leq \sum_{T_{i} \in P_{[q]}} \int_{t_{a}}^{t_{b}} \min \left(0, \gamma\left(T_{i}, t\right)\right) d t
$$

and

$$
\sum_{T_{i} \in P_{[q]}} \mathcal{D}_{i} \cdot \mathrm{e}_{\max }\left(T_{i}\right) \geq \sum_{T_{i} \in P_{[q]}} \int_{t_{a}}^{t_{b}} \max \left(0, \gamma\left(T_{i}, t\right)\right) d t
$$

where $\mathcal{D}_{i}$ denotes the number of desired weight decreases for $T_{i}$ that are both initiated at
or before $t_{b}$ and enacted or canceled after $t_{a}$ and $\mathcal{C}_{i}$ denotes the number of desired weight increases via Case (ii) of Rule P for $T_{i}$ that are both initiated at or before $t_{b}$ and enacted or canceled after $t_{a}$. Notice that $\sum \mathcal{C}_{i}=\mathcal{C}, \sum \mathcal{D}_{i}=\mathcal{D}$, and for any $T_{i}, \mathrm{e}_{\max }\left(T_{i}\right) \leq \mathcal{X}$. Thus, from (4.48) and (4.49), we have

$$
-\mathcal{C} \cdot \mathcal{X} \leq \int_{t_{a}}^{t_{b}} \min \left(0, \frac{1}{\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)}-\frac{1}{\mathcal{T} \mathrm{~S}\left(P_{[q]}, t\right)}\right) d t
$$

and

$$
\mathcal{D} \cdot \mathcal{X} \geq \int_{t_{a}}^{t_{b}} \max \left(0, \frac{1}{\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)}-\frac{1}{\mathcal{T} \mathrm{~S}\left(P_{[q]}, t\right)}\right) d t
$$

which completes the proof.

### 4.9.3 Calculating Drift

Having established Corollary 4.1 and Lemma 4.9 we can now calculate the total drift incurred. Before continuing, we introduce some terminology to facilitate our discussion. Assume that $T_{i}$ is assigned to the processor $P_{[q]}$ over the range $\left[t_{a}, t_{b}\right)$.

- Let $\mathcal{P}_{1}$ denote the number of weight increases via Case (i) of Rule P by $T_{i}$ that were initiated at or before $t_{b}$ and enacted or canceled after $t_{a}$.
- Let $\mathcal{P}_{2}$ denote the number of weight increases via Case (ii) of Rule P by $T_{i}$ that were initiated at or before $t_{b}$ and enacted or canceled after $t_{a}$.
- Let $\mathcal{G}$ denote the number of weight decreases by $T_{i}$ that were initiated at or before $t_{b}$ and enacted or canceled after $t_{a}$.
- Let $\mathcal{C}$ denote the number of weight increases via Case (ii) or Rule P by any task assigned to $P_{[q]}$ that were initiated at or before $t_{b}$ and enacted or canceled after $t_{a}$.
- Let $\mathcal{D}$ denote the number of weight decreases by any task assigned to $P_{[q]}$ that were initiated at or before $t_{b}$ and enacted or canceled after $t_{a}$.
- Let $\mathcal{X}$ denote the maximal execution time of any task assigned to $P_{[q]}$ over the range $\left[t_{a}, t_{b}\right)$.

By Corollary 4.1, the partial drift incurred by the task $T_{i}$ over the range $\left[t_{a}, t_{b}\right)$ is bounded by

$$
\begin{equation*}
-\left(\mathcal{G}+\mathcal{P}_{1}\right) \mathrm{e}_{\max }\left(T_{i}\right) \leq \operatorname{Pdrift}\left(T_{i}, t_{b}\right)-\operatorname{Pdrift}\left(T_{i}, t_{a}\right) \leq\left(\mathcal{P}_{1}+\mathcal{P}_{2}\right) \mathrm{e}_{\max }\left(T_{i}\right), \tag{4.50}
\end{equation*}
$$

assuming that $t_{b}$ and $T_{i}$ satisfy one of Conditions (T-1), ..., (T-3).
In addition, since $T_{i}$ 's allocation over the range $\left[t_{a}, t_{b}\right)$ is given by $\int_{t_{a}}^{t_{b}} \frac{\mathrm{Dwt}\left(T_{i}, t\right)}{\mathcal{T D}\left(P_{[q]}, t\right)} d t$ in the IDEAL schedule (by (4.16)), and by $\int_{t_{a}}^{t_{b}} \frac{\operatorname{Dwt}\left(T_{i}, t\right)}{\mathcal{T S}\left(P_{[q]}, t\right)} d t$ in the PT schedule (by (4.17)), the difference in $T_{i}$ 's allocation in the IDEAL and PT schedules over the range $\left[t_{a}, t_{b}\right)$ is given by

$$
\int_{t_{a}}^{t_{b}} \operatorname{Dwt}\left(T_{i}, t\right)\left(\frac{1}{\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)}-\frac{1}{\mathcal{T} \mathrm{~S}\left(P_{[q]}, t\right)}\right) d t
$$

Since $0 \leq \operatorname{Dwt}\left(T_{i}, t\right) \leq 1, \quad$ and $\min \left(0, \frac{1}{\mathcal{T D}\left(P_{[q]}, t\right)}-\frac{1}{\mathcal{T S}\left(P_{[q]}, t\right)}\right) \leq \frac{1}{\mathcal{T D}\left(P_{[q]}, t\right)}-$ $\frac{1}{\mathcal{T S}\left(P_{[q]}, t\right)} \leq \max \left(0, \frac{1}{\mathcal{T D}\left(P_{[q]}, t\right)}-\frac{1}{\mathcal{T S}\left(P_{[q]}, t\right)}\right)$, it follows that $\int_{t_{a}}^{t_{b}} \min \left(0, \frac{1}{\mathcal{T D}\left(P_{[q]}, t\right)}-\frac{1}{\mathcal{T S}\left(P_{[q]}, t\right)}\right) d t \leq$ $\int_{t_{a}}^{t_{b}} \operatorname{Dwt}\left(T_{i}, t\right)\left(\frac{1}{\operatorname{TD}\left(P_{[q]}, t\right)}-\frac{1}{\operatorname{TS}\left(P_{[q]}, t\right)}\right) d t \leq \int_{t_{a}}^{t_{b}} \max \left(0, \frac{1}{\operatorname{TD}\left(P_{[q]}, t\right)}-\frac{1}{\operatorname{TS}\left(P_{[q]}, t\right)}\right) d t$. Thus, by Lemma 4.9,

$$
\begin{equation*}
-\mathcal{C} \cdot \mathcal{X} \leq \int_{t_{a}}^{t_{b}} \operatorname{Dwt}\left(T_{i}, t\right)\left(\frac{1}{\mathcal{T} \mathrm{D}\left(P_{[q]}, t\right)}-\frac{1}{\mathcal{T} \mathrm{~S}\left(P_{[q]}, t\right)}\right) d t \leq \mathcal{D} \cdot \mathcal{X} \tag{4.51}
\end{equation*}
$$

The total drift incurred by $T_{i}$ within $\left[t_{a}, t_{b}\right)$ is given by $\operatorname{Pdrift}\left(T_{i}, t_{b}\right)-\operatorname{Pdrift}\left(T_{i}, t_{a}\right)+$ $\int_{t_{a}}^{t_{b}} \operatorname{Dwt}\left(T_{i}, t\right)\left(\frac{1}{\mathcal{T D}\left(P_{[q]}, t\right)}-\frac{1}{\mathcal{T S}\left(P_{[q]}, t\right)}\right) d t$. By (4.50) and (4.51), this incurred drift is within the range

$$
\left[-\left(\mathcal{G}+\mathcal{P}_{1}\right) \cdot \mathrm{e}_{\max }\left(T_{i}\right)-\mathcal{C} \cdot \mathcal{X}, \quad\left(\mathcal{P}_{1}+\mathcal{P}_{2}\right) \cdot \mathrm{e}_{\max }\left(T_{i}\right)+\mathcal{D} \cdot \mathcal{X}\right] .
$$

Thus, we have the following lemma.

Lemma 4.10. Let $t_{a}$ and $t_{b}$ be any two times such that $t_{a}<t_{b}$, and the system is not reset over the interval $\left(t_{a}, t_{b}\right)$. Let $\mathcal{Q}$ denote the number of desired weight changes for any task on $P_{[q]}$ that are both initiated at or before $t_{b}$ and enacted or canceled after $t_{a}$, and let $T_{i}$ be assigned to $P_{[q]}$ over the range $\left[t_{a}, t_{b}\right)$. If $T_{i}$ and $t_{b}$ satisfy one of Conditions (T-1), $\ldots,(T-3)$, then the absolute drift incurred by $T_{i}$ within $\left[t_{a}, t_{b}\right)$ is at most $\mathcal{Q} \cdot \mathcal{X}$, where $\mathcal{X}$ is the largest execution time of any task assigned to $P_{[q]}$ over the range $\left[t_{a}, t_{b}\right)$.

### 4.9.4 Incorporating Resets

Having established Lemma 4.10, we now determine the additional drift incurred by resetting the system.

Lemma 4.11. If the system is reset at time $t_{c}$, then the drift incurred by any task $T_{i}$ is within the range $\left[-\mathrm{e}_{\max }\left(T_{i}\right), \mathrm{e}_{\max }\left(T_{i}\right)\right]$.

Proof. Let $T_{i}^{j}$ be the last-released job (if any) of $T_{i}$ before $t_{c}$. Notice that, if $T_{i}^{j}$ is not active at $t_{c}$ or $T_{i}^{j}$ does not exist, then via Case (ii) of Rule $\mathrm{R}, T_{i}$ is simply assigned to a (possibly different) processor and $T_{i}$ 's next job is released at either $t_{c}+\theta\left(T_{i}^{j+1}\right)$ or $t_{c}+\theta\left(T_{i}^{1}\right)$, depending on whether $T_{i}^{j}$ exists. Since the only two potential sources of drift are delays in enacting a desired weight change and halting a job, it follows that no partial drift is incurred in this case. Thus, for the rest of this proof, we assume that $T_{i}^{j}$ is active at $t_{c}$. Thus,

$$
\begin{equation*}
t_{c} \in\left[\mathrm{r}\left(T_{i}^{j}\right), \min \left(\mathrm{r}\left(T_{i}^{j+1}\right), \mathrm{d}\left(T_{i}^{j}\right)\right)\right) . \tag{4.52}
\end{equation*}
$$

Since $T_{i}^{j}$ is active at $t_{c}$, it is reset via Case (i) of Rule R. In this case, $T_{i}$ releases a job immediately with execution time $\operatorname{nextE}\left(T_{i}^{j}, t_{c}\right)$. Moreover, its current job $T_{i}^{j}$ is halted, so it is as though allocation equal to $\mathrm{A}\left(\mathcal{P} T, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t_{c}\right)-\mathrm{A}\left(\mathcal{C} S W, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t_{c}\right)$ is "lost." Notice that, since, $T_{i}^{j}$ becomes inactive at $t_{c}$ (because it is halted), by Lemma 4.3

$$
\begin{equation*}
\mathrm{A}\left(\mathcal{C} S W, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t_{c}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right) \tag{4.53}
\end{equation*}
$$

Also, notice that, if $T_{i}$ does not initiate any desired weight increase over the range $\left[\mathrm{r}\left(T_{i}^{j}\right), t_{c}\right)$ that is not immediately enacted, then for all $t \in\left[\mathrm{r}\left(T_{i}^{j}\right), t_{c}\right), \mathrm{A}\left(\mathcal{P} T, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t\right) \leq$ $\int_{\mathrm{r}\left(T_{i}^{j}\right)}^{t} \mathrm{SGwt}\left(T_{i}, t\right) d t$. Thus, in this case, by Property (D) and (4.52),

$$
\begin{equation*}
\mathrm{A}\left(\mathcal{P} T, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t\right) \leq \mathrm{e}\left(T_{i}^{j}\right) \tag{4.54}
\end{equation*}
$$

Thus, if $T_{i}$ does not initiate any desired weight increase over the range $\left[\mathrm{r}\left(T_{i}^{j}\right), t_{c}\right)$ that is
not immediately enacted, then by (4.53) and (4.54),

$$
\begin{equation*}
-\mathrm{e}_{\max }\left(T_{i}\right) \leq \mathrm{A}\left(\mathcal{P} T, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t_{c}\right)-\mathrm{A}\left(\mathcal{C} S W, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t_{c}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right) . \tag{4.55}
\end{equation*}
$$

In addition, if $T_{i}$ initiates a desired weight increase over the range $\left[\mathrm{r}\left(T_{i}^{j}\right), t_{c}\right)$ that is not immediately enacted, then this change must have been via Case (ii) of Rule P. Notice that, if such a change occurs, then it is possible that the upper bound in (4.55) may not hold, i.e., it is possible that $\mathrm{A}\left(\mathcal{P} T, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t_{c}\right)-\mathrm{A}\left(\mathcal{C} S W, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t_{c}\right)=\mathrm{e}_{\max }\left(T_{i}\right)+X$, where $X>0$. However, if the upper bound in (4.55) is violated by $X$, then it must have been the case that $T_{i}$ incurred at least $X$ units of drift from reweighting events initiated over the range $\left[\mathrm{r}\left(T_{i}^{j}\right), t_{c}\right)$. Moreover, this drift would have been accounted for by Lemma 4.4. Thus, if $\mathrm{A}\left(\mathcal{P} T, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t_{c}\right)-\mathrm{A}\left(\mathcal{C} S W, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t_{c}\right)=\mathrm{e}_{\max }\left(T_{i}\right)+X$ holds, then the maximal amount of additional drift incurred by resetting $T_{i}$ is $\mathrm{A}\left(\mathcal{P} T, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t_{c}\right)-\mathrm{A}\left(\mathcal{C} S W, T_{i}, \mathrm{r}\left(T_{i}^{j}\right), t_{c}\right)-$ $X=\mathrm{e}_{\max }\left(T_{i}\right)$.

Thus, it follows that $T_{i}$ 's drift due to the reset is bounded within $\left[-\mathrm{e}_{\max }\left(T_{i}\right), \mathrm{e}_{\max }\left(T_{i}\right)\right]$.

### 4.9.5 Total Drift Incurred

From Lemmas 4.10 and 4.11, we have the following.

Theorem 4.4. For any task $T_{i}$ and for any interval of time $\left[t_{a}, t_{b}\right)$, let $\mathcal{Q}$ denote the number of system resets plus the number of desired weight changes for any task assigned to the same processor as $T_{i}$ that are both initiated at or before $t_{b}$ and enacted or canceled after $t_{a}$. If $T_{i}$ and $t_{b}$ satisfy one of Conditions (T-1), .., (T-3), then the absolute drift incurred by $T_{i}$ is at most $\mathcal{Q} \cdot \mathcal{X}$, where $\mathcal{X}$ is the largest execution time of any task in the system.

Moreover, if a task is never assigned to an over-utilized processor, then the partial drift of a task always equals its drift. As a result, we can tighten Theorem 4.4.

Theorem 4.5. For any task $T_{i}$ that is never assigned to an over-utilized processor and any interval of time $\left[t_{b}, t_{a}\right)$, let $\mathcal{Q}$ denote the number of system resets plus the number of desired weight changes by $T_{i}$ that are both initiated at or before $t_{b}$ and enacted or canceled after $t_{a}$.

If $T_{i}$ and $t_{b}$ satisfy one of Conditions (T-1), ...,(T-3), then the absolute drift incurred by $T_{i}$ is at most $\mathcal{Q} \cdot \mathrm{e}_{\max }\left(T_{i}\right)$.

### 4.9.6 Modifications for NP-PEDF

Note that delaying the initiation of a reweighting event due to non-preemptivity does not substantially increase the drift incurred per reweighting event, since the longest a reweighting event can be delayed is the execution time of the active job of the task being reweighted.

Suppose that task $T_{i}$ initiates a weight change at time $t_{c}$. If $T_{i}^{j}$ is active at $t_{c}$, and if $T_{i}$ 's reweighting event is delayed until some time $t$ (by a non-preemptive section), then at $t$ either (a) $T_{i}^{j}$ has a non-positive deviance (i.e., $T_{i}^{j}$ completes before its deadline), or (b) $t$ is the first time that $T_{i}^{j}$ becomes inactive (i.e., $t=\min \left(\mathrm{r}\left(T_{i}^{j+1}\right), \mathrm{d}\left(T_{i}^{j}\right)\right)$.)

If Case (a) occurs, then $T_{i}$ is negative-changeable at $t$, and $T_{i}^{j}$ is active at $t$. Hence, if $T_{i}$ increases its weight, then the only drift $T_{i}$ will incur for this reweighting event results from delaying the initiation of the event, i.e., at most $\mathrm{e}_{\max }\left(T_{i}\right)$. If $T_{i}$ decreases its weight, then delaying the reweighting event will not affect partial drift, since the enactment of the reweighting event would occur when $T_{i}^{j}$ becomes inactive, regardless of whether the initiation of the reweighting event was delayed or not.

Example (Figure 4.22). Consider the example in Figure 4.22, which depicts a oneprocessor system scheduled by NP-PEDF with two tasks: $T_{1}$, which has $\mathrm{wt}\left(T_{1}\right)=3 / 10$ and $\mathrm{e}\left(T_{1}\right)=3$; and $T_{2}$, which has $\mathrm{e}\left(T_{2}\right)=2$ and an initial weight of $1 / 5$ and initiates a weight increase to $1 / 2$ at time 4 . Inset (a) depicts the NP-PEDF schedule. Inset (b) depicts the CSW schedule. Inset (c) depicts the IDEAL schedule. Inset (d) depicts $T_{2}$ 's allocations in the CSW and IDEAL schedules. Notice that $T_{2}$ 's weight change is delayed from time 4 to time 5 because $T_{2}$ is non-preemptively executing at time 4 . As a result, $T_{2}$ is negative-changeable at time 5. Also note that, $T_{2}^{2}$ is released when $T_{2}$ 's actual allocation equals its allocation in the CSW schedule at time 7, i.e., when $T_{2}^{2}$ 's deviance equals zero.

If Case (b), mentioned earlier, occurs, then either no job of $T_{i}$ is active at $t$ or $T_{i}^{j+1}$ is active at $t$. If no job of $T_{i}$ is active at $t$, then the change is enacted immediately, and the partial drift that the task incurs from the reweighting event is a result of delaying the initiation of


Figure 4.22: A one-processor example of drift in NP-PEDF, where $T_{2}^{1}$ completes before its deadline. (a) The NP-PEDF schedule. (b) The CSW schedule. (c) The IDEAL schedule. (d) $T_{2}$ 's allocations in the CSW and IDEAL schedules.
the event, i.e., $\mathrm{e}_{\max }\left(T_{i}\right)$. If $T_{i}^{j+1}$ is active at $t$, then since $t=\min \left(\mathrm{r}\left(T_{i}^{j+1}\right), \mathrm{d}\left(T_{i}^{j}\right)\right)$, it must be the case that $\mathrm{r}\left(T_{i}^{j+1}\right)=t$. As a result, the weight change is enacted immediately and $T_{i}^{j+1}$ is released with the new weight. Hence, the only partial drift that is incurred is as a result of delaying the initiation of the reweighting event, i.e., at most $\mathrm{e}_{\max }\left(T_{i}\right)$.

Example (Figure 4.23). Consider the example in Figure 4.23, which depicts a partial NP-PEDF schedule for a task $T_{i}$, which has an initial weight of $1 / 10$ that increases to $1 / 2$ at time $t_{c}$ while the last-released job of $T_{i}$ before $t_{c}, T_{i}^{j}$, is both active and being scheduled. Note that $T_{i}^{j}$ has an execution time of four, and all jobs released after $T_{i}^{j}$ have an execution time of one. Moreover, $T_{i}^{j}$ does not complete execution until after its deadline. Inset (a) depicts the NP-PEDF schedule. Inset (b) depicts the CSW schedule. Inset (c) depicts the IDEAL schedule. Inset (d) depicts $T_{i}$ 's allocation in the CSW and IDEAL schedules. Because $T_{i}^{j}$ is not complete by its deadline, the initiation of the weight change is delayed until $t=\mathrm{d}\left(T_{i}^{j}\right)=\mathrm{r}\left(T_{i}^{j+1}\right)$. Recall that, if a weight change is initiated when $\mathrm{d}\left(T_{i}^{j}\right)=\mathrm{r}\left(T_{i}^{j+1}\right)$, then the weight change is immediately enacted and $T_{i}^{j+1}$ is released with the new weight (even though $T_{i}^{j}$ has not yet completed execution). Thus, the only source of partial drift is because the initiation of the

|  |
| :---: |
|  |  |
|  |  |



Figure 4.23: A partial schedule of a one-processor example of drift in NP-PEDF. $T_{i}^{j}$ completes after its deadline. (a) The NP-PEDF schedule. (b) The CSW schedule. (c) The IDEAL schedule. (d) $T_{i}$ 's allocations in the CSW and IDEAL schedules.
reweighting event is delayed.

In addition, similar reasoning holds for describing the impact of non-preemptability on the difference between a task's allocation in the IDEAL and PT schedules. From this reasoning, the following theorem holds.

Theorem 4.6. In a NP-PEDF scheduled system, for any task $T_{i}$, and for any interval of time $\left[t_{a}, t_{b}\right)$, let $\mathcal{Q}$ denote the number of system resets plus the number of desired weight changes for any task assigned to the same processor as $T_{i}$ that are both initiated at or before $t_{b}$ and enacted or canceled after $t_{a}$. If $T_{i}$ and $t_{b}$ satisfy one of Conditions (T-1), ...,(T-3), then the absolute drift incurred by $T_{i}$ is at most $\mathcal{Q} \cdot \mathcal{X}$, where $\mathcal{X}$ is the largest execution time of any task assigned to $P_{[q]}$ over this range.

It is important to note that Theorem 4.6 still holds even if a task $T_{i}$ is not immediately migrated from $P_{[q]}$ to $P_{[k]}$ at a system reset because its last-released job $T_{i}^{j}$ was executing when the system was reset. The reason why is because $T_{i}$ 's allocations in the IDEAL schedule are based on its assigned processor. Thus, $T_{i}$ does not incur any additional drift by remaining
on its old processor. Similarly, no other task on $P_{[q]}$ or $P_{[k]}$ incurs additional drift by $T_{i}$ remaining assigned to $P_{[q]}$ while $T_{i}^{j}$ completes its execution. It is possible that if $T_{i}$ remains on $P_{[q]}$ while $T_{i}^{j}$ completes execution, then tasks on $P_{[q]}$ will receive a guaranteed weight that is less than their desired weight; however, this difference is already captured by the MROE metric and it is not considered to contribute to drift.

Again, if a task is never assigned to an over-utilized processor, then the partial drift of a task always equals its drift. As a result, we can tighten Theorem 4.6.

Theorem 4.7. In a NP-PEDF scheduled system, for any task $T_{i}$ that is never assigned to an over-utilized processor and for any interval of time $\left[t_{a}, t_{b}\right)$, let $\mathcal{Q}$ denote the number of system resets plus the number of desired weight changes by $T_{i}$ that are both initiated at or before $t_{b}$ and enacted or canceled after $t_{a}$. If $T_{i}$ and $t_{b}$ satisfy one of Conditions (T-1), $\ldots,(T-3)$, then the absolute drift incurred by $T_{i}$ is at most $\mathcal{Q} \cdot \mathrm{e}_{\max }\left(T_{i}\right)$.

### 4.10 Adjusting PEDF for Use with any Metric

In order to allow PEDF to determine the guaranteed weights of tasks via any non-MROE metric, we must make some small changes to our adapative PEDF algorithm. Before continuing, notice that the following property holds for the MROE metric.

QS (queue stability): At any reweighting event on a processor $P_{[q]}$, the guaranteed weight of each non-reweighting task assigned to $P_{[q]}$ changes by the same multiple: $\frac{\text { old }}{\text { new }}$, where old (new) is the total desired weight of all tasks assigned to $P_{[q]}$ immediately before (after) the reweighting event.

Recall that Rules P and N function by changing the future releases and deadlines of a reweighted task. Any currently-queued job of such a task must be reinserted into the scheduler's priority queue. Since QS guarantees that the guaranteed weight of all non-reweighted tasks change by the same multiple, the jobs of these tasks already appear in the queue in the correct order, so they do not have to be reinserted into the priority queue via a rule like Rule P or N . (Moreover, if the concept of virtual time is introduced, then the deadlines of such jobs do not have to be recomputed (Stoica et al., 1996).) In fact, the primary purpose of Rules

P and N is to remove jobs from the scheduler's priority queues as needed and reinsert them into their proper places.

Example. Suppose that four tasks $T_{1}, T_{2}, T_{3}$, and $T_{4}$ with desired weights $0.5,0.2,0.2$, and 0.2 , respectively, are assigned to a processor, and at some time, $T_{4}$ changes its desired weight to 0.3 . Under the MROE metric, $T_{4}$ 's weight change causes $T_{1}$ 's guaranteed weight to change from 0.5/1.1 to $0.5 / 1.2$, and the guaranteed weight of both $T_{2}$ and $T_{3}$ to change from 0.2/1.1 to $0.2 / 1.2$. Thus, the guaranteed weights of $T_{1}, T_{2}$, and $T_{3}$ all change by the same factor, 1.1/1.2 (the old total desired weight divided by the new total desired weight).

Under metrics that are not equivalent to MROE, QS does not hold. Consider the same example as above except that the AROE metric is used. Then, $T_{1}$ 's guaranteed weight changes from $0.5-(1.1-1)=0.4$ to $0.5-(1.2-1)=0.3$, while the guaranteed weights of both $T_{2}$ and $T_{3}$ remain at 0.2 . Thus, $T_{1}$ and $T_{2}$ (as well as $T_{1}$ and $T_{3}$ ) change guaranteed weights by a different multiple. As a result, $T_{1}$ (or both $T_{2}$ and $T_{3}$ ) must change its (their) guaranteed weight by a rule (i.e., Rule P or N) that may remove and reinsert currently-queued jobs into the scheduler's priority queue.

Since, from a reweighting perspective, the primary difference between the MROE metrics and non-MROE metrics is the Property QS, and as we mentioned above, the primary objective of Rules P and N is to remove a job from the scheduler's priority queues and reinsert it into its proper place, is possible to adapt Rules P and N to work for non-MROE metrics by using them whenever a task changes its guaranteed weight (recall that under the MROE metric the Rules P and N are only used when a task changes its desired weight). Thus, in the above example under the AROE metric, when $T_{4}$ changes its desired weight both $T_{1}$ and $T_{4}$ must uses the modified Rules P and N to possibly remove jobs the scheduler's priority queues and reinsert them into their proper places. Under the MROE metric, only $T_{4}$ had to use Rules P and N .

### 4.11 Time Complexity

As noted earlier, the time complexity for PEDF to partition $N$ tasks onto $M$ processors is $O(M+N \log N)$. If we were to implement PEDF using binomial heaps, then the time complexity to make a scheduling decision on a processor $P_{[q]}$ is $O(\log n)$, where $n$ is the number of tasks assigned to $P_{[q]}$. Recall that when a task changes its desired weight using either Rule P or N , one of its jobs may be removed from its processor's priority queue and reinserted. Thus, $O(\log n)$ time is required to change a task's desired weight via Rule P or N using the MROE metric. Under non-MROE metrics, $O(n \log n)$ time is required, due to the potential need to re-enqueue jobs of non-reweighted tasks. Hence, the MROE metric has a clear advantage over the non-MROE metrics.

### 4.12 Conclusion

In this chapter, we presented several different methods for calculating the error on an overutilized processor, presented a variant of the adaptable sporadic task model that allows the guaranteed and desired weights of a task to differ, and presented rules for reweighting a task under the PEDF and NP-PEDF scheduling algorithms. In addition, we proved scheduling correctness and established tardiness and drift bounds for our reweighting rules.

We conclude this chapter with two important remarks. First, it is worth reiterating that even though PEDF does not miss a deadline by our reweighting rules, if a task $T_{i}$ is assigned to an over-utilized processor, then there could possibly be an arbitrarily large difference between $T_{i}$ 's allocations in the actual schedule and a schedule in which it receives a fraction of the system equal to its desired weight at each instant. This differs from GEDF-scheduled systems, in which the difference between $T_{i}$ 's allocations in the actual system and a system in which it receives its desired weight at each instant is bounded (if we defined job deadlines based on desired weights, then PEDF could have unbounded tardiness). Second, we do not claim that our adaptive variant of PEDF is the final word regarding partitioned reweighting schemes. However, we have tried hard to devise reasonable approaches for dealing with the fundamental limitation discussed earlier in Section 4.2 to which such schemes are subject. Thus, we believe
that our adaptive variant of PEDF is a good candidate partitioning approach.

## CHAPTER 5

 $P^{2 *}$In this chapter, we examine the issue of reweighting in Pfair-scheduled systems. Before continuing, we first review the basics of Pfair scheduling.

### 5.1 Preliminaries

As was mentioned in Section 1.2.4, in Pfair-scheduled systems, processor time is allocated in discrete time units, called quanta. The time interval $[t, t+1)$, where $t$ is a nonnegative integer, is called slot $t$. (Hence, time $t$ refers to the beginning of slot $t$.) In this chapter, all time values are assumed to indicate an integral number of quanta, unless specified otherwise.

Recall, from Section 1.2, that the function $\mathrm{A}\left(\mathcal{S}, T_{i}, t_{1}, t_{2}\right)$ denotes the allocations to the task $T_{i}$ in the schedule $\mathcal{S}$ over the range $\left[t_{1}, t_{2}\right)$. Similarly, we use $\mathrm{A}\left(\mathcal{S}, T_{i}^{[j]}, t_{1}, t_{2}\right)$ and $\mathrm{A}\left(\mathcal{S}, \tau, t_{1}, t_{2}\right)$ to denote, respectively, the total allocations to the "subtask" $T_{i}^{[j]}$ (as defined below) and to all tasks in the set $\tau$ over the range $\left[t_{1}, t_{2}\right)$. As a shorthand, we denote $\mathrm{A}\left(\mathcal{S}, T_{i}, t, t+1\right)$ as $\mathrm{A}\left(\mathcal{S}, T_{i}, t\right)$. Let $\mathcal{S}$ be the Pfair schedule of the system $\tau ;$ if $\mathrm{A}\left(\mathcal{S}, T_{i}^{[j]}, t\right)=$ 1 , then we say that $T_{i}^{[j]}$ is scheduled in slot $t$. For reference, all terms used in this chapter are listed in Table 5.1.

### 5.1.1 Periodic Pfair Scheduling

In defining notions relevant to Pfair scheduling, we limit attention (for now) to periodic tasks, all of which begin execution at time 0 , where each task's relative deadline equals its period. A periodic task $T_{i}$ with an integer period $\mathrm{p}\left(T_{i}\right)$ and an integer execution time $\mathrm{e}\left(T_{i}\right)$ has a weight

[^11]| Notation | Definition |
| :---: | :---: |
| $T_{i}^{[j]}$ | The $j^{\text {th }}$ subtask of $T_{i}$. |
| $\mathrm{wt}\left(T_{i}, t\right)$ | Weight of $T_{i}$ at time $t$. |
| wt ( $T_{i}$ ) | Weight of a task $T_{i}$ that does not change its weight. |
| Swt ( $T_{i}, t$ ) | Scheduling weight of $T_{i}$ at time $t$. |
| $\mathrm{r}\left(T_{i}^{[j]}\right)$ | Release time of $T_{i}^{j}$. |
| $\mathrm{d}\left(T_{i}^{[j]}\right)$ | Deadline of $T_{i}^{j}$. |
| $\mathrm{w}\left(T_{i}^{[j]}\right)$ | Window of $T_{i}^{[j]}$, i.e., $\left[\mathrm{r}\left(T_{i}^{[j]}\right), \mathrm{d}\left(T_{i}^{[j]}\right)\right)$. |
| $\mathrm{b}\left(T_{i}^{[j]}\right)$ | b -bit of $T_{i}^{j}$. |
| $\mathrm{C}\left(\mathcal{B}, T_{i}^{[j]}\right)$ | Time $T_{i}^{[j]}$ completes in the schedule $\mathcal{B}$. |
| $\mathrm{D}\left(T_{i}^{[j]}\right)$ | Group deadline of $T_{i}^{j}$. |
| $\mathrm{En}\left(T_{i}, t\right)$ | Last time at or before $t$ that $T_{i}$ was reset. |
| $\operatorname{ld}\left(T_{i}^{[j]}\right)$ | The index of the first subtask of $T_{i}$ after $\operatorname{En}\left(T_{i}, t\right)$. |
| $\omega\left(T_{i}^{[j]}\right)$ | Subtask associated with $\mathrm{D}\left(T_{i}^{[j]}\right)$. |
| $\theta\left(T_{i}^{[j]}\right)$ | IS separation between $T_{i}^{[j-1]}$ and $T_{i}^{[j]}$. |
| SW | Non-clairvoyant scheduling-weight scheduling algorithm. While a task is active, this algorithm allocates the task its scheduling weight at each instant. |
| $\mathcal{S W}$ | SW schedule of a task system $\tau$. |
| CSW | Clairvoyant scheduling-weight scheduling algorithm. Only allocates capacity to non-halted subtasks. |
| CSW | CSW schedule of a task system $\tau$. |
| IDEAL | Ideal scheduling algorithm. While a task is active, this algorithm allocates a task its weight at each instant. |
| I | IDEAL schedule of a task system $\tau$. |
| $\mathcal{S}$ | Actual schedule (i.e., $\mathrm{PD}^{2}$ schedule) of task system $\tau$. |
| $\mathrm{A}\left(\mathcal{B}, T_{i}^{[j]}, t_{1}, t_{2}\right)$ | Allocation to $T_{i}^{[j]}$ in the schedule $\mathcal{B}$ over $\left[t_{1}, t_{2}\right)$. |
| $\mathrm{A}\left(\mathcal{B}, T_{i}, t_{1}, t_{2}\right)$ | Allocation to $T_{i}$ in the schedule $\mathcal{B}$ over $\left[t_{1}, t_{2}\right)$. |
| $\mathrm{drift}\left(T_{i}, t\right)$ | Drift of $T_{i}: \mathrm{A}\left(\mathcal{I}, T_{i}, 0, t\right)-\mathrm{A}\left(\mathcal{C} S W, T_{i}, 0, t\right)$. |
| Ow | Scheduling weight before a reweighting event. |
| Nw | New weight after a reweighting event. |
| $\operatorname{lag}\left(\mathcal{S}, \mathcal{I}, T_{i}, t\right)$ | Lag of $T_{i}$ at time $t: \mathrm{A}\left(\mathcal{I}, T_{i}, 0, t\right)-\mathrm{A}\left(\mathcal{S}, T_{i}, 0, t\right)$. |
| $\operatorname{lag}\left(T_{i}, t\right)$ | Lag of $T_{i}$ at time $t$ when the actual and ideal schedules are implicit. |
| $\operatorname{LAG}(\mathcal{S}, \mathcal{I}, \tau, t)$ | Total lag of all tasks in $\tau$ at time $t$. |
| $\operatorname{LAG}(\tau, t)$ | Total lag of all tasks in $\tau$ at time $t$ when the actual and ideal schedules are implicit. |
| $\mathrm{X}^{(j)}$ | $j^{\text {th }}$ subtask in a chain of displacements. |
| $\left\langle\mathrm{X}^{(j)}, t_{j}, \mathrm{X}^{(j+1)}, t_{j+1}\right\rangle$ | Displacement tuple. |

Table 5.1: Brief description of the notation used in this chapter.
(or utilization) $\mathrm{wt}\left(T_{i}\right)=\mathrm{e}\left(T_{i}\right) / \mathrm{p}\left(T_{i}\right)$, where $0<\mathrm{wt}\left(T_{i}\right)<1$. We say that a task is light if its weight is in the range $(0,1 / 2)$, and heavy if its weight is in the range $[1 / 2,1)$. (For simplicity, we ignore the possibility of a task having a weight of 1 . Such tasks can be included, but at the expense of more complicated notation in the reweighting rules.)

The ideal schedule for a periodic task system allocates $\mathbf{w t}\left(T_{i}\right)$ processing time to each task in each time slot. More specifically, in the ideal schedule, $\mathcal{E}$, of the task system $\tau$, $\mathrm{A}\left(\mathcal{S}, T_{i}, t\right)=\mathrm{wt}\left(T_{i}\right)$ holds for any $T_{i} \in \tau$ and any time $t \geq 0$.

In order to compare the difference in the task $T_{i}$ 's allocations in the actual schedule $\mathcal{S}$ and the ideal schedule $\mathcal{E}$ up to time $t$, we use the function $\operatorname{lag}\left(\mathcal{S}, \mathcal{E}, T_{i}, t\right)=\mathrm{A}\left(\mathcal{E}, T_{i}, 0, t\right)-$ $\mathrm{A}\left(\mathcal{S}, T_{i}, 0, t\right)$. Additionally, we use the function $\operatorname{LAG}(\mathcal{S}, \mathcal{E}, \tau, t)=\sum_{T_{i} \in \mathcal{T}} \operatorname{lag}\left(\mathcal{S}, \mathcal{E}, T_{i}, t\right)$ to compare the differences in allocations for all tasks in the task set $\tau$ in schedules $\mathcal{S}$ and $\mathcal{E}$. We assume $\operatorname{lag}\left(\mathcal{S}, \mathcal{E}, T_{i}, 0\right)=0$. Thus, $\operatorname{LAG}(\mathcal{S}, \mathcal{E}, \tau, t)$ can be rewritten as

$$
\begin{equation*}
\operatorname{LAG}(\mathcal{S}, \mathcal{E}, \tau, t)=\operatorname{LAG}(\mathcal{S}, \mathcal{E}, \tau, t-1)+\mathrm{A}(\mathcal{E}, \tau, t-1)-\mathrm{A}(\mathcal{S}, \tau, t-1) \tag{5.1}
\end{equation*}
$$

For brevity, we denote $\operatorname{lag}\left(\mathcal{S}, \mathcal{E}, T_{i}, t\right)$ as $\operatorname{lag}\left(T_{i}, t\right)$ and $\operatorname{LAG}(\mathcal{S}, \mathcal{E}, \tau, t)$ as $\operatorname{LAG}(\tau, t)$, when $\mathcal{S}$ and $\mathcal{E}$ are well-defined and obvious. (Later, we apply (5.1) in contexts where the ideal schedule is defined to reflect changes caused by reweighting events.)

The schedule $\mathcal{S}$ is Pfair iff $\left(\forall T_{i} \in \tau, t::-1<\operatorname{lag}\left(T_{i}, t\right)<1\right)$. Informally, each task's allocation error must always be less than one quantum. These error bounds are ensured by treating each quantum of a task's execution, henceforth called a subtask, as a schedulable entity. Scheduling decisions are made only at quantum boundaries. The $j^{\text {th }}$ subtask of task $T_{i}$, denoted $T_{i}^{[j]}$, where $j \geq 1$, has an associated pseudo-release

$$
\mathrm{r}\left(T_{i}^{[j]}\right)=\left\lfloor\frac{j-1}{\mathrm{wt}\left(T_{i}\right)}\right\rfloor
$$

and pseudo-deadline

$$
\mathrm{d}\left(T_{i}^{[j]}\right)=\left\lceil\frac{j}{\mathrm{wt}\left(T_{i}\right)}\right\rceil .
$$

(For brevity, we often drop the prefix "pseudo-.") It can be shown that if each subtask $T_{i}^{[j]}$


Figure 5.1: $\mathrm{A}\left(\mathcal{I}, T_{i}^{[j]}, t\right)$ for a (a) periodic and (b) IS task $T_{1}$ of weight $5 / 16$.
is scheduled in the interval $\mathrm{w}\left(T_{i}^{[j]}\right)=\left[\mathrm{r}\left(T_{i}^{[j]}\right), \mathrm{d}\left(T_{i}^{[j]}\right)\right)$, termed its window, then $\left(\forall T_{i} \in \tau, t:\right.$ : $\left.-1<\operatorname{lag}\left(T_{i}, t\right)<1\right)$ is maintained (Baruah et al., 1996).

Example (Figure 5.1). Consider the example in Figure 5.1, which depicts the releases and deadlines for a task $T_{1}$, which has $w t\left(T_{1}\right)=5 / 16$. (This figure also depicts per-slot ideal allocations for each subtask, which are considered below.) In this example, $\mathrm{r}\left(T_{1}^{[2]}\right)=3$, $\mathrm{d}\left(T_{1}^{[2]}\right)=7$, and $\mathrm{w}\left(T_{1}^{[2]}\right)=[3,7)$. Thus, $T_{1}^{[2]}$ must be scheduled in slots 3-6. (Tasks execute sequentially, so if $T_{1}^{[1]}$ is scheduled in slot 3 , then $T_{1}^{[2]}$ must be scheduled in slots 4-6.)

### 5.1.2 The Intra-Sporadic Task Model

The intra-sporadic (IS) task model (Srinivasan and Anderson, 2006) generalizes the wellknown sporadic task model (Mok, 1983) by allowing subtasks to be released late. This extra flexibility is useful in many applications where processing steps may be delayed. Fig. 5.1(b) illustrates the Pfair windows of an IS task of weight $5 / 16$ in which the release of $T_{1}^{[2]}$ is delayed by two quanta and the release of $T_{1}^{[3]}$ is delayed by an additional quantum. Each subtask $T_{i}^{[j]}$ of an IS task has an offset, $\theta\left(T_{i}^{[j]}\right)$, that gives the amount by which its release has been delayed. For example, in Figure 5.1(b), $\theta\left(T_{1}^{[1]}\right)=0, \theta\left(T_{1}^{[2]}\right)=2$, and for $j \geq 3, \theta\left(T_{1}^{[j]}\right)=3$.

The release and deadline of a subtask $T_{i}^{[j]}$ of an IS task $T_{i}$ are defined as

$$
\begin{align*}
\mathrm{r}\left(T_{i}^{[1]}\right) & =\theta\left(T_{i}^{[1]}\right)  \tag{5.2}\\
\mathrm{r}\left(T_{i}^{[j+1]}\right) & =\left(\theta\left(T_{i}^{[j+1]}\right)-\theta\left(T_{i}^{[j+1]}\right)\right)+\mathrm{d}\left(T_{i}^{[j]}\right)-\left\lceil\frac{j}{\mathrm{wt}\left(T_{i}\right)}\right\rceil+\left\lfloor\frac{j}{\mathrm{wt}\left(T_{i}\right)}\right\rfloor  \tag{5.3}\\
\mathrm{d}\left(T_{i}^{[j]}\right) & ==\mathrm{r}\left(T_{i}^{[j]}\right)+\left\lceil\frac{j}{\mathrm{wt}\left(T_{i}\right)}\right\rceil-\left\lfloor\frac{j-1}{\mathrm{wt}\left(T_{i}\right)}\right\rfloor \tag{5.4}
\end{align*}
$$

where the offsets satisfy the property $k \geq j \Rightarrow \theta\left(T_{i}^{[k]}\right) \geq \theta\left(T_{i}^{[j]}\right)$. A subtask $T_{i}^{[j]}$ is active at time $t \operatorname{iff} \mathrm{r}\left(T_{i}^{[j]}\right) \leq t<\mathrm{d}\left(T_{i}^{[j]}\right)$, and a task $T_{i}$ is active at $t$ iff it has an active subtask at $t$. For example, in Figure 5.1(b), $T_{i}$ is active in every slot except slot 4. If $\theta\left(T_{i}^{[j+1]}\right)>\theta\left(T_{i}^{[j]}\right)$, then we say that there is an IS separation between $T_{i}^{[j]}$ and $T_{i}^{[j+1]}$. For example, in Figure 5.1(b), there is an IS separation between $T_{1}^{[1]}$ and $T_{1}^{[2]}$, as well as between $T_{1}^{[2]}$ and $T_{1}^{[3]}$. (Note that an extension of the IS model exists in which a subtask $T_{i}^{[j]}$ can become eligible before $\mathrm{r}\left(T_{i}^{[j]}\right)$ (Srinivasan and Anderson, 2006). All the results of this chapter can be easily extended to such a model, but for clarity, we do not consider this extension to the IS model.)

### 5.1.3 The $\mathrm{PD}^{2}$ Algorithm

The PD ${ }^{2}$ Pfair scheduling algorithm (Srinivasan and Anderson, 2006) is optimal for scheduling IS tasks on an arbitrary number of processors. It prioritizes subtasks on an earliest-pseudo-deadline-first (EPDF) basis, and uses two tie-breaking rules: the b-bit and the group deadline. The b-bit of the subtask $T_{i}^{[j]}$ is defined as

$$
\begin{equation*}
\mathrm{b}\left(T_{i}^{[j]}\right)=\left\lceil\frac{j}{\mathrm{wt}\left(T_{i}\right)}\right\rceil-\left\lfloor\frac{j}{\mathrm{wt}\left(T_{i}\right)}\right\rfloor . \tag{5.5}
\end{equation*}
$$

The group deadline of the subtask $T_{i}^{[j]}$ of a task where $\mathrm{wt}\left(T_{i}\right) \geq 1 / 2$ is defined as

$$
\mathrm{D}\left(T_{i}^{[j]}\right)=\left\{\begin{array}{ll}
0, & \mathrm{wt}\left(T_{i}\right)<1 / 2  \tag{5.6}\\
\theta\left(T_{i}^{[j]}\right)+\left\lceil\frac{\left\lceil\left[\frac{j}{\mathrm{wt}\left(T_{i}\right)}\right] \cdot\left(1-\mathrm{wt}\left(T_{i}\right)\right)\right]}{1-\mathrm{wt}\left(T_{i}\right)}\right], & \mathrm{wt}\left(T_{i}\right) \geq 1 / 2
\end{array} .\right.
$$

$\mathrm{b}\left(T_{i}^{[j]}\right)$ is 1 in a periodic task system (or an IS system where $\theta\left(T_{i}^{[j]}\right)=0$ for every subtask), if $T_{i}^{[j]}$,s window overlaps $T_{i}^{[j+1]}$,s, and is 0 otherwise. For example, in both insets in Figure 5.1, $\mathrm{b}\left(T_{i}^{[j]}\right)=1$ for $1 \leq i \leq 4$ and $\mathrm{b}\left(T_{i}^{[5]}\right)=0$. If two subtasks have equal deadlines, then a subtask with a b-bit of 1 is favored over one with a b-bit of 0 . Notice that, in the absence of IS separations, $\mathrm{r}\left(T_{i}^{[j+1]}\right)=\mathrm{d}\left(T_{i}^{[j]}\right)-\mathrm{b}\left(T_{i}^{[j]}\right)$. For example, in Figure 5.1(a), $\mathrm{r}\left(T_{i}^{[2]}\right)=$ $\mathrm{d}\left(T_{i}^{[1]}\right)-\mathrm{b}\left(T_{i}^{[1]}\right)=4-1=3$, and $\mathrm{r}\left(T_{i}^{[6]}\right)=\mathrm{d}\left(T_{i}^{[5]}\right)-\mathrm{b}\left(T_{i}^{[5]}\right)=16-0=16$. Also, if $\mathrm{b}\left(T_{i}^{[j]}\right)=1, \theta\left(T_{i}^{[j+1]}\right) \geq \theta\left(T_{i}^{[j]}\right)+1$, and $T_{i}^{[j+1]}$ exists, then $\mathrm{r}\left(T_{i}^{[j+1]}\right)=\mathrm{d}\left(T_{i}^{[j]}\right)$. For example, in Figure 5.1(b), $\mathrm{r}\left(T_{1}^{[3]}\right)=\mathrm{d}\left(T_{1}^{[2]}\right)$.

In a periodic task system, the group-deadline of the subtask $T_{i}^{[j]}$ of a heavy task $T_{i}$ represents the slot after which $T_{i}$ will not have two overlapping subtask windows, because some subtask has either a b-bit of 0 or a window length of three. (Note that, by (5.6), if $T_{i}$ is light, then all of its subtasks have a group deadline of zero.) If two subtasks have equal deadlines and b-bits, then subtask with a larger group deadline is favored over a subtask with a smaller group deadline. Further ties are broken arbitrarily. By breaking ties in this manner, the $\mathrm{PD}^{2}$ scheduling algorithm reduces the impact current scheduling decisions have on future ones. Notice that the subtask associated with the group deadline of $T_{i}^{[j]}, T_{i}^{[k]}$, satisfies one of the following two conditions (assuming $T_{i}$ is heavy):
$\left(\omega\right.$-i) $\mathrm{r}\left(T_{i}^{[k]}\right)>\mathrm{r}\left(T_{i}^{[j]}\right)$ and $\mathrm{d}\left(T_{i}^{[k]}\right)-\mathrm{r}\left(T_{i}^{[k]}\right)=3$.
$\left(\omega\right.$-ii) $\mathrm{r}\left(T_{i}^{[k]}\right) \geq \mathrm{r}\left(T_{i}^{[j]}\right), \mathrm{d}\left(T_{i}^{[k]}\right)-\mathrm{r}\left(T_{i}^{[k]}\right)=2$, and $\mathrm{b}\left(T_{i}^{[k]}\right)=0$.

