

The Multi-Scale Systems Center (MuSyC)

An FCRP Multi-University Research Center on Multi-Scale Systems

Center Director: Jan M. Rabaey, University of California, Berkeley

Lead Educational Organization: University of California, Berkeley

Other Participating Academic Institutions: California Institute of Technology; North Carolina State University; Rice University; University of Maryland; Stanford University; University of California, San Diego; University of Illinois, Urbana-Champaign; University of Michigan; University of Southern California.

Executive Summary

The information-technology platform is being radically transformed as we speak. A new generation of applications is emerging that are destined to run in distributed form on a platform that meshes high-performance **compute clusters** with broad classes of **mobiles**, surrounded in turn by even larger **swarms of sensors**. The broad majority of these new applications can be classified as **distributed sense and control systems** that go substantially beyond the “compute” or “communicate” functions traditionally associated with information-technology. They have the potential to radically influence how we deal with a broad range of crucial problems facing our society today: power delivery in emerging micro-grids, emergency response to natural and man-made disasters, wireless healthcare with individualized monitoring, national infrastructural monitoring and adaptation, detection of anomalous events and behaviors in physical or cyberspace for security, or real-time situational awareness on the battlefield, etc. In fact, the opportunities are limited only by our imagination.

These applications often engage all platform components simultaneously **in a closed loop** – data gathered in the sensory swarm may migrate via a hierarchy of feature extraction functions running on mobiles to sophisticated control services executed on the cloud of large-scale servers. They also span **many scales** – they combine the very large with the very small, and the very fast with the very slow, and consist of complex hierarchies of heterogeneous functionalities integrated on a broad range of technologies from the macro- to the nano-scale – leading to ever-more complex systems. Complexity arises from the integration a large number of strongly interacting heterogeneous components within tight constraints on energy, reliability and availability.

Even as Moore’s law has been the driving force in industrial growth and achieving new capabilities, we contend that future growth crucially depends upon **the capability to rapidly design, validate, deploy and manage complex distributed systems**. There are fundamental challenges to the realization and operation of distributed multi-scale systems (MSS) that must be addressed if their enormous capabilities are to be unleashed. A few examples suffice to highlight the concerns: how do we specify guaranteed detection of important events – whether in a city, hospital or battlefield – and suitable response when such data must traverse through a complicated hierarchy of geographically distributed domains? How to ensure that a network of implanted devices administering crucial healthcare will operate correctly over a long time if its energy supply is sparse and variable? How do we guarantee that a complex airborne platform will perform as expected under extreme circumstances? **Failure to address these challenges in a fundamental and comprehensive way will most certainly delay if not prohibit the widespread**

adoption of these exciting and high-impact technologies. To directly address these compelling needs, a multi-university research center on Multi-Scale Systems Design (MuSyC) has been created.

MuSyC addresses the conception, implementation, validation and management of distributed information-technology systems that have important features at multiple scales — which could be *spatial, temporal, functional, or technological*. Linking between scales and taming complexity are the main challenges to be addressed.

To make measurable progress in a distributed center setting requires a clearly identified set of common goals in combination with a focus on some well-defined and relevant application domains. We have selected **smart energy management** as a common tread throughout all the center activities. Never before – since the advent of thermodynamics in late 18th century – has the relationship between work, energy, and heat come under such a scrutiny. Our supply of energy – at all scales – is limited, and failure to address the challenge adequately comes at a tremendous cost to society. A coordinated attack on understanding and optimizing energy use represents both a unique opportunity and formidable challenge for distributed information-technology systems (DITS). Intelligent “energy management, distribution, and utilization” on one hand is one of the most prominent **societal applications** of the distributed IT platform of the future. This applies to nation-wide and metropolitan power grids, data centers, green buildings, complex airborne platforms, and traffic management. It equally holds in the **defense arena**, where intelligent energy management and distribution is of essence in battlefield logistics, broad-scale surveillance, fleets of autonomous vehicles, and personalized information access and dissemination. On the other hand, while distributed IT promises to be a central part of the solution, ironically the aggregate **energy dissipation of these IT platforms themselves is becoming prohibitive**. Today’s IT systems represent 2% of the nations carbon footprint (similar to aviation), and this is expected to double to 4% by 2020. Intelligent solutions that dynamically balance energy availability and demand under varying load conditions and performance considerations through control at all levels of the system hierarchy have the potential of improving energy efficiency of applications and IT systems alike by at least an order of magnitude. Given the broad impact, we have chosen “**energy-smart distributed information-technology systems**” as the unifying theme for the MuSyC center. This is reflected in the center organization and research agenda (selection of the research topics, application-domains being considered, and relevant metrics for optimization).

