Statistical Learning in Chip (SLIC)

(*Invited Paper*)

Ronald D. Blanton, Xin Li, Ken Mai, Diana Marculescu, Radu Marculescu, Jeyanandh Paramesh, Jeff Schneider, and Donald E. Thomas

Electrical and Computer Engineering Department
Carnegie Mellon University
Pittsburgh, PA 15213, USA
{blanton, xinli, kenmai, dianam, radum, paramesh, thomas}@ece.cmu.edu and schneide@cs.cmu.edu

***Abstract*—Despite best efforts, integrated systems are "born" (manufactured) with a unique 'personality' that stems from our inability to precisely fabricate their underlying circuits, and create software a priori for controlling the resulting uncertainty. It is possible to use sophisticated test methods to identify the best-performing systems but this would result in unacceptable yields and correspondingly high costs. The system personality is further shaped by its environment (e.g., temperature, noise and supply voltage) and usage (i.e., the frequency and type of applications executed), and since both can fluctuate over time, so can the system's personality. Systems also "grow old" and degrade due to various wear-out mechanisms (e.g., negative-bias temperature instability), and unexpectedly due to various early-life failure sources. These "nature and nurture" influences make it extremely difficult to design a system that will operate optimally for all possible personalities. To address this challenge, we propose to develop statistical learning in-chip (SLIC). SLIC is a holistic approach to integrated system design based on continuously learning key personality traits on-line, for self-evolving a system to a state that optimizes performance hierarchically across the circuit, platform, and application levels. SLIC will not only optimize integrated-system performance but also reduce costs through yield enhancement since systems that would have before been deemed to have weak personalities (unreliable, faulty, etc.) can now be recovered through the use of SLIC.**

***Keywords*— Integrated system design; low-power design; statistical and machine learning**

I. INTRODUCTION

The most challenging problems in science and engineering are so incredibly complex that many have turned to statistical learning (SL) to derive accurate models from various forms of empirical data. Major advances in SL have resulted in algorithms that can now cope with significant amounts of high-dimensional data, and most importantly, are sufficiently robust to rely upon in critical applications. A popular use of SL is in two-step process optimization. The first step learns a model that approximates the relationship between system parameters and the resulting system performance. This model is constructed on-line from data collected during system operation. The second step uses active learning where the learned performance model is analyzed to determine which parameter settings to try next. Active learning trades off the need to experiment untested areas of the parameter space in order to gain more information for learning, against the objective of selecting parameters that are likely to yield optimal performance. Often the "learner" must accomplish this in the face of non-stationarity.

The design, manufacture and operational characteristics (e.g., yield, performance, reliability, power, security, etc.) of modern integrated systems also exhibit extreme levels of complexity that similarly cannot be easily modeled or predicted from first principles. In this nanoscale era, manufacturers find it increasingly difficult to control fabrication, thus making every aspect of design (circuits, logic, memory, communication networks, cores/uncores, etc.) a grand challenge. Moreover, the operating environment of a system which is characterized by temperature, supply voltage, the amount of noise, etc. also adds a level uncertainty that is extremely difficult to deal with optimally at the time of design. Finally, the fact that the use of an integrated system may vary widely from user to user adds yet another major source of uncertainty. These sources all combine together in the worst possible ways to establish an overall level of uncertainty that leads to systems that exhibit non-optimal performance, or require excessive resources to design and fabricate. For example, modern portable, multimedia devices such as a tablet computer require millions of engineering hours to integrate several SoCs (systems on chip) including multiple radios, DSPs, μPs, application-specific processors, display drivers, and solid-state memories, altogether which execute a variety of applications. The uncertainty exhibited by the integration and use of various heterogeneous sub-systems within diverse environments can be better optimized by learning and then adapting.

II. STATISTICAL LEARNING IN CHIP

We propose to develop new SL algorithms that enable an integrated system to learn and adapt operation across the system stack (i.e., the circuit, platform, and application levels) [1]. Conventional approaches to SL assume learning takes place on server farms characterized by virtually unlimited compute and storage resources. Integrated systems, on the other hand, have stringent constraints on power and security, and thus require a more compressed learning cycle which means a complete re-thinking of SL is necessary for it to be effective within an integrated-system environment.

