T A J A N A   S I M U N I C   R O S I N G
U C S D ,   L A   J O L L A ,   C A
T A J A N A @ U C S D . E D U

Research overview

Research goals: I lead the System Energy Efficiency Lab, SEE Lab (seelab.ucsd.edu), at UCSD. The main focus of my work is on energy efficient computing. Energy consumption has already become one of the most important design parameters in electronic systems. Looking forward to a world where we are surrounded with sensors that we leverage via mobile computing devices, and where cloud computing is pervasive; energy efficiency will be at a premium even more than it is today. In large scale data centers it is because of the high cost of powering and cooling the equipment. Sections 1 & 2 outline related projects. For sensors and mobile systems it is because battery lifetime and energy available from harvesting sources are limited. This work is discussed in more detail in section 3. Whether it is low power devices, or large servers, the goals and methods are still the same – understanding how, where and why energy is spent in computing systems via measurement, characterization and modeling, and then designing and implementing new energy management strategies to ensure excellent performance with low energy consumption.

Publications and advising: In the short period since I started supervising my first graduate student at UCSD in the Fall of 2005 until Fall 2009 I have published 8 journal papers (with three more in review), 2 book chapters, 38 conference papers. Our VLSI-SOC 2009 paper got the **best paper award** [C51] and two others have been nominated for the best paper award [C28 , C50]. A paper was invited as a keynote speech [C30] and three others were invited conference papers [C23, C47, C49]. Over my academic career I have published 14 journal papers, 3 book chapters and 59 conference publications, along with 8 filed patents. One of my papers has been awarded as one of **the best papers in the last 10 years of DATE** conference, a top conference in my research area. Perhaps the most satisfying accomplishment is seeing my first PhD student at UCSD, Ayse Coskun, graduate and start her new job as an assistant professor at Boston University. I have three more PhD students solely advised by me who are expected to graduate during this year. A total of seven MS students have already graduated and went onto jobs in companies such as Cisco, Northrop-Grumman and others.

Technology transfer: My research has also resulted in a number of software and hardware artifacts. My students and I have designed a novel active sensing platform for structural health monitoring, Shimmer (for more details see below) that has been deployed jointly with Los Alamos National Labs on a bridge in New Mexico. LANL is planning to produce approximately 50 of our platforms for a larger scale deployment next year. Our algorithms for energy and thermally aware task scheduling are being included into Open Solaris kernel as a result of the last few years of collaboration with Sun Microsystems.

Funding and collaboration: Since starting my tenure-track position at UCSD, I have been involved in a number of different projects as a PI, co-PI or senior personnel totaling approximately $50M. Most recently, I was asked to head the Large Scale Systems thrust focusing on energy efficient datacenters within $3M MuSyC center. MuSyC center is funded jointly by DARPA and a number of companies. It consists of faculty from 11 top universities in USA, while the LSS thrust I am heading has nine faculty.

I enjoy collaboration, as I find that some of the best ideas come when working with others who have a different background than my own. For example, I have started a joint project between Dept. of CSE and UCSD School of Medicine through CalIT2 focusing on wireless healthcare systems. This work is currently funded by NIH and NSF. I have been an active contributor to Center for Networked Systems (CNS) where I have worked with a number of companies and faculty. I collaborate with San Diego Super Computing Center (SDSC) through MuSyC and HPWREN projects, CalIT2 via number of projects, including energy efficient datacenters work funded by NSF MRI GreenLight and MuSyC, structural engineering faculty and Los Alamos National Lab through CalIT2. My research has also had quite an impact outside of the academic community during this last review period, as can be seen by over 70 invited talks in industry, research labs, universities and abroad. For example, my work on energy and thermal management in multicore systems has been done in collaboration with Sun Microsystems, Qualcomm, Texas Instruments, Google, Yahoo, Intel, IBM research Zurich and TJ Watson, along with EPFL in Switzerland and others.

Service: My service to the broader community consists of two important elements: editorial/reviewer assistance and mentoring roles. Together they have exposed me and my research to the broader community with attendant recognition by prestigious publications and conference venues; and allowed me to lead as a responsible role model to the coming generation of women talent. I have been an Associate Editor for the IEEE Transactions on Mobile Computing since 2008. From 2003 until 2007 I was an Associate Editor for IEEE Transactions on Circuits and Systems. I have been an active member of the technical program committees of prestigious conferences in my area, such as DAC, DATE, and IPSN and have recently been a TPC co-chair for DATE. In addition, I've been a reviewer and moderator for about 20 journals and conferences. In my mentoring and leadership roles, I was an invited panelist at Grace Hopper conference, and a guest speaker at a Workshop for Women in Electronic Design Automation held at DAC.

In the remainder of this overview, I provide highlights and results of some of my recent research contributions. The complete references to the journal (J#) and conference (C#) publications can be find in my bibliography, Section III, part A.I for journals and A.III for the conference publications.

