CALOREE: Learning Control for Predictable Latency and Low Energy

Nikita Mishra

Connor Imes

University of Chicago
nmishra,ckimes@cs.uchicago.edu

John D. Lafferty

Yale University

john.lafferty@yale.edu

Henry Hoffmann

University of Chicago

hankhoffmann@cs.uchicago.edu

Abstract

Many modern computing systems must provide reliable latency with minimal energy. Two central challenges arise when allocating system resources to meet these conflicting goals: (1) *complexity*—modern hardware exposes diverse resources with complicated interactions—and (2) *dynamics*—latency must be maintained despite unpredictable changes in operating environment or input. Machine learning accurately models the latency of complex, interacting resources, but does not address system dynamics; control theory adjusts to dynamic changes, but struggles with complex resource interaction. We therefore propose CALOREE, a resource manager that learns key control parameters to meet latency requirements with minimal energy in complex, dynamic environments. CALOREE breaks resource allocation into two sub-tasks: learning how interacting resources affect speedup, and controlling speedup to meet latency requirements with minimal energy. CALOREE defines a general control system—whose parameters are customized by a learning framework—while maintaining control-theoretic formal guarantees that the latency goal will be met. We test CALOREE's ability to deliver reliable latency on heterogeneous ARM big.LITTLE architectures in both single and multi-application scenarios. Compared to the best prior learning and control solutions, CALOREE reduces deadline misses by 60% and energy consumption by 13%.

CCS Concepts • **Computing methodologies** → **Computational control theory**; **Machine learning**; • **Computer systems organization** → **Heterogeneous (hybrid) systems**; **Embedded systems**; **Real-time system architecture**; • **Hardware** → **Chip-level power issues**;

Keywords machine learning; control theory; real-time systems; energy; heterogeneous architectures; resource allocation

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ASPLOS '18, March 24–28, 2018, Williamsburg, VA, USA

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-4911-6/18/03...\$15.00

<https://doi.org/10.1145/3173162.3173184>

ACM Reference Format:

Nikita Mishra, Connor Imes, John D. Lafferty, and Henry Hoffmann. 2018. CALOREE: Learning Control for Predictable Latency and Low Energy. In *ASPLOS '18: 2018 Architectural Support for Programming Languages and Operating Systems, March 24–28, 2018, Williamsburg, VA, USA*. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3173162.3173184>

1 Introduction

Large classes of computing systems—from embedded to servers—must deliver reliable latency while minimizing energy to prolong battery life or lower operating costs. To address these conflicting requirements, hardware architects expose diverse, heterogeneous resources with a wide array of latency and energy tradeoffs. Software must allocate these resources to guarantee latency requirements are met with minimal energy.

There are two primary difficulties in efficiently allocating heterogeneous resources. The first is *complexity*: resources interact in intricate ways, leading to non-convex optimization spaces. The second is *dynamics*: performance requirements must be met despite unpredictable disturbances; e.g., changes in application workload or operating environment. Prior work addresses each of these difficulties individually.

Machine learning handles complex modern processors, modeling an application's latency and power as a function of resource configurations [5, 12, 15, 31, 52, 57, 58, 66, 85]. These predictions, however, are not useful if the environment changes dynamically; e.g., a second application enters the system. Control theoretic approaches dynamically adjust resource usage based on models of the difference between measured and expected behavior [8, 24, 25, 30, 42, 64, 69, 74, 80, 82]. Control provides formal guarantees that it will meet the latency goal in dynamic environments, but these guarantees are based on ground-truth models relating resources and latency. If these models are not known or there is error between the modeled and actual behavior, the controller will fail to deliver the required latency.

Intuitively, combining learned models of complex hardware resources with control-theoretic resource management should produce predictable latency in complex, dynamic systems. To derive the benefit of both, however, requires addressing two major challenges:

- Dividing resource allocation into sub-problems that suit learning and control's different strengths.
- Defining abstractions that efficiently combine sub-problem solutions, while maintaining control's formal guarantees.

We address the first challenge by splitting resource allocation into two sub-tasks. The first is learning *speedup*—instead of absolute performance—so that all unpredictable external interference is viewed as a change to a *baseline* latency and the relative speedup is independent of these changes. Learning is well-suited to modeling speedups as a function of resource usage and finding Pareto-optimal trade-offs in speedup and energy. The second sub-task is controlling speedup dynamically based on the difference between measured and desired latency. Once the learner has found Pareto-optimal tradeoffs the problem is convex and well-suited to adaptive control solutions which guarantee the required speedup even in dynamic environments. Figure 1 illustrates the intuition: processor complexity creates local optima, where control solutions can get stuck; but learning finds true optimal tradeoffs—“convexifying”—the problem, allowing control techniques to handle dynamics while providing globally optimal energy.

We address the second challenge by defining an interface between learning and control that maintains control's formal guarantees. This interface consists of two parts. The first is a *performance hash table* (PHT) that stores the learned model between configurations and speedup. The PHT allows the controller to find the resource allocation that meets a desired speedup with minimal energy and requires only constant time— $O(1)$ —to access. The second part of the interface is the learned variance. Knowing this value, the controller can adjust itself to maintain formal guarantees even though the speedup is modeled by a noisy learning mechanism at run-time, rather than directly measured offline—as it would be in traditional control design.

Thus, we propose a general methodology where an abstract control system is customized at runtime by a learning mechanism to meet latency requirements with minimal energy. We refer to this approach as CALOREE¹. Unlike previous work on control systems that required numerous user-specified models and parameters [8, 30, 42, 64, 82], CALOREE's learner tunes the control parameters automatically; *i.e.*, it requires no user-level inputs other than latency requirements. We evaluate CALOREE by implementing the learners on an x86 server and the controller on a heterogeneous ARM big.LITTLE device. We compare to state-of-the-art learning (including polynomial regression [15, 66], collaborative filtering—*i.e.*, the Netflix algorithm[3, 12]—and a hierarchical Bayesian model

Figure 1. Learning smooths the controller's domain.

Figure 2. (a) STREAM performance vs. configuration. Darker color means higher performance. (b) Managing STREAM latency: *Learning* handles the complexity, but *control* oscillates.

[52]) and control (including proportional-integral-derivative [24] and adaptive, or self-tuning [41]) controllers. We set latency goals for benchmark applications and measure both the percentage of time the requirements are violated and the energy. We test both *single-app*—where an application runs alone—and *multi-app* environments—where background applications enter the system and compete for resources.

Our results show that CALOREE achieves the *most reliable latency* and *best energy savings*. In the *single-app* case, the best prior technique misses 10% of deadlines on average, while CALOREE misses only 6%. All other approaches miss 100% of deadlines for at least one application, but CALOREE misses, at most, 11% of deadlines. In the *multi-app* case, the best prior approach averages 40% deadline misses, but CALOREE misses just 20%. We evaluate energy by comparing to optimal energy assuming a perfect model of application, system, and future. In the *single-app* case, the best prior approach averages 18% more energy consumption than optimal, but CALOREE consumes only 4% more. In the *multi-app* case, the best prior approach averages 28% more energy than optimal, while CALOREE consumes just 6% more.

In summary, CALOREE is the first work to use learning to customize control systems at runtime, ensuring application latency—both formally and empirically—with no prior knowledge of the controlled application. Its contributions are:

- Separation of resource management into (1) *learning* complicated resource interactions and (2) *controlling* speedup.
- A generalized control design usable with multiple learners.
- A method for guaranteeing latency using learned—rather than measured—models.

2 Background and Motivation

This section illustrates how learning handles complexity, how control handles dynamics, and then describes a key challenge that must be overcome to combine learning and control.

¹Control And Learning for Optimal Resource Energy Efficiency

2.1 Learning Complexity

We demonstrate how well learning handles complex resource interaction for STREAM on an ARM big.LITTLE processor with four big, high-performance cores and four LITTLE, energy efficient cores. The big cores support 19 clock speeds, while the LITTLE cores support 14.

Figure 2a shows STREAM's performance for different resource configurations. STREAM has complicated behavior: the LITTLE cores' memory hierarchy cannot deliver performance. The big cores' more powerful memory system delivers greater performance, with a peak at 3 big cores. At low clockspeeds, 3 big cores cannot saturate the memory bandwidth, while at high clockspeeds thermal throttling creates performance loss. Thus, the peak speed occurs with 3 big cores at 1.2 GHz, and it is inefficient to use the LITTLE cores. STREAM, however, does not have distinct phases, so system dynamics are not an issue in this case.

