

Comments of the
Semiconductor Industry Association (SIA)
And the
Semiconductor Research Corporation (SRC)
On the
Request for Information:
Nanotechnology-Inspired Grand Challenges for the Next Decade
80 Fed. Reg. 34713 (June 17, 2015)

Submitted by email to: NNIChallenges@nnco.nano.gov
July 16, 2015

The Semiconductor Industry Association (SIA) and the Semiconductor Research Corporation (SRC) are pleased to submit these comments jointly in response to the request for information on “Nanotechnology-Inspired Grand Challenges for the Next Decade.”

Background – The Application of Nanotechnology in the Semiconductor Industry

The Nanotechnology-Inspired Grand Challenge information request is particularly significant for the semiconductor industry. This industry has been designing and manufacturing at the nanoscale level for many years, and based on this experience has little doubt that nanotechnology holds great potential to drive future innovation in semiconductor technology.

The features of today’s leading edge transistors are less than 10 nm in dimension, allowing billions of transistors on a single integrated circuit. Manufacturing at this scale requires processes with control at the atomic level. In addition, a primary use of nanomaterials in the semiconductor industry at this time is the use of slurries containing nano-sized particles in the chemical mechanical planarization (CMP) process. The CMP process is an important step in which aqueous slurries containing abrasive nanoparticles are used to polish or smooth a silicon wafer.

Industry experts expect that the drive toward ever more powerful and energy efficient semiconductors increasingly will require the use of nanotechnology. For many years the industry has employed a silicon-based technology known as Complementary Metal Oxide Semiconductor (CMOS) technology, but many experts believe that CMOS is approaching its physical scaling limits. Design and manufacture of CMOS is a highly-automated process of hundreds of carefully controlled steps in which complex equipment applies specific chemicals in precise amounts to a thin, round slice of silicon (known as a “wafer”) to create numerous patterned layers of the integrated circuits. This process is repeated to deposit, modify, or remove materials selectively from the wafer surface. To move beyond CMOS and continue the pace of innovation in the semiconductor industry, new materials, technologies, and processes must be developed. Accordingly, the semiconductor industry undertakes a significant amount of

research on the use of nanotechnology to help advance future breakthroughs; researchers are studying nanomaterials such as nano-silica or nano-clays, nano-metal oxides, self-assembled dielectrics, silicon nanowires, carbon nanotubes, nanodots and nanowires. Yet there is much more to do.

Background on the Submitters of this Response

SIA is the trade association representing leading U.S. companies engaged in the design and manufacture of semiconductors. Semiconductors are the fundamental enabling technology of modern electronics that has transformed virtually all aspects of our economy, ranging from information technology, telecommunications, health care, transportation, energy, and national defense. The U.S. is the global leader in the semiconductor industry, and continued U.S. leadership in semiconductor technology is essential to America's continued global economic leadership. More information about SIA is available at www.semiconductors.org.

SRC is a non-profit consortium of companies having a common interest in accelerating the progress of research in semiconductor science and engineering. SRC seeks to broaden the university base for such research, and increase the supply of qualified personnel for the industry's future workforce. SRC's mission is to assess industry's needs for research, develop strategies to meet these needs, fund research at academic institutions consistent with these strategies, facilitate structured engagement of industry experts with the researchers, and proactively transfer research results to its members. SRC partners with Federal agencies that also fund basic research and have an interest in semiconductor-related science and engineering. The SRC Nanoelectronics Research Initiative (NRI) partners with NIST to fund three multi-university centers and with NSF to fund about a dozen individual projects. STARnet is a program jointly supported by MARCO (a subsidiary of SRC) and DARPA, and supports 6 multi-university centers that involve 42 universities across the country. SRC's core program, Global Research Collaboration, has partnered with NSF on a number of joint programs. SRC was awarded the National Medal of Technology, America's highest recognition for contributions to technology. More information about SRC is available at <https://www.src.org/>.

