This paper received the Best Presentation Award at the 18th International ICIN Conference (ICIN 2015), 17-19 February 2015, Paris, France. It is available from the ICIN 2015 Proceedings on IEEE Xplore.

How Will Rebooting Computing Help IoT?

Bichlien Hoang
IEEE: Technical Activities
Future Directions
Piscataway, NJ, U.S.A.
b.hoang@ieee.org

Sin-Kuen Hawkins
IEEE: Technical Activities
Future Directions
Piscataway, NJ, U.S.A.

***Abstract*— The Internet of Things presents numerous challenges to achieve the vision of interconnecting "all" things, with the need for extreme miniaturization of computing devices that are energy efficient, human-to-computer interfaces that are natural, and processing of layers of intelligence while ensuring data security and privacy. This paper will describe "rebooting computing", an approach to rethinking the computer to address technical challenges of IoT and other key computing trends.**

Keywords—IoT; rebooting computing; security; human-computer interface, energy efficiency; neuromorphic computing, approximate computing; adiabatic computing

I. INTRODUCTION

The Internet of Things (IoT) promises a world where "all" things are interconnected. Some call IoT, the Internet of Everything as it creates a vision of unprecedented era of connectedness [1]. Advancements in networking and computing technologies will be one aspect to enabling IoT to become a reality, in addition to social acceptance and successful business models. This paper will describe "rebooting computing" (RC), an approach to rethinking computing technologies and approaches to address the technical challenges of IoT and other key computing trends.

II. INTERNET OF THINGS

A. Opportunities

Is IoT a buzz phrase or really an opportunity? IoT has been called the most disruptive technological revolution since the advent of the Internet [1]. The Internet has been on a steady path of development and improvement, but IoT is its first real evolution that will drastically change the way people live, learn, work, and entertain themselves [2]. One projection suggests that IoT will have over 50 billion humans and objects, or nodes, interconnected by 2020 [1]. Another projection suggests 1 trillion connected devices by 2025 [3]. A plethora of data can be gathered from multiple devices and applications connected to an always-on network, and then fused and analyzed in real-time. The benefits of IoT seem without limits. IoT may increase the efficiency of a production line, provide 24x7 monitoring of a patient, enable data collection for long-term research, and provide a whole host of new services for consumers. IoT is transforming all industries, including transportation, healthcare and manufacturing with connected cars, connected homes, wearables, and industrial internet [4].

B. Challenges

There are many technical challenges related to implementing a connected "sensorized" world. IoT may become the biggest system that mankind has ever built.

From a networking perspective, there must be measurable service quality with guaranteed connectivity for a large number of mobile objects. An always-on network with a global footprint will require a balance of core communications systems with flexible infrastructure components that can keep a network robust and resilient. However, IoT is a network of networks [2] and any weak "link" can impact its services greatly. Having standard network protocols and interfaces that are agreeable and implemented by all the major manufacturers involved would certainly help drive a unified platform in a world of disparate objects and elements in the network. Standardization would help in the management of this large heterogeneous network, but reaching consensus amongst a group of competitors will not be easy.

From a computing perspective, each of the IoT-connected devices, including the embedded sensors and processors, must be extremely miniaturized, energy efficient, and consume minimum power. These connected devices will capture data and generate data, producing oceans of data that will require a "brain" to analyze and transform the data into useful information. Pockets of personal data repositories everywhere will bring on the important issues of privacy and security. Identity management, access control, and data protection must be in place to ensure sensitive private and personal data do not fall into the wrong hands. Preventing hacking into the IoT connected devices will be a challenge when there are so many potential points of intrusion.

The human-to-computer interfaces also need to be evolved to make the interactions less complex and truly natural. The ability to capture speech, touch, and gestures would enable users to interact with computing devices just as naturally as they interact with each other. Natural language translation and processing within the user's context while leveraging past and current user actions would further improve the user's experience [5].

III. REBOOTING COMPUTING

A. What is it?

"Rebooting Computing" is a new IEEE initiative [6] launched in 2012 by the IEEE Future Directions Committee to rethink the computer from "soup to nuts" including all aspects

from device to user interface. Though the term Rebooting Computing was identified independently as an inspiration for this new initiative, this term was previously coined by IEEE Life Fellow Peter Denning[1] in 2009. Information on the IEEE Rebooting Computing Initiative can be found on the Rebooting Computing portal http://rebootingcomputing.ieee.org/.