Figure 2b shows 20 iterations of existing learning [52] and adaptive control [30] approaches allocating resources to STREAM. The x-axis shows iteration and the y-axis shows latency normalized to the requirement. The *learning* approach estimates STREAM's performance and power for all configurations and uses the lowest energy configuration that delivers the required latency. The *adaptive controller* begins with a generic notion of power/performance tradeoffs. As the controller runs, it measures latency and adjusts both the allocated resources and its own parameters. The adaptive controller dynamically adjusts to non-linearities with a series of linear approximations; however, inaccuracies in the relationship between resources and latency cause oscillations that lead to latency violations. This behavior occurs because the controller's adaptive mechanisms cannot handle STREAM's complexity, a known limitation of adaptive control systems [16, 30, 82]. Hence, the *learner's* ability to model complex behavior is crucial.

2.2 Controlling Dynamics

We now consider a dynamic environment. We begin with x264 running alone on the system. Figure 3a shows x264's behavior. It achieves the best performance on 4 big cores at the highest clockspeed; the 4 LITTLE cores are more energy-efficient but slower. For x264, the challenge is determining how to use both the LITTLE and big cores to conserve energy while still meeting the latency requirements. During execution, we launch a second application—STREAM—on a single big core, dynamically changing available resources.

Figure 3b shows the results. The vertical dashed line at frame 99 shows when the second application begins. At that point, the adaptive controller detects x264's latency spike—rather than detecting the new application specifically—and it increases clockspeed and moves x264 from 4 to 3 big cores. The learner, however, does not have a mechanism to adapt to the altered environment. While we could theoretically add

Figure 3. (a) x264 performance vs. configuration. Darker color means higher performance. (b) Managing x264's latency with another application: *control* adapts to the change (the vertical dashes), but *learning* does not.

feedback to the learner and periodically re-estimate the configuration space, doing so is impractical due to high overhead for learners capable of handling this complexity [12, 13, 52]. Simpler reinforcement learners can adapt, but cannot guarantee reconvergence after the dynamic change [46, 71].

2.3 Challenges Combining Learning and Control

Sections 2.1 and 2.2 motivate splitting the resource allocation problem into modeling—handled by learning—and dynamic management—handled by control. This subsection demonstrates the importance of defining principled techniques for controlling systems using learned models.

The controller's *pole* is a particularly important parameter [41]. Control engineers tune the pole to trade response time for noise sensitivity. Traditionally, the data used to set the pole comes from many observations of the controlled system and is considered *ground truth* [24, 44]. CALOREE, however, must tune the pole based the learner's models, which may have noise and/or errors.

To demonstrate the pole's importance when using learned data, we again control x264, using the adaptive controller from the previous subsection. Instead of using a ground truth model mapping resource usage to performance, we model it using the learner from the first subsection. We compare the results with a carefully hand-tuned pole to those using the default pole provided by the controller developers [30].

As shown in Figure 4, the carefully tuned pole converges. The default pole, however, oscillates around the latency target, resulting in a number of missed deadlines. Additionally,

Figure 4. Comparison of carefully tuned and default poles.

Figure 5. CALOREE overview.

the frames below the desired latency waste energy because they spend more time on the big, inefficient cores. The pole captures the system's *inertia*—dictating how fast it should react to environmental changes. If the learner is noisy or inaccurate, the controller should trust it less and move slowly. Rather than require users with both computing and control knowledge to tune the pole, *CALOREE incorporates the learner's estimated variance to compute a pole that provides probabilistic convergence guarantees*.

3 CALOREE: Learning Control

Figure 5 shows CALOREE's approach of splitting resource management into learning and control tasks and then composing their individual solutions. When a new application enters the system, an adaptive control system allocates resources using a generic model, recording latency and power. The records are sent to a learner, which predicts the application's latency and power in all other resource configurations. The learner extracts those that are predicted to be Pareto-optimal and packages them in a data structure: the performance hash table (PHT). The PHT and the estimated variance are sent to the controller, which sets its pole and selects an energy minimal resource configuration with formal guarantees of convergence to the desired latency. CALOREE's only user-specified parameter is the latency requirement.

Figure 6 illustrates the asynchronous interaction between CALOREE's learner and controller. The controller starts—using a conservative, generic speedup model—when a new application launches. The controller sends the learner the application's name and device type (message 1, Figure 6). The learner determines how many samples are needed for an accurate prediction and sends this number to the controller (message 2). The controller takes these samples and sends the latency and power of each measured configuration to the learner (message 3). The learner may require time to make predictions (*i.e.*, train the model); so, the controller does not wait, but continues with the conservative model. Once the learner predicts the optimal configurations, it sends that

Figure 6. Temporal relationship of learning and control.

data and the variance estimate to the controller (message 4), which uses the learned model from then on.

Figure 6 shows several key points about the relationship between learning and control. First, the controller never waits for the learner: it uses a conservative, less-efficient control specification until the learner produces application-specific predictions. Second, the controller does not continuously communicate with the learner—this interaction happens once at application launch. Third, if the learner crashed, the controller defaults to the generic adaptive control system. If the learner crashed after sending its predictions, the controller does not need to know. Finally, the learner and controller have a clearly defined interface, so they can be run in separate processes or physically separate devices.

We first describe adaptive control. We then generalize this approach, separating out parameters to be learned. Next, we discuss the class of learners that work with CALOREE. Finally, we formally analyze CALOREE's guarantees.

3.1 Traditional Control for Computing

A multiple-input, multiple-output (MIMO) controller manages multiple resources to meet multiple goals. The inputs are measurements, *e.g.*, latency. The outputs are the resource settings to be used at a particular time, *e.g.*, an allocation of big and LITTLE cores and a clockspeed for each.

These difference equations describe a generic MIMO controller managing n resources to meet m goals at time t :²

$$\begin{aligned} \mathbf{x}(t+1) &= \mathbf{A} \cdot \mathbf{x}(t) + \mathbf{B} \cdot \mathbf{u}(t) \\ \mathbf{y}(t) &= \mathbf{C} \cdot \mathbf{x}(t) \end{aligned}, \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^q$ is the controller's *state*, capturing the relationship between resources and goals; q is the controller's *degree*, or complexity of its internal state. $\mathbf{u}(t) \in \mathbb{R}^n$ represents the current resource *configuration*; *i.e.*, the i th vector element is the amount of resource i allocated at time t . $\mathbf{y}(t) \in \mathbb{R}^m$ represents the value of the goal dimensions at time t . The matrices $\mathbf{A} \in \mathbb{R}^{q \times q}$ and $\mathbf{B} \in \mathbb{R}^{q \times n}$ relate the resource configuration to the controller state. The matrix $\mathbf{C} \in \mathbb{R}^{m \times q}$ relates the controller state to the expected behavior. This control definition does not assume the states or the resources are independent, but it does assume a linear relationship.

²We assume discrete time, and thus, use difference equations rather than differential equations that would be used for continuous systems.

For example, in our ARM big.LITTLE system there are four resources: the number of big cores, the number of LITTLE cores, and the speeds for each of the big and LITTLE cores. There is also a single goal: latency. Thus, in this example, $n = 4$ and $m = 1$. The vector $\mathbf{u}(t)$ has four elements representing the resource allocation at time t . q is the number of variables in the controller's state which can vary between 1 to n . The matrices $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{C}$ capture the linear relationship between the control state $\mathbf{x}$, the resource usage $\mathbf{u}$, and the measured behavior. In this example, we know there is a non-linear relationship between the resources. We overcome this difficulty by tuning the matrices at each time step—approximating the non-linear system through a series of changing linear formulations. This approximation is a form of *adaptive* or *self-tuning* control [41]. Such adaptive controllers provide formal guarantees that they will converge to the desired latency even in the face of non-linearities, but they still assume convexity.

This controller has two major drawbacks. First, it requires matrix computation, so its overhead scales poorly in the number of resources and in the number of goals [24, 64]. Second, the adaptive mechanisms require users to specify both (1) starting values of the matrices $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{C}$ and (2) the method for updating these matrices to account for any non-convexity in the relationship between resources and latency [30, 41, 64, 82]. Therefore, typically 100s to 1000s of samples are taken at design time to ensure that the starting matrices are sufficient to ensure convergence [17, 44, 59].

3.2 CALOREE Control System

To overcome the above issues, CALOREE abstracts the controller of Eqn. 1 and factors out those parameters to be learned. Specifically, CALOREE takes three steps to transform a standard control system into one that works without prior knowledge of the application to be controlled:

1. controlling *speedup* (which is an abstraction of latency) rather than resources;
2. turning speedup into a minimal energy *resource schedule*;
3. and exploiting the *problem structure* to solve this scheduling problem in constant time.

