SANDIA REPORT
SAND2013-8594
Unlimited Release
Printed October 2013

Beyond Moore Computing Research Challenge Workshop Report

Mark Huey
John B. Aidun

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: http://www.osti.gov/bridge

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd.
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online order: http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online

SAND2013-8594
Unlimited Release
Printed October 2013

Beyond Moore Computing Research Challenge Workshop Report

Mark Huey
Perspectives, Inc.
9505 Avenida del Oso NE
Albuquerque, NM 87111

John B. Aidun
Advanced Device Technologies
Sandia National Laboratories
P.O. Box 5800
Albuquerque, New Mexico 87185-MS-1322

Abstract

We summarize the presentations and break out session discussions from the in-house workshop that was held on 11 July 2013 to acquaint a wider group of Sandians with the Beyond Moore Computing research challenge.

CONTENTS

NOMENCLATURE

DOE	Department of Energy
SNL	Sandia National Laboratories

1. INTRODUCTION

The aim of the Beyond Moore Computing Research Challenge (BMC) is to identify the devices, architectures, and applications that will succeed CMOS transistors and the microprocessor to provide high energy efficiency while supporting several generations of performance increase that comes from continued scaling of device density, clock speed, and power. Our approach is to bring to bear Sandia's broad expertise in materials, microelectronics, computer science, and computation to design the architecture of future computing systems, as well as operating system software and applications codes having major national security impact. We also plan to demonstrate prototypes of key components in an effort to reduce the risk for industry to undertake manufacture of new architectures using new device technology. The size of the integrated circuit (IC) industry makes it evident that that is where the U.S. government must look to for producing future computing systems for high-performance computing (HPC). However, in addition to making it easier and less costly for industry to do this, the BMC will address the needs of national security computing in the new system designs to help industry identify ways it can support national security needs with the commodity systems they produce. Commodity systems are overwhelmingly driven by the consumer market without regard for the needs of national security computing, which often have competing or much more exacting requirements. The computing systems of interest under this research challenge include those of embedded systems for weapons and space applications, and mobile systems for robotics applications, as well as general purpose computing.

The Beyond Moore Computing Research Challenge Workshop was convened on July 11, 2013 with multiple goals: to disseminate information on the function of Research Challenges in general; to acquaint participants with the current concept for the Beyond Moore Computing Research Challenge; to gather comments and ideas from participants on strategies, gaps, and alternate directions for the BMC; and, perhaps most importantly, to make progress pulling together the technical community that will eventually propel the BMC forward.

The full agenda for the workshop is included at the end of this report.

Participants were assumed to have a basic familiarity with the Beyond Moore Computing Research Challenge. Individuals looking for an in-depth review of the thinking behind Research Challenges and the BMC, in particular, are encouraged to review a presentation by Rob Leland from an earlier town-hall meeting (video stream available).

In brief, BMC is an important investment for Sandia and the nation because:

1) The technical advances that have enabled computing to improve according to Moore's Law are heading toward essentially fundamental limitations. There have been technical challenges to improving performance in the past, but moving past the challenges has not required radical changes to architecture or platform. In this case, it would appear that such changes are likely.
2) National security applications have traditionally driven computing requirements and hardware architectures. However, the ability of the U.S. government to influence computational futures has dramatically lessened in the past two decades, even as the needs of the government persist. There is significant value to the nation in steering a

post-Moore process toward systems that are a fit with important use cases (e.g., national security applications, certain analytic problems).

3) Sandia's vertically integrated set of expertise in computing research puts it in a unique position to make an outsized contribution.

The technical approach envisioned for the BMC is intentionally flexible, organized around a system-level view that will guide selection of viable technology paths to realize various computing systems based on one or more emerging device technologies. A number of different approaches are being pursued by the broader community (potential paths as considered by the ITRS can be found in its roadmapping materials, see here); SNL may be able to make key contributions to one or more of these. Above the device level, contributions are possible in architecture and system design, including understanding the increasingly strong interplay between architecture and applications codes. A path that provides headroom in hardware but cuts off algorithmic advances is not a desirable outcome. The areas where the BMC invests may include these or other areas.