Example (Figure 5.2). Consider the example in Figure 5.1, which depicts the releases and deadlines for the task $T_{1}$, which has $\operatorname{wt}\left(T_{1}\right)=7 / 9$. (Again, the per-slot allocations are considered later.) Notice that, for $j \in\{1,2,3\}, \mathrm{D}\left(T_{1}^{[j]}\right)=5$, and for $j \in\{4,5,6,7\}, \mathrm{D}\left(T_{1}^{[j]}\right)=$ 9. $T_{1}^{[4]}$ is the subtask associated with the group deadline of $T_{1}^{[1]}, \ldots, T_{1}^{[3]}$ and satisfies condition $(\omega-\mathrm{i}) . T_{1}^{[7]}$ is the subtask associated with the group deadline of $T_{1}^{[4]}, \ldots, T_{1}^{[7]}$ and satisfies condition ( $\omega$-ii).

We use $\omega\left(T_{i}^{[j]}\right)$ to denote the subtask that is associated with the group deadline of $T_{i}^{[j]}$.


Figure 5.2: $\mathrm{A}\left(\mathcal{I}, T_{i}^{[j]}, t\right)$ for a periodic task with weight $7 / 9$.

Specifically,

$$
\omega\left(T_{i}^{[j]}\right)= \begin{cases}\text { first subtask that satisfies }\left(\omega \text { -i) or } \left(\omega \text {-ii) for } T_{i}^{j},\right.\right. & \text { if } \mathrm{D}\left(T_{i}^{[j]}\right)>0  \tag{5.7}\\ \text { undefined, } & \text { if } \mathrm{D}\left(T_{i}^{[j]}\right)=0\end{cases}
$$

For example, in Figure 5.2, for $j \in\{1,2,3\}, \omega\left(T_{1}^{[j]}\right)=T_{1}^{[4]}$, and for $j \in\{4,5,6,7\}$, $\omega\left(T_{1}^{[j]}\right)=T_{1}^{[7]}$. Notice that, if a subtask has a b-bit of 0 , then its group deadline is its own deadline (e.g., in Figure 5.2, $\mathrm{D}\left(T_{1}^{[7]}\right)=\mathrm{d}\left(T_{1}^{[7]}\right)=9$ ); however, if a subtask has a window length of three, then its group deadline is after its deadline (e.g., $D\left(T_{1}^{[4]}\right)=9 \geq 6=\mathrm{d}\left(T_{1}^{[4]}\right)$ ).

It is easy to see that for IS task systems, $\omega\left(T_{i}^{[j]}\right)$ is defined if $T_{i}$ is heavy and undefined if $T_{i}$ is light. However, this will not necessarily hold for the adaptable IS task model, which is defined in Section 5.2.

From the definition of a group deadline it is not difficult to show that the following properties holds.
(GD-1) For any $T_{i}^{[j]}$ such that $\mathrm{D}\left(T_{i}^{[j]}\right)>0$, if for all $T_{i}^{[q]} \in\left\{T_{i}^{[j]}, \ldots, \omega\left(T_{i}^{[j]}\right)\right\}, \theta\left(T_{i}^{[q]}\right)=\theta\left(T_{i}^{[j]}\right)$
(i.e., there are no IS separations between subtasks until the group deadline), then:
(i) $\mathrm{r}\left(\omega\left(T_{i}^{[j]}\right)\right)=\mathrm{D}\left(T_{i}^{[j]}\right)-2$;
(ii) either $\mathrm{d}\left(\omega\left(T_{i}^{[j]}\right)\right)=\mathrm{D}\left(T_{i}^{[j]}\right)$ and $\mathrm{b}\left(\omega\left(T_{i}^{[j]}\right)\right)=0$ (i.e., $\omega\left(T_{i}^{[j]}\right)$ has a window length of

```
\(\mathrm{A}\left(\mathcal{I}_{1 \mathrm{~S}}, T_{i}^{[j]}, t\right)\)
    if \(\left(t<\mathrm{r}\left(T_{i}^{[j]}\right)\right) \vee\left(t \geq \mathrm{d}\left(T_{i}^{[j]}\right)\right)\) then
        \(\mathrm{A}\left(\mathcal{I}_{\mathrm{IS}}, T_{i}^{[j]}, t\right):=0\)
    else if \(t=\mathrm{r}\left(T_{i}^{[j]}\right)\) then
        if \(j=1 \vee \mathrm{~b}\left(T_{i}^{[j-1]}\right)=0\) then
            \(\mathrm{A}\left(\mathcal{I}_{1 \mathrm{~S}}, T_{i}^{[j]}, t\right):=\mathrm{wt}\left(T_{i}\right)\)
        else
            \(\mathrm{A}\left(\mathcal{I}_{\mathrm{IS}}, T_{i}^{[j]}, t\right):=\)
                        \(\mathrm{wt}\left(T_{i}\right)-\mathrm{A}\left(\mathcal{I}_{\mathrm{IS}}, T_{i}^{[j-1]}, \mathrm{d}\left(T_{i}^{[j-1]}\right)-1\right)\)
        fi
    else
        \(\mathrm{A}\left(\mathcal{I}_{\mathrm{IS}}, T_{i}^{[j]}, t\right):=\)
            \(\min \left(\operatorname{wt}\left(T_{i}\right), 1-\mathrm{A}\left(\mathcal{I}_{1 S}, T_{i}^{[j]}, 0, t\right)\right)\)
    1: fi
```

Figure 5.3: Pseudo-code defining $\mathrm{A}\left(\mathcal{I}_{\mathrm{I}}, T_{i}^{[j]}, t\right)$.
two and a b-bit of zero) or $\mathrm{d}\left(\omega\left(T_{i}^{[j]}\right)\right)=\mathrm{D}\left(T_{i}^{[j]}\right)+1$ and $\mathrm{b}\left(\omega\left(T_{i}^{[j]}\right)\right)=1$ (i.e., $\omega\left(T_{i}^{[j]}\right)$ has a window length of three).
(GD-2) If $\mathrm{D}\left(T_{i}^{[j]}\right)>0$, then $\mathrm{D}\left(T_{i}^{[j]}\right) \geq \mathrm{d}\left(T_{i}^{[j]}\right)+\mathrm{b}\left(T_{i}^{[j]}\right)$.

### 5.1.4 IS Ideal Schedule

Ideal allocations within the IS task model are defined so that the cumulative allocation for one subtask is one and the total per-slot allocation is at most the weight of the corresponding task (Srinivasan and Anderson, 2006). The total allocation to a task in a given time slot equals the total allocation to all of its subtasks in that slot. Thus, for any ideal schedule $\mathcal{E}$,

$$
\mathrm{A}\left(\mathcal{E}, T_{i}, t\right)=\sum_{T_{i}^{[j]} \in T_{i}} \mathrm{~A}\left(\mathcal{E}, T_{i}^{[j]}, t\right) .
$$

For example, in Figure 5.1(a), $\mathrm{A}\left(\mathcal{E}, T_{i}, 6\right)=\mathrm{A}\left(\mathcal{E}, T_{1}, 6\right)+\mathrm{A}\left(\mathcal{E}, T_{2}, 6\right)+\mathrm{A}\left(\mathcal{E}, T_{3}, 6\right)+\ldots=$ $0+2 / 16+3 / 16+0+\ldots=5 / 16$. Thus, per-task and per-task-set allocations in the ideal schedule $\mathcal{E}$ over an arbitrary interval can be defined by simply defining $\mathrm{A}\left(\mathcal{E}, T_{i}^{[j]}, t\right)$ for an arbitrary subtask $T_{i}^{[j]}$ and time slot $t$.

For an arbitrary IS task system $\tau$, we let $\mathcal{I}_{\text {IS }}$ denote the ideal schedule of $\tau$. $\mathrm{A}\left(\mathcal{I}_{\text {IS }}, T_{i}^{[j]}, u\right)$ can be defined using an arithmetic expression, but we have opted instead for a more intuitive
pseudo-code-based definition in Figure 5.3. The ideal IS schedule allocates each subtask $T_{i}^{[j]}$ some amount of processing time in each slot of its window. For slots other than $\mathbf{r}\left(T_{i}^{[j]}\right)$ and $\mathrm{d}\left(T_{i}^{[j]}\right)-1$, this allocation is $\mathrm{wt}\left(T_{i}\right) . T_{i}^{[j]}$, allocation in slots $\mathrm{r}\left(T_{i}^{[j]}\right)$ and $\mathrm{d}\left(T_{i}^{[j]}\right)-1$ are adjusted so that
(i) $T_{i}^{[j]}$ s entire allocation (across all slots in its window) is one.
(ii) $T_{i}^{[j]}$, s allocation in slot $\mathrm{r}\left(T_{i}^{[j]}\right)$ plus $T_{i}^{[j-1]}$,s allocation in slot $\mathrm{d}\left(T_{i}^{[j-1]}\right)-1$ equals $\mathrm{wt}\left(T_{i}\right)$ (assuming $T_{i}^{[j]}$ and $T_{i}^{[j-1]}$ exist).
(iii) $T_{i}^{[j]}$,s allocation in slot $\mathrm{d}\left(T_{i}^{[j]}\right)-1$ plus $T_{i}^{[j+1]}$,s allocation in slot $\mathrm{r}\left(T_{i}^{[j+1]}\right)$ equals wt $\left(T_{i}\right)$ (assuming $T_{i}^{[j]}$ and $T_{i}^{[j+1]}$ exist).

Examples of such allocations are given in Figure 5.1.

### 5.1.5 Dynamic Task Systems

The dynamic IS task model is an extension of the IS model in which tasks can leave and join by conditions defined in (Srinivasan and Anderson, 2005), which are stated below.
$\mathbf{J}:\left(\right.$ join condition ) A task $T_{i}$ can join at time $t$ iff the sum of the weights of all tasks after joining is at most $M$.

L: (leave condition) Let $T_{i}^{[j]}$ denote the last-scheduled subtask of $T_{i}$. If $T_{i}$ is light, then $T_{i}$ can leave at time $t$ iff $t \geq \mathrm{d}\left(T_{i}^{[j]}\right)+\mathrm{b}\left(T_{i}^{[j]}\right)$. If $T_{i}$ is heavy, then $T_{i}$ can leave at time $t$ iff $t \geq \mathrm{D}\left(T_{i}^{[j]}\right)$.

For example, in Figure 5.1(b), if $T_{1}$ were to leave after $T_{1}^{[1]}$, then $T_{1}$ could not leave until time 5 because $5=\mathrm{d}\left(T_{1}^{[1]}\right)+\mathrm{b}\left(T_{1}^{[1]}\right)=4+1$. Moreover, if $T_{1}$ were to leave after $T_{1}^{[5]}$, then $T_{i}$ could not leave until time 19 because $19=\mathrm{d}\left(T_{1}^{[5]}\right)+\mathrm{b}\left(T_{1}^{[5]}\right)=19+0$.

Theorem 5.1 (From (Srinivasan and Anderson, 2005)). PD ${ }^{2}$ correctly schedules any dynamic IS task system satisfying J and L.

By Theorem 5.1, a task may be reweighted by leaving with its old weight and rejoining with its new weight.

### 5.2 Adaptable Task Model

In this section, we introduce the adaptable IS (AIS) task model. The AIS task model is an extension of IS task model, where the weight of each task $T_{i}, \mathrm{wt}\left(T_{i}, t\right)$, is a function of time $t$. (The AIS task model is similar to the adaptable sporadic task model presented in Chapter 3, except that the AIS task model is designed for Pfair-scheduled systems.) For brevity, we use $\mathrm{wt}\left(T_{i}\right)$ to denote the value of a task $T_{i}$ that never changes its weight.

Before continuing, it is important to mention that while for periodic and IS tasks the weight of a task is bounded by the range $(0,1)$, for AIS tasks, it is possible for a task's weight to equal 0 . When a task's weight equals 0 , we assume, without loss of generality, that it has left the system and will not return.

A task $T_{i}$ changes weight or reweights at time $t+1$ if $\mathrm{wt}\left(T_{i}, t\right) \neq \mathrm{wt}\left(T_{i}, t+1\right)$. If a task $T_{i}$ changes weight at a time $t_{c}$ between the release and the deadline of some subtask $T_{i}^{[j]}$, then the following two actions may occur:
(i) If $T_{i}^{[j]}$ has not been scheduled by $t_{c}$, then $T_{i}^{[j]}$ may be "halted" at $t_{c}$.
(ii) $\mathrm{r}\left(T_{i}^{[j+1]}\right)$ may be redefined to be less than $\mathrm{d}\left(T_{i}^{[j]}\right)-\mathrm{b}\left(T_{i}^{[j]}\right)$.

In addition, if a heavy task, $T_{i}$, changes its weight at time $t_{c}$ and $T_{i}^{[j]}$ is the last-released subtask of $T_{i}$, then the following three additional actions may occur:
(iii) The window length of every subtask released in the range $\left\{r\left(T_{i}^{[j]}\right)+1, \ldots, \mathrm{D}\left(T_{i}^{[j]}\right)-2\right\}$ may have a window length of two, a b-bit of 1 , and a group deadline of $D\left(T_{i}^{[j]}\right)$ regardless of the new weight of the task.
(iv) No subtask of $T_{i}$ will be released at time $\mathrm{D}\left(T_{i}^{[j]}\right)-1$.
(v) If $T_{i}$ decreases its weight, then no task, other than $T_{i}$, can use the "freed" capacity until $\mathrm{D}\left(T_{i}^{[j]}\right)$.

As we discuss shortly, the reason why the above actions may occur is because the releases, deadlines, b-bits, and group deadlines of subtasks may change as a result of a reweighting event. The reweighting rules we present at the end of this section state the conditions under


Figure 5.4: The per-slot SW allocations for three different AIS tasks. (a) $T_{1}$ changes its weight from $3 / 19$ to $2 / 5$ and $T_{1}^{[2]}$ halts. (b) $T_{2}$ changes its weight from $3 / 19$ to $2 / 5$ and $T_{1}^{[2]}$ does not halt. (c) $T_{3}$ does not change its weight from $2 / 5$.
which the above actions may occur and the number of slots before $\mathrm{d}\left(T_{i}^{[j]}\right)-\mathrm{b}\left(T_{i}^{[j]}\right)$ that subtask $T_{i}^{[j+1]}$ can be released.

Halting. If $T_{i}^{[j]}$ is halted before it is scheduled, then it is never scheduled. (Note that a subtask can only be halted if it has not yet been scheduled in the $\mathrm{PD}^{2}$ schedule.) Since a subtask is only halted as a result of a reweighting event, if we do not have a priori knowledge of such events, then we cannot determine whether a released subtask will be halted in the future. It is important to note that we consider a subtask $T_{i}^{[j]}$ to be active at time $t$ only if $\mathrm{r}\left(T_{i}^{[j]}\right) \leq t<\mathrm{d}\left(T_{i}^{[j]}\right)$ and $T_{i}^{[j]}$ has not been halted by $t$.

Example (Figure 5.4). Consider the example in Figure 5.4, which depicts three different tasks. Inset (a) depicts a task $T_{1}$, which has an initial weight of $3 / 19$ and at time 8 both halts its current subtask and changes its weight to $2 / 5$. Inset (b) depicts a task $T_{2}$, which has an initial weight of $3 / 19$ and at time 8 changes its weight to $2 / 5$ but but does not halt the current subtask. Inset (c) depicts a periodic task $T_{3}$, which has $\mathrm{wt}\left(T_{3}\right)=2 / 5$. The dotted window lines indicate that the window would have existed if the subtask task did not reweight. Notice that, in inset (a), we have no knowledge when $T_{1}^{[2]}$ is released at time 6 that it will be halted at time 8. (We will consider the per-slot SW allocations depicted in the figure later.)

Definition 5.1 (Initiated, Enacted, Freed Capacity, and Reset). When a task reweights, there can be a difference between when it "initiates" the change, when the change
is "enacted," and when any newly-available capacity is "freed." The time at which the change is initiated is a user-defined time; the time at which the change is enacted and the capacity is freed (in the case of a weight decrease) are both dictated by a set of conditions discussed shortly. We use the scheduling weight of a task $T_{i}$ at time $t$, denoted $\operatorname{Swt}\left(T_{i}, t\right)$, to represent the "last enacted weight of $T_{i}$." Formally, $\operatorname{Swt}\left(T_{i}, t\right)$ equals $w t\left(T_{i}, u\right)$, where $u$ is the last time at or before $t$ that a weight change was enacted for $T_{i}$. (For the purposes of this definition, we assume an initial weight change occurred for $T_{i}$ when it initially joined the system.) It is important to note that, we henceforth compute subtask deadlines and releases using scheduling weights. If $T_{i}$ decreases its scheduling weight from Ow to Nw , then until the capacity has been freed, no other task can use the capacity Nw - Ow gained by decreasing $T_{i}$ 's scheduling weight. A weight change is finalized by "resetting" the corresponding task. When a task $T_{i}$ is reset at time $t$, its future releases and deadlines are changed so that it is as though the $T_{i}$ joined at time $t$. The times at which a task can be reset are described in the following reweighting rules. We use $\operatorname{En}\left(T_{i}, t\right)$ to denote the last time at or before time $t$ that $T_{i}$ was reset, and $\operatorname{Id}\left(T_{i}^{[j]}\right)$ to denote the smallest index $k$ such that $\operatorname{En}\left(T_{i}, \mathrm{r}\left(T_{i}^{[j]}\right)\right) \leq \mathrm{r}\left(T_{i}^{[k]}\right)$ holds.

Example (Figure 5.4). In Figure 5.4(b), $\operatorname{En}\left(T_{2}, t\right)=0$, for $0 \leq t<8$; for $j \in\{1,2\}$, $\operatorname{Id}\left(T_{2}^{[j]}\right)=1$; for $t \geq 8, \operatorname{En}\left(T_{2}, t\right)=8$; and for $j \in\{3,4,5\}, \operatorname{Id}\left(T_{2}^{[j]}\right)=3$. Note that, if $\operatorname{Id}\left(T_{i}^{[j]}\right)=j$, then $T_{i}^{[j]}$ is the first subtask of $T_{i}$ released at or after a weight change for $T_{i}$ has been enacted.

Example (Figure 5.5). Consider the one-processor example in Figure 5.5, which consists of three tasks: $T_{1}$, which has $\mathrm{wt}\left(T_{1}\right)=1 / 2 ; T_{2}$, which has an initial weight of $1 / 5$ that increases to $3 / 10$ at time 5 ; and $T_{3}$, which has an initial weight of $3 / 10$ that initiates a weight decrease to $1 / 5$ at time 2 that is enacted at time 5 . The capacity gained from $T_{3}$ 's weight decrease is not freed until time 5 . Thus, $T_{2}$ cannot increase its weight to $3 / 10$ until this time.

Definition 5.2 (Complete). If $\mathcal{S}$ is a schedule for the task system $\tau$, then a subtask $T_{i}^{[j]}$ of $T_{i} \in \tau$ is said to have completed by time $t$ in $\mathcal{S}$ iff $t \geq \mathrm{r}\left(T_{i}^{[j]}\right)$ and one of the following holds: (i) $T_{i}^{[j]}$ has been allocated one quantum by $t$ in $\mathcal{S}$; or (ii) $T_{i}^{[j]}$ is halted by time $t$. We use the function $\mathrm{C}\left(\mathcal{S}, T_{i}^{[j]}\right)$ to denote the earliest (integral) time at which $T_{i}^{[j]}$ is complete in $\mathcal{S}$.


Figure 5.5: A one-processor system with three tasks: $T_{1}$, which has $\mathrm{wt}\left(T_{1}\right)=1 / 2 ; T_{2}$, which has an initial weight of $1 / 5$ that increases to $3 / 10$ at time 5 ; and $T_{3}$, which has an initial weight of $3 / 10$ that initiates a weight decrease to $1 / 5$ at time 2 that is enacted at time 5 .

Example (Figure 5.6). Consider the example in Figure 5.6, which depicts a one-processor PD ${ }^{2}$ schedule for two tasks: $T_{1}$, which has $\mathrm{wt}\left(T_{1}\right)=2 / 5$, and $T_{2}$, which has an initial weight of $2 / 5$ that increases to $1 / 2$ at time 3 by halting $T_{2}^{[2]}$. In this example, $T_{1}^{[1]}$ is complete in the $\mathrm{PD}^{2}$ schedule by time 1 because it is scheduled in slot 0 , whereas $T_{2}^{[1]}$ does is not complete in the $\mathrm{PD}^{2}$ schedule until time 2 because it is not scheduled until slot 1 . Notice that, since $T_{2}^{[2]}$ is halted at time 3, it is complete at time 3 even though it is never scheduled.

For an adaptable task, the deadline, b-bit, release, and group deadline of a subtask $T_{i}^{[j]}$, respectively, are defined by (5.8)-(5.12) below, where $z=\operatorname{ld}\left(T_{i}^{[j]}\right)-1, \theta\left(T_{i}^{[j+1]}\right) \geq \theta\left(T_{i}^{[j]}\right) \geq 0$.

$$
\begin{align*}
\mathrm{d}\left(T_{i}^{[j]}\right) & =\mathrm{r}\left(T_{i}^{[j]}\right)+\left\lceil\frac{j-z}{\operatorname{Swt}\left(T_{i}, \mathrm{r}\left(T_{i}^{[j]}\right)\right)}\right\rceil-\left\lfloor\frac{j-z-1}{\operatorname{Swt}\left(T_{i}, \mathrm{r}\left(T_{i}^{[j]}\right)\right)}\right\rfloor  \tag{5.8}\\
\mathrm{b}\left(T_{i}^{[j]}\right) & =\left\lceil\frac{j-z}{\operatorname{Swt}\left(T_{i}, \mathrm{r}\left(T_{i}^{[j]}\right)\right)}\right]-\left\lfloor\frac{j-z}{\operatorname{Swt}\left(T_{i}, \mathrm{r}\left(T_{i}^{[j]}\right)\right)}\right\rfloor \tag{5.9}
\end{align*}
$$



Figure 5.6: A one-processor $\mathrm{PD}^{2}$ schedule for two tasks. The $X$ 's denote where each subtask is scheduled.

$$
\begin{align*}
\mathrm{r}\left(T_{i}^{[1]}\right)= & \theta\left(T_{i}^{[1]}\right)  \tag{5.10}\\
\mathrm{r}\left(T_{i}^{[j+1]}\right)= & \mathrm{d}\left(T_{i}^{[j]}\right)-\mathrm{b}\left(T_{i}^{[j]}\right)+\left(\theta\left(T_{i}^{[j+1]}\right)-\theta\left(T_{i}^{[j]}\right)\right)  \tag{5.11}\\
\mathrm{D}\left(T_{i}^{[j]}\right)= & \begin{cases}0, & \text { if } \operatorname{Swt}\left(T_{i}, \mathrm{r}\left(T_{i}^{[j]}\right)\right)<1 / 2 \\
& {\left[\begin{array}{l}
\left.\left[\frac{j-z}{\operatorname{Swt}\left(T_{i}, r\left(T_{i}^{[j]}\right)\right)}\right] \cdot\left(1-\operatorname{Swt}\left(T_{i}, \mathrm{r}\left(T_{i}^{[j]}\right)\right)\right)\right] \\
1-\operatorname{Swt}\left(T_{i}, \mathrm{r}\left(T_{i}^{[j]}\right)\right) \\
+\left(\theta\left(T_{i}^{[j]}\right)-\theta\left(T_{i}^{[z+1]}\right)+\mathrm{r}\left(T_{i}^{[z+1]}\right)\right)
\end{array},\right.}\end{cases} \tag{5.12}
\end{align*}
$$

It is important to note that, since reweighting events may change a subtask's release time, b-bit, deadline, and group deadline, values obtained from all of these formulas are subject to be changed by the reweighting rules presented in Section 5.4.

The above equations differ (5.2)-(5.6) in two ways. First, (5.8), (5.9), and (5.12) define the deadline, b-bit, and group-deadline of a subtask based on the scheduling weight of the task at the time the subtask is released. Second, after a task enacts a weight change, its release, deadline, b-bit, and group deadline are defined as though a new task with the new weight joined the system. (Recall that a subtask $T_{i}^{[j]}$ is the first-released subtask after the task is reset iff $\operatorname{Id}\left(T_{i}^{[j]}\right)=j$.) For example, in Figure 5.4(a), after $T_{1}$ changes its weight to $2 / 5$, the subtasks $T_{1}^{[3]}-T_{1}^{[5]}$ have similar releases, deadlines, and b-bits as the first three subtasks of the task $T_{3}$ with weight $2 / 5$ in inset (c).

```
A(SW}W,\mp@subsup{T}{i}{[j]},
    if t<r(T\mp@subsup{T}{i}{[j]})\veet\geq\textrm{C}(\mathcal{SW},\mp@subsup{T}{i}{[j]})\mathrm{ then}
        A(SW, Ti
    else if t=r (T\mp@subsup{T}{i}{[j]})\mathrm{ then}
        if j= Id (Tic
            A(SW, T}[i= [j] ,t):= Swt (Ti,t
        else
            A (\mathcal{SW},\mp@subsup{T}{i}{[j]},t):= Swt(T
                                    A}(\mathcal{S}W,\mp@subsup{T}{i}{[j-1]},\textrm{C}(\mathcal{S}W,\mp@subsup{T}{i}{[j-1]})-1
        fi
    else
            A(\mathcal{S}W,T,Ti
                min}(\operatorname{Swt}(\mp@subsup{T}{i}{},t),1-\textrm{A}(\mathcal{S}W,\mp@subsup{T}{i}{[j]},0,t)
11: fi
```

Figure 5.7: Pseudo-code defining the $\mathrm{A}\left(\mathcal{S} W, T_{i}^{[j]}, t\right)$.

### 5.3 SW Scheduling Algorithm

Just as with the adaptable sporadic task model in Chapters 3-4, in order to prove correctness and drift properties, we introduce the scheduling-weight (SW) scheduling algorithm for the AIS task model. Just as with the adaptable sporadic task model, we use $\mathcal{S} W$, to denote the SW schedule for a given system.

As with $\mathcal{I}_{\text {IS }}, \mathrm{A}\left(\mathcal{S} W, T_{i}^{[j]}, t\right)$ can be defined mathematically, but we opt instead for a pseudo-code-based definition, shown in Figure 5.7. There are three differences between the definitions of $\mathrm{A}\left(\mathcal{I}_{\mathrm{IS}}, T_{i}^{[j]}, t\right.$ ) (in Figure 5.3) and $\mathrm{A}\left(\mathcal{S} W, T_{i}^{[j]}, t\right)$ : in lines 5, 7, and 10 (in Figure 5.7), $\operatorname{Swt}\left(T_{i}, t\right)$ is used instead of $\mathrm{wt}\left(T_{i}\right)$; and in lines 1 and 7 (in Figure 5.7), $\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)$ is used instead of $\mathrm{d}\left(T_{i}^{[j]}\right)$. These two changes account for $T_{i}$ 's time-varying weight. The final change is that, in line $4, j=\operatorname{ld}\left(T_{i}^{[j]}\right)$ is used instead of $j=1$. This change causes the per-slot allocations of $T_{i}^{[z]}$, where $z=\operatorname{Id}\left(T_{i}^{[j]}\right)$, to equal that of a task that joins the system at $\mathbf{r}\left(T_{i}^{[z]}\right)$.

Example (Figure 5.4). In Figure 5.4(a), since $\operatorname{ld}\left(T_{1}^{[3]}\right)=3$, by lines 4 and 5, $\mathrm{A}\left(\mathcal{S} W, T_{1}^{[3]}, \mathrm{r}\left(T_{1}^{[3]}\right)\right)=\operatorname{Swt}\left(T_{i}, \mathrm{r}\left(T_{1}^{[3]}\right)\right)=2 / 5$, which is the same per-slot allocation that $T_{3}^{[1]}$ in Figure 5.4(c) receives at time $\mathrm{r}\left(T_{3}^{[1]}\right)$.

Before continuing, there are two important issues to note. First, in the absence of reweighting events, $\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)=\mathrm{d}\left(T_{i}^{[j]}\right)$. Second, when a task is halted via the reweighting rules
given below, it is halted in both the $\mathrm{PD}^{2}$ and $\mathcal{S} W$ schedules. Since $\mathcal{S} W$ is not clairvoyant, it will allocate "normally" to a subtask until that subtask halts, after which the subtask's per-slot allocations are zero, as with $T_{1}^{[2]}$ in Figure 5.4(a). Also note that, in Figure 5.4(b), $T_{2}^{[2]}$ is complete at time 10 , since $\mathrm{A}\left(\mathcal{S} W, T_{2}^{[2]}, 0,10\right)=1$. Several other examples of $\mathcal{S} W$ allocations are also given in Figure 5.4.

### 5.4 Reweighting Rules

In this section, we introduce three reweighting rules that improve upon leave/join reweighting by changing future subtask releases. It is important to note that, in the following rules, for a given subtask $T_{i}^{[j]}$, the value $\mathrm{d}\left(T_{i}^{[j]}\right)$ is used to determine the scheduling priority of $T_{i}^{[j]}$ in the PD ${ }^{2}$ algorithm and does not change once $T_{i}^{[j]}$ has been released. Furthermore, $\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)$ is used for some of these rules to determine the release time of $T_{i}^{[j]}$,s successor, $T_{i}^{[j+1]}$. As mentioned earlier, the completion time of a subtask cannot be accurately predicted without $a$ priori knowledge of weight changes; however, in the reweighting rules below, the completion time of a subtask in $\mathcal{S} W$ is only used after the subtask has completed, and therefore it is well-defined.

Assumptions and definitions. Let $\tau$ be a task system in which some task $T_{i}$ initiates a weight change from weight Ow to weight Nw at time $t_{c}$. If there does not exist a subtask $T_{i}^{[j]}$ of $T_{i}$ such that $\mathrm{r}\left(T_{i}^{[j]}\right) \leq t_{c}$, then the weight change is enacted immediately; otherwise, let $T_{i}^{[j]}$ denote the last-released subtask of $T_{i}$. For simplicity, we assume that the first subtask after a weight change by the corresponding task is released as early as possible. In addition, we assume that for a heavy task all subtasks released before the group deadline are released as early as possible. These assumptions can be removed at the cost of more complex notation.

The choice of which rule to apply depends on the weight of $T_{i}$, whether $t_{c} \leq \mathrm{d}\left(T_{i}^{[j]}\right)$, and whether $T_{i}^{[j]}$ has been scheduled by $t_{c}$. We say that $T_{i}$ is heavy-changeable at time $t_{c}$ from Ow to Nw if $t_{c}<\mathrm{D}\left(T_{i}^{[j]}\right)$ (recall that light tasks have a group deadline of 0 ). If $T_{i}$ is not heavy-changeable at time $t_{c}$ and $\mathrm{d}\left(T_{i}^{[j]}\right) \leq t_{c}$, then the weight change is enacted at time $\max \left(t_{c}, \mathrm{~d}\left(T_{i}^{[j]}\right)+\mathrm{b}\left(T_{i}^{[j]}\right)\right)$. If $\mathrm{D}\left(T_{i}^{[j]}\right) \leq t_{c}<\mathrm{d}\left(T_{i}^{[j]}\right)$ and $T_{i}^{[j]}$ has been scheduled by time $t_{c}$,
then we say that $T_{i}$ is negative-changeable at time $t_{c}$ from weight Ow to Nw ; otherwise, if $T_{i}^{[j]}$ has not been scheduled by $t_{c}$, then we say $T_{i}$ is positive-changeable at time $t_{c}$ from Ow to Nw. ${ }^{1}$

Because $T_{i}$ initiates a weight change at time $t_{c}, \operatorname{wt}\left(T_{i}, t_{c}\right)=\mathrm{Nw}$ holds; however, $T_{i}$ 's scheduling weight does not change until the weight change has been enacted, as specified in the rules below. Note that, if $t_{c}$ occurs between the initiation and enaction of a previous reweighting event of $T_{i}$, then the previous event is canceled, i.e., treated as if it had not occurred. As discussed later, any "error" associated with skipping a reweighting event like this is accounted for when determining drift.

### 5.4.1 Positive- and Negative-Changeable

We first describe the rules for reweighting positive- and negative-changeable tasks, then, in Section 5.4.2, describe the rule for reweighting heavy-changeable tasks. Notice that, for both of the rules below, the capacity gained by decreasing a task's weight is freed when the change is enacted.

Rule P: If $T_{i}$ is positive-changeable at time $t_{c}$ from weight Ow to $\mathrm{Nw}, T_{i}$ has released a subtask prior to $t_{c}$ and $T_{i}^{[j]}$ is the last such subtask, $t_{c}<\mathrm{d}\left(T_{i}^{[j]}\right)$, and $j>1$, then at time $t_{c}$, subtask $T_{i}^{[j]}$ is halted and at time $\max \left(t_{c}, \min \left(\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j-1]}\right), \mathrm{d}\left(T_{i}^{[j-1]}\right)\right)+\mathrm{b}\left(T_{i}^{[j-1]}\right)\right)$, $T_{i}$ 's weight change is enacted, $T_{i}$ is reset, and a new subtask $T_{i}^{[j+1]}$ is released. If $j=1$, then at time $t_{c}, T_{i}^{[j]}$ is halted, $T_{i}$ 's weight change is enacted, $T_{i}$ is reset, and a new subtask $T_{i}^{[j+1]}$ is released.

Rule N: If $T_{i}$ is negative-changeable at time $t_{c}$ from weight Ow to $\mathrm{Nw}, T_{i}$ has released a subtask prior to $t_{c}$ and $T_{i}^{[j]}$ is the last such subtask, and $t_{c}<\mathrm{d}\left(T_{i}^{[j]}\right)$, then one of two actions is taken: (i) if $\mathrm{Nw}>\mathrm{Ow}$, then the weight change is immediately enacted, and at time $\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)+\mathrm{b}\left(T_{i}^{[j]}\right), T_{i}$ is reset and a new subtask $T_{i}^{[j+1]}$ is released; (ii) otherwise, at time $\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)+\mathrm{b}\left(T_{i}^{[j]}\right), T_{i}$ is reset, the change is enacted, and a new subtask $T_{i}^{[j+1]}$ is released.

[^12]

| W | $0 \frac{3}{20} \frac{6}{20} \frac{9}{20} \frac{1}{20}$ | $\frac{15}{20} \frac{18}{20} \frac{20}{20} \frac{20}{20} \frac{20}{20}$ |  |  |  |  |  |  | 2 |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| SW | 0 $\frac{3}{20} \frac{6}{20} \frac{9}{20} \frac{12}{20}$ | $\frac{15}{20} \frac{18}{20} \frac{21}{20} \frac{24}{20} \frac{2}{2}$ |  |  |  |  | $\frac{8}{2}$ |  |  |  |  |
| I | 0 $\frac{3}{20} \frac{6}{20} \frac{9}{20} \frac{12}{20}$ | $\frac{15}{20} \frac{18}{20} \frac{21}{20} \frac{24}{20} \frac{27}{20}$ | 2 |  |  |  | $\frac{8}{2}$ |  |  |  |  |
| rift( $\mathrm{T}_{1}, \mathrm{t}$ ) | $\begin{array}{llllll}0 & 0 & 0 & 0 & 0\end{array}$ | $00 \frac{1}{20} \frac{4}{20} \frac{7}{20}$ | $\frac{1}{2}$ |  |  |  |  |  | 2 |  |  |


(a)

(b)

Figure 5.8: A four-processor illustration of Rule P under $\mathrm{PD}^{2} . C$ is a set of 19 tasks with a weight of $3 / 20$ each. (a) $T_{1}$ leaves at time 8 and $T_{2}$ joins at time 10 . (b) $T_{1}$ reweights to $1 / 2$ via rule P at slot 10 .

Both rules are extensions of Rules L and J given earlier in Section 5.1. However, the rules above exploit the specific circumstances that occur when a task changes its weight to "short circuit" Rules L and J, so that reweighting is accomplished faster. By Rule L, $T_{i}$ can leave at time $\mathrm{d}\left(T_{i}^{[k]}\right)+\mathrm{b}\left(T_{i}^{[k]}\right)$, where $T_{i}^{[k]}$ is its last-scheduled subtask. We can easily extend Rule L to show that $T_{i}$ can leave at time $\mathrm{C}\left(\mathcal{S} W, T_{i}^{[k]}\right)+\mathrm{b}\left(T_{i}^{[k]}\right)$. If task $T_{i}$ (as defined above) is positive-changeable, then its subtask $T_{i}^{[j]}$ has not been scheduled by time $t_{c}$. Such a task can be viewed as having "left" the system at time $\max \left(t_{c}, \mathrm{C}\left(\mathcal{S} W, T_{i}^{[j-1]}\right)+\mathrm{b}\left(T_{i}^{[j-1]}\right)\right)$, in which case, it can rejoin the system immediately.

Example (Figure 5.8). Consider the four-processor example in Figure 5.8, which illustrates Rule P. Inset (a) depicts a set $C$ of 19 tasks each with a weight of $3 / 20$, a task $T_{1}$ with a weight of $3 / 20$ that leaves at time 8 , and $T_{2}$ that joins at time 10 with a weight of $1 / 2$. Inset (b) depicts the same set $C$ plus a task $T_{1}$ that has an initial weight of $3 / 20$ that increases its weight to $1 / 2$ at time 10 via Rule $P$. Notice that $T_{1}$ and $T_{2}$ in inset (a) are scheduled in the same time slots as $T_{1}$ in inset (b). The top of inset (b) depict $T_{1}$ 's per-slot allocations for the SW, CSW, and IDEAL scheduling algorithms as well as $T_{1}$ 's drift. The terms CSW, IDEAL, and drift are defined later in Section 5.6.


Figure 5.9: A four-processor illustration of Rule N under $\mathrm{PD}^{2} . C$ is a set of 19 tasks with a weight of $3 / 20$ each. (a) $T_{1}$ increases its weight from $3 / 20$ to $1 / 2$ at time 10 via rule N. (b) $T_{1}$ decreases its weight from $2 / 5$ to $3 / 20$ via rule N at time 1 .

If $T_{i}$ is negative-changeable, then by Rule L , it may "leave and rejoin" with a new weight at time $\mathrm{d}\left(T_{i}^{[j]}\right)+\mathrm{b}\left(T_{i}^{[j]}\right)$, i.e., its weight change can be enacted at time $\mathrm{d}\left(T_{i}^{[j]}\right)+\mathrm{b}\left(T_{i}^{[j]}\right)$. However, if $\mathrm{C}\left(\mathcal{S} W, T_{j}\right)<\mathrm{d}\left(T_{i}^{[j]}\right)$, then $T_{i}$ may "leave and rejoin" with a new weight at time $\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)+\mathrm{b}\left(T_{i}^{[j]}\right)$.

Example (Figure 5.9). Consider the four-processor example in Figure 5.9, which illustrates Rule N. Inset (a) depicts of a set $C$ of 19 tasks each with a weight of $3 / 20$ and a task $T_{1}$ that increases its weight from $3 / 20$ to $1 / 2$ via Rule N. Inset (b) depicts of the same set $C$ and a task $T_{1}$ that decreases its weight from $2 / 5$ to $3 / 20$. The top of each inset depict $T_{1}$ 's per-slot allocations for the SW, CSW, and IDEAL scheduling algorithms as well as $T_{1}$ 's drift. Again, the terms CSW, IDEAL, and drift are defined later in Section 5.6.

Notice that the difference in Rule N between Cases (i) and (ii) is that, when a task increases its weight, the weight change is immediately enacted, whereas when a task decreases its weight, its weight change is not enacted until time $\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)+\mathrm{b}\left(T_{i}^{[j]}\right)$. Thus, $T_{i}$ 's scheduling weight is redefined at different times.


Figure 5.10: A 35 -processor system consisting of a set $A$ with nine tasks of weight $7 / 9 ; B$ with 35 tasks with an initial weight of $4 / 5$ that decreases to 0 at time 2 ; and $C$ with 35 tasks of weight $4 / 5$ that joins at time 3 . All tie-breaks go against tasks in set $A$.

### 5.4.2 Heavy-Changeable

One of the major complications with changing the weight of a heavy task is that when such a task decreases its weight, capacity cannot be freed until the group deadline of the last-released subtask of that task. If this capacity is freed sooner, then a subtask may miss its deadline.

Example (Figure 5.10). Consider the example in Figure 5.10, which depicts a 35 -processor system that is assigned 79 tasks: set $A$, which consists of nine tasks each with a of weight $7 / 9$; set $B$, which consists of 35 tasks each with an initial weight of $4 / 5$ that decreases to 0 at time 2 , with the corresponding capacity being freed at time 3 ; and set $C$, which consists of 35 tasks each with a weight $4 / 5$ that all join the system at time 3 . All tie-breaks (not resolved by by $\mathrm{PD}^{2}$ ) go against tasks in set $A$. Since the first subtasks of tasks in both $A$ and $B$ have a group deadline at time 5 and all ties are broken against tasks in $A$, the tasks in $B$ are scheduled in the slot 0 . Also, since the capacity gained by decreasing the weight of tasks in set $B$ to 0 has been freed before the group deadline of the tasks in $B$, the tasks in set $C$ can join the system at time 3 . As a result, a task in set $C$ misses its deadline at time 3 . If the capacity gained from decreasing the weight of tasks in set $B$ had not been freed until
time 5 (i.e., the group deadline of tasks in set $B$ ), then all deadline misses would have been avoided.

In light of this complication, we propose the following rule for reweighting heavychangeable tasks.

Rule H: If $T_{i}$ is heavy-changeable at time $t_{c}$ from weight Ow to Nw and $T_{i}^{[j]}$ both exists and is the last-released subtask of $T_{i}$, then the following actions occur.

1. If $\mathrm{Ow}>\mathrm{Nw}$, then the capacity of $\mathrm{Ow}-\mathrm{Nw}$ is not freed until time $\mathrm{D}\left(T_{i}^{[j]}\right)$.
2. If $T_{i}^{[j]}$ has been scheduled by time $t_{c}$, then at time $\max \left(t_{c}, \mathrm{~d}\left(T_{i}^{[j]}\right)+\mathrm{b}\left(T_{i}^{[j]}\right)\right)$, the weight change is enacted, $T_{i}$ is reset, and $T_{i}^{[j]}$ is complete in the SW schedule (i.e., it stops receiving allocations); otherwise, at time $t_{c}, T_{i}^{[j]}$ is halted, and at time $\max \left(t_{c}, \mathrm{~d}\left(T_{i}^{[j-1]}\right)+\mathrm{b}\left(T_{i}^{[j-1]}\right)\right), T_{i}$ is reset, and the weight change is enacted (if $T_{i}^{[j-1]}$ does not exist, at time $t_{c} T_{i}$ is reset and the change is enacted at $t_{c}$ ).
3. Any subtask $T_{i}^{[q]}$ released between $t_{c}$ and time $\mathrm{D}\left(T_{i}^{[j]}\right)-2$ has a window length of two, a b-bit of 1 , a group deadline of $D\left(T_{i}^{[j]}\right)$, and a release time of

$$
\mathrm{r}\left(T_{i}^{[q]}\right)=t_{e}+\left\lfloor\frac{q-1-j}{\mathrm{Nw}}\right\rfloor,
$$

where $t_{e}$ is the time the change was enacted.
4. If $T_{i}^{[\ell]}$ is the last subtask of $T_{i}$ released before $\mathrm{D}\left(T_{i}^{[j]}\right)-1$, then its successor is released and $T_{i}$ is reset at time

$$
\mathrm{r}\left(T_{i}^{[\ell+1]}\right)=\max \left(\mathrm{D}\left(T_{i}^{[j]}\right), t_{e}+\left\lfloor\frac{\ell-j}{\mathrm{Nw}}\right\rfloor\right),
$$

where $t_{e}$ is the time the change was enacted.
Rule H changes the release pattern of subtasks of $T_{i}$ from the time the weight change is enacted until time $\mathrm{D}\left(T_{i}^{[j]}\right)-1$ so that the allocation $T_{i}$ receives over that range of time is commensurate with Nw . It is important to note that the term $\mathbf{r}\left(T_{i}^{[\ell+1]}\right)=$ $\max \left(\mathrm{D}\left(T_{i}^{[j]}\right), t_{e}+\left\lfloor\frac{\ell-j}{N w}\right\rfloor+\theta\left(T_{i}^{[\ell+1]}\right)-\theta\left(T_{i}^{[\ell]}\right)\right)$ in Part 4 is used to guarantee that no subtask of $T_{i}$ is released at time $\mathrm{D}\left(T_{i}^{[j]}\right)-1$. If such a subtask did exist, then it could cause a


Figure 5.11: A one-processor system consisting of a task $T_{1}$, which has $w t\left(T_{1}\right)=1 / 10$, and $T_{2}$, which has an initial weight of $8 / 9$ and initiates a weight increase to $9 / 10$ at time 2. (a) The actual schedule. (b) The allocations in $\mathcal{S} W$ (and $\mathcal{C} S W$ ) and $T_{2}$ 's allocations in $\mathcal{I}$.
deadline to be missed. (Notice that Rule H can cause a $T_{i}$ to be reset twice. Once when the change is enacted and a second time when $T_{i}$ releases its first subtask at or after $\mathrm{D}\left(T_{i}^{[j]}\right)$.)

Example (Figure 5.11). Consider the example in Figure 5.11(a), which depicts a oneprocessor system that is assigned two tasks: $T_{1}$, which has $\operatorname{wt}\left(T_{1}\right)=1 / 10$, and $T_{2}$, which has an initial weight of $8 / 9$ and initiates a weight increase to $9 / 10$ at time 2 . Since $T_{2}^{[2]}$ is the last-released subtask of $T_{2}$ before 2 and it has been scheduled by time 2 , the change is enacted at time $\mathrm{d}\left(T_{2}^{[2]}\right)+\mathrm{b}\left(T_{2}^{[2]}\right)=4$. Moreover, over the time range $[4,8)$, the subtasks of $T_{2}$ are released in a pattern that is commensurate with a task of weight $9 / 10$, and $T_{2}^{[6]}$,s successor is released at $\mathrm{D}\left(T_{2}^{[2]}\right)=9$.

Example (Figure 5.12). Consider the example in Figure 5.12(a), which depicts a oneprocessor system that is assigned two tasks: $T_{1}$, which has an initial weight of $1 / 10$ and initiates and enacts a weight increase to $2 / 3$ at time 9 , and $T_{2}$, which has an initial weight of $8 / 9$ and initiates a weight decrease at time 2 to $1 / 3$. Since $T_{2}^{[2]}$ has been scheduled by time $2, T_{2}$ 's change is enacted at time 4 . Notice that, over the time range $[4,8)$ the subtasks are released in a pattern that is commensurate with a task of weight $1 / 3$, i.e., one subtask


Figure 5.12: A one-processor system consisting of a task $T_{1}$, which has an initial weight of $1 / 10$ and initiates a weight and enacts a weight increase to $2 / 3$ at time 9 , and $T_{2}$, which has an initial weight of $8 / 9$ and initiates a weight decrease at time 2 to $1 / 3$. (a) The actual schedule. (b) The SW (and CSW) schedule.
is released every 3 quanta, even though every subtask has a deadline two quanta after it is released. After time $\mathrm{D}\left(T_{2}^{[2]}\right)$, i.e., time $9, T_{2}$ behaves exactly like a "normal" task of weight $1 / 3$.

It is important to note that while the capacity freed from decreasing a heavy-changeable task $T_{i}$ is not available to the rest of the system until the group deadline of its last-released subtask, it remains available for $T_{i}$.

Example (Figure 5.13). Consider the example in Figure 5.13, which depicts a oneprocessor system that is assigned two tasks: $T_{1}$, which has an initial weight of $1 / 14$ and both initiates and enacts a weight increase to $1 / 4$ at time 14 , and $T_{2}$, which has an initial weight of $13 / 14$, initiates a weight decrease to $1 / 3$ at time 2 , and initiates a weight increase to $3 / 4$ at time 9 . Notice that, while $T_{1}$ cannot increase its weight (by using the capacity gained from $T_{2}$ decreasing its weight at time 2) until time $\mathrm{D}\left(T_{2}^{[2]}\right)=14, T_{2}$ can use this capacity to increase its weight from $1 / 3$ to $3 / 4$ at time 9 .

Because the weight of every task is in the range ( 0,1 ), heavy tasks (and those with a positive group deadline) have a window length of at most three, and because all tasks with a


Figure 5.13: A one-processor system consisting of a task $T_{1}$, which has an initial weight of $1 / 14$ and both initiates and enacts a weight increase to $1 / 4$ at time 14 , and $T_{2}$, which has an initial weight of $13 / 14$, initiates a weight decrease to $1 / 3$ at time 2 , and initiates a weight increase to $3 / 4$ at time 9 .
group deadline of 0 are light tasks, it is not difficult to show that the following property holds (WL) For any subtask $T_{i}^{[j]}$, if $\mathrm{D}\left(T_{i}^{[j]}\right)=0$, then $\mathrm{d}\left(T_{i}^{[j]}\right)-\mathrm{r}\left(T_{i}^{[j]}\right) \geq 3$; otherwise, $2 \leq \mathrm{d}\left(T_{i}^{[j]}\right)-$ $r\left(T_{i}^{[j]}\right) \leq 3$.