While there is a great deal of active research that individually addresses each source of uncertainty within an integrated system, we instead want to tackle them all simultaneously using a universal solution that we call a self-evolving system. A self-evolving system has the ability to

This work has been supported in part by the National Science Foundation of the United States under contract no. 1314876.

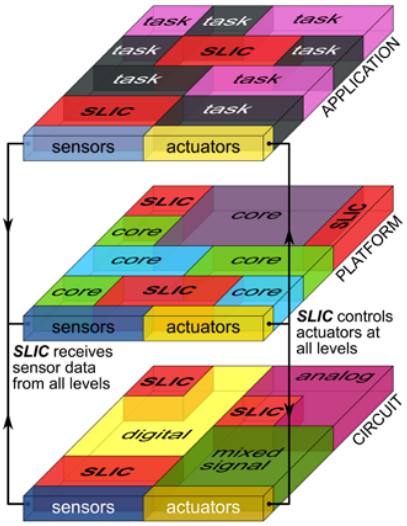

Figure 1: A self-evolving system.

adapt and evolve to changes and unknowns encountered both at the time of manufacturing and over the lifetime of the system. Figure 1 shows a hierarchical view of a self-evolving system that consists of a bottom layer of circuits that together form a variety of cores at the middle (platform) layer, which are used by various tasks at the top layer to execute one or more applications. Sensors at the circuit, platform, and application levels collect various forms of data so that the state of the system, from various perspectives, can be learned. This global view is used to evolve the system into a new state (via the various actuators shown) for improving a wide-variety of system attributes (performance, power, reliability, etc.). At the circuit level, temperature, voltage, and frequency are likely quantities to be sensed. For a platform of heterogeneous cores (e.g., computation, memory, etc.), this same data would be useful but in addition, various real-time measures of workload, queue occupancy, communication among different modules, etc. would be collected as well. Finally, for an application viewed as a plethora of interacting tasks, sensors of various sorts are likely to already exist for enabling the application. For instance, the brain-computer interface (BCI) application discussed later has a sensor array for measuring neural signals. All of this sensed data is provided to the SLIC (statistical learning in chip) cores for learning how to improve the behavior of the entire system which is achieved by providing new parameter values to the level-specific actuators shown in Figure 1. SLIC cores (SCs) perform the learning and are employed at the circuit, platform and application levels, and across levels. SCs are implemented in custom hardware, software, or even "in the cloud" depending on the amount of data and the time allotted for active learning. Actuators take on various forms, but in general they are "control knobs" that allow one or more operational parameters to be fine-tuned. For instance, at the circuit level, these are controls for supply voltage, clock frequency and body bias. At the platform level, actuators can take the form of decisions on what memory accesses to grant [2] [3], what communication policy to invoke [4], and what resources to utilize or avoid to ensure both reliability and availability. At the application level, actuators may naturally exist already such as the sensitivity of a robotic prosthetic controlled by a BCI. In other applications, actuators may not be inherently present, but can always take the form of real-time optimizers that improve application-level performance using various statistical-learning techniques.

There have been many publications focused on using hardware to speedup various learning approaches (neural networks, decision trees, etc.) [5] [6] [7] [8] [9] [10]. It cannot be overly emphasized however, that SLIC does not fall into that category. Our objective instead is to integrate a comprehensive learning capability that applies to all levels (circuit, platform and application) within an integrated system.

III. Learning Background

SLIC will have the flexibility to implement several specific learning algorithms but all will implement the forward model illustrated in Figure 2.

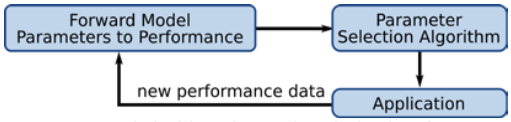

Figure 2: Statistical learning and optimization framework.

The conventional forward model accepts training data in the form of previous parameter settings and the resulting performances. It builds a model that predicts what performance will be achieved by a hypothetical, new parameter selection. A key element of these predictions is that they come with uncertainty estimates derived from the variance in the training data and the amount of relevant training data for each prediction. A trivial example of learning is a simple linear regression. In this case, learning is simply the process of fitting the regression parameters, and the "training data" is just the data used to perform the fit. "Testing data" refers to any data used later to query the fitted model for checking how well its predictions perform relative to the true values in the data. As described next, we propose more sophisticated methods that can fit non-linear and discrete models, but the basic concepts are the same.