Semiconductor Industry Efforts on Addressing Grand Challenges

In March 2015, the semiconductor industry hosted a workshop, "Rebooting the IT Revolution," to consider the breakthrough innovations in semiconductors needed to continue to advance the information technology revolution. This workshop was attended by representatives of industry, academia, and government experts from NSF, NIST, DARPA, DOE, and OSTP. The workshop identified the need for robust research agendas to be developed in the following areas:

- Energy efficient sensing and computing
- Data storage
- Real-time communications ecosystem

- Multi-level and scalable security
- New manufacturing paradigm to overcome material limitations

Two of the areas for future research identified by the workshop participants overlap with the examples of potential nanotechnology-inspired grand challenges identified by OSTP for the next decade. Specifically, both the workshop and the grand challenges identified the need for the following:

1. Development of new sensors and computing and communications technology that can enable a *secure* “Internet of Things” (IoT). This area overlaps with the proposed Grand Challenge to “create devices no bigger than a grain of rice that can sense, compute, and communicate without wires or maintenance for 10 years, enabling an ‘internet of things’ revolution.”
2. More energy efficient semiconductors that can enable more energy efficient computing. This topic is related to the proposed Grand Challenge to “create computer chips that are 100 times faster yet consume less power.”

Research in both these areas, as well as the other areas addressed in the workshop report, is essential to maintaining U.S. economic and technology leadership. Other countries, such as China, are actively investing in the development of an indigenous semiconductor industry. In the face of this challenge, it is imperative that the U.S. invest in these research initiatives and not abdicate U.S. leadership in the semiconductor industry. The U.S. semiconductor industry looks forward to collaborating with federal scientific agencies and academic institutions to pursue these and other grand challenges, and we urge federal agencies to consider funding for these research areas in its budget preparations.

I. Grand Challenge – Nano-devices to Enable the IoT

Enabling the IoT has the potential to have an immense positive impact on society. Addressing this grand challenge will contribute to maintaining U.S. technology and economic leadership, while also addressing significant societal challenges. The following are brief responses to the questions posed in the RFI with regard to the Grand Challenge to:

Create devices no bigger than a grain of rice that can sense, compute, and communicate without wires or maintenance for 10 years, enabling an “internet of things” revolution.

1. *What is the audacious yet achievable goal proposed?*

The semiconductor industry proposes to achieve low cost, high performance, low energy semiconductors and sensors that can drive the IoT. To attain the promise of IoT fully, new devices for computation and memory must be demonstrated and developed – devices capable of operating at power supply voltages much lower than will be possible

with foreseeable, incremental advances in today's transistors. Ongoing research shows the feasibility of such devices, but does not yet address the many challenges of integration into systems such as the distributed sensor-processor nodes of the IoT. Another key component to enabling the IoT is the sensor/actuator systems that generate measured data and/or respond to actuation commands. A challenge that must be addressed is reducing the costs of manufacture and packaging for sensor/actuator systems. Today, the fabrication processes and the packages for sensors are typically specific-sensor-dependent and this drives cost. Moreover, remotely-located sensors will need to be securely accessible for calibration and validation checks, and if they are stand-alone, they must have mechanisms for energy collection and efficient utilization. Finally, a major aspect of this challenge is providing appropriate levels of security, assurance, and privacy at levels from the edge nodes, through gateways and to the cloud.

2. *Why is it important for the Federal government and others to invest in solving this challenge?*

The IoT has the potential to transform virtually every aspect of the economy and government, including health care, transportation, energy production, the environment, education, and national security. Taken as a whole, these technologies have the potential to deliver essential government services and perform critical functions in a less expensive, more efficient manner. According to one recent report, as smart phones, networked sensors, smart grids and other ICT devices become faster, cheaper and more available globally, they have the potential to deliver profound environmental, economic and social benefits, including a 20 percent reduction of global carbon emissions by 2030, over \$11 trillion dollars in new economic benefits, the ability to extend e-healthcare to an additional 1.6 billion more people worldwide, and an estimated 30% increase in agriculture yields. See Global e-Sustainability Initiative, SMARTer 2030 (June 2015).