Funding specific to Research Challenges is currently relatively limited, so BMC needs to be lean and efficient. That said, SNL is already funding significant work in this area: FY13 projects focused on Beyond Moore Computing total approximately $10M (including a number of Research Foundation LDRDs as well as the Extreme-scale Computing Grand Challenge - XGC).

Looking forward, relevant funding vehicles include:

- FY14 CTO late-start reserves, which are tentatively allocated for Research Challenges (with 1-2 potential additional 1- or 2-year projects per Research Challenge)
- Existing LDRD Investment Areas, especially CIS, Materials, NTM.
- The FY15 LDRD call, which will include funding for RC-related projects. Workshop participants are encouraged to contribute ideas for the call language (call language will be completed by December).
- Exploratory Express funding (or other similar low-dollar / fast turn-around funding pools), which may be a good fit for some BMC-related efforts.

The Research Challenges are an important part of SNL's ST&E strategy and, as the currently, most mature Research Challenge, BMC has a chance to help define what is possible and how to coordinate the combined activities of a large set of talented participants representing a number of Centers. The goal of the team shepherding the BMC is to build and focus the community that will meet the challenge of helping define, design, and demonstrate key components for a Beyond Moore computing platform.

2. Charge and Summary Response

Participants were asked to consider a number of questions regarding the Research Challenge:

- Do you like the basic direction?
- Do any major structural elements need to be added?

- Are any important technical approaches or themes missing?
- Which (if any) additional communities within the lab should be engaged in the BMC?
- How best to engage the broader community?

Workshop attendees did not, of course, speak with a single voice on these or other issues, and several questions were largely deferred for future consideration. The group was, collectively, very positively disposed toward the basic notion and direction of the Research Challenge, as reflected by the strong (80+) person attendance and the many conversations that continued well after the official close of the session. Feedback in the following sections of this workshop report suggest that participants' primary concerns relate to resourcing (especially the basic challenge of influencing a very large industry with a very small seed investment) and balance (participant questions and breakout group comments included a sizable number relating to the need to balance device hardware with other historically important aspects – architecture / software / algorithms / application codes).
Identified next steps include follow-on workshops (e.g., "Beyond Moore using two-state logic," looking into the physical devices, algorithms, and architectures that would support two-state logic; or "Beyond Moore using neuromorphic logic, "looking at similar points with respect to the latter topic). The aim is to start a discussion on the integrated aspects of a system; the mixture of device physics, architecture, and algorithms / applications.

3. Questions Raised by Participants

- *What is the time horizon for the Research Challenge?* Research Challenges are envisioned to be 5- to 10-year long endeavors composed of a number of shorter-duration projects.
- *Should certain additional areas be considered in scope?* There are strong arguments for inclusion of a number of areas that are not currently defined as in-scope (e.g., interconnects, mobile computing for robotics or other applications). However, there is a clear tradeoff to broadening the focus, especially given very limited resources. Even the entirety of SNL's research investment is tiny relative to the IC industry (e.g., Intel accounts for $10B+ in R&D). If BMC hopes to make a meaningful impact, the management team's conclusion is that the RC must downselect from the full problem to a subset of areas. Management is taking responsibility for making some of these downselect decisions.
- *Should interconnects and mobile computing be considered in-scope?* These areas were previously noted as out-of-scope, but will be included in-scope going forward.
- *Is a von Neumann architecture assumed?* No, the RC is not intended to be fixed to von Neumann approaches. Indeed, it is expected that gaining full advantage of new device technologies may drive designs to non-von Neumann architectures.
- *Does a BMC platform need to be a general purpose machine?* No. Traditional application drivers (national security applications, data analytics problems, fusion applications, climate, etc.) will continue to motivate the high-performance computing

community. The computing industry is responding to well-characterized demand (e.g., for mobile devices). The team shepherding the BMC is focused on using SNL's leverage to influence the outcome such that it supports a variety of constituencies, and, while it might anticipate that a mix of evolution (for instance, improved processors and memory) and radical departures (e.g., quantum or brain-inspired computing) is likely and that systems deployed in 15-25 years will be a mix of general and special purpose computing, it is very early for predictions. Success for the BMC is not predicated on a particular outcome in this respect.