Summary of Major Research Areas

1. Power/Energy Management

Funded by: NSF, MARCO-GSRC, MARCO-MuSyC, Sun Microsystems, Cisco, CNS, UC Micro
In collaboration with: Sun Microsystems, Qualcomm, Cisco, Intel, TI, Google, Yahoo, SDSC, IBM TJ Watson

Online learning for adaptive power management:

A number of dynamic power management (DPM) policies have been proposed in the past with the goal of lowering the energy consumption by transitioning devices into sleep when idle. The key to a quality decision is modeling and predicting when an idle period is long enough to compensate for the cost of transition between power states. We found out from the measurement and characterization that workload inter-arrival times to various devices, ranging from hard disks, to wireless LAN cards, follow heavy-tailed distribution. Based on this I developed a new class of optimal event-driven power management policies [J2] that are capable of meeting performance constraints while minimizing the energy consumption. The algorithm is light weight, and has been implemented in both hardware (using only 900 gates) and software. When running within Windows OS with a goal of power managing the hard disk drive of a desktop computer, our algorithm got energy savings that are within 11% of maximum possible with realistic workloads. A benefit of the algorithm over previous heuristic and predictive policies is the ability to guarantee a desired level of performance. A key assumption made when deriving this optimal policy is the stationary of workload in a statistical sense. No single policy can guarantee optimality with non-stationary workloads common in real life systems.

We took an entirely different approach with our work presented in [J13] - instead of designing a new power management policy that performs reasonably well across many different types of ever changing workloads, we created a library of policies and a control algorithm that selects the best suited one for the current workload. The selection is done among a set of experts, which refers to a set of DPM policies and voltage-frequency settings, leveraging the fact that different experts outperform each other under different workloads and device leakage characteristics. The online-learning algorithm adapts to changes in the workload characteristics while guaranteeing fast convergence to the best-performing expert. In our evaluation, we perform experiments on a hard disk drive (HDD) and Intel PXA27x core (CPU) with real-life workloads. Our results show that our algorithm adapts really well and achieves an overall performance comparable to the best-performing expert at any point in time, with energy savings as high as 61% and 49% for HDD and CPU, respectively. Since it is so lightweight, we were able to deploy it in real systems with negligible overhead.

Energy management using virtualization:

The most recent thrust in SEE Lab's power management research is the focus on large scale data centers and the design of energy management algorithms using virtualization as the base technology. Modern data centers use virtualization (eg. Xen and VMware) to get better fault and performance isolation, improved system manageability and reduced infrastructure cost through resource consolidation and live migration. Consolidating multiple servers running in different virtual machines (VMs) on a single physical machine (PM) can increase the overall utilization and efficiency of the equipment across the whole deployment.

Through our involvement with NSF MRI GreenLight project that has funded deployment of two data center containers along with all the equipment inside them, we have been able to instrument and measure the effect various workloads have on energy and performance efficiency in larger scale deployments. We found that symbiotically placing VMs on individual servers can have a beneficial effect on both energy consumption and performance. As a result we designed vGreen [C50], a new a multi-tiered software system for energy efficient computing in virtualized environments. It uses hierarchical metrics to capture power and performance characteristics of virtual and physical machines and then to schedule the VMs accordingly. We show through real life implementation on a state of the art testbed that vGreen is able to dynamically characterize and schedule the VMs across the PM cluster, improving the overall performance and energy efficiency by 20% and 15% respectively compared to state of the art VM scheduling policies even when running at very high utilization levels. Furthermore, vGreen is lightweight with negligible runtime overhead. Our vGreen paper [C50] has been nominated for the best paper award at ISLPED 2009.

As energy efficiency in the datacenters is a strong function of how much power is spent on both powering and cooling systems, we plan to extend vGreen's capabilities beyond monitoring performance, estimating power consumption and controlling power states along with the workload placement, to also measuring temperature and controlling cooling system (e.g. fan speeds). Our recent work on thermal management, reliability and cooling control shows that the potential benefits of the combined approach are large, as is discussed in the next section.

2. Thermal and cooling management in computing systems (in both 2 and 3D integration)

Funded by: NSF, MARCO-GSRC, MARCO-MuSYC, Sun Microsystems, Cisco, CNS, UC Micro
In collaboration with: Sun Microsystems, Qualcomm, Cisco, Intel, Texas Instruments, IBM Zurich & TJ Watson, EPFL

Thermal management in 2D: Managing power states (DPM and DVFS) is only one part of the overall energy consumption. Another just as important aspect is thermal management, whose goal it is to ensure that temperature profiles are kept under critical levels, and are balanced in time and space. Thermal hot spots, spatial and temporal temperature gradients degrade reliability and performance, and increase cooling costs and leakage power. Since temperature is a strong function of power density, it would be tempting to think that minimizing power consumption in a system automatically solves any thermal issues. In our recent work [J12] we formulate integer linear programs (ILPs) whose goal is to first minimize energy consumption under performance constraints, and then to also meet thermal constraints when running known workloads on a multicore processor. Our implementation shows that even in today's processors minimizing energy consumption does not solve thermal problems, and can, in fact, increase them. When thermal constraints are considered, our ILP solution can reduce the frequency of hot spots by 35%, spatial gradients by 85% and thermal cycles by 61% in comparison to the ILP for minimizing energy, thus resulting in lower cooling costs and better reliability.

