



U.S. DEPARTMENT OF
ENERGY

Office of
Science

Neuromorphic Computing

-Beyond CMOS Emerging Computing Architecture-
Basic Research Needs

Dr. Robinson Pino

Program Manager

Advanced Scientific Computing Research (ASCR)

Robinson.Pino@science.doe.gov

April 4, 2016



BRIEFING ROOM

ISSUES

THE ADMINISTRATION

PARTICIPATE

1600 PENN



HOME · BLOG

A Nanotechnology-Inspired Grand Challenge for Future Computing

OCTOBER 20, 2015 AT 6:00 AM ET BY LLOYD WHITMAN, RANDY BRYANT, AND TOM KALIL



Summary: Today, the White House is announcing a grand challenge to develop transformational computing capabilities by combining innovations in multiple scientific disciplines.

In June, the Office of Science and Technology Policy issued a [Request for Information](#) seeking suggestions for *Nanotechnology-Inspired Grand Challenges for the Next Decade*. After considering over 100 responses, OSTP is excited to announce the following grand challenge that addresses three Administration priorities—the [National Nanotechnology Initiative](#), the [National Strategic Computing Initiative](#) (NSCI), and the [BRAIN initiative](#):

Create a new type of computer that can proactively interpret and learn from data, solve unfamiliar problems using what it has learned, and operate with the energy efficiency of the human brain.

<https://www.whitehouse.gov/blog/2015/10/15/nanotechnology-inspired-grand-challenge-future-computing>




U.S. DEPARTMENT OF
ENERGY

Office of
Science

Organizing Committee

- **Chairs**
 - Rick Stevens
 - Argonne National Lab and University of Chicago
 - Ivan K. Schuller
 - University of California, San Diego
- **DOE Contacts**
 - Robinson Pino
 - Advanced Scientific Computing Research (ASCR)
 - Michael Pechan
 - Basic Energy Sciences (BES)



http://science.energy.gov/~media/ascr/pdf/programdocuments/docs/Neuromorphic-Computing-Report_FNLBLP.pdf

Participants

Nathan Baker, *PNNL*

Jim Brase, *LLNL*

Hans Christen, *ORNL*

Mike Davies, *Intel Corporation*

Massimiliano Di Ventra, *UCSD*

Supratik Guha, *ANL*

Helen Li, *University of Pittsburg*

Wei Lu, *University of Michigan*

Robert Lucas, *University of Southern California*

Matt Marinella, *SNL*

Jeff Neaton, *LBNL*

Stuart Parkin, *Max Planck Institute
of Microstructure Physics, IBM*

Thomas Potok, *ORNL*

John Sarrao, *LANL*

Catherine Schuman, *ORNL*

Narayan Srinivasa, *HRL*

Stan Williams, *HP*

Philip Wong, *Stanford University*



Objective

- **Joint activity sponsored by ASCR and BES**
 - **Goal:** Evaluate both advanced materials and scientific computing research opportunities to support development of a new paradigm for extreme and self-reconfigurable computing architectures that go beyond Moore's Law and mimic neuro-biological computing architectures
- **The event took place during Oct 29-30, 2015**
- **Approximately 20 experts in emerging computer architectures and related materials science research topics were be invited**

What is a Neuromorphic System?

- Biological information-processing systems operate on completely different principles from those with which most engineers are familiar:
 - For many problems, in which the input data are ill-conditioned, biological solutions are many orders of magnitude more effective than digital methods
 - This advantage can be attributed principally to the use of elementary physical phenomena as computational primitives
 - Representation of information by the relative values of analog signals, rather than by the absolute values of digital signals
 - Requires adaptive techniques to mitigate the effects of component variation
 - This kind of adaptation leads naturally to systems that learn about their environment
- Large-scale adaptive analog systems are more robust to component degradation and failure than are more conventional systems, and they use far less power.
- Recently, emerging technologies such as memristive systems and their potential integration with CMOS has ignited renewed interest in the field

