

# Faculty Portfolio: Promotion and Tenure

---

Date: 11/7/2013

Name: Phillip H. Jones III

Department: Electrical and Computer Engineering

Current Rank: Assistant Professor

## I. SELF ASSESSMENT OF SCHOLARSHIP

### A. Summary Statement of Accomplishments in and Impact of Scholarship

The evolution of computing technology is a major driving force in society. Even only looking back 20 years, rapid advancements in computing has had transformational impact on areas of health, energy, and the environment. For example, the idea of doctors in U.S. cities operating on patients in remote rural areas has moved from the realm of science fiction to reality. Cyber physical system is a computing field that looks to enable and further improve the safety, reliability, and performance of such society impacting technology. As an Assistant Professor within Iowa State University's Electrical and Computer Engineering department, my research has focused on the following computing areas:

**Hardware support for embedded Cyber Physical Systems (CPS):** CPS can be viewed as systems that have a tight coupling between their cyber (i.e. computing) and physical components. Often these systems have real-time and controls aspects. My research in this area has targeted the investigation, development and evaluation of computing hardware mechanisms to support these real-time and control aspects [2, 3, 8, 12, 18, 31].

**Reconfigurable computing:** In this area, I have used reconfigurable computing technology as a research medium to explore the idea of using temperature feedback to improve computing performance [1, 13, 17, 19, 20, 21, 22, 23, 30], and leveraging the reconfigurable nature of this medium to improve a design's resilience to faults within its computing fabric [16].

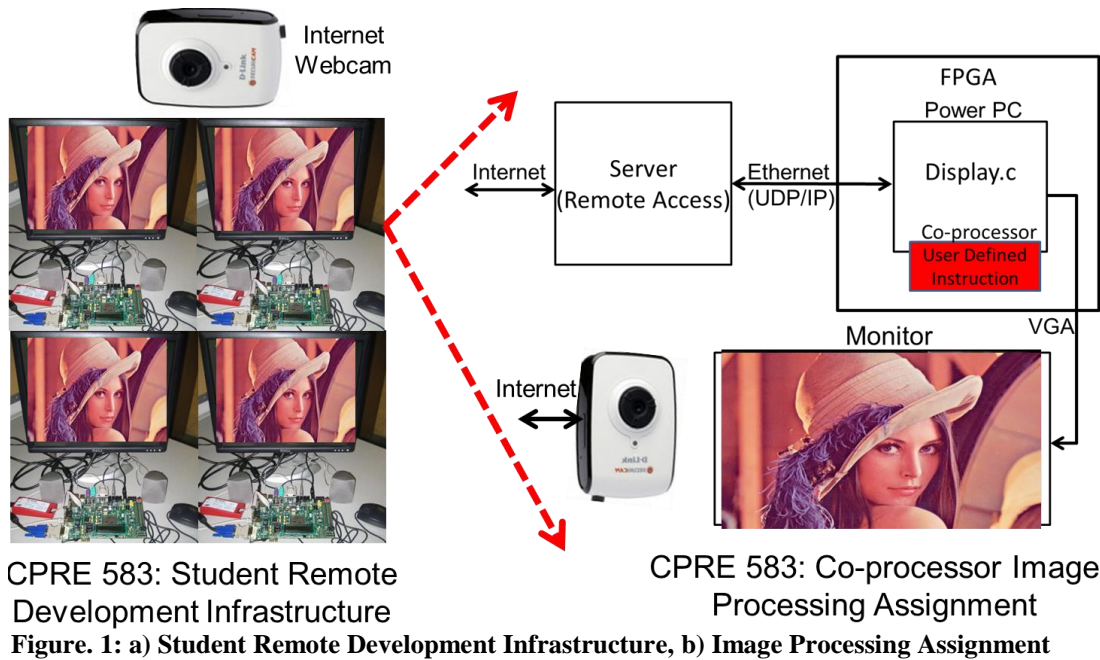
**Application specific hardware architectures:** I have an interest in investigating hardware architectures for accelerating applications. In addition to the success I have found with accelerating sensor processing, controls, and real-time scheduling computations [3, 12, 18], I have found some initial success in the area of bioinformatics [10], and the area of sparse-matrix acceleration [4].

### B. Specific Statements of Accomplishments and Impact

#### 1. Self Assessment of Accomplishments in and Impact of Teaching and Learning

I have enjoyed the opportunity to innovate within the curriculum during my time at Iowa State University. (ISU) Thus far, I have focused on four courses: 1) Reconfigurable Computing (CPRE 583), 2) Models and Techniques in Embedded Systems (CPRE 584), 3) Senior Design (CPRE 491/492), and Introduction to Embedded System (CPRE 288).

**Reconfigurable Computing (CPRE 583):** This course helps train graduate students in conducting research in the area of computer architecture and reconfigurable computing. It is also offered as a distance education course to help service the needs of regional companies (e.g. Rockwell Collins and John Deere) in advancing their employees' skill set. Before my arrival at ISU, this was already a successful course, however it was primarily lecture based. I developed lab infrastructure that allowed both on-campus and distance students to gain equivalent hands on experiences.



Remotely accessible servers (as shown in Figure 1a) allowed distance students access to same industrial strength tools as their on campus counter parts, local networking allowed the distance students to interact with Reconfigurable Computing platforms, and Web cameras allowed them to remotely observe simple debugging mechanisms such as LEDs and small LCD screens.

In addition, hands-on assignments were carefully designed to give distance students an equivalent experience as on-campus students. For example, Web cameras directed at computer monitors allowed students to observe the behavior of their image processing hardware, and network processing assignments allowed students to analysis and debug their hardware design by monitoring the local network traffic using tools such as TCPdump (Figure 1b).

*Two unexpected outcomes* of my efforts in developing hands-on components for this course were 1) the on-campus students who had physical access to the lab setup began to abandon the physical lab in favor of using the distance student infrastructure. This could be considered the ultimate evidence that the distance infrastructure was fulfilling its intent to give on-campus and distance students an equivalent experience, 2) during my first semester at Iowa State (Fall of 2008), a major challenge was finding a way to allow distance students to remotely use the design tools, which are intensely graphical, over their low bandwidth internet connections. The solution that I came across was a remote access tool called NX. NX was so successful with my small class of 20 students that the department's computer support team has adopted it as the primary tool for having student's remotely access computing resources, impacting on the order of 1,000 students.

I have an *ASEE* [11], and *MSE* [15] publication based off of an undergraduate course (primarily driven by Joseph Zambreno) that leveraged the infrastructure I developed for CPRE 583.

**Models and Techniques in Embedded Systems (CPRE 584):** I co-developed CPRE 584 with Dr. Joseph Zambreno with the goal of training new graduate students for conducting research in the area of embedded computing systems and software/hardware codesign.

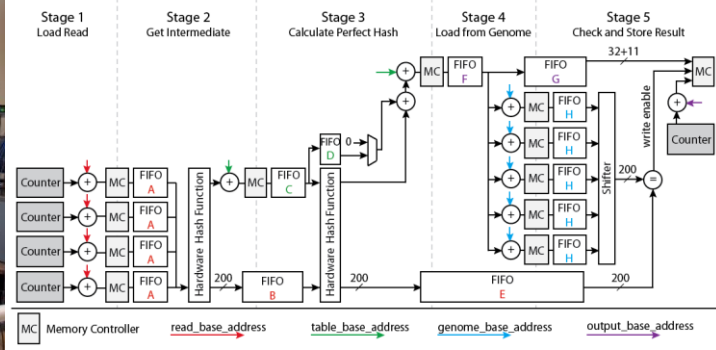



Figure 2: a) CPRE 584:MemoCODE team members (I am at the far right), b) winning hardware accelerated DNA mapping design [10]

**Diving into research.** A key idea behind this class was recreating the experience a new graduate student encounters in the field of embedded systems research. As part of the first third of the semester, students are exposed to a number of fairly complex design tools and are expected to ramp themselves up quickly. Since most of these tools have hundreds of pages of documentation, students are forced to learn to extract pertinent information efficiently. When students hit walls, I step in to help them navigate roadblocks. In order to show tool proficiency, they must demonstrate the implementation of a baseline experimentation infrastructure that can be used to evaluate research ideas. A common objective has been implementing an infrastructure that can be used to design, implement, and evaluate image processing algorithms (e.g. Sobel edge detection).

**Teamwork and generating excitement.** Synchronizing this class with an international design competition has been a successful innovation. The ACM-IEEE International Conference on Formal Methods and Models for Codesign (MemoCODE) design competition has been an excellent vehicle for testing the skills students develop during the first third of the course. On March 1<sup>st</sup>, MemoCODE releases a hardware/software codesign challenge, and teams have until March 31<sup>st</sup> to complete a fully functional solution. Winners are chosen in two categories: 1) Fastest, and 2) Most efficient use of computing resources. Since 2009, my class has had a few 2<sup>nd</sup> place finishes and in 2012 won 1<sup>st</sup> in the category of Fastest (20X faster than 2<sup>nd</sup> place), and took second in computing efficiency. Two team members are shown in Figure 2a, and the winning design’s high-level architecture is shown in Figure 2b. Further detail can be found in the MemoCode conference proceedings [10].