We next designed new reliability-aware job scheduling and power management algorithms for chip multiprocessors. Accurate evaluation of these policies requires a novel simulation framework that can estimate architecture-level effects over tens of seconds or longer, while also capturing thermal interactions among cores resulting from dynamic scheduling policies. Using our new framework and a set of novel thermal management policies, our work [C43] shows that techniques that offer similar performance, energy, and even peak temperature can differ significantly in their effects on the expected processor lifetime. Furthermore, we develop a method for finding the policy capable of giving minimum expected energy consumption under reliability and performance constraints [J11]. Our results show that on a quad core embedded processor we obtain 40% improvement in energy consumption in tandem with meeting a reliability constraint for all operating temperatures of the chip multiprocessor.

Proactive thermal management: Our implementation on multicore CPUs shows that reactive strategies, such as voltage scaling when temperature threshold is reached, are effective in lowering the temperature but come at a high performance cost. With that in mind, we developed a new set of proactive management policies [J14,C33, C36, C44], which predict the future temperature based on currently sensed data and adjusts the thread allocation on the multicore CPU to minimize the impact of thermal hot spots and temperature variations without degrading performance. The experimental results using real workloads on an actual multicore system show that our proactive technique is able to dramatically reduce the adverse effects of temperature by 60% while saving energy.

Thermal management in 3D systems: We have also studied how our thermal and energy management techniques can be generalized to 3D integration [C41, C42]. A benefit of 3D stack architectures is the reduction in the length of the wires going across the chip, but it comes at a cost of serious thermal challenges due to the high power density resulting from placing computational units on top of each other. We both study how 2D thermal and power management techniques perform in 3D and also propose a dynamic thermally-aware job scheduling technique for 3D systems to reduce the thermal problems at very low performance cost. In higher end processors the thermal problems become critical in 3D integration. As a result, we have looked at how liquid cooling can be used to reduce the thermal hot spots [C51, C58]. For that work we got a **best paper award**. Our proactive management technique provides an additional 75% reduction in hot spots in comparison to applying only liquid cooling on a 3D chip multiprocessor. Furthermore, for systems capable of varying the coolant flow rate at runtime, our feedback controller increases the improvement to 95% on average.

Cooling and workload scheduling: Most recently we have studied the effect of combining forced air cooling and job scheduling. Most work to date reports results of thermal management while assuming fixed ambient temperature conditions. This clearly is not the case in real systems due to non-uniform rate of heat dissipation in packages and cooling via fans. We designed a cooling aware dynamic workload management strategy [C53, C55] that is significantly more energy efficient than state-of-the art solutions in multi-socket CPU systems because it performs workload scheduling in tandem with controlling socket fan speeds. Our experimental results indicate that applying our scheme gives average fan energy savings of 85% concurrently with reducing the maximum fan speed by 53%, thus leading to lower vibrations and noise levels.

Our work to date indicates that significant energy savings are possible if we identify and eliminate large sources of inefficiency present in various layers of the computing system hierarchy. Thus, we plan to develop a multi-scale crosslayer distributed and hierarchical management scheme that ensures that energy is only consumed when needed in computing systems. This work will be done as a part of a Large Scale Systems thrust of MuSyC center that I'm heading, and will leverage two modular data centers complete with racks full of servers and high speed optical interconnect infrastructure we obtained at UCSD as a part of NSF GreenLight project.

3. Energy management in heterogeneous wireless networks

Funded by: NSF, NIH, DOE & LANL, CNS, Qualcomm, Sun Microsystems, HP Labs
In collaboration with: UCSD School of Medicine, SDSC, LANL, Sun Microsystems, Qualcomm, Seacoast Science, HP Labs, University of Bologna, Stanford University, Georgia Tech

The energy and thermal management issues that plague large scale computing systems are also present in battery operated devices. For example, our work on minimizing the energy consumption of mobile phones has covered multiple aspects, ranging from lowering the energy consumption due to wireless communication [J8], software energy profiling and optimization [J3,J5,J7], operating system schedulers to ensure energy efficient operation [J13], and thermal management techniques for heterogeneous multicore processors used in mobile phones whose goal is to control the temperature distribution so that packaging costs can be kept to minimum necessary [C54]. Five patents have been filed related to this work; the students involved have since graduated and went to academia (Politecnico di Torino), MIT research lab and various companies.