The **grand goal of the MuSyC** is to create a comprehensive and systematic solution to the distributed multi-scale system design challenge. While addressing the full portfolio of needs, we have specifically selected as **grand challenge** the development of “energy-smart” distributed systems: that is, distributed systems that are deeply aware of the balance between energy availability and demand, and adjust their behavior in response through dynamic and adaptive optimization through all scales of the design hierarchy. Common application drivers for which energy is of maximum relevancy will drive research.

The ultimate success of the Center hinges upon the development of a general design methodology for DITS that covers a comprehensive set of needs and that scales gracefully with the exponential increase in complexity we anticipate over the coming decades. Nevertheless, innovative technology solutions are often application and domain-dependent. To address this duality while ensuring cohesiveness and focus, we have organized MuSyC along a mixed horizontal/vertical format as shown in Figure 1. Overlaying it all is a **unified methodology** addressing the specification, optimization, synthesis and run-time management of **complex distributed sense-and-control systems** such that long-term reliable and

efficient operation is ensured. To explore the various dimensions of the multi-scale space, the development of specific technologies and solutions is divided over two vertical themes, called “**large-scale systems**” and “**small-scale systems**” respectively. The former addresses “**energy-intensive**” applications covering wide spatial and temporal dimensions (such as data centers). The latter focuses on “**energy-frugal**” systems (such as human-centered networks for augmented sensing). Those domains were chosen because they address opposing corners of the distributed IT platform, present different challenges, and most likely require different solutions to problems such as intelligent energy

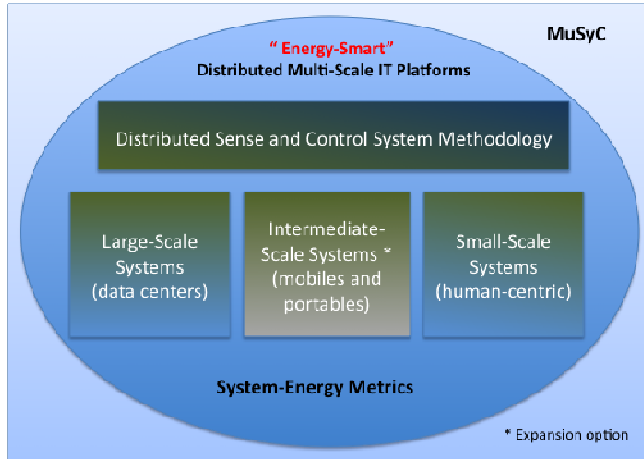


Figure 1: Organization of MuSyC with focus on “energy-smart” distributed IT systems. The Center combines a horizontal methodology theme with vertical application domain thrusts. Common system-energy metrics will be used to measure impact.

management. At the intersection of the two are “**intermediate-scale systems**” such as mobiles and portables, which are “**energy-bounded**”. We have selected this area as an option for future expansion in light of budget limitations and the need for focus in the center

We believe that division along these lines in combination with the overlaying energy focus presents a synergistic and collaborative framework for focusing the Center’s research activities. A properly coordinated effort among these forms a system-driven ecosystem, where energy – measured by some relevant system-level metric – is balanced against viability (metrics such as cost/schedule, performance guarantees, and reliability). The vision and potential outcomes of the individual themes can now be summarized as below.

A concise formulation of the vision for each of themes helps to understand how we plan to accomplish our vision through a combination of general methodology and focus on specific application spaces.