**Introduction to Embedded Systems (CPRE 288):** This is the first undergraduate-level course that I recently began teaching. I find I greatly enjoy teaching undergrads, and have already made significant enhancements to lecture slides. Currently, I am helping upgrade the lab’s infrastructure. This course has also motivated me to help with the Junior-level version of the course (CPRE 488).

## 2. Self Assessment of Accomplishments in and Impact of Research and Creative Activities

By their nature, cyber-physical systems (CPS) require examining the interaction of computing and physical components of a system in a holistic and cross disciplinary manner. Often these systems have real-time and controls aspects. The grand vision of my research is to drive the state-of-the-art in systems architecture and analysis techniques for engineer applications using such a cross disciplinary and holistic systems view. In particular, I have conducted investigations into cyber-physical interactions that occur at the intersection of hardware architecture, real-time systems, and controls design. Figure 3 captures my vision of the type of resource-constrained CPS I would like to help enable and an example of high-level architecture modifications that move toward this goal.

## Resource constrained CPS



- Relatively easy to program and analysis
- Inefficient sequential execution
- Targets supporting CPS parallelism
- Challenges: Controls & Real-time Analysis

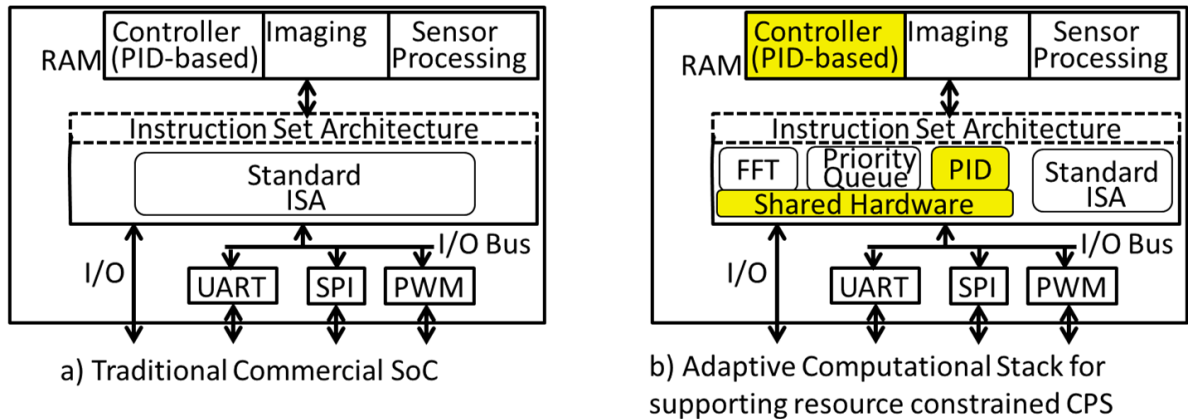
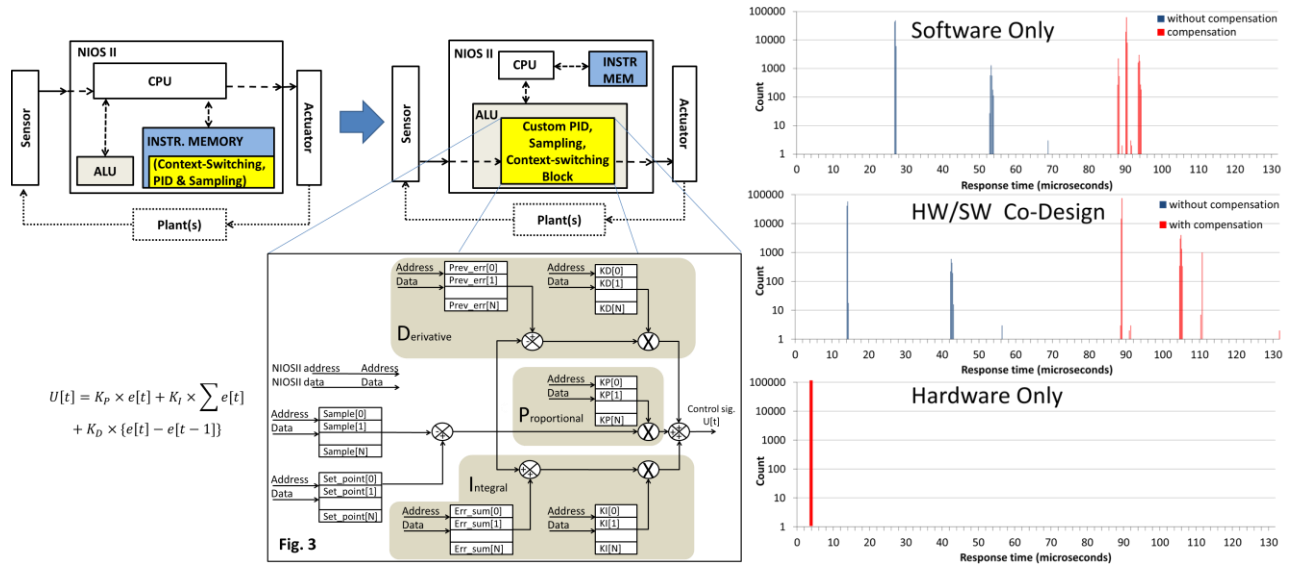


Figure 3: Vision for hardware architectural support of resource constrained CPS (Cyber Physical Systems)

In summary, the following presents three projects that reflect my efforts thus far in this endeavor. **Project 1** examines the use of hardware to assist in the implementation and design of **controls** aspects of embedded systems. **Project 2** explores the use of hardware to assist in reducing **real-time** scheduling overheads and improving execution time determinism. **Project 3** focuses on the specific **physical property of temperature** with respect to reliable measurement in reconfigurable fabrics to support thermal feedback for application performance improvement. These projects integrate into the computational stack hardware features that support controls and real-time aspects of CPS to better serve the dynamics of CPS applications.

### Project 1: Controls Systems Hardware Support

Embedded systems and digital control theory have independently developed into mature fields, despite the clear connection between controllers and the embedded platforms that host them. The first digital controllers were implemented on dedicated microprocessor, thus helping maintain certain traditional simplifying assumptions for controls design (e.g. constant sensor sample rates, and constant computing time). However, increasing demands for tighter system integration, smaller form-factors, and lower-power design has played a large role in making these simplifying assumptions less valid. In modern embedded systems, it is common for control tasks and non-control tasks to share a single processor via a scheduling algorithm or Operating System (OS), which in turn contributes to variations in task execution time. The **2 primary goals** of this project are: 1) characterizing the degree to which traditional controller design simplifying assumptions hold true and evaluating hardware mechanisms to help maintain these simplifying controller design assumptions [3,18], and 2) developing a design approach using a Plant-on-Chip (PoC) hardware architecture to evaluate controllers on a given embedded platform earlier in the design process [2].



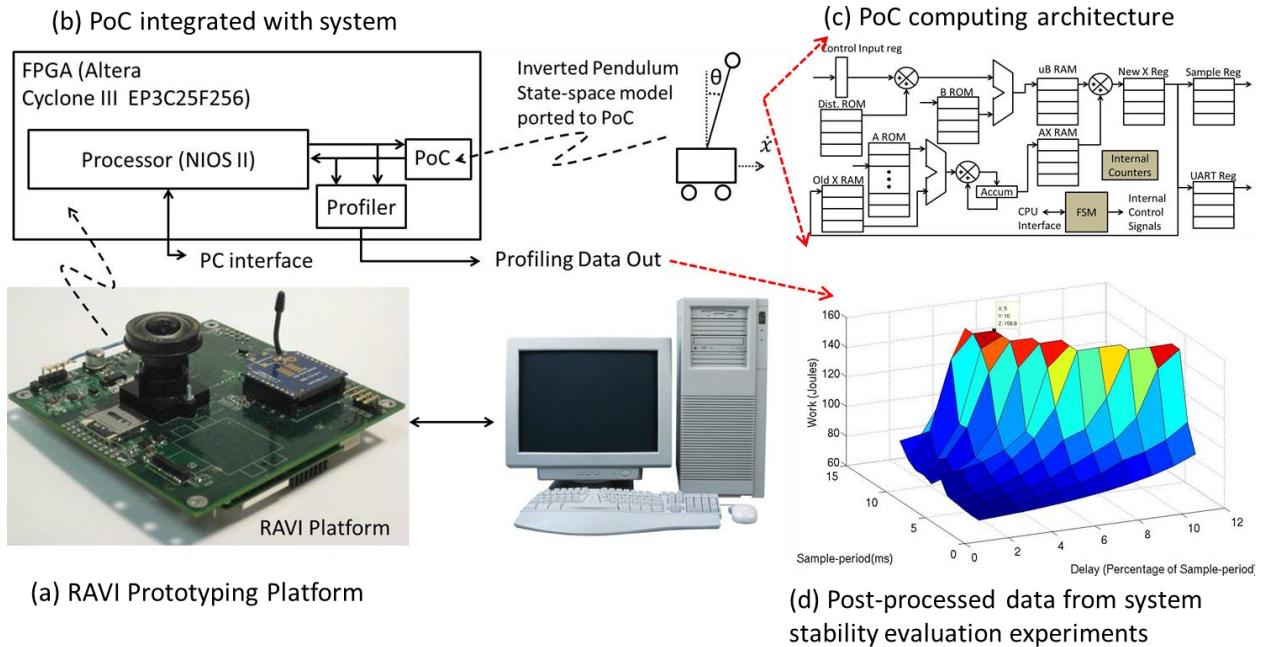
**Figure 4: a) FPGA-based evaluation setup for moving functionality between software and hardware, and b) histograms representing the variation in response-time for executing a PID controller for several SW/HW distributions of functionality.**