These steps assume a separate learner has produced predictions of how resource usage affects latency and power. The result is that CALOREE's controller runs in constant time without requiring any user-specified parameters.

3.2.1 Controlling Speedup

CALOREE converts Eqn. 1 into a single-input (latency), single-output (speedup) controlling using $\mathbf{A} = 0$, $\mathbf{B} = b(t)$, $\mathbf{C} = 1$, $\mathbf{u} = \text{speedup}$, and $y = \text{perf}$; where $b(t)$ is a time-varying parameter representing the application's *base speed*—the speed when all resources are available—and *perf* is the measured latency. Using these substitutions, we eliminate $\mathbf{x}$ from

Eqn. 1 to relate speedup to latency:

$$\text{lat}(t) = 1/(b(t) \cdot \text{speedup}(t - 1)) \quad (2)$$

While $b(t)$ is application-specific. CALOREE assumes base speed is time-variant as applications will transition through phases and it estimates this value online using the standard technique of Kalman filter estimation [75].

CALOREE must eliminate the error between the target latency and the goal: $\text{error}(t) = \text{goal} - 1/\text{lat}(t)$. Given Eqn. 2, CALOREE uses the integral control law [24]:

$$\text{speedup}(t) = \text{speedup}(t - 1) - \frac{1 - \rho(t)}{b(t)} \cdot \text{error}(t) \quad (3)$$

which states that the speedup at time t is a function of the previous speedup, the error at time t , the base speed $b(t)$, and the controller's *pole*, $\rho(t)$. Standard control techniques statically determine the pole and the base speed, but CALOREE *dynamically sets the pole and base speed to account for error in the learner's predictions—an essential modification for providing formal guarantees of the combined control and learning systems*. For stable control, CALOREE ensures $0 \leq \rho(t) < 1$. Small values of $\rho(t)$ eliminate error quickly, but make the controller more sensitive to the learner's inaccuracies. Larger $\rho(t)$ makes the system more robust at the cost of increased convergence time. Section 3.5 describes how CALOREE sets the pole, but we first address converting speedup into a resource allocation.

3.2.2 Converting Speedup to Resource Schedules

CALOREE must map Eqn. 3's speedup into a resource allocation. On our example big.LITTLE architecture an allocation includes big and LITTLE cores as well as a speed for both. The primary challenge is that speedups in real systems are discrete non-linear functions of resource usage, while Eqn. 3 is a continuous linear function. We bridge this divide by assigning time to resource allocations such that the average speedup over a control interval is that produced by Eqn. 3.

The assignment of time to resource configurations is a *schedule*; e.g., spending 10 ms on the LITTLE cores at 0.6 GHz and then 15 ms on the big cores at 1 GHz. Typically many schedules can deliver a particular speedup and CALOREE must find one with minimal energy. Given a time interval T , the $\text{speedup}(t)$ from Eqn. 3, and C different resource configurations, CALOREE solves:

$$\underset{\tau \in \mathbb{R}^C}{\text{minimize}} \quad \sum_{c=0}^{C-1} \tau_c \cdot p_c \quad (4)$$

$$\text{s.t.} \quad \sum_{c=0}^{C-1} \tau_c \cdot s_c = \text{speedup}(t)T \quad (5)$$

$$\sum_{c=0}^{C-1} \tau_c = T \quad (6)$$

$$0 \leq \tau_c \leq T, \quad \forall c \in \{0, \dots, C-1\} \quad (7)$$

Figure 7. Data structure to efficiently convert required speedup into a resource configuration.

where p_c and s_c are configuration c 's estimated *powerup*—analogous to speedup—and speedup; τ_c is the time to spend in configuration c . Eqn. 4 is the objective: minimizing energy (power times time). Eqn. 5 states that the average speedup must be maintained, while Eqn. 6 requires the time to be fully utilized. Eqn. 7 simply avoids negative time.

3.3 Exploiting Problem Structure for Fast Solutions

By encoding the learner's predictions in the performance hash table, CALOREE solves Eqns. 4–7 in constant time.

Kim et al. analyze the problem of minimizing energy while meeting a latency constraint and observe that there must be an optimal solution with the following properties [34]:

- At most two of τ_c are non-zero, meaning that at most two configurations will be used in any time interval.
- If you chart the configurations in the power and speedup tradeoff space (e.g., the top half of Figure 7) the two configurations with non-zero τ_c lie on the lower convex hull of the points in that space.
- The two configurations with non-zero τ_c are adjacent on the convex hull: one above the constraint and one below.

The PHT (shown in Figure 7) provides constant time access to the lower convex hull. It consists of two arrays. The first is pointers into the second, which stores resource configurations the learner estimates to be on the lower convex hull sorted by speedup. Recall speedups are computed relative to the base speed, which uses all resources. The largest estimated speedup is therefore 1. The first array of pointers has a *resolution* indicating how many decimal points of precision it captures and it is indexed by speedup. The example in Figure 7 has a resolution of 0.1. Each pointer in the first array points to the configuration in the second array that has the largest speedup less than or equal to the index.

CALOREE computes $speedup(t)$ and uses the PHT to convert speedup into two configurations: hi and lo . To find the

hi configuration, CALOREE clamps the desired speedup to the largest index lower than $speedup(t)$, indexes into the configuration array, and then walks forward until it finds the first configuration with speedup higher than $speedup(t)$. To find lo , it clamps the desired speedup to the smallest index higher than $speedup(t)$, indexes into the configuration array, and then walks backwards until it finds the configuration with the largest speedup less than $speedup(t)$.

For example, consider the PHT in Figure 7 and a $speedup(t) = .65$. To find hi , CALOREE indexes at .6 and walks up to find $c = 2$ with $s_c = .8$, setting $hi = 2$. To find lo , CALOREE indexes the table at .7 and walks backward to find $c = 1$ with $s_c = .2$, setting $lo = 1$.

CALOREE sets τ_{hi} and τ_{lo} by solving:

$$T = \tau_{hi} + \tau_{lo} \quad (8)$$

$$speedup(t) = \frac{s_{hi} \cdot \tau_{hi} + s_{lo} \cdot \tau_{lo}}{T} \quad (9)$$

where the controller provides $speedup(t)$ and the learner predicts s_c . By solving Eqns. 8 and 9, CALOREE has turned the controller's speedup into a resource schedule using predictions stored in the PHT.

3.4 CALOREE Learning Algorithms

The previous subsection describes a general, abstract control system, which can be customized with a number of different learning methods. The requirements on the learner are that it must produce 1) predictions of each resource configuration's speedup and powerup and 2) estimate of its own variance σ^2 . This section describes the general class of learning mechanisms that meet these requirements.

We refer to application-specific predictors as *online* because they work for the current application, ignoring knowledge of other applications. We refer to general predictors as *offline* as they use prior observations of other applications to predict the behavior of a new application. A third class of *transfer learning* combines information from the previously seen applications and current application to model the future behavior of the current application [56]. Transfer learning produces highly accurate models since it augments online data with offline information from other applications. CALOREE uses transfer learners because CALOREE's separation of learning and control makes it easy to incorporate data from other applications—the learner in Figure 6 can simply aggregate data from multiple controllers. We describe two examples of appropriate transfer learning algorithms.

Netflix Algorithm: The Netflix problem is a famous challenge to predict users' movie preferences. The challenge was won by realizing that if 2 users both like some movies, they might have similar taste in other movies [3]. This approach allows learners to borrow large amounts of data from other applications to answer questions about a new application. One formulation of this problem is to assume the matrix of resource-vs-speedup is low-rank and solve the problem

while True do

Measure application latency
 Compute required speedup (Equation (2))
 Lookup s_{hi} and s_{lo} with PHT
 Compute τ_{hi} and τ_{lo} (Equations 8 & 9)
 Configure to system to hi & sleep τ_{hi} .
 Configure to lo & sleep τ_{lo} .

end while

Algorithm 1: CALOREE's runtime control algorithm.

using mathematical optimization techniques. The Netflix approach has been used to predict application response to heterogeneous resources in data centers [12, 13].