Why the need to rethink the computer now? – because there is a general consensus that the primary technology driver for almost 5 decades, Moore's Law for scaling of integrated circuits is finally ending. Performance had been increasing exponentially every 18 months with the number of transistors on a chip doubling, drastically decreasing the size of typical computing devices (see Fig. 1). However, CMOS appears to be reaching physical limits, including size and power density, and there is presently no technology available that can take its place [7].

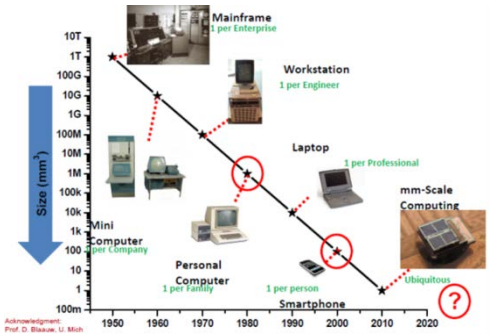

Fig. 1. Scaling of computing devices in successive generations [7].

Rebooting Computing is rethinking the entire approach to computation, from design, function, performance, and energy efficiency, starting from a clean slate. RC takes a holistic look that addresses all aspects of computing, and as such, this is an inter-society effort with 9 IEEE Societies and Councils contributing. The RC committee consists of representatives and volunteers from IEEE Circuits and Systems Society, Computer Society, Council on Electronic Design Automation, Council on Superconductivity, Electron Devices Society, Magnetics Society, Nanotechnology Council, Reliability Society, and Solid-State Circuits Society. In addition, the RC Initiative collaborates with other organizations involved in looking at future trends in the technology of computing, such as the International Technology Roadmap for Semiconductors[2] (ITRS).

The RC Committee organized a series of Rebooting Computing Summits (RCS), each with a theme, to discuss the future of computing with invited thought leaders, experts, and decision makers from government, industry, and academia. The first summit (RCS 1, December 2013) established the initial foundations with the 3 pillars of computing. The second summit (RCS 2, May 2014) and the third summit (RCS 3, October 2014) drilled down further and focused on the engines and structures of future computation.

This paper presents some of the concepts discussed in the three summits that are relevant and transferrable to supporting IoT design and development. In particular, information from the RCS 1 [8], RCS 2 [7] and RCS 3 [9] reports were used throughout this paper.

B. Vision of Future Computing – Three Pillars

A vision identified in RCS 1 [8] consisted of ubiquitous computing that is fully integrated into the lives of people at all levels of society. Future generations of smartphones and networked sensors are connected via wireless broadband links to the Internet and with large computing engines in the "Cloud". This is very much in line with the vision of IoT.

To achieve this vision and other future visions of computing, three "pillars" must be addressed as the initial foundation: Energy Efficiency, Human-Computer Interface, and Security.

1) Energy Efficiency

Energy efficiency is important so that recharging is relatively infrequent when there are potentially billions and even trillions of smart devices. Ideally, these smart devices at the periphery of the Internet are energy source aware, and consume power in proportion to the level of performance that they deliver, and only as much as their energy harvesting rate or expected charging power allows them [8].

The information and communication infrastructure (ICI) provides a core of large and powerful datacenters to the billions and possibly trillions of smart devices at the edge (see Fig. 2).

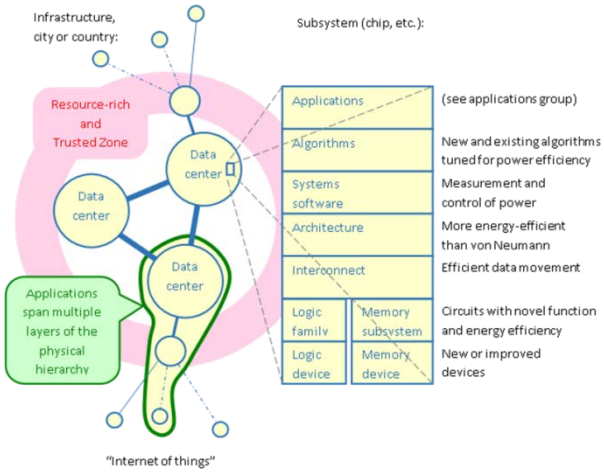

Fig. 2. A conceptual diagram of cloud computing system (data centers) surrounded by a large number of edge devices (the IoT) [8].