- *How will the different – in some cases quite different – device classes be compared?* Cross-technology comparisons are currently lacking, not least because they are very challenging. Current BMC efforts are focused on developing useful measures for cross-technology comparisons.
- *Discussions seemed dominated by hardware. Is this the focus?* No. BMC-relevant investments showcased at the session were focused on hardware, but materials and the device layer are not intended to be the primary focus. Execution model, runtime, and application drivers are recognized as extremely important. As noted elsewhere in this report, a big advantage that Sandia has in this RC is a nuanced understanding of the interplay between applications and system design and hardware. The BMC management team fully expects to see this vertically integrated model reflected in the portfolio of funded projects.
- *How to make progress on the above-hardware elements of the stack without knowing what devices will actually be used?* It may be that there is significant room for improvement just through getting people talking and through defining various options in potential "strawman" hardware stacks. It may also be that an interim down-select (to a subset of most-credible technical paths) at some to-be-determined point midway through the endeavor will be a key bet by the BMC.
- *Does the mission community see this as a crisis? For instance, does the embedded computing community see a fall-off from Moore's Law performance gains as severely problematic?* This question (or the natural follow-on question of how the mission community can be helped to see the pertinence of the issue) is worthy of future scrutiny. In particular, there is likely a rich discussion to be pursued on the drivers and requirements for embedded systems. In terms of other mission drivers, it was pointed out by workshop participants that the DOE exascale community has already invested significant effort in detailing the case for dramatic increases in computing power.
- *What funding is available for the Research Challenge?* See the introductory discussion above.
- *Resilience is potentially a highly limiting factor for HPC (e.g., double-bit memory errors that are rare but cause huge problems, continued productive work even in light of errors). Does BMC anticipate contributions here?* Resiliency is important to key constituencies and is certainly worthy of further discussion. It was also a topic of conversation in breakouts and is mentioned below.
- *Is the technical investment path set for the Research Challenge?* The team shepherding BMC through the process thus far is extremely motivated to refine and improve the

current plan. The Workshop was an initial step towards broadening the pool of participants. The management team hopes that line technical staff, in particular, will take a strong role influencing the path forward for the Research Challenge through the follow-on workshops and LDRD proposals.

4. Breakout Commentary

The **architecture and systems breakout** keyed in on the following points:

- Where does architecture begin? Is it at the device level / gate level / atomic level? The group consensus was that architecture considerations start at the level of two devices. Defining this explicitly will be important to guiding efforts within BMC.
- The architecture community, at Sandia and more generally, looks at this next generation of computing with excitement, but also a high degree of trepidation, based on past conditioning that each new generation of, generally leakier, devices inevitably makes computer design harder. The notion of addressing computational throughput from an overall perspective for BMC, *e.g.*, not making architecture subservient to device design, is tantalizing. But given past experience, enacting this new model will involve cultural adaptation and may require working through a period of skepticism.
- There is significant work to be done in defining and framing metrics (performance, power, scaling, density, resiliency, reliability, etc.) and connecting them to particular technical decisions. A "building blocks" approach may be beneficial. In terms of power, this problem may be better represented as two different problems; one for the device and one for the interconnect /interfaces. In terms of resilience and reliability, basic work in defining metrics is needed.
- The individuals participating in the architecture breakout encourage the BMC to avoid stove-piping, especially in terms of the different layers of the system technology stack. Information flow up and down the stack is vital to viability and optimization. An alternative approach might be to weaken the rigid hierarchy of how systems are historically defined: if the next platform ends up not being general purpose or based on von Neumann, some of the historical motivation for the layered approach is arguably less persuasive.

The **application drivers and algorithm development breakout** covered points including:

- While BMC is explicitly noted as not being a hardware-centric endeavor, the discussion and current materials appear to some participants to reflect hardware-centric thinking. The group asked the community to keep in mind that history is more of a triangle, with important contributions and tensions coming from hardware, algorithms, and applications and to think further about how to make sure the Research Challenge as carried out reflects this.
- Moving to a Beyond Moore computing model will require, along with hardware advances, a coincident – and similarly scaled – investment in algorithms and applications software. Some in the algorithms community wish to emphasize more explicitly why this

is appropriate: i.e., that software and algorithmic contributions to computing performance are, over the past 30+ years, on par with, or perhaps greater than, the improvements that resulted from hardware advances. There is a sense that it is easy to discount or underestimate the degree to which moving Beyond Moore requires a large investment for application and OS software, a tendency that stands in marked contrast to (what the applications community perceives are) prevailing attitudes toward investments in hardware.

