

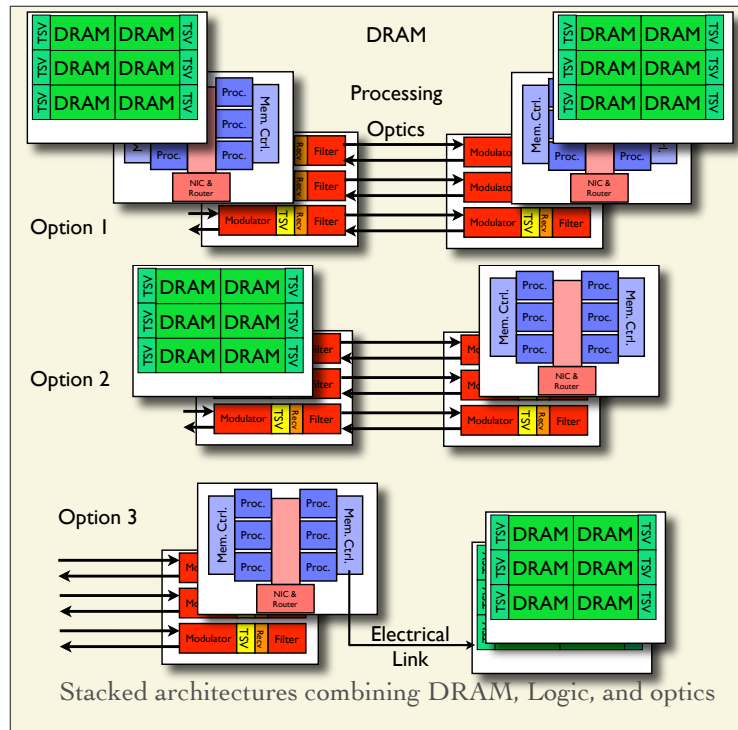
DATA MOVEMENT DOMINATES



BERKELEY LAB

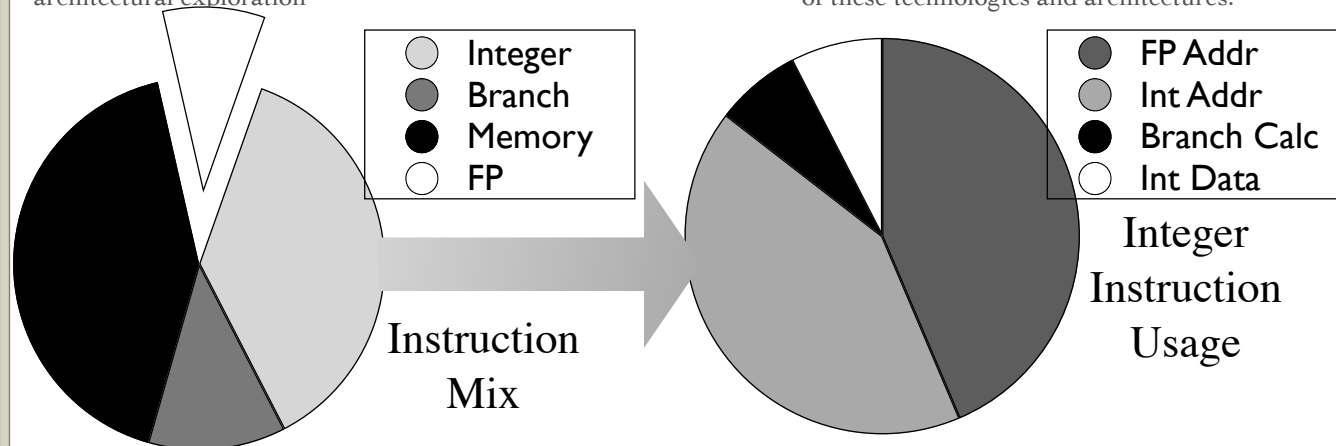


Energy is the fundamental barrier to exascale computing, and is dominated by the cost of moving data, not computation. Further, data movement, not computation, dominates the performance of real applications in HPC environments. This project will address the problems of data movement by examining three critical technologies: 3D integration, optical chip-to-chip communication and hardware support for logic operations in the memory system.



