# Integrated Circuit and Systems Sciences
# Research Needs: Integrated System Design
# June 2013

The Semiconductor Research Corporation (SRC) Global Research Collaboration (GRC) program member companies are pleased to release this document that describes the research needs in the thrust of Integrated System Design including that for Texas Analog Center of Excellence (TxACE). Incorporated into this document are the Grand Challenges from the International Technology Roadmap for Semiconductors (http://public.itrs.net) as well as needs identified through SRC GRC Executive Technical Advisory Board (ETAB) priorities, the Integrated Circuit and Systems Sciences (ICSS) strategic planning process, TxACE workshops and joint discussions.

The system design research needs of the members are described in nine major categories:

- Early Design Space Exploration at System Level
- System Level Architecture and Implementation
- Design Robustness
- System Power Optimization
- Multi-core SoC Design
- Large Scale Computing Applications
- Algorithms for Sensing and Real Time Embedded Control/Communication Applications
- SW/HW Co-development
- Security and Anti-Counterfeiting

Each of these major categories is broken down into several sub-categories which describe the need in more detail. Even so, these are written to be broad in nature to not restrict the investigator's approach.

The needs in the system design space cover a broad range of applications, including high performance processors for data centers, low power processors for mobile computing and communication, healthcare devices, and efficient energy usage and management systems. Investigators are encouraged to link the results of their work with a potential application to help describe the relevance of the proposed work.

White Papers for all the categories below will be considered for funding. Note that there is no priority assigned to the order of the major categories or the sub categories. Investigators are limited to participation in two white papers in this Integrated System Design solicitation. Submissions should highlight which needs are addressed, such as "S2.1".

**CONTRIBUTORS**


| | |
|---|---|
| AMD | Srilatha Manne, Steve Kosonocky, Gabriel Loh |
| Freescale | Gary Morrison, Doug Garrity, Colin McAndrew |
| IBM | Prabhakar Kudva |
| Intel | Robert Nguyen |
| TI | Anand Dabak, Steven Bartling |

# 2013 Integrated System Design Needs Categories

| | |
|---|---|
| **S1** | **Early Design Space Exploration at System Level** |
| *The increasing inter-dependencies of primary system design objectives such as power, performance, reliability, security, and programmability are simultaneously complicated by technology trends toward new memory devices (e.g., NVRAM), new packaging options (e.g., die stacking), new interconnect technologies (e.g., photonics), and the inherent complexity from continuously increasing the number of devices per system. Research needs include developing modeling methods, tools, and algorithms, and exploring systems, architectures, and sub-system to enable effective systems that exploit (or mitigate) technology trends.* | |
| S1.1 | High level executable models of complex systems and their use for design and validation |
| S1.2 | Algorithms and methods for quantifying performance-power reliability trade-off in early design |
| S1.3 | Workload-adaptive/reconfigurable heterogeneous multicore systems (processors, accelerators), SoC, NoC, and I/O architectures |
| S1.4 | High level exploration for systems utilizing 3-D integration and heterogeneous technologies |
| S1.5 | High level modeling, design and synthesis flows to enable early analog system design and tradeoff analysis for reconfigurable analog sub-systems |
| S1.6 | Exploration of architecture and systems that incorporate novel memory types such as MRAM, PRAM and RRAM in addition to or as replacements of traditional memory such as SRAM and e-DRAM |
| S1.7 | Early-architectural trade-off analysis of fault-detection, -correction, -containment, and –repair architectures |
| **S2** | **System Level Architecture and Implementation** |
| *Demanding applications from data centers to smart devices continue to drive for systems with higher performance with respect to all merits. Of high value are research for novel architectures, communication and microarch for processor cores, memory subsystems, and analog/RF subsystems, in both server and SoC domains, that significantly boost performance while effectively managing power and energy consumption, robustness and cost. Of equal importance are research to address effective implementation, optimization and validation of large mixed signal heterogeneous systems, possibly with reconfigurable subsystems and emerging technologies.* | |
| S2.1 | System level validations of performance and robustness including both hardware and software |
| S2.2 | Analog/RF/mixed-signal sub-systems architectures, including ultra-low power applications |
| S2.3 | Effective and fast turn-around design flows for SoC/SiP/3D, including high performance IP blocks (e.g. processor cores) |
| S2.4 | System architectures having access to significantly higher off-chip bandwidth (e.g. optical) |
| S2.5 | Microarchitecture synthesis including customized reconfigurable accelerators |
| S2.6 | Latency insensitive and elastic designs |
| S2.7 | System correctness/validation and optimization in aggressive power and energy management schemes |
| S2.8 | Novel memory sub systems and architectures enabled by the inclusion of emerging non-volatile memory technologies (related to S1.6) |
| S2.9 | High-performance processor microarchitectures |
| S2.10 | Architectural options for SoCs integrating reconfigurable hardware and interconnect at various levels of granularity |
| **S3** | **Design Robustness** |
| *Reliability continues to be a significant system-design concern. Continuing primary research foci include estimation/modeling, avoidance, detection, recovery, correction, containment, and repair of soft/transient and hard/permanent failures, plus failures from gradual deterioration. Emphasis should be upon less-than-100%-redundancy mechanisms, and upon reducing susceptibility to known failure mechanisms in the first place.* | |
| S3.1 | Low-power, low-cost fault-detection, -correction, -containment, and –repair micro/architecture |
| S3.2 | System-repair architectures with quantifiable performance degradation, at reasonable power and cost overhead |
| S3.3 | Highly available architectures (uninterrupted execution upon failure), at substantially reduced power and cost overhead |
| S3.4 | Design for correctness and completeness of fault-tolerant architectures |
| S3.5 | Novel architecture, microarchitecture, software, and communication protocols that tolerate excessive and mismatch process variation in very deep sub-micron nodes. To control yield and cost, design methodologies are needed to deliver and verify system correctness at time zero, without heavy dependence upon redundancy/repair techniques or variation control. |
| S3.6 | Much less than 100% redundancy fault tolerance schemes |