Light and heavy tasks. Recall from Section 5.1.3 that for IS light tasks, if $D\left(T_{i}^{[j]}\right)=0$, then $T_{i}$ is light, and if $\mathrm{D}\left(T_{i}^{[j]}\right)>0$, then $T_{i}$ is heavy. We can see that for AIS tasks, because of Rule H, the terms "light" and "heavy" are more difficult to distinguish. Specifically, it is possible for a task $T_{i}$ to have a scheduling weight less than $1 / 2$ and still release subtasks with a group deadline greater than zero. For this reason, for the remainder of this chapter, we do not distinguish between light and heavy tasks, but rather between tasks that release subtasks with a group deadline of zero and those that release subtasks with a group deadline greater than zero.

Before concluding this section, we make one final observation. When a heavy-changeable
task decreases its weight, it is possible that a subtask will receive allocations in the SW schedule beyond its deadline. For example, in Figure 5.12(b), $T_{2}^{[3]}$ and $T_{2}^{[4]}$ both receive allocations in $\mathcal{S} W$ after their deadlines. This behavior may occur because Rule H "artificially" decreases the window length of all subtasks until the group deadline of the last-released subtask to two; however, this does not negatively impact the scheduling correctness or drift bounds of PD ${ }^{2}$.

Throughout this chapter we use PD-PNH (respectively, PD-LJ) to refer to reweighting via Rules P, N, and H (respectively, the Rules L and J) under PD ${ }^{2}$. Unless otherwise specified, throughout this chapter, we use $\mathcal{S}$ to denote the PD-PNH schedule of a system. Since these rules change the ordering of a task in the priority queues that determine scheduling, the time complexity for reweighting one task is $O(\log N)$, where $N$ is the number of tasks in the system (assuming that binomial heaps are used to implement needed priority queues).

### 5.5 Scheduling Correctness

In this section, we prove that, in PD-PNH-scheduled systems, deadlines are not missed. The properties used throughout this section are summarized in Table 5.2. We begin with a property pertaining to cancelled weight-change events.
(C) If $T_{i}$ initiates two consecutive weight-change events at $t_{c}$ and $t_{c}^{\prime}$, where $t_{c}<t_{c}^{\prime}<t_{e}$, and $t_{e}$ denotes the time at which the change initiated at $t_{c}$ would have been enacted in the absence of other reweighting events, then $t_{e}^{\prime} \leq t_{e}$, where $t_{e}^{\prime}$ denotes the time at which the change initiated at $t_{c}$ would have been enacted in the absence of other reweighting events.

Proof of (C). Assume that $t_{c}, t_{e}, t_{c}^{\prime}$, and $t_{e}^{\prime}$ are as defined in Property (C). Notice that, if $\mathrm{d}\left(T_{i}^{[j]}\right) \leq t_{c}^{\prime}$, where $T_{i}^{[j]}$ is the last-released subtask of $T_{i}$ at $t_{c}$ (as defined in Rules P, N , and H ), then the change initiated at $t_{c}^{\prime}$ is enacted by $t_{c}^{\prime}+1$. Since $t_{c}^{\prime}<t_{e}$, this implies that $t_{e}^{\prime} \leq t_{e}$ holds. In the rest of the proof, we assume $t_{c}^{\prime}<\mathrm{d}\left(T_{i}^{[j]}\right)$ and consider the different types of reweighting events initiated at $t_{c}$.

| Property | Definition |
| :---: | :---: |
| $\begin{gathered} (\mathbf{C}) \\ {[\text { Page 199] }} \end{gathered}$ | If $T_{i}$ initiates two consecutive weight-change events at $t_{c}$ and $t_{c}^{\prime}$, where $t_{c}<t_{c}^{\prime}<t_{e}$, and $t_{e}$ denotes the time at which the change initiated at $t_{c}$ would have been enacted in the absence of other reweighting events, then $t_{e}^{\prime} \leq t_{e}$, where $t_{e}^{\prime}$ denotes the time at which the change initiated at $t_{c}$ would have been enacted in the absence of other reweighting events. |
| $\begin{gathered} \text { (X1) } \\ {[\text { Page 203] }} \\ \hline \end{gathered}$ | If a task $T_{i}$ initiates a weight change at time $t_{c}$ and $T_{i}$ releases a subtask $T_{i}^{[j]}$ at or after $t_{c}$, then the change initiated at $t_{c}$ is enacted no later than time $\mathrm{r}\left(T_{i}^{[j]}\right)$. |
| (X2) [Page 203] | If a weight change is enacted over the range $\left(\mathrm{r}\left(T_{i}^{[\ell]}\right), \mathrm{d}\left(T_{i}^{[\ell]}\right)\right]$, then that change must have been initiated over the range $\left[\mathrm{r}\left(T_{i}^{[\ell]}\right), \mathrm{d}\left(T_{i}^{[\ell]}\right)\right]$. |
| $\begin{gathered} \text { (W) } \\ {[\text { Page 203] }} \end{gathered}$ | For any time $t, \sum_{T_{i} \in \tau} \operatorname{Swt}\left(T_{i}, t\right) \leq M$, where $M$ is the number of processors. |
| (V) [Page 203] | For any two subtasks, $T_{i}^{[j]}$ and $T_{i}^{[j+1]}$, if $\mathrm{r}\left(T_{i}^{[j+1]}\right)<\mathrm{d}\left(T_{i}^{[j]}\right)-\mathrm{b}\left(T_{i}^{[j]}\right)$, then $\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right) \leq \mathrm{r}\left(T_{i}^{[j+1]}\right)$ and $\mathrm{C}\left(\mathcal{S}, T_{i}^{[j]}\right) \leq \mathrm{r}\left(T_{i}^{[j+1]}\right)$. |
| $\begin{gathered} \hline \text { (RW) } \\ {[\text { Page 210] }} \end{gathered}$ | Suppose $T_{i}$ initiates a weight change at time $t_{c} \geq \mathrm{r}\left(T_{i}^{[1]}\right)$ and $T_{i}^{[j]}$ is the last-released subtask of $T_{i}$ at $t_{c}$. If $\mathrm{r}\left(T_{i}^{[j]}\right) \leq t_{c}<\mathrm{d}\left(T_{i}^{[j]}\right)$, then $T_{i}$ is either positive-, negative-, or heavy-changeable at $t_{c}$. |
| $\begin{gathered} (\mathbf{G V}) \\ {[\text { Page 211] }} \end{gathered}$ | For the subtasks $T_{i}^{[j]}$ and $T_{i}^{[k]}$, where $j<k$, if $\mathrm{r}\left(T_{i}^{[k]}\right)<\mathrm{d}\left(T_{i}^{[j]}\right)-\mathrm{b}\left(T_{i}^{[j]}\right)$, then $\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right) \leq \mathrm{r}\left(T_{i}^{[k]}\right)$ and $\mathrm{C}\left(\mathcal{S}, T_{i}^{[j]}\right) \leq \mathrm{r}\left(T_{i}^{[k]}\right)$. |
| $\begin{gathered} \text { (GD-1) } \\ {[\text { Page 182] }} \end{gathered}$ | For any $T_{i}^{[j]}$ such that $\mathrm{D}\left(T_{i}^{[j]}\right)>0$, if for all $T_{i}^{[q]} \in\left\{T_{i}^{[j]}, \ldots, \omega\left(T_{i}^{[j]}\right)\right\}, \theta\left(T_{i}^{[q]}\right)=\theta\left(T_{i}^{[j]}\right)$ (i.e., there are no IS separations between subtasks until the group deadline), then:(i) $\mathrm{r}\left(\omega\left(T_{i}^{[j]}\right)\right)=\mathrm{D}\left(T_{i}^{[j]}\right)-2$; and (ii) either $\mathrm{d}\left(\omega\left(T_{i}^{[j]}\right)\right)=\mathrm{D}\left(T_{i}^{[j]}\right)$ and $\mathrm{b}\left(\omega\left(T_{i}^{[j]}\right)\right)=0$ or $\mathrm{d}\left(\omega\left(T_{i}^{[j]}\right)\right)=\mathrm{D}\left(T_{i}^{[j]}\right)+1$ and $\mathrm{b}\left(\omega\left(T_{i}^{[j]}\right)\right)=1$. |
| $\begin{gathered} \text { (GD-2) } \\ {[\text { Page 182] }} \end{gathered}$ | If $\mathrm{D}\left(T_{i}^{[j]}\right)>0$, then $\mathrm{D}\left(T_{i}^{[j]}\right) \geq \mathrm{d}\left(T_{i}^{[j]}\right)+\mathrm{b}\left(T_{i}^{[j]}\right)$. |
| $\begin{gathered} \text { (GD-3) } \\ {[\text { Page 211] }} \end{gathered}$ | For the subtasks $T_{i}^{[j]}$ and $T_{i}^{[k]}$, where $j<k$, if $\mathrm{r}\left(T_{i}^{[k]}\right)<\mathrm{d}\left(T_{i}^{[j]}\right)-\mathrm{b}\left(T_{i}^{[j]}\right)$, then $\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right) \leq \mathrm{r}\left(T_{i}^{[k]}\right)$ and $\mathrm{C}\left(\mathcal{S}, T_{i}^{[j]}\right) \leq \mathrm{r}\left(T_{i}^{[k]}\right)$. |
| $\begin{gathered} \text { (AF1) } \\ {[\text { Page 211] }} \end{gathered}$ | For all $t \geq 0, \mathrm{~A}\left(\mathcal{S} W, T_{i}, t\right) \leq \operatorname{Swt}\left(T_{i}, t\right)$. |
| $\begin{gathered} \text { (AF2) } \\ {[\text { Page 211] }} \end{gathered}$ | For any present subtask $T_{i}^{[j]}$ of the task $T_{i} \in \tau$ and its successor $T_{i}^{[k]}$, if $\mathrm{b}\left(T_{i}^{[j]}\right)=$ 1 and $\mathrm{r}\left(T_{i}^{[k]}\right) \geq \mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)$, then $\mathrm{A}\left(\mathcal{S} W, T_{i}, \mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)-1, \mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)+1\right) \leq$ $\operatorname{Swt}\left(T_{i}, \mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)\right)$. |
| (AF3) <br> [Page 211] | For any subtask $T_{i}^{[j]}, \mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right) \leq \mathrm{d}\left(T_{i}^{[j]}\right)$. |
| $\begin{gathered} \text { (AF4) } \\ {[\text { Page 212] }} \end{gathered}$ | For any subtask $T_{i}^{[j]}$ and any time $t$, if $t<\mathrm{r}\left(T_{i}^{[j]}\right)$ or $t \geq \mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)$, then $\mathrm{A}\left(\mathcal{S} W, T_{i}^{[j]}, t\right)=0$. |
| $\begin{gathered} \text { (AF5) } \\ {[\text { Page 212] }} \end{gathered}$ | For any present subtask $T_{i}^{[j]}$, let $T_{i}^{[k]}$ be its next present successor $T_{i}^{[k]}$ (if it exists). If $T_{i}^{[k]}$ does not exist, and $\mathrm{D}\left(T_{i}^{[j]}\right)>0$, then for any $t$ such that $\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)-1 \leq t \leq$ $\mathrm{D}\left(T_{i}^{[j]}\right), \mathrm{A}\left(\mathcal{S} W, T_{i}, \mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)-1, t+1\right) \leq \operatorname{Swt}\left(T_{i}, t\right)$ holds. Otherwise, if $T_{i}^{[k]}$ exists, $\mathrm{D}\left(T_{i}^{[j]}\right)>0$, and $\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right) \leq \mathrm{r}\left(T_{i}^{[k]}\right)$, then for any $t$ such that $\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)-1 \leq$ $t \leq \min \left(\mathrm{r}\left(T_{i}^{[k]}\right), \mathrm{D}\left(T_{i}^{[j]}\right)-1\right), \mathrm{A}\left(\mathcal{S} W, T_{i}, \mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)-1, t+1\right) \leq \operatorname{Swt}\left(T_{i}, t\right)$ holds. |
| $\begin{gathered} \hline \text { (T1) } \\ {[\text { Page 213] }} \end{gathered}$ | $\tau$ misses a deadline under PD-PNH at $t_{d}$. |
| $\begin{gathered} \text { (T2) } \\ {[\text { Page 214] }} \end{gathered}$ | No task system satisfying (T1) has fewer present subtasks in $\left[0, t_{d}\right.$ ) than $\tau$. |

Table 5.2: Summary of properties used in Section 5.5.

Positive-changeable. We begin by considering the case wherein $T_{i}$ is positivechangeable at $t_{c}$. In this case, the change initiated at $t_{c}$ is enacted at time $t_{e}=$ $\max \left(t_{c}, \min \left(\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j-1]}\right), \mathrm{d}\left(T_{i}^{[j-1]}\right)\right)+\mathrm{b}\left(T_{i}^{[j-1]}\right)\right)$. If $t_{e}=t_{c}$, then since $t_{c}<t_{c}^{\prime}<t_{e}, t_{c}^{\prime}$ cannot exist. Thus, we assume $t_{c}<t_{e}$, which implies $t_{e}=\min \left(\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j-1]}\right), \mathrm{d}\left(T_{i}^{[j-1]}\right)\right)+\mathrm{b}\left(T_{i}^{[j-1]}\right)>$ $t_{c}$. Since $T_{i}$ is positive-changeable at $t_{c}, T_{i}^{[j]}$ is halted at $t_{c}$ and no successor subtask can be released until the change initiated at $t_{c}$ (or a future change) has been enacted. Hence, since $t_{c}^{\prime}<$ $t_{e}, T_{i}^{[j+1]}$ is not released until at or after $t_{c}^{\prime}$. Thus, since $t_{c}<t_{c}^{\prime}, T_{i}^{[j]}$ is the last-released subtask of $T_{i}$ at $t_{c}^{\prime}$. Because, by Rule P, $T_{i}^{[j]}$ was halted at $t_{c}<t_{c}^{\prime}$ and (as we assumed at the beginning of the proof) $t_{c}^{\prime}<\mathrm{d}\left(T_{i}^{[j]}\right), T_{i}$ is therefore positive-changeable at $t_{c}^{\prime}$. Thus, by Rule P , the change initiated at $t_{c}^{\prime}$ is enacted at time $t_{e}^{\prime}=\max \left(t_{c}^{\prime}, \min \left(\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j-1]}\right), \mathrm{d}\left(T_{i}^{[j-1]}\right)\right)+\mathrm{b}\left(T_{i}^{[j-1]}\right)\right)$. Again, since $t_{c}^{\prime}<t_{e}$ and $t_{e}=\min \left(\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j-1]}\right), \mathrm{d}\left(T_{i}^{[j-1]}\right)\right)+\mathrm{b}\left(T_{i}^{[j-1]}\right)$, it follows that $t_{e}^{\prime}=\min \left(\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j-1]}\right), \mathrm{d}\left(T_{i}^{[j-1]}\right)\right)+\mathrm{b}\left(T_{i}^{[j-1]}\right)=t_{e}$.

Decreasing-weight negative changeable. We next consider the case wherein $T_{i}$ is decreasing-weight negative-changeable at $t_{c}$. By Rule N , such a change initiated at $t_{c}$ will be enacted at time $t_{e}=\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)+\mathrm{b}\left(T_{i}^{[j]}\right)$. Since $T_{i}$ is negative-changeable at $t_{c}$, no subtask can be released until the change that was initiated at $t_{c}$ (or a future change) has been enacted. Hence, since $t_{c}^{\prime}<t_{e}, T_{i}^{[j+1]}$ is not released until at or after $t_{c}^{\prime}$. Thus, since $t_{c}<t_{c}^{\prime}, T_{i}^{[j]}$ is the last-released subtask of $T_{i}$ at $t_{c}^{\prime}$. Because $T_{i}^{[j]}$ is scheduled before $t_{c}<t_{c}^{\prime}$, and (as we assumed at the beginning of the proof) $t_{c}^{\prime}<\mathrm{d}\left(T_{i}^{[j]}\right), T_{i}$ is negative-changeable at $t_{c}^{\prime}$. If the change at $t_{c}^{\prime}$ is a decreasing-weight event, then by Rule N , it is enacted at time $t_{e}^{\prime}=\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)+\mathrm{b}\left(T_{i}^{[j]}\right)=t_{e}$.

Increasing-weight negative-changeable. Notice that, if $T_{i}$ at $t_{c}$ is increasing-weight negative-changeable, then the change initiated at $t_{c}$ is enacted immediately, i.e., $t_{c}=t_{e}$. Thus, since $t_{c}<t_{c}^{\prime}<t_{e}$, if $T_{i}$ at $t_{c}$ is increasing-weight negative-changeable, then $t_{c}^{\prime}$ cannot exist.

Heavy-changeable. Notice that, if $T_{i}$ is heavy-changeable at $t_{c}$, then $\mathrm{D}\left(T_{i}^{[j]}\right)>t_{c}$ holds, and, by Part 2 of Rule H , if $T_{i}^{[j]}$ has been scheduled by $t_{c}$, then $t_{e}=\max \left(t_{c}, \mathrm{~d}\left(T_{i}^{[j]}\right)+\mathrm{b}\left(T_{i}^{[j]}\right)\right)$,
and if $T_{i}^{[j]}$ has not been scheduled by $t_{c}$, then $t_{e}=\max \left(t_{c}, \mathrm{~d}\left(T_{i}^{[j-1]}\right)+\mathrm{b}\left(T_{i}^{[j-1]}\right)\right)\left(\right.$ or $t_{e}=t_{c}$ if $T_{i}^{[j-1]}$ does not exist). If $t_{c}=t_{e}$, then since $t_{c}<t_{c}^{\prime}<t_{e}, t_{c}^{\prime}$ cannot exist. Thus, assume that $t_{c}<t_{e}$.

We first consider the case where $T_{i}^{[j]}$ has been scheduled before $t_{c}$. In this case, since we have assumed that $t_{c}<t_{e}$, it follows that $t_{e}=\mathrm{d}\left(T_{i}^{[j]}\right)+\mathrm{b}\left(T_{i}^{[j]}\right)>t_{c}$. Thus, in this case,

$$
\begin{equation*}
t_{c}^{\prime} \in\left(t_{c}, \mathrm{~d}\left(T_{i}^{[j]}\right)+\mathrm{b}\left(T_{i}^{[j]}\right)\right), \tag{5.13}
\end{equation*}
$$

and (by Part 3 of Rule H ), $\mathrm{r}\left(T_{i}^{[j+1]}\right) \geq t_{e}$. Thus, it follows that at $t_{c}^{\prime}>t_{c}, T_{i}^{[j]}$ is the last-released subtask of $T_{i}$.

By Property (GD-2), $\mathrm{D}\left(T_{i}^{[j]}\right) \geq \mathrm{d}\left(T_{i}^{[j]}\right)+\mathrm{b}\left(T_{i}^{[j]}\right)$. Thus, since (as we assumed at the beginning of the proof) $t_{c}^{\prime}<\mathrm{d}\left(T_{i}^{[j]}\right)$, it follows that $t_{c}^{\prime}<\mathrm{d}\left(T_{i}^{[j]}\right) \leq \mathrm{D}\left(T_{i}^{[j]}\right)$. Thus, since $T_{i}^{[j]}$ is the last-released subtask of $T_{i}$ at $t_{c}^{\prime}$ and $t_{c}^{\prime}<\mathrm{D}\left(T_{i}^{[j]}\right), T_{i}$ is heavy-changeable at $t_{c}^{\prime}$. Hence, by Step 2 of Rule $\mathrm{H}, t_{e}^{\prime}=\max \left(\mathrm{d}\left(T_{i}^{[j]}\right)+\mathrm{b}\left(T_{i}^{[j]}\right), t_{c}^{\prime}\right)$. Since, by (5.13), $t_{c}^{\prime}<\mathrm{d}\left(T_{i}^{[j]}\right)+\mathrm{b}\left(T_{i}^{[j]}\right)$, it follows that $t_{e}^{\prime}=\mathrm{d}\left(T_{i}^{[j]}\right)+\mathrm{b}\left(T_{i}^{[j]}\right)=t_{e}$.

We now consider the case where $T_{i}^{[j]}$ has not been scheduled before $t_{c}$. Since we have assumed that $t_{c}<t_{e}$, if $T_{i}^{[j]}$ has not been scheduled by $t_{c}$, then $t_{c}<t_{e}=\mathrm{d}\left(T_{i}^{[j-1]}\right)+\mathrm{b}\left(T_{i}^{[j-1]}\right)$. Thus, in this scenario,

$$
\begin{equation*}
t_{c}^{\prime} \in\left(t_{c}, \mathrm{~d}\left(T_{i}^{[j-1]}\right)+\mathrm{b}\left(T_{i}^{[j-1]}\right)\right), \tag{5.14}
\end{equation*}
$$

and (by Step 3 of Rule H$), \mathrm{r}\left(T_{i}^{[j+1]}\right) \geq t_{e}=\mathrm{d}\left(T_{i}^{[j-1]}\right)+\mathrm{b}\left(T_{i}^{[j-1]}\right)$. Thus, it follows that at $t_{c}^{\prime}>t_{c}, T_{i}^{[j]}$ is the last-released subtask of $T_{i}$.

By Property (GD-2), $\mathrm{D}\left(T_{i}^{[j]}\right) \geq \mathrm{d}\left(T_{i}^{[j]}\right)+\mathrm{b}\left(T_{i}^{[j]}\right)$. Thus, since (as we assumed at the beginning of the proof) $t_{c}^{\prime}<\mathrm{d}\left(T_{i}^{[j]}\right)$, it follows that $t_{c}^{\prime}<\mathrm{D}\left(T_{i}^{[j]}\right)$. Thus, since $T_{i}^{[j]}$ is the lastreleased subtask of $T_{i}$ at $t_{c}^{\prime}$ and $t_{c}^{\prime}<\mathrm{D}\left(T_{i}^{[j]}\right), T_{i}$ is heavy-changeable at $t_{c}^{\prime}$. Moreover, since $T_{i}^{[j]}$ had not been scheduled at $t_{c}$, it follows by Part 2 of Rule H , that $T_{i}^{[j]}$ was halted at $t_{c}$. Thus, $T_{i}^{[j]}$ is not scheduled by $t_{c}^{\prime}>t_{c}$. Thus, by Part 2 of Rule $\mathrm{H}, t_{e}^{\prime}=\max \left(\mathrm{d}\left(T_{i}^{[j-1]}\right)+\mathrm{b}\left(T_{i}^{[j-1]}\right), t_{c}^{\prime}\right)$. Since, by (5.14), $t_{c}^{\prime}<\mathrm{d}\left(T_{i}^{[j-1]}\right)+\mathrm{b}\left(T_{i}^{[j-1]}\right)$, it follows that $t_{e}^{\prime}=\mathrm{d}\left(T_{i}^{[j-1]}\right)+\mathrm{b}\left(T_{i}^{[j-1]}\right)=t_{e}$.

Initiation and enactment properties. We now state two properties about the relationship between the initiation and enactment of reweighting events.
(X1) If a task $T_{i}$ initiates a weight change at time $t_{c}$ and $T_{i}$ releases a subtask $T_{i}^{[j]}$ at or after $t_{c}$, then the change initiated at $t_{c}$ is enacted no later than time $r\left(T_{i}^{[j]}\right)$.
(X2) If a weight change is enacted over the range $\left(\mathrm{r}\left(T_{i}^{[\ell]}\right), \mathrm{d}\left(T_{i}^{[\ell]}\right)\right]$, then that change must have been initiated over the range $\left[\mathrm{r}\left(T_{i}^{[\ell]}\right), \mathrm{d}\left(T_{i}^{[\ell]}\right)\right]$.

Given that Property (C) guarantees that no sequence of reweighting events can delay the next weight-change enactment after a weight change has been initiated, and that the Rules $\mathrm{N}, \mathrm{P}$, and H guarantee a subtask is released within one quantum of a weight change enactment, the subtask $T_{i}^{[j]}$ (as defined in Property (X1)) will eventually be released, if it exists. Thus, from the definitions of Rules $\mathrm{P}, \mathrm{N}$, and H, Property (X1) should be fairly intuitive. Property (X2) is implied by Property (X1).

When Srinivasan and Anderson (Srinivasan and Anderson, 2005) proved the scheduling correctness of PD-LJ for an IS task system, they assumed that the weight of all tasks is at most $M$ and utilized the property that in an IS task system the windows for any subtask $T_{i}^{[j]}$ and its successor $T_{i}^{[j+1]}$ do not "overlap" by more than $\mathrm{b}\left(T_{i}^{[j]}\right)$ quanta, i.e., $\mathrm{d}\left(T_{i}^{[j]}\right)-\mathrm{b}\left(T_{i}^{[j]}\right) \leq \mathrm{r}\left(T_{i}^{[j+1]}\right)$. However, this property can be weakened without affecting most of their proof, so that their proof can be applied to an AIS task system. Specifically, their proof can be used to establish the scheduling correctness of PD-PNH for any AIS task system $\tau$, if the following conditions hold, which parallel the assumption that the weight of all tasks is at most $M$ and the property that in IS system $\mathrm{d}\left(T_{i}^{[j]}\right)-\mathrm{b}\left(T_{i}^{[j]}\right) \leq \mathrm{r}\left(T_{i}^{[j+1]}\right)$. (In these properties, we denote the PD-PNH schedule of $\tau$ as $\mathcal{S}$.)
(W) For any time $t, \sum_{T_{i} \in \tau} \operatorname{Swt}\left(T_{i}, t\right) \leq M$, where $M$ is the number of processors.
(V) For any two subtasks, $T_{i}^{[j]}$ and $T_{i}^{[j+1]}$, if $\mathrm{r}\left(T_{i}^{[j+1]}\right)<\mathrm{d}\left(T_{i}^{[j]}\right)-\mathrm{b}\left(T_{i}^{[j]}\right)$, then $\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right) \leq$ $\mathrm{r}\left(T_{i}^{[j+1]}\right)$ and $\mathrm{C}\left(\mathcal{S}, T_{i}^{[j]}\right) \leq \mathrm{r}\left(T_{i}^{[j+1]}\right)$.

Since Property (W) can be satisfied by policing weight-change requests, we focus our attention on showing that $\mathcal{S}$ and $\mathcal{S} W$ satisfy Property (V).

Example (Figures 5.8 and 5.9). In Figure 5.8(b), $\mathbf{d}\left(T_{1}^{[2]}\right)+\mathbf{b}\left(T_{1}^{[2]}\right)=15$ and $\mathbf{r}\left(T_{1}^{[3]}\right)=10$. Notice that, by Rule P, $T_{1}^{[2]}$ is halted (and hence complete) in both the SW and PD ${ }^{2}$ schedules when the change is enacted at time $10=\mathrm{r}\left(T_{1}^{[3]}\right)$. In Figure 5.9(a), $\mathrm{d}\left(T_{1}^{[2]}\right)+\mathrm{b}\left(T_{1}^{[2]}\right)=15$ and $\mathrm{r}\left(T_{1}^{[3]}\right)=12$. Notice that, $T_{1}^{[2]}$ has been scheduled (and is therefore complete) before time $10<\mathrm{r}\left(T_{1}^{[3]}\right)$. Further, $T_{1}^{[3]}$ is released once $T_{1}^{[2]}$ has received one unit of allocation in the SW schedule, and as a result $T_{1}^{[2]}$ is complete in the SW schedule by $\mathrm{r}\left(T_{1}^{[3]}\right)$.

Proof of (V). Referring to Property (V), let $T_{i}^{[j]}$ be some subtask such that $\mathrm{d}\left(T_{i}^{[j]}\right)-$ $\mathrm{b}\left(T_{i}^{[j]}\right)>\mathrm{r}\left(T_{i}^{[j+1]}\right)$. By the definition of a subtask release, if $\mathrm{d}\left(T_{i}^{[j]}\right)-\mathrm{b}\left(T_{i}^{[j]}\right)>\mathrm{r}\left(T_{i}^{[j+1]}\right)$, then $T_{i}$ enacted a weight change $t_{e} \in\left(\mathrm{r}\left(T_{i}^{[j]}\right), \mathrm{r}\left(T_{i}^{[j+1]}\right)\right]$; otherwise, we would have $\mathrm{d}\left(T_{i}^{[j]}\right)-\mathrm{b}\left(T_{i}^{[j]}\right) \leq$ $\mathrm{r}\left(T_{i}^{[j+1]}\right)$, by (5.11). Since a weight change is enacted in the range $\left(\mathrm{r}\left(T_{i}^{[j]}\right), \mathrm{r}\left(T_{i}^{[j+1]}\right)\right]$, by Property (X2) a change must have been initiated in the range $\left[\mathrm{r}\left(T_{i}^{[j]}\right), \mathrm{r}\left(T_{i}^{[j+1]}\right)\right]$. Without loss of generality, let $t_{c}$ be the earliest time in this range that $T_{i}$ initiates a weight change.

At time $t_{c}, T_{i}$ is positive-, negative-, or heavy-changeable. If at $t_{c}, T_{i}$ is positivechangeable, then by Rule $\mathrm{P}, T_{i}^{[j]}$ is complete by $t_{c} \leq t_{e} \leq \mathrm{r}\left(T_{i}^{[j+1]}\right)$ in both $\mathcal{S}$ and $\mathcal{S} W$. If $T_{i}$ is negative-changeable at $t_{c}$, then $T_{i}^{[j]}$ has been scheduled in $\mathcal{S}$ before $t_{c}$, and hence, $T_{i}^{[j]}$ is complete by $t_{c} \leq \mathrm{r}\left(T_{i}^{j+1}\right)$ in $\mathcal{S}$. Furthermore, in this case $T_{i}^{[j+1]}$ is not released until time $\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)+\mathrm{b}\left(T_{i}^{[j]}\right)$. If at $t_{c}, T_{i}$ is heavy-changeable and $T_{i}^{[j]}$ has not been scheduled by $t_{c}$, then by Rule $\mathrm{H}, T_{i}^{[j]}$ is complete by $t_{c} \leq t_{e} \leq \mathrm{r}\left(T_{i}^{[j+1]}\right)$ in both $\mathcal{S}$ and $\mathcal{S} W$. If at $t_{c}, T_{i}$ is heavy-changeable and $T_{i}^{[j]}$ has been scheduled by $t_{c}$, then by Rule $\mathrm{H}, T_{i}^{[j]}$ is complete by $t_{e} \leq \mathrm{r}\left(T_{i}^{[j+1]}\right)$ in both $\mathcal{S}$ and $\mathcal{S} W$.

We now prove that PD-PNH correctly schedules any AIS task system that satisfies Property (W). Note that, this proof is only a slight modification of the correctness proof for PD-LJ originally presented by Srinivasan and Anderson in (Srinivasan and Anderson, 2005).

Before proving that PD-PNH correctly schedules any AIS task system, we introduce some basic concepts and properties that are useful in the proof. We begin by introducing the "adaptive generalized intra-sporadic" task model. After this, we introduce the notion of a "displacement." Lastly, we introduce some properties and definitions pertaining to PD-PNH.


Figure 5.14: An illustration of the AGIS task model. The dashed window lines indicate that a subtask is absent.

### 5.5.1 The AGIS Task Model

To prove the scheduling correctness of PD-PNH in an AIS system, we consider an extension of the AIS task model called the adaptive generalized intra-sporadic (AGIS) task model. The AGIS model generalizes the AIS model by allowing some subtasks to be "absent." An absent subtask is never scheduled; however, such subtasks are considered to be part of a given task system, and as such, they have both releases and deadlines. If a subtask is not absent then we say that it is present. In an AGIS task system, $T_{i}^{[j]}$ is $T_{i}^{[k]}$,s predecessor (and $T_{i}^{[k]}$ is $T_{i}^{[j]}$,s successor) iff $T_{i}^{[j]}$ and $T_{i}^{[k]}$ are both present and there are no present subtasks that have an index between $j$ and $k$. The per-slot allocations to a subtask $T_{i}^{[j]}$ in the AGIS variant of an $\mathcal{S} W$ schedule are the same as in the AIS variant, except that if a subtask is absent, then its per-slot allocation is zero in all slots.

Example (Figure 5.14). Consider the example depicted in Figure 5.14, which consists of one task, $T_{1}$, which has $\operatorname{wt}\left(T_{1}\right)=5 / 16$, and has one absent subtask, $T_{1}^{[3]} . T_{1}^{[2]}$ is $T_{1}^{[4]}$,s predecessor, and $T_{1}^{[4]}$ is $T_{1}^{[2]}$ s successor. The per-slot allocations to all subtasks except $T_{1}^{[3]}$ are the same as in an AIS system, and $T_{1}^{[3]}$,s per-slot allocation is zero for each slot.

Example (Figure 5.15). As a second example, consider Figure 5.15, which consists of one task, $T_{1}$, which has $w t\left(T_{1}\right)=7 / 9$. In inset (a), no tasks are absent. In inset (b), $T_{1}^{[2]}, T_{1}^{[3]}$, and $T_{1}^{[5]}$ are absent.

Throughout this section, we use $\operatorname{LAG}(\tau, t)$ to denote $\operatorname{LAG}(\mathcal{S}, \mathcal{S} W, \tau, t)$ and $\operatorname{lag}\left(T_{i}, t\right)$ to denote $\operatorname{lag}\left(\mathcal{S}, \mathcal{S} W, T_{i}, t\right)$, where $\mathcal{S}$ is the PD-PNH schedule of a task system $\tau$.


Figure 5.15: $\mathrm{A}\left(\mathcal{S} W, T_{i}^{[j]}, t\right)$ for a periodic task with weight 7/9. (a) No subtasks are absent. (b) $T_{1}^{[2]}, T_{1}^{[3]}$, and $T_{1}^{[5]}$ are absent.

Completed. Since, under the AGIS task model, absent tasks never receive any allocations, by the definition of completed, presented earlier, such a subtask would never complete unless it was halted. Therefore, we amend the definition of completed so that an absent subtask $T_{i}^{[j]}$ is considered to be complete in all schedules as soon as it is released, i.e., $\mathrm{C}\left(\mathcal{S}, T_{i}^{[j]}\right)=$ $\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)=\mathrm{r}\left(T_{i}^{[j]}\right)$. For example, in Figure 5.14, $\mathrm{C}\left(\mathcal{S} W, T_{1}^{[3]}\right)=\mathrm{r}\left(T_{1}^{[3]}\right)=7$.

Absent subtasks and reweighting. Notice that, if a task $T_{q}$ initiates a weight change at time $t_{c}$, and the last-released subtask of $T_{q}, T_{q}^{[k]}$, is absent, then that subtask has not yet been scheduled. Therefore, if $t_{c}<\mathrm{d}\left(T_{q}^{[k]}\right)$ and $T_{q}$ is not heavy-changeable at $t_{c}$, then $T_{q}$ is positive-changeable at $t_{c}$. Similarly, if $T_{q}$ is heavy-changeable at $t_{c}$, then $T_{q}^{[k]}$ is "halted" and its successor (which may also be absent) is released at the appropriate time. In such cases, an absent subtask is considered to be "halted," even though it was never eligible to be scheduled.

Example (Figure 5.16). Consider the example in Figure 5.16, which depicts the impact of the reweighting rules on an AGIS task $T_{1}$, which changes its weight from $3 / 19$ to $2 / 5$. In


Figure 5.16: The impact of reweighting on an AGIS task. (a) $T_{1}$ changes its weight via Rule P. (b) $T_{1}$ changes its weight via Rule N. (c) $T_{1}^{[2]}$ is absent when $T_{1}$ is changed causing it to change via Rule P .
inset (a), $T_{2}$ changes its weight via Rule P causing $T_{1}^{[2]}$ to be halted at time 8 . In inset (b), $T_{1}$ changes its weight via Rule N causing $T_{1}^{[2]}$ to be complete at time 10 . In inset (c), $T_{1}^{[2]}$ is absent, and so $T_{1}$ is positive-changeable at time 8 . Notice that, in inset (a), the subtask $T_{1}^{[2]}$ is halted at time 8 , so $\mathrm{C}\left(\mathcal{S} W, T_{1}^{[2]}\right)=\mathrm{r}\left(T_{1}^{[2]}\right)=8$. In contrast, in inset $(\mathrm{b})$, the subtask $T_{i}^{[2]}$, is never halted, and therefore $\mathrm{C}\left(\mathcal{S} W, T_{1}^{[2]}\right)=10$. Also note that, in inset (c), when $T_{1}$ changes its weight via Rule P at time $8, T_{1}^{[2]}$ is considered to be halted, even though it is absent.

### 5.5.2 Displacements

We now introduce the notions of an "instance" of a task system and a task "displacement." An instance of a task system is obtained by specifying a unique assignment of release times for each subtask and weight changes for each task.


Figure 5.17: A one-processor system consisting of a task $T_{1}$ with an initial weight of $1 / 9$ that increases its weight at time 9 to $2 / 3$, and a task $T_{2}$ that has an initial weight of $8 / 9$ that decreases to $1 / 3$. (a) $T_{2}$ 's weight change is initiated at time 2. (b) Subtasks $T_{2}^{[3]}, T_{2}^{[4]}, T_{2}^{[6]}$, $T_{2}^{[7]}$ are absent and $T_{2}$ 's weight change is initiated at time 9.

Equivalent instances. Before continuing, it is worth pointing out that since halted subtasks are never scheduled, and Rules P, N, and H behave the same whether the last-released subtask is absent or not, PD-PNH produces the same schedule regardless of whether a halted subtask is absent or present. Thus, we assume that in every task instance presented in the remainder of this section, if a subtask is halted, then it is absent.

Similarly, notice that, Rule H can be emulated by delaying the initiation of a weight change until the group deadline of the last-released subtask and selectively choosing some subtasks to be absent.

Example (Figure 5.17). Consider the example in Figure 5.17, which depicts a oneprocessor system that contains two tasks: $T_{1}$, which has an initial weight of $1 / 9$ that increases to $2 / 3$ at time 9 ; and $T_{2}$, which has an initial weight of $8 / 9$ that decreases to $1 / 3$. In inset (a), $T_{2}$ initiates its weight change at time 2 . In inset (b), subtasks $T_{2}^{[3]}, T_{2}^{[4]}, T_{2}^{[6]}, T_{2}^{[7]}$ are absent, and $T_{2}$ initiates its weight change at time $\mathrm{D}\left(T_{2}^{[2]}\right)=9$. Notice that, these two systems have exactly the same schedule.


Figure 5.18: An illustration of displacements. (a) The original system. (b) The system with $T_{1}^{[1]}$ removed.

Given the equivalence demonstrated above, in every AGIS task instance considered in the remainder of this section, we assume that a heavy task only initiates a weight change at or after the group deadline of the last-released subtask. (This assumption removes the possibility that a subtask of a heavy-changeable task could receive allocations in the SW schedule after its deadline.)

Definition 5.3 (Removal and Displacements (Srinivasan, 2003)). By definition, the removal of a subtask (i.e., changing a subtask from present to absent) from one instance of an AGIS task system results in another instance. (Note that, only present subtasks can be removed.) Let $\mathrm{X}^{(j)}$ denote a subtask of any task in an AGIS task system $\tau$. Let $\mathcal{S}$ denote the PD-PNH schedule of $\tau$. Assume that removing $\mathrm{X}^{(1)}$ scheduled at slot $t_{1}$ in $\mathcal{S}$ causes $\mathrm{X}^{(2)}$ to shift from slot $t_{2}$ to $t_{1}$, where $t_{1} \neq t_{2}$, which in turn may cause other shifts. We call this shift a displacement and represent it by the four-tuple $\left\langle\mathrm{X}^{(1)}, t_{1}, \mathrm{X}^{(2)}, t_{2}\right\rangle$. A displacement $\left\langle\mathrm{X}^{(1)}, t_{1}, \mathrm{X}^{(2)}, t_{2}\right\rangle$ is valid iff $\mathrm{r}\left(\mathrm{X}^{(2)}\right) \leq t_{1}$. Because there can be a cascade of shifts, we may have a chain of displacements.

Example (Figure 5.18). Consider the two-processor example in Figure 5.18, which depicts the schedule for a system with four task $T_{1}, \ldots, T_{4}$, each of which has a weight of $3 / 7$, and $T_{5}$,
which has $w t\left(T_{5}\right)=1 / 7$. Inset (a) depicts the system in which all tasks are present. Inset (b) depicts the system where $T_{1}^{[1]}$ has been removed causing a chain of displacements.

Removing a subtask may also lead to slots in which some processors are idle. In a schedule $\mathcal{S}$, if $k$ processors are idle in slot $t$, then we say that there are $k$ holes in $\mathcal{S}$ in slot $t$. Note that, holes may exist because of late subtask releases or absent subtasks, even if total utilization is $M$. We now present three lemmas that describe the relationship among subtasks in a chain of displacements. These three lemmas were originally proven by Srinivasan and Anderson for the "generalized IS task model," i.e., an IS task system, where subtasks can be absent (Srinivasan and Anderson, 2006). Since the logic of their proof holds for AGIS task systems, we state these lemmas without proof.

Lemma 5.1 (From (Srinivasan, 2003)). Let $\mathbf{X}^{(1)}$ be a subtask that is removed from $\tau$, where all halted subtasks are absent, and let the resulting chain of displacements in $\mathcal{S}$ be $C=\Delta_{1}, \Delta_{2}, \ldots, \Delta_{k}$, where $\Delta_{j}=\left\langle\mathrm{X}^{(j)}, t_{j}, \mathrm{X}^{(j+1)}, t_{j+1}\right\rangle$. Then $t_{j+1}>t_{j}$, for all $j \in\{1, \ldots, k\}$

Lemma 5.2 (From (Srinivasan, 2003)). Let $\Delta=\left\langle\mathrm{X}^{(1)}, t_{1}, \mathrm{X}^{(2)}, t_{2}\right\rangle$ be a valid displacement in $\mathcal{S}$, in which all halted subtasks are absent. If $t_{1}<t_{2}$ and there is a hole in slot $t_{1}$ in that schedule, then $\mathrm{X}^{(2)}$ is $\mathrm{X}^{(1)}$ 's successor in $\tau$.

Lemma 5.3 (From (Srinivasan, 2003)). Let $\Delta=\left\langle\mathrm{X}^{(1)}, t_{1}, \mathrm{X}^{(2)}, t_{2}\right\rangle$ be a valid displacement in $\mathcal{S}$, in which all halted subtasks are absent. If $t_{1}<t_{2}$ and there is a hole in slot $t^{\prime}$ such that $t_{1} \leq t^{\prime}<t_{2}$ in that schedule, then $t^{\prime}=t_{1}$ and $\mathbf{X}^{(2)}$ is the successor of $\mathbf{X}^{(1)}$ in $\tau$.

### 5.5.3 Reweighting Properties

In this section, we introduce some properties that are necessary to prove scheduling correctness.

Basic properties. The following simple property follows directly from the definition of PD-PNH.
(RW) Suppose $T_{i}$ initiates a weight change at time $t_{c} \geq \mathrm{r}\left(T_{i}^{[1]}\right)$ and $T_{i}^{[j]}$ is the last-released subtask of $T_{i}$ at $t_{c}$. If $\mathrm{r}\left(T_{i}^{[j]}\right) \leq t_{c}<\mathrm{d}\left(T_{i}^{[j]}\right)$, then $T_{i}$ is either positive-, negative-, or heavy-changeable at $t_{c}$.

In Section 5.5, we presented Property (V), which relates the release times, completion times, and deadlines of two subtasks $T_{i}^{[j]}$ and $T_{i}^{[j+1]}$. In the following proof, it is useful to extend this property to relate the release times, completion times, and deadlines of two subtasks $T_{i}^{[j]}$ and $T_{i}^{[k]}$, where $j<k$.
(GV) For the subtasks $T_{i}^{[j]}$ and $T_{i}^{[k]}$, where $j<k$, if $\mathrm{r}\left(T_{i}^{[k]}\right)<\mathrm{d}\left(T_{i}^{[j]}\right)-\mathrm{b}\left(T_{i}^{[j]}\right)$, then $\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right) \leq \mathrm{r}\left(T_{i}^{[k]}\right)$ and $\mathrm{C}\left(\mathcal{S}, T_{i}^{[j]}\right) \leq \mathrm{r}\left(T_{i}^{[k]}\right)$.

The proof of Property (GV) follows directly from Property (V). Property (GV) holds even if $T_{i}^{[j]}$ or $T_{i}^{[j]}$ is absent.

Group deadlines of light tasks. Notice that, if a present subtask $T_{i}^{[j]}$ has a group-deadline of 0 , then it must be the case that $\operatorname{Swt}\left(T_{i}, \mathrm{r}\left(T_{i}^{[j]}\right)\right)<1 / 2$. Thus, unless $T_{i}$ enacts a weight change by time $\mathrm{r}\left(T_{i}^{[k]}\right)$, where $T_{i}^{[k]}$ is $T_{i}^{[j]}$,s successor, then $\mathrm{D}\left(T_{i}^{[k]}\right)=0$ must hold. Moreover, by Rules P and N , if $T_{i}^{[j]}$ is scheduled in slot $\mathrm{d}\left(T_{i}^{[j]}\right)-1, \mathrm{~b}\left(T_{i}^{[j]}\right)=1$, and $\mathrm{r}\left(T_{i}^{[k]}\right) \leq \mathrm{d}\left(T_{i}^{[j]}\right)-1$, then it is not possible for $T_{i}$ to enact a weight change before $T_{i}^{[k]}$ is released. Thus, we have the following property
(GD-3) For any subtask $T_{i}^{[j]}$, and its successor $T_{i}^{[k]}$, if $\mathrm{D}\left(T_{i}^{[j]}\right)=0, \mathrm{~b}\left(T_{i}^{[j]}\right)=1, T_{i}^{[j]}$ is scheduled in slot $\mathrm{d}\left(T_{i}^{[j]}\right)-1$, and $\mathrm{r}\left(T_{i}^{[k]}\right) \leq \mathrm{d}\left(T_{i}^{[j]}\right)-1$, then $\mathrm{D}\left(T_{i}^{[k]}\right)=0$.

Per-slot allocation properties. We now introduce five properties about the per-slot allocations of a task and the completion time of a subtask in an AGIS system $\tau$ that are useful in the correctness proof. (Recall that we assume that in $\tau$ all halted subtasks are absent and no heavy task initiates a weight change before the group deadline of its last-released subtask.)
(AF1) For all $t \geq 0, \mathrm{~A}\left(\mathcal{S} W, T_{i}, t\right) \leq \operatorname{Swt}\left(T_{i}, t\right)$.
(AF2) For any present subtask $T_{i}^{[j]}$ of the task $T_{i} \in \tau$ and its successor $T_{i}^{[k]}$, if $\mathrm{b}\left(T_{i}^{[j]}\right)=$ 1 and $\mathrm{r}\left(T_{i}^{[k]}\right) \geq \mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)$, then $\mathrm{A}\left(\mathcal{S} W, T_{i}, \mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)-1, \mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)+1\right) \leq$ $\operatorname{Swt}\left(T_{i}, C\left(\mathcal{S} W, T_{i}^{[j]}\right)\right)$.
(AF3) For any subtask $T_{i}^{[j]}, \mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right) \leq \mathrm{d}\left(T_{i}^{[j]}\right)$.
(AF4) For any subtask $T_{i}^{[j]}$ and any time $t$, if $t<\mathrm{r}\left(T_{i}^{[j]}\right)$ or $t \geq \mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)$, then $\mathrm{A}\left(\mathcal{S} W, T_{i}^{[j]}, t\right)=0$.
(AF5) For any present subtask $T_{i}^{[j]}$, let $T_{i}^{[k]}$ be its next present successor $T_{i}^{[k]}$ (if it exists). If $T_{i}^{[k]}$ does not exist, and $\mathrm{D}\left(T_{i}^{[j]}\right)>0$, then for any $t$ such that $\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)-1 \leq t \leq$ $\mathrm{D}\left(T_{i}^{[j]}\right), \mathrm{A}\left(\mathcal{S} W, T_{i}, \mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)-1, t+1\right) \leq \operatorname{Swt}\left(T_{i}, t\right)$ holds. Otherwise, if $T_{i}^{[k]}$ exists, $\mathrm{D}\left(T_{i}^{[j]}\right)>0$, and $\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right) \leq \mathrm{r}\left(T_{i}^{[k]}\right)$, then for any $t$ such that $\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)-1 \leq t \leq$ $\min \left(\mathrm{r}\left(T_{i}^{[k]}\right), \mathrm{D}\left(T_{i}^{[j]}\right)-1\right), \mathrm{A}\left(\mathcal{S} W, T_{i}, \mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)-1, t+1\right) \leq \operatorname{Swt}\left(T_{i}, t\right)$ holds.