The parameter-selection method uses the forward model to select parameters for the next learning cycle. Selection typically involves addressing an exploration/exploitation problem where a tradeoff has to be made by selecting parameters that are expected to perform well versus those that have uncertainty, and thus could provide useful data for improving the forward model. The selected parameters are then applied in the target system and new data is collected on the resulting performance that is used to update the forward model (i.e., included in the training data for the model). The algorithms developed for SLIC depend on features of the target application. Here we outline some scenarios. The highest-level distinction between scenarios is whether the

parameters represent a continuous-valued space of possible policies, or if they are a discrete set of alternative policies.

Discrete Policy Choices. Figure 3 shows the performance of a hypothetical forward model where there are four alternative policies the system can choose. This model does not attempt to generalize what it learns between policies. The mean and confidence intervals for the performance of each policy are simply tabulated based on the data collected when the corresponding policy is employed. In this form, the problem is a classic multi-armed bandit [11]. Each alternative is treated independently, and we want to learn the performance of each while concentrating most of the trials on the best performers. For instance, when dealing with memory accesses, policy choices include first-come first-serve (FCFS), round-robin, or row-buffer hit-first strategies to optimize the overall system performance. There are two common methods for parameter selection. The first is the UCB1 (Upper Confidence Bounds) algorithm [12]. It provides regret bounds for arbitrary reward distributions and can be computed quickly based on the means and numbers of samples from each policy. The second is Gittins indices [13]. These have the benefit of yielding optimal (in expectation) choices for the time-discounted case when the distributional assumptions on reward are correct. These two algorithms are ideal for SLIC since they both require little overhead for implementation.

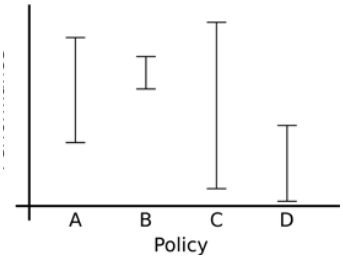

Figure 3: A forward model for four alternative policies.

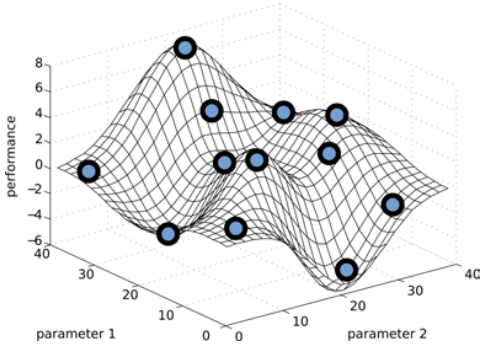

Figure 4: A hypothetical forward model with two continuous parameters.

Continuous Parameters. Figure 4 is a second hypothetical forward model where there are two parameters that control a policy. The circles represent data points collected by choosing specific parameter settings and observing the resulting performance. The surface is a current estimate of the function that maps parameter values to performance. While there are many possible function approximators that could produce this estimate, Gaussian processes are quite interesting [14] since they provide the ability to model nonlinear functions and also yield confidence intervals (not shown in the figure) on their estimates. For example, we could use these and similar approaches to tune the transistor back-bias voltage to balance performance and leakage of an arithmetic-logic unit (ALU) or SRAM (static random access memory). Or in a system that implements dynamic voltage frequency scaling (DVFS), we can use this approach to identify the optimal voltage and frequency combination for a time segment of the specific application workload [15]. For a BCI application, these techniques can be used to deal with neuroplasticity.

It has been shown that good empirical performance can be achieved by modifying bandit selection algorithms for the continuous case even when the theoretical assumptions are violated [16]. With that in mind, the UCB1 and Gittins methods can be extended from the discrete alternatives case for use with continuous parameters. One approach simply evaluates a point in the continuous space as if it were a discrete alternative by retrieving its mean, variance, and effective number of supporting data points from the function approximator. Then a set of candidate points are evaluated as discrete alternatives. Generalization happens through the function approximator, but is not explicitly considered during selection. In higher-dimensional parameter spaces, it may be difficult to explicitly build or search the forward model described here. This is especially true in non-stationarity systems where the performance data becomes stale long before sufficient data has been collected to build a reasonable model. For these scenarios, a "model free" approach can be implemented based on gradient ascent [17]. The only information stored by the forward model will be an estimate of the gradient derived from the most recent data points. The policy selection method will simply choose gradient steps. This method however is subject to being trapped in local minima, but is simpler and faster than the methods based on a full forward model.