3. *What would success look like? How would you know the challenge has been met? For the examples provided, are the proposed end points appropriate and ambitious yet achievable?*

Success would be the creation of cost and energy-efficient sensor nodes that can be remotely-located to enable the secure collection and transfer of data in support of the IoT. The proposed Grand Challenge, with the additional need for energy efficient devices to be "secure and trustworthy", (i.e. to create such sensor nodes no bigger than a grain of rice that can sense, compute, and communicate without wires or maintenance for 10 years, enabling a *secure and trustworthy* 'internet of things' revolution), is extremely ambitious. In order to be able to function in 10 years, significant advances in energy harvesting are required, no matter how low the energy requirements can be made. Though difficult, the goal is feasible and would accelerate progress in the key areas of energy efficient computing and communications, and energy harvesting and management.

4. *What would be potential nanotechnology solutions to the challenge and what intermediate steps and activities are necessary to develop those solutions?*

In the case of small-scale sensor nodes that are not connected to external power, the available volume for on-board energy storage is very limited. Enabling the IoT will require large advances in the energy efficiency and energy harvesting capacity of nanoscale devices for sensing and computing. An individual sensor node should include components for transduction, computation, memory, communication, and power supply. These nodes may need to act autonomously, be programmable via the network to increase flexibility, and support communication with other network nodes—all while maintaining security. In many applications, the nodes will be embedded in the environment, should be very small and should not interfere with the things that are sensed. The smallest node that has been reported is $\sim 10 \text{ mm}^3$, and the challenging target of demonstrating a fully operational node 1 mm^3 in size has yet to be achieved.

5. *What potential metrics and milestones could be used to measure intermediate progress toward solving the challenge?*

The milestone for system performance and reliability would be performance of a defined set of sensing, computing, and data storage tasks for an indefinitely long time period using only power scavenged from the local environment. Specific metrics will target large reductions in power for computing and communication *without* comparable reductions in performance as measured by rates of data transfer and speed of computation. Additional metrics would focus on the efficiency with which ambient energy could be harvested, stored, and managed.

6. *Can the challenge be achieved in the next decade? If not, how long will it take?*

While it is difficult to determine the success rate for any research program, this Grand Challenge may be achievable in the next decade. The challenge must be met to continue development of an IoT that will lead to sustained economic, environmental, national security, health, and social benefits. However, it is not yet known how sensors will advance to meet the full-functionality demands required for an IoT to accomplish such benefits.

7. *Why is this challenge worth pursuing now? What recent advances, trends, or research point to this challenge being solvable in the proposed time frame?*

The IoT is in its infancy. Inventors are just at the beginning of making items “smart” and connecting devices to the Internet and smartphones. In order to truly achieve the benefits of the IoT, we will need to have ubiquitous computing, inexpensive digital storage, and constant connectivity. Many of these needs will be met by continued advances in semiconductors. Through increased research and new breakthroughs, the U.S. can achieve the economic and technological benefits of leading in this area.

8. *What opportunities are there for partnerships between the Federal government, State and regional governments, foundations, industry, and academia to support the solution of the challenge?*

The semiconductor industry has a long history of partnering with the federal government and universities in advancing research that has resulted in significant benefits for the country and society as a whole. Industry envisions a cost-sharing partnership with government and academia to support research that will address the challenge. Existing programs such as the Semiconductor Research Corporation (SRC) Nanoelectronics Research Initiative (NRI) and the Semiconductor Technology Advanced Research Network (STARnet) are already funding some of the appropriate university research, and provide good models for an expanded research effort.

9. *Why do you expect this challenge to capture the public's imagination?*

Major corporations already advertise their pursuit of the IoT. To capture the public's imagination, researchers must present a vision of positive societal impact that goes well beyond what everyone already "knows" will happen. For example, the IoT could revolutionize agriculture with sensors that detect precisely when individual plants need watering, fertilization, or harvesting. It could optimize health and wellness by alerting individuals of health-threatening conditions long before they feel ill. Such striking examples would generate significant interest and support by the public.

II. Grand Challenge – Creation of computer chips that are 100 times faster yet consume less power

The last ten years have seen a slowing of improvements in the performance of microprocessors or computer chips. The primary cause of this slowdown is the energy cost of processing and moving data. While transistors continue to shrink to nanometer scale dimensions in size, the voltages at which they operate cannot be similarly reduced. To keep power and heat generation within acceptable bounds, system designers must throttle computing speed. On fundamental grounds, this power and heat bottleneck cannot be broken unless the transistor is replaced by a new, far more power efficient device. At the same time that performance gains are slowing, computers must handle ever more data – from widely distributed sensors, from social media, from exponentially growing genetic databases, and much more. Today, more than half of the energy used by processors is consumed by the task of moving data, e.g. between memory and logic. The exponential rate of growth in the volume of data, as well as the changing nature of data types and workloads, cannot be handled effectively or efficiently by the same fundamental computing architecture that has been in place for over 50 years.