**Characterizing traditional simplifying assumptions for controller design.** In [3] and [18], hardware mechanisms were designed and evaluated on an FPGA-based system-on-chip platform (called RAVI, see Figure 5a) to support sensor processing and proportional-integration-derivative (PID) controller computations. Figure 4a illustrates the extreme configures for evaluating the impact of migrating controller functionality between a complete software implementation (far left of Figure 4a) and a complete hardware implementation (far right of Figure 4a). Figure 4b provides a flavor of the type of data collect from performing controller response-time performance evaluations.

Maintaining consistency is a key factor when comparing software/hardware tradeoffs for controller implementation. The RAVI development board allowed the use of a single platform for developing and evaluating each of the four software/hardware controller variations evaluated. The FPGA was used to implement the NIOS-II (Altera’s soft-processor) based system-on-chip setup shown in the upper part of Figure 4a.

The metrics of interest for our evaluation were 1) response time (defined to be the time to service all plants once), and 2) response-time jitter. The system variables varied to evaluate these metrics were: (1) the architecture (I:Full SW, II: PID core in HW, III: PID core + sensor sampling in HW, IV: PID core + sensor sampling + plant context switching in HW), (2) the number of plants controlled (10,100, 1000), (3) the processor interrupt timer (1 ms, 100 ms), (4) the sensor sampling rate, in samples per second (SPS) (No Delay, 200KSPS, 819SPS), (5) software-implemented jitter compensation (used or not used).

**Result Summary.** First, while keeping a fixed sampling rate and a fixed number of plants, and moving from Case I (completely software) to Case IV (completely hardware) we observed improved response times. Most interesting, was the large improvement in response-time performance when moving plant context switching functionality from software to hardware. It makes sense that hardware performs context switching more efficiently since it stores all context information locally



**Figure 5: (a) RAVI FPGA-based prototyping platform, (b) PoC substituted for an inverted pendulum plant for evaluating system stability for variation in controller response time, (c) hardware computing architecture of the PoC, (d) example post-processed data showing actuator effort vs. sensors sample period and controller computation delay.**

in BlockRAM (i.e. on-chip memory), which can be accessed in a single clock cycle. In comparison, the main processor potentially has to fetch context information from main memory (taking many and a non-deterministic number clock cycles).

*Contribution Summary.* The primary contributions of this part of the project ([3] and [18]) were 1) the tight integration of a time multiplexed hardware PID controller within an embedded processor, 2) the characterization of the PID controller architecture for several alternative hardware/software hybrid-designs with respect to response time and jitter, and 3) the tight integration of a hardware-based sensor processing unit (SPU) within an embedded processor and its evaluation with respect to the software implementation of common sensor processing tasks in terms of response time.

***Plant-on-Chip(PoC) approach for controller system- integration.*** As the assumption of deterministic controller execution on shared embedded platforms has continued to become less valid, interest has sparked in both the field of controls theory and the field of real-time scheduling to address this concern. Experts in real-time scheduling have explored new scheduling techniques to help maintain the assumption of periodic sample rates (i.e. reducing response-time jitter), and the controls community has developed advanced control algorithms and analysis techniques to better tolerate response time-jitter. Ideally the goal is to maximize CPU utilization, while maintaining a physical system's stability needs.

Additionally, a number of software tools (e.g. Jitter-Bug, TrueTime) have been developed to help analysis the impact of non-deterministic execution time on system stability. Though these analytical and simulation-based design tools provide a good approximation of a system's robustness to sample-period and delay, they work in environments and under assumptions that delay can be modeled as a probability distribution function. Research shows this to be not realistic and that computer elements (cache misses, bus arbitration) cause non-deterministic time variation in delay and sample-period.

*Overview.* In this work [2] an approach for controller system-integration is presented that directly interfaces a Plant-on-Chip (PoC) emulator with an actual processor under test, the setup inherently incorporates the target system's sources of non-determinism, thus giving a more accurate result when characterizing the system's robustness against sample period and delay variation. Figure 5b illustrates the high-level system organization of the PoC directly interfaced to the embedded system on a single chip. Figure 5c provides an architectural view of the PoC, which implements a generic four state-space calculator, and Figure 5d provides an example of post-processed data used to relate system stability to various average sample times, and controller response times.

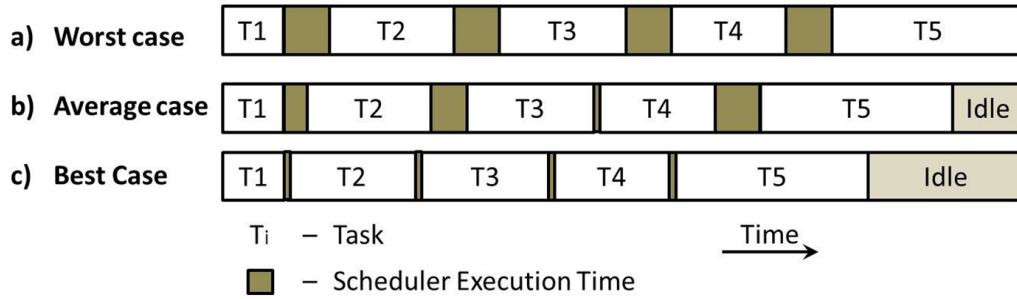
*Contribution Summary.* The experimental results indicate that this proposed framework both safely and accurately captures the non-deterministic effects of modern processor architecture on a physical plant. Comparing our results with those obtained by Jitter Bug (a popular tool for evaluating system stability); we demonstrate a more accurate representation of a real system.

## **Project 2: Real-time Systems Hardware Support**

Deploying increasing amounts of computation into smaller form factor devices is required to keep pace with the ever increasing needs of real-time and embedded system applications. The area of micro Unmanned Ariel Vehicles (UAVs) is an example of where such need exists. The size of these vehicles has rapidly decreased, while the capabilities users wish to deploy continue to explode. In June of 2011, the New York Times published several articles on the amazing work being pursued by Wright Patterson Air Force Base to develop micro-drones to aid soldiers on the battlefield. In February of 2011, the DARPA funded Nano Air Vehicle (NAV) program demonstrated a humming bird form-factor UAV weighing less than 20 grams (e.g. less than an AA battery) with video streaming capabilities. These real-time embedded applications can no longer rely on manufacturing advances to provide computing performance at Moore's law rates, due transistors approaching atomic scales and thermal constraints. Thus, more efficient use of the transistors available is needed.

In this project [12, 31] we assert the boundaries of software and hardware must be reexamined and believe a fruitful realm for research is the hardware-software co-design of real-time system functionality that has been traditionally implemented in software. Such a co-design is needed to balance the cost of dedicating limited silicon resources for high-performance fixed hardware functionally, with the flexibility and scalability offered by software. Additionally, we claim seamless migration between software and hardware implemented functionality is required to allow systems to adapt to the dynamic needs of applications. Specifically, this project examines hybrid architecture for priority queue management and evaluates this architecture within a real-time scheduling context.

The following motivates the importance of low processing overhead and timing predictably to a real-time scheduler's performance. Real-time operating systems (RTOSs) are designed to execute tasks within given timing constraints. An important characteristic of an RTOS is predictable response under all conditions. The core of the RTOS is the scheduler, which ensures tasks are completed by their deadline. The choice of a scheduling algorithm is crucial for a real-time application. Online scheduling algorithms incur overhead, as the task queues must be updated regularly. This action is typically paced using a timer that generates periodic interrupts. The scheduler overhead generally increases with the number of tasks. A high resolution timer is required to distribute CPU load accurately based on a scheduling discipline in real-time systems, but such fine-grain time management increases the operating system overhead.