APPROACH

Simulation of the proposed systems will be accomplished by merging and improving several existing simulation models: the PhoenixSim optical interconnect simulator; the DRAMsim advanced memory simulator; and the Structural Simulation Toolkit (SST), which will provide processor and I/O models as well as a parallel simulation and power analysis infrastructure. This unified simulation infrastructure will provide accurate physical layer device models as well as more abstract designs for architectural exploration



Memory instructions and address computation dominate HPC applications

3D INTEGRATION


This project will evaluate potential exascale designs with enabling optical and 3D packaging technologies. The use of 3D integration enables high bandwidth, low-power communication over short distances between heterogeneously fabricated devices. This is critical because logic devices (such as processors) are fabricated on different processes than DRAM or optical devices. 3D integration provides a path for fabricating optics and DRAM in a simpler processes and binding it them logic devices for high-speed low power communication.

ARCHITECTURE

We plan to study the architectural innovation enabled by these technologies. In particular, remote memory accesses become much less expensive. Further, there are opportunities to study message-driven computation, where rather than moving or copying data between nodes, work is moved to the node holding the data. This has significant opportunity to consume less system bandwidth and save power.

GOALS


The DMD project aims to produce results that will allow DOE to select technology and architecture investments for an exascale system before the end of the decade. This approach is vendor agnostic, facilitating the broad adoption of these technologies and architectures.



MEMORY ARCHITECTURE


Close connection of logic and memory has long been a goal of computer architecture. Low latency, high bandwidth connections between processing elements and main memory storage would eliminate the von Neumann bottleneck. Unfortunately, attempts to integrate memory and logic onto the same die have met with limited success due to differences in the fabrication processes for DRAM and high performance logic. Hybrid approaches such as eDRAM typically sacrifice memory density and have a high cost. The separation between logic and memory is the cause of the “memory wall” and is the dominant factor in node-level performance.

The use of 3D stacking will be the most fundamental change to main memory systems since the invention of DRAM. Its most important feature is the ability to integrate dense DRAM memory with high performance CMOS logic parts in the same package. By connecting logic and memory together in close proximity it is possible to move processing, data handling, and other tasks closer to the memory, reducing latency and power. Quantifying this savings will be a major focus of this project.



PHOTONIC COMMUNICATION

Chip-scale nano-photonics is poised to uniquely address challenges in high performance computing due to its high IO bandwidth density and distance-independent signaling. However, important system-level questions must be answered before making use of this game-changing technology, such as how physical characteristics and constraints affect the architectural design and performance. Our combined experience with photonic device test and modeling along with our continued system-level simulation efforts will lead us towards realistic and feasible designs for enabling significant power and performance benefits with photonics for exascale computing.



COLLABORATION

The Data Movement Dominates project is actively seeking feedback and collaboration from the community. We are especially eager to work with application groups on understanding the characteristics and requirements of their codes. We are building a detailed modeling infrastructure, and plan to share this infrastructure.

Our goal is to provide useful feedback to the DoE and vendors, to make technology investment choices for the 2015 timeframe.

PROJECT CONTACTS

Arun Rodrigues PI Sandia	afrodri@sandia.gov (505) 284-6090
Richard Murphy Co-PI Sandia	rcmurph@sandia.gov (505) 844-7122
Paul Hargrove Co-PI LBNL	PHHargrove@lbl.gov (510) 495-2352
John Shalf LBNL	jshalf@lbl.gov (510)-486-4508
Keren Bergman Co-PI Columbia	bergman@ee.columbia.edu (212) 854-2280
Bruce Jacob Co-PI Maryland	blj@umd.edu (301) 405-0432