Bayesian Predictors: A hierarchical Bayesian model (HBM) provides a statistically sound framework for learning across applications and devices [21, 54]. In the HBM, each application has its own model, allowing specificity, but these models are conditionally dependent on some underlying probability distribution with a hidden mean and co-variance. In practice, an HBM predicts behavior for a new application using a small number of observations and combining those with the large number of observations of other applications. Rather than over-generalizing, the HBM uses only similar applications to predict new application behavior. The HBM's accuracy increases as more applications are observed because increasingly diverse behaviors are represented in the pool of prior knowledge [52]. Of course, the computational complexity of learning also increases with increasing applications.

3.5 Formal Analysis

Control System Complexity CALOREE's control system (see Algorithm 1) runs on the local device along with the application under control, so its overhead must be minimal. In fact, each controller invocation is $O(1)$. The only parts that are not obviously constant time are the PHT lookups. Provided the PHT resolution is sufficiently high to avoid collisions, then each PHT lookup requires constant time.

Control Theoretic Formal Guarantees The controller's pole $\rho(t)$ is critical to providing control theoretic guarantees in the presence of learned—rather than directly measured—data. CALOREE requires any learner estimate not only speedup and powerup, but also the variance σ . CALOREE uses this information to derive a lower bound for the pole which guarantees probabilistic convergence to the desired latency. Specifically, we prove that with probability 99.7% CALOREE converges to the desired latency if the pole is

$$[1 - \lfloor \max(\hat{s}) / (\min(\hat{s}) - 3\sigma) \rfloor_0]_0 \leq \rho(t) < 1,$$

where $\lfloor x \rfloor_0 = \max(x, 0)$ and $\hat{s}$ is the estimated speedup. See appendix A for the proof. Users who need higher confidence can set the scalar multiplier on σ higher; e.g., using 6 provides a 99.99966% probability of convergence.

Thus we provide a lower-bound on the value of $\rho(t)$ required for confidence that CALOREE converges to the desired latency. This pole value only considers latency, and not energy efficiency. In practice, we find it better to use a higher pole based on the *uncertainty* between the controller's observed energy efficiency and that predicted by the learner. We follow prior work [72] in quantifying uncertainty as $\beta(t)$, and setting the pole based on this uncertainty:

$$\begin{aligned} \beta(t) &= \exp\left(-\left(\left|\frac{\bar{s}(t)}{\bar{p}(t)} - \frac{\hat{s}(t)}{\hat{p}(t)}\right|\right)/5\right) \\ \rho(t) &= \frac{1-\beta(t)}{1+\beta(t)} \end{aligned} \quad (10)$$

where $\bar{s}$ and $\bar{p}$ are the measured values of speedup and powerup and $\hat{s}$ and $\hat{p}$ are the estimated values from the learner. This measure of uncertainty captures both power and latency. We find that it is generally higher than the pole value given by our lower bound, so in practice CALOREE sets the pole dynamically to be the higher of the two values and CALOREE makes spot updates to the estimated speedup and power based on its observations.

4 Experimental Setup

4.1 Platform and Benchmarks

We run applications on an ODROID-XU3 with a Samsung Exynos 5 Octa processor (an ARM big.LITTLE architecture), running Ubuntu 14.04. The 4 big cores support 19 clockspeeds, the 4 LITTLE ones have 13. An on-board power meter updated at 0.25s intervals captures core, GPU, network, and memory. We allocate cores using thread affinity and set speeds using `cpufrequtils`. The ODROID has no screen, but recent trends in mobile/embedded processor design and workloads have seen processor power become the dominant factor in energy consumption [23]. We note that the underlying system automatically performs thermal throttling at high clockspeeds, reducing clockspeed when temperature becomes too high. We cannot disable this feature and it actually creates a challenge for the learners, as they must accurately estimate when high clockspeeds will actually reduce performance due to this thermal throttling behavior. We run the learners on an Intel server with E5-2690 processors. The ODROID and the server are connected with Gigabit Ethernet.

We use 12 benchmarks representing embedded and mobile sensor processing. These include video encoding (x264), video analysis (bodytrack), image similarity search (ferret), and animation (facesim) from PARSEC [4]; medical imaging (heartwall, leukocyte), image processing (srad), and machine learning (kmeans) from Rodinia [7]; security (sha) from ParMiBench [32]; memory intensive processing (stream) [49]; and synthetic aperture radar (radar) [26]. These benchmarks are representative of either existing embedded/mobile workloads (video encoding, radar processing) or examples of the emerging class of learning and analysis applications that are being increasingly pushed to edge devices (clustering, video analysis).

Figure 8. Lack-of-fit for performance vs clock-speed. Lower lack-of-fit indicates a more compute-bound application, higher values indicate a memory-bound one.

Figure 8 shows the variety of workloads indicated by the *lack-of-fit*—the absence of correlation between frequency and performance. Applications with high lack-of-fit do not speed up with increasing frequency—typical of memory bound applications. Applications with low lack-of-fit increase performance with increasing clock speed. Applications with intermediate lack-of-fit tend to improve with increasing clock speed up to a point and then stop. Each application has an outer loop which processes one input (e.g., a point for kmeans or a frame for x264). The application signals the completion of an input using a standard API [27]. Performance requirements are specified as latencies for these inputs.

4.2 Evaluation Metrics

For each application, we measure its worst-case execution time (wcet) running without management; *i.e.*, the highest latency for any input. We set a latency goal—or *deadline*—for each input equal to its wcet; the standard approach for ensuring real-time latency guarantees or maximum responsiveness [6]. We quantify performance reliability by measuring the missed deadlines. If the application processes n total inputs and m exceeded the target latency the deadline misses are:

$$\text{deadline misses} = 100\% \cdot \frac{m}{n}. \quad (11)$$

We evaluate energy savings by running every application in every resource configuration and recording performance and power for every input. By post-processing this data we determine the minimal energy resource configuration that meets the latency for each input. To compare across applications, we normalize energy:

$$\text{normalized energy} = 100\% \cdot \left(\frac{e_{\text{measured}}}{e_{\text{optimal}}} - 1 \right) \quad (12)$$

where e_{measured} is measured energy and e_{optimal} is the optimal energy. We subtract 1, so that this metric shows the percentage of energy over optimal.

4.3 Points of Comparison

We compare to existing learning and control approaches:

1. *Race-to-idle*: This well-known heuristic allocates all resources to the application to complete each input as fast as possible, then idles until the next input is available

[34, 36, 53]. This heuristic is a standard way to meet hard deadlines, but it requires conservative resource allocation [6].

2. *PID-Control*: a standard single-input (performance), multiple-output (big/LITTLE core counts and speeds) proportional-integral-controller representative of several that have been proposed for computer resource management [24, 64]. This controller is tuned to provide the best average case behavior across all applications and targets.
3. *Online*: measures a few sample configurations then performs polynomial multivariate regression to estimate unobserved configurations' behavior [43, 52, 58].
4. *Offline*: does not observe the current application—instead using previously observed applications to estimate power and performance as a linear regression [37, 39, 81, 85].
5. *Netflix*: a matrix completion algorithm for the Netflix challenge. Variations of this approach allocate heterogeneous resources in data centers [12, 13].
6. *HBM*: a hierarchical Bayesian learner previously used to allocate resources to meet performance goals with minimal energy in server systems [52].
7. *Adaptive-Control*: a state-of-the-art, adaptive controller that meets application performance with minimal energy [30]. This approach requires a user-specified model relating resource configuration to performance and power. For this paper, we use the *Offline* learner's model.

We compare the above baselines to:

1. *CALOREE-NoPole*: uses the HBM learner, but sets the pole to 0, which shows the importance of incorporating the learned variance into control. All other versions of CALOREE set the pole according to Section 3.5.
2. *CALOREE-online*: uses the online learner.
3. *CALOREE-Netflix*: uses the Netflix learner.
4. *CALOREE-HBM*: uses the HBM learner.

We use leave-one-out cross validation: to test application x , we train on all other applications, then test on x .

5 Experimental Evaluation

5.1 Performance and Energy for Single App

Figure 9a summarizes the average error across all targets for the single application scenario. The figure shows deadline misses in the top chart and energy over optimal in the bottom. The dots show the average, while the error bars show the minimum and maximum values.

Race-to-idle meets all deadlines, but its conservative resource allocation has the highest average energy consumption. Among the prior approaches HBM has the lowest average deadline misses (9%) and lowest energy (20% more than optimal). CALOREE with no pole misses 15% of all deadlines, which is worse than prior approaches. Note that all prior approaches—other than racing—have at least one application that misses all deadlines. In many cases these approaches

Figure 9. Summary data for (a) single- and (b) multi-app scenarios. The top row shows deadline misses, the bottom energy consumption.

Figure 10. Comparison of application performance error and energy for single application scenario.