**Figure 6: In order to allow analytical analysis of schedule feasibility, worst-case execution time (WCET) typically needs to be assumed. Thus, scheduler execution time variations that cause large differences between WCET and typical case execution time reduce utilization of system computing resources.**

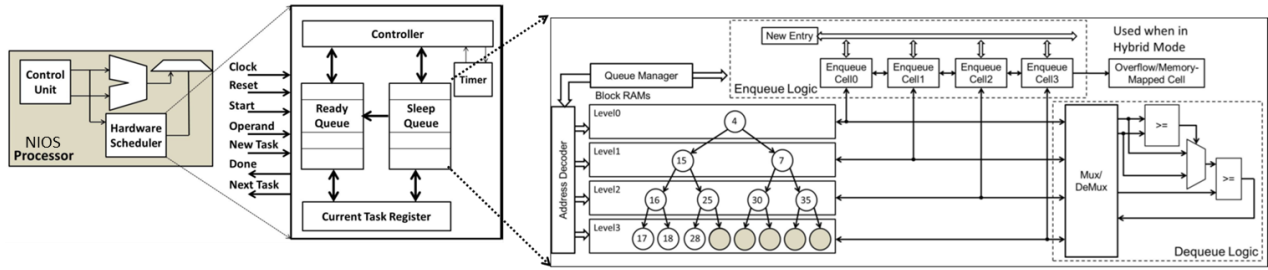
The extent to which a scheduler can ideally implement a given scheduling paradigm (e.g. Earliest Deadline First (EDF), Rate Monotonic (RM)), and thus provide the guarantees associated with that paradigm, is in part dependent on its timing determinism. A metric for helping quantify the amount of non-determinism that is introduced to the system by the scheduler is the variation in execution time among individual scheduler invocations. This can be roughly summarized by noting its best-case and worst-case execution times. Variations in scheduler execution time can be caused by system factors such as changes in task set composition, cache misses, etc. Reducing the scheduler's timing sensitivity to such factors can help increase deterministic behavior, which in turn allows the scheduler to better model a given scheduling paradigm.

Figure 6 illustrates how the variation in scheduler overhead affects processor utilization. To ensure that tasks meet their deadlines, the scheduler's worst-case execution times are often overestimated. This can cause a system to be underutilized and wastes CPU resources. In this project, we examine how the scheduler overhead and its variation can be reduced by migrating scheduling functionality (along with time-tick interrupt processing) to hardware logic. The expected results of our efforts are increased CPU utilization, better system predictability, and finer schedule and timing resolution.

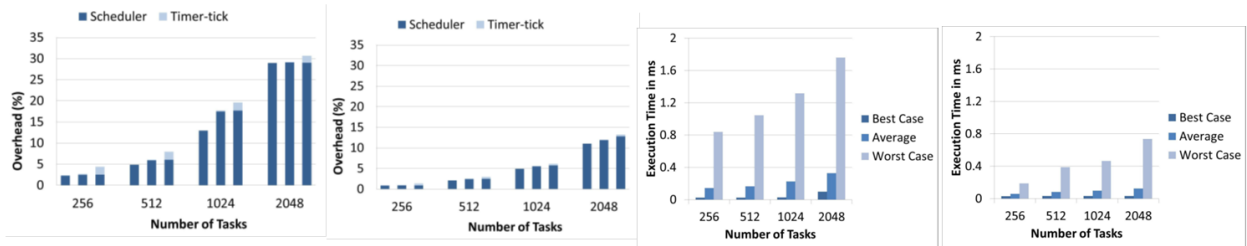
*Overview.* Figure 7a provides an overview of our software/hardware priority queue co-design implementation that supports up to 255 elements while under hardware management, and transitions to software-hybrid mode when this limit is surpassed, placing additional elements into main memory. A memory mapped interface provides software with access to priority-queue-structured on-chip memory, which enables quick and low overhead transitions between hardware and software management. As an application of this hybrid architecture, we implemented and evaluated a scalable task scheduler for real-time systems that reduces scheduler processing overhead and improves timing determinism of the scheduler.

*Contribution Summary.* The primary contributions made in this area have been 1) a hardware accelerated binary min heap that supports enqueue and peek operations in  $O(1)$  time, returns the top-priority element in  $O(1)$  time, and completes a dequeue operation in  $O(\log(n))$  time, 2) a scalable hardware-software priority queue architecture that enables fast and low-overhead transitions of queue management between hardware and hybrid modes of operation, and 3) A hybrid scheduler architecture that reduces scheduling overhead and improves predictability [12, 31].





(a) SW/Hardware co-design of a Priority Queue implemented as a custom instruction



(b) Performance of a software implemented EDF scheduler compared to one based on SW/HW co-designed priority queue.

Figure 7

*Results Summary.* A new hybrid priority queue architecture has been implemented, which can be managed in hardware and/or software. The priority queue when managed in hardware supports enqueue and peek operations in  $O(1)$  time, returns the top-priority element in  $O(1)$  time, and completes a dequeue operation in  $O(\log(n))$  time. The design enables quick and low overhead transition between hardware and software management. It utilizes hardware logic to enhance the performance of queue operations even when managing the priority queue in software. As an application of the proposed priority queue architecture, a scalable hybrid scheduler is implemented that supports 255 tasks in hardware mode and up to an arbitrarily large number of tasks in hybrid mode. Figure 7b provides a brief summary of the results of the performance evaluation of our Field Programmable Gate Array (FPGA)-based system-on-chip prototype. It shows a 90% reduction in scheduling overhead and a 98% decrease in scheduler execution time variation, when the scheduler is managed by hardware as compared to software, thus giving more predictable execution times, which is necessary in high-performance real time systems.

*Future Research.* Avenues of future work include, 1) reducing the rate of performance degradation as the queue overflows into software, 2) evaluating the use of the hybrid priority queue in discrete event simulation and network optimization algorithms and 3) integrating the hybrid scheduler with Real-time Linux and characterizing the scheduler performance.

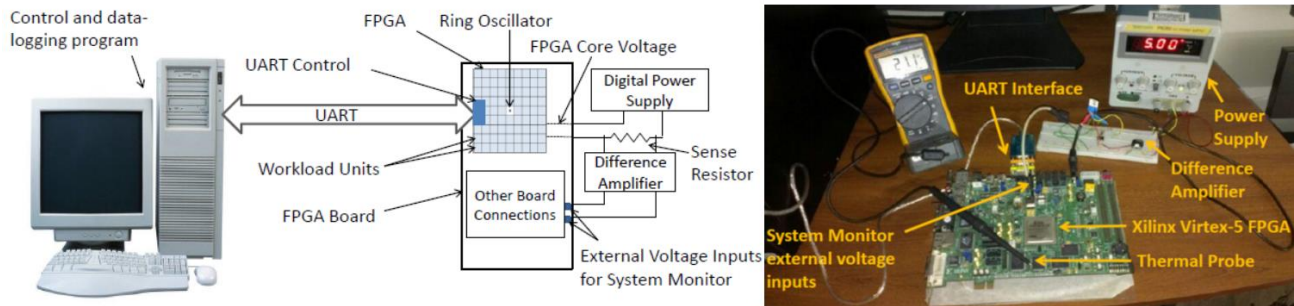


Figure 8: Experimentation setup for monitoring FPGA current draw to track workload intensity