More recently, I've also worked on energy management in heterogeneous wireless sensor networks, consisting of devices with and without energy harvesting. The key differentiation of our work is the focus on more computationally intensive devices, and the resulting hierarchical sensing systems. This has been motivated by the fact that the increase in the performance/Watt in computation has been much more rapid than the corresponding increase in bandwidth/Watt of communication. Thus, instead of creating sensor networks whose mode of operation is to gather the data and send to the back end for storage and processing, we designed smart sensor nodes capable of significant in-situ processing, sensing and actuation (e.g. Shimmer node described below). Furthermore, we designed energy management algorithms for such hierarchical networks that can accurately predict energy availability from energy harvesting sources, and based on the estimated energy available can perform energy efficient task scheduling across a set of devices (Shimmer and CitiSense projects), along with communication scheduling and routing (HPWREN project).

CitiSense: The CitiSense project funded recently by NSF represents a next step in networked sensing systems by combining in a new way large scale cell phone networks with environmental and body area sensor nets. The goal of CitiSense is to build and show how a such an infrastructure can make it possible for various interested parties to understand the best strategies to prevent disease development and to help manage chronic illnesses, such as asthma and diabetes, which develop as a function of genetics, environment and individual choice. CitiSense is envisioned as an end-to-end health and environmental information system with near real-time data streams and feedback loops from the system to the sensing, processing, and actuation infrastructure. A key to success of CitiSense is the ability to perform sophisticated processing in the field so that quality feedback can be provided to people using the network on the daily basis. For this we need adaptive task and communication management algorithms that handle allocation of processing, sensing, and communication tasks to devices of the system. There are several dynamically changing characteristics in the system, such as various task parameters (execution time, arrival rate, and output data), the available battery capacity in various devices, and varying wireless channel conditions, network load, task arrival rates, and processing complexity. Our dynamic task assignment algorithm [C46, C56] improves the overall system lifetime under varying dynamic conditions by up to 35% with respect to an optimal static design time assignment. Going forward, CitiSense will leverage smart sensor nodes such as Shimmer for intelligent environmental monitoring with energy harvesting, and will use the large scale environmental monitoring network (HPWREN) already deployed in San Diego area.

Shimmer: Shimmer is a wireless platform for active sensing that combines localized processing with energy harvesting [C31]. This platform was initially designed for structural health monitoring, but since the design is modular, it can easily be adapted to any type of more sophisticated sensing/actuation systems requiring up to 600 MIPS of processing power. Unlike other sensor networks that periodically monitor a structure and route information to a base station, our device sends a wave via piezoelectric sensor (PZT) with a peak-to-peak amplitude of +/-15V, acquires data via another PZT at the rate of 10M samples/sec, and processes it locally either in steady-state operation mode, or after being radio-triggered by an external agent. Analysis of up to 120 combinations of PZT paths can be done in minutes, enabling not only damage detection, but also localization, while relying only on energy harvesting via solar power and large capacitors (supercaps) for energy storage; there are no batteries. This completely changes the sensing paradigm – instead of assuming the network consists of small, low-power and low computational capability sensors that route the data to the back end, we now have a network of intelligent devices capable of orders of magnitude more computation, with much higher bandwidth and fidelity in sensing and actuation, while running just off energy harvesting. An invention disclosure has been filed, and the device has already been deployed on a bridge in New Mexico as a part of a joint project with Los Alamos National Lab. As a part of this project, we developed a new set of task management algorithms that use accurate estimates of available energy from supercaps and harvesting sources to then enable tradeoffs in task accuracy, time to completion and available response time. This work thus far resulted in a patent filing, seven publications [J10, C30-31, C47-49, C57, B.I.1], of which two were invited papers [C47, C49], and one was an invited keynote [C30].

HPWREN: High Performance Wireless Research and Educational Network (HPWREN) provides wireless connectivity to a number of different sensors with varying resource requirements, such as large bandwidth requirements of the Palomar observatory, low bandwidth but tight real-time traffic deadlines of seismic sensor nodes, and long battery lifetime requirements of small and remotely deployed weather stations. It covers an area of approx 20,000 sq. miles via a hierarchical structure consisting of sensors as the bottom layer, sensor node cluster heads in the middle and high speed wireless mesh network at the top. Given many nodes at the sensor and cluster head layer are battery powered or use energy harvesting, the ability to dynamically trade off priority of various data traffic sources with energy consumption is key to the long term operation of this network. QoS methodologies used for the internet do not apply due to stringent energy and computational requirements of the sensor nodes, highly variable wireless channel conditions, and the continually changing network topology. We address this problem by looking at scheduling and routing data at sensor node cluster heads, so that QoS requirements are met while the cost of energy spent in communication is lowered by maximizing the time each node spends asleep. Our new scheduling and routing algorithm [C27, C35, C52, journal paper in review at ACM TOSN] operates above MAC layer thus allowing us to use off-the-shelf wireless connectivity, while reducing the power consumption by 60% and increasing data throughput by 10%.