Carver Mead, Invited Paper, Neuromorphic Electronic Systems, Proc of IEEE, 78 (1990) 1629-1636



U.S. DEPARTMENT OF
ENERGY

Office of
Science

Neuromorphic Computing, Memristor Science

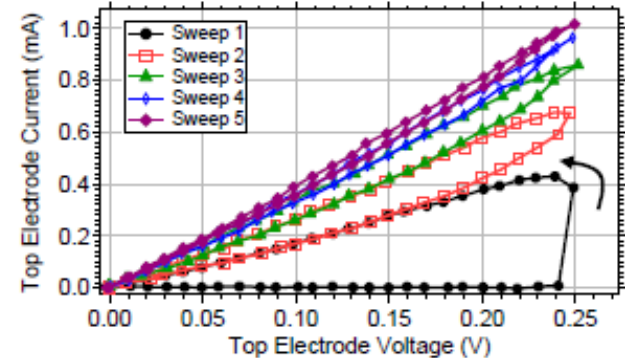
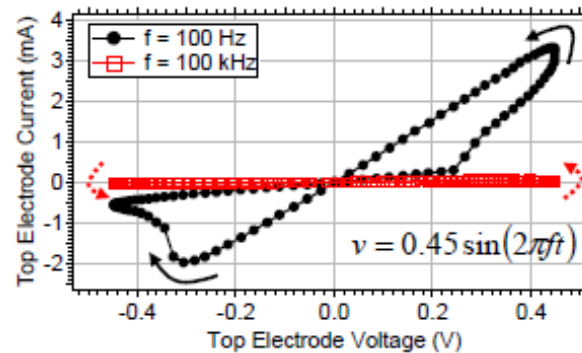

- Demonstrate integrated system beyond Von Neumann architecture

What Is The Memristor?

- Memory-Resistor
- Two-terminal passive device
- Theory postulated in 1971
- (Re-)Discovered in 2008

Significance

- Large-scale parallel computing
- High-density non-volatile memory
- Inherent radiation hardness
- Optimum power/size/weight



The memristor--the functional equivalent of a synapse--could revolutionize circuit design

L. Chua, **Memristor - The Missing Circuit Element**, *IEEE Trans on Circuits Theory*, 18 (1971) 507– 519


R. Stanley Williams, **How We Found the Missing Memristor**, *IEEE Spectrum*, Dec 2008

Dmitri B. Strukov, Gregory S. Snider, Duncan R. Stewart & R. Stanley Williams, **The missing memristor found**, *Nature* 453, 80-83

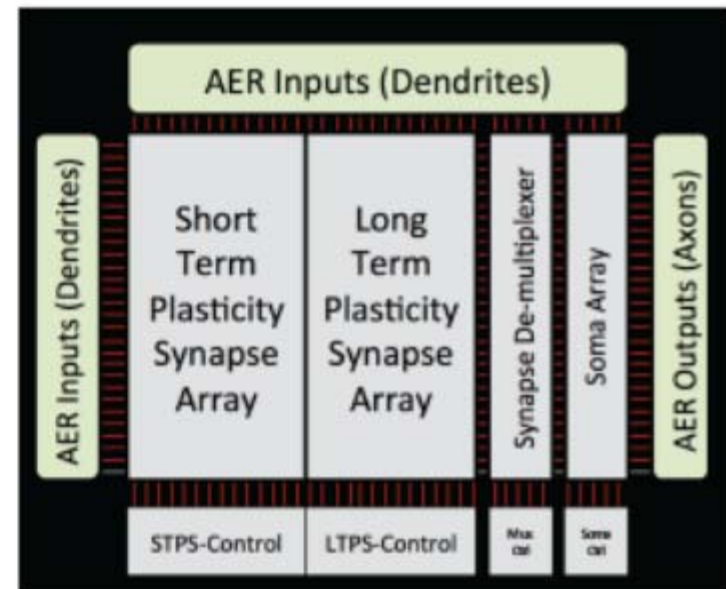
Kristy A. Campbell, **Silver Chalcogenide Based Memristor Devices**, Boise State University