Figure 11. Comparison of application performance error and energy for multiple application scenario.

are close to the latency (within 10%), but not close enough to deliver reliable performance.

When CALOREE adaptively tune its pole, the results greatly improve. The best combination is CALOREE-HBM, which averages 6.0% missed deadlines, while consuming just 4.3% more energy than optimal. Thus, CALOREE-HBM reduces

average deadline misses by 65% and energy consumption by 13% compared to the best prior approach. The error bars on the CALOREE-HBM approach demonstrate that it is the only approach—besides racing—that handles every test application; all others see at least 100% deadline misses for one test case. Yet, CALOREE-HBM reduces energy consumption by

27% compared to race-to-idle. The energy savings comes because most inputs are not worst case, leaving slack for smart resource allocators to save energy. *Among many smart approaches CALOREE-HBM provides highly reliable performance with very low energy.*

Figure 10 presents a detailed, per-application comparison between CALOREE-HBM and selected prior approaches which have performed well in other scenarios: race-to-idle, Netflix, HBM, and adaptive control. Other data has been omitted for space. The benchmarks are shown on the x-axis; the y-axis shows the number of deadline misses and the normalized energy, respectively.

We thoroughly evaluate sensitivity to the latency goal in Appendix B. In brief, we find that these general trends are true across a wide range of latency goals. Furthermore, while it is beyond the scope of this paper, we have ported CALOREE to an Intel server with many more states and, despite the larger search space, we find that CALOREE still produces better results than learning or control alone [50].

5.2 Performance and Energy for Multiple Apps

We again launch each benchmark with a goal of meeting its worst case latency. One quarter of the way through execution, we start another application randomly drawn from our benchmark set—bound to one big core—which interferes with the original application. Delivering the required latency tests the ability to react to environmental changes.

Figure 9b shows the average number of deadline misses and energy over optimal for all approaches. Some targets are unachievable for some applications; specifically, bodytrack, hear twall, and sha. Due to these unachievable targets, both optimal and race-to-idle show some deadline misses. Race-to-idle misses more deadlines than optimal because it cannot make use of LITTLE cores to do some work, it simply continues using all big cores despite the degraded performance due to the second application. Most approaches do badly in this scenario—even adaptive control has 40% deadline misses. CALOREE-HBM produces the lowest deadline misses with an average of 20%, which is only 2 points more than optimal. It also produces the lowest energy, just 6% more than optimal. Figure 11 shows the detailed results.

The multi-application scenario demonstrates that CALOREE can adapt to large, unpredictable changes. Neither CALOREE's learner nor controller are ever explicitly aware that a new application has entered the system. CALOREE, however, immediately detects the disturbance as a change in the observed latency and then adjusts resource allocation to bring the application back to its target performance.

These results assume that single application is the highest priority in the system. CALOREE “protects” that single application from interference by other applications. In future work, we will investigate applying CALOREE to competing applications by extending prior control work that addresses

competing application needs while assuming fully accurate models of all applications that might run together [48].

5.3 Adapting to Workload Changes

Figure 12. Controlling x264 through scene changes.

We compare CALOREE and HBM reacting to input variations. Figure 12 shows the x264 video encoder with 2 different phases caused by a scene at the 500th frame. The first scene is difficult, the second much easier. In the first, both the HBM and CALOREE find a configuration that achieves the latency target (1 in the figure) with minimal energy. When the input changes, CALOREE initially misses the latency, then adjusts to an optimal configuration. In contrast, the HBM does not find a new configuration, but idles more. During the second scene, CALOREE operates at 1.7W, while the HBM is at 2W. Here, CALOREE's use of learning and control reduces energy by 14% compared to learning alone.

5.4 The Pole's Importance

Section 3.5 argues that tuning the controller to learned variance prevents oscillation and provides probabilistic guarantees despite using noisy, learned data to control unseen applications. We demonstrate this empirically by showing srdd using both CALOREE-NoPole and CALOREE-HBM. Figure 13 shows time on the x-axis and normalized latency and power on the y-axes. CALOREE-NoPole oscillates and causes wide power fluctuations. In contrast, CALOREE provides reliable performance and saves tremendous energy because it avoids oscillation, using a mixture of big and LITTLE cores to minimize energy.

Figure 13. Comparison of learned and default poles.

Figure 14. Estimation accuracy versus sample size.

5.5 Sensitivity to the Measured Samples

We show how the number of samples affects model accuracy for the Online, Netflix, and HBM learners. We quantify accuracy as how close the learner is to ground truth (found through exhaustive exploration), with 1 meaning the learner perfectly models the real performance or power. Accuracy matters because the fewer the samples, the faster the controller switches to the learner’s application-specific model.

Figure 14 shows the accuracy vs sample count for both performance (top) and power (bottom). The HBM incorporates prior knowledge and its accuracy uniformly improves with more samples—exceeding 0.9 after 20 samples. The Online approach needs at least 7 samples to even generate a prediction. As Online receives more samples, its accuracy improves but never exceeds HBM’s for the same sample count. Netflix is very noisy for small sample sizes, but after about 50, it is competitive with HBM. These results not only demonstrate the sensitivity to sample size, they show why CALOREE-HBM achieves the best results.

5.6 Overhead

CALOREE’s main overhead is sampling, where the controller tests a few configurations before CALOREE can reliably estimate the entire power and performance frontier. The sampling cost can be distributed across devices by asking each of them to contribute samples for estimation. Once the sampling phase is over, the HBM generates an estimate in at most 500 ms, which is significantly smaller than the time required to run any of our applications. In the worst case (facesim), the controller sends 320B of sample data to the learner, which sends back 1KB. In this case, the sampling overhead and communication cost is less than 2% of total execution time. CALOREE’s asynchronous communication means that the controller never waits for the learner. For all other benchmarks it is lower, and for most it is negligible.

The controller requires only a few floating point operations to execute, plus the table lookups in the PHT. To evaluate its overhead, we time 1000 iterations. We find that it is under 2 microseconds, which is significantly faster than we can change any resource allocation on our system; the

controller has negligible impact on performance and energy consumption of the controlled device.

6 Related Work

Energy has long been an important resource for mobile and embedded computing. Several OSs make energy an allocatable resource [62, 65, 66]. Others have specialized OS constructs to monitor [20] and reduce [18, 25, 40, 74, 80] energy for mobile and embedded applications. We examine related work applying learning and control to energy management. **Offline Learning** approaches build predictors before deployment and then use those fixed predictors to allocate resources [2, 9, 14, 37, 39, 79]. The training requires both many samples and substantial computation. Applying the predictor online, however, is low overhead. The main drawback is that the predictions are not updated as the system runs: a problem for adapting workloads. Carat is an offline learner that aggregates data across multiple devices to generate a report for human users about how to reconfigure their device for energy savings [55]. While both Carat and CALOREE learn across devices, they have very different goals. Carat returns very high-level information to human users; e.g., update a driver to extend battery life. CALOREE automatically builds and applies low-level predictions to save energy.

Online Learning techniques observe the current application to tune system resource usage for that application [1, 38, 43, 57, 58, 68]. For example, Flicker is a configurable architecture and optimization framework that uses online prediction to maximize performance under a power limitation [57]. Another example, ParallelismDial, uses online adaptation to tailor parallelism to application workload [68].

Hybrid Approaches combine offline predictions with online adaptation [10, 15, 62, 66, 76, 78, 83]. For example, Dubach et al. use a hybrid scheme to optimize the microarchitecture of a single core [15]. Such predictors have also been employed at the operating system level to manage system energy consumption [62, 66, 78]. Other approaches combine offline prediction with online updates [5, 25, 31]. For example, Bitirgen et al use an artificial neural network to allocate resources to multiple applications in a multicore [5]. The neural network is trained offline and then adapted online to maximize performance but without considering energy.

Control solutions can be thought of as a combination of offline prediction with online adaptation. Their formal properties make them attractive for managing resources in operating systems [24, 33, 69]. The offline phase involves substantial empirical measurement that is used to synthesize a control system [8, 30, 42, 60, 61, 67, 77, 80, 82]. Control solutions work well over a narrow range of applications, as the rigorous offline measurement captures the general behavior of a class of application and require negligible online overhead. This focused approach is extremely effective

for multimedia applications [18, 19, 35, 47, 74, 80] and web-servers [29, 45, 70] because the workloads can be characterized ahead of time to produce sound control.