- Address the challenges in complex distributed control systems by employing **structured and formal design methodologies** that seamlessly and coherently combine the various dimensions of the multi-scale design space, and that provide the appropriate abstractions to manage the inherent complexity. Central to this approach will be system-level metrics that weigh energy efficiency versus design cost, complexity, reliability and trust.
- Realize distributed closed-loop **power-management strategies** that result in **large-scale systems** to be orders of magnitude more energy-efficient, while ensuring that mission-critical goals (such as computational throughput, latency, reliability, longevity) are met. This can be accomplished by employing a holistic multi-scale solution including all components of the system at multiple hierarchy levels. Employing structured and formal design methodologies at this scale is crucial in managing hard- and software complexity. The focus of this theme is energy-balanced data centers of the future.
- Explore the **absolute bounds of energy-efficiency and miniaturization** in “**energy-frugal**” **human-centric distributed IT systems** such as smart objects and advanced human-environment interfaces exploiting augmented senses. The essence of our solution is again a distributed management strategy, that dynamically and adaptively selects the correct operational point corresponding to the varying application needs in terms of accuracy or resolution, crossing all the layers of the design scales from system over algorithm and architecture to technology.

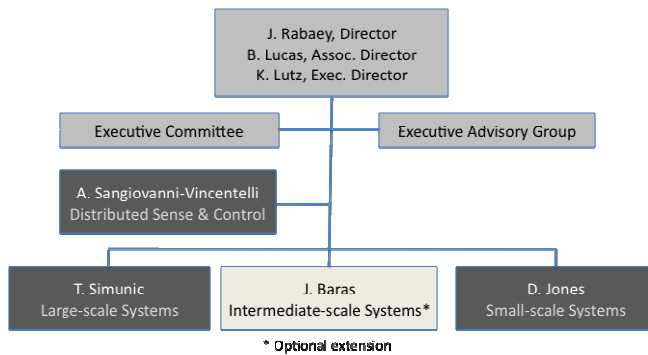


Figure 2: MuSyC Center organization

MuSyC is headquartered in Berkeley. Its management structure is shown in Figure 2. The total research team consists of **19 principal investigators (PIs) from 10 US universities**, each of them a world-leader in their respective fields. The primary responsibility of setting the research directions and priorities resides with the **executive committee**, consisting of the Center’s management (including the theme leaders), as well as representation from other FCRP Centers. In addition, a well-organized

administrative structure is set in place to facilitate open communications and result dissemination, and assure sound fiscal responsibility.

Given the broad range of challenges and the large space covered by the MuSyC research agenda, collaboration between the themes and with other FCRP centers is of absolute importance. The deliberate commonality of some of the research challenges and needs, as identified in the Executive Summary, will ensure collaboration and resource sharing between the different themes. Having researchers being part of multiple themes is another powerful mechanism. It is our intention to organize *high-visibility workshops and roundtables* on the application domains and specific challenge topics in the first year to ensure that the Center hits the ground running. In addition, as explicitly requested in the Research Announcement, MuSyC is structured such that **interactions and cooperation with the other FCRP Centers are an integral part of its approach to research**, as is detailed in the Synergy section. Special cross-center themes have been created jointly between the 6 centers on topics such as connectivity, 3D Integration, and memory.

In summary, it is our belief that the proposed center agenda and assembled team will bring to the domain of multi-scale systems the same spirit of disruptive innovation and collaborative mindshare that has become the hallmark of the FCRP program.

MuSyC Tasks

Theme 1: Distributed Sense and Control Systems (SSC) - Lead: Sangiovanni

6.1.1. Modeling for Distributed Systems [Lee, Sangiovanni]

Task 6.1.1.1: Multi-modeling [Lee, Sangiovanni]. We will develop methods for composing distinct models, including those of physical dynamics, control logic, energy, networking behavior, fault models, and computation. Hierarchically heterogeneous model composition will be supported through the development and refinement of abstract semantics and interface theories. Finally, systematic techniques supporting model transformations will be designed to convert models of one type into another.

Task 6.1.1.2: Cyberphysical Models [Lee, Sangiovanni]. We will develop methods for joint design of computational components, networking, and physical dynamics. Our approach will be to specify executable semantics for discrete events and continuous dynamics.

Task 6.1.1.3: Fault models [Lee, Sangiovanni]. We propose to incorporate into our modeling framework faults as first-class citizens to allow for fault analysis at multiple levels of abstractions. In our research we will include also models of aircraft energy generation equipment with the related fault models.