Von Neumann *versus* Neuromorphic

von Neumann Architecture



Neuromorphic Architecture



AER: Address Event Representation

http://science.energy.gov/~media/ascr/pdf/programdocuments/docs/Neuromorphic-Computing-Report_FNLBLP.pdf



U.S. DEPARTMENT OF
ENERGY

Office of
Science

Advantages of Parallel Architectures

- **Perform N computations simultaneously**
 - Move from serial to parallel computation
- **Traditionally, software has been written for serial computation:**
 - Instructions are executed one after another
 - Ideally, we would like to perform many at once
- **Disadvantages**
 - Excessive communication between computing nodes will degrade performance
- **Nanoscale modeling and simulation**
 - Non-linear behavior, long and short range statistical variation
“Molecular and nanodevices will remain novelty toys for nanodevice specialists unless they possess realistic nonlinear circuit models so that future nano circuit designers can simulate their exotic designs as easily and accurately as current CMOS circuit designers,” Chua 2003.

L.O. Chua, “Nonlinear circuit foundations for nanodevices. I. The four-element torus” Proceedings of the IEEE, 91 (2003) 1830



Why Neuromorphic Computing

- **Computers have become essential in all aspects of modern life—from process controls, engineering, and science to entertainment and communications—and are omnipresent all over the globe**
 - Currently, about 5–15% of the world’s energy is spent in some form of data manipulation, transmission, or processing
- **Conventional computing devices fail in some of the most basic tasks that biological systems have mastered, such as language and vision understanding**
 - Hence the idea that taking cues from biology might lead to fundamental improvements in computational capabilities



Why Neuromorphic Computing (Cont.)

- **Well-supported predictions, based on solid scientific and engineering data, indicate that conventional approaches to computation will hit a wall and become critical in the next 10 years**
 - 1) Fundamental (atomic) limits exist beyond which devices cannot be miniaturized
 - 2) Local energy dissipation limits device packing density
 - 3) Increase and lack of foreseeable limit in overall energy consumption are becoming prohibitive
- **Novel approaches and new concepts are needed in order to achieve the goals of developing increasingly capable computers that consume decreasing amounts of power**



The Need for Enhanced Computing

- **The DOE has charted a path to Exascale computing by early in the next decade.**
 - Exascale machines will be orders of magnitude faster than the most powerful machines today.
 - Even though they will be incredibly powerful, these machines will consume between 20 and 30 megawatts of power and will not have intrinsic capabilities to learn or deal with complex and unstructured data.
- **It has become clear that the mission areas of DOE in national security, energy sciences, and fundamental science will need even more computing capabilities than what can be delivered by Exascale class systems.**
 - Some of these needed capabilities will require revolutionary approaches for data analysis and data understanding.
- **Neuromorphic computing systems are aimed at addressing these needs**



The Need for Enhanced Computing (Cont.)

- **Data science supports research involving complex and unstructured data**
 - Machine learning is applied where traditional methods are inadequate
 - Analyze data produced from climate models, *in search of complex patterns not obvious to humans*
 - Recognize features in large-scale cosmology data, *where the data volumes are too large for human inspection*
 - Predict maintenance needs for accelerator magnets, *so they can be replaced before they fail, and to search for rare events in high-energy physics experiments*
 - Biological research from searching for novel features in genomes to predicting which microbes are more likely to be in a given environment at a given time
 - Gaining traction in designing materials and predicting faults in computer systems, especially in the so-called “materials genome” initiative
- **Machine learning has affected nearly every major DOE research area**
 - Today applications run on existing parallel computers
 - Scientist are still pursuing new opportunities leveraging deep learning
 - Neuromorphic computing could significantly impact and accelerate scientific discovery



The Need for Enhanced Computing (Cont.)