An edge device is connected to the core via an intermediate edge infrastructure that is a trusted zone and resource-rich. Smart devices can access the intermediate edge and the core to enhance their own functionality, resiliency, and performance. For example, a device can harvest energy from this resource-

[1] Rebooting Computing" was coined by IEEE Life Fellow Peter Denning as part of his National Science Foundation-sponsored initiative to revamp computing education.

[2] ITRS is sponsored by five leading chip manufacturing regions in the world: Europe, Japan, Korea, Taiwan, and the United States. The objective of ITRS is to ensure cost-effective advancements in the performance of the integrated circuit and advanced products and applications that employ such devices. For more information, visit http://www.itrs.net/.

rich zone and be fully replenished in a seamless manner. A higher concentration of connectivity is expected where there are larger concentrations of people and thus, a greater need for energy harvesting [8].

Figure 2 also shows that the computational infrastructure within each chip will have advances over today's CMOS technology. The technology stack from devices to applications may be redesigned for power efficiency [8].

Sensors in these edge devices may be able to leverage the environment as potential sources of energy. Environmental elements such as vibrations, light, and airflow may generate sufficient electricity to enable the sensors to be self-sustaining [2].

Some computing/processing can be shifted to the Cloud to lessen the need of power for the edge devices. Data at the edges may be sent to the Cloud for more energy efficient processing. At the same time, energy efficiency in the core needs to be addressed or it might become a major cost of operation. Hardware and specialized computational devices embedded throughout the environment must be more robust and last longer in the future [8].

2) Human-Computer Interface

The earliest computers were large mainframes with batch processing used for scientific calculations in the late 40's. The personal computer revolution started in late 60's with the perspective changing to a one-to-one paradigm. Ubiquitous computing began 20 years later as microprocessors became integrated in a variety of devices [9].

We now move to the 4th generation of computing, known as "complementary computing", which further integrates computing devices into the human environment and blurs the distinction between humans and computers. The 4G computer network brings an increased level of individual self-sufficiency where any individual can have complete information on himself/herself, his environment, and the world [9]. As such, the human-to-computer interface (HCI) should be as natural as possible. Improved natural language processing and ability to capture multisensory input will be important aspects of advanced HCI [5].

The human-to-computer interaction is similar to having a best friend. The computer doesn't simply answer questions but anticipates needs as it learns and grows from the evolving history of interaction with the human [8]. It is also contextual aware to better aid the user [5]. The computer of the future acts as an "aug-mentor", a computer mentor that augments human capabilities. The user never needs to ask for updates or think about synching devices. Information is automated and readily available to the user, and functions that are not needed or inefficient will be disabled when necessary [8].

The HCI would need to maintain a body of information for interaction with the user, such as user's preferences for various items and activities. The information is indexed and the user can ask questions of the computer or give commands. Having the user's evolving history of interaction with the computer and other data sources indexed in real time would create a higher computational load. This higher computational load might not be achievable without advances in rebooting computing technology [7].

3) Dynamic Security

If the computer of the future acts as an "aug-mentor" to the user, the concept of security and trust must go hand-in-hand. The user is trusting that the device has his or her best interests in mind and doing the right thing for the user [8]. In addition, with data moving freely between the worldwide network of devices and computers, security measures and solutions must be in place so that users can operate without fear of data diversion and eavesdropping [7].

Security cannot be an after-thought but be a key component and building block in the design and implementation process in all devices and networks. Security must be pervasive in every aspect of the components, systems, and networks. The number of intrusion points in a heterogeneous network with billions and possibly trillions of edge devices and core systems is staggering, and must be protected.

In a world where ubiquitous computing is fully integrated into the lives of people at all levels of society, systems are expected to work correctly at all times despite possible hardware faults, software errors, and malicious intrusions [8]. Details need to be given at both hardware and software levels, with solutions that are dynamic and flexible to react and respond to new threats [9]. Special attention to mitigate high consequence failures must be taken into consideration as well [8].

Establishing forward-looking security standards and practices will help establish and implement a universal secured and trusted network of systems and devices. Current security systems represent a patchwork of solutions for different kinds of systems. Authenticating identity and privilege in platforms and applications and computation should be done universally and automatically. When there are flaws and exploitations, standard protocols must be devised for automatic updating of security software in new systems as well as legacy systems to minimize further damages [9].