The need for good predictions is the central tension in developing control for computing systems. It is always possible to build a controller for a specific application and system by specializing for that pair. Prior work addresses the need for accurate predictions in various ways. Some provides control libraries that require user-specified models [22, 30, 61, 67, 82]. Others automatically synthesize both a predictor and a controller for either hardware [59] or software [16, 17]. JouleGuard combines learning for energy efficiency with control for managing application parameters [25]. In JouleGuard, a learner adapts the controller's coefficients to uncertainty, but JouleGuard does not produce a new set of predictions. JouleGuard's computationally efficient learner runs on the same device as the controlled application, but it cannot identify correlations across applications or even different resource configurations. CALOREE is unique in that a separate learner generates an application-specific predictions automatically. By offloading the learning task, CALOREE (1) combines data from many applications and systems and (2) applies computationally expensive, but highly accurate learning techniques.

Combining Learning and Control Two recent projects explore such a combination. Recht et al have proposed several approaches for combining statistical learning models with optimal control theory [11, 73]. Simultaneously, Hoffmann et al have developed OS- [25] and hardware-level resource management systems [63, 84] that combine learning and control to provide both energy and latency guarantees in dynamic environments. This prior work, however, still requires expertise in both learning and control methods to effectively deploy the proposed solution. CALOREE, however, defines abstractions that allow a number of AI and learning techniques to be combined with an adaptive controller, maintaining control-theoretic formal guarantees. CALOREE requires no user specified parameters, other than the goal, allowing it to be used by non-experts.

7 Conclusion

Much recent work builds systems to support learning, CALOREE uses learning to build better systems. CALOREE is a resource manager that meets application latency requirements with minimal energy, even without prior knowledge of the application. CALOREE is the first work that provides formal guarantees that it will converge to the required latency despite not having prior knowledge. CALOREE achieves this breakthrough by using learning to model complex resource interaction and control theory to manage system dynamics. CALOREE proposes foundational techniques that allow control to be applied using noisy learned models—instead of ground truth models—while maintaining formal guarantees. We demonstrate CALOREE's effectiveness with a case study

using embedded applications on a heterogeneous processor. Compared to prior learning and control approaches, CALOREE is the only approach that provides reliable latency for all applications with near minimal energy. This ability to meet goals for applications without prior knowledge is a key-enabler for self-aware computing systems [28].

Acknowledgments

We thank the anonymous reviewers for their insightful feedback. We thank Shan Lu and Fred Chong for improving early drafts of the manuscript. This research is supported by NSF (CCF-1439156, CNS-1526304), and generous support from the CERES Center for Unstoppable Computing. Additional support comes from the Proteus project under the DARPA BRASS program and a DOE Early Career award.

A Probabilistic Convergence Guarantees

Theorem A.1. *Let s_c and $\hat{s}_c$ denote the true and estimated speedups of various configurations in set C as $\mathbf{c} \in \mathbb{R}^{|C|}$. Let σ denote the estimation error for speedups such that, $\hat{s}_i \sim N(s_i, \sigma^2) \forall i$. We show that with probability greater than 99.7%, the pole $\rho(t)$ can be chosen to lie in the range, $[1 - \lfloor \max(\hat{s}) / (\min(\hat{s}) - 3\sigma) \rfloor_0, 1)$, where $\lfloor x \rfloor_0 = \max(x, 0)$.*

Proof. Let Δ denote the multiplicative error over speedups, such that $\hat{s}_c \Delta = s_c$. To guarantee convergence the value of pole, $\rho(t)$ can vary in the range $[1 - \frac{2}{\Delta}]_0, 1$ [16]. The lower $\rho(t)$, the faster the convergence. Equations 8 & 9 show that any $s(t)$ is a linear combination of two speedups:

$$s(t) = \hat{s}_{hi} \cdot \tau_{hi} + \hat{s}_{lo} \cdot (T - \tau_{hi}) \quad (13)$$

$$\hat{s}(t) = s_{hi} \cdot \tau_{hi} + s_{lo} \cdot (T - \tau_{hi}) \quad (14)$$

We can upper bound and lower bound each of these terms,

$$s(t) \leq T\hat{s}_{hi} \text{ and } \hat{s}(t) \geq Ts_{lo} \quad (15)$$

The speedup estimates are close to the actual speedups since $\hat{s} \sim N(s, \sigma^2)$, therefore with probability greater than 99.7% and the speedups can be given by, $s_{lo} \geq \hat{s}_{lo} - 3\sigma$. Hence, $\hat{s}(t) \geq T(\hat{s}_{lo} - 3\sigma)$. Since, over all configurations, $\Delta \leq \lfloor \max(\hat{s}) / (\min(\hat{s}) - 3\sigma) \rfloor_0$, we can choose the pole from the range, $([1 - \lfloor \max(\hat{s}) / (\min(\hat{s}) - 3\sigma) \rfloor_0, 1)$. $\square$

B Sensitivity to Latency Target

The results in the main body of the paper set a latency target that is equivalent to the worst observed latency for any input. In this section we explore the sensitivity of our results to the latency target itself. In general, we want to answer the question of whether the results are still good when less aggressive latency targets are used. Therefore, we set a range of performance targets from 1.1–3.0× the worst case latency. (Note that the earlier results set the latency target equal to the worst case). As before, we measure deadline misses and energy over optimal for all points of comparison. Figure 15

Figure 15. Summary data for single-app scenario averaging across many different latency targets from 1.1 to 3.0× the worst case latency.

represents the summary results as an average error across all targets for the single application scenario. This figure shows two charts with the percentage of deadline misses in the top chart and the energy over optimal in the bottom. The dots show the average for each technique, while the error bars show the minimum and maximum values.

Not surprisingly, race-to-idle meets all deadlines, but its conservative resource allocation has the highest average

energy consumption. Among the prior learning approaches Netflix has the lowest average deadline misses (11%), but with high energy (40% more than optimal), while the HBM has higher deadline misses (17%) but with significantly lower energy consumption (16%). Adaptive control achieves similar deadline misses (14%) with lower average energy than any of the prior learning approaches (12%). CALOREE with no pole misses 45% of all deadlines, which is clearly unacceptable.

When we allow CALOREE to adaptively tune its pole, however, we see greatly improved results. The best combination is CALOREE with the HBM, which misses only 5.5% of deadlines on average, while consuming just 4.4% more energy than optimal. These numbers represent large improvements in both performance reliability and energy efficiency compared to prior approaches. The other learners paired with CALOREE achieve similar results to the prior adaptive control approach.

This data confirms that the trends reported in the body of the paper hold across a range of deadlines. The major difference between this data and that in the main body of the paper is that even relaxing deadlines slightly makes it much less likely that an approach will completely fail to meet the deadlines. Detailed results for individual applications are available in an extended version of this paper [50, 51].