Given the examples in Figure 5.14 and Figure 5.16, both Properties (AF1) and (AF4) should be fairly intuitive. As for Property (AF2), notice that, in Figure 5.14, $\mathrm{A}\left(\mathcal{S} W, T_{1}, \mathrm{C}\left(\mathcal{S} W, T_{1}^{[1]}\right)-1, \mathrm{C}\left(\mathcal{S} W, T_{1}^{[1]}\right)+1\right)=1 / 16+4 / 16 \leq \operatorname{Swt}\left(T_{1}, 4\right)=5 / 16$ and $\mathrm{A}\left(\mathcal{S} W, T_{1}, \mathrm{C}\left(\mathcal{S} W, T_{1}^{[4]}\right)-1, \mathrm{C}\left(\mathcal{S} W, T_{1}^{[4]}\right)+1\right)=4 / 16+0 \leq \operatorname{Swt}\left(T_{1}, 14\right)=5 / 16$. Also notice that, in Figure 5.16(b), which depicts a task $T_{i}$ that changes its weight from 3/19 to $2 / 5$ via Rule N at time $8, \mathrm{~A}\left(\mathcal{S} W, T_{i}, \mathrm{C}\left(\mathcal{S} W, T_{i}^{[2]}\right)-1, \mathrm{C}\left(\mathcal{S} W, T_{i}^{[2]}\right)+1\right)=32 / 95+0 \leq \operatorname{Swt}\left(T_{i}, 10\right)=$ $2 / 5$. Also note that, in Figure 5.16(c), which depicts a task $T_{1}$ that changes its weight from $3 / 19$ to $2 / 5$ via Rule P at time $8, \mathrm{~A}\left(\mathcal{S} W, T_{i}, \mathrm{C}\left(\mathcal{S} W, T_{i}^{[1]}\right)-1, \mathrm{C}\left(\mathcal{S} W, T_{i}^{[1]}\right)+1\right)=1 / 19+0 \leq$ $\operatorname{Swt}\left(T_{i}, 7\right)=3 / 19$.

As for Property (AF3), recall that (for systems in which no heavy task initiates a weight change before the group deadline of its last-released subtask) in the absence of reweighting events, $\mathrm{d}\left(T_{i}^{[j]}\right)=\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)$. To increase the completion time of $T_{i}^{[j]}$ (and hence $\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)$ ) in $\mathcal{S} W, T_{i}$ would have to enact a weight change in the range $\left(\mathrm{r}\left(T_{i}^{[j]}\right), \mathrm{d}\left(T_{i}^{[j]}\right)\right)$ that decreases the weight of $T_{i}$, without halting $T_{i}^{[j]}$. However, by Property (X2), a change enacted in the range $\left(\mathrm{r}\left(T_{i}^{[j]}\right), \mathrm{d}\left(T_{i}^{[j]}\right)\right)$ must have been initiated in the range $\left[\mathrm{r}\left(T_{i}^{[j]}\right), \mathrm{d}\left(T_{i}^{[j]}\right)\right)$. Thus, by Property (RW) when such a change is initiated, $T_{i}$ is either positive, negative-, or heavy-changeable. Since only Rule N can enact a weight change before $\mathrm{d}\left(T_{i}^{[j]}\right)$ without halting the last-released subtask (i.e., $T_{i}^{[j]}$ ), when such a change is initiated, $T_{i}$ must be negative-changeable. However, by Rule N , no weight decrease that is initiated in the range $\left[\mathrm{r}\left(T_{i}^{[j]}\right), \mathrm{d}\left(T_{i}^{[j]}\right)\right)$ can be enacted before time $\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)$. Thus, the completion time for a subtask in $\mathcal{S} W$ is upper bounded by the deadline of the subtask.

As for Property (AF5), notice that, in Figure 5.15 (a), as successive subtasks of a task $T_{i}$ are released, the allocations to each of $T_{i}$ 's subtasks in the SW schedule in the first slot of each subtask's window decreases up until the group deadline. For example, $\mathrm{A}\left(\mathcal{S} W, T_{1}^{[1]}, r\left(T_{1}^{[1]}\right)\right)=7 / 9$, $\mathrm{A}\left(\mathcal{S} W, T_{1}^{[2]}, \mathrm{r}\left(T_{1}^{[2]}\right)\right)=5 / 9, \mathrm{~A}\left(\mathcal{S} W, T_{1}^{[3]}, \mathrm{r}\left(T_{1}^{[3]}\right)\right)=3 / 9$, and $\mathrm{A}\left(\mathcal{S} W, T_{1}^{[4]}, \mathrm{r}\left(T_{1}^{[4]}\right)\right)=1 / 9$. As a result, if $T_{i}$ has no present subtasks between $T_{i}^{[j]}$ and $T_{i}^{[k]}$ and $r\left(T_{i}^{[k]}\right) \leq \omega\left(T_{i}^{[j]}\right)$, then the allocation over the range $\left[\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)-1, \mathrm{r}\left(T_{i}^{[k]}\right)\right)$ is at most the scheduling weight of $T_{i}$. For example, in Figure $5.15(\mathrm{~b}), \mathrm{A}\left(\mathcal{S} W, T_{1}, \mathrm{C}\left(\mathcal{S} W, T_{1}^{[1]}\right)-1, \mathrm{r}\left(T_{1}^{[4]}\right)\right)=2 / 9+1 / 9=$ $3 / 9 \leq \operatorname{Swt}\left(T_{1}, 4\right)=7 / 9$ and $\mathrm{A}\left(\mathcal{S W}, T_{1}, \mathrm{C}\left(\mathcal{S} W, T_{1}^{[4]}\right)-1, \mathrm{r}\left(T_{1}^{[6]}\right)+1\right)=1 / 9+4 / 9 \leq 5 / 9=$ $\operatorname{Swt}\left(T_{1}, 14\right)$.

It is worthwhile to note that the proofs of Properties (AF1), (AF3), and (AF5) follow directly from corresponding IS properties that were proven in (Srinivasan, 2003).

As a consequence of Properties (AF1) and (W), LAG can only increase over a slot if there is a hole in that slot. Hence, the lemma below follows.

Lemma 5.4. If $\operatorname{LAG}(\tau, t)<\operatorname{LAG}(\tau, t+1)$, then there is a hole in slot $t$.

### 5.5.4 Correctness Proof

Having defined the AGIS task model, displacements, and some basic properties, we can now prove the following theorem.

Theorem 5.2. Under PD-PNH, no subtask is scheduled after its deadline, provided that Property (W) holds.

To prove of Theorem 5.2 , suppose that it does not hold. Then, there exists a time $t_{d}$ and a task system $\tau$ as given in the definitions below.

Definition $5.4\left(t_{d}\right) . t_{d}$ is the earliest time at which any AGIS task system instance misses a deadline under PD-PNH.

Definition $5.5(\tau$ and $\mathcal{S}) . \tau$ is an instance of an AGIS task system with the following properties.
(T1) $\tau$ misses a deadline under PD-PNH at $t_{d}$.
(T2) No task system satisfying (T1) has fewer present subtasks in $\left[0, t_{d}\right)$ than $\tau$.
In the remainder of this proof, we let $\mathcal{S}$ denote PD-PNH schedule of $\tau$.

By (T1), (T2), and the definition of $t_{d}$, exactly one subtask in $\tau$ misses its deadline at $t_{d}$ : if several subtasks miss their deadlines, then all but one can be removed and the remaining subtask will still miss its deadline, contradicting (T2). We now prove several properties about $\mathcal{S}$.

Lemma 5.5. The following properties hold for $\tau$ and $\mathcal{S}$, where $T_{i}^{[j]}$ is any subtask in $\mathcal{S}$.
(a) For any present subtask $T_{i}^{[j]}, \mathrm{d}\left(T_{i}^{[j]}\right) \leq t_{d}$.
(b) There are no holes in slot $t_{d}-1$.
(c) $\operatorname{LAG}\left(\tau, t_{d}\right)=1$.
(d) $\operatorname{LAG}\left(\tau, t_{d}-1\right) \geq 1$.
(e) No present subtask halts.

Proof of (a). Suppose that $\tau$ contains a subtask $T_{i}^{[j]}$ with a deadline greater than $t_{d} . T_{i}^{[j]}$ can be removed without affecting the scheduling of higher-priority subtasks with earlier deadlines. Thus, if $T_{i}^{[j]}$ is removed, then a deadline still missed at $t_{d}$. This contradicts (T2).

Proof of (b). If there were a hole in slot $t_{d}-1$, then the subtask that misses its deadline at $t_{d}$ would have been scheduled there, which is a contradiction. (Note that, by the minimality of $t_{d}$, its predecessor meets its deadline at or before $t_{d}-1$ and hence is not scheduled in slot $t_{d}-1$.)

Proof of (c). By (5.1), we have

$$
\operatorname{LAG}\left(\tau, t_{d}\right)=\mathrm{A}\left(\mathcal{S} W, \tau, 0, t_{d}\right)-\mathrm{A}\left(\mathcal{S}, \tau, 0, t_{d}\right) .
$$

In the above equation, the term $\mathrm{A}\left(\mathcal{S} W, \tau, 0, t_{d}\right)$ equals the total number of present subtasks in $\tau$. The second term corresponds to the number of subtasks scheduled by PD-PNH in $\left[0, t_{d}\right)$.


Figure 5.19: The task sets $A, B$, and $I$.

Since exactly one subtask misses its deadline, the difference between these two terms is one, i.e., $\operatorname{LAG}\left(\tau, t_{d}\right)=1$.

Proof of (d). By (b), there are no holes in slot $t_{d}-1$. Hence, by Lemma 5.4, $\operatorname{LAG}\left(\tau, t_{d}-1\right) \geq$ $\operatorname{LAG}\left(\tau, t_{d}\right)$. Therefore, by $(\mathrm{c}), \operatorname{LAG}\left(\tau, t_{d}-1\right) \geq 1$.

Proof of (e). If $T_{i}^{[j]}$ is a halted subtask, then it is never scheduled. Hence, it can be removed and a deadline will still be missed at $t_{d}$, contradicting (T2).

Definition $5.6\left(\mathrm{~A}_{t}, \mathrm{~B}_{t}\right.$, and $\left.\mathrm{I}_{t}\right) . \mathrm{A}_{t}, \mathrm{~B}_{t}$, and $\mathrm{I}_{t}$ are all defined with respect to schedule $\mathcal{S}$ and some time $t$. $\mathrm{A}_{t}$ denotes the set of tasks that have a subtask scheduled at $t$. $\mathrm{B}_{t}$ denotes the set of tasks that are not scheduled at $t$, and receive some allocation in $\mathcal{S} W$ at slot $t$, i.e., $\mathrm{A}\left(\mathcal{S} W, T_{i}, t\right)>0$ for $T_{i} \in \mathrm{~B}_{t} . \mathrm{I}_{t}$ denotes the set of all tasks that are in the system at $t$ but are not in $\mathrm{A}_{t}$ or $\mathrm{B}_{t} . \mathrm{A}_{t}, \mathrm{~B}_{t}$, and $\mathrm{I}_{t}$ are illustrated in Figure 5.19. Notice that, there are two types of tasks in set $\mathbf{I}_{t}$ : tasks that have a deadline at or before $t$ (denoted $\mathbf{I}_{t}^{1}$ ) and tasks that have a deadline after $t$ but complete in $\mathcal{S} W$ at or before $t$ (denoted $\mathbf{I}_{t}^{2}$ ).

Displacement-based proofs. We now prove several lemmas about tasks in $A_{t}$ and $B_{t}$ when there is a hole at time $t$. The proofs of the following lemmas all use the same basic technique. First, we assume, to derive a contradiction, that the lemma does not hold for some subtask, $T_{u}^{[j]}$. Then, we show that $T_{u}^{[j]}$ can be removed from $\tau$ without causing a chain of displacements that extends beyond time $t$. Thus, the system without $T_{u}^{[j]}$ misses a deadline at time $t_{d}$. This, in turn, implies that $\tau$ violates property (T2), which is a contradiction.


Figure 5.20: An illustration of the chain of displacements that occurs by removing $\mathrm{X}^{(1)}=T_{u}^{[j]}$ in Lemma 5.6.

Lemma 5.6. Let $t<t_{d}-1$ be a time at which there is a hole in $\mathcal{S}$. Let $T_{u}$ be any task in $\mathrm{B}_{t}$ or $\mathrm{A}_{t}$. Let $T_{u}^{[j]}$ be the subtask with the largest index such that $\mathrm{r}\left(T_{u}^{[j]}\right) \leq t<\mathrm{d}\left(T_{u}^{[j]}\right)$ and $T_{u}^{[j]}$ is scheduled at or before $t$. Then, $\mathrm{d}\left(T_{u}^{[j]}\right)=t+1 \wedge \mathrm{~b}\left(T_{u}^{[j]}\right)=1$.

Proof. Let $t$ and $T_{u}$ be as defined in the statement of the lemma. If $T_{u} \in \mathrm{~A}_{t}$, then since $t<t_{d}-1$ and $T_{u}^{[j]}$ is scheduled at or before $t, \mathrm{~d}\left(T_{u}^{[j]}\right) \geq t+1$. If $T_{u} \in \mathrm{~B}_{t}$, then since $\mathrm{A}\left(\mathcal{S} W, T_{u}, t\right)>0$ it follows that $T_{u}$ is not complete in $\mathcal{S} W$ before $t+1$. Thus, by Property (AF3), $\mathrm{d}\left(T_{u}^{[j]}\right) \geq \mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right) \geq t+1$. Thus, if we can show that a contradiction follows from

$$
\begin{equation*}
\mathrm{d}\left(T_{u}^{[j]}\right)>t+1 \text { or } \mathrm{d}\left(T_{u}^{[j]}\right)=t+1 \wedge \mathrm{~b}\left(T_{u}^{[j]}\right)=0, \tag{5.15}
\end{equation*}
$$

then the proof is complete.
We now show that if (5.15) holds, then $T_{u}^{[j]}$ can be removed and a deadline will still be missed at $t_{d}$, contradicting Property (T2). (Before continuing, notice that, since $T_{u}^{[j]}$ is scheduled, it is present, and as a result, $T_{u}^{[j]}$ can be removed.) Let the chain of displacements caused by removing $T_{u}^{[j]}$ be $\Delta_{1}, \Delta_{2}, \ldots, \Delta_{k}$, where $\Delta_{i}=\left\langle\mathrm{X}^{(i)}, t_{i}, \mathrm{X}^{(i+1)}, t_{i+1}\right\rangle$ and $\mathrm{X}^{(1)}=T_{u}^{[j]}$. By Lemma 5.1, $t_{i+1}>t_{i}$, for $1 \leq i \leq k$. (This chain of displacements is illustrated in Figure 5.20.)

Note that, at slot $t_{i}$, the priority of $\mathbf{X}^{(i)}$ is at least that of $\mathbf{X}^{(i+1)}$, because $\mathrm{X}^{(i)}$ was chosen over $\mathbf{X}^{(i+1)}$ in $\mathcal{S}$. Thus, because $\mathbf{X}^{(1)}=T_{u}^{[j]}$, by (5.15), for each subtask $1 \leq i \leq k+1$, either $\mathrm{d}\left(\mathrm{X}^{(i)}\right)>t+1$ or $\mathrm{d}\left(\mathrm{X}^{(i)}\right)=t+1 \wedge \mathrm{~b}\left(\mathrm{X}^{(i)}\right)=0$. We now show that the displacements do not
extend beyond slot $t$. Assume to the contrary that $t_{k+1}>t$. Consider $h \in\{2, \ldots, k+1\}$ such that $t_{h}>t$ and $t_{h-1} \leq t$. Such an $h$ exists because $t_{1} \leq t<t_{k+1}$. Because there is a hole in slot $t$ and $t_{h-1} \leq t<t_{h}$, by Lemma 5.3, $t_{h-1}=t$ and $\mathbf{X}^{(h)}$ is $\mathbf{X}^{(h-1)}$ 's successor. Since a subtask cannot be scheduled before it is released, $\mathrm{r}\left(\mathrm{X}^{(h)}\right)<t+1$. Since $h-1 \leq k$, either $\mathrm{d}\left(\mathrm{X}^{(h-1)}\right)>t+1$ or $\mathrm{d}\left(\mathrm{X}^{(h-1)}\right)=t+1 \wedge \mathrm{~b}\left(\mathrm{X}^{(h-1)}\right)=0$ holds. Therefore, since $\mathrm{r}\left(\mathrm{X}^{(h)}\right)<t+1$, $\mathrm{d}\left(\mathrm{X}^{(h-1)}\right)-\mathrm{b}\left(\mathrm{X}^{(h-1)}\right)>\mathrm{r}\left(\mathrm{X}^{(h)}\right)$ holds. Thus, by Property (GV), $\mathrm{X}^{(h-1)}$ is scheduled before $\mathrm{r}\left(\mathrm{X}^{(h)}\right) \leq t$ in $\mathcal{S}$, contradicting our assumption that $\mathrm{X}^{(h-1)}$ is scheduled in slot $t$.

Thus, the displacements do not extend beyond slot $t$. Hence, no subtask scheduled after $t$ is "left-shifted." Thus, a deadline is still missed at time $t_{c}$, contradicting Property (T2). Hence, $\mathrm{d}\left(T_{u}^{[j]}\right)=t+1 \wedge \mathrm{~b}\left(T_{u}^{[j]}\right)=1$.

Lemma 5.7. Let $t<t_{d}-1$ be a time at which there is a hole in $\mathcal{S}$. Let $T_{u}$ be any task in $\mathrm{B}_{t}$. Let $T_{u}^{[j]}$ be the subtask with the largest index such that $\mathrm{r}\left(T_{u}^{[j]}\right) \leq t<\mathrm{d}\left(T_{u}^{[j]}\right)$ and $T_{u}^{[j]}$ is scheduled before $t$. Then, there exists a subtask $T_{a}^{[b]}$ scheduled in slot $t$ such that $\mathrm{D}\left(T_{a}^{[b]}\right) \leq \mathrm{D}\left(T_{u}^{[j]}\right)$.

Proof. Let $t$ and $T_{u}$ be as defined in the statement of the lemma. Suppose that all subtasks scheduled in slot $t$ have a group deadline greater than $\mathrm{D}\left(T_{u}^{[j]}\right)$. We now show that $T_{u}^{[j]}$ can be removed and a deadline will still be missed at $t_{d}$, contradicting Property (T2). (Before continuing, notice that, since $T_{u}^{[j]}$ is scheduled, it is present, and as a result, $T_{u}^{[j]}$ can be removed.) Let the chain of displacements caused by removing $T_{u}^{[j]}$ be $\Delta_{1}, \Delta_{2}, \ldots, \Delta_{k}$, where $\Delta_{i}=\left\langle\mathbf{X}^{(i)}, t_{i}, \mathrm{X}^{(i+1)}, t_{i+1}\right\rangle, \mathrm{X}^{(1)}=T_{u}^{[j]}$. By Lemma 5.1, $t_{i+1}>t_{i}$, for $1 \leq i \leq k$.

Note that, at slot $t_{i}$, the priority of $\mathrm{X}^{(i)}$ is at least that of $\mathrm{X}^{(i+1)}$, because $\mathrm{X}^{(i)}$ was chosen over $\mathrm{X}^{(i+1)}$ in $\mathcal{S}$. Thus, because $\mathrm{X}^{(1)}=T_{u}^{[j]}$ and, by Lemma 5.6, $\mathrm{d}\left(T_{u}^{[j]}\right)=t+1 \wedge \mathrm{~b}\left(T_{u}^{[j]}\right)=1$, we have the following property.
(Q) For any value of $i$ such that $1 \leq i \leq k+1$, if $\mathrm{d}\left(\mathbf{X}^{(i)}\right)=t+1$ and $\mathrm{b}\left(\mathrm{X}^{(i)}\right)=1$, then $\mathrm{D}\left(\mathrm{X}^{(i)}\right) \leq \mathrm{D}\left(T_{u}^{[j]}\right)$.

We now show that the displacements do not extend beyond slot $t$. Assume to the contrary that $t_{k+1}>t$. Consider $h \in\{2, \ldots, k+1\}$ such that $t_{h}>t$ and $t_{h-1} \leq t$. Such an $h$ exists because $t_{1} \leq t<t_{k+1}$. Because there is a hole in slot $t$ and $t_{h-1} \leq t<t_{h}$, by Lemma 5.3, $t_{h-1}=t$. Thus, $\mathbf{X}^{(h-1)}$ is scheduled in slot $t$ and, by extension, $\mathbf{X}^{(h-1)}$ has the largest
index of any subtask of its task that is scheduled at or before $t$. Thus, by Lemma 5.6, $\mathrm{d}\left(\mathrm{X}^{(h-1)}\right)=t+1 \wedge \mathrm{~b}\left(\mathrm{X}^{(h-1)}\right)=1$. Thus, by Property (Q), $\mathrm{D}\left(\mathrm{X}^{(h-1)}\right) \leq \mathrm{D}\left(T_{u}^{[j]}\right)$, which contradicts our assumption all subtasks scheduled in slot $t$ have a group deadline greater than $\mathrm{D}\left(T_{u}^{[j]}\right)$.

Thus, the displacements do not extend beyond slot $t$. Hence, no subtask scheduled after $t$ is "left-shifted." Thus, a deadline is still missed at time $t_{d}$, contradicting Property (T2). Hence, there exists a subtask $T_{a}^{[b]}$ scheduled in slot $t$ such that $\mathrm{D}\left(T_{a}^{[b]}\right) \leq \mathrm{D}\left(T_{u}^{[j]}\right)$.

Lemma 5.8. Let $t<t_{d}-1$ be a slot in which there is a hole in $\mathcal{S}$. Let $T_{u}$ be any task in $\mathrm{A}_{t}$. Let $T_{u}^{[j]}$ be the subtask of $T_{u}$ that is scheduled in slot $t$. Then, $T_{u}^{[j]}$,s successor is released at time $t$.

Proof. Let $t$ and $T_{u}$ be as defined in the statement of the lemma. By Lemma 5.6, $\mathrm{d}\left(T_{u}^{[j]}\right)=t+1$ and $\mathrm{b}\left(T_{u}^{[j]}\right)=1$. Since $T_{u}^{[j]}$ is scheduled at time $t, T_{u}^{[j]}$ is not complete by time $t$ in $\mathcal{S}$. Thus, by Property (GV), $T_{u}^{[j]}$, s successor cannot be released before $\mathrm{d}\left(T_{u}^{[j]}\right)-\mathrm{b}\left(T_{u}^{[j]}\right)=t$.

We now show that $T_{u}^{[j]}$,s successor must be released by time $t$, which completes the proof. Suppose, to derive a contradiction, that $T_{u}^{[j]}$,s successor is released after time $t$. We now show that $T_{u}^{[j]}$ can be removed and a deadline will still be missed at $t_{d}$, contradicting Property (T2). (Before continuing, notice that, since $T_{u}^{[j]}$ is scheduled, it is present, and as a result, $T_{u}^{[j]}$ can be removed.) Let the chain of displacements caused by removing $T_{u}^{[j]}$ be $\Delta_{1}, \Delta_{2}, \ldots, \Delta_{k}$, where $\Delta_{i}=\left\langle\mathrm{X}^{(i)}, t_{i}, \mathrm{X}^{(i+1)}, t_{i+1}\right\rangle, \mathrm{X}^{(1)}=T_{u}^{[j]}$. By Lemma 5.1, $t_{i+1}>t_{i}$, for $1 \leq i \leq k$.

We now show that the displacements do not extend beyond slot $t$. Assume to the contrary that $t_{k+1}>t$. Consider $h \in\{2, \ldots, k+1\}$ such that $t_{h}>t$ and $t_{h-1} \leq t$. Such an $h$ exists because $t_{1} \leq t<t_{k+1}$. Because there is a hole in slot $t$ and $t_{h-1} \leq t<t_{h}$, by Lemma 5.3, $t_{h-1}=t, \mathbf{X}^{(h)}$ is $\mathrm{X}^{(h-1)}$ 's successor. Moreover, by Lemma 5.1, $\mathrm{X}^{(h-1)}=T_{u}^{[j]}$. Notice that the displacement $\left\langle\mathrm{X}^{(h-1)}, t_{h-1}, \mathrm{X}^{(h)}, t_{h}\right\rangle$ is valid iff $\mathrm{r}\left(\mathrm{X}^{(h)}\right) \leq t$, which implies that $T_{u}^{[j]}$,s successor is released by time $t$. However, this contradicts our assumption that $T_{u}^{[j]}$,s successor is released after time $t$.

Thus, the displacements do not extend beyond slot $t$. Hence, no subtask scheduled after $t$ is "left-shifted." Thus, a deadline is still missed at time $t_{d}$, contradicting Property (T2). Hence, $T_{u}^{[j]}$,s successor is released at time $t$.

Consecutive holes. Notice that, if there are two consecutive slots with holes, $t_{a}$ and $t_{b}$, then by Lemma 5.8, for any subtask $T_{u}^{[j]}$ scheduled in slot $t_{a}, T_{u}^{[j]}$,s successor, $T_{u}^{[q]}$, is released at time $t_{a}$. As a result, since $\mathrm{r}\left(T_{u}^{[q]}\right)=t_{a}<t_{b}$, there is a hole in slot $t_{b}$, and $T_{u}^{[j]}$ is scheduled in slot $t_{a}$, it follows that $T_{u}^{[q]}$ is scheduled in slot $t_{b}$. Thus, we have the following lemma.

Lemma 5.9. Let $t_{a}$ and $t_{b}$ be two consecutive slots such that $t_{a}=t_{b}-1<t_{d}-1$ and both $t_{a}$ and $t_{b}$ contain holes in the schedule $\mathcal{S}$. Let $T_{u}$ be any task in $\mathrm{A}_{t_{a}}$. Let $T_{u}^{[j]}$ be the subtask of $T_{u}$ that is scheduled in slot $t_{a}$. Then, $T_{u}^{[j]}$, s successor is scheduled in slot $t_{b}$.

Lemma 5.10. Let $t_{a}$ and $t_{b}$ be two consecutive slots such that $t_{a}=t_{b}-1<t_{d}-1$ and both $t_{a}$ and $t_{b}$ contain holes in the schedule $\mathcal{S}$. Let $T_{u}$ be any task in $\mathrm{A}_{t_{b}}$. Let $T_{u}^{[j]}$ be the subtask of $T_{u}$ that is scheduled in slot $t_{b}$. Then, $T_{u}^{[j]}$,s predecessor is scheduled in slot $t_{a}$.

Proof. Let $t_{a}, t_{b}$, and $T_{u}$ be as defined in the statement of the lemma. By Lemma 5.6, $\mathrm{d}\left(T_{u}^{[j]}\right)=t_{b}+1$ and $\mathbf{b}\left(T_{u}^{[j]}\right)=1$. Thus, by Property (WL), $\mathbf{r}\left(T_{u}^{[j]}\right) \leq \mathrm{d}\left(T_{u}^{[j]}\right)-2=t_{b}-1=t_{a}$. Since there is a hole at $t_{a}$ and $\mathrm{r}\left(T_{u}^{[j]}\right) \leq t_{a}, T_{u}^{[j]}$ would be scheduled in slot $t_{a}$, unless its predecessor was scheduled there. Since $T_{u}^{[j]}$ is scheduled in slot $t_{b}$, its predecessor must be scheduled in slot $t_{a}$.

From Lemmas 5.9 and 5.10, the subsequent corollary follows.
Corollary 5.1. Let $t_{a}$ and $t_{b}$ be two slots such $t_{a} \leq t_{b} \leq t_{d}-1$ and there is a hole in every slot in the range $\left\{t_{a}, \ldots, t_{b}\right\}$. If a task $T_{x}$ is not scheduled in slot $t_{a}$, then it is not scheduled in any slot over the range $\left\{t_{a}, \ldots, t_{b}\right\}$.

Notice that, by Lemma 5.9, any task $T_{u}$ that is scheduled in the first hole in a sequence of slots with holes must be scheduled in every slot in the sequence. Moreover, by Lemma 5.8, a subtask of $T_{u}$ must be released at every slot in this sequence. Thus, there can be no IS separations between subtasks until the sequence of holes is over. In addition, every subtask of $T_{u}$ that is released over the range $\left\{t_{a}, t_{b}-2\right\}$ will be scheduled in the second slot of its window. Thus, we have the following corollaries (depicted in Figure 5.21).

Corollary 5.2. Let $t_{a}$ and $t_{b}$ be two slots such that $t_{a} \leq t_{b} \leq t_{d}-1$ and there is a hole in every slot in the range $\left\{t_{a}, \ldots, t_{b}\right\}$. If $T_{u}^{[j]}$ is scheduled at time $t_{a}$, then for $q \in\left\{j, \ldots, j+1+t_{b}-t_{a}\right\}$, $\theta\left(T_{u}^{[q]}\right)=\theta\left(T_{u}^{[j]}\right)$, i.e., there are no IS separations between any two subtasks of $T_{u}$.


Figure 5.21: An illustration of Corollaries 5.2 and 5.3.

Corollary 5.3. Let $t_{a}$ and $t_{b}$ be two slots such $t_{a}<t_{b} \leq t_{d}-1$ and there is a hole in every slot in the range $\left\{t_{a}, \ldots, t_{b}\right\}$. If $T_{u}^{[j]}$ is scheduled at time $t_{a}$, then for any subtask $T_{u}^{[q]}$ such that $t_{a} \leq \mathrm{r}\left(T_{u}^{[q]}\right) \leq t_{b}-1, T_{u}^{[q]}$ is scheduled in slot $\mathrm{r}\left(T_{u}^{[q]}\right)+1$.

Time $t_{\mathcal{H}}$. Since, by part (d) of Lemma 5.5, $\operatorname{LAG}\left(\tau, t_{d}-1\right) \geq 1$, and by the definition of $\operatorname{LAG}, \operatorname{LAG}(\tau, 0)=0$, there exists a time $t_{\mathcal{H}}$ such that

$$
\begin{equation*}
0 \leq t_{\mathcal{H}}<t_{d}-1 \wedge \operatorname{LAG}\left(\tau, t_{\mathcal{H}}\right)<1 \wedge \operatorname{LAG}\left(\tau, t_{\mathcal{H}}+1\right) \geq 1 . \tag{5.16}
\end{equation*}
$$

Without loss of generality, let $t_{\mathcal{H}}$, be the latest such time, i.e., for all $u$ such that $t_{\mathcal{H}}<u \leq$ $t_{d}-1, \operatorname{LAG}(\tau, u) \geq 1$. We now show that such a $t_{\mathcal{H}}$ cannot exist, thus contradicting our starting assumption that $t_{d}$ and $\tau$ exist. For brevity, we use $A$ to denote $\mathrm{A}_{t_{\mathcal{H}}}, B$ to denote $\mathrm{B}_{t_{\mathcal{H}}}$, and $I$ to denote $\mathrm{I}_{t_{\mathcal{H}}}$.

Lemma 5.11. $B$ is non-empty

Proof. Let the number of holes in slot $t_{\mathcal{H}}$ be $h$. Then, $\mathrm{A}\left(\mathcal{S}, \tau, t_{\mathcal{H}}\right)=M-h$. By (5.1), $\operatorname{LAG}\left(\tau, t_{\mathcal{H}}+1\right)=\operatorname{LAG}\left(\tau, t_{\mathcal{H}}\right)+\mathrm{A}\left(\mathcal{S} W, \tau, t_{\mathcal{H}}\right)-\mathrm{A}\left(\mathcal{S}, \tau, t_{\mathcal{H}}\right)$ (recall that for this section we have assumed that $\operatorname{LAG}(\tau, t)=\operatorname{LAG}(\mathcal{S}, \mathcal{S} W, \tau, t))$. Thus, because $\operatorname{LAG}\left(\tau, t_{\mathcal{H}}+1\right)>$ $\operatorname{LAG}\left(\tau, t_{\mathcal{H}}\right)$ (by (5.16)), we have $\mathrm{A}\left(\mathcal{S} W, \tau, t_{\mathcal{H}}\right)>M-h$. Since, for every $T_{v} \notin A \cup B$, $\mathrm{A}\left(\mathcal{S} W, T_{v}, t_{\mathcal{H}}\right)=0$, it follows that $\mathrm{A}\left(\mathcal{S} W, A \cup B, t_{\mathcal{H}}\right)>M-h$. Therefore, by Property (AF1), $\sum_{T_{i} \in A}\left(\operatorname{Swt}\left(T_{i}, t_{\mathcal{H}}\right)\right)+\sum_{T_{i} \in B}\left(\mathrm{~A}\left(\mathcal{S} W, T_{i}, t_{\mathcal{H}}\right)\right)>M-h$. Because the number of tasks
scheduled in slot $t_{\mathcal{H}}$ is $M-h,|A|=M-h$. Because $\operatorname{Swt}\left(T_{i}, t\right) \leq 1$, for any task $T_{i}$ at any time $t, \sum_{T_{i} \in A} \operatorname{Swt}\left(T_{i}, t_{\mathcal{H}}\right) \leq M-h$. Thus, $\sum_{T_{i} \in B} \mathrm{~A}\left(\mathcal{S} W, T_{i}, t_{\mathcal{H}}\right)>0$. Hence, $B$ is not empty.
(TK) Let $T_{u}$ be any task in $B$ and let $T_{u}^{[j]}$ be the subtask of $T_{u}$ with the largest index such that $\mathrm{r}\left(T_{u}^{[j]}\right) \leq t_{\mathcal{H}}<\mathrm{d}\left(T_{u}^{[j]}\right)$ and $T_{u}^{[j]}$ is scheduled before $t_{\mathcal{H}}$. Then, no present subtask of $T_{u}$ with an index greater than $j$ (including $T_{u}^{[j]}$,s successor) is released before $\mathrm{d}\left(T_{u}^{[j]}\right)$.

Notice that Lemma 5.6 implies $\mathrm{d}\left(T_{u}^{[j]}\right)=t_{\mathcal{H}}+1$. Thus, Property (TK) easily follows from the fact that there is a hole in slot $t_{\mathcal{H}}$ and no subtask of any task in $B$ is scheduled in slot $t_{\mathcal{H}}$.

Lemma 5.12. Let $T_{u}$ be any task in $B$. Let $T_{u}^{[j]}$ be the subtask of $T_{u}$ with the largest index such that $\mathrm{r}\left(T_{u}^{[j]}\right) \leq t_{\mathcal{H}}<\mathrm{d}\left(T_{u}^{[j]}\right)$ and $T_{u}^{[j]}$ is scheduled before $t_{\mathcal{H}}$. Then, $\mathrm{C}\left(\mathcal{S} W, T_{u}^{[j]}\right)=t_{\mathcal{H}}+1$.

Proof. Let $T_{u}$ and $T_{u}^{[j]}$ be defined as in the statement of the lemma. Since $T_{u}^{[j]}$ is the subtask of $T_{u}$ with the largest index such that $\mathrm{r}\left(T_{u}^{[j]}\right) \leq t_{\mathcal{H}}<\mathrm{d}\left(T_{u}^{[j]}\right)$ and $T_{u}^{[j]}$ is scheduled before $t_{\mathcal{H}}$ (by the definition of $B$, no subtask of $T_{u} \in B$ is scheduled in slot $t_{\mathcal{H}}$ ), by Lemma 5.6, $\mathrm{d}\left(T_{u}^{[j]}\right)=t_{\mathcal{H}}+1$. Thus, by Property (AF3), $\mathrm{C}\left(\mathcal{S} W, T_{u}^{[j]}\right) \leq t_{\mathcal{H}}+1$.

We now show that for $\ell \neq j, \mathrm{~A}\left(\mathcal{S} W, T_{u}^{[\ell]}, t_{\mathcal{H}}\right)=0$. First, we consider $\ell>j$. From the definition of $T_{u}^{[j]}$, at least one of the following three conditions must hold: (i) $\mathrm{r}\left(T_{u}^{[\ell]}\right)>t_{\mathcal{H}}$; (ii) $\mathrm{d}\left(T_{u}^{[\ell]}\right) \leq t_{\mathcal{H}}$; or (iii) $T_{u}^{[\ell]}$ is not scheduled by time $t_{\mathcal{H}}$. If Conditions (i) or (ii) hold, then by Properties (AF3) and (AF4), $\mathrm{A}\left(\mathcal{S} W, T_{u}^{[\ell]}, t_{\mathcal{H}}\right)=0$. If Condition (iii) holds, then by Property (TK), either $\mathrm{r}\left(T_{u}^{[\ell]}\right)>\mathrm{d}\left(T_{u}^{[j]}\right)$, in which case Condition (i) holds, or $T_{u}^{[\ell]}$ is not present, in which case $\mathrm{A}\left(\mathcal{S} W, T_{u}^{[\ell]}, t_{\mathcal{H}}\right)=0$. Thus, for $\ell>j, \mathrm{~A}\left(\mathcal{S} W, T_{u}^{[\ell]}, t_{\mathcal{H}}\right)=0$. Now, consider $\ell<j$. By Lemma 5.6, $\mathrm{d}\left(T_{u}^{[j]}\right)=t_{\mathcal{H}}+1 \wedge \mathrm{~b}\left(T_{u}^{[j]}\right)=1$. By Property (WL), every subtask has a window length of at least two. Hence, $r\left(T_{u}^{[j]}\right) \leq t_{\mathcal{H}}-1$. By Property (GV), if $\mathrm{d}\left(T_{u}^{[\ell]}\right)-\mathrm{b}\left(T_{u}^{[\ell]}\right)>\mathrm{r}\left(T_{u}^{[j]}\right)$ holds, then $\mathrm{C}\left(\mathcal{S} W, T_{u}^{[\ell]}\right) \leq \mathrm{r}\left(T_{u}^{[j]}\right)$ holds. Thus, either $\mathrm{C}\left(\mathcal{S} W, T_{u}^{[\ell]}\right) \leq \mathrm{r}\left(T_{u}^{[j]}\right)$ or $\mathrm{d}\left(T_{u}^{[\ell]}\right)-\mathrm{b}\left(T_{u}^{[\ell]}\right) \leq \mathrm{r}\left(T_{u}^{[j]}\right)$ holds. Since $\mathrm{r}\left(T_{u}^{[j]}\right) \leq t_{\mathcal{H}}-1$, this implies that either $\mathrm{C}\left(\mathcal{S} W, T_{u}^{[\ell]}\right) \leq \mathrm{r}\left(T_{u}^{[j]}\right) \leq t_{\mathcal{H}}-1$ or $\mathrm{d}\left(T_{u}^{[\ell]}\right) \leq \mathrm{r}\left(T_{u}^{[j]}\right)+\mathrm{b}\left(T_{u}^{[\ell]}\right) \leq t_{\mathcal{H}}-1+\mathrm{b}\left(T_{u}^{[\ell]}\right) \leq t_{\mathcal{H}}$ holds. Since, by Property (AF3), $\mathrm{C}\left(\mathcal{S} W, T_{u}^{[\ell]}\right) \leq \mathrm{d}\left(T_{u}^{[\ell]}\right), \mathrm{C}\left(\mathcal{S} W, T_{u}^{[\ell]}\right) \leq t_{\mathcal{H}}$ holds in either case. Thus, by Property (AF4), A $\left(\mathcal{S} W, T_{u}^{[\ell]}, t_{\mathcal{H}}\right)=0$.

Thus, the allocation to each subtask of $T_{u}$, except $T_{u}^{[j]}$, in $\mathcal{S} W$ in slot $t_{\mathcal{H}}$ is zero. By the definition of $B, \mathrm{~A}\left(\mathcal{S} W, T_{u}, t_{\mathcal{H}}\right)>0$. Thus, since $\mathrm{A}\left(\mathcal{S} W, T_{u}, t_{\mathcal{H}}\right)=\sum_{T_{u}^{[i]} \in T_{u}} \mathrm{~A}\left(\mathcal{S} W, T_{u}^{[i]}, t_{\mathcal{H}}\right)$,


Figure 5.22: An illustration of the proof of Lemma 5.13, where there is a hole in time slots $t_{\mathcal{H}}$ and $t_{\mathcal{H}}+1$.
there must exist at least one present subtask with a positive allocation in $\mathcal{S} W$ in the slot $t_{\mathcal{H}}$. Thus, $\mathrm{A}\left(\mathcal{S} W, T_{u}^{[j]}, t_{\mathcal{H}}\right)>0$. By Property (AF4), this implies that $\mathrm{C}\left(\mathcal{S} W, T_{u}^{[j]}\right) \geq t_{\mathcal{H}}+1$. Since we have already established that $\mathrm{C}\left(\mathcal{S} W, T_{u}^{[j]}\right) \leq t_{\mathcal{H}}+1$, we have $\mathrm{C}\left(\mathcal{S} W, T_{u}^{[j]}\right)=t_{\mathcal{H}}+1$.

Lemma 5.13. Let $T_{u}$ be a task in $A$ and let $T_{u}^{[j]}$ be the subtask of $T_{u}$ that is scheduled in slot $t_{\mathcal{H}}$. If $\mathrm{D}\left(T_{u}^{[j]}\right)=0$, then there is no hole in slot $t_{\mathcal{H}}+1$.

Proof. Let $T_{u}$ and $T_{u}^{[j]}$ be as defined in the statement of the lemma. Since tasks $t_{\mathcal{H}}<t_{d}$, $T_{u}^{[j]}$ is scheduled before its deadline, i.e., $t_{\mathcal{H}}<\mathrm{d}\left(T_{u}^{[j]}\right)$. Moreover, since tasks are scheduled sequentially, it follows that $T_{u}^{[j]}$ has the largest index of any subtask of $T_{u}$ scheduled at or before $t_{\mathcal{H}}$. ( $T_{u}$ is illustrated in Figure 5.22.)

Thus, by Lemma 5.6, $\mathrm{d}\left(T_{u}^{[j]}\right)=t_{\mathcal{H}}+1 \wedge \mathrm{~b}\left(T_{u}^{[j]}\right)=1$. In addition, since there is a hole in $t_{\mathcal{H}}$, by Lemma 5.8, $T_{u}^{[j]}$,s successor, $T_{u}^{[j+1]}$, is released at time $t_{\mathcal{H}}$, i.e., $\mathrm{r}\left(T_{u}^{[j+1]}\right)=t_{\mathcal{H}}$.

We now assume that there is a hole in slot $t_{\mathcal{H}}+1$ to derive a contradiction. By this assumption and Lemma 5.9, it follows that $T_{u}^{[j+1]}$ is scheduled in slot $t_{\mathcal{H}}+1$. Since $T_{u}^{[j+1]}$ is scheduled at $t_{\mathcal{H}}+1<t_{d}, T_{u}^{[j+1]}$ does not miss its deadline. Hence, $\mathrm{d}\left(T_{u}^{[j+1]}\right) \geq t_{\mathcal{H}}+2$. Since subtasks are scheduled in index order, $T_{u}^{[j+1]}$ has the largest index of any subtask of $T_{u}$ scheduled at or before $t_{\mathcal{H}}+1$. Thus, by Lemma 5.6, $\mathrm{d}\left(T_{u}^{[j+1]}\right)=t_{\mathcal{H}}+2 \wedge \mathrm{~b}\left(T_{u}^{[j+1]}\right)=1$. Thus, because (as was already established) $\mathrm{r}\left(T_{u}^{[j+1]}\right)=t_{\mathcal{H}}$, it follows that $\mathrm{d}\left(T_{u}^{[j+1]}\right)-\mathrm{r}\left(T_{u}^{[j+1]}\right)=2$.

Since $T_{u}^{[j]}$ is scheduled in slot $t_{\mathcal{H}}, \mathrm{b}\left(T_{u}^{[j]}\right)=1, \mathrm{D}\left(T_{u}^{[j]}\right)=0$, and $\mathrm{r}\left(T_{u}^{[j+1]}\right)=t_{\mathcal{H}}$, by Property (GD-3), $\mathrm{D}\left(T_{u}^{[j+1]}\right)=0$. Thus, by Property (WL), $\mathrm{d}\left(T_{u}^{[j+1]}\right)-\mathrm{r}\left(T_{u}^{[j+1]}\right) \geq 3$; however, this contradicts our earlier assertion that $\mathrm{d}\left(T_{u}^{[j+1]}\right)-\mathrm{r}\left(T_{u}^{[j+1]}\right)=2$. Thus, there cannot exist a hole in slot $t_{\mathcal{H}}+1$, which completes the proof.

Corollary 5.4. Let $T_{b}$ be a task in $B$ and let $T_{b}^{[c]}$ be the subtask of $T_{b}$ with the largest index such that $\mathrm{r}\left(T_{b}^{[c]}\right) \leq t_{\mathcal{H}}<\mathrm{d}\left(T_{b}^{[c]}\right)$. If $\mathrm{D}\left(T_{b}^{[c]}\right)=0$, then there is no hole in slot $t_{\mathcal{H}}+1$.

Proof. Let $T_{b}$ and $T_{b}^{[c]}$ be as defined in the statement of the lemma. Since by the statement of a lemma, $\mathrm{D}\left(T_{b}^{[c]}\right)=0$, it follows, by Lemma 5.7, that there is a subtask, $T_{u}^{[j]}$, of a task $T_{u}$ in $A$ such that $\mathrm{D}\left(T_{u}^{[j]}\right)=0$ and $T_{u}^{[j]}$ is scheduled at time $t_{\mathcal{H}}$. Thus, by Lemma 5.13, there is no hole in slot $t_{\mathcal{H}}+1$.

Lemma 5.14. Let $T_{u}$ be a task in $B$ and let $T_{u}^{[j]}$ be the subtask of $T_{u}$ with the largest index such that $\mathrm{r}\left(T_{u}^{[j]}\right) \leq t_{\mathcal{H}}<\mathrm{d}\left(T_{u}^{[j]}\right)$. If $\mathrm{D}\left(T_{u}^{[j]}\right)>0$, then there exists a slot with no holes in $\left[\mathrm{d}\left(T_{u}^{[j]}\right), \min \left(\mathrm{D}\left(T_{u}^{[j]}\right), t_{d}\right)\right)$.

Proof. By Lemma 5.6, $\mathrm{d}\left(T_{u}^{[j]}\right)=t_{\mathcal{H}}+1 \wedge \mathrm{~b}\left(T_{u}^{[j]}\right)=1$. By (5.16), $t_{\mathcal{H}}<t_{d}-1$. Thus, $\mathrm{d}\left(T_{u}^{[j]}\right)<t_{d}-1$. Thus, by Lemma $5.5(\mathrm{~b})$, if $\mathrm{D}\left(T_{u}^{[j]}\right) \geq t_{d}-1$, then the proof is complete. Therefore, for the remainder of the proof, we assume that

$$
\begin{equation*}
\mathrm{D}\left(T_{u}^{[j]}\right)<t_{d}-1 . \tag{5.17}
\end{equation*}
$$

Since $\mathrm{D}\left(T_{u}^{[j]}\right)>0$ and $\mathrm{b}\left(T_{u}^{[j]}\right)=1$, by Property (GD-2), we have that $\mathrm{D}\left(T_{u}^{[j]}\right)>\mathrm{d}\left(T_{u}^{[j]}\right)$. Thus,

$$
\begin{equation*}
\mathrm{D}\left(T_{u}^{[j]}\right) \geq \mathrm{d}\left(T_{u}^{[j]}\right)+1=t_{\mathcal{H}}+2 . \tag{5.18}
\end{equation*}
$$

By Lemma 5.7, we have the following property:
(E-1) There exists a subtask $T_{x}^{[z]}$ scheduled in slot $t_{\mathcal{H}}$ such that $\mathrm{D}\left(T_{x}^{[z]}\right) \leq \mathrm{D}\left(T_{u}^{[j]}\right)$.
( $T_{x}$ is illustrated in Figure 5.23.) By Lemma 5.13, if $\mathrm{D}\left(T_{x}^{[z]}\right)=0$, then there is no hole in slot $t_{\mathcal{H}}+1$. Since, we have assumed that $\mathrm{D}\left(T_{u}^{[j]}\right)<t_{d}-1$ and, by (5.18), $\mathrm{D}\left(T_{u}^{[j]}\right) \geq t_{\mathcal{H}}+1=\mathrm{d}\left(T_{u}^{[j]}\right)$, it follows that if $\mathrm{D}\left(T_{x}^{[z]}\right)=0$, then there exists a slot with no holes in $\left[\mathrm{d}\left(T_{u}^{[j]}\right), \min \left(\mathrm{D}\left(T_{u}^{[j]}\right), t_{d}\right)\right)$. Thus, for the remainder of the proof we assume that

$$
\begin{equation*}
\mathrm{D}\left(T_{x}^{[z]}\right)>0 . \tag{5.19}
\end{equation*}
$$

Since $T_{x}^{[z]}$ is scheduled in slot $t_{\mathcal{H}}$, and $t_{\mathcal{H}}<t_{d}$, it follows $T_{x}^{[z]}$ is scheduled before its
deadline. Thus, since a subtask must be released before it is scheduled, $\mathrm{r}\left(T_{x}^{[z]}\right) \leq t_{\mathcal{H}}<\mathrm{d}\left(T_{x}^{[z]}\right)$. Moreover, since subtasks are scheduled sequentially, $T_{x}^{[z]}$ has the largest index of any subtask of $T_{x}$ that is scheduled at or before $t_{\mathcal{H}}$ such that $\mathrm{r}\left(T_{x}^{[z]}\right) \leq t_{\mathcal{H}}<\mathrm{d}\left(T_{x}^{[z]}\right)$. Thus, by Lemma 5.6, $\mathrm{d}\left(T_{x}^{[z]}\right)=t_{\mathcal{H}}+1 \wedge \mathrm{~b}\left(T_{x}^{[z]}\right)=1$. Thus, by (5.19) and Properties (E-1) and (GD-2), it follows that

$$
\begin{equation*}
t_{\mathcal{H}}+2 \leq \mathrm{D}\left(T_{x}^{[z]}\right) \leq \mathrm{D}\left(T_{u}^{[j]}\right) . \tag{5.20}
\end{equation*}
$$

To derive a contradiction, we assume that every slot in the set $\left\{t_{\mathcal{H}}, \ldots, \mathrm{D}\left(T_{u}^{[j]}\right)-1\right\}$ contains a hole. Since by $(5.20), \mathrm{D}\left(T_{x}^{[z]}\right) \leq \mathrm{D}\left(T_{u}^{[j]}\right)$, if there is a hole in every slot in the set $\left\{t_{\mathcal{H}}, \ldots, \mathrm{D}\left(T_{u}^{[j]}\right)-1\right\}$, then we have the following property.
(E-2) Every slot in the set $\left\{t_{\mathcal{H}}, \ldots, \mathrm{D}\left(T_{x}^{[z]}\right)-1\right\}$ contains a hole.
(Notice that the set $\left\{t_{\mathcal{H}}, \ldots, \mathrm{D}\left(T_{x}^{[z]}\right)-1\right\}$ is non-empty since by (5.20), $t_{\mathcal{H}}+2 \leq \mathrm{D}\left(T_{x}^{[z]}\right)$.) By Lemma 5.9 and Properties (E-1) and (E-2), it follows that $T_{x}$ is scheduled in every slot in the set $\left\{t_{\mathcal{H}}, \ldots, \mathrm{D}\left(T_{x}^{[z]}\right)-1\right\}$. Thus, by Corollary 5.2 , we have the following property.
(E-3) For all $q \in\left\{z, \ldots, z+\left(\mathrm{D}\left(T_{x}^{[z]}\right)-t_{\mathcal{H}}\right)\right\}, \theta\left(T_{x}^{[z]}\right)=\theta\left(T_{x}^{[q]}\right)$.
Let $k$ equal the index of the subtask $\omega\left(T_{x}^{[z]}\right)$. Notice that $\mathrm{D}\left(T_{x}^{[z]}\right)-t_{\mathcal{H}}+z=k$ since the number of $T_{x}$ 's subtasks that are released after $\mathrm{r}\left(T_{x}^{[z]}\right)$ but before $\mathrm{D}\left(T_{x}^{[z]}\right)$ equals $\mathrm{D}\left(T_{x}^{[z]}\right)-t_{\mathcal{H}}$. Thus, by Property (E-3), it follows that for all $T_{x}^{[q]} \in\left\{T_{x}^{[z]}, \ldots, \omega\left(T_{x}^{[z]}\right)\right\}, \theta\left(T_{x}^{[z]}\right)=\theta\left(T_{x}^{[q]}\right)$. Thus, by Property (GD-1), the subtask $\omega\left(T_{x}^{[z]}\right)$, has the following two properties
(Y-1) $\mathrm{r}\left(\omega\left(T_{x}^{[z]}\right)\right)=\mathrm{D}\left(T_{x}^{[z]}\right)-2$.
(Y-2) Either $\mathrm{d}\left(\omega\left(T_{x}^{[z]}\right)\right)=\mathrm{D}\left(T_{x}^{[z]}\right)$ and $\mathrm{b}\left(\omega\left(T_{x}^{[z]}\right)\right)=0$ or $\mathrm{d}\left(\omega\left(T_{x}^{[z]}\right)\right)=\mathrm{D}\left(T_{x}^{[z]}\right)+1$ and $\mathrm{b}\left(\omega\left(T_{x}^{[z]}\right)\right)=1$.