### Project 3: Temperature Feedback in Reconfigurable Computing

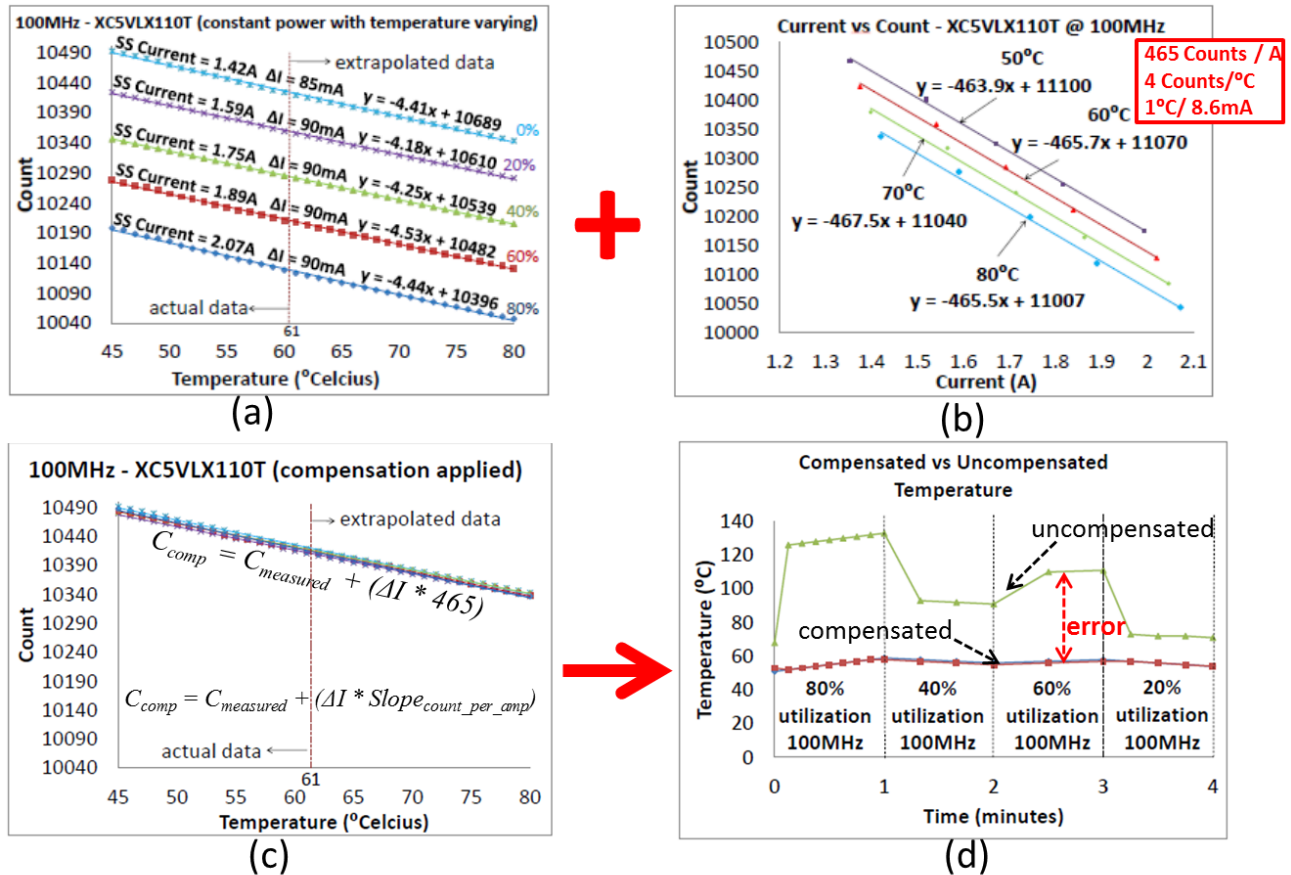
Embedded systems can be found in harsh and dynamic thermal environments. Environments such as deserts and outer space inherently exhibit large temperature changes over time. In general, embedded computers may need to operate in diverse thermal environments due the mobility of their host, or different instances of a system being deployed in significantly different thermal environments. Even large high-performance systems located in environmentally controlled rooms can be exposed to dynamic thermal conditions. For instance, a blade based computing system is one example, where densely packed computing nodes can cause constrained airflow and the thermal output of a node can impact its neighbors.

Systems with highly variable workload intensities can be thermally constrained by worst-case workload conditions. This is especially true of application specific integrated circuits (ASICs) that implement highly parallel and computationally intense applications. These applications can have many parallel modules active simultaneously under worst case workloads, while under light workloads most modules may be idle. Workload variation is an even larger concern for platforms that use reconfigurable devices. The diverse range of functionality that can be mapped onto these platforms, overtime, creates a wide variation in the amount heat that can be generated by the system.

*Project Goals.* This project investigates the use of temperature feedback for improving FPGA-based system performance. Specifically **two** aspects that have been explored are 1) obtaining reliable on-chip temperature measurements [13, 30], and 2) quantifying the degree to which thermal feedback can be used to improved system performance.

#### **Reliable Temperature Measurement.**

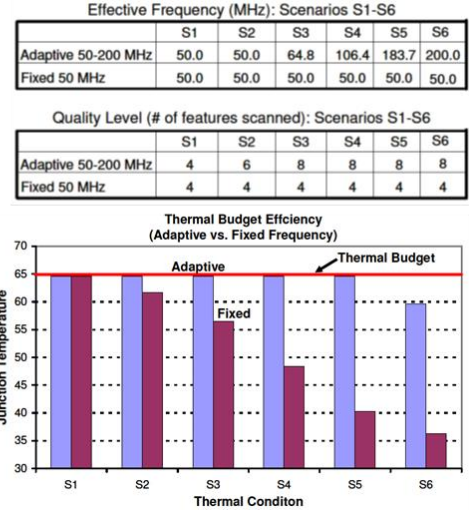
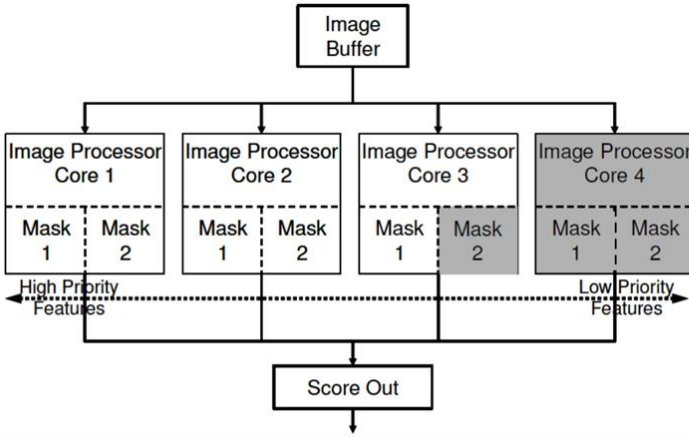
*Problem Investigated.* During my previous work with using temperature feedback for improving the performance of FPGA-based systems [1, 19, 20, 21, 22, 23], it was observed that for variations in workload intensity that the output vs. temperature relationship of an on-chip ring oscillator based thermometer varies (Figure 9a). For applications that have extremely dynamic workload this can make using ring-oscillators based thermometers a challenge. Referring back to Figure 9a, varying FPGA utilization from 0% to 40% induces an error in temperate measurement of ~35 degrees. Ideally one wants the workload intensity vs. ring-oscillator frequency to be constant (Figure 9c).



**Figure 9:** a) Observation that ring oscillator frequency vs. temperature relation varies with the workload intensity of the FPGA (0% - 80% utilization), b) results of experiments to characterize FPGA current draw vs. ring oscillator frequency for lines of constant temperature, c) the application of the results of characterization experiments to (a) to compensate for non-thermal effects of workload intensity on ring oscillator frequency, d) experimental evaluation of compensation technique, as can be seen the error between the compensated ring-oscillator thermometer and actual temperature is not noticeable compared to the uncompensated ring-oscillator thermometer as workload intensity is varied. .

*Approach.* The approach taken to mitigate this problem was to first characterize the non-thermal impact that workload intensity variation had on ring-oscillator frequency (typically caused by workload variation inducing microvolt level variation in the FPGA power source). Figure 8, shows the experimentation setup to perform this characterization, and Figure 9b provides the results of this characterization, a linear relation can be observed. This relation was then applied to the measured ring oscillator frequency to compensate for the non-thermal effects of workload intensity variation (Figure 9c). Figure 9d shows the resulting increase in measurement accuracy under workload variation.

*Results Summary.* For the specific FPGA characterized in Figure 9b, it was found that every 8.6mA change in workload current draw resulted in ~1 C error in temperature measurement (using a base workload as a reference). It should be noted that this linear relationship held when applied to a much larger FPGA hosted on a completely board, just with a different constant for slope. As can be seen in Figure 9d, before applying compensation, workload variation caused up to a 74 degree error in temperature measurement. The proposed compensation reduced the maximum error to 2 degrees.



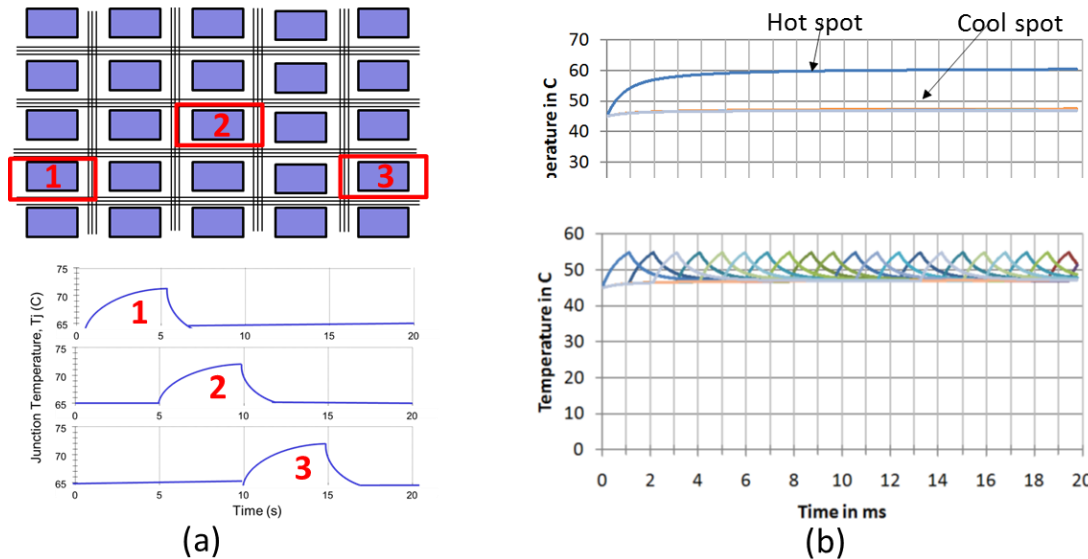
**Figure 10:** a) Run-time parameterizable FPGA-based image processing application, where application clock frequency, and number of Masks applied to an image (i.e. features scanned for) adjust to maximize performance for a given “Thermal Budget”, b) Summary of experimental results that quantify performance improvement when using temperature feedback to dynamically adjust application parameters, as compared to a static design that must meet thermal constraints under worst-case operating conditions. Scenario S1 represents worst-case operation conditions (e.g. broken fan, high ambient temperature), while S6 represents best-case operating conditions (e.g. operational fans, and low ambient temperature) [19].