| S4 | System Power Optimization |
|---|---|
| *One of the major concerns in integrated system design is power management and energy minimization. Traditional performance-power trade-off techniques are inadequate to address future technology leakage power and computational demands in a reliable manner. Multicore and heterogeneous architectures will be the norm. Future power gains will likely be made by the addition of dedicated accelerators and reconfigurable hardware. Software and hardware cooperation is needed to optimize use of resources in different parts of the system. Novel global optimization techniques will be needed at every design step, top to bottom, to achieve the targeted solution, addressing power for computation, memory and on-die and on-package communication.* | |
| S4.1 | Efficient power management to enable high performance, highly integrated, multi-function SoC |
| S4.2 | Novel techniques for reducing on-die data communication power |
| S4.3 | Reducing di/dt through microarchitectural design and software techniques |
| S4.4 | Novel techniques for resource scheduling and workload balancing |
| S4.5 | Dynamic and/or static power management and optimization for micro watt to high power systems |
| S4.6 | SoC/SiP/3D microarchitectures that support multiple power grids, multiple VRM, multiple system power states/clocks |
| S4.7 | Reconfigurable-hardware systems for power reduction |
| **S5** | **Multi-core SoC Design** |
| *Multi-core, many-core, and heterogeneous processor designs provide ways of continuing performance scaling while leaving within the limitations of fixed power, energy, and thermal budgets. New architectures and design methodologies for such processors will be needed to continue providing more performance across a wide range of form factors, performance targets, and power constraints. In addition to novel processor organizations and accelerators, further research is needed in scalable and flexible interconnection fabrics and corresponding protocols to tie together the many different components of future SoCs. At the same time, emerging die-stacking and system-in-package technologies provide new opportunities (mixed process technologies, dense integration, high chip-to-chip bandwidths, IP reuse) to innovate novel processor and interconnect architectures, and research is required to determine the most promising approaches.* | |
| S5.1 | Multi-core or many-core (homogeneous and/or heterogeneous) architectures and design methodologies for micro watt to high power systems |
| S5.2 | Design and validation of on-chip communication fabrics and protocols |
| S5.3 | SiP/3D chip-to-chip architectures/interfaces and design methodologies |
| **S6** | **Large Scale Computing Applications** |
| *Co-design of hardware (compute, memory, network, storage), software and middleware for large scale data center and workload optimized system design for performance, power and reliability Architecture development of both centralized and decentralized large scale data centric (big data ) and compute processing driven systems, with different response time requirements (from mission critical, real-time to longer time aggregation and analysis). Novel architectures which incorporate a combination of new chips, embedded software, middleware and data center software system management. Agility and flexibility through virtualization, programmability and dynamic hardware and software methods.* | |
| S6.1 | Large scale system (including datacenter) aware co-design of hardware, embedded software, middleware and systems for power, performance and reliability |
| S6.2 | Virtualization, programmability, runtimes and dynamic modification of large scale hardware and software systems |
| S6.3 | Modeling and simulation of large scale computing platforms based on application domains |
| S6.4 | Novel system architectures enabled by the inclusion of emerging technologies: non-volatile memories, high speed IO and networking, non von-Neumann and cognitive hardware |
| S6.5 | IT system architectures for big data with decentralized processing approaches in addition to centralized data centers |
| S6.6 | Processor, accelerator and system design for real-time interactive, mission-critical and deep learning and long-term analytics |