We now show that $\omega\left(T_{x}^{[z]}\right)$ cannot satisfy both Properties (Y-1) and (Y-2), which contradicts Property (E-2) and completes the proof.

Since $t_{\mathcal{H}}<\mathrm{D}\left(T_{x}^{[z]}\right)-1($ by $(5.20))$, and $\mathrm{D}\left(T_{x}^{[z]}\right)-1<t_{d}-1$ (by 5.17 ), by Properties (E-1) and (E-2), we have the following.
(E-4) $t_{\mathcal{H}}<\mathrm{D}\left(T_{x}^{[z]}\right)-1 \leq t_{d}-1$, where there is a hole in every slot in the range $\left\{t_{\mathcal{H}}, \ldots, \mathrm{D}\left(T_{x}^{[z]}\right)-\right.$ $1\}$, and $T_{x}^{[z]}$ is scheduled in slot $t_{\mathcal{H}}$.


Figure 5.23: An illustration of the two possibilities for $T_{x}$ in Lemma 5.14.
Recall that, by Property $(\mathrm{Y}-1), \mathrm{r}\left(\omega\left(T_{x}^{[z]}\right)\right)=\mathrm{D}\left(T_{x}^{[z]}\right)-2$. Thus, by Corollary 5.3 and Property (E-4), it follows that $\omega\left(T_{x}^{[z]}\right)$ is scheduled in slot $\mathrm{r}\left(\omega\left(T_{x}^{[z]}\right)\right)+1=\mathrm{D}\left(T_{x}^{[z]}\right)-1$.

Since $\omega\left(T_{x}^{[z]}\right)$ is scheduled at time $D\left(T_{x}^{[z]}\right)-1$, by Property (E-2), there is a hole at the time it is scheduled. Thus, by Lemma 5.6, it follows that $\mathrm{b}\left(\omega\left(T_{x}^{[z]}\right)\right)=1$ and $\mathrm{d}\left(\omega\left(T_{x}^{[z]}\right)\right)=\mathrm{D}\left(T_{x}^{[z]}\right)$; however, this contradicts Property (Y-2).

The following lemma contradicts our choice of $t_{\mathcal{H}}$ as the last slot such that $\operatorname{LAG}\left(\tau, t_{\mathcal{H}}\right)<1$.
Lemma 5.15. Let $T_{b}$ be a task in $B$ and let $T_{b}^{[c]}$ be the subtask of $T_{b}$ with the largest index such that $\mathrm{r}\left(T_{b}^{[c]}\right) \leq t_{\mathcal{H}}<\mathrm{d}\left(T_{b}^{[c]}\right)$. If $\mathrm{D}\left(T_{b}^{[c]}\right)=0$, then $\operatorname{LAG}\left(\tau, t_{\mathcal{H}}+2\right)<1$.

Proof. Let the number of holes in slot $t_{\mathcal{H}}$ be $h$. Assume that $T_{b}^{[c]}$ exists and is defined as in the statement of the lemma. We now derive some properties about the per-slot allocations to tasks in the $\mathcal{S} W$ schedule in slots $t_{\mathcal{H}}$ and $t_{\mathcal{H}}+1$.

By the definition of $I$, if task $T_{y}$ is in $I$, then $\mathrm{A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}\right)=0$. Since $\tau=A \cup$ $B \cup I, \sum_{T_{y} \in \tau} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}\right)=\sum_{T_{y} \in A \cup B} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}\right)$. Since $\operatorname{Swt}\left(T_{y}, t\right) \leq 1$, for any task $T_{y}$ and any time $t$, we have $\sum_{T_{y} \in A} \operatorname{Swt}\left(T_{y}, t_{\mathcal{H}}\right) \leq|A|$. Thus, by Property (AF1), $\sum_{T_{y} \in A} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}\right) \leq|A|$. Because there are $h$ holes in slot $t_{\mathcal{H}}, M-h$ tasks are scheduled at $t_{\mathcal{H}}$, i.e., $|A|=M-h$. Thus, $\sum_{T_{y} \in A} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}\right) \leq M-h$, and hence

$$
\begin{equation*}
\sum_{T_{y} \in \tau} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}\right) \leq M-h+\sum_{T_{y} \in B} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}\right) . \tag{5.21}
\end{equation*}
$$

Let $C$ denote the set of tasks that receive a positive allocation in $\mathcal{S} W$ in slot $t_{\mathcal{H}}+1$ and are not in $B$. Then, the set of tasks that receive a positive allocation in $\mathcal{S} W$ is a subset of $C \cup B$. Thus, by Property (W) in Section 5.5,

$$
\begin{equation*}
\sum_{T_{y} \in C \cup B} \operatorname{Swt}\left(T_{y}, t_{\mathcal{H}}+1\right) \leq M \tag{5.22}
\end{equation*}
$$

Also, $\sum_{T_{y} \in \mathcal{T}} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}+1\right)=\sum_{T_{y} \in C \cup B} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}+1\right)$. By Property (AF1), this implies that $\sum_{T_{y} \in \tau} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}+1\right) \leq \sum_{T_{y} \in C} \operatorname{Swt}\left(T_{y}, t_{\mathcal{H}}+1\right)+\sum_{T_{y} \in B} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}+1\right)$. Thus, by (5.21),

$$
\begin{equation*}
\sum_{T_{y} \in \tau} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}, t_{\mathcal{H}}+2\right) \leq M-h+\sum_{T_{y} \in C} \operatorname{Swt}\left(T_{y}, t_{\mathcal{H}}+1\right)+\sum_{T_{y} \in B} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}, t_{\mathcal{H}}+2\right) \tag{5.23}
\end{equation*}
$$

Consider $T_{u} \in B$. Let $T_{u}^{[j]}$ be the subtask of $T_{u}$ with the largest index such that $\mathrm{r}\left(T_{u}^{[j]}\right) \leq$ $t_{\mathcal{H}}<\mathrm{d}\left(T_{u}^{[j]}\right)$ that is scheduled before $t_{\mathcal{H}}$. Let $D$ denote the set of such subtasks for all tasks in $B$. Then, by Lemmas 5.6 and 5.12,

$$
\begin{equation*}
\text { for all } T_{u}^{[j]} \in D, \mathrm{C}\left(\mathcal{S} W, T_{u}^{[j]}\right)=\mathrm{d}\left(T_{u}^{[j]}\right)=t_{\mathcal{H}}+1 \wedge \mathrm{~b}\left(T_{u}^{[j]}\right)=1 . \tag{5.24}
\end{equation*}
$$

By (TK), $T_{u}^{[j]}$ s successor $T_{u}^{[k]}$ (if it exists) is not released until at or after $t_{\mathcal{H}}+1 \geq \mathrm{C}\left(\mathcal{S} W, T_{u}^{[j]}\right)$. Since $\mathbf{r}\left(T_{u}^{[k]}\right) \geq \mathrm{C}\left(\mathcal{S} W, T_{u}^{[j]}\right)$ and $\mathbf{b}\left(T_{u}^{[j]}\right)=1$, by (AF2), $\mathrm{A}\left(\mathcal{S} W, T_{u}, t_{\mathcal{H}}, t_{\mathcal{H}}+2\right) \leq \operatorname{Swt}\left(T_{u}, t_{\mathcal{H}}+\right.$ 1). Thus, $\sum_{T_{y} \in B} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}, t_{\mathcal{H}}+2\right) \leq \sum_{T_{y} \in B} \operatorname{Swt}\left(T_{y}, t_{\mathcal{H}}+1\right)$.

By (5.23), this implies that

$$
\sum_{T_{y} \in \tau} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}, t_{\mathcal{H}}+2\right) \leq M-h+\sum_{T_{y} \in C \cup B} \operatorname{Swt}\left(T_{y}, t_{\mathcal{H}}+1\right) .
$$

Thus, from (5.22) it follows that

$$
\begin{equation*}
\sum_{T_{y} \in \tau} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}, t_{\mathcal{H}}+2\right) \leq M-h+M . \tag{5.25}
\end{equation*}
$$

Notice that, by Corollary 5.4, the existence of $T_{b}^{[c]}$ (from the statement of the lemma) im-
plies that there is no hole in slot $t_{\mathcal{H}}+1$. Thus, since there are $h$ holes in slot $t_{\mathcal{H}}$, we have $\mathrm{A}\left(\mathcal{S}, \tau, t_{\mathcal{H}}, t_{\mathcal{H}}+2\right)=M-h+M$. Hence, by (5.25), $\sum_{T_{y} \in \tau} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}, t_{\mathcal{H}}+2\right) \leq$ $\sum_{T_{y} \in \tau} \mathrm{~A}\left(\mathcal{S}, \tau, t_{\mathcal{H}}, t_{\mathcal{H}}+2\right)$. Using this relation in (5.1) we obtain, $\operatorname{LAG}\left(\tau, t_{\mathcal{H}}+2\right)=\operatorname{LAG}\left(\tau, t_{\mathcal{H}}\right)+$ $\sum_{T_{y} \in \tau} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}, t_{\mathcal{H}}+2\right)-\sum_{T_{y} \in \tau} \mathrm{~A}\left(\mathcal{S}, T_{y}, t_{\mathcal{H}}, t_{\mathcal{H}}+2\right)$ (recall that for this section, we have assumed that $\operatorname{LAG}(\tau, t)=\operatorname{LAG}(\mathcal{S}, \mathcal{S} W, \tau, t))$. Since, $\operatorname{LAG}\left(\tau, t_{\mathcal{H}}\right)<1$, we have $\operatorname{LAG}\left(\tau, t_{\mathcal{H}}+\right.$ $2)<1$.

Lemma 5.16. If $t_{\mathcal{N}}$ is the first time after $t_{\mathcal{H}}$ such that there are no holes in the schedule $\mathcal{S}$, then $\operatorname{LAG}\left(\tau, t_{\mathcal{N}}+1\right)<1$.

Proof. By Corollary 5.4, if there exists a task $T_{b}$ in $B$ with a subtask $T_{b}^{[c]}$, where $\mathrm{D}\left(T_{b}^{[c]}\right)=0$, and $T_{b}^{[c]}$ has the largest index of any task in $T_{b}$ such that $\mathrm{r}\left(T_{b}^{[c]}\right) \leq t_{\mathcal{H}}<\mathrm{d}\left(T_{b}^{[c]}\right)$ holds, then there is no hole in slot $t_{\mathcal{H}}+1$ (i.e., $t_{\mathcal{N}}=t_{\mathcal{H}}+1$ ), and by Lemma 5.15, $\operatorname{LAG}\left(\tau, t_{\mathcal{N}}+1\right)<1$. Thus, for the remainder of this proof we assume that no such task is in $B$.

Let the number of holes in slot $t_{\mathcal{H}}$ be $h$. By Corollary 5.1, only tasks in $A$ are scheduled in slots in the range $\left\{t_{\mathcal{H}}, \ldots, t_{\mathcal{N}}-1\right\}$, and by Lemma 5.9, every task in $A$ is scheduled in every slot in this range. Thus, we have the following property
(H) There are $h$ holes in every slot in the range $\left\{t_{\mathcal{H}}, \ldots, t_{\mathcal{N}}-1\right\}$.

We now derive some properties about the per-slot allocations to tasks in the $\mathcal{S} W$ schedule in the slots $\left\{t_{\mathcal{H}}, \ldots, t_{\mathcal{N}}\right\}$.

Allocations in $\mathcal{S} W$ in slot $t_{\mathcal{H}}$. By the definition of set $I$, if a task $T_{y}$ is in $I$, then $\mathrm{A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}\right)=0$. Since $\tau=A \cup B \cup I, \sum_{T_{y} \in \tau} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}\right)=\sum_{T_{y} \in A \cup B} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}\right)$. Since $\operatorname{Swt}\left(T_{y}, t\right) \leq 1$ for any task $T_{y}$, we have $\sum_{T_{y} \in A} \operatorname{Swt}\left(T_{y}, t_{\mathcal{H}}\right) \leq|A|$. Thus, by Property (AF1), $\sum_{T_{y} \in A} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}\right) \leq|A|$. Because there are $h$ holes in slot $t_{\mathcal{H}}, M-h$ tasks are scheduled at $t_{\mathcal{H}}$, i.e., $|A|=M-h$. Thus, $\sum_{T_{y} \in A} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}\right) \leq M-h$, and hence

$$
\begin{equation*}
\sum_{T_{y} \in \mathcal{T}} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}\right) \leq M-h+\sum_{T_{y} \in B} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}\right) . \tag{5.26}
\end{equation*}
$$

Allocations in $\mathcal{S} W$ between slots $t_{\mathcal{H}}$ and $t_{\mathcal{N}}$. By Corollary 5.1, only tasks in $A$ release subtasks in the slots $\left\{t_{\mathcal{H}}+1, \ldots, t_{\mathcal{N}}-1\right\}$. (If a subtask of a task not in $A$ was released within
this range, then it would be scheduled within this range since there are holes in every slot.) Moreover, by Lemma 5.6, any task that is released at or before $t_{\mathcal{H}}$ and is not part of $A$ must have a deadline no later than $t_{\mathcal{H}}+1$. Thus, by Properties (AF3) and (AF4), the allocations to all tasks in $\mathcal{S} W$, not in $A$, in the slots $\left\{t_{\mathcal{H}}+1, \ldots, t_{\mathcal{N}}-1\right\}$ is zero. Moreover, reasoning as above, $\sum_{T_{y} \in \tau} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t\right) \leq M-h$ holds for any slot $t$ in this range. Therefore,

$$
\begin{equation*}
\sum_{T_{y} \in \tau} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}+1, t_{\mathcal{N}}\right) \leq\left(t_{\mathcal{N}}-t_{\mathcal{H}}-2\right) \cdot(M-h) \tag{5.27}
\end{equation*}
$$

Allocations in $\mathcal{S} W$ in slot $t_{\mathcal{N}}$. Let $C$ denote the set of tasks that receive a positive allocation in $\mathcal{S} W$ in slot $t_{\mathcal{N}}$ and are not in $B$. Then, the set of tasks that receive a positive allocation in $\mathcal{S} W$ is a subset of $C \cup B$. By Property (W) in Section 5.5,

$$
\begin{equation*}
\sum_{T_{y} \in C \cup B} \operatorname{Swt}\left(T_{y}, t_{\mathcal{N}}\right) \leq M . \tag{5.28}
\end{equation*}
$$

Also, $\sum_{T_{y} \in \tau} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{N}}\right)=\sum_{T_{y} \in C \cup B} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{N}}\right)$. By Property (AF1), this implies that

$$
\begin{equation*}
\sum_{T_{y} \in \tau} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{N}}\right) \leq \sum_{T_{y} \in C} \operatorname{Swt}\left(T_{y}, t_{\mathcal{N}}\right)+\sum_{T_{y} \in B} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{N}}\right) . \tag{5.29}
\end{equation*}
$$

By (5.26), (5.27), and (5.29), the total SW allocation in the slots $\left\{t_{\mathcal{H}}, \ldots, t_{\mathcal{N}}\right\}$ is

$$
\begin{align*}
\sum_{T_{y} \in \tau} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}, t_{\mathcal{N}}\right) \leq & \sum_{T_{y} \in C} \operatorname{Swt}\left(T_{y}, t_{\mathcal{N}}\right)+\left(t_{\mathcal{N}}-t_{\mathcal{H}}-1\right) \cdot(M-h)  \tag{5.30}\\
& +\sum_{T_{y} \in B} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}, t_{\mathcal{N}}+1\right) .
\end{align*}
$$

Constructing an upper bound for $\sum_{T_{y} \in B} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}, t_{\mathcal{N}}+1\right)$. Consider $T_{u} \in B$. Let $T_{u}^{[j]}$ be the subtask of $T_{u}$ with the largest index such that $\mathrm{r}\left(T_{u}^{[j]}\right) \leq t_{\mathcal{H}}<\mathrm{d}\left(T_{u}^{[j]}\right)$ that is scheduled before $t_{\mathcal{H}}$. Let $D$ denote the set of such subtasks for all tasks in $B$. Notice that, by the assumption we made at the beginning of this proof, for any subtask $T_{u}^{[j]}$ in $D$,

$$
\begin{equation*}
\mathrm{D}\left(T_{u}^{[j]}\right)>0 \tag{5.31}
\end{equation*}
$$

and $\omega\left(T_{u}^{[j]}\right)$ is defined. By Lemmas 5.6 and 5.12,

$$
\begin{equation*}
\mathrm{C}\left(\mathcal{S} W, T_{u}^{[j]}\right)=\mathrm{d}\left(T_{u}^{[j]}\right)=t_{\mathcal{H}}+1 \wedge \mathrm{~b}\left(T_{u}^{[j]}\right)=1 . \tag{5.32}
\end{equation*}
$$

In addition, since, by the statement of the lemma $t_{\mathcal{H}}<t_{\mathcal{N}}$, it follows from (5.32) that

$$
\begin{equation*}
\mathrm{C}\left(\mathcal{S} W, T_{u}^{[j]}\right)=t_{\mathcal{H}}+1 \leq t_{\mathcal{N}} . \tag{5.33}
\end{equation*}
$$

Moreover, by Lemma $5.14, t_{\mathcal{N}}$ is before the earliest group deadline of any subtask in $D$. Thus,

$$
\begin{equation*}
t_{\mathcal{N}}<\mathrm{D}\left(T_{u}^{[j]}\right) . \tag{5.34}
\end{equation*}
$$

Let $T_{u}^{[f]}$ be $T_{u}^{[j]}$, s next present successor. We now show that, if $T_{u}^{[f]}$ does not exist, then by Property (AF5), $\mathrm{A}\left(\mathcal{S} W, T_{u}, t_{\mathcal{H}}, t_{\mathcal{N}}+1\right) \leq \operatorname{Swt}\left(T_{u}, t_{\mathcal{N}}\right)$. In order to apply (AF5) the following conditions must hold: (i) $\mathrm{D}\left(T_{u}^{[j]}\right)>0$; (ii) $\mathrm{C}\left(\mathcal{S} W, T_{u}^{[j]}\right)-1 \leq t_{\mathcal{N}}$; and (iii) $t_{\mathcal{N}} \leq \mathrm{D}\left(T_{u}^{[j]}\right)$. Condition (i) holds by (5.31). Condition (ii) holds by (5.33). Condition (iii) holds by (5.34). Thus, Property (AF5) applies, and $\mathrm{A}\left(\mathcal{S} W, T_{u}, t_{\mathcal{H}}, t_{\mathcal{N}}+1\right) \leq \operatorname{Swt}\left(T_{u}, t_{\mathcal{N}}\right)$.

We now consider the possibility where $T_{u}^{[f]}$ exists. By Corollary 5.1, since no subtask of any task in $B$ is scheduled in slot $t_{\mathcal{H}}$, no subtask of any task in $B$ is scheduled over the range $\left\{t_{\mathcal{H}}, \ldots, t_{\mathcal{N}}-1\right\}$. Thus, since there are holes in every slot in this range, it follows that no task in $B$ releases a present subtask at a slot within the range $\left\{t_{\mathcal{H}}, \ldots, t_{\mathcal{N}}-1\right\}$. Thus, if $T_{u}^{[f]}$ exists, then $t_{\mathcal{N}} \leq \mathrm{r}\left(T_{u}^{[f]}\right)$. Hence, by (5.34),

$$
\begin{equation*}
t_{\mathcal{N}} \leq \min \left(\mathrm{r}\left(T_{u}^{[f]}\right), \mathrm{D}\left(T_{u}^{[j]}\right)\right) . \tag{5.35}
\end{equation*}
$$

We now apply Property (AF5) to show that $\mathrm{A}\left(\mathcal{S} W, T_{u}, t_{\mathcal{H}}, t_{\mathcal{N}}+1\right) \leq \operatorname{Swt}\left(T_{u}, t_{\mathcal{N}}\right)$. In order for Property (AF5) to apply, the following conditions must hold: (i) $\mathrm{D}\left(T_{u}^{[j]}\right)>0$; (ii) $\mathrm{C}\left(\mathcal{S} W, T_{u}^{[j]}\right) \leq \mathrm{r}\left(T_{u}^{[f]}\right)$; (iii) $\mathrm{C}\left(\mathcal{S} W, T_{u}^{[j]}\right)-1 \leq t_{\mathcal{N}}$; and (iv) $t_{\mathcal{N}} \leq \min \left(\mathrm{r}\left(T_{u}^{[f]}\right), \mathrm{D}\left(T_{u}^{[j]}\right)-1\right)$. Condition (i) holds by (5.31). By (5.35) and (5.33), $\mathrm{C}\left(\mathcal{S} W, T_{u}^{[j]}\right)=t_{\mathcal{H}}+1 \leq t_{\mathcal{N}} \leq$ $\min \left(\mathrm{r}\left(T_{u}^{[f]}\right), \mathrm{D}\left(T_{u}^{[j]}\right)\right)$. This implies that $\mathrm{C}\left(\mathcal{S} W, T_{u}^{[j]}\right)-1 \leq \mathrm{r}\left(T_{u}^{[f]}\right)$ holds, which satisfies Condition (ii). Similarly, (5.33) implies that Condition (iii) holds, and (5.35) implies that Condition
(iv) holds. Thus, by Property (AF5), it follows that $\mathrm{A}\left(\mathcal{S} W, T_{u}, t_{\mathcal{H}}, t_{\mathcal{N}}+1\right) \leq \operatorname{Swt}\left(T_{u}, t_{\mathcal{N}}\right)$.

Thus, regardless of whether $T_{u}^{[f]}$ exists or not, $\sum_{T_{y} \in B} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}, t_{\mathcal{N}}+1\right) \leq$ $\sum_{T_{y} \in B} \operatorname{Swt}\left(T_{y}, t_{\mathcal{N}}\right)$.

By (5.30), this implies that

$$
\sum_{T_{y} \in \tau} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}, t_{\mathcal{N}}+1\right) \leq\left(t_{\mathcal{N}}-t_{\mathcal{H}}-1\right) \cdot(M-h)+\sum_{T_{y} \in C \cup B} \operatorname{Swt}\left(T_{y}, t_{\mathcal{N}}\right) .
$$

Thus, from (5.28) it follows that

$$
\begin{equation*}
\sum_{T_{y} \in \tau} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}, t_{\mathcal{N}}+1\right) \leq\left(t_{\mathcal{N}}-t_{\mathcal{H}}-1\right) \cdot(M-h)+M \tag{5.36}
\end{equation*}
$$

Completing the proof. Since, by the statement of the lemma, there is no hole in slot $t_{\mathcal{N}}$ and by Property (H), there are $h$ holes in every slot in the range $\left\{t_{\mathcal{H}}, \ldots, t_{\mathcal{N}}-1\right\}$, we have

$$
\mathrm{A}\left(\mathcal{S}, \tau, t_{\mathcal{H}}, t_{\mathcal{N}}+1\right)=\left(t_{\mathcal{N}}-t_{\mathcal{H}}-1\right) \cdot(M-h)+M
$$

Hence, by (5.36), $\sum_{T_{y} \in \tau} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}, t_{\mathcal{N}}+1\right) \leq \sum_{T_{y} \in \tau} \mathrm{~A}\left(\mathcal{S}, \tau, t_{\mathcal{H}}, t_{\mathcal{N}}+1\right)$. By (5.1), $\operatorname{LAG}\left(\tau, t_{\mathcal{N}}+1\right)=\operatorname{LAG}\left(\tau, t_{\mathcal{H}}\right)+\sum_{T_{y} \in \tau} \mathrm{~A}\left(\mathcal{S} W, T_{y}, t_{\mathcal{H}}, t_{\mathcal{N}}+1\right)-\sum_{T_{y} \in \tau} \mathrm{~A}\left(\mathcal{S}, T_{y}, t_{\mathcal{H}}, t_{\mathcal{N}}+1\right)$ (recall that for this section, we have assumed that $\operatorname{LAG}(\tau, t)=\operatorname{LAG}(\mathcal{S}, \mathcal{S} W, \tau, t))$. Since, $\operatorname{LAG}\left(\tau, t_{\mathcal{H}}\right)<1$, we obtain $\operatorname{LAG}\left(\tau, t_{\mathcal{N}}+1\right)<1$.

It follows, by Lemma 5.16, there exists a time, $t$ after $t_{\mathcal{H}}$ such that $\operatorname{LAG}(\tau, t)<1$. This contradicts (5.16), which implies that $t_{d}$ does not exist. Thus, Theorem 5.2 holds.

### 5.6 Drift

We now turn our attention to the issue of measuring drift under PD-PNH. In order to measure the drift of a task system $\tau$, we introduce two additional theoretical scheduling algorithms: the clairvoyant scheduling-weight (CSW) scheduling algorithm; and the ideal (IDEAL) scheduling algorithm. The CSW scheduling algorithm is the same as the SW scheduling algorithm except that the CSW scheduling algorithm is "clairvoyant" so that it does not allocate capacity to
tasks that will halt. Under the IDEAL scheduling algorithm, at each instant $t$, each task $T_{i}$ in $\tau$ is allocated a share equal to its weight $\mathrm{wt}\left(T_{i}, t\right)$. Hence, over the interval $\left[t_{1}, t_{2}\right)$, the task $T_{i}$ is allocated $\mathrm{A}\left(\mathcal{I}, T_{i}, t_{1}, t_{2}\right)=\int_{t_{1}}^{t_{2}} \mathrm{wt}\left(T_{i}, u\right) d u$ time. We use $\mathcal{C} S W$ and $\mathcal{I}$ to denote, respectively, CSW and IDEAL schedules for a given system.

Using the definition of $\mathcal{S} W$, we can simply define $\mathcal{C} S W$ as follows:

$$
\mathrm{A}\left(\mathcal{C} S W, T_{i}^{[j]}, t\right)= \begin{cases}\mathrm{A}\left(\mathcal{S} W, T_{i}^{[j]}, t\right), & \text { if } T_{i}^{[j]} \text { never halts } \\ 0, & \text { otherwise. }\end{cases}
$$

For example, in Figure 5.4(a) $\mathrm{A}\left(\mathcal{S} W, T_{i}^{[j]}, t\right)=\mathrm{A}\left(\mathcal{C} S W, T_{i}^{[j]}, t\right)$ except for $T_{1}^{[2]}$, where $\mathrm{A}\left(\mathcal{C} S W, T_{1}^{[2]}, t\right)=0$, for all $t$. We use $\mathcal{C} S W$ and $\mathcal{I}$ to denote, respectively, CSW and IDEAL schedules for a given system.

For the remainder of this section, we assume that every subtask in $T_{i}$ is released as early as possible. This assumption can be removed at the cost of more complex notation. If we did not make this assumption, then the allocation function for $\mathcal{I}$ would equal zero between active subtasks.

Comparing IDEAL to SW and CSW. $\mathcal{I}$ is similar to $\mathcal{S} W$ and $\mathcal{C} S W$, with three major exceptions: (i) tasks in $\mathcal{I}$ continually receive allocations, whereas tasks in $\mathcal{S} W$ and $\mathcal{C} S W$ receive allocations only at quantum boundaries; (ii) under $\mathcal{I}$, each task receives an allocation equal to its weight, whereas under $\mathcal{S} W$ and $\mathcal{C} S W$, each task receives allocations according to its scheduling weight; and (iii) the total allocation each task receives in $\mathcal{S} W$ and $\mathcal{C} S W$ is calculated based on the releases and completion times of its active subtasks, whereas allocations in $\mathcal{I}$ are independent of subtask releases and completion times. Hence, even if all active subtasks of a given task are halted, $\mathcal{I}$ still allocates capacity to that task.

Example (Figure 5.24). Consider the example in Figure 5.24, which depicts the allocations in the schedules $\mathcal{C} S W$ and $\mathcal{I}$ (insets (a) and (b), respectively) to a task $T_{1}$ that has an initial weight of $3 / 19$ that increases to $2 / 5$ (via Rule N) at time 8 . Notice that, over the range $[8,10$ ) in $\mathcal{I}, T_{1}$ receives an allocation equal to its weight at every instant (for a total allocation of $4 / 5$ over $[8,10))$. Compare this to $\mathcal{C} S W$, in which $T_{1}$ receives only an allocation of $44 / 95$ over


Figure 5.24: Allocations for a task $T_{1}$ with an initial weight of $3 / 19$ that changes to $2 / 5$. (a) The value of $\mathrm{A}\left(\mathcal{C} S W, T_{1}^{[j]}, u\right)$ for each slot and subtask. (b) The allocations to $T_{1}$ in $\mathcal{I}$ at each instant.
the same range.
As was the case for the adaptable sporadic task model and the modified adaptable sporadic task model, the drift of a task $T_{i}$ is calculated as the difference between $T_{i}$ 's allocations in the $\mathcal{C} S W$ and the $\mathcal{I}$ schedules. Formally, under PD-PNH, the drift of a task $T_{i}$ is defined as

$$
\begin{equation*}
\operatorname{drift}\left(T_{i}, t\right)=\mathrm{A}\left(\mathcal{I}, T_{i}, 0, t\right)-\mathrm{A}\left(\mathcal{C} S W, T_{i}, 0, t\right) \tag{5.37}
\end{equation*}
$$

Example (Figure 5.8). For example, in Figure 5.8(b), the drift of task $T_{1}$ at time $t=9$ is $\mathrm{A}\left(\mathcal{I}, T_{1}, 0,9\right)-\mathrm{A}\left(\mathcal{C} S W, T_{1}, 0,9\right)=27 / 20-20 / 20=7 / 20$, whereas at time $t=10$, the drift of $T_{i}$ is $\mathrm{A}\left(\mathcal{I}, T_{1}, 0,10\right)-\mathrm{A}\left(\mathcal{C S} W, T_{1}, 0,10\right)=3 / 2-1=1 / 2$. Notice that, since $T_{1}^{[2]}$ is halted at time $10, \mathrm{~A}\left(\mathcal{C S W}, T_{1}^{[2]}, 0,10\right)=0$.

We say that a reweighting algorithm is fine-grained iff there exists some constant value $c$ such that the drift per weight change is less than $c$. We say that a reweighting algorithm is coarse-grained otherwise.

### 5.6.1 PD-LJ is Not Fine-Grained

We now prove that PD-LJ is not fine-grained. (The definition of drift for PD-LJ is the same as the definition of drift for PD-PNH, except that $\mathcal{C} S W$ is determined by using PD-LJ.)


Figure 5.25: A four-processor example illustrating why PD-LJ is coarse-grained. $A$ is a set of 35 tasks each with a weight $1 / 10 . T_{1}$ has a weight of $1 / 10$ that increases to $1 / 2$ at time 4 .

Example (Figure 5.25). Consider the four-processor example in Figure 5.25, which depicts the PD-LJ schedule for a system that consists of a set $A$ of 35 tasks with weight $1 / 10$ and a task $T_{1}$ with weight $1 / 10$ that increases to $1 / 2$ at time 4 . By Rule $\mathrm{L}, T_{1}$ cannot "leave" until time 10. Hence, the change is not enacted until time 10 . Thus, over the range $[4,10)$, $T_{i}$ receives a $1 / 10$ per-slot allocation in $\mathcal{C} S W$ and $1 / 2$ in $\mathcal{I}$. Hence, $T_{i}$ 's drift reaches a value of $24 / 10$ at time 10 . This example can be generalized by decreasing the weights of the tasks in set $A$ and the initial weight of $T_{1}$ to $\frac{1}{10 c}$ and increasing the number of tasks in $A$ to $35 c$, where $c$ is a positive integer, in which case $\operatorname{drift}\left(T_{1}, \mathrm{~d}\left(T_{1}^{[1]}\right)\right)=5 c-3+\frac{2}{5 c}$.

From the generalization of Figure 5.25, the theorem below follows.

Theorem 5.3. PD-LJ is not fine-grained.

### 5.6.2 All EPDF Scheduling Algorithms Incur Drift

Next, we show that any EPDF scheduling algorithm incurs some drift.

Example (Figure 5.26). Consider the example in Figure 5.26, which depicts a twoprocessor system that consists of a set $A$ of 10 tasks with weight $1 / 7$ that leave at time 7, a set $B$ of two tasks with weight $1 / 6$ that leave at time 6 , a set $C$ of two tasks with weight $1 / 14$ that join at time 6 , and a set $D$ of five tasks with a weight of $1 / 21$ that increases to $1 / 3$ at time 7. The projected deadlines of tasks in $D$ based on their IDEAL allocation are labeled above the schedule, $\mathcal{I}$. With subtask deadlines defined by $\mathcal{I}$, the deadline for each task in set $D$ changes at time 7 from 21 to 9 . The tasks in $D$ have an original deadline of 21


Figure 5.26: A two-processor system illustrating that all EPDF algorithms incur drift.
because that is the projected time at which their $\mathcal{I}$ allocations will equal one if their weights do not change. These tasks change their deadlines to 9 at time 7 because the new weight, $1 / 3$, changes the projected time by which their allocations in $\mathcal{I}$ will equal one to time 9 . Note that any EPDF algorithm will not schedule the tasks in $D$ until time 7 . As a result, a deadline is missed. Notice also that any EPDF algorithm would need to use projections for determining subtask deadlines if we assume no prior knowledge of weight changes. To prevent a deadline miss, the lag-bound range must be shifted, thus incurring drift.

From Figure 5.26, the theorem below follows.

Theorem 5.4. All EPDF algorithms can incur non-zero drift per reweighting event.

### 5.6.3 PD-PNH is fine-grained

Finally, we show that PD-PNH is fine-grained. We first show that the rules for reweighting non-heavy-changeable tasks are fine-grained, and then we show that the rules for reweighting heavy-changeable tasks are fine grained.

Non-heavy-changeable tasks. By the definition of drift, in order to prove that PD-PNH is fine-grained, for non-heavy-changeable tasks, we merely need to consider wether a subtask has been halted and the window placement of a task after it is reweighted. Suppose that a non-heavy-changeable task $T_{i}$ initiates a weight change at $t_{c}$. Let $t_{e}$ be the next time at which $T_{i}$ enacts a change at or after $t_{c}$, assume that $T_{i}$ releases a subtask at or before $t_{c}$, and let
$T_{i}^{[j]}$ be the last-released such subtask. (If $T_{i}$ does not release a subtask at or before $t_{c}$, then $T_{i}$ 's drift does not change when reweighted at $t_{c}$.)

We now show that if the change initiated at $t_{c}$ is enacted at $t_{e}$, then absolute drift increases by at most two. We begin by observing that there are three sources of drift: drift incurred because $T_{i}^{[j]}$ is halted; drift incurred because $t_{c}<t_{e}$; and drift that is incurred because $T_{i}^{[j+1]}$ is not released at time $t_{e}$. (Throughout this discussion we assume that a reweighting event is not canceled, since canceling a reweighting event would decrease the amount of drift incurred.)

Example (Figure 5.27). Consider the four-processor example in Figure 5.27, which consists of a set $C$ of 19 tasks, each with a weight of $3 / 20$, and a task $T_{1}$ that increases its weight from $3 / 20$ to $1 / 2$ via Rule P (tie-breaks not resolved by $\mathrm{PD}^{2}$ go against task $T_{1}$ ). Inset (a) depicts the PD-PNH schedule. Inset (b) depicts the CSW schedule. Inset (c) depicts the IDEAL schedule. Notice that $T_{1}$ incurs $1 / 2$ of a quantum of drift over the range $\left[6,10\right.$ ) because $T_{1}^{[2]}$ receives no allocation in the CSW schedule.

Example (Figure 5.28). Consider the four-processor example in Figure 5.28, which consists of a set $C$ of 19 tasks each with a weight of $3 / 20$ and a task $T_{1}$ that increases its weight from $3 / 20$ to $1 / 2$ via Rule N (tie-breaks not resolved by $\mathrm{PD}^{2}$ go against tasks in $C$ ). Inset (a) depicts the PD-PNH schedule. Inset (b) depicts the CSW schedule. Inset (c) depicts the IDEAL schedule. Notice that $T_{1}$ incurs $1 / 2$ of a quantum of drift over the range $[11,12$ ) because $T_{1}^{[3]}$, release is delayed.

Example (Figure 5.29). Consider the four-processor example in Figure 5.29, which consists of a set $C$ of 19 tasks each with a weight of $3 / 20$ and a task $T_{1}$ that decreases its weight from $2 / 5$ to $3 / 20$ via Rule N at time 1. Inset (a) depicts the PD-PNH schedule. Inset (b) depicts the CSW schedule. Inset (c) depicts the IDEAL schedule. Notice that $T_{1}$ incurs $-3 / 20$ of a quantum of drift over the range $[1,4)$ because the enactment of the weight change occurs after the change is initiated.

Positive-changeable. We now consider the amount of drift that is incurred based on whether $T_{i}$ is positive- or negative-changeable at $t_{c}$. If $T_{i}$ is positive-changeable

| X | Allocation to a task (or a task set) | $\downarrow$ Subtask released | $\wedge$ Subtask deadline |
| :---: | :---: | :---: | :---: |
|  | eighting event enacted --- Subt | out without reweight | Reweighting even |


(a)

(b)

(c)

Figure 5.27: A four-processor illustration of Rule P under $\mathrm{PD}^{2} . C$ is a set of 19 tasks with weight of $3 / 20 . T_{1}$ increases its weight from $3 / 20$ to $1 / 2$ at time 10 via rule P. (Tie-breaks not resolved by PD ${ }^{2}$ go against $T_{1}$.) (a) The PD-PNH schedule. (b) The CSW schedule. (c) The IDEAL schedule.


Figure 5.28: A four-processor illustration of Rule N under $\mathrm{PD}^{2}$. $C$ is a set of 19 tasks with weight of $3 / 20$. (Tie-breaks not resolved by $\mathrm{PD}^{2}$ go against tasks in $C$.) $T_{1}$ increases its weight from $3 / 20$ to $1 / 2$ at time 10 via rule N. (a) The PD-PNH schedule. (b) The CSW schedule. (c) The IDEAL schedule.


Figure 5.29: A four-processor illustration of Rule N under $\mathrm{PD}^{2}$. $C$ is a set of 19 tasks with weight of $3 / 20 . T_{1}$ decreases its weight from $2 / 5$ to $3 / 20$ at time 1 via rule N. (a) The PD-PNH schedule. (b) The CSW schedule. (c) The IDEAL schedule.
at $t_{c}$, then it changes its weight by Rule P. Thus, $T_{i}^{[j]}$ is halted and at time $\max \left(t_{c}, \min \left(\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j-1]}\right), \mathrm{d}\left(T_{i}^{[j-1]}\right)\right)+\mathrm{b}\left(T_{i}^{[j-1]}\right)\right)$, the change is enacted and $T_{i}^{[j+1]}$ is released. Thus, if $T_{i}$ is positive-changeable, then it may incur drift because $T_{u}^{[j]}$ is halted and because the change is not immediately enacted.

Since it is trivial to show that $T_{i}$ incurs at most one quantum of drift because $T_{u}^{[j]}$ is halted, we focus on showing that $T_{i}$ incurs at most one quantum of absolute drift because the change is not enacted immediately. Notice that, if $\max \left(t_{c}, \min \left(\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j-1]}\right), \mathrm{d}\left(T_{i}^{[j-1]}\right)\right)+\mathrm{b}\left(T_{i}^{[j-1]}\right)\right) \leq$ $t_{c}$, then the change is immediately enacted and the only drift $T_{i}$ incurs is a result of $T_{i}^{[j]}$ halting. Thus, we assume that $t_{c}<\max \left(t_{c}, \min \left(\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j-1]}\right), \mathrm{d}\left(T_{i}^{[j-1]}\right)\right)+\mathrm{b}\left(T_{i}^{[j-1]}\right)\right)$.

Since we have assumed that $T_{i}^{[j]}$ is the last-released subtask of $T_{i}$ at $t_{c}$, it follows that $\mathrm{r}\left(T_{i}^{[j]}\right) \leq t_{c}$. Thus, since we are assuming that $t_{c}<\max \left(t_{c}, \min \left(\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j-1]}\right), \mathrm{d}\left(T_{i}^{[j-1]}\right)\right)+\right.$ $\left.\mathrm{b}\left(T_{i}^{[j-1]}\right)\right)$, we have

$$
\mathrm{r}\left(T_{i}^{[j]}\right) \leq t_{c}<\min \left(\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j-1]}\right), \mathrm{d}\left(T_{i}^{[j-1]}\right)\right)+\mathrm{b}\left(T_{i}^{[j-1]}\right) \leq \mathrm{d}\left(T_{i}^{[j-1]}\right)+\mathrm{b}\left(T_{i}^{[j-1]}\right)
$$

We now show that the range $\left[t_{c}, \min \left(\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j-1]}\right), \mathrm{d}\left(T_{i}^{[j-1]}\right)\right)+\mathrm{b}\left(T_{i}^{[j-1]}\right)\right)$ is at most two. Notice that, in order for the range $\left[t_{c}, \min \left(\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j-1]}\right), \mathrm{d}\left(T_{i}^{[j-1]}\right)\right)+\mathrm{b}\left(T_{i}^{[j-1]}\right)\right)$ to be greater than two, $\mathrm{r}\left(T_{i}^{[j]}\right)<\mathrm{d}\left(T_{i}^{[j-1]}\right)-\mathrm{b}\left(T_{i}^{[j-1]}\right)$ must hold. However, by Property (V), if $\mathrm{r}\left(T_{i}^{[j]}\right)<\mathrm{d}\left(T_{i}^{[j-1]}\right)-\mathrm{b}\left(T_{i}^{[j-1]}\right)$, then $\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j-1]}\right) \leq \mathrm{r}\left(T_{i}^{[j]}\right)$. This, in turn, implies that the range $\left[t_{c}, \min \left(\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j-1]}\right), \mathrm{d}\left(T_{i}^{[j-1]}\right)\right)+\mathrm{b}\left(T_{i}^{[j-1]}\right)\right)$ is at most one. Thus, the range $\left[t_{c}, \min \left(\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j-1]}\right), \mathrm{d}\left(T_{i}^{[j-1]}\right)\right)+\mathrm{b}\left(T_{i}^{[j-1]}\right)\right)$ is at most two. This implies that

$$
\begin{equation*}
\max \left(t_{c}, \min \left(\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j-1]}\right), \mathrm{d}\left(T_{i}^{[j-1]}\right)\right)+\mathrm{b}\left(T_{i}^{[j-1]}\right)\right)-t_{c}=t_{e}-t_{c} \leq 2 \tag{5.38}
\end{equation*}
$$

Since the maximal weight for any non-heavy changeable task is less than $1 / 2$, and $t_{e}-t_{c} \leq$ 2 , it follows that $\mathrm{A}\left(\mathcal{C} S W, T_{i}, t_{c}, t_{e}\right)<1$ and $\mathrm{A}\left(\mathcal{I}, T_{i}, t_{c}, t_{e}\right)<1$. Thus, since the absolute drift incurred over the range $\left[t_{c}, t_{e}\right)$ is given as $\left|\mathrm{A}\left(\mathcal{I}, T_{i}, t_{c}, t_{e}\right)-\mathrm{A}\left(\mathcal{C} S W, T_{i}, t_{c}, t_{e}\right)\right|$, it follows that $T_{i}$ incurs up to one additional quantum of absolute drift waiting for the change initiated at $t_{c}$ to be enacted. Thus, combined with the additional quantum of drift incurred by halting $T_{i}^{[j]}$, if $T_{i}$ is positive-changeable at $t_{c}$, then it incurs at most two quanta of absolute drift.

Negative-changeable. If $T_{i}$ is negative-changeable at $t_{c}$, then it changes its weight by Rule N . Therefore, $T_{i}^{[j]}$ does not halt and, as a result, $T_{i}^{[j]}$ receives the same allocation in $\mathcal{S} W$ and $\mathcal{C} S W$. Thus,

$$
\begin{equation*}
\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)=\mathrm{C}\left(\mathcal{C} S W, T_{i}^{[j]}\right) \tag{5.39}
\end{equation*}
$$

If $T_{i}$ increases its weight at $t_{c}$, then the change is immediately enacted and $T_{i}^{[j+1]}$ is released at time $\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)+\mathrm{b}\left(T_{i}^{[j]}\right)$. Thus, by (5.39), $\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)+\mathrm{b}\left(T_{i}^{[j]}\right)=\mathrm{C}\left(\mathcal{C} S W, T_{i}^{[j]}\right)+\mathrm{b}\left(T_{i}^{[j]}\right)$ Since the weight change is immediately enacted, $T_{i}$ receives the same allocation in $\mathcal{C} S W$ and $\mathcal{I}$ in every time slot in the range $\left[t_{c}, \mathrm{C}\left(\mathcal{C S} W, T_{i}^{[j]}\right)-1\right.$ ). (Specifically, if $T_{i}$ increases its weight to Nw at time $t_{c}$, then in every slot in the range $\left[t_{c}, \mathrm{C}\left(\mathcal{C} S W, T_{i}^{[j]}\right)-1\right), T_{i}$ is allocated Nw in both $\mathcal{C S W}$ and $\mathcal{I}$.) Thus, since Rule N does not halt $T_{i}^{[j]}$, the only source of drift that can be incurred is over the range $\left[\mathrm{C}\left(\mathcal{C} S W, T_{i}^{[j]}\right)-1, \mathrm{r}\left(T_{i}^{[j+1]}\right)\right)$. Since $\mathrm{r}\left(T_{i}^{[j+1]}\right)=$ $\mathrm{C}\left(\mathcal{C} S W, T_{i}^{[j]}\right)+\mathrm{b}\left(T_{i}^{[j]}\right)$, the length of this interval is at most two. Since the weight of a non-heavy-changeable task is less than $1 / 2$, the increase in drift is at most $2 \cdot 1 / 2$.

For example, in Figure 5.9(a), $\mathrm{A}\left(\mathcal{C} S W, T_{1}, 10,12\right) \quad=\mathrm{A}\left(\mathcal{C} S W, T_{1}, 0,12\right)-$ $\mathrm{A}\left(\mathcal{C S W}, T_{1}, 0,10\right)=4 / 2-3 / 2=1 / 2<1=5 / 2-3 / 2=\mathrm{A}\left(\mathcal{I}, T_{1}, 0,12\right)-\mathrm{A}\left(\mathcal{I}, T_{1}, 0,10\right)=$ $\mathrm{A}\left(\mathcal{I}, T_{1}, 10,12\right)$.