In sum, improvements in the performance and energy-efficiency of computing systems are slowing and will soon cease to drive the ongoing revolution in information technology. But fundamental advances in computing technology – new nanoscale

devices and circuits for computation and communication and new system architectures that optimally exploit the characteristics of the devices – can maintain the vitality and deliver the benefits that are envisioned and benefits not yet imagined. The Grand Challenge that will enable this revolution is:

Computer chips that are 100 times faster and use less power.

The following are brief responses to the questions posed in the RFI with regard to this grand challenge:

1. *What is the audacious yet achievable goal proposed?*

The semiconductor industry proposes to achieve low cost, high performance, low energy devices for computation that will drive the IoT and all facets of the use of information technology in society. Today, computing systems, from smart phones to scientific and engineering supercomputers, are limited in their performance by energy inefficiencies that result in overheating and thermal management issues. A new computing paradigm must be achieved.

To achieve computing systems that are 100 times faster and use less power, new nanoscale electronic devices must be conceived, demonstrated and developed – devices that are far more energy efficient in their operation than today's (and tomorrow's) transistors, and capable of switching much more rapidly. Furthermore, these new devices ultimately must be integrated into a system architecture that optimally exploits the characteristics of the new devices. Ideally, the system should:

- be scalable from mobile platforms to data centers,
- perform programmed algorithmic calculations with orders of magnitude more energy efficiency at two orders of magnitude higher speeds than today, and
- integrate a cognitive learning and predictive capability onto the same physical technology platform.

2. *Why is it important for the Federal government and others to invest in solving this challenge?*

Discovering and developing new devices and associated architectures for computing requires the exploration of a vast range of scientific and engineering possibilities – a task beyond the capabilities of even the largest and most competent industrial research organizations. Thus the development of more energy efficient devices for computation should be an important priority for investment by the Federal government and others. In order to achieve the next generation of high performance computers, and to enable the full benefits envisioned for the Internet of Things (IoT), the device technology and architecture needs to be several orders of magnitude more energy efficient than best current estimates for mainstream digital semiconductor technology. Otherwise, energy consumption and heat generation will put ever more severe limits on performance.

“The essential engine that made exponential growth possible is now in considerable danger. The implications of a dramatic slowdown in how quickly computer performance is increasing – for our economy, our military, our research institutions, and our way of life – are substantial.” (Quoted from ‘The Future of Computing Performance: Game Over or Next Level?’; National Academy of Sciences, 2011)

3. *What would success look like? How would you know the challenge has been met? For the examples provided, are the proposed end points appropriate and ambitious yet achievable?*

In order for any new computing paradigm to be adopted and widely deployed by industry it has to meet several critical parameters. First and foremost, it has to address the needs of the widest range of customer applications cost-effectively. Appropriate benchmarks have been established by consensus of the industry to measure performance for the existing computing architectures, and revised benchmark metrics will be created to test and compare any new proposed platform for the broadest range of projected future application and workload environments. Likewise, reliability and security benchmarks will also be applied to determine viability and robustness. The industry itself will determine the economic viability and sustainability of any new platform before it can be adopted as ‘the new paradigm’.

According to one report, a computer system in 1978 operated at 1,400 instructions per second per watt, whereas in 2008 it operated at 40 million instructions per second per watt, an increase of 2,857,000 percent. See “A Smarter Shade of Green,” ACEEE, Report for the Technology CEO Council (2008). In the future, these efficiencies can drive significant improvements throughout the economy. According to a report commissioned by SIA, the U.S. could decrease the amount of electricity used annually by 1.2 trillion KWh – approximately 27% less electricity consumed than the reference case, and 11% less than today, even though the economy will be about 70% larger – by accelerating the adoption of semiconductor enabled technologies. ACEEE, “*Semiconductor Technologies: The Potential to Revolutionize U.S. Energy Productivity*” (May 2009).