C. Engines of Computation

Mainstream and alternative computing technologies for future computing were discussed at RCS 2 and RCS 3. More R&D is needed for some of these technologies, but they have potential application in the development of IoT.

1) Augmenting CMOS

The performance of CMOS devices and systems have been following Moore's law (doubling in performance every year or two) for the past several decades, together with device scaling to smaller dimensions and integration to larger scales. CMOS, however, appears to be reaching physical limits, including size and power density, but there is presently no technology available that can take its place [7].

The primary focus of the semiconductor industry has been on scaling digital integrated circuits. However, Moore's Law can be pursued at the system level by the use of "orthogonal

scaling". One can focus on scaling the package and the board, instead of the chip [7]. One example of "orthogonal scaling" is when multiple memory chips are stacked (see Fig. 3).

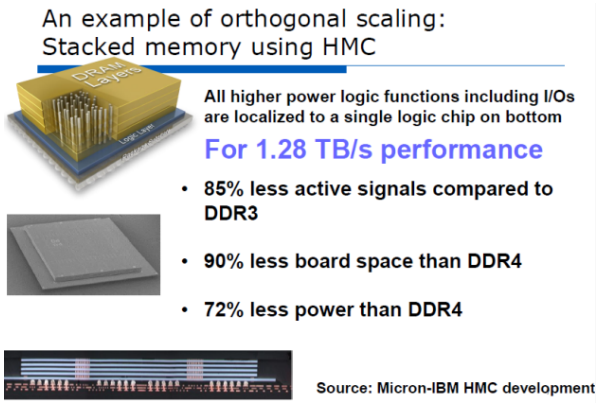

Fig. 3. Illustration of "orthogonal scaling" that continues to yield increased device density and performance even as Moore's Law device scaling seems to be ending [7].

This change in thinking in terms of scaling is in alignment with the work being developed by ITRS. ITRS is helping to guide the development of System in Package (SiP) technologies as a path to improve computing performance, power, cost, and size without dependence on conventional CMOS scaling. ITRS has a living document on SiP that is kept for use as a continuing reference to the state of the art in SiP. Stacked memory and logic devices and small modules have been the broadest adoption of SiP to date [10].

2) Neuromorphic Computing

Neuromorphic Computing is very different from conventional computing which is based on the classic von Neumann architecture where there is a separate arithmetic logic unit and a memory unit, and data is shuttled between these two units. In neuromorphic computing, logic and memory are closely integrated in the same basic device, the neuron, and in connections between neurons known as synapses. Neuromorphic computing leverages the organization of the brain (neurons, connecting synapses, hierarchies and levels of abstraction, etc.) to identify those features (massive device parallelism, adaptive circuitry, content addressable distributed memory) that may be emulated in electronic circuits [7].

The SyNAPSE project is a DARPA effort in neuromorphic computing. The SyNAPSE architecture (see Fig. 4) shows an example of the type of thinking behind neuromorphic computing, where CMOS circuits and adaptive crossbar junctions emulate the dynamic behavior of biological neurons and synapses [7].

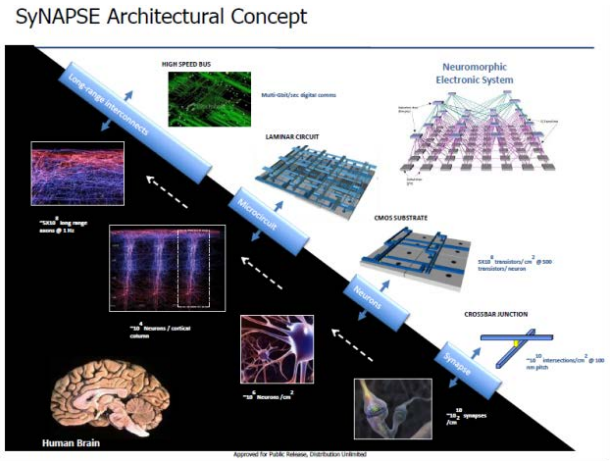

Fig. 4. Memory/processing hierarchy in neural systems and in neuromorphic circuits [7].