If $T_{i}$ decreases its weight at $t_{c}$ via Rule N , then $T_{i}^{[j+1]}$ is released at time $\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)+$ $\mathrm{b}\left(T_{i}^{[j]}\right)$. Thus, by (5.39), $T_{i}^{[j+1]}$ is released at time $\mathrm{C}\left(\mathcal{C} S W, T_{i}^{[j]}\right)+\mathrm{b}\left(T_{i}^{[j]}\right)$. Since $T_{i}$ decreases its weight, over the range $\left[t_{c}, \mathcal{C}\left(\mathcal{C} S W, T_{i}^{[j]}\right)\right), T_{i}$ is allocated at most one quantum more in $\mathcal{C} S W$ than in $\mathcal{I}$. Furthermore, over the range $\left[\mathrm{C}\left(\mathcal{C} S W, T_{i}^{[j]}\right), \mathrm{C}\left(\mathcal{C} S W, T_{i}^{[j]}\right)+\mathrm{b}\left(T_{i}^{[j]}\right)\right), T_{i}$ is allocated less than $1 / 2$ quanta more in $\mathcal{I}$ than in $\mathcal{C} S W$, since the length of this range is at most one and the weight of a non-heavy changeable task is less than $1 / 2$. Thus, the maximal possible decrease in drift is one and the maximal possible increase in drift is $1 / 2$. For example, in Figure 5.9(b), the drift incurred by changing the weight of $T_{1}$ from $2 / 5$ to $3 / 20$ is $-3 / 20$, i.e., $\mathrm{A}\left(\mathcal{C} S W, T_{1}, 1,4\right)=\mathrm{A}\left(\mathcal{C} S W, T_{1}, 0,4\right)-\mathrm{A}\left(\mathcal{C} S W, T_{1}, 0,1\right)=5 / 5-2 / 5=3 / 5>9 / 20=$ $17 / 20-2 / 5=\mathrm{A}\left(\mathcal{I}, T_{1}, 0,4\right)-\mathrm{A}\left(\mathcal{I}, T_{1}, 0,1\right)=\mathrm{A}\left(\mathcal{I}, T_{1}, 1,4\right)$.

Heavy-changeable tasks. We now show that the drift of a heavy-changeable task is at most five. By the definition of drift, in order to prove that PD-PNH is fine-grained, for heavy-
changeable tasks, we only need to consider the window placement of a task $T_{i}$, from the reweighting initiation, at time $t_{c}$, until $\mathrm{r}\left(T_{i}^{[k]}\right)$, where $T_{i}^{[k]}$ is the first subtask of $T_{i}$ released at or after $\mathrm{D}\left(T_{i}^{[j]}\right)$ and $T_{i}^{[j]}$ is the last-released subtask of $T_{i}$ at or before time $t_{c}$. (If $T_{i}$ does not release a subtask at or before $t_{c}$ or $\mathrm{D}\left(T_{i}^{[j]}\right) \leq t_{c}$, then $T_{i}$ is not heavy-changeable.) We bound the drift by showing that the maximal absolute difference between $\mathcal{C} S W$ and $\mathcal{I}$ in allocating to $T_{i}$ over the interval $\left[t_{c}, \mathrm{r}\left(T_{i}^{[j+1]}\right)\right)$ is at most four and that one additional quantum of drift is incurred in the slot $\mathrm{D}\left(T_{i}^{[j]}\right)-1$. (Notice that, even if the reweighting event initiated at $t_{c}$ were canceled, then the absolute drift per reweighting event would still be at most five.)

If $\mathrm{d}\left(T_{i}^{[j]}\right) \leq t_{c}$, then, by Part 2 of Rule H , the weight change is enacted within one quantum. Moreover, by Part 3 of Rule $\mathrm{H}, T_{i}^{[j+1]}$ is released when the change is enacted. Thus, the range $\left[t_{c}, \mathrm{r}\left(T_{i}^{[j+1]}\right)\right)$ is at most one quantum long. Since the weight of a heavy-changeable task is less than one, the maximal increase in the absolute value of drift in such a case is less than one over the time range $\left[t_{c}, \mathrm{r}\left(T_{i}^{[j+1]}\right)\right)$. If $\mathrm{d}\left(T_{i}^{[j]}\right)>t_{c}$, then, again by Part 2 of Rule H , the weight change is enacted within four quanta since the maximal window length of any heavychangeable task is three (and the enactment may be delayed by an additional quantum if the b-bit is one). Moreover, by Part 3 of Rule $\mathrm{H}, T_{i}^{[j+1]}$ is released when the change is enacted. Thus, the range $\left[t_{c}, \mathrm{r}\left(T_{i}^{[j+1]}\right)\right)$ is at most four quantum long. As a result, since the weight of a heavy-changeable subtask is less than one, the maximal increase in the absolute drift in such a case is less than four over the time range $\left[t_{c}, \mathrm{r}\left(T_{i}^{[j+1]}\right)\right)$. For example, in Figure $5.11(\mathrm{~b}), T_{2}$ receives an allocation $18 / 10$ in the schedule $\mathcal{I}$ over the range $[2,4$ ), whereas it receives only $2 / 9$ over this same range in the schedule $\mathcal{C} S W$. Thus, $T_{2}$ incurs $18 / 10-2 / 9 \approx 1.58$ units of drift over this region.

Notice that, over the range $\left[\mathrm{r}\left(T_{i}^{[j+1]}\right), \mathrm{r}\left(T_{i}^{[k]}\right)\right)$, with one possible exception in the slot $\mathrm{D}\left(T_{i}^{[j]}\right)-1$, the allocation to $T_{i}$ in $\mathcal{I}$ and $\mathcal{C} S W$ are the same (assuming no additional reweighting events are initiated), despite the fact that the window length of each subtask of $T_{i}$ released over this time range is two (where $T_{i}^{[k]}$ is, as above, the first subtask of $T_{i}$ released at or after $\left.\mathrm{D}\left(T_{i}^{[j]}\right)\right)$. The reason for this behavior is two-fold. First, the deadline of a subtask is not used in the definition of $\mathcal{C} S W$, which is based on the pseudo-code in Figure 5.7. Hence, $T_{i}$ 's allocation in the CSW schedule is not affected by the fact that its subtasks may have had
their window lengths "artificially" shrunk to two by Rule H. For example, in Figure 5.11, $T_{2}$ receives an allocation of $36 / 10$ in both the CSW and IDEAL schedules over the range $[4,8)$. Second, by Rule H, $T_{i}$ releases subtasks with the same frequency as a "normal" task. For example, in Figure 5.12 , over the range $[4,10), T_{2}$ releases one subtask every three quantum, which is exactly the same frequency at which "normal" tasks with a weight of $1 / 3$ release subtasks.

As we mentioned in the previous paragraph, it is possible that in the slot $\mathrm{D}\left(T_{i}^{[j]}\right)-1, T_{i}$ 's allocation may differ in $\mathcal{I}$ and $\mathcal{C} S W$. The reason for this behavior is that, by Part 4 of Rule H , no subtask of $T_{i}$ is released at time $\mathrm{D}\left(T_{i}^{[j]}\right)-1$. Thus, if the definition of release given in Part 3 of Rule $H$ specifies that a job should be released in slot $D\left(T_{i}^{[j]}\right)-1$, then drift is incurred since that subtask will not be released. Since the maximal weight of a heavy-changeable task is less than 1 , in the slot $\mathrm{D}\left(T_{i}^{[j]}\right)-1$, less than one additional unit of drift may be incurred.

Also, notice that after $\mathrm{r}\left(T_{i}^{[k]}\right), T_{i}$ behaves as a normal task with its new weight. Thus, the maximal absolute drift that can be incurred by a heavy-changeable task changing its weight is five. For example, in Figure 5.11 (b), $T_{2}$ receives an allocation of $9 / 10$ in the schedule $\mathcal{I}$ over the range $[8,9)$, whereas it receives only $4 / 10$ over this same range in the schedule $\mathcal{C} S W$. Thus, $T_{2}$ incurs $9 / 10-4 / 10=1 / 2$ units of drift over this interval. Thus, the total drift incurred as a result of this change is $\left(\mathrm{A}\left(\mathcal{I}, T_{2}, 2,4\right)-\mathrm{A}\left(\mathcal{C} S W, T_{2}, 2,4\right)\right)+\left(\mathrm{A}\left(\mathcal{I}, T_{2}, 4,8\right)-\right.$ $\left.\mathrm{A}\left(\mathcal{C} S W, T_{2}, 4,8\right)\right)+\left(\mathrm{A}\left(\mathcal{I}, T_{2}, 8,9\right)-\mathrm{A}\left(\mathcal{C} S W, T_{2}, 8,9\right)\right)=(18 / 10-2 / 9)+0+(9 / 10-4 / 10) \approx$ 2.08.

One final note: if $T_{i}$ initiates an additional reweighting event over the range $\left[\mathrm{r}\left(T_{i}^{[j+1]}\right), \mathrm{D}\left(T_{i}^{[j]}\right)-1\right.$ ), then $T_{i}$ may experience three additional quanta of drift (as opposed to four quanta). The reason why it experiences an additional three quanta of drift is because the window length of a subtask in the range $\left[\mathrm{r}\left(T_{i}^{[j+1]}\right), \mathrm{D}\left(T_{i}^{[j]}\right)-1\right)$ is at most 2. Thus, this enactment may be delayed for up to three quantum. Since such a change would be an additional reweighting event, this does not impact the per-reweighting-event measurement of drift.

Theorem 5.5. PD-PNH is fine-grained; moreover, the absolute value of the per-event drift under PD-PNH is at most two for non-heavy-changeable tasks and at most five for heavy-
changeable tasks.

### 5.7 Lost Utilization

As we mentioned earlier, when a task decreases its weight, there is a delay between when the weight change is initiated and when the capacity gained by this decrease is freed. When a non-heavy-changeable initiates a weight decrease, the capacity is freed when the change is enacted. For example, in Figure 5.9(b), only at or after time 4 could another task could use the capacity of $2 / 5-3 / 20$ gained by $T_{2}$ decreasing its weight. For heavy-changeable tasks, the capacity is freed at the group deadline of the last-released subtask. For example, in the system depicted in Figure 5.12, $T_{1}$ cannot use the additional capacity gained by $T_{2}$ 's weight decrease until time 9 , which is the group deadline of $T_{1}^{[2]}$.

This delay is tantamount to "idling" a faction of the system, which could cause some utilization to be lost with respect to a system in which weight changes could be enacted instantly. The lost utilization caused by $T_{i}$ initiating a weight decrease at time $t_{c}$ is defined as

$$
\begin{equation*}
\int_{t_{c}}^{t_{f}} \mathrm{LC}\left(T_{i}, u\right) d u \tag{5.40}
\end{equation*}
$$

where $\mathrm{LC}\left(T_{i}, t\right)$ is the capacity that has not been freed by $T_{i}$ at time $t$, and $t_{f}$ is the earliest time $t$ such that $T_{i}$ frees its capacity, $\mathrm{wt}\left(T_{i}, t\right) \geq \mathrm{Ow}$, or $T_{i}$ initiates another weight decrease.

Example (Figures 5.13 and 5.29). Notice that in Figure 5.29, over the range [1, 4), the lost utilization is $2 / 5-3 / 20$ at each instant. Thus, the total lost utilization as the result of this weight change is. $\int_{1}^{4}(2 / 5-3 / 20) d t=3 \cdot 5 / 20=15 / 20$. Also notice that, in Figure 5.13, $T_{2}$ incurs $13 / 14-1 / 3$ of lost utilization at each instant over the range $[2,9$ ), and $13 / 14-3 / 4$ over the range $[9,14)$. Thus, the total lost utilization incurred by $T_{2}$ decreasing its weight at time 2 is $\int_{2}^{14}\left(13 / 14-\operatorname{wt}\left(T_{2}, t\right)\right) d t=\int_{2}^{9}(13 / 14-1 / 3) d t+\int_{9}^{14}(13 / 14-3 / 4) d t \approx 5.059$. Notice that the weight increase at time 9 does not stop $T_{2}$ from accruing lost utilization because the new weight after the change (i.e., 3/4) is less than the original scheduling weight (i.e., 13/14).

Non-heavy-changeable tasks. Notice that, if a task is not heavy-changeable when it initiates a weight change, then the time the capacity is freed is the time the weight change is enacted. Thus, by (5.40), if $T_{i}$ is non-heavy-changeable when it initiates a weight decrease to Nw at time $t_{c}$, then the lost utilization caused by this weight change equals

$$
\begin{equation*}
\int_{t_{c}}^{t_{e}} \mathrm{LC}\left(T_{i}, u\right) d u \tag{5.41}
\end{equation*}
$$

where $t_{e}$ is the time that the weight change is enacted or canceled. Notice that, if the change initiated at $t_{c}$ is canceled at $t_{e}$, then this implies that either $\mathrm{wt}\left(T_{i}, t_{e}\right) \geq \mathrm{Ow}$ holds or $T_{i}$ initiated another weight decrease at $t_{e}$. Thus, for non-heavy-changeable tasks (5.41) is equivalent to (5.40).

Let $T_{i}^{[j]}$ denote the last-released subtask (if any) of $T_{i}$ at $t_{c}$. If $T_{i}^{[j]}$ does not exist, then the change is enacted immediately, and the lost utilization is zero. If $T_{i}^{[j]}$ exists but $\mathrm{d}\left(T_{i}^{[j]}\right) \leq t_{c}$, then the change is enacted within one quantum. Moreover, since the weight of a non-heavychangeable task is less than $1 / 2, \mathrm{LC}\left(T_{i}, t\right)<1 / 2$ for any time $t$. Thus, by (5.41), the maximal amount of lost utilization is 1 .

Positive-changeable. If $T_{i}$ is positive-changeable at $t_{c}$, then the change is initiated via Rule P. Thus, the change is enacted at $\max \left(t_{c}, \min \left(\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j-1]}\right), \mathrm{d}\left(T_{i}^{[j-1]}\right)\right)+\mathrm{b}\left(T_{i}^{[j-1]}\right)\right)$. As we proved in (5.38), $\max \left(t_{c}, \min \left(\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j-1]}\right), \mathrm{d}\left(T_{i}^{[j-1]}\right)\right)+\mathrm{b}\left(T_{i}^{[j-1]}\right)\right)-t_{c} \leq 2$. Thus, since the weight of a non-heavy-changeable task is less than $1 / 2$, by (5.41), the lost utilization $T_{i}$ is negative-changeable at $t_{c}$ is less than one.

Negative-changeable. If $T_{i}$ is negative-changeable at $t_{c}$, then by Rule N , the weight change is enacted or canceled by $\mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)+\mathrm{b}\left(T_{i}^{[j]}\right)$, i.e., $t_{e} \leq \mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)+\mathrm{b}\left(T_{i}^{[j]}\right)$. By the definition of completed, $\int_{\mathbf{r}\left(T_{i}^{[j]}\right)}^{\mathcal{C}\left(\mathcal{S} W T_{i}^{[j]}\right)} \operatorname{Swt}\left(T_{i}, u\right) d u=1$. Since $T_{i}^{[j]}$ is the last-released subtask of $T_{i}$ at $t_{c}$, $\mathrm{r}\left(T_{i}^{[j]}\right) \leq t_{c}$. Thus, since $t_{e}-\mathrm{b}\left(T_{i}^{[j]}\right) \leq \mathrm{C}\left(\mathcal{S} W, T_{i}^{[j]}\right)$, Swt $\left(T_{i}, t\right)<1 / 2$ for every $t \in\left[t_{c}, t_{e}\right)$, and by assumption $t_{c} \leq t_{e}$,

$$
\begin{equation*}
\int_{t_{c}}^{t_{e}} \operatorname{Swt}\left(T_{i}, u\right) d u=\int_{t_{c}}^{t_{e}-\mathbf{b}\left(T_{i}^{[j]}\right)} \operatorname{Swt}\left(T_{i}, u\right) d u+\int_{t_{e}-\mathbf{b}\left(T_{i}^{[j]}\right)}^{t_{e}} \operatorname{Swt}\left(T_{i}, u\right) d u<1+1 / 2 \tag{5.42}
\end{equation*}
$$

Since $T_{i}$ decreases its weight at $t_{c}$, it follows that for every $t \in\left[t_{c}, t_{e}\right), \operatorname{LC}\left(T_{i}, t\right) \leq$ $\operatorname{Swt}\left(T_{i}, t\right)$. Thus, by (5.42), $\int_{t_{c}}^{t_{e}} \mathrm{LC}\left(T_{i}, u\right) d u<3 / 2$. Thus, by (5.41), if $T_{i}$ is negativechangeable at $t_{c}$, then the lost utilization is less than $3 / 2$. Notice that, for non-heavychangeable tasks, the lost utilization is closely related to the drift incurred.

Heavy-changeable tasks. For a heavy-changeable task, the amount of lost utilization could be substantially larger than the drift incurred. Let $T_{i}^{[j]}$ denote the last-released subtask of a heavy-changeable task $T_{i}$ at $t_{c}$. (Notice that, $T_{i}^{[j]}$ must exist, because by the definition of heavy-changeable, a $T_{i}$ must have released a subtask at or before $t_{c}$ that had a group deadline after $t_{c}$.) Notice that, by Part 1 of Rule H , if $T_{i}$ decreases its weight from Ow to Nw at $t_{c}$, then the capacity is freed at time $\mathrm{D}\left(T_{i}^{[j]}\right)$. Thus, (5.40), i.e., the lost utilization incurred by this change, is upper-bounded by

$$
\begin{equation*}
\int_{t_{c}}^{\mathrm{D}\left(T_{i}^{[j]}\right)}(\mathrm{Ow}-\mathrm{Nw}) d t \tag{5.43}
\end{equation*}
$$

Before continuing, notice that, a heavy task must have at least one group deadline every period. As a result, $\mathrm{D}\left(T_{i}^{[j]}\right)-t_{c} \leq \mathrm{p}\left(T_{i}\right)$. Moreover, recall that the weight of $T_{i}$ is given by $\mathrm{wt}\left(T_{i}\right)=\mathrm{e}\left(T_{i}\right) / \mathrm{p}\left(T_{i}\right)$. Thus, (5.43) can be upper-bounded by

$$
(\mathrm{Ow}-\mathrm{Nw}) \cdot\left(\mathrm{D}\left(T_{i}^{[j]}\right)-t_{c}\right) \leq \mathrm{Ow} \cdot \mathrm{p}\left(T_{i}\right) \leq \mathrm{e}_{\max }\left(T_{i}\right)
$$

The above bound still holds even if $T_{i}$ initiates additional changes over the range $\left[t_{c}, \mathrm{D}\left(T_{i}^{[j]}\right)\right)$.
For example, in the system depicted in Figure 5.12, the amount of lost utilization caused by $T_{2}$ decreasing its weight from $8 / 9$ is $(\mathrm{Ow}-\mathrm{Nw}) \cdot\left(\mathrm{D}\left(T_{2}^{[2]}\right)-t_{c}\right)=(8 / 9-1 / 3) \cdot(9-2)=$ $35 / 9 \approx 3.88$. Since under Pfair scheduling, all execution times and periods are an integral number of quanta, the only way $T_{2}$ could have a weight of $8 / 9$ is if its execution time is at least 8 . Thus, the amount of lost utilization caused by this decrease is at least $\approx 4.12$ less than the execution time for $T_{2}$.

Theorem 5.6. The lost utilization per reweighting event initiation under under PD-PNH is at most $3 / 2$ for non-heavy-changeable tasks and at most $\mathrm{e}_{\max }\left(T_{i}\right)$ for heavy-changeable tasks.

### 5.8 Conclusion

In this chapter, we presented the AIS task model as well as the rules for reweighting a task under the $\mathrm{PD}^{2}$ scheduling algorithm. In addition, by using techniques borrowed from (Srinivasan and Anderson, 2005), we proved that no subtask misses its deadline using our reweighting rules. Also, we proved that the absolute drift that can be incurred per reweighting event is at most five. Finally, we proved that the amount of lost utilization caused by a weight decrease is at most $3 / 2$ for non-heavy changeable tasks and at most $\mathrm{e}_{\max }\left(T_{i}\right)$ for heavy-changeable tasks.

## CHAPTER 6

## AGEDF*

In this section, we present the adaptable GEDF (AGEDF) scheduling algorithm, which extends the GEDF algorithm by using feedback techniques in order to determine when a task should enact a weight change.

### 6.1 Adaptable Service Level Tasks

Before discussing the AGEDF scheduling algorithm, we first introduce the adaptable servicelevel task model, which is based on a task model presented in (Lu et al., 2002) and extends the notion of a sporadic task system in two major ways. First, worst-case execution times are not assumed. Second, each task $T_{i}$ has a set of service levels, denoted $\operatorname{SL}\left(T_{i}\right)$, each of which represents a different level of QoS for $T_{i}$, and a weight translation function, denoted $\mathrm{g}\left(T_{i}, e, k, q\right)$, which is used to compute the weight of $T_{i}$ at different service levels, as explained below.

The $k^{\text {th }}$ service level of $\operatorname{SL}\left(T_{i}\right)$ is defined by an importance value, $\mathrm{v}\left(T_{i}, k\right)$, a period, $\mathrm{p}\left(T_{i}, k\right)$, and a code segment. Without loss of generality, we assume that the service levels in $\mathrm{SL}\left(T_{i}\right)$ are indexed from 1 to $\left|\mathrm{SL}\left(T_{i}\right)\right|$ by increasing importance value, where $\left|\mathrm{SL}\left(T_{i}\right)\right|$ is the number of elements in $\operatorname{SL}\left(T_{i}\right)$. The importance value represents some user-defined notion of "goodness," where 0.0 represents a service level that has no value and 1.0 represents the maximal possible value associated with any service level of any task in the system.

At any point in time $t$, one service level in $\operatorname{SL}\left(T_{i}\right)$ is said to be the functional service level

[^13]of $T_{i}$. The index of the functional service level of $T_{i}$ at time $t$ is denoted $\mathfrak{f}\left(T_{i}, t\right)$. For now, we assume that the functional service level of a task $T_{i}$ does not change within $\left(\mathrm{r}\left(T_{i}^{j}\right), \mathrm{d}\left(T_{i}^{j}\right)\right)$ for any job $T_{i}^{j}$ of $T_{i}$. In Section 6.2.3, we discuss how to change the functional service level of a task at any time. If $k$ is the functional service level at $\mathrm{r}\left(T_{i}^{j}\right)$, then $T_{i}^{j}$ is said to be functioning at service level $k$. If $T_{i}^{j}$ is functioning at service level $k$, then both $\mathrm{r}\left(T_{i}^{j+1}\right) \geq \mathrm{r}\left(T_{i}^{j}\right)+\mathrm{p}\left(T_{i}, k\right)$ and $\mathrm{d}\left(T_{i}^{j}\right)=\mathrm{r}\left(T_{i}^{j}\right)+\mathrm{p}\left(T_{i}, k\right)$ hold. We consider a task $T_{i}$ to be active at time $t$ if there exists a job $T_{i}^{j}$ (called $T_{i}$ 's active $j o b$ ) such that $t \in\left[r\left(T_{i}^{j}\right), \mathrm{d}\left(T_{i}^{j}\right)\right)$. We use $\operatorname{ACT}(t)$ to denote the set of active jobs at time $t$.

The code segment associated with the $k^{\text {th }}$ service level is the code segment that a job $T_{i}^{j}$ will execute if $T_{i}^{j}$ is functioning at service level $k$. Depending on the specific application, there are numerous different methods for defining such code segments. For some applications, each service level may execute the same code segment, and the only difference between service levels is the period. For other applications, the difference between service levels may be something as simple as the number of iterations in a loop, while for others, each service level may use entirely different code. As we discuss in Section 6.2, how the code segment is implemented will impact the efficacy of AGEDF at adapting tasks.

Just as for sporadic tasks, the value of $\operatorname{Ae}\left(T_{i}^{j}\right)$ denotes the amount of time for which $T_{i}^{j}$ is actually scheduled. The actual weight of a job $T_{i}^{j}$, denoted $\operatorname{Aw}\left(T_{i}^{j}\right)$, represents the actual fraction of a processor that $T_{i}^{j}$ requires and is defined by $\operatorname{Aw}\left(T_{i}^{j}\right)=\operatorname{Ae}\left(T_{i}^{j}\right) / \mathrm{p}\left(T_{i}^{j}, k\right)$, where $T_{i}^{j}$ is functioning at service level $k$. Just as with sporadic tasks, we assume that the value of $\mathrm{Ae}\left(T_{i}^{j}\right)$ (and by extension $\left.\mathrm{Aw}\left(T_{i}^{j}\right)\right)$ is not known until $T_{i}^{j}$ finishes execution.

As we discuss in Section 6.2.1, since the actual weight of a job is not known until it completes, AGEDF uses an estimated weight for incomplete jobs, denoted Ew $\left(T_{i}^{j}\right)$. When AGEDF calculates the estimated weight for a job $T_{i}^{j}$, it does so for a specific service level (typically, the same service level at which $T_{i}^{j-1}$ was functioning). The weight translation function is used to map the estimated weight as calculated by AGEDF for a specific job $T_{i}^{j}$ functioning at a specific service level to what the estimated weight of $T_{i}^{j}$ would be if it functioned at a different service level. Specifically, if $e$ is the estimated weight of $T_{i}^{j}$ assuming that $T_{i}^{j}$ is functioning at service level $k$, then the weight translation function, $\mathrm{g}\left(T_{i}, e, k, q\right)$,


Figure 6.1: Estimated weight vs. importance value/service level for two tasks. (a) $e=0.1$. (b) $e=0.2$.
returns the estimated weight of $T_{i}^{j}$ if it were to be functioning at $q^{\text {th }}$ service level instead of the $k^{\text {th }}$ service level.

Example (Figure 6.1). Consider the example in Figure 6.1, which depicts the estimated weight vs. importance value/service level for two tasks: $T_{1}$ and $T_{2}$, each of which have three service levels with importance values of $0.1,0.2$, and 0.3 . For $T_{1}, \mathrm{~g}\left(T_{1}, e, 1,2\right)=2 e$ and $\mathrm{g}\left(T_{1}, e, 1,3\right)=3 e$, and for $T_{2}, \mathrm{~g}\left(T_{2}, e, 1,2\right)=e^{1 / 4}$ and $\mathrm{g}\left(T_{2}, e, 1,3\right)=e^{1 / 8}$, where $e$ is the estimated weight while functioning at service level one. Inset (a) depicts the scenario where $e=0.1$ for both tasks. Inset (b) depicts the scenario where $e=0.2$ for both tasks. Notice that, in inset (a), if $e=0.1, k=1$, and $q=3$, then $g\left(T_{1}, e, k, q\right)=0.3$, which is the estimated weight of a job of $T_{1}$ if it had been calculated for the third service level instead of the first. Also, for $T_{2}, \mathrm{~g}\left(T_{2}, e, k, q\right) \approx 0.75$, which is the estimated weight of a job of $T_{2}$ if it had been calculated for the third service level instead of the first.

As we discuss in Section 6.2.2, the weight translation function is used to determine the effect on the system caused by changing the functional service level of a task. We make only two assumption about the behavior of $\mathrm{g}\left(T_{i}, e, k, q\right)$ : if $q<k$, then $\mathrm{g}\left(T_{i}, e, k, q\right) \leq e$; and if $\mathrm{g}\left(T_{i}, e_{1}, k, q\right)=e_{2}$, then $\mathrm{g}\left(T_{i}, e_{2}, q, k\right)=e_{1}$. It is important to note that the func-
tion $\mathrm{g}\left(T_{i}, e, k, q\right)$ can return approximate values; however, the accuracy of $\mathrm{g}\left(T_{i}, e, k, q\right)$ will impact the performance of AGEDF's optimizer component, which determines the functional service level of each task. Like service levels and code segments, the weight translation function is defined by the application developer and can be determined empirically.

The primary difference between the task model presented in (Lu et al., 2002) and our task model is that in (Lu et al., 2002) each service level of a task $T_{i}$ has a static notion of "estimated weight" that represents the nominal fraction of a processor required by $T_{i}$. Statically assigning an estimated weight to a task implies that the task has a typical behavior and that if it requires a smaller or larger fraction of a processor, then such a scenario is an anomaly. While this may be true for many applications, for systems like Whisper and VEC, predetermining the nominal weight of a task can be difficult if not impossible. Thus, as we will discuss in Section 6.2, rather than statically determining estimated weights, AGEDF will dynamically calculate the estimated weight for each job.

### 6.2 The AGEDF Scheduling Algorithm

We now present the AGEDF scheduling algorithm and its three components: the predictor (Section 6.2.1), which uses feedback-based techniques to estimate the actual weights of future jobs; the optimizer (Section 6.2.2), which given estimated job weights, attempts to determine an optimal set of functional service levels; and several reweighting rules (Section 6.2.3), which are used to change the functional service level of a task to match that chosen by the optimizer. In the following, we assume that AGEDF is used on an $M$-processor system.

The major components of AGEDF are depicted in Figure 6.2. At a high level, these components function as follows.

- At each instant, the $M$ pending jobs with the smallest deadlines are scheduled.
- At $T_{i}^{j}$ 's completion, the predictor is used to estimate the weight for the next job release of $T_{i}$. If maintaining a constant weight is important, then the reweighting rules may change $T_{i}^{j+1}$ 's functional service level.
- After some user-specified threshold, the optimization component is run to determine


Figure 6.2: (a) The AGEDF scheduling algorithm. (b) The model of AGEDF's feedback component.
new service levels for each task. Then, the following two steps are performed. First, if some tasks require an estimated weight decrease, then the reweighting rules are used to change the service levels of those tasks. This creates spare capacity in the system. Second, as the spare capacity created by weight decreases becomes available, if some tasks require an estimated weight increase, then the reweighting rules are used to change the service levels of those tasks.

It is worthwhile to note that the optimization component (and hence large-scale changes to task functional service levels) is only executed after some user-specified threshold. We offer some guidelines for choosing this threshold in Section 6.2.4.

### 6.2.1 The Feedback Predictor

Before continuing, we briefly review the basics of feedback systems (a thorough review of feedback systems can be found in Section 2.4). Most feedback systems consist of the following components, which are labeled in the model of our system in Figure 6.2(b): the input value, the output value, the actuator, the error, the plant, and the controller. The input value is the reference value for the system, while the output value is value computed by the system.


Figure 6.3: An example of an over-damped, under-damped, and critically-damped feedback system responding to a step input.

The actuator calculates the error by subtracting the output from the input. The plant is the system we wish to control. The controller modifies the input to change the behavior of the output.

The performance of a feedback system is measured in terms of transient response, steadystate error, and stability. The transient response of a system is the initial output of the system to a change in input, as depicted in Figure 6.3. The steady-state error denotes the difference between the output and the input of the system as time increases (also depicted in Figure 6.3). A system is considered to be stable if every bounded input causes the system's steady-state error to be bounded.

It is worthwhile to note that while feedback-based techniques are primarily used to control the behavior of a plant for which the (reference) input is known, another viable use for such techniques is to predict future values of a changing and unknown input. The design of such a system is exactly the same as the typical feedback system, except that the feedback loop does not directly impact the behavior of the system. In such a system, the transient response describes the initial accuracy of predictions after there has been a change in the input, and the steady-state error describes the difference between the predicted and actual values as system time increases.

The feedback predictor. Since the predictor in AGEDF uses feedback-based techniques to predict the weight of future jobs instead of using a simpler approach, such as setting
$\mathrm{Ew}\left(T_{i}^{j}\right)=\operatorname{Aw}\left(T_{i}^{j-1}\right)$, the predictor both produces values of $\mathrm{Ew}\left(T_{i}^{j}\right)$ that are less susceptible to ephemeral fluctuations in the workload and is capable of closely tracking trends in the actual weight (e.g., when the actual weight of the task changes at a constant rate). Using a feedback loop to predict the weight of future jobs is similar to the approach in (Abeni et al., 2002), described earlier in Section 2.4.

As depicted in Figure 6.2(a), in the predictor, each task has its own feedback loop. Also, as depicted in Figure 6.2(b), for each feedback loop, the input is the actual weight; the output is the estimated weight; the error is the actual weight minus the estimated weight; and the controller is a proportional-integral (PI) controller that uses information about the current error and the sum of all previous errors in order to calculate a new estimated weight. Specifically, the controller is defined as

$$
\begin{equation*}
\operatorname{Ew}\left(T_{i}^{j+1}\right)=a \cdot \epsilon\left(T_{i}^{j}\right)+b \sum_{k=1}^{k=j-1} \epsilon\left(T_{i}^{k}\right), \tag{6.1}
\end{equation*}
$$

where $\operatorname{Ew}\left(T_{i}^{1}\right)=0, \epsilon\left(T_{i}^{j}\right)=\operatorname{Aw}\left(T_{i}^{j}\right)-\operatorname{Ew}\left(T_{i}^{j}\right)$, and both $a$ and $b$ are user-defined values that we discuss shortly. Taking the Z-transform of (6.1) and rearranging, we get

$$
\begin{equation*}
\mathrm{G}(z)=\frac{a(z-c)}{z(z-1)}, \tag{6.2}
\end{equation*}
$$

where $c=(a-b) / a$. As discussed in Section 2.4, in control theory parlance, (6.2) is called the open-loop transfer function because it represents the behavior of the controller ignoring the feedback loop. The closed-loop transfer function, which incorporates both the behavior of the controller and feedback loop is given by

$$
\begin{equation*}
\mathrm{H}(z)=\frac{\mathrm{G}(z)}{1+\mathrm{G}(z)}=\frac{a(z-c)}{z^{2}+(a-1) z-a c} . \tag{6.3}
\end{equation*}
$$

From the above equation, the predictor has a closed-loop zero (the value of $z$ for which $H(z)=0$ ) at $z=c$, and closed-loop poles (the values of $z$ for which $H(z)$ is undefined) at

$$
\begin{equation*}
\frac{(1-a) \pm \sqrt{(a-1)^{2}+4 a c}}{2} \tag{6.4}
\end{equation*}
$$

Because the predictor has two closed-loop poles, it is a second-order system. This is the reason why we chose to use a PI controller instead of a proportional-integral-derivative (PID) controller. Since PID controllers are third-order systems, i.e., such systems have three poles, the transient response analysis is substantially more complex. In fact, the typical means for determining the transient response of a third-order system is to approximate it as a secondorder system.

Feedback characteristics. We use standard techniques, which are discussed in detail in Section 2.4, for analyzing the feedback characteristics of our system. We begin by rewriting (6.4) as

$$
\begin{align*}
& \mathcal{P}_{1}=\frac{(1-a)+\sqrt{(a-1)^{2}+4 a c}}{2}  \tag{6.5}\\
& \mathcal{P}_{2}=\frac{(1-a)-\sqrt{(a-1)^{2}+4 a c}}{2} \tag{6.6}
\end{align*}
$$

We let $\mathcal{P}_{m}$ denote the pole from (6.5) and (6.6) that is the farthest from the origin. Also, we use $\mathcal{R}(\mathcal{P})$ and $\theta(\mathcal{P})$ to denote, respectively, the radius and angle (in radians) of the pole $\mathcal{P}$ in polar-complex form.

Stability. By using both the open- and closed-loop transfer functions and the closed-loop poles, we can discuss how setting the values of $a$ and $c$ impact stability, transient response, and steady-state error. First, we address the stability of the system. The system is stable if both closed-loop poles are within the unit circle in the complex plane, i.e., $\mathcal{R}\left(\mathcal{P}_{m}\right)<1$. The system is unstable if either closed-loop pole is outside the unit circle, i.e., $\mathcal{R}\left(\mathcal{P}_{m}\right) \geq 1$. The system is in a state called marginally stable, in which case the output neither converges nor diverges, if one pole is on the unit circle and the other is within it, i.e., $\mathcal{R}\left(\mathcal{P}_{m}\right)=1$.

Transient response. The transient response is usually evaluated by the behavior of the output when the system incurs a step input, i.e., the input suddenly increases to a given value. Since feedback systems use previous results to predict future results, a step input represents the worst-case scenario - a sudden change from one value to a substantially different value.

The transient response of a second-order system is characterized by both the settling time (i.e., the time it takes for the output to attain and stay within $2 \%$ of its steady-state value), and whether the system is over-damped, under-damped, or critically-damped (depicted in Fig. 6.3). The settling time (where time is measured in terms of job releases) of the system is given by the standard formula

$$
\begin{equation*}
\left\lceil\frac{-4}{\ln \left(\mathcal{R}\left(\mathcal{P}_{m}\right)\right)}\right\rceil . \tag{6.7}
\end{equation*}
$$

A lower settling time will improve the system's capacity to respond to sudden changes in the execution time of a task; however, it will also make the system more susceptible to ephemeral fluctuations in the execution time of tasks. As a result, a low settling time may be undesirable for the purposes of predicting future values.

If a system is over-damped, then the output will never "overshoot" the input for a step input. If a system is under-damped, then the output will overshoot the input for a step input. For under-damped systems, the percent overshoot is an additional characteristic of transient response. If the system is critically-damped, then the settling time is as small as possible without causing the output to overshoot the input. Whether a system is under-, over-, or critically-damped depends on the location of the closed loop poles in (6.5) and (6.6). If both poles are unique and both poles are real then the system is over-damped. If both poles have the same radius and are real, then the system is critically-damped. Otherwise, the system is under-damped.

For under-damped systems, the percent overshoot is given by

$$
\begin{equation*}
e^{-\left(\zeta \pi / \sqrt{1-\zeta^{2}}\right)} \cdot 100 \tag{6.8}
\end{equation*}
$$

where $\zeta$ is a value called the damping ratio and is given by

$$
\begin{equation*}
\zeta=\frac{-\ln \left(\mathcal{R}\left(\mathcal{P}_{m}\right)\right)}{\sqrt{\theta\left(\mathcal{P}_{m}\right)^{2}+\ln ^{2}\left(\mathcal{R}\left(\mathcal{P}_{m}\right)\right)}} \tag{6.9}
\end{equation*}
$$

Steady-state error. The steady-state error of a system is measured based on the system's response to a step and/or a ramp input. The ramp input simulates an input that constantly
increases by a rate of $\mathcal{T}$ per job release. The steady-state error is based on the system type, which is given by the power to which $(z-1)$ is raised in the denominator of (6.2). Since our system has a system type of one, the steady-state error for the step input is zero, and the steady-state error of the ramp input is given by

$$
\begin{equation*}
\frac{\mathcal{T}}{\lim _{z \rightarrow 1}(z-1) \mathrm{G}(z)}=\frac{\mathcal{T}}{a(1-c)} \tag{6.10}
\end{equation*}
$$

The fact that a PI controller has zero steady-state error for a step response is the reason why we chose a PI controller instead of a proportional-derivative (PD), which would have a superior transient response but would have a non-zero value for a step input (i.e., if the actual weight is constant, a PD controller would still have error).

Putting it together. Now that we have established formulas for stability, transient response, and steady-state error, it is possible to choose values for $a$ and $c$ (and thus implicitly set the value of $b$ ) that satisfy our design objectives. Suppose, for example, that we wish to construct a critically-damped system with a settling time of five job releases. From the definitions of critically-damped and settling time, it is not difficult to calculate that if $a=0.10206228$ and $c=-1.975$, then these two design objectives are achieved. Specifically, in this case, by (6.5) and (6.6), the closed-loop poles are

$$
\begin{aligned}
& \mathcal{P}_{1}=\frac{(1-a)+\sqrt{(a-1)^{2}+4 a c}}{2} \approx 0.449 \\
& \mathcal{P}_{2}=\frac{(1-a)-\sqrt{(a-1)^{2}+4 a c}}{2} \approx 0.449
\end{aligned}
$$

Thus, $\mathcal{P}_{1} \approx \mathcal{P}_{2}$, which implies that the system is critically-damped (or at least close to it). In addition, by (6.7), the settling time (in terms of number of jobs) is

$$
\left\lceil\frac{-4}{\ln \left(\mathcal{R}\left(\mathcal{P}_{m}\right)\right)}\right\rceil=\lceil 4.997\rceil=5 .
$$

Moreover, by (6.10), we can calculate that the steady-state error of this system for a ramp input that increases at a constant rate of $\mathcal{T}$ per job release is

$$
\frac{\mathcal{T}}{\lim _{z \rightarrow 1}(z-1) \mathrm{G}(z)}=\frac{\mathcal{T}}{a(1-c)} \approx 3.29 \cdot \mathcal{T} .
$$

However, if we wish to construct an under-damped system with a settling time of five job releases, and a percent overshoot of approximately $10 \%$, then it is not difficult to show that $a=1.4008$ and $c=-0.1439$ satisfy these design objectives. Specifically, by (6.5) and (6.6), the closed-loop poles are

$$
\begin{aligned}
& \mathcal{P}_{1}=\frac{(1-a)+\sqrt{(a-1)^{2}+4 a c}}{2} \approx-0.2004+0.4018 i \\
& \mathcal{P}_{2}=\frac{(1-a)-\sqrt{(a-1)^{2}+4 a c}}{2} \approx-0.2004-0.4018 i .
\end{aligned}
$$

Because $\mathcal{P}_{1}$ and $\mathcal{P}_{2}$ are complex values, it follows that the system is under-damped. In addition, because both $\mathcal{P}_{1}$ and $\mathcal{P}_{2}$ are equidistance from the origin, either $\mathcal{P}_{m} \approx-0.2004+$ $0.4018 i$ or $\mathcal{P}_{m} \approx-0.2004-0.4018 i$ holds (the results are the same either way). Because $\mathcal{R}\left(\mathcal{P}_{m}\right)$ and $\theta\left(\mathcal{P}_{m}\right)$ denote, respectively, the radius and angle (in radians) of the pole $\mathcal{P}_{m}$ in polar-complex form, if we set $\mathcal{P}_{m} \approx-0.2004+0.4018 i$, then $\mathcal{R}\left(\mathcal{P}_{m}\right) \approx 0.449$ and $\theta\left(\mathcal{P}_{m}\right) \approx$ -1.108 .

Thus, by (6.7), the settling time (in terms of number of jobs) is given by

$$
\left\lceil\frac{-4}{\ln \left(\mathcal{R}\left(\mathcal{P}_{m}\right)\right)}\right\rceil=\lceil 4.995\rceil=5 .
$$

Furthermore, by (6.9), we can calculate the damping ratio as

$$
\zeta=\frac{-\ln \left(\mathcal{R}\left(\mathcal{P}_{m}\right)\right)}{\sqrt{\theta\left(\mathcal{P}_{m}\right)^{2}+\ln ^{2}\left(\mathcal{R}\left(\mathcal{P}_{m}\right)\right)}} \approx 0.5857 .
$$

Thus, by (6.8), we can calculate the percent overshoot as

$$
e^{-\left(\zeta \pi / \sqrt{1-\zeta^{2}}\right)} \cdot 100 \approx 10.33 \% .
$$

Moreover, by (6.10), we can calculate that the steady-state error of this system for a ramp input that increases at a constant rate of $\mathcal{T}$ per job release as

$$
\frac{\mathcal{T}}{\lim _{z \rightarrow 1}(z-1) \mathrm{G}(z)}=\frac{\mathcal{T}}{a(1-c)} \approx 0.624 \cdot \mathcal{T} .
$$

### 6.2.2 Optimization

As mentioned above, the optimization component of AGEDF uses the estimated weights of tasks in order to choose service levels for each task. There are a variety of different methods for implementing this component depending on what metric the user wants to optimize and the behavior of $\mathrm{g}\left(T_{i}, e, k, q\right)$.

For example, suppose the objective is to optimize the total importance value in the system. In this case, if the relationship between the importance value and weight is linear (like $T_{1}$ in Figure 6.1), then an approximate solution for this objective could be achieved by assigning the highest service level possible to those tasks with the highest value density, as given by,

$$
\begin{equation*}
\frac{\mathrm{v}\left(T_{i},\left|\mathrm{SL}\left(T_{i}\right)\right|\right)-\mathrm{v}\left(T_{i}, 1\right)}{\mathrm{g}\left(T_{i}, \operatorname{Ew}\left(T_{i}^{j}\right), k,\left|\operatorname{SL}\left(T_{i}\right)\right|\right)-\mathrm{g}\left(T_{i}, \operatorname{Ew}\left(T_{i}^{j}\right), k, 1\right)}, \tag{6.11}
\end{equation*}
$$

while ensuring at least every task receives its minimum service level and the system is not over-utilized. In this approach, the value $\mathrm{v}\left(T_{i},\left|\mathrm{SL}\left(T_{i}\right)\right|\right)-\mathrm{v}\left(T_{i}, 1\right)$ denotes by how much $T_{i}$ 's importance value improves by changing from the lowest service level to the highest service level. Additionally, the value $\mathrm{g}\left(T_{i}, \operatorname{Ew}\left(T_{i}^{j}\right), k,\left|\mathrm{SL}\left(T_{i}\right)\right|\right)-\mathrm{g}\left(T_{i}, \operatorname{Ew}\left(T_{i}^{j}\right), k, 1\right)$ represents by how much $T_{i}$ 's weight would have to be changed to improve its service level from the lowest service level to the highest service level. Notice that, for two tasks $T_{1}$ and $T_{2}$, if $T_{1}$ has a larger value for (6.11) than $T_{2}$, and both tasks had their estimated weight increase by the same amount, then $T_{1}$ 's importance value would improve more than $T_{2}$ 's. This approach is similar to the highest-value-density-first approach used in (Lu et al., 2002).

Example (Figure 6.4). Consider the example in Figure 6.4, which depicts the estimated weight vs. importance value/service level for two tasks: $T_{1}$ and $T_{2}$, each of which has three service levels with importance values of $0.1,0.2$, and 0.3 . Inset (a) depicts the scenario


Figure 6.4: Estimated weight vs. importance value/service level for two tasks. (a) This relationship is linear for both tasks. (b) This relationship is linear for $T_{1}$ and non-linear for $T_{2}$.
where both tasks have a linear relationship between the estimated weight and the importance value/service level. Inset (b) depicts the scenario where the estimated weight and the importance value/service level relationship is linear for $T_{1}$ and non-linear for $T_{2}$. Notice that, in inset (a), the value densities for $T_{1}$ and $T_{2}$ are, respectively, $\frac{0.2}{0.2}=1$ and $\frac{0.2}{0.6}=\frac{1}{3}$. Thus, improving $T_{1}$ 's service level requires less weight. Hence, by the highest-value-density-first rule, the service level of $T_{1}$ is improved before $T_{2}$.

On the other hand, if the relationship between the importance value and weight is nonlinear (like $T_{2}$ in Figure 6.1), then an approximate solution for this objective could be achieved by using nonlinear programming techniques such as steepest descent or Newton's method (Bertsekas, 1999). If exact solutions are required, then techniques like branch-andbound can be used offline, and the optimization component could then switch between several predetermined system states. (In Section 7.5, we discuss how we implemented the optimizer for both Whisper and the VEC.)

Example (Figure 6.4). Notice that, in Figure 6.4(b), the relationship between the the estimated weight and the importance value/service level for $T_{2}$ is non-linear. Moreover, the value density for $T_{1}$ and $T_{2}$ are, respectively, $\frac{0.2}{0.2}=1$ and $\frac{0.2}{0.65} \approx 0.308$. Thus, if the highest-
value-density-first approach were used to determine the service levels of these two tasks, then $T_{1}$ would be improved before $T_{2}$. However, improving $T_{2}$ 's service level from Level 2 to Level 3 requires less weight than improving $T_{1}$ 's service level by one level. Thus, the highest-value-density-first approach would not produce an optimal distribution of weights for these two tasks.

It is important to note that the use of weight translation functions is the reason why the optimization component is extensible because it allows any optimizing function to assess the impact of changing the functional service level. In prior work on adaptive real-time systems, the two primary methods for optimizing service levels have been to assume either that each service level has a "nominal" utilization (Lu et al., 2002) or the relationship between the service level and importance value is linear (Marti et al., 2004). As we discussed in Section 6.1, the problem with the first approach is that assessing a meaningful "nominal" utilization may be difficult if not impossible for many applications. The problem with the second approach is that there exist applications for which linearity cannot be assumed. For example, consider any video application in which each service level corresponds to a different resolution. Typically, in such a system, as the service level (and by extension the resolution) increases, the amount of benefit to user perception per pixel added decreases. It is easy to see that in such a scenario, the relationship between importance value and estimated weight is nonlinear.

### 6.2.3 Reweighting

Whenever a task is reweighted (i.e., changes its functional service level) either by the optimization component or by the main AGEDF algorithm, its code segment and/or period may change. If no job of a task, $T_{i}$, is active when $T_{i}$ changes its functional service level from the $\ell_{0}^{\text {th }}$ to $\ell_{1}^{\text {th }}$ service level at time $t$, then the change is simple - the next released job of $T_{i}$ has the period and code segment associated with the $\ell_{1}^{\text {th }}$ service level. If a job of $T_{i}$ is active at $t$, then the situation is more complicated. For the remainder of this section, let $T_{i}^{j}$ denote the active job of $T_{i}$ at $t$. Recall from Chapter 3 that, when a task with an active job reweights, there can be a difference between when it "initiates" the change and when the change is "enacted." The time at which the change is initiated is defined externally to the reweighting component
(by either the optimization component or the main AGEDF algorithm); the time at which the change is enacted, i.e., the functional service level is changed, is dictated by reweighting rules.