### Improving System Performance.

*Problems Investigated.* Two brute force approaches typically used to address thermal issues are: 1) overprovisioning through the use of heatsinks and fans to account for worst-case thermal operating conditions, and 2) constraining an application performance so that thermal thresholds will not be surpassed while operating under such conditions by limiting operating frequency, or the amount of parallelism (area) the application is allowed to utilize. Over provisioning a system can be expensive in terms of cost, performance, power, and weight (a major concern in micro UAVs). Constraining application performance to meet worst case thermal conditions can cause a system to perform significantly below its realizable potential during typical operating conditions.

Here I give a brief overview of two approaches that I have investigated to make use of temperature feedback to increase FPGA-based application performance for a given thermal constraint (i.e. “Thermal Budget”), as compared to a static (or fixed) implementation design to meet temperature constraints under worst-case thermal conditions.

*Adapting Application parameters using temperature feedback[19].* Figure 10a depicts an image processing application that can scan images for up to 8 features (one Mask per feature). The number of features scanned in parallel is run-time parameterizable based on temperature feedback. Additionally, the frequency at which the design runs is a function of temperature. Figure 10b summarizes the performance evaluation result when comparing this adaptive architecture to a static version that meets temperature constraints under worst thermal operating conditions. As can be seen, the adaptive architecture manages excess temperature margins to significantly increase application performance while operating in better than worst-case thermal operating conditions. Under best-case thermal operating conditions, we see a 4x factor increase in image throughput (operating at 50 MHz vs. 200 MHz), and a 2x improvement in the number of features that could be scanned in parallel (4 vs. 8).



**Figure 11: a) Illustration of FPGA junction temperature as a computing module is moved from location 1 to 2 to 3, migration is triggered based on temperature reaching a specified threshold, b) “HotSpot” based simulation results, showing about a 5 degree decrease in maximum hot spot temperature, as compared to a design with a static layout [17].**


*Mitigating Hot Spots*[17]. I have proposed, implemented, and evaluated a method for mitigating hotspots using the dynamic partial self-reconfiguration capabilities available in modern FPGAs. This approach swaps “hot” and “cool” modules at run-time using a temperature threshold to trigger reconfiguration, thereby allowing an FPGA floor plan to adjust to its current thermal conditions (Figure 11a). Using “Hotspot” to run simulation evaluation experiments, this approach shows a reduction in maximum hotspot temperature by up to 8 °C (Figure 11b), and an increase in application throughput by up to a factor of 2-3 times that of a fixed design for a given thermal budget. [17]

### **3. Self Assessment of Accomplishments in and Impact of Institutional Service**

I have been given a number of opportunities to contribute to the Electrical and Computer Engineering Department. I have served in a number of departmental administrative positions, volunteered to advise 2-3 Senior Design projects each year, and have been Iowa State University’s proctor for IEEE’s Xtreme Programming competition.

#### **Departmental Administration.**

*Curriculum Committee.* I have served on the department’s Undergraduate Curriculum Committee over the last four years. During my tenure on this committee, I have participated in discussions with respect to performing a complete redesign of the curriculum, with an end goal of deploying a curriculum that more tightly couples Electrical and Computer engineering content. The vision is to give both EE and CPRE students the same course work for their Freshman and Sophomore year, therefore not having to commit to Computer Engineering or Electrical Engineering until the start of their Junior year, and to give them flexibility to look as much as an EE or CPRE as they want for their career goals.



**Figure 12:** Since 2008 I have advised a Unmanned Aerial Vehicle (UAV) design team. The 2012-2013 team is shown in the upper left, a high-level system view of their project is given in the upper right, and the lower figure shows the UAV in flight. Note, the RAVI board that was developed for CPS research is being utilized for UAV control in combination with a high-speed camera system.

*Computer and Networking Systems (CNS) Group Chair.* I have served as the chair of my research-area group since Spring of 2011. As chair, my responsibilities included 1) driving group teaching assignments, 2) managing the scheduling and staffing of PhD qualifying exams, and 3) representing my group's interests in the departmental Graduate Committee. Within the Graduate Committee my responsibilities have included giving input on Graduate Curriculum policies and evaluation students for College level recognition for excellence in graduate research, and excellence in teaching.

*Other Administrative Service.* For the department most recent ABET accreditation, I served a small role in helping gather evidence of compliance for the undergraduate embedded systems course (CPRE 288) that I teach. I have also served on the department's Graduate Admission committee for two semesters.

**Senior Design Advising.** I believe faculty involvement with undergraduate students outside of class is important not only in helping student development, but also as a means of direct feedback for how well the department's curriculum is meeting its goals. Since arriving at Iowa State University, I have been committed to hands-on mentoring of 2-3 Senior Design Projects each year. I have been most intensely involved with the ever evolving Microcontroller Controlled Aerial Robot Team (MicroCART) project. Over the last 5 years, I have co-advised students with Dr. Nicola Elia (an expert in controls systems) to develop an indoor test infrastructure for developing autonomous flying battery powered vehicles. A photograph of my most recent team can be found in the upper left part of Figure 12. The remainder of Figure 12 shows this year's team high-level system organization and their (Unmanned Aerial Vehicle) UAV in flight.

This project exposes teams composed of both electrical and computer engineers to real world issues that must be accounted for when applying knowledge from the class room to a complex system. In



**Figure 13: Over the last 3 years, I (upper right) have volunteered to proctor the IEEE 24 hour programming competition (IEEE Xtreme Programming Competition). My primary role is to make sure students follow the IEEE Code of Ethics during the competition, and that they stay energized and fed through the non-stop 24 hour period. This is a fantastic event for students to learn teamwork and perseverance in the face of challenging problems.**

In addition to learning the value of team work, students learn (often the hard way through spectacular UAV crashes) the value of testing software and hardware, appropriately characterizing noise associated with signals being processed, and systematically approaching and analyzing problems.

As the UAV infrastructure has become more refined, it has been leveraged to have senior design teams use the infrastructure to design and implement solutions to problems using autonomous vehicles. For example, this past year's team was able to have their UAV follow a small robot moving on the ground. This was demonstrated during the annual engineering open house (VEISHEA) that Iowa State University hosts each Spring.

**24-hour IEEE Xtreme Programming Competition.** As the name indicates, this is a non-stop 24 hour programming competition, sponsored by the IEEE Society. In summary, across the globe students form team of 4 students or less and over a 24-hour period 1-2 challenging problems are released every hour. The goal is to solve and upload the solution to as many problems as possible. Over the last 3 years, I have served as the Iowa State University proctor. My primary responsibilities have been to ensure teams comply with the IEEE code of ethics, stay well nourished, and keep their energy levels up. There are typically over 300 teams worldwide that compete, and the problems given are very challenging. From my proctoring experience, it is not uncommon for teams to go 5 or more hours before completing their first solution. I have found it to be a great opportunity for students to learn the importance of team work and perseverance. I look forward to proctoring this event each year; I use the student's energy to help me have a productive 24-hour Xtreme Research session.

#### **4. Self Assessment of Accomplishments in and Impact of Public, Extension and Professional Service**

While conducting professional level service has not been my focus since arriving at Iowa State University, I have been able to find small ways to serve.

*International Conference on Contemporary Computing (IC3).* In the Spring of 2011, I took on a leadership role as co-chair of the Systems track for this conference. It was eye opening how much work goes into identifying and recruiting potential reviewers. I also gained a greater appreciation for how valuable to the community volunteering to participate in the peer review process is.

*Academic Paper Reviewer.* I have been involved in reviewing papers for major conferences (e.g. IPDPS, ISCAS) and journals (e.g. IEEE Transactions on Computers, TRET, TECS, TPDS). Typically I review two batches of conference papers a year, and a few journal manual scripts.

*National Science Foundation (NSF) panel member.* Serving on a number of NSF grant panels has been among the most rewarding of my service experiences. I have enjoyed getting to meet new colleagues. Also gaining an understanding of what goes on behind the scenes during the review processes has made a profound impact on how I approach writing grants.