Section III - Bibliography

A. PRIMARY PUBLISHED OR CREATIVE WORK

I. **Refereed Journal Articles** -

1. T. Simunic, J. Rozenblit, J. Brews: "VLSI Interconnect Design Automation Using Qualitative and Symbolic Techniques"; IEEE Transactions on Components, Packaging, and Manufacturing Technology Part B: Advanced Packaging, pp.803-812, November 1996.

2. T. Simunic, L. Benini, P. Glynn, G. De Micheli: "Event-Driven Power Management", IEEE Transactions on CAD, pp.840-857, July 2001.

3. T. Simunic, L. Benini, G. De Micheli: "Energy-Efficient Design of Battery-Powered Embedded Systems", Special Issue of IEEE Transactions on VLSI, pp. 18-28, May 2001.

4. T. Simunic, M. Smith: "Dynamic Power Management at HP", Invited Paper in the Special Issue of IEEE Design and Test of Computers, March-April 2001.

5. A. Peymandoust, T. Simunic, G. De Micheli: "Complex Instruction and Software Library Mapping for Embedded Software Using Symbolic Algebra," Special Issue of IEEE Transactions on CAD, pp.964-975, August 2003.

6. T. Simunic, S. Boyd, P. Glynn: "Managing Power Consumption in Networks on Chips," IEEE Transactions on VLSI, pp. 96- 107, Jan 2004.

7. B. Delaney, T. Simunic, N. Jayant: "Power Aware Distributed Speech Recognition for Wireless Mobile Devices," Special Issue on Embedded Systems for Multimedia, IEEE Design & Test of Computers, 2004.

8. A. Acquaviva, T. Simunic, V. Deolalikar, S. Roy: "Remote Power Control of Wireless Network Interfaces", Special Issue of Journal of Embedded Computing, No. 3, 2004.

9. A. Coskun, T. Simunic Rosing, K. Mihic, G. De Micheli, Y. Leblebici, "Analysis and Optimization of MPSoC Reliability," Invited paper to Journal of Low-Power Electronics, April 2006. \

---

10. Gyuhae Park, **Tajana Rosing**, Michael D. Todd, Charles R. Farrar, Wiliam Hodgkiss, "Energy Harvesting for Structural Health Monitoring in Sensor Networks," ASCE Journal, 2007. **NEW**

11. **Tajana Simunic Rosing**, Kresimir Mihic, Giovanni De Micheli, "Power and reliability management of SOCs," IEEE Transactions on VLSI, 2007. **NEW**

12. Ayse Kivilcim Coskun, **Tajana Simunic Rosing**, Keith A. Whisnant, Kenny C. Gross, "Static and dynamic temperature-aware scheduling for multiprocessor SoCs," IEEE TVLSI, 2008. **NEW**

13. Gaurav Dhiman, **Tajana Simunic Rosing**, "System-Level Power Management Using Online Learning," IEEE TCAD, 2009. **NEW**

14. Ayse Kivilcim Coskun, **Tajana Simunic Rosing**, Kenny C. Gross, "Utilizing Predictors for Efficient Thermal Management in Multiprocessor SoCs," IEEE TCAD, 2009. **NEW**

II. **Books and Book Chapters**

1. T. Simunic: "Dynamic Management of Power Consumption" in Power Aware Computing, Edited by R. Graybill, R. Mehlem, Kluwer Academic Publishers pp.102-125, 2002.

---

2. Jihong Kim, **Tajana Simunic Rosing**, "Power-aware resource management techniques for low-power embedded systems," in Handbook of Real-Time and Embedded Systems, Edited by S. H. Son, I. Lee, J. Y-T Leung, Taylor-Francis Group LLC, 2006. **NEW**

3. Yung-Hsiang Lu, Eui-Young Chung, **Tajana Simunic**, Luca Benini, Giovanni De Micheli: "Quantitative Comparison of Power Management Algorithms", in *The Most Influential Papers of 10 Years DATE*, Edited by Lauwereins, Rudy; Madsen, Jan, Springer-Verilag, 2008. **NEW**

III. **Refereed Conference Proceedings**

1. T. Simunic, P. Hsu, J. Rozenblit, C. Wolff, J. Prince, A. Cangellaris: "An Integrated Framework for Modeling and Simulation of Electronic Packaging," Proceedings of TECHCON, 1993.

2. T. Simunic, J. Rozenblit: "Reduction of Signal Delay and Crosstalk in Electronic Packaging Systems," Proceedings of Second Topical Meeting on Electrical Performance of Electronic Packages, 1993.

3. J. Rozenblit, T. Simunic: "Techniques for Intelligent VLSI Interconnect Design," Proceedings of NFS Design and Manufacturing Grantees Conference, 1994.

4. Y. Lu, T. Simunic, G. De Micheli: "Software Controlled Power Management", CODES 1999.

5. T. Simunic, L. Benini, G. De Micheli: "Cycle-Accurate Simulation of Energy Consumption in Embedded Systems", DAC 1999.