Changing the period. In order to change the period of a job, we use the GEDF reweighting Rules P and N, presented in Section 3.4.1. For brevity, we do not review these rules here.

Changing the code segment. Whether the code segment of the task $T_{i}$ that released $T_{i}^{j}$ can change depends on the implementation of $T_{i}$. For example, if the $\ell_{0}^{\text {th }}$ and $\ell_{1}^{\text {th }}$ service levels have substantially different code segments, then $T_{i}^{j}$ cannot change its code segment. On the other hand, suppose that the difference between the code segments for the $\ell_{0}^{t h}$ and $\ell_{1}^{t h}$ service levels is simply the number of iterations in a loop. Then, as long as $T_{i}^{j}$ is not complete and this change would not cause either $\operatorname{Ew}\left(T_{i}^{j}\right)>1$ or $\sum_{T_{a}^{b} \in \operatorname{ACT}(t)} \operatorname{Ew}\left(T_{a}^{b}\right)>M, T_{i}^{j}$ can change its code segment immediately. Moreover, if the code segment is changed, then $\mathrm{Ew}\left(T_{i}^{j}\right)$ is changed to

$$
\begin{equation*}
\frac{\max \left(\mathrm{Nw} \cdot \mathrm{p}\left(T_{i}, \ell_{1}\right), \mathrm{A}\left(\mathcal{S}, T_{i}^{j}, \mathrm{r}\left(T_{i}^{j}\right), t\right)\right)}{\mathrm{p}\left(T_{i}, \ell_{0}\right)} \tag{6.12}
\end{equation*}
$$

where $\mathcal{S}$ is the AGEDF schedule, and $\mathrm{Nw}=\mathrm{g}\left(T_{i}, \operatorname{Ew}\left(T_{i}^{j}\right), \ell_{0}, \ell_{1}\right)$. Notice that the estimated amount of time for which $T_{i}^{j}$ will execute as a consequence of changing its code segment is the larger of the amount of time it has already been schedule by time $t$, i.e., $\mathrm{A}\left(\mathcal{S}, T_{i}^{j}, \mathrm{r}\left(T_{i}^{j}\right), t\right)$, and the amount of time that $T_{i}^{j}$ would have been scheduled if the $\ell_{1}^{t h}$ service level was the functional service level at $\mathrm{r}\left(T_{i}^{j}\right), \mathrm{Nw} \cdot \mathrm{p}\left(T_{i}, \ell_{1}\right)$. Thus, the estimated weight of $T_{i}^{j}$ is the estimated amount of time that $T_{i}^{j}$ will be scheduled divided by $\mathrm{p}\left(T_{i}, \ell_{0}\right)$. (In Section 7.5, we discuss how we implemented the code segment for both Whisper and the VEC.)

Example (Figure 6.5). Consider the example in Figure 6.5, which depicts a three-processor system scheduled by AGEDF with three tasks, all of which have a period of 7 , an estimated weight of $3 / 7$, and in the absence of a weight change, would be scheduled for 3 time units. At time 1, all three tasks experience a service level change that changes the code segment for each job. Moreover, we assume that $\mathrm{g}\left(T_{1}, \operatorname{Ew}\left(T_{1}^{1}\right), \ell_{1,0}, \ell_{1,1}\right)=4 / 7, \mathrm{~g}\left(T_{2}, \operatorname{Ew}\left(T_{2}^{1}\right), \ell_{2,0}, \ell_{2,1}\right)=2 / 7$, and $\mathrm{g}\left(T_{3}, \operatorname{Ew}\left(T_{3}^{1}\right), \ell_{3,0}, \ell_{3,1}\right)=1 / 14$, where $\ell_{i, 0}$ is the initial service level of $T_{i}$ and $\ell_{i, 1}$ is the service level that $T_{i}$ changes to at time 1 . Thus, as a result of the change, $T_{1}^{1}$ executes for 4


Figure 6.5: An illustration of changing the code segment.
time units, $T_{2}^{1}$ executes for 2 time units, and $T_{3}^{1}$ executes for 1 time unit. Notice that $\operatorname{Ew}\left(T_{3}^{1}\right)$ is changed to $1 / 7$ even though $g\left(T_{3}, 3 / 7, \ell_{3,0}, \ell_{3,1}\right)=1 / 14$. The reason for this is that $1=$ $\mathrm{A}\left(\mathcal{S}, T_{3}^{1}, 0,1\right)>\mathrm{Nw} \cdot \mathrm{p}\left(T_{i}, \ell_{1}\right)=\frac{1}{14} \cdot 7=1 / 2$. Thus, by $(6.12), \operatorname{Ew}\left(T_{3}^{1}\right)=\frac{\max (1,1 / 2)}{7}=1 / 7$.

### 6.2.4 User-Defined Threshold

Choosing a specific user-defined threshold for invoking the optimizer will depend largely on the targeted application. Some possible thresholds could include a duration of time, a substantial change in the estimated weight for one task, or a substantial change in the total estimated weight for all tasks. While running the optimizer more frequently will increase the accuracy of the system, it will also increase the amount of time the scheduler is active with the system not producing "useful" work. Additionally, as we discussed in Section 6.2.3, the reweighting rules cannot always be enacted immediately. Thus, if the optimizer is called before all changes have been enacted, then it may produce an inaccurate result. Notice that, if the weight translation function is accurate, then after all reweighting events have been enacted, the system will remain in an "optimal" state, unless the actual weight of a task changes. Thus, if the separation between optimizer invocations is sufficiently large for all tasks to enact their functional service level changes (i.e., at least the largest period of a job in the system), then it is possible to guarantee that no task will unnecessarily "thrash" between service levels. (In Section 7.5, we discuss how we chose the user-defined threshold for both Whisper and the VEC.)

### 6.3 Conclusion

In this chapter, we presented the adaptable service-level task model, which is based on a task model presented in (Lu et al., 2002). In addition, we presented the AGEDF scheduling framework, in which tasks are scheduled by GEDF augmented with three components to facilitate adaption: a feedback predictor, an optimizer, and a task reweighter. It is worth noting that these components are modular. As a result, a developer could modify each of these components to improve the performance for a specific application.

## CHAPTER 7

## IMPLEMENTATION and EXPERIMENTS

In this chapter, we present two sets of experiments. First, we present simulations in which Whisper and VEC are scheduled by our adaptive variants of GEDF, NP-GEDF, PEDF, NP-PEDF, and PD ${ }^{2}$. Second, we present experiments conducted using the real-time Linux testbed LITMUS ${ }^{\text {RT }}$ to evaluate the performance of AGEDF when running the core operations of Whisper and VEC.

Unfortunately at this time, it is not feasible to produce experiments involving a complete implementation of either Whisper or VEC, for two reasons. First, both the existing Whisper and VEC designs are single-threaded (and non-adaptive) and consist of several thousands of lines of code. Converting each implementation to a multi-threaded implementation is a nontrivial task. Indeed, because of this, it is essential that we first understand the scheduling and resource-allocation trade-offs involved. The development of our various adaptive algorithms can be seen as an attempt to articulate these tradeoffs. Second, support for task synchronization is required, and while there has been work on real-time task synchronization on multiprocessors (Block et al., 2007; Brandenburg et al., 2008; Brandenburg and Anderson, 2008; Chen and Tripathi, 1994; Devi et al., 2006; Gai et al., 2003; Holman and Anderson, 2006; Lopez et al., 2004; Rajkumar, 1991; Sha et al., 1990), applying such work in the context of adaptive scheduling algorithms is non-trivial. For these reasons, we have chosen to conduct our evaluation using both simulations and an implementation of the core operations for Whisper and VEC, i.e., correlation computations for Whisper and bilateral filters for VEC.

We begin this chapter with brief descriptions of Whisper (Section 7.1), VEC (Section 7.2), and LITMUS ${ }^{\text {RT }}$ (Section 7.3). Then, in Section 7.4, we present our simulations of Whisper


Figure 7.1: The Whisper system.
and VEC when scheduled via the adaptive variants of GEDF, NP-GEDF, PEDF, NP-PEDF, and $\mathrm{PD}^{2}$. In Section 7.5, we present our evaluation of the AGEDF algorithm when implemented under LITMUS ${ }^{R T}$ and scheduling the core operations of both Whisper and VEC. (We emphasize that this set of experiments involved running real code on a real OS kernel and are not merely simulations.) Finally, we conclude in Section 7.6.

### 7.1 Whisper

As depicted in Figure 7.1, Whisper tracks users via speakers that each emit a unique sound wave and are attached to each user's hands, feet, and head. Microphones located on the wall and ceiling receive these signals and a tracking computer calculates (via a speed-of-sound computation) each speaker's position by measuring signal delays. Whisper is able to compute the signal delay between the transmitted and received versions of the sound by performing a correlation calculation on the most recent set of samples. Because correlations are computationally intensive, Whisper uses a Kalman filter to decrease the number of correlations required to track a user.

We begin this section by reviewing the concepts of correlation (Section 7.1.1) and the Kalman filter (Section 7.1.2). Next, we discuss the impact of occluding objects (Section 7.1.3). We conclude this section with a discussion of Whisper's real-time characteristics (Section 7.1.4).

```
\(\operatorname{Cor}(x\) : array \([0 \ldots m-1]\) of doubles, \(t\) : array \([0 \ldots n-1]\) of doubles):
    array \([0 \ldots m-n-1]\) of doubles
    \(i, j\) : integer;
    \(y\) : array \([0 \ldots m-n-1]\) of doubles
    for \(i:=0\) to \(m-n-1\) do
        \(y[i]:=0 ;\)
        for \(j:=0\) to \(n-1\) do
            \(y[i]:=y[i]+x[i+j] \cdot t[j]\)
        od
    od;
    return \(y\)
```

Figure 7.2: Pseudo-code defining correlation. $x$ is the received signal, $t$ is the target signal, and $y$ is cross-correlation signal.

### 7.1.1 Correlation

Correlation is a signal-processing technique for locating a known waveform in a signal. In this section, we briefly review the central concepts behind correlation computations. A more detailed discussion can be found in (Smith, 1997).

As an input, correlation takes two discrete signals, $t$ and $x$, where $t$ is the known waveform, called the target signal, that contains $n$ samples, and $x$ is the received signal that contains $m$ samples all of which have some level of white noise. As an output, correlation produces a discrete signal, $y$, called the cross-correlation signal, of $m-n$ samples, where $y[i]=\sum_{0 \leq j<n}(t[j] \cdot x[i+j])$. (Pseudo-code for the correlation computation is given in Figure 7.2.) The cross-correlation signal has the property that the value of $i$ for which $y[i]$ is the maximal value in $y$ denotes the index at which the signal $t$ likely appears in $x$.

Example (Figure 7.3). Consider the example in Figure 7.3, which illustrates the received signal with noise $x$, known waveform $t$, and output of the correlation $y$. Notice that the value of $y[2]$ is substantially larger than any other value in $y$. Thus, it is easy to see that the waveform $t$ begins at $x[2]$.

Whisper uses correlation computations to determine the number of samples that have elapsed from the time a signal is emitted by a speaker to the time it is received by a microphone. Whisper is capable of making such a calculation because the microphones and speakers are synchronized. So, in Figure 7.3, if $x$ is the signal received by a microphone from


Figure 7.3: An illustration of correlation.
a speaker starting at time 0 , and $t$ is the signal sent by that speaker at time 0 , then two samples would have elapsed between the time when the signal was sent to the time it was received.

Signal-to-noise ratio. It is important to note that the ability of a correlation computation to determine the location of the target signal in the received signal is directly related to the signal-to-noise ratio. This behavior occurs because, as the signal-to-noise ratio decreases, the relative difference between the maximal value in the cross-correlation value and the other values in the cross-correlation signal decreases. It is possible to compensate for a decreasing signal-to-noise ratio by increasing the number of samples in the target signal; however, this increases the computation time of a task.

Speed of sound. Notice that, once Whisper has used a correlation computation to determine the location of the target signal in the received signal, then (given the number of samples emitted/received per second and the speed of sound) it is a simple matter to compute the distance between a speaker/microphone pair. Specifically, if a speaker/microphone pair emits/receives $k$ samples per second, and from a correlation computation it is determined that $\ell$ samples have elapsed from the time the signal was sent from the speaker to the time
it was received by the microphone, then the total time it took for the sound to reach the microphone from the speaker is $\ell / k$. Thus, the distance between the speaker/microphone pair is

$$
c \cdot \frac{\ell}{k}
$$

where $c$ is the speed of sound.

Example. Consider a scenario where a speaker/microphone pair emits/receives 1,000 samples per second. If a correlation computation determines that 50 samples have elapsed between the time a signal was sent from the speaker to the time it was received by the microphone, then the sound from the speaker took 0.05 seconds to reach the microphone. Thus, given that the speed of sound is approximately $343 \mathrm{~m} / \mathrm{s}$, the distance between the microphone speaker pair is $0.05 \mathrm{~s} \cdot 343 \mathrm{~m} / \mathrm{s} \approx 17.15 \mathrm{~m}$.

### 7.1.2 Kalman Filter

One of the drawbacks to using correlation to locate a known waveform is that it is computationally expensive. Specifically, the cost of using a correlation to locate a target signal $t$ of length $m$ in a received signal $x$ of length $n$ is $O(m \cdot n)$. In order to reduce this cost, Whisper attempts to estimate the location of $t$ in $x$ by using a Kalman filter. The correlation computation merely verifies the estimated location of the target rather than searching through the entire received signal.

The Kalman filter is a recursive mathematical process that combines multiple measurements and the error associated with those measurements to produce an estimated value that is more accurate than any previous measurement. Because understanding the Kalman filter requires knowledge of digital signal processing that is will beyond the scope of this dissertation, we present an intuitive example here and refer the reader to (Welch and Bishop, 1995) for a more detailed discussion.

Example (Figure 7.4). Consider the following example (originally presented in (Maybeck, 1979)), depicted in Figure 7.4, in which two sailors attempt to find their location in onedimension by using the stars. By using his sextant, the first sailor estimates the ship's


Figure 7.4: An illustration of the Kalman Filter.
position as $z_{1}$; however, due to both human and equipment error, the standard deviation for this measurement is $\sigma_{1}$. A second sailor measures their position as $z_{2}$, with a standard deviation of $\sigma_{2}$, which is less than $\sigma_{1}$. The Kalman filter combines both measurements to produce a measurement, $\mu$, with a smaller deviance, $\sigma$, as given by the formulas

$$
\begin{gathered}
\mu=\frac{\sigma_{2}^{2} \cdot z_{1}}{\sigma_{1}^{2}+\sigma_{2}^{2}}+\frac{\sigma_{1}^{2} \cdot z_{2}}{\sigma_{1}^{2}+\sigma_{2}^{2}} \\
\frac{1}{\sigma^{2}}=\frac{1}{\sigma_{1}^{2}}+\frac{1}{\sigma_{2}^{2}} .
\end{gathered}
$$

Notice that, while the value $\mu$ is between the two values $z_{1}$ and $z_{2}$, it is closer to $z_{2}$ since $\sigma_{2}$ is less than $\sigma_{1}$. Also notice that the standard deviation produced by the Kaman filter, i.e., $\sigma$, is less than both $\sigma_{1}$ and $\sigma_{2}$.

The Kalman filter and Whisper. In Whisper, the Kalman filter is used both before and after the correlation computation. Before Whisper runs the correlation computation, the Kalman filter uses the previous position and velocity of a tracked object to produce an estimate of its position. The second time the Kalman filter is used, the position of an object, produced by the correlation computation, is fed back into the Kalman filter to refine its internal state. This loop is depicted in Figure 7.5. It is important to note that the computation time associated with the Kalman filter is dwarfed in comparison to the computation time required for the correlations.


Figure 7.5: The core loop in Whisper.


Figure 7.6: An illustration of an occluding object.

### 7.1.3 Occlusion

One of the advantages of using sound to track a user is that it can bend around objects. As a result, if a user's torso, arms, head, or any other body part obstruct the path between a speaker/microphone pair, then it is still possible to estimate the distance between the pair. This being said, occlusions impact the performance of Whisper in two ways. First, an occluding object acts as a low-pass filter since high frequencies are attenuated as they pass around objects. As a result, the correlation computation will not be as precise. Second, as depicted in Figure 7.6, occluding objects increase the perceived distance between the speaker/microphone pair. The problem with increasing the perceived distance is that it is extremely difficult to correct because the system cannot distinguish between interference by an occluding object or an increase in distance. As a result, this introduces an additional source of error in Whisper's measurements. Whisper handles the additional source of error by estimating the maximal amount of occlusion that an object is likely to cause in order to calculate an estimated upper bound on the error associated with any measurement. Then, this information is used in the Kalman filter, which can partially mitigate the error associated with occlusion.

### 7.1.4 Real-time Characteristics

As noted above, when the signal-to-noise ratio decreases, correlation become less effective at finding the target signal. To compensate for a decreasing signal-to-noise ratio, the accuracy of Whisper can be improved at the expense of additional computation by either increasing the number of samples in the target signal or increasing the number of position updates per second. In addition, notice that the signal-to-noise ratio is inversely related to the distance between a speaker/microphone pair (since as the distance between a speaker/microphone pair grows, the signal becomes weaker). As a result, when the distance between a speaker/microphone pair is large, the signal-to-noise ratio will be small, which implies that the task associated with this pair will need more computation time to compensate. Thus, as users move around in a virtual environment, the processor shares of tasks assigned to different speaker/microphone pairs must change to compensate for dynamic signal-to-noise ratios.

In addition, since Whisper continuously performs calculations on incoming data, at any point in time, it does not have a significant amount of "useful" data stored in cache. As a result, migration/preemption costs in Whisper are fairly small (at least, on a tightly-coupled system, as assumed here, where the main cost of a preemption or migration is a loss of cache affinity). In addition, fairness and real-time guarantees are important due to the inherent "tight coupling" among tasks required to accurately perform triangulation calculations.

### 7.2 VEC

In this section, we provide a brief introduction to VEC. Just as with Whisper, a detailed discussion of VEC would involve aspects of multimedia systems that are well beyond the scope of this dissertation. Thus, we refer readers to (Bennett, 2007) for a detailed description of VEC. We begin our discussion of VEC by first reviewing some basics of videography and providing an overview of the system.

Video as a collection of frames. All video is a collection of still images called frames. Associated with each pixel in a video frame are luminance and chrominance values. The luminance value denotes the brightness of the pixel (the higher the value, the brighter the
pixel) and the chrominance values denote the color. The luminance of a frame can be increased by lengthening the time the camera's shutter is open, called the exposure time. Frames with faster exposure times capture moving objects with more detail, while frames with slower exposure times are brighter. If a frame is underexposed (i.e., the exposure time is too fast), then the image can be too dark to discern any object.

VEC corrects underexposed video while maintaining the detail captured by faster exposure times by combining the information of multiple frames. To intuitively understand how VEC achieves this behavior, consider the following example. If a camera, $\mathbf{A}$, has an exposure time of $1 / 30^{\text {th }}$ of a second, and a second camera, $\mathbf{B}$, has an exposure time of $1 / 15^{\text {th }}$ of a second, then for every two frames shot by camera $\mathbf{A}$, the shutter is open for the same time as one frame shot by B. VEC is capable of exploiting this observation in order to allow camera $\mathbf{A}$ to shoot frames with the detail of $1 / 30^{\text {th }}$ of a second exposure time but the brightness of $1 / 15^{\text {th }}$ of a second exposure time.

### 7.2.1 Bilateral Filter

The primary complication with adjusting the luminance values of pixels to correctly expose a frame is the presence of noise in an image. As a result, a pixel may have different luminance values across multiple frames even if the image being recorded is static, and adjacent pixels may have different luminance values even if all represent the same object.

Example (Figure 7.7). Consider the example in Figure 7.7, which illustrates a $5 \times 5$ pixel region of a video frame. The black pixels all represent the same dark object and would all have a luminance value of 20 , if not for noise. Thus, in the absence of noise the luminance of $p 13$ is 20 . Inset (a) depicts the luminance value of each pixel. Insets (b), (c), and (d) are discussed later. Notice that, because of noise, pixels that represent the same object may have subtle variations in luminance intensity.

One approach for removing the luminance noise from a pixel, henceforth referred to as the origin pixel, is to change its luminance to be a weighted ${ }^{1}$ average of every pixel within

[^14]| $\begin{gathered} \text { p1 } \\ \text { l=20 } \end{gathered}$ | $\underset{\mathrm{l}=10}{\mathrm{p} 2}$ | $\underset{\mathrm{l}=10}{\mathrm{p} 3}$ | $\begin{gathered} \text { p4 } \\ \mathrm{l}=10 \end{gathered}$ | $\underset{\mathrm{l}=200}{\mathrm{p} 5}$ |
| :---: | :---: | :---: | :---: | :---: |
| $\begin{gathered} p 6 \\ l=10 \end{gathered}$ | $\underset{l=10}{p 7}$ | $\begin{gathered} \mathrm{p} 8 \\ \mathrm{l}=20 \end{gathered}$ | $\underset{\mathrm{l}=10}{\mathrm{p} 9}$ | $\underset{\mathrm{l}=190}{\substack{\text { p10 }}}$ |
| $\begin{array}{r} \text { p11 } \\ \mathrm{l}=20 \end{array}$ | $\begin{gathered} \text { p12 } \\ \mathrm{l}=20 \end{gathered}$ | $\begin{array}{r} \text { p13 } \\ \mathrm{l}=22 \end{array}$ | $\begin{array}{r} \text { p14 } \\ \mathrm{l}=10 \end{array}$ | $\underset{\mathrm{l}=200}{\mathrm{p} 15}$ |
| $\underset{\mathrm{l}=20}{\mathrm{p} 16}$ | $\begin{array}{r} \text { p17 } \\ \text { l=10 } \end{array}$ | $\underset{\substack{\text { p18 }}}{\substack{ \\\hline}}$ | $\begin{gathered} \text { p19 } \\ \mathrm{l}=20 \end{gathered}$ | $\underset{\mathrm{I}=190}{\substack{\text { p20 }}}$ |
| $\underset{\mathrm{l}=200}{\mathrm{p} 21}$ | $\underset{\mathrm{l}=190}{\substack{\text { p22 }}}$ | $\underset{\mathrm{l}=200}{\mathrm{p} 23}$ | $\underset{\substack{\text { p } \\ \hline \\ \hline}}{ }$ | $\underset{\mathrm{l}=190}{\mathrm{p} 25}$ |

(a)

| p1 | p2 | $p 3$ | $p 4$ | p5 |
| :---: | :---: | :---: | :---: | :--- |
| 0.017 | 0.017 | 0.017 | 0.017 | 0.017 |
| $p 6$ | $p 7$ | $p 8$ | $p 9$ | $p 10$ |
| 0.017 | 0.076 | 0.076 | 0.076 | 0.017 |
| p11 | $p 12$ | $p 13$ | $p 14$ | $p 15$ |
| 0.017 | 0.076 | 0.125 | 0.076 | 0.017 |
| p16 | $p 17$ | $p 18$ | $p 19$ | $p 20$ |
| 0.017 | 0.076 | 0.076 | 0.076 | 0.017 |
| $p 21$ | $p 22$ | $p 23$ | $p 24$ | $p 25$ |
| 0.017 | 0.017 | 0.017 | 0.017 | 0.017 |

(c)

| p1 | p2 | p3 | p4 | p5 |
| :--- | :--- | :--- | :--- | :--- |
| 0.04 | 0.04 | 0.04 | 0.04 | 0.04 |
| p6 | p7 | p8 | p9 | p10 |
| 0.04 | 0.04 | 0.04 | 0.04 | 0.04 |
|  |  |  |  |  |
| p11 | p12 | p13 | p14 | p15 |
| 0.04 | 0.04 | 0.04 | 0.04 | 0.04 |
| p16 | p17 | p18 | p19 | p20 |
| 0.04 | 0.04 | 0.04 | 0.04 | 0.04 |
| p21 | p22 | p23 | p24 | p25 |
| 0.04 | 0.04 | 0.04 | 0.04 | 0.04 |

(b)

| p1 | p2 | $p 3$ | $p 4$ | p5 |
| :---: | :---: | :---: | :---: | :--- |
| 0.019 | 0.019 | 0.019 | 0.019 | 0.004 |
| $p 6$ | $p 7$ | $p 8$ | $p 9$ | $p 10$ |
| 0.019 | 0.085 | 0.076 | 0.085 | 0.005 |
| $p 11$ | $p 12$ | $p 13$ | $p 14$ | $p 15$ |
| 0.019 | 0.086 | 0.141 | 0.085 | 0.004 |
| p16 | $p 17$ | $p 18$ | $p 19$ | $p 20$ |
| 0.019 | 0.085 | 0.086 | 0.086 | 0.005 |
| $p 21$ | $p 22$ | $p 23$ | $p 24$ | $p 25$ |
| 0.004 | 0.005 | 0.004 | 0.005 | 0.005 |

(d)

Figure 7.7: Example of different methods for removing noise from a pixel.
a nearby proximity. (This distance can be calculated either spatially, i.e., the space between two pixels on the same frame, or temporally, i.e., the number of frames between two pixels on different frames.) While such an approach will remove the noise from the origin pixel, it is not immediately obvious what is the best method for determining how much each surrounding pixel should contribute to the origin pixel's final luminance value. The simplest approach is to value all pixels within a given distance equally. The problem with this approach is that it does not account for the fact that the closer a pixel is to the origin pixel the more likely both pixels represents the same object.

Example (Figure 7.7(b)). Such a solution is illustrated in Figure 7.7(b). In this inset (as well as insets (c) and (d)), the values depict the contribution associated with each pixel when computing a weighted average for pixel $p 13$, e.g., in this inset, the weighted average for $p 13$ is $0.04 \cdot I_{p 1}+0.04 \cdot I_{p 2}+0.04 \cdot I_{p 3}+\ldots 0.04 \cdot I_{p 25}$, where $I_{p}$ is the luminance intensity of the pixel $p$. The weighted average of $p 13$ 's luminance computed in this way is 79.68 , which is nearly four-times its actual noise-free luminance of 20 . The reason why this weighted average calculation is so inaccurate is because it values the luminance of all pixels equally, regardless of how close they are to the origin pixel, e.g., pixel $p 18$ has the same weight as pixel $p 23$, even though pixel $p 18$ is closer to pixel $p 13$.

An alternative approach is to value pixels that are closer (again, either temporally or spatially) over those that are far away. One method of determining a pixel's contribution to the weighted average is to use a Gaussian distribution. Specifically, if $x$ is the temporal or spatial distance between a pixel $p$ from the origin pixel $s$, and $\sigma$ is used to determine the rate of fall off (i.e., the higher the value of $\sigma$ the steeper the Gaussian distribution), then $p$ 's contribution to the weighted average of $s$ is

$$
\begin{equation*}
g(x, \sigma)=\frac{e^{\frac{-x^{2}}{2 \sigma^{2}}}}{\sigma \sqrt{2 \pi}} \tag{7.1}
\end{equation*}
$$

The problem with this approach is that nearby pixels may not represent the same object.
Example (Figure 7.7(c)). Such a solution is illustrated in Figure 7.7(c). In this case, in this inset, the weighted average of $p 13$ 's luminance is 43.42 , which is nearly two-times its
actual noise-free luminance of 20 . This measurement is more accurate than that given by Figure 7.7 (b) because the closer a pixel is to $p 13$ the more it contributes to the weighted average; however, this approach still produces an inaccurate result because nearby pixels may not represent the same object. To observe this behavior, notice that, the luminance of $p 11$ should have a greater impact on the final value of $p 13$ than $p 15$ (since $p 11$ and $p 13$ represent the same object in the video, while $p 15$ represents a different object); however, since $p 15$ and $p 11$ are the same distance from $p 13$, a purely distance-based approach values both pixels equally.

A third approach, called the bilateral filter, involves taking a weighted average that not only gives preference to closer pixels but also pixels that have a similar luminance. By adding similarity in luminance as an additional constraint, such a technique would not be as prone to interference by abutting objects.

Example (Figure 7.7(d)). Such a solution is illustrated in Figure 7.7(d). Notice that, in this inset, the weighted average of $p 13$ 's luminance is 23.23 , which nearly equals its actual noise-free luminance of 20 .

Formally, the bilateral filter of the pixel $s$, with a spatial/temporal fall off of $\sigma_{h}$ and a luminance similarity fall off of $\sigma_{i}$, is given by the formula

$$
\begin{equation*}
B\left(s, \sigma_{h}, \sigma_{i}\right)=\frac{\sum_{p \in N_{s}} g\left(\|p-s\|, \sigma_{h}\right) \cdot g\left(D(p, s), \sigma_{i}\right) \cdot I_{p}}{\sum_{p \in N_{s}} g\left(\|p-s\|, \sigma_{h}\right) \cdot g\left(D(p, s), \sigma_{i}\right)}, \tag{7.2}
\end{equation*}
$$

where $I_{p}$ denotes the luminance intensity of the pixel $p,\|p-s\|$ denotes either the spatial or temporal distance between pixels $p$ and $s, D(p, s)$ is the difference in luminance intensity between pixels $p$ and $s$, and $N_{s}$ is called the kernel and denotes the space of pixels that could contribute to $s$.

### 7.2.2 A Few Observations

Before continuing, we make a few observations about the bilateral filter and image processing. First, notice that if an object is not moving, then applying the temporal bilateral filter over
the same pixel in multiple frames will likely produce better results than the spatial filter, since each pixel on each frame represents the same object.

Second, because there is a base level of noise in every pixel, the signal-to-noise ratio will be smaller for objects with lower luminance than those with higher luminance (i.e., dark pixels are nosier than light pixels). In addition, humans are more capable of perceiving minor differences between pixels with low luminance than those with high luminance (Pappas and Safranek, 2000). As a result (of both the signal-to-noise and human perception issues), pixels with a low luminance require more noise correction than those with high luminance.

Third, while the bilateral filter (and by extension VEC) can trade accuracy for computational intensity by increasing the number of pixels used to correct a single pixel, the exact number of pixels required to correct a single pixel $s$ is a function of both $s$ 's luminance and the similarity of luminance of pixels that surround $s$. For example, in Figure 7.7(a), Pixel $p 13$ will require fewer pixels than $p 19$ since all pixels that are adjacent to $p 13$ have similar luminance values. On the other hand, the pixel $p 19$ has only three adjacent pixels with similar luminance.

### 7.2.3 VEC's Algorithm

VEC consists of two phases (illustrated in Figure 7.8). First, VEC uses both spatial and temporal bilateral filters to correct the noise of each pixel. Second, VEC applies a technique called tone mapping to each noise-reduced frame to change the luminance levels of each pixel into a range that is more palatable for human perception. Since tone mapping requires relatively little processing time in comparison to removing the noise from a frame, and since understanding the process by which VEC applies a tone mapping requires knowledge of image processing techniques, which are beyond the scope of this dissertation, we focus on VEC's noise-removal technique and refer the reader to (Bennett, 2007) for a discussion of tone mapping.

In order to remove the noise from a frame, VEC first calculates the luminance level of each pixel. Next, VEC calculates the gain factor for each pixel. The gain factor for a pixel represents the amount of noise correction that pixel requires. The higher the gain factor,


Figure 7.8: The flow diagram of the VEC.
the more noise correction that pixel requires. As we observed earlier, the lower a pixel's luminance, the more error correction it will require. As a result, pixels with a low luminance will have a high gain factor and pixels with a high luminance will have a low gain factor. We denote the gain factor for the pixel $s$ as $\lambda_{s}$. $\lambda_{s}$ is linearly proportional to the ratio between the output tone-mapped luminance and the input luminance at $s .{ }^{2}$

As we observed in the previous section, some pixels contribute more than others when correcting noise. To formalize this notion, we say that a pixel $p$ when used to correct a pixel $s$ has a vote of $g\left(\|p-s\|, \sigma_{h}\right) \cdot g\left(D(p, s), \sigma_{i}\right)$ (where $g(x, \sigma)$ is as defined in (7.1)). Moreover, the total amount of pixel votes that are required to correct the pixel $s$ equals $\lambda_{s} \cdot g\left(0, \sigma_{h}\right) \cdot g\left(0, \sigma_{i}\right)$. (Notice that, by (7.1), the larger the value of $x$, the smaller the value of $g(x, \sigma)$. Thus, since the vote of a task is defined as $g\left(\|p-s\|, \sigma_{h}\right) \cdot g\left(D(p, s), \sigma_{i}\right)$, the maximal vote of a pixel is $\left.g\left(0, \sigma_{h}\right) \cdot g\left(0, \sigma_{i}\right).\right)$

After the gain factor has been computed for each pixel, VEC then runs a temporal bilateral filter over a predetermined range of frames. If the total votes of all pixels used by the temporal bilateral filter, $\gamma_{s}$, is less than $\lambda_{s} \cdot g\left(0, \sigma_{h}\right) \cdot g\left(0, \sigma_{i}\right)$, then VEC runs a spatial bilateral filter over a circular area around the pixel $s$ with a radius approximately $3 \sigma_{i}$. If VEC runs both spatial and temporal bilateral filters on the pixel $s$, then the luminance value of $s$ is a weighted average of the values produced by the spatial and temporal bilateral filters, where the value of each filter is weighted based on the number of votes its pixels contributed. (Notice that, if the total number of votes is still less than $\lambda_{s} \cdot g\left(0, \sigma_{h}\right) \cdot g\left(0, \sigma_{i}\right)$ after the spatial and temporal filters have been run, then the pixel will still have some noise.)

### 7.2.4 Real-time Characteristics

The most straightforward method for using real-time tasks in VEC is to assign each task a region of each frame to correct, as depicted in Figure 7.9. As a result, since darker objects require more computation than lighter objects to correct, as dark objects move in the video, the processor shares of the tasks assigned to process different areas of the video will

[^15]| $\mathrm{T}_{1}$ | $\mathrm{~T}_{2}$ | $\mathrm{~T}_{3}$ | $\mathrm{~T}_{4}$ |
| :---: | :---: | :---: | :---: |
| $\mathrm{~T}_{5}$ | $\mathrm{~T}_{6}$ | $\mathrm{~T}_{7}$ | $\mathrm{~T}_{8}$ |
| $\mathrm{~T}_{9}$ | $\mathrm{~T}_{10}$ | $\mathrm{~T}_{11}$ | $\mathrm{~T}_{12}$ |
| $\mathrm{~T}_{13}$ | $\mathrm{~T}_{14}$ | $\mathrm{~T}_{15}$ | $\mathrm{~T}_{16}$ |

Figure 7.9: The VEC system divided into real-time tasks.
change. Hence, tasks will need to adjust their weights as quickly as an object can move across the screen. Since VEC continuously performs calculations based on previous frames, it performs best when a substantial amount of "useful" data is stored in cache. As a result, migration/preemption costs in VEC are fairly high. In addition, while strong real-time and fairness guarantees would be desirable in VEC, they are not as important here as in Whisper, because tasks can function more independently in VEC.

### 7.3 LITMUS $^{\text {RT }}$

In this section, we discuss the LITMUS ${ }^{R T}$ testbed. Since LITMUS ${ }^{R T}$ is a joint effort by our entire research group with work that has spanned multiple publications (Block et al., 2008c; Brandenburg et al., 2008; Brandenburg et al., 2007; Brandenburg and Anderson, 2008; Calandrino et al., 2006), an in depth description of it would be outside of the scope of this dissertation. Instead, we provide a brief overview of LITMUSRT and refer the reader to the aforementioned papers for a more detailed discussion.

LITMUS ${ }^{\text {RT }}$ is an extension of Linux that supports a variety of real-time multiprocessor scheduling policies. In its current state, it is most useful as a testbed within which different scheduling policies can be implemented and empirically evaluated. LITMUS ${ }^{\mathrm{RT}}$ is designed in such a way that adding support for additional scheduling policies is straightforward.

LITMUS ${ }^{\text {RT }}$ was implemented by modifying the Linux 2.6 .20 kernel ${ }^{3}$ configured to run on a symmetric multiprocessor (SMP) architecture. Our particular development platform is an SMP consisting of four 32 -bit $\operatorname{Intel}(\mathrm{R})$ Xeon(TM) processors running at 2.70 GHz , with 8 K instruction and data caches, and a unified 512K L2 cache per processor, and 2 GB of main memory.

Why provide real-time support in Linux? We chose to create our testbed by modifying Linux instead of an existing real-time operating system (RTOS) for two reasons. First, Linux is free, open-source software that is easy to obtain and modify, and is widely accepted by both developers and end users. Second, the potential client base for LITMUS ${ }^{\text {RT }}$ as it evolves will include many real-time graphics and multimedia applications developed within our own department. The developers of those applications actually prefer Linux as a development platform.

We acknowledge that producing system designs in any Linux-based system in which realtime correctness is guaranteed with certainty is not feasible. Therefore, we expect systems to be provisioned in LITMUS ${ }^{\text {RT }}$ using experimentally-determined worst-case (average-case) values for execution costs and system overheads in the hard (soft) real-time case, instead of using analytically-determined, verified values. Thus, in LITMUS ${ }^{\text {RT }}$, the term "hard realtime" should really be interpreted to mean that deadlines are almost never missed, and "soft real-time" to mean that deadline tardiness almost always remains within some bound, even if individual tasks misbehave. These are stronger guarantees than provided by most real-time Linux variants in commercial use today.

### 7.3.1 The Design of LITMUS ${ }^{\text {RT }}$

LITMUS ${ }^{\text {RT }}$ has been implemented via changes to the Linux kernel and the creation of userspace libraries. Since LITMUS ${ }^{\text {RT }}$ is concerned with real-time scheduling, most kernel changes affect the scheduler and timer interrupt code. The kernel modifications can be split into roughly three components. The core infrastructure consists of modifications to the Linux

[^16]scheduler, as well as support structures and services such as tracing and sorted run queues that can be used by scheduler plugins. The scheduler plugins encapsulate the available real-time scheduling algorithms by providing functions that implement the methods of the scheduler plugin interface. Finally, a collection of system calls provides a user-space API for real-time tasks to interact with the kernel. In the following subsections, we describe each component in turn.

Note that, in the discussion that follows, the term real-time task means tasks that are scheduled by LITMUS ${ }^{\text {RT }}$. Normal Linux tasks that run with a static priority from the "POSIX real-time range" are not considered to be real-time tasks in LITMUS ${ }^{\text {RT }}$. Since they do not follow the sporadic task model, they are considered to be just best-effort tasks with a high static priority.

### 7.3.2 Core Infrastructure

To facilitate the releasing and queuing of real-time tasks, LITMUS ${ }^{\text {RT }}$ provides the abstraction of a real-time domain, which consists of a ready queue and a release queue (as well as one lock per queue). When a real-time domain is instantiated, it is parametrized with an order function that is used to sort tasks in the ready queue (the release queue is ordered by ascending release time). Wrapper functions are provided in the real-time domain for operations such as queuing, dequeuing, and inspecting designated queue elements. This removes the need for list-handling in most scheduler plugins, thereby reducing development effort (and also removing a common source of bugs).

Scheduling quanta are defined to be the intervals between local timer interrupts. To realize aligned quanta, LITMUS ${ }^{R T}$ synchronizes timer interrupts during boot across all processors. This is done by having each processor disable its local timer within the local timer interrupt handler, enter a barrier, and restart its timer immediately afterward. When all processors reach the barrier, they will be simultaneously released, resulting in all processors restarting their timers at approximately the same time. Using this method, we have been able to achieve aligned quanta with an error of at most $10 \mu s$ on our test platform - in some cases, error is as low as $1-2 \mu s$.

The core LITMUS ${ }^{\text {RT }}$ infrastructure also includes an implementation of the MCS queue lock (Mellor-Crummey and Scott, 1991). Ideally, deterministic locking primitives should be used throughout the kernel. Unfortunately, the Linux kernel uses non-FIFO spin locks, and it is not currently feasible to replace all kernel spin locks with queue locks. Thus, we must be aware of their potential impact on the real-time guarantees that can be made.

At the heart of LITMUS ${ }^{R T}$, the core infrastructure is also responsible for interfacing with the rest of Linux. It initializes a real-time scheduler plugin (based on a kernel command-line parameter) during system boot. To pass control to the plugin, it hooks into the Linux scheduler_tick() and schedule() functions. Overriding the Linux scheduler works as follows. Real-time tasks are assigned the highest static Linux scheduling priority upon creation. However, they are not kept in the standard Linux run queues. Instead each plugin is responsible for managing its own run queue. (Similarly, time-slice management is also delegated to plugins for real-time tasks.) When schedule() is invoked, control is passed to the current scheduler plugin. If it selects a real-time task to be scheduled on the local processor, then the task is inserted into the run queue and the Linux scheduler is bypassed. When a real-time task is preempted, it is removed again from the run queue, thereby taking it out of the reach of the Linux scheduler.

LITMUS ${ }^{\text {RT }}$ has two modes of operation, real-time and non-real-time. When started, the system is initially in non-real-time mode. Real-time tasks are not scheduled as long as the system is in non-real-time mode. This feature allows complete task systems to be set up before they are scheduled, thereby allowing for the synchronous release of the first jobs of all tasks.

### 7.3.3 Scheduler Plugins

As mentioned before, real-time scheduling policies are implemented as scheduler plugins. Such plugins are realized similarly to other pluggable components in Linux such as file systems. To create a scheduler plugin, functions that realize several methods ${ }^{4}$ of the plugin interface must be implemented and registered to the LITMUS ${ }^{R T}$ core. These methods include adding

[^17]a task to the set, "tearing-down" a task, and scheduling a task. ${ }^{5}$

### 7.3.4 System Call API

LITMUS ${ }^{R T}$ introduces a number of new system calls to Linux. While some of these system calls can be used directly, most of them are intended to be used by liblitmus, a user-space library that provides higher-level abstractions. The introduced system calls are organized by purpose into five groups: managing real-time tasks, querying state information, controlling job releases, system setup, and synchronization. ${ }^{6}$ Real-time management APIs handle setting up a task and adding it to the task set. State information APIs are used to query a task about their real-time characteristics, e.g., WCET and period. Job control APIs are invoked when a job completes. System setup APIs are used to configure scheduler-specific settings. Finally, synchronization APIs are used to synchronize data across tasks. It is important to note that, currently, synchronization APIs are only implemented for non-adaptive tasks.

### 7.4 Comparison

In this section, we discuss the first set of experiments we conducted, in which we evaluated the performance of the adaptive algorithms proposed in Chapters 3, 4, and 5 when scheduling Whisper- and VEC-like tasks on a simulated four-processor system. While these experiments are just simulations, most of the parameters used here were obtained by implementing and timing the scheduling algorithms discussed in this dissertation and some of the signal-processing and video-enhancement code in Whisper and VEC, respectively, on a real multiprocessor testbed. Thus, the behaviors in these simulations should fairly accurately reflect what one would see in a real Whisper or VEC implementation.

For both Whisper and VEC, the simulated platform was assumed to be a shared-memory multiprocessor, with four $2.7-\mathrm{GHz}$ processors and a 1-ms quantum, like our test platform. All simulations were run 100 times. Both systems were simulated for 10 secs. (Note that decreasing and increasing the simulation time gives similar results.) We implemented and

[^18]timed each scheduling scheme considered in our simulations on our test platform and found that all scheduling and reweighting computations could be completed within $5 \mu \mathrm{~s}$. We considered this value to be negligible in comparison to a $1-\mathrm{ms}$ quantum and thus did not consider scheduling overheads in our simulations. For both Whisper and VEC, we conducted two types of experiments: (i) all preemption and migration costs were the same and corresponded to a loss of cache affinity; and (ii) the preemption cost was set to a fixed value and the migration cost was varied. If a task was preempted and then migrated, we assumed that it incurred the maximum of the two costs. Based on measurements taken on our testbed system, we estimated Whisper's migration/preemption cost as $2 \mu \mathrm{~s}-10 \mu \mathrm{~s}$, and VEC's as $50 \mu \mathrm{~s}-60 \mu \mathrm{~s}$. While we believe that these costs may be typical for a wide range of systems, in our experiments we varied the preemption/migration cost over a slightly larger range. (It is worth noting that, in related work by our research group (Calandrino et al., 2006), the average-case preemption (migration) cost for a 4 KB working set size, i.e., applications that randomly access 4 KB of data, was found to be $15.70 \mu \mathrm{~s}(15.80 \mu \mathrm{~s})$, and the worst-case preemption (migration) cost for a 4 KB working set size as $42.00 \mu \mathrm{~s}(44.00 \mu \mathrm{~s})$. This research was conducted using the same test platform as considered here.) For all experiments, the maximum execution time was 7 ms for PEDF and NP-PEDF and 5 ms for GEDF and NP-GEDF. These values were determined by profiling each system beforehand to determine the "best" compromise of accuracy and performance.

While the ultimate metric for determining the efficacy of both systems would be user perception, such a metric would require a full implementation of both systems (which as we discussed at the beginning of this chapter is not currently feasible). Therefore, we compared each of the tested schemes by comparing allocations against each algorithm's respective notion of an IDEAL schedule. In particular, we measured both the "average under-allocation" and "fairness factor" for each task set at the end of each simulation (i.e., 10 secs.). The average under-allocation (UA) is the average amount each task is behind its IDEAL allocation (this value is defined to be nonnegative, i.e., for a task that is not behind its IDEAL, this value is zero). The fairness factor (FF) of a task set is the largest deviance from the allocations in IDEAL between any two tasks (e.g., if a system has three tasks, one that deviates from its

IDEAL allocation by -10 , another by 20 , and the third by 50 , then the FF is $50-(-10)=60$ ). The FF is a good indication of how fairly a scheme allocates processing capacity. A lower FF means the system is more fair. For applications like Whisper, where the output generated by multiple tasks is periodically combined, a low FF is important, since if any one task is "behind," then performance of the entire system is impacted; however, for applications like VEC, where tasks are more independent, a high FF does not affect the system performance nearly as much. These metrics should provide us with a reasonable impression of how well the tested schemes will perform when Whisper and VEC are fully re-implemented.

### 7.4.1 Whisper Experiments

In our Whisper experiments, we simulated three speakers (one per object) revolving around pole in a $1 \mathrm{~m} \times 1 \mathrm{~m}$ room with a microphone in each corner, as shown in Figure 7.10-the results of these simulations appear in Figure 7.11 and Figure 7.12. The pole creates potential occlusions. One task is required for each speaker/microphone pair, for a total of 12 tasks. In each simulation, the speakers were evenly distributed around the pole at an equal distance from the pole, and rotated around the pole at the same speed. The starting position for each speaker was set randomly. As mentioned in Section 7.1.1, as the distance between a speaker and microphone changes, so does the amount of computation necessary to correctly track the speaker. This distance is (obviously) impacted by a speaker's movement, but is also lengthened when an occlusion is caused by the pole. The range of weights of each task was determined (as a function of a tracked object's position) by implementing and timing the basic computation of the correlation algorithm (an accumulate-and-multiply operation) on our testbed system.

In the Whisper simulations, we made several simplifying assumptions. First, all objects are moving in only two dimensions. Second, there is no ambient noise in the room. Third, no speaker can interfere with any other speaker. Fourth, all objects move at a constant rate. Fifth, the weight of each task changes only once for every 5 cm of distance between its associated speaker and microphone. Sixth, all speakers and microphones are omnidirectional. Finally, all tasks have a minimum weight based on measurements from our testbed system


Figure 7.10: The simulated Whisper system.
and a maximum weight of 1.0. A task's current weight at any time lies between these two extremes and depends on the corresponding speaker's current position. Even with theses assumptions, frequent share adaptations are required.

We conducted Whisper experiments in which the tracked objects were sampled at a rate of $1,000 \mathrm{~Hz}$, the distance of each object from the room's center was set at 50 cm , the speed of each object was set at $5 \mathrm{~m} / \mathrm{sec}$. (this is within the speed of human motion), and the maximum execution cost, migration, and preemption cost were varied.

The first set of graphs in Figure 7.11 and Figure 7.12 show the result of the Whisper simulations conducted to compare PD ${ }^{2}$, PEDF, NP-PEDF, GEDF, and NP-GEDF. Figure 7.11 depicts the average UA and FF, respectively, for each scheme, where the preemption cost is varied from 0 to $100 \mu$ s and the migration cost equals the preemption cost. Figure 7.12 depicts the average UA and FF, respectively, for each scheme, where the preemption cost is set at $10 \mu$ s (the maximum expected preemption cost for Whisper) and the migration cost is varied from 0 to $100 \mu \mathrm{~s}$. There are five things worth noting here. First, when the preemption/migration cost is varied over the range 2 to $10 \mu$ s (the expected range for Whisper, as noted on each graph), the UA is about the same for all schemes (Figure 7.11(a)); however, $\mathrm{PD}^{2}$ has the best FF (Figure 7.11(b)). Second, while GEDF and NP-GEDF do not have the best UA for the expected preemption/migration costs for Whisper, for higher preemption/migration costs, i.e., preemption/migration costs larger than $10 \mu \mathrm{~s}$, GEDF and NP-GEDF both have a substantially better UA than $\mathrm{PD}^{2}$ and better FF than either PEDF or NP-PEDF. Third, as the migration cost (but not preemption cost) of a task increases, the UA of PEDF and NP-PEDF increases slowly (Figure 7.12(a)). However the performance of the other three schemes decays quickly.


