

Integrated Circuit and Systems Sciences

Research Needs: Integrated System Design

December 2011

The SRC GRC member companies are pleased to release this document that describes their research needs in the thrust of Integrated System Design. 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 ETAB priorities, the ICSS strategic planning process, workshops and joint discussions.

The ITRS cites design cost due to exploding silicon and system complexities as a key factor with the potential to limit the semiconductor roadmap. These complexities are particularly notable given that ITRS predicts that by the end of the next decade, new non-CMOS devices will be needed to augment the capabilities of the CMOS process, effectively marking the end of traditional CMOS scaling. Thus, while creating the system function, designers must contend with all underlying technology challenges and approaching physical limits. These needs categories were formulated in anticipation of these upcoming challenges and also recognizing that many design degrees of freedom are available for continuing the pace of system capability.

Design automation has led to quantum gains in design capability and productivity; design cycle times have not grown as rapidly as the sizes of design tasks. However, while CAD is necessarily related to system design, the main objectives of the 2011 needs categories are innovations in systems architectures and algorithms to execute high-level tasks within the context of technology capability. Any proposed CAD research must develop new capabilities in the context of a system design proposal and have an emphasis on electronic system level descriptions and models. Proposals with CAD tools as their main deliverable are beyond the scope of this solicitation.

A designer-driven, semi-automatic optimization, verification and synthesis flow from high-level platform descriptions into micro-architectures (RTL) and analog/RF blocks is of high importance to the member companies. Back annotation of technology and implementation specific constraints (delays, current/power, area, cross-coupling between signals, reliability, etc.) into the high-level descriptions is important to achieve adequate accuracy.

As power and energy efficiency become the main driver for system design, novel processing architectures and on-chip data storage and communication is needed to achieve system-level performance targets within global power and resources constraints. Scalable, modular SoC architectures based on heterogeneous cores and customized accelerators to achieve performance, power and cost effectiveness has potential. Some examples of research required to enable such SoC designs include early design platform with software and analog/RF components exploration (virtual platforms), data/application-driven synthesis of (reconfigurable) accelerators, and heterogeneous, coherent and globally optimized communication fabrics (for network-on-chip and memory).

One of the major concerns in integrated system design is power management and minimization. Traditional performance-power trade-off techniques are inadequate to address future technology leakage power in a reliable manner. Novel global optimization techniques will be needed at every design step, top to bottom, to achieve the targeted solution.

Variability of technological parameters is a growing concern, as is reliability. Major emphasis is needed in the area of design robustness to address these issues. For instance, these needs include approaches for analysis, reduction, and avoidance of coupling effects; design methods for achieving acceptable performance from systems in which some components are not functional, or which may fail either transiently or permanently; quantification of both power/area impact of robust design, and of the certainty of fault detection, containment, and recovery.

As the emphasis on system optimization research has increased, the importance of software as a part of the overall system has also increased. The optimization of embedded systems often includes operating system and application software as part of the overall power management, reliability, and quality-of-service strategy. In addition, multi-core systems performance may be limited by the underlying architecture and the software programming concurrency model. Thus, research into software as it relates to system optimization is included in this needs document.

White Papers for all the categories below will be considered for funding. Please note that while the order of the categories does not imply priority level, the sub-categories in bold were highlighted by member companies for this 2011 solicitation as being particularly valuable to the industry.

2011 Integrated System Design Needs Categories

S1	Early Design Space Exploration at System Level
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 trade-off in early design steps with reasonable accuracy
S1.3	Workload-adaptive/reconfigurable processors and/or accelerators
S1.4	Workload-adaptive/reconfigurable SoC, NoC, and I/O architectures
S1.5	High level exploration for systems utilizing 3-D integration and heterogeneous process technologies
S1.6	Exploration of reconfigurable accelerators
S1.7	Reconfigurable analog circuits and/or digital and associated high level modeling, design and synthesis flows
S1.8	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
S1.9	Early-architectural trade-off analysis of fault-detection, -correction, -containment, and –repair architectures
S2	System Level Architecture and Implementation
S2.1	System level validations of performance and robustness including both hardware and software
S2.2	Analog/RF/mixed-signal sub-systems, platform based design, 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	Algorithmic level design including hardware, software and memory trade-offs
S2.7	Latency insensitive and elastic designs
S2.8	System design and partitioning for improved verification of architectural correctness in the presence of aggressive power management techniques
S2.9	Architectural options for SoCs integrating reconfigurable hardware and interconnect at various levels of granularity
S3	Design Robustness
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	Non-stop 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	Deep sub-micron aware microarchitectures, accounting for noise, power, timing, interconnects, etc.
S4	System Power Optimization
S4.1	Efficient power management to enable high performance, highly integrated, multi-function SoC
S4.2	Novel cache architectures: Efficient and reliable leakage control for large arrays
S4.3	Efficient control of active power and elimination of hot spots
S4.4	Reducing di/dt through microarchitectural design and software techniques
S4.5	Novel scheduling of resources, balancing workload, optimizing microcode
S4.6	Dynamic and/or static power management and optimization for micro watt to high power systems
S4.7	SoC/SiP/3D microarchitectures that support multiple power grids, multiple VRM, multiple system power states/clocks
S4.8	Methods for power reduction in reconfigurable-hardware systems
S5	Multi-core SoC Design
S5.1	Design methodology for many core of homogeneous and/or heterogeneous cores for micro watt to high power systems
S5.2	Multi-core or many core architectures (homogeneous and/or heterogeneous) for micro watt to high power systems
S5.3	Design and validation of on-chip communication fabrics and protocols
S5.4	SiP/3D chip-to-chip interfaces

2011 Integrated System Design Needs Categories, continued

S6	Algorithms for High Performance Computing Applications
S6.1	Multi-core algorithms, coding, software tools and applications
S6.2	Virtualization for hardware and software systems
S6.3	Novel memory sub systems and architectures enabled by the inclusion of emerging memory technologies such as MRAM, PRAM, RRAM, in addition to or replacements of traditional memory such as SRAM and e-DRAM
S6.4	High-performance processor microarchitectures
S6.5	Novel system architectures enabled by the inclusion of emerging technologies
S.7	Algorithms for Sensing and Real Time Embedded Control/Communication Applications
S7.1	Smart grid applications
S7.2	Body area networks, mesh networks, sensor networks
S7.3	Algorithms for medical imaging applications including optical, ultra-sound
S7.4	Sensors and algorithms/IP cores enabling combining purposely different forms of sensor data (e.g. stereoscopic imaging + time-of-flight phase info + THz imaging)
S7.5	Multi-core algorithms, coding, and software tools for embedded ultra low power sensing and control applications
S7.6	Algorithms and architectures for energy efficient embedded vision and auditory sensing
S7.7	Algorithms and architectures for energy efficient motor control/drive
S8	SW/HW Co-design and Co-optimization
S8.1	Optimizing SW-HW interfaces
S8.2	Dynamic scheduling and thread migration
S8.3	Communications of information between OS/run time/application and the HW platform
S8.4	System level validation across SW-HW boundary
S8.5	Quality assurance and validation of firmware