6.1.2: Verification [Lee, Murray, Sangiovanni]

Task 6.1.2.1: Robustness and Verification [Lee, Murray]. We will develop the foundations of a modern framework for testing the robustness of distributed control systems.

Task 6.1.2.2: Abstraction, Modeling, and Interface Specification [Lee, Murray, Sangiovanni]. We will develop new automatic abstraction methods for multi-scale distributed systems. Our approach will be to use a combination of algorithmic verification and statistical learning for inferring interface specifications and generating environment models automatically.

Task 6.1.2.3: Automated Diagnostics [Lee]. We will develop algorithmic techniques for detection, isolation and diagnosis of faults in the system. We will develop a hierarchical, model-based approach to diagnosis of distributed control systems.

6.1.3. Distributed Control Algorithms [Baras, Lee, Martins, Murray, Sangiovanni]

Task 6.1.3.1: Distributed Real Time Control [Lee, Murray, Sangiovanni]. We will develop methods for orchestrating distributed computer-controlled actions. Our approach will be to use concurrent models of computation with timed semantics, together with distributed and partially-ordered models of time.

Task 6.1.3.2: Distributed estimation with communication costs [Baras, Martins]. We will develop on-line algorithms for improving overall sensing performance, while balancing this improvement with the costs of communication in dynamic architectures.

Task 6.1.3.3: Taxonomy of structure versus behavior [Baras]. We will develop a taxonomy of communication architectures that enable high performance in various distributed sense and control functions. We will analyze expander graphs and other graph types. We will develop multi-scale analogs via hierarchies. We will develop dynamic self-organization algorithms for large-scale sense and control networks based on these principles.

6.1.4. Security and Trust [Baras, Martins]

Task 6.1.4.1: Coalitional Security [Baras, Martins]. We will develop models for security of coalitions that are particularly suited for multi-scale systems. The models will account for members (subsystems) that might participate in several coalitions simultaneously, as well as dynamic coalitions.

Task 6.1.4.2: Composite trust and its effects on distributed sense and control performance [Baras]. We will develop models, methods, and algorithms to investigate the effects of trust across networks –

from sensor nets to communication nets and the reverse. We will develop 'trust aware' algorithms and protocols. We will develop 'local' trust evaluation methods in distributed sense and control systems, introduce algorithms that incorporate these evaluations in weighted multi-graphs and extensions, and investigate tradeoffs between performance and trust.

Task 6.1.4.3: Physical layer authentication and compositional security [Baras]. We will develop combinations of physical layer methods that can be used to strengthen authentication, trust and establish universal compositional security. We will develop methods to guarantee component-based security for distributed sense and control systems.

Task 6.1.4.4: Policies, semi-rings and vulnerabilities [Baras]. We will extend the weighted multi-graph models to incorporate policies, and develop a methodology to analyze distributed sense and control systems with policies as distributed hybrid systems. We will extend our methods based on ordered semi-rings to model and investigate the effect of trust on various distributed inference and decision-making algorithms on graphs, and develop methods to automatically analyze and discover vulnerabilities in distributed sense and control systems due to mistrust.

Task 6.1.4.5: Energy versus security and trust tradeoffs [Baras, Martins]. We will investigate the tradeoffs between higher level of security and trust and the energy costs associated with these schemes.

6.1.5. Reliable and Robust Distributed Systems Architectures [Rabaey, Sangiovanni]

Task 6.1.5.1: Dynamic resource brokerage [Rabaey]. We propose to develop a structured software framework that allows various agents (representing resources) to discover availability, share the information with other agents through a hierarchical repository, find the optimal solution given the demand and availability, and configure the system accordingly, all of this in light of the reigning security settings. The resulting architecture will be applied to some of the application drivers of the Center – with energy considered as the most precious resource to be traded.

Task 6.1.5.2: Architectural selection [Rabaey, Sangiovanni] We will develop algorithms and models to select the architecture (e.g., type of objects, number of objects, and their locations) of a distributed system by minimizing an appropriate set of metrics that may include power consumed, monetary costs, reliability, and accuracy. Particular attention will be given to the selection of protocols, physical interconnects including wired and wireless solutions, buffers and gateways for the communication infrastructure.

