Disruptive Solutions for Energy Efficient ICT

WRITTEN CONTRIBUTIONS

Expert Consultation Workshop
8 & 9 February 2010

••• Future and Emerging Technologies

Proactive

European Commission
Information Society and Media

Written Contributions - Expert Consultation Workshop
Disruptive Solutions for Energy Efficient ICT
Brussels, 8-9 February 2010

Table of contents

Guiding Questions

1. Challenge / Motivation

What are the characteristics of the challenge? What is the motivation behind it? What is the contribution to transforming ICT and ICT research with respect to energy efficiency? Will it profit from new insights gained through multidisciplinary perspectives?

2. Suggested Approach / Specific Topics

What are the foundational transformative research topics? What approaches can be suggested? List research topics that would be relevant and together bring ICT forward with respect to energy efficiency.

3. Target Outcome / Expected Impact

What would be the specific outcome of research supported in this area? What would be the expected impact of a breakthrough in this area in a time frame of 10 to 20 years? What are the success measures / indicators?

4. Suitability for ICT and FET / Long-term Vision

In what way and how far does it go beyond the state of the art? What makes this suitable for ICT-FET as opposed to mainstream ICT? Is it vision-driven and high-risk, embryonic or foundational? Is there already a (pre-)established research road map for this topic(s)? What would be the added value of FET funding?

5. Communities

Is this topic addressing an existing or new community? What areas of expertise are to be involved? Is it a multidisciplinary area, and if so is this a new or an existing mix of disciplines? What would be the critical mass of European researchers needed to carry out research in this area? What is the current situation?

Written Contributions

Luca Benini,
Università di Bologna

1. Challenge / Motivation

Characteristics of the challenge (why now? what is the motivation?):

1. Moore's law benefits in terms of energy have reached a region of diminishing returns in terms of energy efficiency. Addressing the energy efficiency challenge just through the benefits of device scaling is now impossible. More radical innovation is therefore needed.

2. ICT-based solution are becoming increasingly widespread and more critical for well-being and economic sustainability of advanced countries (EU, USA, JAPAN), and the quest for higher-quality quality ICT-enhanced solutions is not slowing down (e.g. 3D-TV, augmented reality, pervasive connectivity, internet of things), thereby imposing exponentially tighter GOPS/W requirements

3. With Cindia economies growing fast, more than 2B people are going to create a huge demand for ICT products and solutions. Hence, any saturation in highly-developed countries will be more than compensated by the growth in developing economies

4. ICT is becoming the cornerstone for sustainable development: most of the green economy innovation is fueled by ICT (energy-efficient cars, buildings, etc.), thanks to embedded intelligence as well as computational techniques for modeling and designing energy efficient artifact and supporting sustainable development policies (computational sustainability)

2. Suggested Approach / Specific Topics

From lower to higher abstraction level:

a. technology

a.1. energy-efficient devices beyond pure scaling - not only smaller.

a.1.1 new devices with lower Watt/Area (lower power density)

a.1.2 new devices with higher Watt(idle)/Watt(active) (better power controllability)

a.2 energy-efficient communication

a.2.1 energy-efficient interconnects - hybrid communication and computation (smart wires), beyond ohmic interconnects (RF, optical)

a.2.2 3D integration as enabler for energy efficiency (shorter interconnects, heterogeneous integration)

a.3 energy-efficient storage technologies

a.4 energy-efficient organic computing fabrics

b. circuits (possibly leveraging disruptive technology innovations, but also in mainstream technology)

b.1 ultra-low S/N ratio circuit (e.g. Ultra-low voltage CMOS with noise tolerance)

b.2 energy-recycling logic (adiabatic circuits and beyond)

b.3 unregulated supply digital and analog circuits enabling direct coupling between computing fabric and energy harvesters)

b.4 energy efficient circuit templates for organig and post-silicon devices

b.5 highly-integrated, efficient environmental energy conversion devices

c. architecture

c.1 teraops/watt architectures - thousand cores and beyond (average power, power delivery, thermal issues)

c.2 energy-neutral computing fabrics - computing without batteries and power grid

c.3 architectures for ultra-energy efficient using non conventional computing paradigms (e.g. soft computing architectues - neural, fuzzy, etc.)