Figure 7.11: (a) The average UA and (b) FF for Whisper as a function of preemption/migration cost, as scheduled by each tested algorithm. The key in each graph is in the order that the schemes appear in that graph at $100 \mu$ s. Standard deviations are shown. Note that, in (b), GEDF and NP-GEDF are indistinguishable from each other.


Figure 7.12: (a) The average UA and (b) FF for Whisper as a function of migration cost (preemption cost is fixed at $10 \mu \mathrm{~s}$ ), as scheduled by each tested algorithm. The key in each graph is in the order that the schemes appear in that graph at $100 \mu \mathrm{~s}$. Standard deviations are shown. Note that, in (b), GEDF and NP-GEDF are indistinguishable from each other.


Figure 7.13: The simulated VEC system.

Fourth, standard deviations of the FF for GEDF, NP-GEDF, and PD ${ }^{2}$ are smaller than for PEDF and NP-PEDF, since GEDF, NP-GEDF, and PD ${ }^{2}$ have better accuracy. Fifth, as seen in Figure 7.12, PD $^{2}$ 's and GEDF's UA and FF do not appreciably increase until the migration cost exceeds $10 \mu \mathrm{~s}$. This is because, until the migration cost is $10 \mu \mathrm{~s}, \mathrm{PD}^{2}$ and GEDF incur the maximum of the migration or preemption cost, which is $10 \mu \mathrm{~s}$.

### 7.4.2 VEC Experiments

In our VEC experiments, we simulated a $640 \times 640$-pixel video feed where a grey square that is $160 \times 160$ pixels moves around in a circle with a radius of 160 pixels on a white background. This is illustrated in Figure 7.13. The grey square makes one complete rotation every ten seconds. The position of the grey square on the circle is random. Each frame is divided into sixteen $160 \times 160$-pixel regions; each of these regions is corrected by a different task. A task's weight is determined by whether the grey square covers its region. By analyzing VEC's code, we determined that the grey square takes three times more processing time to correct than the white background. Hence, if the grey square completely covers a task's region, then its weight is three times larger than that of a task with an all-white region. The video is shot at a rate of 25 frames per second, and as a result, each frame has an exposure time of 40 ms .

Figures 7.14 and 7.15 show the results of the VEC simulations conducted to compare the five tested scheduling algorithms. Figure 7.14 depicts the average UA and FF, for each scheme, where the preemption cost is varied from 0 to $100 \mu \mathrm{~s}$ and the migration cost equals the preemption cost. Figure 7.15 depicts the average UA and FF for each scheme, where the
preemption cost is set at $60 \mu \mathrm{~s}$ (the maximum expected preemption cost for VEC) and the migration cost is varied from 0 to $100 \mu$ s. There are two things worth noting here. First, when the preemption/migration cost is varied over the range 50 to $60 \mu \mathrm{~s}$ (the expected range for VEC, as noted on each graph), NP-PEDF and PEDF have the smallest UA (Figure 7.14(a)); however, GEDF and NP-GEDF both have a UA that is competitive with both PEDF and NP-PEDF (Figure 7.14(a)) and have a substantially smaller FF (Figure 7.14(b)). Second, as seen in Figure 7.15, $\mathrm{PD}^{2}$ 's and GEDF's UA and FF do not appreciably increase until the migration cost equals $60 \mu \mathrm{~s}$. This occurs for the same reason that $\mathrm{PD}^{2}$ and GEDF did not noticeably increase until $10 \mu$ s in Figure 7.12.

### 7.5 AGEDF Implementation and Evaluation

In this section, we first describe our implementation of AGEDF in the LITMUS ${ }^{\mathrm{RT}}$ framework. Next, we describe the experiments that we used to evaluate our implementation.

### 7.5.1 Implementation

Because LITMUS ${ }^{\text {RT }}$ was designed for sporadic tasks provisioned using WCETs, several modifications were needed to support adaptable sporadic tasks. These included: adjusting the internal structure of the task control block to allow each task to have multiple service levels; disabling the enforcement of WCETs to allow tasks to overrun their expected allocation; and modifying LITMUS ${ }^{\mathrm{RT}}$ to allow task statistics such as actual execution times to be gathered.

After making these changes, we implemented the AGEDF framework by changing the GEDF scheduling algorithm (which had already been implemented in LITMUS ${ }^{\mathrm{RT}}$ ) in two ways. First, we introduced a system call to query the kernel in order for a task to determine its current service level. Second, we implemented the feedback, optimization, and reweighting components in kernel space. Since in Linux floating point operations cannot be used in kernel space, these components were implemented using fixed-point calculations instead.


Figure 7.14: (a) The average UA and (b) FF for VEC as a function of preemption/migration cost, as scheduled by each tested algorithm. The key in each graph is in the order that the schemes appear in that graph at $100 \mu \mathrm{~s}$. Standard deviations are shown. Note that, in (b), GEDF and NP-GEDF are indistinguishable from each other.


Figure 7.15: (a) The average UA and (b) FF for VEC as a function of migration cost (preemption cost is fixed at $60 \mu \mathrm{~s}$ ), as scheduled by each tested algorithm. The key in each graph is in the order that the schemes appear in that graph at $100 \mu \mathrm{~s}$. Standard deviations are shown. Note that, in (b), GEDF and NP-GEDF are indistinguishable from each other.

### 7.5.2 Evaluation

The development platform used in our experiments is an SMP consisting of four 32-bit $\operatorname{Intel}(\mathrm{R})$ Xeon(TM) processors running at 2.7 GHz , with 8 K L1 instruction and data caches, and a unified 512 K L2 cache per processor, and 2 GB of main memory (this is the same system that was used for the experiments in Section 7.4). For each task, each job was implemented as a loop in which the core operations of Whisper and VEC are performed iteratively. The exact manner in which jobs behave is discussed below. In implementing the optimizing component, we assumed that there is a linear relationship between importance value and estimated weight, and attempted to maximize the total importance value for all tasks via the highest-value-density-first rule (described in Section 6.2.2). The optimizer was configured to run at least once every second, and also whenever the estimated weight of a task changed by at least $50 \%$, or upon a job completion, the total estimated system weight exceeded four. However, it was constrained to run at most once every 200 ms . Note that, in full Whisper and VEC implementations, these choices could possibly be improved upon by carefully considering human-factors issues of relevance to virtual-reality or night-vision systems.

In all experiments, we defined the PI controller using $a=0.102$ and $c=-1.975$ (see Section 6.2.1 for a discussion of the feedback characteristics associated with these values). We chose these values because we believe that they represent a good tradeoff between transient response and steady-state error (for a ramp input).

Whisper experiments. In our Whisper experiments, we simulated three speakers (one per object) revolving at a speed of $2 \mathrm{~m} / \mathrm{s}$ (this is within the speed of human motion) in a 10 m $\times 10 \mathrm{~m}$ room with a microphone in each corner, as shown in Figure 7.16. Tracked objects were sampled at a rate of $2,000 \mathrm{kHz}$, and the distance of each object from the room's center was set at 5 m . While this test scenario may seem simple (since the path of each object is simple and pre-determined), it is actually a challenging test case for Whisper. This is because objects moving at a relatively high speed of $2 \mathrm{~m} / \mathrm{s}$ require significant computational resources to track. Moreover, while it is possible to simulate objects that start, stop, and change directions, such scenarios actually require less computational resources and adapt


Figure 7.16: The simulated Whisper system.
task service levels less frequently, because user motion is typically slower when motion is not continuous.

In the above scenario, one task is required per speaker/microphone pair, for a total of 12 tasks. Each task was configured to have three service levels, with periods/importance values of $66 \mathrm{~ms} / 0.25,33 \mathrm{~ms} / 0.5,22 \mathrm{~ms} / 0.75$, respectively, and $\mathrm{g}\left(T_{i}, e, 1,2\right) \approx 5 e$ and $\mathrm{g}\left(T_{i}, e, 1,3\right) \approx 9 e$ ( $\mathrm{g}\left(T_{i}, e, \ell_{1}, \ell_{2}\right)$ is defined in Section 6.1). The importance values were selected somewhat arbitrarily after some trial-and-error experimentation; in an actual deployment, user studies would be required to assess the impact of different settings. Since the weight/importance value relationship is linear, we used the approach in Section 6.2.2 to optimize the system. The other parameters were selected based upon the existing Whisper implementation. As we discussed in Section 7.1, in Whisper, the QoS provided is directly related to the number of correlation computations (CCs) performed per second. When the signal-to-noise ratio decreases, the number of CCs must be increased to maintain the same QoS. Similarly, the QoS provided can be increased by increasing the number of CCs per second. A change in the functional service level of a Whisper task changes the number of CCs per second. We estimated that the existing Whisper implementation, if implemented on our test platform, would perform approximately $27,600,000$ CCs per second in the average case. The task periods and $\mathrm{g}\left(T_{i}, e, \ell_{1}, \ell_{2}\right)$ values given above were defined so that the average number of CCs per second for the second service level matches this rate. Note that, because the code segments of the three service levels differ only in the number of CCs performed, the code segment of an active job can be changed.

The first experiment we discuss was conducted to see if adaptivity is even needed in implementing Whisper. In this experiment, we ran each of the twelve speaker/microphone
pair tasks individually as a normal Linux task for 20 seconds at all three service levels and measured their actual weight. The results for one of the twelve tasks is shown in Figure 7.17 (the other eleven tasks have a similar behavior). After 5.5 and 12.3 seconds, the system experiences 480 ms of noise that doubles the number of correlation computations required per job. The average weight of the task at the first, second, and third service level is, respectively, $0.05,0.25$, and 0.45 . Notice that, for a system with 12 tasks, operating each task at its highest service level and allocating it a processor share based on its worst-case weight (which is 1.0) gives a total actual weight of 12 , which substantially over-utilizes the system. Even with an average-case provisioning, the system is still over-utilized, as the total actual weight is 5.4 in this case. On the other hand, configuring each task to run at its second service level using its average-case weight gives a total actual weight of 3.0 , which does not over-utilize the system; however, using a constant average-case allocation would likely cause the system to be over-utilized when noise is encountered. Thus, from this experiment, we can infer that, in order for Whisper to schedule tasks at any service level higher than the lowest one, adaptive scheduling is needed (as is a multiprocessor).

In the second experiment, we ran all 12 Whisper tasks on LITMUS ${ }^{R T}$, scheduled by AGEDF, for 20 seconds with 480 ms bursts of ambient noise after 5.5 and 12.3 seconds that double the number of CCs required to maintain the same QoS. Figure 7.18 shows the total actual weight and the total importance value of the system as a function of time. Figures 7.197.24 depict the actual and estimated weights and error (defined as the difference between the actual and estimated weight) for all twelve tasks as a function of time. They also show the functional service level for each task as a function of time. There are several interesting things to notice about these graphs. First, for the tasks depicted in Figure 7.19-7.24, error is typically within the range $[-0.05,0.05]$. Second, whenever the functional service level changes or the system encounters noise, error briefly spikes but quickly falls back within the range [-0.05, 0.05]. Third, when the task in Figure 7.19(a) encounters noise at time 5.5, its service level is changed and there is a substantial drop in its weight; however, when it encounters noise at time 12.3, its service level is not decreased because its actual weight is so low. Note that its low weight at this time is coincidental: its weight varies between times 8 and 20


Figure 7.17: The actual weight of a Whisper task at three different service levels over a 20-second run with two bursts of noise at approximately times 6 and 13.
as depicted because of the movement of the corresponding object, and that object happens to be closest to the microphone for which this task is defined at approximately time 14. Fourth, when a job of the task in Figure 7.19(b) completes after the noise at time 12.3, the total estimated weight is greater than four, so the optimizer is invoked causing this task to decrease its service level. This is why the actual weight of this task is briefly greater than one. Fifth, the total utilization of the system is typically close to four, and the system is briefly over-utilized when noise is encountered. Because the total actual weight is always close to four, this system would not be schedulable using a partitioning approach. Sixth, the total importance value of the system is typically in the range [7.5, 8] and drops below 6.0 only when noise is encountered. In contrast, if tasks were statically assigned their second service level and scheduled by GEDF, then the total importance value would never exceed 6.0.


Figure 7.18: Total actual weight and importance value as a function of time, when executing 12 Whisper tasks for 20 seconds.

VEC experiments. In the VEC experiments, we considered a $320 \times 320$-pixel video feed where a grey square that is $80 \times 80$ pixels moves around in a circle with a radius of 80 pixels on a white background, as illustrated in Figure 7.25. Each frame is divided into sixteen $80 \times 80$ pixel regions (as depicted in Figure 7.9); each of these regions is corrected by a different task. The amount of execution time a task requires is determined by whether the grey square covers its region. We assumed that the grey square is sufficiently dark that it takes three times more processing time to correct than the white background. Hence, if the grey square completely covers a task's region, then its weight is three times larger than that of a task with an allwhite region. Moreover, occasionally the video briefly becomes dark, doubling the execution time for each job. For each experiment, we considered two different sets of three service levels-one in which periods change and one in which code segments change (by varying the number of iterations described earlier in Section 6.2.3). The periods/importance levels were set at $66 \mathrm{~ms} / 0.25,33 \mathrm{~ms} / 0.5$, and $22 \mathrm{~ms} / 0.75$, respectively. Furthermore, we do not allow the code segment of an active job to change. Also, $\mathrm{g}\left(T_{i}, q, 1,2\right)=2 q$ and $\mathrm{g}\left(T_{i}, q, 1,3\right)=3 q$


Figure 7.19: Results from executing 12 Whisper tasks for 20 seconds. Actual and estimated weights and error for (a) $T_{1}$ and (b) $T_{2}$ as a function of time. The service level for each task is depicted across the top of each inset. In both insets, the error line is centered around zero, and the actual and estimated lines are often indistinguishable.


Figure 7.20: Results from executing 12 Whisper tasks for 20 seconds. Actual and estimated weights and error for (a) $T_{3}$ and (b) $T_{4}$ as a function of time. The service level for each task is depicted across the top of each inset. In both insets, the error line is centered around zero, and the actual and estimated lines are often indistinguishable.


Figure 7.21: Results from executing 12 Whisper tasks for 20 seconds. Actual and estimated weights and error for (a) $T_{5}$ and (b) $T_{6}$ as a function of time. The service level for each task is depicted across the top of each inset. In both insets, the error line is centered around zero, and the actual and estimated lines are often indistinguishable.


Figure 7.22: Results from executing 12 Whisper tasks for 20 seconds. Actual and estimated weights and error for (a) $T_{7}$ and (b) $T_{8}$ as a function of time. The service level for each task is depicted across the top of each inset. In both insets, the error line is centered around zero, and the actual and estimated lines are often indistinguishable.


Figure 7.23: Results from executing 12 Whisper tasks for 20 seconds. Actual and estimated weights and error for (a) $T_{9}$ and (b) $T_{10}$ as a function of time. The service level for each task is depicted across the top of each inset. In both insets, the error line is centered around zero, and the actual and estimated lines are often indistinguishable.


Figure 7.24: Results from executing 12 Whisper tasks for 20 seconds. Actual and estimated weights and error for (a) $T_{11}$ and (b) $T_{12}$ as a function of time. The service level for each task is depicted across the top of each inset. In both insets, the error line is centered around zero, and the actual and estimated lines are often indistinguishable.


Figure 7.25: The simulated VEC system.
( $\mathrm{g}\left(T_{i}, q, \ell_{1}, \ell_{2}\right)$ is defined in Section 6.1).
Video is shot at a rate of 30 frames per second, and as a result, each frame has an exposure time of approximately 33 ms . In the experiment that we report on here, we ran VEC for 20 seconds, the dark square makes one revolution approximately every 2.2 seconds, and after 8.9 seconds, the system encountered $1,100 \mathrm{~ms}$ of darkness that doubled the number of computations required by each task.

In the first set of VEC experiments, we determined if adaptivity is needed in implementing VEC. In this experiment, we ran each of the sixteen VEC tasks individually as normal Linux tasks for 20 seconds at all three service levels and measured their actual weight. The results for $T_{1}, T_{2}$, and $T_{6}$ are given in Figures $7.26(\mathrm{a}), 7.26(\mathrm{~b})$, and 7.27 , respectively. (As seen in Figure 7.9, the corner tasks, $T_{4}, T_{13}$, and $T_{16}$, have the same behavior as $T_{1}$; the side tasks, $T_{3}, T_{5}, T_{8}, T_{9}, T_{12}, T_{14}$, and $T_{15}$, have the same behavior as $T_{2}$; and the center tasks, $T_{7}, T_{10}$, and $T_{11}$, have the same behavior as $T_{6}$.) Notice that, if tasks are statically assigned their maximum weight and either their second or third service level, then the total weight of all tasks would exceed four. On the other hand, if tasks are statically assigned their average weight, then the total weight of all tasks at their highest service level would not exceed four (the corner, side, and center tasks in this case are assigned weights that are approximately, $0.17,0.19$, and 0.22 , respectively). However, in this case, each task would be incapable of correcting its associated region when either the grey square was present or the system incurred darkness. Thus, from this experiment, we can infer that, in order for VEC to fully utilize the system, adaptive scheduling is needed (as is a multiprocessor).


Figure 7.26: The actual weight of a VEC task at three different service levels. (a) $T_{1}$. (b) $T_{2}$.


Figure 7.27: The actual weight of a VEC task, $T_{6}$, at three different service levels.

In the second set of experiments, we ran all sixteen tasks as real-time tasks in LITMUST ${ }^{\text {RT }}$ at the same time. Figure 7.28 shows the total actual weight and total importance value of the system as a function of time. Figures 7.29-7.36 depicts the actual and estimated weights and error for all of the tasks as a function of time. (As before, the mapping of tasks to regions of the screen is as given in Figure 7.9. Thus, $T_{1}$ corrects the upper-left corner of the screen, and $T_{16}$ corrects the lower-right corner of the screen.) There are a few interesting things to notice about these graphs. First, the error of the task in Figures 7.29-7.36 is typically within the range $[-0.08,0.08]$, except when either the system starts or a major change to the system occurs, i.e., darkness. Second, the total importance value of the system has less jitter than for Whisper. This is because in our experimental set-up, fewer tasks are changing at any given point in time. Third, the weights of tasks correcting the corners of the screen (i.e., $T_{1}$, $T_{4}, T_{13}$, and $T_{16}$ ) do not change that much (except when darkness is incurred). The reason why this behavior occurs is because the dark square barley enters the four corners. On the other hand, the weights of the center four tasks (i.e., $T_{6}, T_{7}, T_{10}$, and $T_{11}$ ) have substantial


Figure 7.28: Results from executing 16 VEC tasks for 20 seconds. Total actual weight and importance value as a function of time.
variance. This behavior occurs because the dark square frequently covers a large fraction of the center four regions.

One final note. Notice that, in the above evaluations of AGEDF, we used real code from both Whisper and VEC in simulations of deterministic scenarios. Since we are primarily interested in the behavior of AGEDF rather than the behavior of either Whisper or VEC, we have chosen to simulate simple environments. These simulations (elements moving around in a circle with occasional bursts of noise or darkness) provide a sufficient evaluation of AGEDF since they allow us to measure its behavior when scheduling systems that require a substantial amount of adaption (in both systems, tasks continually change their execution time and there are occasional periods of stress). As a result, if we were to simulate more complex scenarios (objects moving in random or different deterministic patterns), then the results would be similar. In addition, by evaluating simple scenarios, we can more easily discern the efficacy of AGEDF when scheduling either Whisper or VEC.


Figure 7.29: Results from executing 16 VEC tasks for 20 seconds. Actual and estimated weights and error for (a) $T_{1}$ and (b) $T_{2}$ as a function of time. The service level for each task is depicted across the top of each inset. In both insets, the error line is centered around zero, and the actual and estimated lines are often indistinguishable.


Figure 7.30: Results from executing 16 VEC tasks for 20 seconds. Actual and estimated weights and error for (a) $T_{3}$ and (b) $T_{4}$ as a function of time. The service level for each task is depicted across the top of each inset. In both insets, the error line is centered around zero, and the actual and estimated lines are often indistinguishable.


Figure 7.31: Results from executing 16 VEC tasks for 20 seconds. Actual and estimated weights and error for (a) $T_{5}$ and (b) $T_{6}$ as a function of time. The service level for each task is depicted across the top of each inset. In both insets, the error line is centered around zero, and the actual and estimated lines are often indistinguishable.


Figure 7.32: Results from executing 16 VEC tasks for 20 seconds. Actual and estimated weights and error for (a) $T_{7}$ and (b) $T_{8}$ as a function of time. The service level for each task is depicted across the top of each inset. In both insets, the error line is centered around zero, and the actual and estimated lines are often indistinguishable.


Figure 7.33: Results from executing 16 VEC tasks for 20 seconds. Actual and estimated weights and error for (a) $T_{9}$ and (b) $T_{10}$ as a function of time. The service level for each task is depicted across the top of each inset. In both insets, the error line is centered around zero, and the actual and estimated lines are often indistinguishable.


Figure 7.34: Results from executing 16 VEC tasks for 20 seconds. Actual and estimated weights and error for (a) $T_{11}$ and (b) $T_{12}$ as a function of time. The service level for each task is depicted across the top of each inset. In both insets, the error line is centered around zero, and the actual and estimated lines are often indistinguishable.


Figure 7.35: Results from executing 16 VEC tasks for 20 seconds. Actual and estimated weights and error for (a) $T_{13}$ and (b) $T_{14}$ as a function of time. The service level for each task is depicted across the top of each inset. In both insets, the error line is centered around zero, and the actual and estimated lines are often indistinguishable.


Figure 7.36: Results from executing 16 VEC tasks for 20 seconds. Actual and estimated weights and error for (a) $T_{15}$ and (b) $T_{16}$ as a function of time. The service level for each task is depicted across the top of each inset. In both insets, the error line is centered around zero, and the actual and estimated lines are often indistinguishable.

### 7.6 Conclusion

In this section, we presented two sets of experiments. First, we presented a simulation-based comparison of our adaptive variants of GEDF, NP-GEDF, PEDF, NP-PEDF, and PD ${ }^{2}$. Second, we presented an implementation and evaluation of our AGEDF framework using LITMUS ${ }^{\text {RT }}$. The results of our simulation-based comparison suggest the following: first, when it is critical that every task make its deadline and migration/preemption costs are low (i.e., systems like Whisper), $\mathrm{PD}^{2}$ is the best choice. Second, when preemption/migration costs are high (i.e., either Whisper or VEC as implemented on a system where the processors are not as tightly integrated), average case performance is of the utmost importance, and fairness and timeliness are less important, then either PEDF or NP-PEDF may be the best choice. Third, when migration/preemption costs are high and a good mix of average-case performance and fairness factors is beneficial (i.e., systems like VEC), then either GEDF or NP-GEDF may the best choice. In addition, our evaluation of the AGEDF framework shows that it is capable of enacting needed adaptions in a way that enhances overall QoS for both Whisper and VEC.

## CHAPTER 8

## CONCLUSION AND FUTURE WORK

In research on real-time systems, multiprocessor platforms are of growing importance, due to both hardware trends such as the emergence of multicore technologies and the prevalence of computationally-intensive applications for which single-processor designs are insufficient. While research on real-time systems has traditionally focused on applications with static timing constraints, for many applications, such constraints can change at run time. For these applications, adaptive real-time scheduling techniques are needed. This dissertation has focused on developing multiprocessor adaptive real-time scheduling techniques and studying the usefulness of such techniques when developing computationally-intensive multimedia applications such as human-tracking and night-vision systems.

### 8.1 Summary of Results

In this dissertation, we examined the thesis that multiprocessor real-time scheduling algorithms can be made more adaptive by allowing tasks to reweight between job releases. Feedback and optimization techniques can be used to determine at run time which reweighting events are needed. The accuracy of such an algorithm can be improved by allowing more frequent task migrations and preemptions; however, this accuracy comes at the expense of higher migration and preemption costs, which impacts average-case performance. Thus, there is a tradeoff between accuracy and average-case performance that will be dependent on the frequency of task migrations/preemptions and their cost.

The main difficulty in constructing an adaptable system is that it is often necessary to delay enacting weight changes that have been initiated. (If weight changes are always immediately enacted, then it is possible for a task to artificially boost its weight, possibly
causing over-utilization, by aggressively requesting weight changes that, if delayed appropriately, would not cause over-utilization.) As a result, fundamentally, whenever a task's weight changes, there may be some "loss" to the system relative to an "ideal" system in which each weight change can be enacted as soon it is initiated. The allocation difference between the ideal and actual systems that arises for a single task as a result of one weight change is called drift. Before the research presented in this dissertation, the only available method for changing the weight of a task on a multiprocessor was for the task to "leave" with its old weight and "rejoin" with its new weight. The problem with this method is that drift can be arbitrarily large. (This is because a task is permitted to leave only after its next deadline. Since deadlines can be arbitrarily large, drift can be as well.) As a result, any adaptive algorithm that changes weights by such a method would be unresponsive to the frequent and substantial changes that occur in applications such as Whisper and VEC. Moreover, before the research presented in this dissertation, the only methods for multiprocessor systems for detecting when task weights should change, and the extent of change required, made assumptions about the behavior of tasks that are too conservative for these applications. In this section, we review the multiprocessor framework that we have presented and implemented for scheduling tasks that require adaptation, and our evaluation of this framework by using the core operations of Whisper and VEC (correlation computations for Whisper and bilateral filters for VEC). While this framework was designed for Whisper and VEC, it is general enough to be used for any real-time application that has a workload that is both intensive and time-varying.

Adaptive algorithm for restricted global systems. As stated in Chapter 3 under restricted global scheduling algorithms (i.e., GEDF and NP-GEDF), jobs may miss their deadlines by a bounded amount. In Chapter 3, we presented rules that allow a task to change its weight, even if that task has expired deadlines.

Adaptive algorithm for partitioned systems. In Chapter 4, we presented for partitioned scheduling algorithms (i.e., PEDF and NP-PEDF) rules for changing the weight of a task that are similar to the rules for restricted global scheduling. However, because in partitioned systems each task is assigned to a specific processor, it is possible that when task
weights change, a single processor may become overloaded, i.e., the sum of the weights of all tasks assigned to it is greater than one. To resolve this issue, we presented two techniques in Chapter 4. In the first, the guaranteed weight each task receives is proportional to its "desired" weight relative to the "desired" weight of all tasks assigned to the same processor. By assigning task weights in this manner (i.e., allocating to a task a guaranteed weight that is less then its desired weight), it is possible to guarantee that jobs do not miss deadlines even when a processor is overloaded. (While this technique will guarantee that deadlines are not missed, this behavior only occurs because tasks that are assigned to an overloaded processor receive a guaranteed weight less than their desired weights. If tasks were to receive their desired weights, then, in such a case, deadline tardiness would grow unboundedly.) In the second technique, whenever a processor becomes overloaded by some user-defined value, the entire system is repartitioned. By initiating such repartitioning events, it can be guaranteed that no processor is overloaded for "too long."

Adaptive algorithm for unrestricted global systems. One of the most important unrestricted global scheduling algorithms is the Pfair algorithm $\mathrm{PD}^{2} . \mathrm{PD}^{2}$ is the most efficient known algorithm (in terms of achievable utilization) for scheduling a set of hard real-time sporadic tasks that fully utilize a multiprocessor system. In Chapter 5, we presented a modification of $\mathrm{PD}^{2}$ that allows a task's weight to change at run-time with only a small, constant amount of drift.

Empirical comparison. In Section 7.4, we presented an experimental comparison of our adaptive variants of GEDF, NP-GEDF, PEDF, NP-PEDF, and PD ${ }^{2}$. The results of our experimental comparison suggest the following: first, when it is critical for every task to meet its deadlines and migration/preemption costs are low (i.e., systems like Whisper), PD ${ }^{2}$ is the best choice. Second, when preemption/migration costs are high (i.e., either Whisper or VEC as implemented on a system where the processors are not tightly integrated), average case performance is of the utmost importance, and both fairness and timeliness are less important, then either PEDF or NP-PEDF may be the best choice. Third, when migration/preemption costs are high and a good mix of average-case performance and fairness is beneficial (i.e.,
systems like VEC), then either GEDF or NP-GEDF may the best choice.

Multiprocessor feedback-controlled adaptive algorithm. One method for determining task weights is to assume that each task has multiple service levels, each of which represents a different level of QoS and a different processor weight. When a task has multiple service levels, there may exist times when it must be forced to change its service level because of system constraints. For example, in VEC, if one region is particularly dark, then tasks associated with other lighter regions may be forced to reduce their weights in order to provide more resources to the task assigned to the dark region. Forced changes in service levels present two challenges: how does the system detect when the service level of the system should change, and how are the service levels of tasks determined? In Chapter 6 we presented the adaptable GEDF (AGEDF) framework, which attempts to solve this problem by using feedback control and optimization techniques. In feedback-controlled systems, prior states of the system are used to predict the future state of the system. By employing feedback-control techniques, it is possible to determine needed service-level changes. In addition to constructing this algorithm, in Section 7.5, we presented an implementation of the AGEDF framework on LITMUS ${ }^{\text {RT }}$, a real-time multiprocessor testbed developed by our research group, and evaluated its performance when performing the core operations of both Whisper and VEC. This evaluation showed that, by using a feedback-controlled mechanism, the QoS of both Whisper and VEC can be improved.

### 8.2 Other Related Work

In this section, we briefly discuss other contributions by the author to the field of real-time systems that are outside of the scope of this dissertation.

Quick-release fair scheduling. One drawback to $\mathrm{PD}^{2}$ is that it is not work conserving, i.e., a processor can become idle even if there exists pending work. This can cause task response times to be unnecessarily long. In prior work on introducing work conserving behavior to $\mathrm{PD}^{2}$, techniques were used that can cause a task that consumes otherwise-idle processing capacity
to be "unfairly" penalized later. In joint work with Anderson and Srinivasan, a technique called quick-release fair scheduling was developed that ensures that such tasks are treated fairly (Anderson et al., 2003).

LITMUS ${ }^{\text {RT }}$. In order to better understand how PD $^{2}$, GEDF, NP-GEDF, PEDF, and NP-PEDF would behave in practice, our research group constructed the aforementioned LITMUS ${ }^{\text {RT }}$ real-time multiprocessor testbed. Using this system, we evaluated the performance of PD² ${ }^{2}$ GEDF, NP-GEDF, PEDF, and NP-PEDF. We found that, for hard real-time systems, if tasks have relatively small weights, PEDF performs better in terms of scheduability; however, if task weights are relatively large, then $\mathrm{PD}^{2}$ has superior performance. On the other hand, if bounded deadline tardiness is acceptable, then GEDF and NP-GEDF have superior performance regardless of any task's weight (Brandenburg et al., 2007; Brandenburg et al., 2008; Calandrino et al., 2006).

Multiprocessor synchronization. In work on multiprocessor real-time systems, there is a dearth of research on task synchronization techniques for use in scheduling algorithms where task priorities can change at run time (e.g., PD ${ }^{2}$, GEDF, NP-GEDF, PEDF, and NP-PEDF). To remedy this situation, members of our research group developed the flexible multiprocessor locking protocol (FMLP), which is capable of synchronizing tasks via either semaphores or nonpreemptable queue locks in globally-scheduled systems. We also implemented this protocol in LITMUS ${ }^{\text {RT }}$ and compared both the semaphore and queue-lock versions of the FMLP with each other, and, when implementing shared data objects, with lock-free and wait-free algorithms. We found that for simple data structures, lock-free and wait-free approaches are superior to locking approaches; however, for more complex data structures, non-preemptable queue locks provide superior performance. Interestingly, we found that the use of semaphores always results in worse scheduability than non-preemptable queue locks (Block et al., 2007; Brandenburg et al., 2008).

### 8.3 Future Work

One limitation of this dissertation is that we were not able to fully re-implement Whisper and VEC. In order to do this, four objectives must be achieved. First, an adaptive synchronization protocol must be devised. Second, AGEDF should be extended to more complex machine models. Third, tools must be developed that allow software engineers to easily specify desired adaptive behaviors. Fourth, user studies are needed in order to optimize the performance of both Whisper and VEC. Below, we briefly outline a plan for achieving these three objectives.

An adaptive synchronization protocol. Such a protocol could be devised by modifying the aforementioned FMLP. While an adaptive variant of this protocol will probably be similar to the non-adaptive version, determining the protocol's impact on scheduability will likely be complex since the non-adaptive protocol's scheduability analysis is based on the assumption that task weights do not change. Another complication with the construction of an adaptive FMLP is that critical sections impose additional constraints on the times at which weight changes can be enacted. As a result, introducing synchronization into an adaptive system will likely increase the maximal possible drift because of synchronization-related delays in enacting weight changes.

More complex machine models. In addition to constructing an adaptive synchronization protocol, it would be interesting to extend AGEDF to incorporate more complex machine models in which various factor (e.g., caching, page faults, TLB misses, etc.) that affect execution costs are directly considered. One complication with such machine models is that it becomes more difficult to predict the future weight of a task based solely on its prior weights. For this reason, in more complex machine models, it may be worthwhile for applications to provide additional information about the expected weight of a task to assist in the calculation of future task weights.

Developer tools. One of the advantages of AGEDF is that there are several parameters that can be adjusted for each task in order to control how an application responds to workload changes. In order for a user to take advantage of this flexibility, it would be desirable to have
a set of tools that assist a developer in choosing these parameters under different types of system load. In addition, constructing multithreaded code where each thread (i.e., task) has multiple service levels is a non-trivial issue. Such a toolset could be configured to include three utilities. The first would allow a developer to express a desired system response to different types of workload changes, and from this information, compute appropriate values for the AGEDF parameters. The second would consist of a simulator that allows a developer to quickly assess the behavior of AGEDF for a given set of parameters. The final utility would assist the developer in constructing multithreaded code where each thread may have multiple versions and periods.

User studies. Since both Whisper and VEC are multimedia applications, the ultimate metric of success is user perception. Thus, while many of the parameters of AGEDF can be chosen by simulations conducted by the developer, fine-tuning the parameters for Whisper and VEC will involve conducting user studies. Upon completing these three objectives, we will have built a system for which applications can be easily constructed to fully utilize a multiprocessor system while at the same time providing real-time guarantees.

## BIBLIOGRAPHY

Abeni, L. and Buttazzo, G. (1998). Integrating multimedia applications in hard real-time systems. In Proceedings of the19th IEEE Real-Time Systems Symposium, pages 4-13. IEEE Computer Society.

Abeni, L., Palopoli, L., Lipari, G., and Walpole, J. (2002). Analysis of a reservation-based feedback scheduler. In Proceedings of the 23rd IEEE Real-Time Systems Symposium, pages 71-80. IEEE Computer Society Press.

Al-Omari, R., Manimaran, G., Salapaka, M. V., and Somani., A. K. (2003). Novel algorithms for open-loop and closed-loop scheduling of real-time tasks in multiprocessor systems based on execution time estimation. In Proceedings of the 2003 International Parallel and Distributed Processing Symposium, pages 7-14. IEEE Computer Society Press.

Anderson, J., Block, A., and Srinivasan, A. (2003). Quick-release fair scheduling. In Proceedings of the 24th IEEE Real-Time Systems Symposium, pages 130-141. IEEE Computer Society Press.

Baruah, S., Cohen, N., Plaxton, C., and Varvel, D. (1996). Proportionate progress: A notion of fairness in resource allocation. Algorithmica, 15:600-625.

Bennett, E. (2007). Computational Video Enhancement. PhD thesis, University of North Carolina at Chapel Hill.

Bennett, E. and McMillan, L. (2005). Video enhancement using per-pixel virtual exposures. ACM Transactions on Graphics (SIGGRAPH), 24(3):845-852.

Bertsekas, D. (1999). Nonlinear Programming. Athena Scientific, second edition.
Block, A. and Anderson, J. (2006). Accuracy versus migration overhead in multiprocessor reweighting algorithms. In Proceedings of the 12th International Conference on Parallel and Distributed Systems, pages 355-364. IEEE Computer Society Press.

Block, A., Anderson, J., and Bishop, G. (2008a). Fine-grained task reweighting on multiprocessors. Journal of Embedded Computing, Special Issue on Multiprocessor Real-Time Scheduling (to appear).

Block, A., Anderson, J., and Devi, U. (2008b). Task reweighting under global scheduling on multiprocessors. Real-Time Systems, Special Issue on Selected Papers from the 18th Euromicro Conference on Real-Time Systems, 39:123-167.

Block, A., Brandenburg, B., Anderson, J., and Quint, S. (2008c). Feedback-controlled adaptive multiprocessor real-time systems. In Proceedings of the 20th Euromicro Conference on Real-Time Systems, pages 23-33. Kluwer Academic Publishers.

Block, A., Leontyev, H., Brandenburg, B., and Anderson, J. (2007). A flexible real-time locking protocol for multiprocessors. In Proceedings of the 13th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications, pages 47-57. IEEE Computer Society Press.

Brandenburg, B. and Anderson, J. (2008). A comparison of the M-PCP, D-PCP, FMLP, on LITMUS ${ }^{\text {RT }}$ (in submission).

Brandenburg, B., Block, A., Calandrino, J., Devi, U., Leontyev, H., and Anderson, J. (2007). LITMUS ${ }^{\text {RT: A status report. In Proceedings of the 9th Real-Time Linux Workshop, }}$ pages 107-123. The Real-Time Linux Foundation.

Brandenburg, B., Calandrino, J., Block, A., Leontyev, H., and Anderson, J. (2008). Real-time synchronization on multiprocessors: To block or not to block, to suspend or spin? In Proceedings of the 14th IEEE Real-Time and Embedded Technology and Applications Symposium, pages 342-353. IEEE Computer Society Press.

Brandt, S. A., Banachowski, S., Lin, C., and Bisson, T. (2003). Dynamic integrated scheduling of hard real-time, soft real-time and non-real-time processes. In Proceedings of the 24th IEEE International Real-Time Systems Symposium, pages 396-407. IEEE Computer Society Press.

Calandrino, J., Leontyev, H., Block, A., Devi, U., and Anderson, J. (2006). LITMUS ${ }^{\text {RT }}$ : A testbed for empirically comparing real-time multiprocessor schedulers. In Proceedings of the 27th IEEE International Real-Time Systems Symposium, pages 111-126. IEEE Computer Society Press.

Carpenter, J., Funk, S., Holman, P., Srinivasan, A., Anderson, J., and Baruah, S. (2004). Handbook of Scheduling: Algorithms, Models, and Performance Analysis, chapter A Categorization of Real-time Multiprocessor Scheduling Problems and Algorithms, pages 30.1-30.19. Chapman and Hall/CRC.

Chen, C. and Tripathi, S. (1994). Multiprocessor priority ceiling based protocols. CS-TR3252, University of Maryland.

Cucinotta, T., Palopoli, L., Marzario, L., Lipari, G., and Abeni, L. (2004). Adaptive reservations in a linux environment. In Proceedings of the 10th IEEE Real-Time and Embedded Technology and Applications Symposium, pages 238-245. IEEE Computer Society Press.

Devi, U. and Anderson, J. (2008). Tardiness bounds under global EDF scheduling on a multiprocessor. Real-Time Systems, 38(2):133-189.

Devi, U., Leontyev, H., and Anderson, J. (2006). Efficient synchronization under global EDF scheduling on multiprocessors. In 18th Euromicro Conference on Real-Time Systems, pages 75-84. Kluwer Academic Publishers.

Gai, P., Natale, M. D., Lipari, G., Ferrari, A., Gabellini, C., and Marceca, P. (2003). A comparison of MPCP and MSRP when sharing resources in the Janus multiple processor on a chip platform. In Proceedings of the 9th IEEE Real-Time and Embedded Technology and Applications Symposium, pages 189-198. IEEE Computer Society Press.

Hamadoui, M. and Ramanathan, P. (1995). A dynamic priority assignment technique for streams with ( $m, k$ )-firm deadlines. IEEE Transactions on Computers, 44(12):14431451.

Holman, P. and Anderson, J. (2006). Locking under Pfair scheduling. ACM Transactions on Computer Systems, 24(2):140-170.

Lopez, J. M., Diaz, J. L., and Garcia, D. F. (2004). Utilization bounds for EDF scheduling on real-time multiprocessor systems. Real-Time Systems, 28(1):39-68.

Lu, C., Stankovic, J., Abdelzaher, T., Gang, T., Son, S., and Marley, M. (2000). Performance specifications and metrics for adaptive real-time systems. In Proceedings of the 21st IEEE Real-Time Systems Symposium, pages 13-23. IEEE Computer Society Press.

Lu, C., Stankovic, J., Son, S., and Tao, G. (2002). Feedback control real-time scheduling: Framework, modeling, and algorithms. Real-Time Systems, 23(1-2):85-126.

Lu, C., Stankovic, J., Tao, G., and Son, S. (1999). Design and evaluation of a feedback control EDF scheduling algorithm. In Proceedings of the 20th IEEE Real-Time Systems Symposium, pages 56-67. IEEE Computer Society Press.

Marti, P., Lin, C., Brandt, S., Velasco, M., and Fuertes, J. (2004). Optimal state feedback based resource allocation for resource-constrained control tasks. In Proceedings of the 25th IEEE International Real-Time Systems Symposium, pages 161-172. IEEE Computer Society Press.

Maybeck, P. (1979). Stochastic models, estimation, and control, volume 141 of Mathematics in Science and Engineering. Academic Press, Inc.

Mellor-Crummey, J. and Scott, M. (1991). Algorithms for synchronization on shared-meory multiprocessors. ACM Transactions on Computer Systems, 9(1):21-65.

Mok, A. K. (1983). Fundamental design problems of distributed systems for the hard-real-time environment. Technical report, Massachusetts Institute of Technology.

Nise, N. (2004). Control Systems Engineering. Wiley and Sons, fourth edition.
Pappas, T. and Safranek, R. (2000). Handbook of Image and Video Processing, chapter Perceptual Criteria for Image Quality Evaluation, pages 669-684. Academic Press, Inc.

Rajkumar, R. (1991). Synchronization in Real-Time Systems: A Priority Inheritance Approach. Kluwer Academic Publishers.

Sahoo, D., Swaminathan, S., Al-Omari, R., Salapaka, M., Manimaran, G., and Somani, A. (2002). Feedback control for real-time scheduling. In Proceedings of the 21st American Control Conference, Volume 2, pages 1254-1259. IEEE Computer Society Press.

Sha, L., Rajkumar, R., and Lehoczky, J. (1990). Priority inheritance protocols: An approach to real-time synchronization. IEEE Transactions on Computers, 39(9):1175-1185.

Smith, S. (1997). The Scientist and Engineer's Guide to Digital Signal Processing. California Technical Publishing. Available at www.dspguide.com.

Srinivasan, A. (2003). Efficient and Flexible Fair Scheduling of Real-time Tasks on Multiprocessors. PhD thesis, University of North Carolina at Chapel Hill.

Srinivasan, A. and Anderson, J. (2005). Fair scheduling of dynamic task systems on multiprocessors. Journal of Software Systems, 77(1):67-80.

Srinivasan, A. and Anderson, J. (2006). Optimal rate-based scheduling on multiprocessors. Journal of Computer and System Science, 72(6):1094-1117.

Stoica, I., Abdel-Wahab, H., Jeffay, K., Baruah, S., Gehrke, J. E., and Plaxton, C. G. (1996). A proportional share resource allocation algorithm for real-time, time-shared systems. In Proceedings of the 17th IEEE Real-Time Systems Symposium, pages 288-299. IEEE Computer Society Press.

Vallidis, N. (2002). WHISPER: A Spread Spectrum Approach to Occlusion in Acoustic Tracking. PhD thesis, The University of North Carolina at Chapel Hill, North Carolina.

Welch, G. and Bishop, G. (1995). An introduction to the Kalman filter. Technical Report TR95-041, Department of Computer Sceience. The Univerity of North Carolina at Chapel Hill.


[^0]:    ${ }^{1}$ In prior work (Block and Anderson, 2006; Block et al., 2008a; Block et al., 2008b), we referred to VEC as ASTA, which stands for Adaptive Saptio-Temporal Accumulation Filter; however, VEC is more technically

[^1]:    ${ }^{2}$ Specifically, the first correct feedback algorithm was proposed in (Lu et al., 2000). The original system, FC-EDF, proposed in (Lu et al., 1999) could not satisfy its design specification because it was possible for the controller to become saturated, thus rendering it unable to correctly adjust the system.

[^2]:    ${ }^{3}$ A tardiness of 0 is only guaranteed to hold for a given task $T_{i}$ with an execution cost of $a$ and a weight of $b / c$ only if there exists some integer $n$ such that $b \cdot n=a$; otherwise the tardiness of a task may be up to one quantum.

[^3]:    ${ }^{1}$ In the literature on proportional share scheduling, weights can be arbitrary positive values; however in this dissertation, we will continue to view weights as desired processor shares.

[^4]:    ${ }^{2}$ The FCS framework can be used for aperiodic tasks; however, since the focus of this dissertation is on periodic/sporadic systems, we focus exclusively on this aspect of the FCS framework.
    ${ }^{3}$ The FCS framework can be extended to work with non-EDF based approaches; however, since those are beyond the focus of this dissertation, we do not discuss such extensions here.

[^5]:    * Contents of this chapter previously appeared in preliminary form in the following paper: Block, A., Anderson J., and Devi, U. (2008b). Task reweighting under global scheduling on multiprocessors. Real-Time Systems, Special Issue on Selected Papers from the 18th Euromicro Conference on Real-Time Systems, 39:123-167.

[^6]:    ${ }^{1}$ Since reweighting events may modify the actual execution time of a job, removing this assumption would entail having notation to distinguish between redefined execution times as the result of reweighting and redefined execution times that occur due to a task executing for less than its specified execution time.

[^7]:    ${ }^{2}$ In (Devi and Anderson, 2008), an effective task is referred to as an active task. We use this alternative term here to avoid conflicts in terminology.
    ${ }^{3}$ Here again, to avoid conflicting terminology, we differ from the term used in (Devi and Anderson, 2008). Stopping is referred to as halting in (Devi and Anderson, 2008).
    ${ }^{4}$ Since each extended sporadic task has only one weight, in an SW schedule the extended sporadic task $T_{i}$ is allocated $\mathrm{wt}\left(T_{i}\right)$ at each instant it is effective.

[^8]:    ${ }^{5}$ This usage of the term "subtask" should not be confused with that used in work on Pfair scheduling.

[^9]:    * Contents of this chapter previously appeared in preliminary form in the following paper:

    Block, A. and Anderson J. (2006). Accuracy versus migration overhead in multiprocessor reweighting algorithms. In Proceedings of the 12th International Conference on Parallel and Distributed Systems, pages 355-364.
    ${ }^{1}$ A reweighting algorithm is optimal if each task can always be granted a guaranteed weight equal to its desired weight, provided the sum of all desired weights is at most the number of available processors.

[^10]:    ${ }^{2}$ In Table 4.3, the formula for minimizing the MAOE only holds if $\frac{\omega\left(P_{[q]}, t\right)}{\operatorname{ASSN}\left(P_{[q]}, t\right) \mid} \leq \operatorname{Dwt}\left(T_{z}, t\right)$ for every $T_{z} \in \operatorname{ASSN}\left(P_{[q]}, t\right)$. Additionally, the formula for minimizing the AROE only holds if $\omega\left(P_{[q]}, t\right) \leq \operatorname{Dwt}\left(T_{H}, t\right)$.

[^11]:    * Contents of this chapter previously appeared in preliminary form in the following paper:

    Block, A., Anderson J., and Bishop, G. (2008a). Fine-Grained task reweighting on multiprocessors. Journal of Embedded Computing, (to appear).

[^12]:    ${ }^{1}$ Originally, in (Block et al., 2008a), a negative-changeable task was called ideal-changeable and a task that was positive-changeable was called omission-changeable. The names have been changed here to be consistent with the reweighting rules for GEDF and PEDF.

[^13]:    * Contents of this chapter previously appeared in preliminary form in the following paper:

    Block, A., Brandenburg, B., Anderson J., and Quint, S. (2008). An Adaptive Framework for Multiprocessor Real-Time Systems. In Proceedings of the 20th Euromicro Conference on Real-Time Systems, pages 23-33.

[^14]:    ${ }^{1}$ The usage of the term "weight" here should not be confused with that elsewhere in this dissertation, where this term is used to indicate a processor share.

[^15]:    ${ }^{2}$ Since understanding $\lambda_{s}$ 's exact formula in detail requires knowledge of image processing techniques that are beyond the scope of this dissertation, we refer the reader to (Bennett, 2007) for a complete discussion of this topic.

[^16]:    ${ }^{3}$ This is true of the LITMUS ${ }^{\text {RT }}$ release that was used in performing the experiments in this chapter. Recently, LITMUS ${ }^{\text {RT }}$ was re-based to Linux 2.6.24 and a number of improvements were made.

[^17]:    ${ }^{4}$ Sometimes also called "operations" or "callbacks."

[^18]:    ${ }^{5} \mathrm{~A}$ complete list of all 13 methods can be found in (Brandenburg et al., 2007).
    ${ }^{6} \mathrm{~A}$ complete discussion of the system call APIs can be found in (Brandenburg et al., 2007).