6.1.6 Avionics Test Bed [Murray, Sangiovanni]

Task 6.1.6.1: Integrated power management in aircraft [Murray, Sangiovanni]. We will apply the methodologies and algorithms, developed in the other tasks, to the power generation and fly-by-wire subsystems in modern aircraft. In particular, the verification and validation problem involving multiple conditions involving distributed architectures for control of more electric aircraft will be addressed by using a combination of complexity management techniques such as abstraction, decomposition and stochastic modeling of uncertain environment behavior.

Task 6.1.6.2: Dynamic configuration of aircraft for energy-efficiency [Murray]. We will explore dynamical reconfiguration and coordination of subsystems using awareness of current and predicted environment, operations and constraints. In current systems, these tradeoffs are largely performed at design time, with sufficient redundancy to provide fault-tolerance. We plan to develop techniques for future systems that will reconfigure their operations in real-time, requiring significantly more sophisticated architectures to insure high confidence, robust operations while at the same time substantially increasing efficiency and operability.

Theme 2: Large-Scale Systems (LSS) - Lead: Simunic-Rosing

6.2.1. Software energy management [Snavelly, Draper, Sarkar]

Task 6.2.1.1: Extrapolation of future workload requirements [Snavelly, Draper]

- Provide application context and boundary conditions for how applications will exercise future systems
- Extrapolate application requirements out to 2019. Provide models and analysis for the growth in data storage, compute, network, and other resources that future applications will require.
- Contrast projected application requirements against technology trends to predict the changing system balance anticipated over the course of the next decade.

Task 6.2.1.2: Automated Modeling and Management of Energy in Managed Runtime Systems [Sarkar]

- Design automatic energy characterization methodologies for managed run-time systems in the context of Java Virtual Machine and .NET frameworks, enabling systems to construct energy models without any prior hardware knowledge.
- Study component-wise profiling of applications based on run-time systems; design adaptive mechanisms and policies for run-time systems to fine-tune for efficiency based on application profiles
- Implement interfaces between a run-time system and its host operating/virtualization system; devise mechanisms for coordinated energy management and provide automated techniques to generate optimized policies.
- Demonstrate prototypes on servers that are a part of the BlackBox test environment.

6.2.2. System management in multi-scale computing systems [Rosing, Draper, Lucas, Gupta]

Task 6.2.2.1: System level energy management [Simunic-Rosing]

- Develop software interface to sensors and actuators in data-center components. Monitor and model energy consumption across different system components while running realistic workloads. Compare the accuracy of performance and energy predictions to system measurements.
- Design novel, proactive energy and thermal management algorithms capable of exploiting heterogeneous HW/SW architectures.
- Develop distributed management policies that utilize information from individual VMs to guide the system-wide management.
- Design cross-data-center energy management and workload allocation strategies. Understand how this affects the overall building management.
- Deploy in a distributed data center container testbed connected with ultra-high speed optical links.

Task 6.2.2.2: Energy management via aggressive duty-cycling [Gupta]

- Architectural design of systems that incorporate physical or logical heterogeneous components that enable aggressive duty-cycling of systems. Design component interfaces, novel states and data exchange methods to enable usable duty-cycling across the system hierarchy.
- Algorithms for optimal online operation of duty-cycled systems that have provable bounds on availability and reliability. Design pre-wake up or polling strategies that reach identified identifiable limits on energy efficiency.
- Develop methods to capture, convey and update relevant semantic information that can be used for managing power states. Develop appropriate meta-data and meta-data handling via reflection and introspection through the system hierarchy.
- Schemes to transport workloads across diverse physical and administrative domains efficiently in dynamic environments. Virtualization methods that enable distribution of work and sleep to achieve maximal energy utilization against performance and availability constraints.

Task 6.2.2.3: Managing Resilience [Lucas]

- Devise an API for communicating an application's requirements for arithmetic precision to the computing system and an error-handling API that allows an application to reason about an error that has been detected, attempt repairs if possible, and continue if feasible.
- Explore the performance and energy tradeoff of multi-media extensions for pairing (or even TMR) to ensure correct arithmetic results versus using these same resources to maximize throughput and then checking the result.