c.4 architectural templates for energy-efficient non-silicon devices and circuits

d. systems

d.1 toward instantly-available (on-demand) computing with zero idle power at multiple scales (deeply embedded devices, personal devices, cloud computing)

d.2 predictive and pro-active environmentally powered systems (energy management and provisioning)

d.3 energy-neutral and energy positive embedded devides (they harvest enough energy to power them selves or they save more energy in the system they are embedded in than the energy required for they operation)

e. Modeling and design tools and methods

e.1 Computational sustainability: Modeling and multi-objective optimization of very-large scale, dynamic, uncertain systems-of-systems requiring an intelligent orchestration of contributions coming from applied mathematics, artificial intelligence, statistics, dynamic systems.

e.2 Holistic system level methods and tools: cross-layer approaches (from power supply to technology tuning) that go significantly beyond the "state-of-the-art" design tool layering and interfacing to enable strong co-design sinergies between various susystems (horizontal) and design abstraction layes (vertical).

3. Target Outcome / Expected Impact

Quantitative metrics.

a. energy efficiency GOPS/W

b. energy sustainability net energy flow (J intake)

Major breakthrough (goal):

energy-neutral embedded computing

Widespread usage of computational approaches for addressing global sustainability challenges (i.e. computer-aided tools developing for sustainable policies, programs and their assessment, energy-footprint aware decision support systems)

4. Suitability for ICT and FET / Long-term Vision

Moving beyond silicon roadmaps

Emphasize stronger links on between system and architecture circuits and technologies beyond market segmentation - demonstrate the advantages of a thoroughly holistic approach.

Reference roadmaps

ITRS

SRAs: Artemis, Eniac

Catrene

ITEA

5. Communities

The focus would be on creating truly multi-disciplinary teams, a preliminary condition for creating a new community.

Mix of disciplines: material science, chemistry and physics, electrical, power and computer engineering, system design, economics, systems biology.

Current situation: links between fundamental sciences and applied sciences is still weak - different terminologies different goals.

Mats Brorsson,
KTH School of Information and Communication Technology, Stockholm

1. Challenge / Motivation

The use of ICT services has been steadily and ever-faster increasing since the advent of Internet some 40 years ago. The general public started using it about 15 years ago and Internet-based services are now expected both between citizens and authorities, citizens and private enterprises and between businesses (B2B).

Power consumption in ICT-equipment has been a primary focus in the research community for about 10 years and has targeted all levels from applications programming down to CMOS circuit level. For mobile equipment, the driving force has been to extend battery life and in stationary equipment the main problem has been to master the heat dissipation. One of the most power-hungry components in ICT-hardware is the main processing unit, the microprocessor. We have until some five years ago been able to meet the increasing demand on advanced ICT services by increasing the clock frequency on microprocessor-based components. The price for this has been a galloping power consumption and since about 2005 the maximum clock frequency does no longer increase. Instead the possibility for performance improvement comes from multicore solutions where two or more processing cores can be used to speed up applications which can be parallelised.

Unfortunately, development of ICT services, which is largely a software development business, has not yet embraced taking power (and thus also energy consumption) into account. It is till mainly viewed as a heat dissipation problem (except for larger server farms where the electricity bill is non-negligible) even for hand-held mobile devices where the hand comfort is in fact the limiting factor for the power consumption setting it to be about 3 W. The service providers are most often not the operators of hardware and other infrastructures such as networks and can not directly influence their usage and have no incentives in minimising power consumption. There is much to do in order to let energy consumption for ICT services be a first order design parameter.

On the other hand, it does not make sense to set a cap on the ICT energy consumption. In fact, its share of the total energy consumption should grow as we have reason to believe that ICT services are part of the solution to lower the global energy consumption. Some examples of this from diverse business segments are:

- Algorithms to optimise logistics in order to reduce transportation and storage needs (WalMart),
- Collaborative tools and systems to replace and complement physical meetings,
- Engine control systems to reduce combustion engine emissions and to make fuel usage more efficient,
- SmartGRIDs where clever adaption to electricity usage can help produce and transport the electricity more efficiently, and
- Clever coordination of nearby base stations in mobile phone networks to emit the radio power more efficiently to the mobile phones nearby.