*Future Plans.* As I continue to advance in my academic career, actively seeking more opportunities for professional service is something that I am interested. For example, given the success that I have had with mentoring students for the ACM-IEEE sponsored MemoCODE design competition; this would be an activity that I would be interested in taking the lead running for a year or two.

#### **5. Self Assessment of Accomplishments in Impact of Technology Transfer**

N/A

#### **6. Self Assessment of Accomplishments in and Impact of Outreach Activities and Efforts to Improve Intellectual, Gender and Ethnic Diversity**

Since early in my career at Iowa State University, I have made efforts to support on-campus programs for increasing diversity and participation in STEM fields of study. As my career progresses, I would like to take on larger roles in support of this endeavor. The follow gives a brief description of four programs that I have been involved, and the manner in which I have participated.

*Digital Women.* The official stated charter of this University group is to: 1) encourage, support, and retain women in Computer Science, Software Engineering, Computer Engineering, Electrical Engineering, MIS, and any other field involved with electrical technology, 2) provide students with meaningful opportunities and social activities that promote positive relationships among students and faculty, and 3) encourage members of Digital Women to become an influence for good on campus and in the community.

This past year Digital Women officially became recognized as a University Student group, making me their first official mentor. I came into this role when several female students who had recently taken my undergraduate Embedded Systems course (CPRE 288) indicated that they enjoyed the





**Figure 14: Recent Digital Women group photograph. Being the mentor of this newly recognized university group has done much to increase my awareness of female student concerns within the Electrical and Computer Engineering department.**

manner in which I conducted lecture, and felt that I would make a good mentor for their group. My primary role as their mentor is to share my experiences as an engineer, and to convey to them a Faculty member's perspective on the topics they discuss. Attending their group meetings has given me an awareness of the needs and struggles that females in engineering majors face that I had never considered. I have also gained unexpected insight with respect to how they interact with one another. As compared to participating in group meetings that are primarily or entirely made up of male students, I have found Digital Women meetings have more of a social and collaborative quality to them, but still result in making great progress on a given topic.

***Leadership through Engineering Academic Diversity (LEAD).*** This College of Engineering program aims to enhance the experience of undergraduate multicultural and women students in the College of Engineering through events and services geared toward Academic, Professional, and Social development and success. LEAD programs and services strive to increase the number of successful multicultural and female engineering graduates by helping students get connected, become leaders, improve grades, graduate sooner, and become better prepared for engineer careers.

My participation with LEAD has included: 1) speaking with students to share my experiences as a member of an underrepresented group while making my way through an undergraduate degree, 2) answering questions during round-table discussions with respect to what I do for research and what opportunities I see in industry after they complete a degree in Electrical or Computer engineering, and 3) participating in recruitment dinners for highly qualified high school Juniors and Seniors.

***IT Olympics.*** This two day annual capstone event involves hundreds of high school students throughout Iowa for the purpose of 1) celebrating their year-long efforts toward developing IT based community service projects, and 2) competing in real-time challenges in the areas of Video Game Design, Robot Sumo, and Cyber Defense. The goal of the larger year-long "IT Adventures" program is to increase participation in technology throughout Iowa by getting high school students directly involved and using their projects to infuse technology into their local community.

I have helped organize the Game Design Competition during IT Olympics every year that I have been at Iowa State University. My responsibilities have included defining exciting real-time problems to keep students engaged during the event and judging their IT projects.

***Make-to-Innovate (M:2:I)***. This is a program within the Aerospace Engineering Department that engages students in hands-on projects to augment their understanding of engineering fundamentals. It encourages participation of students from all engineering areas who have an interest in pursuing projects in spacecraft design, control systems, embedded systems, or robotics. By giving students the opportunity to build (and break) their creations, they are introduced to the complexity of systems and the concept of design optimization, concepts they will carry forth throughout their career.

This program has a legacy of mentoring student driven projects, and encouraging student leadership. In addition, it has successfully attracted and retained female undergraduate participation in engineering projects. I have been a supporter of M:2:I, formally known as the Space Systems and Controls Laboratory, for most of my time at ISU. I have helped mentor their micro-satellite project called CySAT, and have participated in technology open houses that they have organized for local elementary schools.

## **II. Student, Peer, External and Other Assessments of Scholarship**

### **A. Assessments of Teaching and Learning**

#### **1. Student Ratings of Teaching Effectiveness**

- 4.67/5.00: Spring 2013, CPRE 584 Models and Techniques in Embedded Systems (10 students)
- 4.01/5.00: Spring 2013, CPRE 288 Embedded Systems (142 students)
- 4.48/5.00: Fall 2012, CPRE 288 Embedded Systems (48 students)
- 4.75/5.00: Spring 2012, CPRE 584 Models and Techniques in Embedded Systems (8 students)
- 4.42/5.00: Fall 2011, CPRE 288 Embedded Systems (45 students)
- 5.00/5.00: Fall 2011, CPRE 583 Reconfigurable Computing (15 students: 6 of which distance)
- 4.63/5.00: Spring 2011, CPRE 288 Embedded Systems: Lab Instructor (29 students)
- 4.91/5.00: Fall 2010, CPRE 583 Reconfigurable Computing (20 students: 6 of which distance)
- 4.89/5.00: Spring 2010, CPRE 584X Models and Techniques in Embedded Systems(10 students)
- 4.58/5.00: Fall 2009, CPRE 583 Reconfigurable Computing (16 Students: 2 of which distance)
- 4.83/5.00: Spring 2009, CPRE 594 Embedded Systems Research Skills (9 students)
  - Co-instructed with Joseph Zambreno
- 4.18/5.00: Fall 2008, CPRE 583 Reconfigurable Computing (21 students: 10 of which distance)

## **2. Peer Evaluations**

In the Spring of 2010, two of my peers (Morris Chang and Dionysios Aliprantis) evaluated my teaching of CPRE584 (Models and Techniques in Embedded Systems).

*Methodology.* This evaluation examined three aspects of my teaching: 1) course teaching material, 2) my classroom presence during lecture, and 3) student's perspective of my teaching through interviewing one student from class with respect to my teaching effectiveness.

*Summary of Results.*

1) Course material: It was concluded that the course material was well organized and available on the course website. This material acted as the primary content for the course, as there was no text book.

2) Classroom observation: It was observed that there was regular interaction between myself and the students. For improving my teaching effectiveness, it was suggested I proactively look for feedback from the students when moving between the use of PowerPoint slides and the dry-erase board with respect to readability (e.g. size of writing, brightness of markers).

3) Student interview: The student described the course as being composed of challenging projects given throughout the semester. The student felt that the interactions and discussions throughout the course were "excellent". With respect to improving the course, the student suggested adding more Architectural design aspects to the class.

*Conclusion.* Overall it was concluded that I was an effective teacher in terms of the topics covered in the course and the quality of my explanations.

## **3. Impacts**

Iowa State University is a "Land-Grant University", and as such its primary directive is the teaching of practical agriculture, science, and engineering. During my time at Iowa State University, I have taken to heart the importance of conducting lectures and developing assignments that make knowledge accessible to my students. As I look back at the impact of my teaching, there are four areas where I see that my striving for excellence in teaching has yielded fruitful results: 1) academic publications, 2) research infrastructure, 3) graduate student recruitment, and 4) direct student impact.

*Academic Publications.* The effort I put forth when first arriving at Iowa State University to enable distance-students to have the same experience as on-campus students in my Reconfigurable Computing (CPRE 583) course helped put in place the framework that led to two publications, one in the *Proceedings of the Annual Conference of American Society for Engineering Education (ASEE)* [11], and one in the *IEEE International Conference on Microelectronic Systems Education (MSE)* [15].

My Models and Techniques in Embedded Systems (CPRE 584) class' annual participation in the ACM/IEEE MemoCODE design challenge has led to a number of publications [10, 27, 28], and a few of my student's Reconfigurable Computing (CPRE 583) class projects have led to publications [9, 12, 14]. Additionally, winning MemoCODE in 2012 led to a couple of press releases.

**Research Infrastructure.** My teaching efforts in CPRE 583 attracted a highly capable Master student (Matthew Clausman) to work with me for a semester. During that semester, a reconfigurable development platform called RAVI was designed and fabricated targeting research for CPS. This board has since been used to provide evidence of capability for obtaining a funded NSF-EAGER and AFOSR grant, has been used in a number of publications[2, 3, 12, 31], and has additionally been used by my MicroCART Senior Design team for controlling a UAV.

Another student from my first offering of CPRE583 (Kent Vander Velden), who has since received a PhD in computational biology, was asked by XtremeData Corporation (a high-end reconfigurable computing company) to be allowed to port and deploy his final class project (a hardware accelerated bioinformatics application) during their demo at the 2010 ACM/IEEE Super Computing conference. This led to interactions with another high-end reconfigurable computing startup named Convey Computing. In turn my student set up initial communications between myself and Convey, which resulted in them providing their state of the art development system at a greatly reduced price with \$57,000 in donated parts from Xilinx Corporation. The industrial ties with Convey and the platform they have provided has increased the department's position for pursuing high quality research in the realm of high performance application acceleration using reconfigurable hardware technology.