- There was a call for application or algorithm designers to advise hardware designers on valuable additions. But culturally, things have evolved differently and the different communities are not well equipped to collaborate in this manner. The applications representatives do not have the domain expertise or perhaps even the lexicon to articulate what they want out of a new device in a technically meaningful manner. A useful aspect of BMC may be working to establish translation mechanisms from hardware to applications ("given a component or architecture with different functionality, what can then be done at the application level?") and vice versa ("if you could make an architecture that could do A, we could do X within the application"). One example from this breakout group's conversation where members were already able to specify connections: resiliency, where parallelism to the transistor level could essentially obviate the potential for silent corruption (e.g., catastrophic double-bit errors).
- Maturing the dynamic of how novel hardware translates into capabilities would be of great value, but may be fundamentally challenging. For instance, designing hardware is already sufficiently challenging given a relatively small number of constraints (e.g., faster, cheaper, lower power). Adding more voices and constraints is not guaranteed to lead to better outcomes.
- BMC may represent a very strong argument for a comprehensive simulation system. This would (in theory) allow mathematically expressed constraints on a hardware device to guide exploration of design at the physics level and would help the applications community to talk about more than just what does not work.
- Co-design appears to be the path forward that addresses many of the issues identified in the algorithms and applications breakout. Indeed, a Grand Challenge-like co-design program could be an excellent mechanism to not only advance the Beyond Moore technology base, but also put SNL in position to command the co-design space.
- This breakout group also noted that co-design does not necessarily connect directly to algorithms; traditionally, the algorithms perspective is tied in indirectly via the applications. The algorithms point of view perhaps should be more explicitly connected into future co-design discussions.
- Given the many unknowns and uncertainties of how a Beyond Moore technology stack may come together, this breakout group encourages openness to serendipity.

The **device physics breakout group** noted the following:

- This breakout group called attention to the importance of manufacturability and industrialization in thinking through potential novel devices and materials. Especially if the notion is for a prototype in the relatively near term (e.g., 10 years), it is crucial to

consider what will go into enabling commercial-scale production in evaluating potential paths. A tremendous amount of time and effort has gone into silicon production infrastructure. Importantly, these sorts of investments are not necessarily easily accelerated (i.e., doubling the budget does not translate into getting to a given point in half the time).

- The group emphasized the value of the competitive tension between the various technical paths, both the set of technologies discussed at the workshop as well as the many other devices and mechanisms and classes (e.g., optical). It will be important to ensure representation from each of the various classes in making plans and decisions about how to steer BMC.

Addendum: Workshop Agenda

2:30 Welcome, Introductions, Workshop Goals (John Aidun)

2:40 Overview of Motivation and Goals for the BMC (Rob Leland)

2:55 Questions, Attendee Input on Approaches to BMC Goals

3:10 Potential Funding for R&D in Support of the BMC (John Aidun)

3:15 Approach to date, Initial LDRD Recap, R&D Needs (Erik DeBenedictis)

3:30 Q&A

3:40 Technical summaries of selected BMC-relevant projects

- Carbon Nanotube Transistors, Francois Leonard (8656)
- Nanowires, Alec Talin (8656)
- Piezoelectronics, Jon Ihlefeld (1816)
- Reversible Computing with Superconducting Electronics, Nancy Missert (1114)

4:10 Q&A

4:40 Breakout Sessions

- Device Physics
- Application Drivers and Algorithm Development
- Architecture and Systems

5:15 Outbriefs from Breakout Sessions

5:30 Adjourn

Addendum: Workshop Participants

The Workshop was attended by 80+ individuals, with participation from a number of different Sandia organizations.

Org	**# of Participants**
1100	9
1400	35
1700	9
1800	5
5000	3
6000	3
7900	4
8600	7
8900	6
9000	2

DISTRIBUTION

1	MS0351	Julia M. Phillips	7000 (electronic copy)
1	MS0359	Benjamin K. Cook	7911 (electronic copy)
1	MS0899	Technical Library	9536 (electronic copy)