Reinforcement Learning. The examples presented so far have assumed that the credit assignment between a newly chosen set of parameters and newly observed performance data is immediate. When time delays in the system mean that performance is the cumulative result of earlier choices, reinforcement learning will be used to deal with the credit-assignment problem. The forward modeling described above may be used to learn value functions and the parameter-selection algorithms may be adapted for Q-learning [18].

IV. SLIC PROJECTS

To demonstrate the viability of SLIC, several projects have been initiated [19-32]. A brief overview of some of these projects, which span from applications, architectures, and circuits, are presented here in this section.

Core Power. While historically the major goal of processor designers was to gain better performance by continuously shrinking device size or speeding up the clock speed, the power wall was eventually reached and energy has become the main design constraint. As a result, improving performance under the Thermal Design Power (TDP) constraint becomes one of the main directions in power/performance optimization. Many algorithms have been proposed to find near-optimal DVFS control solutions in polynomial time; however they suffer from unawareness of future machine state and excessive budget overshoot, and may only be efficient for small-scale multi-core systems, rather than systems with hundreds of cores. By exploiting both spatial and temporal hierarchies, we propose an On-line Distributed Reinforcement Learning (OD-RL) method [29] that is able to improve the performance with much less TDP overshoot, higher relative performance improvement, and smaller runtime overhead.

Figure 5 shows the hierarchical structure of OD-RL. At the finer grain, a per-core reinforcement learning method is used to learn the optimal control policy of the voltage/frequency levels in a system model-free manner. At the coarser grain, an efficient global power budget reallocation algorithm is used to maximize the overall performance. Experiments demonstrate that compared to the state-of-the-art algorithms: 1) OD-RL produces up to 98% less budget overshoot, 2) up to 44.3x better throughput per over-the-budget energy and up to 23% higher energy efficiency, and 3) two orders of magnitude speedup over state-of-the-art techniques for systems with hundreds of cores.

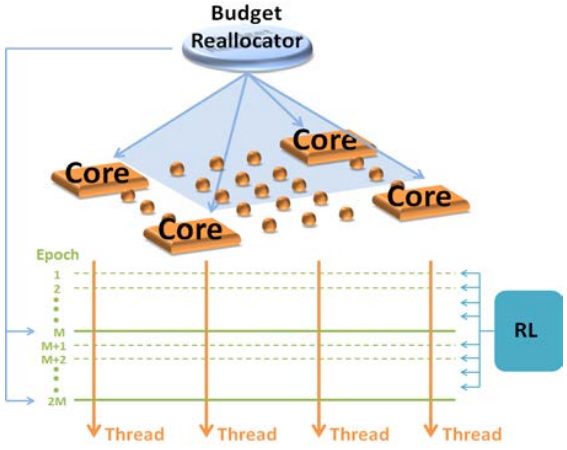

Figure 5: Hierarchical structure of OD-RL [29].

JTAG Protection. IEEE 1149.1, commonly known as JTAG (joint test action group), is the standard for implementing a serial test access port for ICs. JTAG is primarily utilized at the time of IC fabrication but is also employed in the field, giving access to internal sub-systems of the IC during operation, or for failure analysis and debugging. Because JTAG is left intact for post-fabrication use, it inevitably provides a "backdoor" that can be exploited to undermine the security of the chip. Potential attackers can therefore use JTAG to dump critical data or reverse engineer intellectual-property cores. Because an attacker uses JTAG differently from a legitimate user, it is possible to detect an unauthorized access using customized machine-learning algorithms. Specifically, a JTAG protection scheme, termed SLIC-J, is proposed to monitor JTAG activity, detect malicious accesses, and ultimately protect the JTAG from being misused [31]. SLIC-J characterizes user behavior with respect to a set of specially-defined features, and makes online prediction using a classifier implemented in hardware. Further, due to the variance that naturally occurs within both legitimate uses and attacks of the JTAG, we have developed a feature-revision mechanism, which delays the labeling of the JTAG operation until ample evidence is gathered.