Under DARPA's SyNAPSE, IBM researchers have designed a brain-inspired chip that has more than 5 billion transistors with a million electronic "neurons" and over 250 million "synapses" between neurons. Biological systems consume much less energy than current computers performing the same tasks. Each SyNAPSE-developed chip consumes less than 100 milliWatts of electrical power during operation. In one comparison using tasks of pattern recognition, the SyNAPSE-developed chip used two orders of magnitude less in energy than the traditional computing systems. This high energy efficiency is achieved partly by the chip running in an asynchronous manner similar to how the brain works, and processing and transmitting data only when needed. This new chip design and high energy efficiency makes it ideal for remote sensors where electrical power is limited [11].

Though we don't fully understand true "intelligence" and how the brain works yet, short-term progress can continue to be made in this field. Near-term efforts can focus on dynamic (unsupervised) machine learning, preferably in a low-power system such as a smartphone. This type of learning would have real application in image recognition, virtual reality, and anomaly detection [7].

Interdisciplinary collaboration with neuroscientists such as those working on the U.S. Brain Initiative would also be helpful [7].

3) Approximate Computing

Traditionally, computing hardware and software were designed for numerical calculations requiring a high degree of precision, but for many applications today such as image processing and data mining, a sufficiently good answer is enough. Conventional logic circuits are also highly sensitive to bit errors, but as devices get smaller, the likelihood of random errors increases. Approximate computing represents a variety of software and hardware approaches that seek to trade off accuracy for speed, efficiency, and error-tolerance [7].

Approximate computing is very similar to how the human brain works in terms of its ability to scale the degree of accuracy depending on the given task. Asking someone

whether dividing 500 by 21 is greater than one requires a lot less time and effort than asking if the answer is greater than 23 [12].

Since an approximate result can be obtained with decreased computation time, this translates to reduced power on devices. A computation that is only as accurate as needed will be lower in energy (and time) than one which always achieves maximum accuracy. This is important for edge devices that need to be energy efficient [7].

Researchers have shown to apply approximate computing to programmable processors. The interface between the software and hardware is altered to allow the software to tell the hardware the level of accuracy needed for a given task [12].

To progress further, the computer culture needs to adopt randomness/approximation more widely and permits it to be easily implemented. Embracing error holistically across the system stack, and in communication and storage, provide more opportunities to exploit the concept of approximation for more efficient and higher performance systems [9].

4) Adiabatic/Reversible Computing

One of the primary sources of power dissipation in digital circuits is associated with switching of transistors and other elements. Adiabatic and reversible computing describe a class of approaches to reducing power dissipation on the circuit level by minimizing and reusing switching energy, and applying supply voltages only when necessary [7].

Reductions in power can come from slowing down the clock speed of the circuits. This reduction in speed, which is opposite the traditional direction, can be compensated by increasing the integration scale and applying massive parallelism (see Table 1). A computer could be scaled from a 2D array of 10^8 gates in 1 cm^2, to a 3D array of 10^{15} gates in 1 cm^3. As reflected in Table 1, this type of scaling could lead to a power-efficiency improvement of 3000. While this is likely not possible with conventional CMOS and more research is needed, one might envision a 3D assembly with integrated logic and memory [7].

TABLE I. HOW TO GET ENHANCED COMPUTING PERFORMANCE USING ADIABATIC COMPUTING AND 3D DEVICE PACKAGING [7]

Timeframe	Today	Changes	Tomorrow
Integration scale	10^8 logic transistors	$\times 10^7$	10^{15} logic transistors
Clock speed	3 GHz	3000× slower	1 MHz
Performance	Chip is 2D comprised of 100 nm^2 gates.	3000× reduction in joules/op OR 3000× increase in energy efficiency	Chip is 3D comprised of 100 nm^3 gates.

On the software side, improved algorithms are an area that can provide energy-efficiency gains. Various algorithms may solve a particular computational problem, but they might utilize differing amount of time, memory, and energy. A reversible algorithm is one that can run forward or backward, and the theory is that when an algorithm is run backward, no energy is lost during the computation. For the new algorithms to work, they would need to run on a reversible computer with reversible components [13].

A tangential benefit to adiabatic/reversible computing is security. Energy is more uniformly distributed in this type of computing. Computer attacks that rely on figuring out how much and where energy is consumed will be tougher to achieve. Besides energy improvement, security will be another reason for companies to implement adiabatic/reversible computing in their infrastructure [13].