References

- [1] Jason Ansel, Maciej Pacula, Yee Lok Wong, Cy Chan, Marek Olszewski, Una-May O'Reilly, and Saman Amarasinghe. 2012. Siblingrivalry: online autotuning through local competitions. In *CASES*.
- [2] Jason Ansel, Yee Lok Wong, Cy Chan, Marek Olszewski, Alan Edelman, and Saman Amarasinghe. 2011. Language and compiler support for auto-tuning variable-accuracy algorithms. In *CGO*.
- [3] R. M. Bell, Y. Koren, and C. Volinsky. 2008. *The BellKor 2008 solution to the Netflix Prize*. Technical Report. ATandT Labs.
- [4] C. Bienia, S. Kumar, J. P. Singh, and K. Li. 2008. The PARSEC Benchmark Suite: Characterization and Architectural Implications. In *PACT*.
- [5] Ramazan Bitirgen, Engin Ipek, and Jose F. Martinez. 2008. Coordinated management of multiple interacting resources in chip multiprocessors: A machine learning approach. In *MICRO*.
- [6] Giorgio C Buttazzo, Giuseppe Lipari, Luca Abeni, and Marco Caccamo. 2006. *Soft Real-Time Systems: Predictability vs. Efficiency: Predictability vs. Efficiency*. Springer.
- [7] Shuai Che, Michael Boyer, Jiayuan Meng, David Tarjan, Jeremy W. Sheaffer, Sang-Ha Lee, and Kevin Skadron. 2009. Rodinia: A Benchmark Suite for Heterogeneous Computing. In *IISWC*.
- [8] Jian Chen and Lizy Kurian John. 2011. Predictive coordination of multiple on-chip resources for chip multiprocessors. In *JCS*.
- [9] Jian Chen, Lizy Kurian John, and Dimitris Kaseridis. 2011. Modeling Program Resource Demand Using Inherent Program Characteristics. *SIGMETRICS Perform. Eval. Rev.* 39, 1 (June 2011), 1–12.
- [10] Ryan Cochran, Can Hankendi, Ayse K. Coskun, and Sherief Reda. 2011. Pack & Cap: adaptive DVFS and thread packing under power caps. In *MICRO*.
- [11] Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. 2017. *On the Sample Complexity of the Linear Quadratic Regulator*. Technical Report 1710.01688v1. arXiv.
- [12] Christina Delimitrou and Christos Kozyrakis. 2013. Paragon: QoS-aware Scheduling for Heterogeneous Datacenters. In *ASPLOS*.
- [13] Christina Delimitrou and Christos Kozyrakis. 2014. Quasar: Resource-efficient and QoS-aware Cluster Management. In *ASPLOS*.
- [14] Zhaoxia Deng, Lunkai Zhang, Nikita Mishra, Henry Hoffmann, and Fred Chong. 2017. Memory Cocktail Therapy: A General Learning-Based Framework to Optimize Dynamic Tradeoffs in NVM. In *MICRO*.
- [15] Christophe Dubach, Timothy M. Jones, Edwin V. Bonilla, and Michael F. P. O'Boyle. 2010. A Predictive Model for Dynamic Microarchitectural Adaptivity Control. In *MICRO*.
- [16] Antonio Filieri, Henry Hoffmann, and Martina Maggio. 2014. Automated design of self-adaptive software with control-theoretical formal guarantees. In *ICSE*.
- [17] Antonio Filieri, Henry Hoffmann, and Martina Maggio. 2015. Automated multi-objective control for self-adaptive software design. In *FSE*.
- [18] J. Flinn and M. Satyanarayanan. 1999. Energy-aware adaptation for mobile applications. In *SOSP*.
- [19] Jason Flinn and M. Satyanarayanan. 2004. Managing battery lifetime with energy-aware adaptation. *ACM Trans. Comp. Syst.* 22, 2 (May 2004).
- [20] Rodrigo Fonseca, Prabal Dutta, Philip Levis, and Ion Stoica. 2008. Quanto: Tracking Energy in Networked Embedded Systems. In *OSDI*.
- [21] Andrew Gelman, John B Carlin, Hal S Stern, David B Dunson, Aki Vehtari, and Donald B Rubin. 2013. *Bayesian data analysis*. CRC press.
- [22] Ashvin Goel, David Steere, Calton Pu, and Jonathan Walpole. 1998. SWIFT: A Feedback Control and Dynamic Reconfiguration Toolkit. In *2nd USENIX Windows NT Symposium*.
- [23] Matthew Halpern, Yuhao Zhu, and Vijay Janapa Reddi. [n. d.]. Mobile CPU's rise to power: Quantifying the impact of generational mobile CPU design trends on performance, energy, and user satisfaction. In *HPCA*.
- [24] Joseph L. Hellerstein, Yixin Diao, Sujay Parekh, and Dawn M. Tilbury. 2004. *Feedback Control of Computing Systems*. John Wiley & Sons.
- [25] Henry Hoffmann. 2015. JouleGuard: energy guarantees for approximate applications. In *SOSP*.
- [26] Henry Hoffmann, Anant Agarwal, and Srinivas Devadas. 2012. Selecting Spatiotemporal Patterns for Development of Parallel Applications. *IEEE Trans. Parallel Distrib. Syst.* 23, 10 (2012).
- [27] Henry Hoffmann, Jonathan Eastep, Marco D. Santambrogio, Jason E. Miller, and Anant Agarwal. 2010. Application Heartbeats: a generic interface for specifying program performance and goals in autonomous computing environments. In *ICAC*.
- [28] Henry Hoffmann, Jim Holt, George Kurian, Eric Lau, Martina Maggio, Jason E. Miller, Sabrina M. Neuman, Mahmut Sinangil, Yildiz Sinangil, Anant Agarwal, Anantha P. Chandrakasan, and Srinivas Devadas. 2012. Self-aware computing in the Angstrom processor. In *DAC*.
- [29] T. Horvath, T. Abdelzaher, K. Skadron, and Xue Liu. 2007. Dynamic Voltage Scaling in Multitier Web Servers with End-to-End Delay Control. *Computers, IEEE Transactions on* 56, 4 (2007).
- [30] Connor Imes, David H. K. Kim, Martina Maggio, and Henry Hoffmann. 2015. POET: A Portable Approach to Minimizing Energy Under Soft Real-time Constraints. In *RTAS*.
- [31] Engin Ipek, Onur Mutlu, José F. Martínez, and Rich Caruana. 2008. Self-Optimizing Memory Controllers: A Reinforcement Learning Approach. In *ISCA*.
- [32] Syed Muhammad Zeeshan Iqbal, Yuchen Liang, and Hakan Grahm. 2010. ParMiBench - An Open-Source Benchmark for Embedded Multiprocessor Systems. *IEEE Comput. Archit. Lett.* 9, 2 (July 2010).
- [33] C. Karamanolis, M. Karlsson, and X. Zhu. 2005. Designing controllable computer systems. In *HotOS*. Berkeley, CA, USA.
- [34] David H. K. Kim, Connor Imes, and Henry Hoffmann. 2015. Racing and Pacing to Idle: Theoretical and Empirical Analysis of Energy Optimization Heuristics. In *CPSNA*.
- [35] Minyoung Kim, Mark-Oliver Stehr, Carolyn Talcott, Nikil Dutt, and Nalini Venkatasubramanian. 2013. xTune: A Formal Methodology for Cross-layer Tuning of Mobile Embedded Systems. *ACM Trans. Embed. Comput. Syst.* 11, 4 (Jan. 2013).
- [36] Etienne Le Sueur and Gernot Heiser. 2011. Slow Down or Sleep, That is the Question. In *Proceedings of the 2011 USENIX Annual Technical Conference*. Portland, OR, USA.
- [37] B.C. Lee, J. Collins, Hong Wang, and D. Brooks. 2008. CPR: Composable performance regression for scalable multiprocessor models. In *MICRO*.
- [38] Benjamin C. Lee and David Brooks. 2008. Efficiency Trends and Limits from Comprehensive Microarchitectural Adaptivity. In *ASPLOS*.
- [39] Benjamin C. Lee and David M. Brooks. 2006. Accurate and Efficient Regression Modeling for Microarchitectural Performance and Power Prediction. In *ASPLOS*.
- [40] Matthew Lentz, James Litton, and Bobby Bhattacharjee. 2015. Drowsy Power Management. In *SOSP*.
- [41] W.S. Levine. 2005. *The control handbook*. CRC Press.
- [42] Baochun Li and K. Nahrstedt. 1999. A control-based middleware framework for quality-of-service adaptations. *IEEE Journal on Selected Areas in Communications* 17, 9 (1999).
- [43] J. Li and J.F. Martinez. 2006. Dynamic power-performance adaptation of parallel computation on chip multiprocessors. In *HPCA*.
- [44] Lennart Ljung. 1999. *System Identification: Theory for the User*. Prentice Hall PTR, Upper Saddle River, NJ, USA.
- [45] C. Lu, Y. Lu, T.F. Abdelzaher, J.A. Stankovic, and S.H. Son. 2006. Feedback Control Architecture and Design Methodology for Service Delay Guarantees in Web Servers. *IEEE TPDS* 17, 9 (September 2006), 1014–1027.
- [46] Martina Maggio, Henry Hoffmann, Alessandro V. Papadopoulos, Jacopo Panerati, Marco D. Santambrogio, Anant Agarwal, and Alberto Leva. 2012. Comparison of Decision-Making Strategies for Self-Optimization in Autonomic Computing Systems. *ACM Trans. Auton. Adapt. Syst.* 7, 4, Article 36 (Dec. 2012), 32 pages. <https://doi.org/10.1145/2187888>