SLIC-J is implemented within the JTAG of the OpenSPARC T2 which is a 64-bit 8-core microprocessor (Figure 6). To validate the effectiveness of SLIC-J, both legitimate uses and attacks of the JTAG, consisting of a 110 programs in total, are emulated. By using the feature-revision mechanism, the overall accuracy of detecting malicious accesses is 99.2%, while the overall escape rate (i.e., the percentage of attacks that escape detection) of 0.8%.

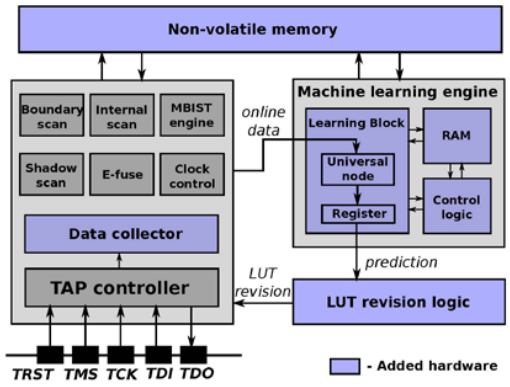

Figure 6: SLIC-J is integrated with the JTAG of the OpenSPARC T2. The JTAG, equipped with SLIC-J, adds only 2% to the original chip area [31].

On-Chip Classifier. In this project, we consider a case study of linear discriminant analysis (LDA) for binary classification [23]. We found that rounding error incurred from fixed-point arithmetic can significantly distort the classification output. We therefore propose a new LDA algorithm for fixed-point computation (LDA-FP). LDA-FP is formulated as a mixed integer programming problem with consideration of the non-idealities (i.e., rounding and overflow) posed by fixed-point arithmetic. Furthermore, a branch-and-bound method with several efficient heuristics is developed to find the globally optimal classification boundary of LDA-FP. With our re-designed training algorithm, LDA-FP can be efficiently implemented with extremely small word length for on-chip low-power operation. Experiment results show that LDA-FP is able to reduce the word length by up to 3x (i.e., equivalent to 9x power reduction) compared to the conventional LDA algorithm, without surrendering any classification accuracy.

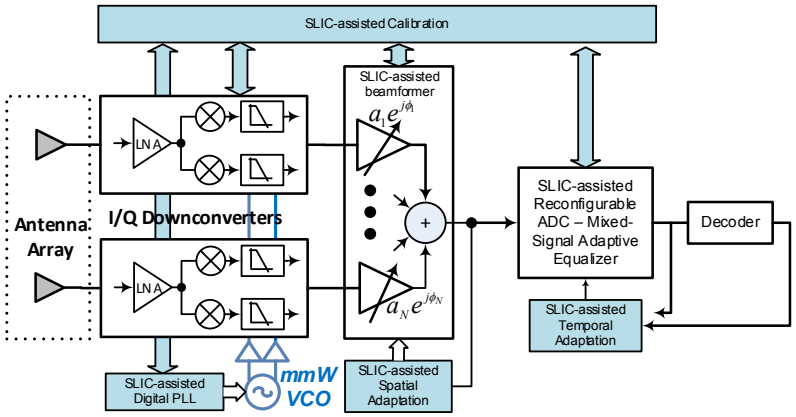

Figure 7: A future mm-wave multi-antenna receiver that utilizes SLIC-based concepts, namely: Signal processing circuitry and adaptation paths in the spatial and temporal loops; all-digital phase-locked loop frequency synthesizer with SLIC-enhancement for accurate time-to-digital conversion; SLIC-assisted calibration of mismatches in RF front-end and analog-digital interface circuits (for e.g., [30]); and built-in self-test of key sub-systems including the oscillator and mixed-domain beamformer.

Receiver Calibration. The use of mm-wave frequencies is emerging as a viable solution to the extreme paucity of bandwidth available for wireless communication at low gigahertz frequencies. Compared to low-GHz transceivers, millimeter-wave transceivers must incorporate vastly greater functionality and complexity spanning the system, signal processing and circuit levels. In order to meet this goal at low cost with low power consumption, SLIC-based approaches are being investigated. In particular:

- Spatial-domain signal processing becomes imperative due to the high directionality of millimeter-wave links, over and above temporal (or frequency-domain) signal processing. Such processing is accomplished through the use of multi-antenna transmitters and receivers. Current mm-wave transceivers perform algorithmically rudimentary spatial signal processing (i.e., non-adaptive phased arrays). More sophisticated spatial signal processing algorithms such as adaptive beam-steering are essential in mm-wave cellular links. These algorithms are based on learning and adaptation, but the power consumption of simple-minded digital-domain realizations are prohibitively expensive, thus mandating the development of architectures that partition the requisite signal processing optimally across the analog, digital and RF (mm-wave) domains.
- The aforementioned challenges are exacerbated by two other factors: (1) the wide bandwidth of mm-wave signals leads to high power consumption in the temporal signal processing circuits (such as the DSP, the ADC/DAC interfaces and frequency synthesis circuits), and (2) the quest for high performance in the underlying wide bandwidth and high carrier frequency degrades circuit robustness in the face of process, voltage, temperature and aging-related variations.