5) Parallelism

Current approaches in parallelism and supercomputing to improve computing performance can be limiting. High performance computing is based on massive parallelism which is inefficient and expensive due to the high cost of needed power and electricity. For example, present supercomputer hardware that is based on massive parallelism uses ~25 MW of power and costs about $25M. Programming of parallel computers is also inefficient due to the unpredictable effects of inhomogeneous heating. Temperatures can vary greatly which causes variance in clock speeds and uncontrolled latency [9].

To address some of these issues, new tools that will measure and predict distributions in latency and processor/memory execution will be helpful. Having software that provides users access to dynamically control processors and memory may better allow parallelism to work. Integrated processor and memory can also provide more effective parallelism with less data to move. Certainly, lower-power technology will be key to enabling even more massive parallelism [9].

Major hurdles will need to be overcome to achieve exascale computing, which is 10^{18} calculations per second. In a hearing held by the House Science Committee on America's Next Generation Supercomputer: The Exascale Challenge, one researcher/professor noted that the five major challenges to reaching exascale computing include: reducing power consumption by at least a factor of 50, increasing the parallelism of applications software and operating systems by at least a factor of 1000, developing new programming methods, improving memory performance and cost by a factor of 100, and improving systems reliability by at least a factor of 10. All components including processors, storage, wiring, and interconnections must be redesigned to achieve these factors [14].

IV. PATH TO INTERNET OF THINGS

Several of the concepts discussed at the RC Summits have potential application to the development of an IoT environment. The three pillars of computing, which are

Energy Efficiency, Human-Computer Interface, and Security are important foundation in the Internet of Things. Similarly, the computing technologies and approaches discussed at the RC Summits may help in the maturation of the IoT.

Each of the IoT-connected devices must be energy efficient and energy aware. The concept of the edge mobile devices harvesting energy from an intermediate edge infrastructure could be an architectural implementation design to be considered for IoT. Similarly, the oceans of data that need to be processed could be off-loaded to the core of large and powerful datacenters with advanced computing and processing to further minimizing power consumption of the edge devices so that energy harvesting is sufficient. At the same time, energy efficiency in the core or grid still need to be addressed or it might become a major cost of operation. "Orthogonal Scaling" and adiabatic computing with massive parallelism might help drive power reductions in the circuits in devices and systems not only in the core/grid but throughout the IoT environment.

Applying the concept of dynamic and unsupervised machine learning in neuromorphic computing to the HCI and applications will most likely provide a better user-to-computer experience for the IoT edge device users. Some IoT applications can leverage approximate computing when an accurate answer is not necessary but a sufficiently good answer is enough to minimize energy consumption from user and network devices. Neuromorphic computing approaches offer great promises for efficiently analyzing large amounts of unstructured information under dynamic conditions that IoT demand.

Circuit designs for the devices and systems in the IoT can leverage the approach of augmenting CMOS to enhance performance through "orthogonal scaling" of memory storage and processing power.

Security is an important aspect in IoT as it may be the biggest system that mankind has ever built. With billions and perhaps trillions of nodes, security against unauthorized access or control must be distributed throughout the system. No centralized master protection system would be effective. Establishing forward-looking security standards and practices will help establish and implement a universal secured and trusted network of systems and devices. Each node must contain its own standardized security module that can automatically be updated with the latest protection algorithm. Furthermore, privacy must be maintained, which requires adequate encryption of data signals, as well as proper verification of authorized users.

Below are some key elements to consider in the development of a fully secure IoT.

- Low power security
- Neuromorphic computing
- Machine learning
- Separation of devices/networks
- Fail properly techniques
- Anti-hacker techniques
- Leverage the physical in cyber/physical
- Aggressive response to attack (antibodies)
- Learn from fault tolerance
- Formal methods for simple systems
- Aggregation of data issues

V. Conclusion

Further R&D is needed for many of these promising Rebooting Computing technologies and approaches. They need to be further developed and matured before supplanting conventional CMOS and existing computing platforms, but their benefits in performance, power consumption, security, and human interaction are promising. At the same time, broader collaboration and partnership with organizations such as ITRS and the BRAIN initiative may help bring new ideas to the forefront to continue to further future computing technologies.