- //doi.org/10.1145/2382570.2382572
- [47] M. Maggio, H. Hoffmann, M. D. Santambrogio, A. Agarwal, and A. Leva. 2013. Power Optimization in Embedded Systems via Feedback Control of Resource Allocation. *IEEE Transactions on Control Systems Technology* 21, 1 (Jan 2013).
 - [48] Martina Maggio, Alessandro Vittorio Papadopoulos, Antonio Filieri, and Henry Hoffmann. 2017. Automated Control of Multiple Software Goals Using Multiple Actuators. In *ESEC/FSE*.
 - [49] John D. McCalpin. 1995. Memory Bandwidth and Machine Balance in Current High Performance Computers. *IEEE TCCA Newsletter* (Dec. 1995), 19–25.
 - [50] Nikita Mishra. 2017. *Statistical Methods for Improving Dynamic Scheduling and Resource Usage in Computing Systems*. Ph.D. Dissertation. <https://search.proquest.com/docview/1928485902?accountid=14657>
 - [51] Nikita Mishra, Connor Imes, Huazhe Zhang, John D Lafferty, and Henry Hoffmann. 2016. *Big Data for LITTLE Cores: Combining Learning and Control for Mobile Energy Efficiency*. Technical Report TR-2016-10. University of Chicago, Dept. of Comp. Sci.
 - [52] Nikita Mishra, Huazhe Zhang, John D. Lafferty, and Henry Hoffmann. 2015. A Probabilistic Graphical Model-based Approach for Minimizing Energy Under Performance Constraints. In *ASPLOS*.
 - [53] Akihiko Miyoshi, Charles Lefurgy, Eric Van Hensbergen, Ram Rajamony, and Raj Rajkumar. 2002. Critical Power Slope: Understanding the Runtime Effects of Frequency Scaling. In *ICS*.
 - [54] Carl N Morris. 1983. Parametric empirical Bayes inference: theory and applications. *J. Amer. Statist. Assoc.* 78, 381 (1983), 47–55.
 - [55] Adam J. Oliner, Anand P. Iyer, Ion Stoica, Eemil Lagerspetz, and Sasu Tarkoma. 2013. Carat: Collaborative Energy Diagnosis for Mobile Devices. In *Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems (SenSys '13)*. ACM, New York, NY, USA, Article 10, 14 pages. <https://doi.org/10.1145/2517351.2517354>
 - [56] Sinno Jialin Pan and Qiang Yang. 2010. A Survey on Transfer Learning. *IEEE Trans. on Knowl. and Data Eng.* 22, 10 (Oct. 2010), 1345–1359. <https://doi.org/10.1109/TKDE.2009.191>
 - [57] Paula Petrica, Adam M. Izraelevitz, David H. Albonesi, and Christine A. Shoemaker. 2013. Flicker: A Dynamically Adaptive Architecture for Power Limited Multicore Systems. In *ISCA*.
 - [58] Dmitry Ponomarev, Gurhan Kucuk, and Kanad Ghose. 2001. Reducing Power Requirements of Instruction Scheduling Through Dynamic Allocation of Multiple Datapath Resources. In *MICRO*.
 - [59] Raghavendra Pothukuchi, Amin Ansari, Petros Voulgaris, and Josep Torrellas. 2016. Using Multiple Input, Multiple Output Formal Control to Maximize Resource Efficiency in Architectures. In *ISCA*.
 - [60] Ramya Raghavendra, Parthasarathy Ranganathan, Vanish Talwar, Zhikui Wang, and Xiaoyun Zhu. 2008. No "power" struggles: coordinated multi-level power management for the data center. In *ASPLOS*.
 - [61] R. Rajkumar, C. Lee, J. Lehoczkzy, and Dan Siewiorek. 1997. A resource allocation model for QoS management. In *RTSS*.
 - [62] Arjun Roy, Stephen M. Rumble, Ryan Stutsman, Philip Levis, David Mazières, and Nikolai Zeldovich. 2011. Energy Management in Mobile Devices with the Cinder Operating System. In *EuroSys*.
 - [63] Muhammad Husni Santriaji and Henry Hoffmann. 2016. GRAPE: Minimizing energy for GPU applications with performance requirements. In *MICRO*.
 - [64] Akbar Sharifi, Shekhar Srikantaiah, Asit K. Mishra, Mahmut Kandemir, and Chita R. Das. 2011. METE: meeting end-to-end QoS in multicores through system-wide resource management. In *SIGMETRICS*.
 - [65] Kai Shen, Arrvinth Shriraman, Sandhya Dwarkadas, Xiao Zhang, and Zhuan Chen. 2013. Power Containers: An OS Facility for Fine-grained Power and Energy Management on Multicore Servers. *SIGPLAN Not.* 48, 4 (March 2013), 65–76. <https://doi.org/10.1145/2499368.2451124>
 - [66] David C. Snowdon, Etienne Le Sueur, Stefan M. Petters, and Gernot Heiser. 2009. Koala: A Platform for OS-level Power Management. In *EuroSys*.
 - [67] Michal Sojka, Pavel Pisa, Dario Faggioli, Tommaso Cucinotta, Fabio Checoni, Zdenek Hanzálek, and Giuseppe Lipari. 2011. Modular software architecture for flexible reservation mechanisms on heterogeneous resources. *Journal of Systems Architecture* 57, 4 (2011).
 - [68] Srinath Sridharan, Gagan Gupta, and Gurindar S. Sohi. 2013. Holistic Run-time Parallelism Management for Time and Energy Efficiency. In *ICS*.
 - [69] David C. Steere, Ashvin Goel, Joshua Gruenberg, Dylan McNamee, Calton Pu, and Jonathan Walpole. 1999. A Feedback-driven Proportion Allocator for Real-rate Scheduling. In *Proceedings of the Third Symposium on Operating Systems Design and Implementation (OSDI '99)*. USENIX Association, Berkeley, CA, USA, 145–158. <http://dl.acm.org/citation.cfm?id=296806.296820>
 - [70] Q. Sun, G. Dai, and W. Pan. 2008. LPV Model and Its Application in Web Server Performance Control. In *ICSSSE*.
 - [71] G. Tesauro. 2007. Reinforcement Learning in Autonomic Computing: A Manifesto and Case Studies. *IEEE Internet Computing* 11 (2007). Issue 1.
 - [72] Michel Tokic. 2010. Adaptive ϵ -Greedy Exploration in Reinforcement Learning Based on Value Differences. In *KI*.
 - [73] Stephen Tu and Benjamin Recht. 2017. *Least-Squares Temporal Difference Learning for the Linear Quadratic Regulator*. Technical Report 1712.08642v1. arXiv.
 - [74] Vibhore Vardhan, Wanghong Yuan, Albert F. Harris III, Sarita V. Adve, Robin Kravets, Klara Nahrstedt, Daniel Grobe Sachs, and Douglas L. Jones. 2009. GRACE-2: integrating fine-grained application adaptation with global adaptation for saving energy. *IJES* 4, 2 (2009).
 - [75] Greg Welch and Gary Bishop. [n. d.]. *An Introduction to the Kalman Filter*. Technical Report TR 95-041. UNC Chapel Hill, Department of Computer Science.
 - [76] Jonathan A. Winter, David H. Albonesi, and Christine A. Shoemaker. 2010. Scalable thread scheduling and global power management for heterogeneous many-core architectures. In *PACT*.
 - [77] Qiang Wu, Philo Juang, Margaret Martonosi, and Douglas W. Clark. 2004. Formal online methods for voltage/frequency control in multiple clock domain microprocessors. In *ASPLOS*.
 - [78] Weidan Wu and Benjamin C Lee. 2012. Inferred models for dynamic and sparse hardware-software spaces. In *Microarchitecture (MICRO), 2012 45th Annual IEEE/ACM International Symposium on*. IEEE, 413–424.
 - [79] Joshua J. Yi, David J. Lilja, and Douglas M. Hawkins. 2003. A Statistically Rigorous Approach for Improving Simulation Methodology. In *HPCA*.
 - [80] Wanghong Yuan and Klara Nahrstedt. 2003. Energy-efficient soft real-time CPU scheduling for mobile multimedia systems. In *SOSP*.
 - [81] Huazhe Zhang and Henry Hoffmann. 2016. Maximizing Performance Under a Power Cap: A Comparison of Hardware, Software, and Hybrid Techniques. In *ASPLOS*.
 - [82] R. Zhang, C. Lu, T.F. Abdelzaher, and J.A. Stankovic. 2002. ControlWare: A middleware architecture for Feedback Control of Software Performance. In *ICDCS*.
 - [83] Xiao Zhang, Rongrong Zhong, Sandhya Dwarkadas, and Kai Shen. 2012. A Flexible Framework for Throttling-Enabled Multicore Management (TEMM). In *ICPP*.
 - [84] Yanqi Zhou, Henry Hoffmann, and David Wentzlaff. 2016. CASH: Supporting IaaS Customers with a Sub-core Configurable Architecture. In *ISCA*.
 - [85] Yuhao Zhu and Vijay Janapa Reddi. 2013. High-performance and energy-efficient mobile web browsing on big/little systems. In *HPCA*.