4. *What would be potential nanotechnology solutions to the challenge and what intermediate steps and activities are necessary to develop those solutions?*

New nanoscale devices, operating on physical principles that are fundamentally different from today’s field effect transistors, will be absolutely necessary to meeting this challenge. For example, new “steep slope” (low-voltage) devices might operate at a tenth of a volt compared to roughly 1 volt for today’s transistors. All else being equal, this 10x voltage reduction would translate to a 100x reduction in the energy required to perform a computation or communicate a block of data. A greater challenge will be to demonstrate such a device that is also very fast in its operation. As another example, emerging nanomagnetic devices may enable the merging of memory and logic functions in ways that would greatly reduce the need to move data. As a third example, novel neuro-inspired, nanoscale functional elements integrated onto the same computational

platform might enable the system to learn and predict, and configure itself for optimum energy efficiency and performance of whatever computational task is at hand.

5. *What potential metrics and milestones could be used to measure intermediate progress toward solving the challenge?*

SRC's Nanoelectronics Research Initiative has developed a rigorous methodology for benchmarking the relative performance and energy efficiency of a wide range of exploratory devices and associated circuits and architectures. With this physics-based methodology, the ultimate performance potential of a device can be roughly gauged, even if that device only exists as a crude laboratory prototype or a theoretical model. It can be used to focus and guide the research toward device concepts with the potential to support the aggressive performance requirements of the grand challenge.

6. *Can the challenge be achieved in the next decade? If not, how long will it take?*

It is obviously difficult (and risky) to predict the future of any research program, much less one as aggressive as the one envisioned here, but there can be milestones put in place to determine when sufficient progress has been made to drive towards industry adoption of this new paradigm. Recent progress in post-silicon device research, neuroscience, and computational science supports new approaches and has raised the enthusiasm and confidence level of the technical community. This self-perpetuating progress will continue to accelerate with adequate support. Certainly within a decade after such a comprehensive and concerted effort is launched, it is very possible that a new paradigm will emerge that meets the goals of this Grand Challenge.

7. *Why is this challenge worth pursuing now? What recent advances, trends, or research point to this challenge being solvable in the proposed time frame?*

The silicon field effect transistor, the key building block for all of today's computing systems, is rapidly approaching fundamental limits to further improvement. Recent breakthroughs and the convergence among nanotechnology, neuroscience, and computational science point to new directions and approaches for the Grand Challenge. Better understanding, thanks in part to programs such as the BRAIN Initiative, in how the brain processes information offers novel strategies for future computing architecture. Exploratory device research, for example under the Nanoelectronics Research Initiative (NRI) supported by NSF and NIST along with a consortium of semiconductor companies, has illuminated many new possibilities for nanoscale devices and associated architectures for computing.

8. *What opportunities are there for partnerships between the Federal government, State and regional governments, foundations, industry, and academia to support the solution of the challenge?*

The semiconductor industry has a long history of partnering with the federal government and universities in advancing research that has resulted in significant benefits for the country and society as a whole. Recent discussions and planning sessions between the Semiconductor Industry Association (SIA), Semiconductor Research Corporation (SRC), and various government research funding agencies such as NSF, DARPA, etc. have provided an initial plan from which new programs can be launched and extended towards critical mass if significantly expanded resources are applied.

SRC, a nonprofit industry consortium, has ongoing partnerships with NSF, DARPA and NIST whereby government and industry jointly identify areas of research of mutual interest and together sponsor fundamental university research programs to address those.

9. *Why do you expect this challenge to capture the public's imagination?*

The systems of today have delivered tremendous business and societal benefits by automating tabulation and harnessing computational processing and programming to deliver enterprise and personal productivity. The systems of tomorrow will forever change the way we interact with computing systems to help us extend our expertise across any domain of knowledge and make complex decisions involving extraordinary volumes of fast moving data. These future systems offer the possibility of transforming our society and improving our daily lives, whether it is in health care, weather and storm prediction, transportation or finance, to name only a few areas.

There is no greater or more worthy challenge to embark on now than to create a new engine to drive the next generations of human experience, progress, security, and sustainability.