- **Neuromorphic computing may even play a role in replacing existing numerical methods**
 - Where lower power functional approximations are used
 - could directly augment planned Exascale architectures
- **Important neuromorphic computing questions for the future**
 - Which areas of science are most likely to be impacted?
 - What will the requirements for may be?
- **This roundtable did not focus on an applications**
 - It is important to identify these areas and to further understand how neuromorphic hardware might address them



Current State of the Art Systems

- **IBM's TrueNorth chip announce in 2015**
 - 256 neurons and $(65-256) \times 10^3$ synapses
 - Research groups experimenting
 - No, on-chip learning
- **SpiNNaker hybrid CPU/Neuron, University of Manchester and FACETS, University of Heidelberg**
 - 180K neurons and 4×10^7 synapses per wafer
 - Over 10^4 synapses per neuron, increase biological fidelity
 - Yes, on-chip learning
- **Commons architectures arrange synapses in crossbars**
 - A challenge for such systems is fan-in/out ratios
 - Biology can have up to 2×10^4 synaptic connections per neuron
 - Current state of the art in CMOS is 256; however, it is not yet understood such why biology favors such highly dense connectivity



Potential Envisioned Research Path

- **Research, Develop, and Identify:**
 - Potential neuromorphic architectures inspired by biology
 - Materials, devices, and models to emulate nonlinear circuits
 - New tools and techniques for fabrication and integration with CMOS
 - Enable “accelerator” platforms or prototypes
 - Large scale scientific user facilities
 - Implement small systems or co-processors
 - Software for extreme scale simulation and verification
 - Libraries, programming environments, compilers, programming languages, and runtime systems specific to neuromorphic
 - Scale promising architectures into end-to-end systems from/to
 - Hundreds of neurons / thousands of synapses
 - Thousands of neurons / millions of synapses
 - Millions of neurons / billions of synapses, etc.



Main conclusions of the roundtable

- **Creating a new computational system will require developing new system architectures to accommodate all needed functionalities**
- **One or more reference architectures should be used to enable comparisons of alternative devices and materials**
- **The devices to be used in these new computational systems require the development of novel nano and meso structured materials**
 - Unlocking the properties of quantum materials based on new materials physics
 - Designing systems and materials that exhibit self- and external-healing, three-dimensional reconstruction, distributed power delivery, fault tolerance, co-location of memory and processors



Main conclusions of the roundtable (Cont.)

- **The development of a new brain-like computational system will not evolve in a single step**
 - It is important to implement well defined intermediate steps that give useful scientific and technological information
- **Successfully addressing these challenges will lead to a new class of computers and systems architectures.**
 - These new systems will exploit massive:
 - Enable the near real-time analysis of extreme-scale data
 - Learn from examples
 - Compute with the power efficiency approaching that of the brain
- **Future computing systems with these capabilities will offer considerable scientific, economic, and social benefits.**

Summary

- High pay-off research and development
- Foundational technology for:
 - ✓ Systems able to perform autonomous operations
 - ✓ Sense-making for scientific discovery
 - ✓ Accelerated and reconfigurable in situ analysis
- Plenty of opportunity to make substantial contribution in the area of information processing and understanding
- Brain-inspired architectures in software or hardware offer a path for much needed technological evolution
- It is truly a multi- and cross-disciplinary effort in materials, physics, chemistry, biology, mathematics, engineering, computer science, neuroscience, etc., to ensure success in this field



Upcoming Event: Neuromorphic Computing Workshop

Focus: Architectures, Models, and Applications

Venue: Oak Ridge National Laboratory, Oak Ridge, TN

When: June 29 - July 1, 2016

Link: <http://ornlcda.github.io/neuromorphic2016/>

POC: Dr. Thomas Potok, potokte@ornl.gov

Thank you!

Dr. Robinson Pino

Robinson.Pino@science.doe.gov

<http://science.energy.gov/ascr/>

“If we are to achieve results never before accomplished, we must expect to employ methods never before attempted.”

Sir Francis Bacon



U.S. DEPARTMENT OF
ENERGY

Office of
Science