6. T. Simunic, L. Benini, G. De Micheli: "Energy-efficient design of Battery-Powered Embedded Systems", ISLPED 1999.

7. T. Simunic, L. Benini, G. De Micheli: "Event-driven Power Management of Portable Systems", ISSS 1999.

8. Y. Lu, E. Chung, T. Simunic, L. Benini, G. De Micheli: "Quantitative Comparison of Power Management Algorithms", DATE 2000. Selected for publication in ***The Most Influential Papers of 10 Years DATE***, Edited by Lauwereins, Rudy; Madsen, Jan, 2008.

9. T. Simunic, L. Benini, P. Glynn, G. De Micheli: "Dynamic Power Management of Laptop Hard Disk", DATE 2000.

10. T. Simunic, H.Vikalo, P. Glynn, G. De Micheli: "Energy Efficient Design of Portable Wireless Systems", ISLPED 2000.

11. T. Simunic, L. Benini, G. De Micheli, M.Hans: "Source Code Optimization and Profiling of Energy Consumption in Embedded Systems", ***Invited paper*** at ISSS 2000.

12. T. Simunic, L. Benini, P. Glynn, G. De Micheli: "Dynamic Power Management of Portable Systems", MOBICOM 2000.

13. T. Simunic, L. Benini, A. Acquaviva, P. Glynn, G. De Micheli: "Dynamic Voltage Scaling for Portable Systems", DAC 2001.

14. B. Delaney, N. Jayant, M. Hans, T. Simunic, A. Acquaviva: "Low-Power Fixed-Point Front-End Feature Extraction for Distributed Speech Recognition", ICASSP 2002.

15. A. Peymandoust, T. Simunic, G. De Micheli: "Low Power Embedded Software Optimization using Symbolic Algebra", DATE 2002.

16. T. Simunic, S. Boyd: "Managing Power Consumption in Networks on Chips", DATE 2002.

17. A. Peymandoust, T. Simunic, G. De Micheli: "Complex Software Library Element Mapping with Symbolic Algebra", DAC 2002.

18. B. Delaney, N. Jayant, T. Simunic: “A WLAN Scheduling Algorithm to Reduce the Energy Consumption of a Distributed Speech Recognition Front-End”, ESTIMedia 2003.

19. A. Acquaviva, T. Simunic, V. Deolalikar, S. Roy: “Remote Power Control of Wireless Network Interfaces”, PATMOS 2003.

20. W. Quadeer, T. Simunic, J. Ankcorn, V. Krishnan, G. De Micheli, “Heterogeneous wireless network management”, PACS 2003.

21. O. Celebican, T. S. Rosing, V. J. Mooney: “Energy estimation of peripheral devices in embedded systems,” GLVLSI 2004.

22. G. Manjunath, V. Krishnan, T. Simunic, J. Tourrilhes, A. McReynolds, D. Das, V. Srinivasmurthy, A. Srinivasan: “Smart Edge Server – going beyond a wireless access point,” WMASH’04.

23. T. Simunic, K. Mihic, G. De Micheli: “Reliability and Power Management of Integrated Systems,“ ***Invited paper*** at DSD’04.

24. T. Simunic, W. Quadeer, G. De Micheli: “Managing heterogeneous wireless environments via Hotspot servers, “ MMCN’05.

25. T. Simunic, K. Mihic, G. De Micheli: “Optmization of Reliability and Power Consumption in Systems on a Chip,“ PATMOS’05.

26. A. Coskun, T. Simunic Rosing, “A Simulation Methodology for Reliability Analysis in Multi-Core SoCs,” GVLSI’06.

---

27. Daeseob Lim, Jaewook Shim, **Tajana Simunic Rosing**, Tara Javidi, “Scheduling data delivery in heterogeneous wireless sensor networks,” ISM’06. **NEW**

28. Gaurav Dhiman, **Tajana Simunic Rosing**, “Dynamic Power Management Using Machine Learning,” ***Nominated for the best paper award*** at ICCAD’06 **NEW**

29. Gaurav Dhiman, **Tajana Simunic Rosing**, “Dynamic Voltage Frequency Scaling for Multi-Tasking Systems using Machine Learning,” ISLPED’07. **NEW**

30. Michael Todd, David Mascarenas, Eric Flynn, **Tajana Rosing**, Ben Lee, Daniele Musiani, Sanjoy Dasgupta, Samori Kpotufe, Daniel Hsu, Rajesh Gupta, Gyuhae Park, Tim Overly, Matt Nothnagel, Chuck Farrar, "A different approach to sensor networking for SHM: Remote powering and interrogation with unmanned aerial vehicles", ***Keynote*** at Workshop on Structural Health Monitoring, 2007. **NEW**

31. Daniele Musiani, Kaisen. Lin, **Tajana Simunic Rosing**, “An Active Sensing Platform for Wireless Structural Health Monitoring Application,” IPSN-SPOTS’07. **NEW**