Despite the apparent saturation of traditional Moore's Law device scaling, it seems clear that the exponential growth in numbers of devices connected to the Internet will continue for the foreseeable future. This is the vision of the Internet of Things (IoT), with seamless integration among edge devices, communication platforms, and high-performance computing engines. For these to work effectively, edge devices must be low-power or no-power (energy harvesting), while data centers must operate efficiently. The networks must operate autonomously, without centralized control or constant human monitoring. Furthermore, the networks must maintain security and foster trust. Finally, the IoT must be configured and programmed so as to improve the quality of life (health, information, education, entertainment, transportation, employment) for all citizens.

We hope that the work done in the IEEE Rebooting Computing Initiative will contribute to the materialization of an effective, energy-efficient and secure IoT.

Acknowledgment

We thank the IEEE Rebooting Computing Steering Committee for allowing the use of the material from the Rebooting Computing Summits in this paper.

References

[1] J. R. de Marca, "Coming Next: The Internet of Everything," IEEE *the institute*, 7 March 2014. [Online]. Available: http://theinstitute.ieee.org/opinions/presidents-column/whats-coming-next-the-internet-of-everything [Accessed 28 October 2014].

[2] D. Evans, "The Internet of Things, How the Next Evolution of the Internet is Changing Everything," Cisco Internet Business Solutions Group (IBSG), April 2011. [Online]. Available: http://www.cisco.com/web/about/ac79/docs/innov/IoT_IBSG_0411FINAL.pdf [Accessed 8 January 2015].

[3] S. Solur, "Internet of Things and Wearables," December 2013. [Online]. Available:

http://h30614.www3.hp.com/collateral/Barcelona2013/presentations/IT3112.pdf [Accessed 8 January 2015].

[4] Goldman Sachs Global Investment Research, "What is the Internet of Things?," Goldman Sachs, September 2014. [Online]. Available: http://www.goldmansachs.com/our-thinking/outlook/iot-infographic.html [Accessed 8 January 2015].

[5] IEEE Computer Society, "IEEE CS 2022 Report," IEEE, February 2014. [Online]. Available: http://www.computer.org/web/computingnow/2022-Report

[6] K. Pretz, "The Future of Computing," IEEE *the institute*, 8 March 2013. [Online]. Available: http://theinstitute.ieee.org/technology-focus/technology-topic/the-future-of-computing [Accessed 28 October 2014].

[7] A. Kadin and the IEEE Rebooting Computing Committee, "RCS 2 2nd Rebooting Computing Summit Summary Report," IEEE, June 2014. [Online]. Available: http://sites.ieee.org/rcsummit/files/2014/07/RCS2ReportFinal.pdf

[8] R. Wedewer, "Rebooting Computing Summit - Event Summary," IEEE, December 2013. [Online]. Available: http://rebootingcomputing.ieee.org/images/files/images/pdf/ieee_rebooting_computing_summary.pdf

[9] A. Kadin and the IEEE Rebooting Computing Committee, "RCS 3 3rd Rebooting Computing Summit "Rethinking Structures of Computing" Summary Report," IEEE, Nov 2014. [Online] Available: http://rebootingcomputing.ieee.org/images/files/images/pdf/rcs3reportfinaldraft9dec14.pdf

[10] International Technology Roadmap for Semiconductors, "The Next Step in Assembly and Packaging: System Level Integration in the Package," 2007. [Online]. Available: http://www.itrs.net/Links/2007ITRS/LinkedFiles/AP/AP_Paper.pdf . [Accessed 9 January 2015].

[11] DARPA, "SyNAPSE Progam Develops Advanced Brain-inspired Chip," DARPA, 7 August 2014. [Online]. Available: http://www.darpa.mil/newsevents/releases/2014/08/07.aspx [Accessed 9 January 2015].

[12] E. Venere, "'Approximate computing' improves efficiency, saves energy," Purdue University, 17 December 2013. [Online]. Available: http://www.purdue.edu/newsroom/releases/2013/Q4/approximate-computing-improves-efficiency,-saves-energy.html [Accessed 10 January 2015].

[13] N. Stauffer, "Energy-efficient Computing," MIT Energy Initiative, 20 June 2013. [Online]. Available: http://mitei.mit.edu/news/energy-efficient-computing [Accessed 11 January 2015].

[14] I. Wladawsky-Berger, "Parallel Supercomputing: Past, Present and Future," CIO Journal, 2 August 2013. [Online]. Available: http://blogs.wsj.com/cio/2013/08/02/parallel-supercomputing-past-present-and-future/ [Accessed 11 January 2015].