The Convey Platform has since been used to win the MemoCODE design competition in 2012, aid in a number of publications, and now several PhD students are using this as their primary platform for conducting research.

**Graduate Student Recruitment.** My teaching of CPRE 583, CPRE 584, CPRE 288, and mentoring of Senior Design has resulted in a number of students deciding to attend graduate school, and/or continuing after their MS to pursue a PhD. Currently I have a couple of undergraduate students conducting undergraduate research with me, motivated by their experience in my undergraduate embedded systems course (CPRE 288), who will now likely continue their education at Iowa State University for a MS degree, and I am hoping to have at least one stay to pursue a PhD.

**Direct Student Impact.** Throughout my time at Iowa State University, I have had a number of students send me emails and letters in appreciation of my teaching and mentoring efforts. These communications have made clear the real-world impacts that my duties as an educator have on students' lives. I would like to share some of these comments.

Xin Zhao (email from a former CPRE 583 student, now at SanDisk): "I am starting to use FPGAs to do ASIC prototyping verifications... I am using 3 to 4 FPGA because the design is too big... I am using the muxing I/O technique I learned in your class... Thank you again".

Christy and Matt (letter from student leads of the IEEE Xtreme Programming Challenge): "Thank you for helping us again with IEEEExtreme. We truly appreciate it! Enjoy the pie!"

Cimone Wright (letter from student mentor within the LEAD program): "Thank you for representing the CprE Dept. at the LEAD round table... The students appreciate your time."

Cimone Wright (letter): "Thank you for the writing book you gave me. It has helped tremendously!"

Sarah, Max, Tommy, Dillon (letter from Freshman within "Take a Professor to lunch program"): "Thank you for meeting with us. Your answers to our questions is an invaluable resource"

#### **4. Other Assessments**

During the 2009-2010 academic year, I successfully completed the **Center for Excellence in Learning and Teaching (CELT) Teaching Partners Program**. In this program, I was paired with a senior lecturer (Cinzia Cervato, Full Professor from Geology) and another Assistant Professor. I attended regular meetings with my senior partner to discuss her experiences in teaching, observed her teaching, received feedback from her on my teaching, and attended a number of CELT classes related to improving teaching effectiveness.

#### **B. Assessments of Research and Creative Activities**

##### **1. Summary of Citations for up to 10 Publications**

Citations from: "Publish or Perish" 1/2013, with self-references manually removed

Impact factors from: Journal Citation Reports (JCR) 2009

**H-index = 5**

The Effects of an ARMOR-based SIFT Environment on the Performance and Dependability of User Applications,

K. Whisnant, R.K. Iyer, Z.T. Kalbarczyk, P.H. Jones III, D.A. Rennels, R. Some;  
*IEEE Transactions on Software Engineering (TSE)*, Volume 30, Issue 4, April 2004, Pages 257 - 277. (13 citations, Impact factor = 3.57)

Adaptive Thermoregulation for Applications on Reconfigurable Devices,

by Phillip H. Jones, James Moscola, Young H. Cho, and John W. Lockwood;  
*IEEE International Conference on Field Programmable Logic and Applications (FPL)*,  
Amsterdam, Netherlands, Aug 27-29, 2007.  
(6 citations)

A Thermal Management and Profiling Method for Reconfigurable Hardware Applications,

by Phillip H. Jones, John W. Lockwood, and Young H. Cho;  
*IEEE International Conference on Field Programmable Logic and Applications (FPL)*,  
Madrid, Spain, Aug 28-30, 2006.  
(5 citations)

Dynamically Optimizing FPGA Applications by Monitoring Temperature and Workloads,

by Phillip H. Jones, Young H. Cho, and John W. Lockwood;  
*IEEE International Conference on VLSI Design (VLSI Design)*, Bangalore, India, Jan 6-10,  
2007.  
(5 citations)

Liquid Architecture,

by Phillip Jones, Shobana Padmanabhan, Daniel Rymarz, John Maschmeyer, David V. Schuehler, John W. Lockwood, and Ron K. Cytron;  
*IEEE International Parallel and Distributed Symposium (IPDPS), Next Generation Software (NGS) Workshop*, Santa Fe, New Mexico, April 26, 2004.  
(5 citations)

Hotspot Mitigation using Dynamic Partial Reconfiguration for Improved Performance ,  
by Adwait Gupte, and Phillip H. Jones;  
*IEEE International Conference on Reconfigurable Computing and FPGAs (Reconfig)*,  
Cancun, Mexico, Dec 9-11, 2009.  
(4 citation)

Extracting and Improving Microarchitecture Performance on Reconfigurable Architectures,  
by Shobana Padmanabhan, Phillip Jones, David V. Schuehler, Scott J. Friedman, Praveen  
Krishnamurthy, Huakai Zhang, Roger Chamberlain, Ron K. Cytron, Jason Fritts, and John  
W. Lockwood;  
*International Journal of Parallel Programming*, Volume 33, Issue 2 - 3, June 2005, Pages  
115 - 136. (3 citations, Impact factor = .82)

Use of a Soft-Core Processor in a Hardware/Software Codesign Laboratory,  
by Roger Chamberlain, John Lockwood, Saurabh Gayen, Richard Hough, and Phillip Jones;  
*IEEE Intl. Conf. on Microelectronic Systems Education*, June, 2005.  
(3 citations)

Cycle-Accurate Microarchitecture Performance Evaluation,  
by Richard Hough, Phillip Jones, Scott Friedman, Roger Chamberlain, Jason Fritts, John  
Lockwood, Ron Cytron;  
*IEEE Workshop on Introspective Architecture (WISA)*, Austin, TX, February 2006.  
(2 citations)

## **2. Evidence of Impact on Society of Research and Creative Activities**

I would not classify my contributes to date as having significant impact at the research society level. However, I believe as I move toward more broadly generalizable results in the field of Cyber Physical Systems that some of those results will have impact at the research society level.

## **3. Other Assessments**

### **Best paper award VLSI Design 2007:**

Dynamically Optimizing FPGA Applications by Monitoring Temperature and Workloads,  
by Phillip H. Jones, Young H. Cho, and John W. Lockwood;  
*IEEE International Conference on VLSI Design (VLSI Design)*, Bangalore, India, Jan 6-10, 2007.

### **1<sup>st</sup> Place, 2012 ACM-IEEE MemoCODE design competition:**

Shepard: A Fast Exact Match Short Read Aligner,  
by C. Nelson, K. Townsend, B S. Rao, P. Jones and J. Zambreno;  
*Proceedings of the International Conference on Formal Methods and Models for Codesign (MEMOCODE)*, July, 2012

## **C. Assessments of Extension/Professional Practice and Service**

N/A

### III. Future Plans

As I continue with my academic career, I have a number of near and long-term plans I would like to follow for both research and education.

**Research.** The field of Cyber Physical Systems (CPS) is an exciting and growing field where I look forward to continuing to employ reconfigurable hardware to bring innovations to the computing stack. I have just scratched the surface of this research front with respect to hw/sw co-design approaches to support controls and real-time aspects of dynamic CPS applications.

Over the short term, I plan to have my two current PhD students bring their research to closure in the areas of hardware support for control algorithms, and hardware support for real-time scheduling. In the longer term, I will delve deeper into developing generalizations of their work in collaboration with colleges at Iowa State University who are more focused on control theory (e.g. Dr. Nicola Elia) to more formally unify hw/sw codesign trade-offs at the embedded system platform level with system stability concerns at the control theory level. Likewise, for the real-time scheduling aspects of my current work, I will pursue more generalizable contributions to better unify hw/sw codeisgn trade-offs at the embedded systems platform level with real-time schedule feasibility concerns.

**Education.** I am excited to continue contributing to innovating the courses and curriculum at Iowa State University. Over the short term, I am helping revamp two undergraduate embedded system courses (CPRE 288, and CPRE 488). The largest innovations will be to CPRE 488, where some of the outcomes from the UAV Senior design project that I advise will be applied. The end goal will be to bring more real world issues to the forefront with respect to controlling a physical system by having their labs involve integrating a mico-UAV into their embedded system.

In the longer term, I feel fortunate that I am at this department at a time when it is working to more holistically integrate the Electrical Engineering and Computer Engineering curriculum. Given my mix of training as an Electrical Engineering Undergraduate and Masters student, and my training as a Computer Engineer during my PhD, I believe I will be able to uniquely contribute to this process.

**Closing Thoughts.** I can honestly say that these five years at Iowa State University have been among the most rewarding of my life. The challenges that I have faced, while progressing through the tenure process, have forced me to mature and grow in ways that I do not think any other experience would have achieved. I am thankful for the journey I have taken thus far, and look forward to continuing to grow and contribute to society through academia.