32. Ayse Kivilcim Coskun, **Tajana Simunic Rosing**, “Temperature-aware task scheduling in MPSoCs,” DATE’07. **NEW**

33. Ayse Kivilcim Coskun, **Tajana Simunic Rosing**, Kenny Gross, “Proactive temperature balancing for low cost thermal management in MPSOCs,” ICCAD’08. **NEW**

34. Gaurav Dhiman, Kishore Kumar Pusukuri, **Tajana Simunic Rosing**, “Analysis of Dynamic Voltage Scaling for System Level Energy Management,” USENIX-HotPower’08. **NEW**

35. E. Regini, D. Lim, **T. Simunic Rosing**, “Distributed scheduling for heterogeneous wireless sensor networks,“ IASTEAD’08. **NEW**

36. Ayse Kivilcim Coskun, **Tajana Simunic Rosing**, Kenny Gross, "Proactive temperature management in MPSOCs," ISLPED'08. **NEW**

37. Ayse Kivilcim Coskun, **Tajana Simunic Rosing**, Kenny Gross, "Temperature management in MPSOCs using online learning," DAC'08. **NEW**

38. Ayse Kivilcim Coskun, **Tajana Simunic Rosing**, "Temperature-aware MPSOC scheduling for reducing hot spots and gradients," ASPDAC'08. **NEW**

39. Shervin Sharifi, **Tajana Simunic Rosing**, "An analytical model for the upper bound on temperature differences on a chip," GLVLSI'08. **NEW**

40. Shervin Sharifi,ChunChen Liu, **Tajana Simunic Rosing**, "Accurate temperature estimation for efficient thermal management," ISQED'08. **NEW**

41. Ayse K. Coskun, Jose L. Ayala, David Atienza, **Tajana Simunic Rosing,** Yusuf Leblebici. "Dynamic Thermal Management in 3D Multicore Architectures," DATE 2009.

42. Ayse K. Coskun, Andrew B. Kahng, **Tajana Simunic Rosing**, "Temperature- and Cost-Aware Design of 3D Multiprocessor Architectures", DSD'09. **NEW**

43. Ayse K. Coskun, Richard Strong, Dean M. Tullsen, and **Tajana Simunic Rosing**, "Evaluating the Impact of Job Scheduling and Power Management on Processor Lifetime for Chip Multiprocessors," SIGMETRICS'09. **NEW**

44. Raid Ayoub, **Tajana Simunic Rosing**, "Predict and Act: Dynamic Thermal Management for Multi-core Processors," ISLPED'09. **NEW**

45. Gaurav Dhiman, Raid Ayoub, **Tajana Rosing**, "PDRAM: A hybrid PRAM DRAM main memory system", DAC'09. **NEW**

46. Priti Aghera, **Tajana Simunic Rosing,** Diana Fang, Kevin Patrick "Energy management in wireless healthcare systems," IPSN'09. **NEW**

47. Jamie B. Steck, **Tajana S. Rosing**, "Adapting Performance in Energy Harvesting Wireless Sensor Networks for Structural Health Monitoring Applications," ***Invited paper*** at IWSHM'09. **NEW**

48. Jamie Bradely Steck, **Tajana Simunic Rosing**, "Adapting Task Utility in Externally Triggered Energy Harvesting Wireless Sensing Systems," INSS'09. **NEW**

49. Joaquin Recas Piorno, Carlo Bergonzini, David Atienza, **Tajana Simunic Rosing**, "Prediction and Management in Energy Harvested Wireless Sensor Nodes," ***Invited paper*** at Wireless VITAE'09. **NEW**

50. Gaurav Dhiman, Giacomo Marchetti, **Tajana Simunic Rosing**, "vGreen: A System for Energy Efficient Computing in Virtualized Environments," ***Nominated for the best paper award*** at ISLPED'09. **NEW**

51. Ayse K. Coskun, Jose L. Ayala, Daivd Atienza, **Tajana Simunic Rosing**, "Modeling and Dynamic Management of 3D Multicore Systems with Liquid Cooling," ***Best paper award*** at VLSI-SOC 2009. **NEW**

52. E. Regini, T. Simunic Rosing, "An Energy Efficient Wireless Communication Mechanism for Sensor Node Cluster Heads," to appear at ISSNIP'09. **NEW**

53. R. Ayoub, T. Simunic Rosing, "Cool and Save: Cooling Aware Dynamic Workload Scheduling in Multi-socket CPU Systems," to appear at ASPDAC'10. **NEW**

54. S. Sharifi, A. Coskun, T. Simunic Rosing, "Hybrid Dynamic Energy and Thermal Management in Heterogeneous Embedded Multiprocessors," to appear at ASPDAC'10. **NEW**

55. R. Ayub, S. Sharifi, T. Simunic Rosing, "GentleCool: cooling aware proactive workload scheduling in multi-machine systems," to appear at DATE'10. **NEW**

56. P. Aghera, A. Coskun, D. Fang, D. Krishnaswamy, T. Simunic Rosing, "DynAHeal: Dynamic energy efficient task assignment for wireless healthcare systems," to appear at DATE'10. **NEW**