| S7 | Algorithms for Sensing and Real Time Embedded Control/Communication Applications |
|---|---|
| *This category pertains to signal processing algorithms/data aggregation and their efficient implementation with constraints of low power, low cost, improved performance depending on the application need. More specifically this area involves applications in smart grid – communications/metering/structural health monitoring/gas or water flow sensing, autonomous vehicles, medical applications including imaging/body area networks, and motor control. The smart grid area includes electric power grid, water, gas utilities, and monitoring of structures like buildings, bridges, wind mills etc. The body area networks include but are not limited to wireless ambulatory medical monitoring or personal health/fitness, home and industrial monitoring. The autonomous vehicles include but are not limited to automobiles, robotic applications in industrial environment and other areas.* | |
| S7.1 | Smart grid applications |
| S7.2 | Body area networks, mesh networks, sensor networks |
| S7.3 | Sensors and algorithms/IP cores enabling combining purposely different forms of sensor data (e.g. stereoscopic imaging + time-of-flight phase info + imaging (mm Wave or THz)) |
| S7.4 | Multi-core algorithms, coding, and software tools for embedded ultra-low power sensing and control applications |
| S7.5 | Algorithms and architectures for energy efficient embedded vision (including 3D) and auditory sensing |
| S7.6 | Algorithms and architectures for energy efficient motor control/drive |
| S7.7 | Structural health monitoring applications for infrastructure including bridges, buildings, pipes |
| S7.8 | Algorithms for non-invasive sensing |
| **S8** | **SW/HW Co-development** |
| *SW/HW co-development includes environment and techniques to design, optimize and validate system function, performance and security across the SW/HW boundary. Hardware assistance for RTOS functions can improve interrupt response time, provide very-low-jitter compute and I/O, and reduce RTOS overhead.* | |
| S8.1 | Optimizing SW-HW interfaces |
| S8.2 | Dynamic scheduling and thread migration (e.g. real time OS, power/performance/thermal optimization) |
| S8.3 | Communications of information between OS/run time/application and the HW platform |
| S8.4 | System level validation across the SW-HW boundary |
| S8.5 | Hardware schedulers and hardware acceleration for RTOS functionality, such as for reduced interrupt response time, jitter and RTOS overhead |
| **S9** | **Security and Anti-Counterfeiting** |
| *Security concerns and counterfeiting are increasingly undermining the hard-won reputations of SoC manufacturers and their customers, and sometimes even compromising bodily safety. Formal analysis is emerging to join simulation-based techniques for detection of attack vectors and estimation of their risk levels. Proposals must involve Security/anti-counterfeiting mechanisms that have hardware-design ramifications. Example application areas may include but are not limited to authentication, smart grid, automotive, medical patient data access, medical implants, and point of sales.* | |
| S9.1 | Formal verification of security - formal proof that a given path cannot be a security leak |
| S9.2 | Security breach detection, mitigation, and recovery techniques |
| S9.3 | Methods and tools for measuring system security risk |
| S9.4 | Mitigating safety risk due to tampering with safety-critical systems |
| S9.5 | Anti-counterfeiting approaches utilizing/building on PUF and/or non-PUF techniques |