Task 6.2.2.4: Balancing Energy and Resilience [Draper]

- Evaluate environmental event models, such as noise models, to assess their ability for relating to memory cell reliability measures for future silicon fabrication technologies.
- Develop new environmental event models as necessary and evaluate baseline SRAM performance.
- Characterize the trade-off space of temporal and spatial redundancy of resilient SRAM designs and develop a framework for resilient SRAM design.
- Assess efficacy of radiation-tolerant designs for providing resilience in the context of other environmental events.
- Design a memory system that can adapt energy and time consumed to maintain a specified bit-error rate. This should vary on a page-by-page basis, depending on the type of data being stored.

6.2.3. Infrastructure energy management [Katz, Ousterhout, Vahdat]

Task 6.2.3.1: Energy Scalable Networks [Vahdat]

- Design scheduling algorithms to account for path diversity in a highly scalable fat-tree network topology. Model and verify system scalability, latency, and memory consumption. Implement scheduling algorithm heuristics on fat-trees, balancing responsiveness with communication, memory, and computation overhead.
- Complete design of fault-tolerant, scalable, layer-2 forwarding schemes. Implement MAC address rewriting to support positional Pseudo MAC architecture. Implement a fabric manager to maintain connectivity in the face of link or switch failures.
- Instrument for energy measurements and provide energy management controls. Provide inputs and controls needed to interact with SmartGrid.
- Complete hardware and software prototype of scalable switch architecture in the BlackBox.

Task 6.2.3.2: Efficient storage with RAMCloud [Ousterhout]

- Create protocols and system software to enable low-latency access to RAMCloud storage from application servers in the same data center.
- Develop and implement algorithms that provide a high level of data durability and availability for information stored primarily in DRAM.
- Investigate how RAMCloud techniques can be applied to other memory technologies such as flash.
- Evaluate performance and energy efficiency.
- Demonstrate RAMCloud as a part of the BlackBox; release in open source.

Task 6.2.3.3: Network Architectures for Localized Electrical Energy Reduction, Generation and Sharing [Katz]

- Develop initial machine-room-scale energy monitoring infrastructure to support system-level energy measurement and modeling;
- Design and construct "SmartGrid"-compatible system components: processor, network, and storage nodes, with embedded energy storage; sensors and actuators for "SmartGrid"-compatible facility components, renewable energy sources (Wind mills and solar panels) and buffers (batteries, mechanical energy storage). Deploy and experiment with SmartGrid-compatible components.
- Design energy exchange protocols between renewable grid components and adaptive data center nodes/loads.

- Complete experiments and validate models and mechanisms for data center energy reduction, generation and sharing.

Theme 3: Small-Scale Systems (SSS) – Lead: Jones

6.3.1 Utility maximization [Jones and SSS Team]

Task 6.3.1.1: Utility maximization for microscopic sensing system design [Jones]

A general methodology for joint design optimization of all system components, parameters, and algorithms maximizing end-to-end system utility within energy and size constraints will be developed.

Task 6.3.1.2: Run-time dynamic system optimization via utility maximization [Jones]

Stochastic utility maximization will be applied to dynamic run-time optimization of scalable system parameters to maximize lifetime total expected utility within energy, bandwidth, and other conservable and non-conservable system resources.

Task 6.3.1.3: Utility metrics for microscopic sensing system applications [Jones, SSS Team]

New classes of utility metrics relevant to microscopic sensing systems will be developed: e.g. utility-weighted mean-squared error, event detection probability, and response-time-weighted metrics.

6.3.2 Attention-optimized multi-scale systems [Carmena, Jones, Baras, Martins, Murray]

Task 6.3.2.1: Stochastic feedback control methods for multiscale systems [Carmena]

- Low-cost dynamic system-adjustment algorithms based on feedback and inhibition
- Optimal feedback delivery paradigms based on stochastic control and feedback information theory

Task 6.3.2.2: Biologically-inspired attentional-adaptive schemes [Murray, Martins, Baras]

We will develop efficient adaptive mechanisms to control the attention and selection of sensors.