57. A. Sitaraman, D. Dondi , T. Simunic Rosing, "DVFS Based Task Scheduling in a Harvesting WSN for Structural Health Monitor," to appear at DATE'10. **NEW**

58. A. Coskun, D. Atienza, T. Simunic Rosing, "Energy-efficient variable-flow liquid cooling in 3D stacked architectures," to appear at DATE'10. **NEW**

IV. Patents

1. T. Simunic, N. Mehta, C. Crome : "Method for Test Vector Analysis"; USA Patent No. 5,923,567, granted 2005.

2. V. Deolalikar, T. Simunic: "Method and system for power control", USA Patent No. 7,190,980, granted 2007.(Patent changed name slightly when granted)

3. T. Simunic, A. Acquaviva, L. Benini, "Application-driven method and apparatus for limiting power consumption in a processor-controlled hardware platform," USA Patent No. 7,272,730, granted 2007. .(Patent changed name slightly when granted)

4. T. Simunic, A. Acquaviva, "Server-driven Power Management," filed 2003.

5. T. Simunic, O. Celebican: "Estimation and optimization of system I/O device power consumption," filed 2004.

6. T. Simunic, "Wireless resource management in heterogeneous networks," filed 2004.

---

7. T. Simunic, N. Mehta, C. Crome : "Device for Test Vector Analysis"; USA Patent No. 6,197,605, granted 2005.(Originally submitted with Patent 1. as one patent, two were ultimately granted)

8. D. Musiani, K. Lin, T. Simunic Rosing, D. Mascarenas, E. Flynn, C. Farrar, G. Park, "Wireless node for active sensing in SHM applications," filed 2006. **NEW**

B. OTHER WORK

I. **Other Conference Proceedings**

1. J. Recas, C. Bergonzini, B. Lee, T. Simunic Rosing, "Solar energy harvesting prediction algorithm," Energy Harvesting Workshop'09. **NEW**

II. **Abstracts of Non-refereed Conference Proceedings**

III.**Other Presentations/Works**

Theses (supervised & published)

1. T. Simunic, "Energy efficient system design and utilization," PhD 2001. **(Previously listed as B III 1)**

2. T. Simunic, "VLSI interconnect design automation using qualitative and quantitative techniques," MS 1993. **(Previously listed as B III 2)**

---

3. D. Lim, "Distributed proxy-layer scheduling in heterogeneous wireless networks," MS 2007. **NEW**

4. D. Musiani, "Design of an active sensing platform for wireless structural health monitoring," MS 2007. **NEW**

5. E. Regini, "Resource management in heterogeneous wireless sensor networks," MS 2009. **NEW**

6. J. Steck, “Energy and task management in energy harvesting wireless sensor networks for structural health monitoring,” MS 2009. **NEW**

7. C. Bergonzini, “Management of solar harvested energy in actuation based embedded systems,” MS 2009. **NEW**

8. A. K. Coskun, “Efficient Thermal Management for Multiprocessor Systems,” PhD 2009. **NEW**

Technical Reports (Previously listed as B III 3 – 9)

1. B. Delaney, M. Hans, T. Simunic, A. Acquaviva: “Low-Power, Fixed-Point, Front-End Feature Extraction for a Distributed Speech Recognition System”, HP Labs Technical Report, 2001.

2. CT Mar, M. Hans, M. Smith, T. Simunic, R. Schafer: “ A High-Quality, Energy Optimized, Real-Time Sampling Rate Conversion Library for the StrongARM Microprocessor”, HP Labs Technical Report, 2002. .

3. A. Acquaviva, T. Simunic, L. Benini, “Application-driven power-aware RTOS enhancements,” HP Labs Technical Report, 2003.

4. A. Acquaviva, T. Simunic, V. Deolalikar, S. Roy: “Server-driven Power Management,” HP Labs Technical Report, 2003.

5. W. Quadeer, T. Simunic, J. Ankcorn, V. Krishnan, G. De Micheli, “Heterogeneous wireless network management”, HP Labs Technical Report, 2003.

6. O. Celebican, T. Simunic, “Energy Estimation of Peripheral Devices in Embedded Systems,” HP Labs, Technical Report 2003.

7. B. Delaney, T. Simunic, N. Jayant, “Power Aware Distributed Speech Recognition for Wireless Mobile Devices,” HP Labs Technical Report, 2004.

C. WORK IN PROGRESS

1. E. Regini, D. Lim, T. Simunic Rosing, “Energy management in heterogeneous wireless sensor networks,” under review in ACM TOSN, 2009.

2. J. Recas, C. B, T. Simunic Rosing, D. Atienza, “Energy management in SHM”, under review for a special issue of JIMSS, 2009.

3. S. Sharifi, T. Simunic Rosing, “Accurate direct and indirect on-chip temperature sensing for efficient dynamic thermal management,” under review in IEEE TCAD, 2009.