SLIC concepts and circuits are being "sprinkled" throughout the system to incorporate algorithmically sophisticated spatial-temporal signal processing while reducing energy consumption and cost, and increasing robustness. In Figure 6, several instances of SLIC-assisted and SLIC-enhanced functionalities are identified.

V. SUMMARY

Statistical learning in chip (SLIC) applied to the design and on-line operation of an integrated system has great potential to have significant impact on a number of areas:

- **Design** – It is challenging to design an integrated system so that all of its possible personalities can be seamlessly handled. With SLIC, the burden on the designer is eased since the ever-changing personality of the system can instead be learned and adapted to.
- **Yield** – Currently, an integrated system that does not meet specifications is discarded. With SLIC, it will be possible to increase yield since some flaws in the system personality will be compensated based on learning.
- **Test** –With SLIC, stress testing can be mitigated since changes in operation due to a subtle flaw can be detected and compensated for by learning a model of normal/expected operation.
- **Performance** – SLIC allows the performance optimization across the system stack, allowing the unique system personality to be exploited for maximum gain. This capability is not only critical for mobile integrated systems but will also be quite beneficial for server farms since power for such entities is also paramount.
- **Individualization** – Since SLIC learns the habits of the user, applications that were before learning-agnostic can now be fine-tuned to enhance the overall experience of every individual user. For learning-inherent applications,

especially from the medical field, SLIC promises to usher in a new field of personalized medical instrumentation.

REFERENCES

[1] D. Ricketts et al., "Enhancing CMOS using Nanoelectronic Devices, a Perspective on Hybrid Integrated Systems," *Proceedings of the IEEE*, pp. 2061-2075, Dec. 2010.

[2] E. Ipek et al., "Self-Optimizing Memory Controllers: A Reinforcement Learning Approach," *International Symposium on Computer Architecture*, pp. 39-50, 2008.

[3] J. Martinez and E. Ipek, "Dynamic Multicore Resource Management: A Machine Learning Approach," *IEEE Micro*, Sept./Oct. 2009.

[4] C. -L. Chou and R. Marculescu, "User-Aware Dynamic Task Allocation in Networks-on-Chip," *ACM/IEEE Design, Automation and Test in Europe*, pp. 1232-1237, 2009.

[5] V. Prabha and E. Moine, "Hardware Architecture of Reinforcement Learning Scheme for Dynamic Power Management in Embedded Systems," *EURASIP Journal on Embedded Systems*, 2007.

[6] Z. Baker and V. Prasanna, "N Architecture for Efficient Hardware Data Mining using Reconfigurable Computing Systems," *IEEE Symposium on Field-Programmable Custom Computing Machines*, pp. 67-75, 2006.

[7] Z. Baker and V. Prasanna, "Efficient Hardware Data Mining with the Apriori Algorithm on FPGAs," *IEEE Symposium on Field-Programmable Custom Computing Machines*, pp. 18-20, 2005.

[8] R. Narayanan et al., "An FPGA Implementation of Decision Tree Classification," in *ACM/IEEE Design, Automation and Test in Europe*, pp. 16-20, 2007.

[9] Y. Luo, K. Xiang, and S. Li, "Acceleration of Decision Tree Searching for IP Traffic Classification," *ACM/IEEE Symposium on Architectures for Networking and Communications Systems*, pp. 40-49, 2008.

[10] S. Chkrabarty and G. Cauwenberghs, "Sub-Microwatt Analog VLSI Trainable Pattern Classifier," *IEEE Journal of Solid-State Circuits*, May 2007.

[11] D. Berry and B. Fristedt, *Bandit Problems: Sequential Allocation of Experiments*.: Chapman and Hall, 1985.