Task 6.3.2.3: Real-Time information flow management [Murray]

We will develop methods of capturing the utility of information and use it via feedback to improve performance, taking into account the global topology of the information flow, and retransmissions.

Task 6.3.2.4: Hugely scalable adjustable-attention signal-processing algorithms [Carmena, Jones]

- Multiscale hierarchical distributed detection algorithms scalable across several orders of magnitude in computation/power-consumption
- Signal detection and -sorting methods for brain-machine interface applications scalable across several orders of magnitude in computation/power-consumption
- Attention-adaptive joint sensing and processing strategies for energy/performance management

6.3.3 Hugely scalable platforms for microscopic systems [Rabaey, Blaauw, Franzon]

Task 6.3.3.1: Platform strategy for ULE microscopic systems [Rabaey]

We will develop a re-usable and modular platform strategy for ultra-low energy (ULE) microscopic systems. This would include libraries of scalable components, and a composition and integration methodology.

Task 6.3.3.2: Integrated 3D packaging for microscopic systems [Franzon]

- Energy-efficiency-optimized 3D system integration and packaging for heterogeneous integration
- Ultra-dense, ultra-low-power cross-talk-minimizing packaging approaches
- Jointly optimized 3D packaging for energy-harvesting microscopic sensing systems

Task 6.3.3.3: Hugely power/performance-scalable system design [Blaauw]

- Processing systems that achieve near-optimal performance across orders-of-magnitude scaling (sub-kHz to 100's of MHz)
- Ultra-pipelined signal-processing implementations operating below 250mV for significantly improved energy efficiency

Task 6.3.3.4: Hugely-scalable ULE RF wireless links [Rabaey] in collaboration with IFC

- Ultra-low-energy pulse-based proximity communication for implantable applications
- End-to-end trade-off analysis of hugely scalable wireless link options

Task 6.3.3.5: Energy scavenging and wireless power for microscale devices [Rabaey]

- Optimal distributed RF remote-power solutions for microscale sensor nodes (over varying node sizes and communication environment)
- Complete end-to-end energy-harvesting system including energy conversion and storage, based on distributed system-level adaptive energy management strategy.

Task 6.3.3.6: Microscale distributed sensors [Rabaey]

- Explore and develop concepts of passive embedded sensor arrays
- Distributed RF interrogation technology for addressing and sensing from distributed microsphere array
- Ultra-efficient RF array design and power management for passive distributed microsensor arrays

6.3.4 Multi-Scale Small-Scale Sensing System Demonstrators [Carmena, Blaauw, Carmena, Franzon, Jones, Rabaey, *In collaboration with IFC, C2S2 and GSRC*]

Task 6.3.4.1: Microscopic system platform demonstrator [Carmena, Jones, Martins, Rabaey]

The following elements will be included in this multi-scale BMI system integration:

- A task-specific, dynamic utility metric that adjusts to the varying accuracy, precision, and latency requirements of a deployed BMI
- A stochastic adaptive feedback control algorithm that dynamically optimizes system performance within bandwidth and energy constraints and adjusts itself as a result of learning and adaptation
- Personal-area network management with dynamic attentional adaptation
- An end-to-end experimental BMI system employing a variety of sensory inputs at different resolutions controlling diverse actuators (prosthetics and micro-stimulators)

Task 6.3.4.2: Enhanced human-centric microscopic platform demonstrator [Rabaey, Blaauw, Franzon, Jones] in collaboration with GSRC (Shanbhag)

We will develop an ultra-low-energy *brain-machine system* instance that monitors a distributed array of microsphere neural firing sensors, performs scalable local processing at the monitoring microscopic implant, and transmits an optimized information stream to an array of interrogators. A microscopic prototype system will be integrated combining both fabricated and off-the shelf components (as resulting from the previous tasks). A joint optimization over the complete system space will be performed, guaranteeing the absolute lowest possible energy consumption in correspondence with demanded functionality. System elements include hugely performance-scalable processor, stochastic scalable co-processor (in collaboration with GSRC), passive micro-sensors and array-based RF interrogation, hugely scalable RF communication link, remote powering system, and system-optimized packaging solution.