[12] P. Auer, N. Cesa-Bianchi and P. Fischer, "Finite-time Analysis of the Multiarmed Bandit Problem," *Machine Learning*, 2002.

[13] J. Gittins, *Multi-Armed Bandit Allocation Indices*.: Wiley, 1989.

[14] C. Rasmussen and C. Williams, *Gaussian Processes for Machine Learning*.: MIT Press, 2006.

[15] D. C. Juan et al., "Learning the Optimal Operating Point for Many-Core Systems with Extended Range Voltage/Frequency Scaling," CODES+ISSS, Oct. 2013.

[16] J. Schneider and A. Moore, "Active Learning in Discrete Input Spaces," *Interface Symposium*, 2002.

[17] P. C. Pendharkar, "A Comparison of Gradient Ascent, Gradient Descent and Genetic-Algorithm-Based Artificial Neural Networks for the Binary Classification Problem," *Expert Systems*, pp. 65-86, May 2007.

[18] R. Sutton and A. Barto, *Reinforcement Learning: An Introduction*.: MIT Press, 1998.

[19] H. Peng and R. Marculescu, "Identifying Dynamics and Collective Behaviors in Microblogging Traces," *IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, pp. 846-853, 2013.

[20] G. Liu, J. Park, and D. Marculescu, "Dynamic Thread Mapping for High-performance, Power-efficient Heterogeneous Many-core Systems," *IEEE International Conference on Computer Design*, pp. 54-61, 2013.

[21] H. Peng and R. Marculescu, "ASH: Scalable Mining of Collective Behaviors in Social Media using Riemannian Geometry," *ASE BIGDATA/SOCIALCOM/CYBERSECURITY Conference*, 2014.

[22] M. Li and X. Li, "Verification based ECG Biometrics with Cardiac Irregular Conditions using Heartbeat Level and Segment Level Information Fusion," *IEEE International Conference on Acoustics, Speech and Signal Processing,* pp. 3769-3773, 2014.

[23] H. Albalawi, Y. Li, and X. Li, "Computer-aided Design of Machine Learning Algorithm: Training Fixed-point Classifier for On-chip Low-power Implementation," *ACM/IEEE Design Automation Conference*, pp. 1-6, 2014

[24] M. Won, H. Albalawi, X. Li, and D. E. Thomas, "Low-power Hardware Implementation of Movement Decoding for Brain Computer Interface with Reduced-resolution Discrete Cosine Transform," *International Conference of the IEEE Engineering in Medicine and Biology Society*, pp. 1626-1629, 2014.

[25] S. Kundu et al., "A 1.2 V 2.64 GS/s 8bit 39 mW Skew-tolerant Time-interleaved SAR ADC in 40 nm Digital LP CMOS for 60 GHz WLAN," *IEEE Proceedings of the Custom Integrated Circuits Conference,* pp. 1-4, 2014.

[26] H. Goncalves et al., "DALM-SVD: Accelerated Sparse Coding through Singular Value Decomposition of the Dictionary," *IEEE International Conference on Image Processing,* pp. 4907-4911, 2014.

[27] X. Li et al., "Ultra-low-power Biomedical Circuit Design and Optimization: Catching the Don't Cares," *International Symposium on Integrated Circuits,* pp. 115-118, 2014.

[28] H. Peng and R. Marculescu, "Multi-Scale Compositionality: Identifying the Compositional Structures of Social Dynamics Using Deep Learning," *PloS One*, vol. 10, no. 4, p. e0118309, 2015.

[29] Z. Chen and D. Marculescu, "Distributed Reinforcement Learning for Power Limited Many-core System Performance Optimization," *Design, Automation and Test in Europe*, pp. 1521-1526, 2015.

[30] S. Kundu and J. Paramesh, "DAC Mismatch Shaping for Quadrature Sigma-Delta Data Converters," *IEEE Midwest Symposium on Circuits and Systems*, 2015.

[31] X. Ren, V. G. Tavares, and R. D. Blanton, "Detection of Illegitimate Access to JTAG via Statistical Learning in Chip," *ACM/IEEE Design, Automation Test and Europe,* pp. 109-114, 2015.

[32] X. Ren, M. Martin, and R. D. Blanton, "Improving Accuracy of On-chip Diagnosis via Incremental Learning," *IEEE VLSI Test Symposium,* pp. 1-6, 2015.