The main challenge is to make energy consumption visible across all interfaces from hardware, networks, processors, operating systems, low-level software and service applications to service interactions.

2. Suggested Approach / Specific Topics

There are two ways to proceed. On the one hand, we should seek ICT technologies that have the potential to significantly reduce the energy use in other segments of society as outlined above. My belief is that this approach has the highest possibility to achieve significant energy reductions. On the other hand, with this approach, the ICT-share of energy consumption increases and there is a need for techniques that also significantly reduce this. A lot of effort has been spent during the last ten years to reduce, in particular, power consumption at device level, architectural level and at the lower levels of system software. Examples of the latter are low-power communications protocols, OS scheduling for low power and compilation optimisations for low power. A reduction in power consumption most often also leads to a reduction in energy consumption if the performance is kept at a similar level as before the power reduction.

In higher levels of the software stack, however, the energy consumption (or even performance to some extent) aspect is largely absent.

If we should focus on the possibilities to reduce power consumption in the ICT area itself, the following areas need attention:

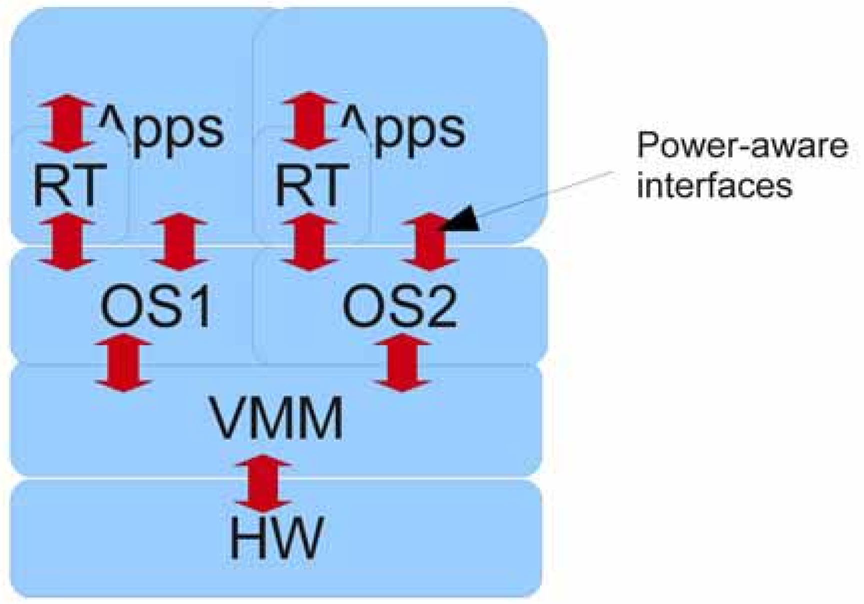

Figure 1: A typical software stack.

ICT-associated equipment:

Displays, hard disk drives, printers etc. Each of these needs to be designed with low energy consumption as first-order priority in even higher degree than already done. The effort has to be spread across device technology to life-cycle analysis.

Computational platforms:

This involves energy reduction in components such as processors, DRAM memory devices and associated components.

Software development:

Development of software can contribute much more to energy conservation that what is the case today. I explain more later.

Research into energy reduction in devices and components is largely independent of other approaches except in one area. Recent development has brought forward a multitude of control points in processors and DRAM devices that gives software the possibility to control the trade-offs between power consumption and performance. This has largely been underutilized and we have an opportunity to make significant progress in taking advantage of this possibility.

In order to be able to take advantage of the capabilities in the hardware and to make all software-developers energy-aware, I propose that we take a holistic approach with standardised interfaces between all layers in the software stack. These interfaces should be double directed so that information from lower levels can come up all the way to the application/service developer, if needed, and information from the application can go down to lower-levels in order to provide information for better, e.g., scheduling decisions.

Figure 1 shows the various layers in a typical software stack. In this example, on top of the hardware runs a Virtual Machine Monitor (VMM) which provides some rudimentary services in order to provide virtualization of the hardware so that two or more operating system instances (OS1 and OS2) can co-exist. On top of the operating systems run applications, possibly with extensive runtime (RT) system support. Other software stacks could also include managed software systems between the OS and the application. The hardware layer is increasingly parallel and the amount of parallelism and heterogeneity in the hardware layer will only increase.

The figure also show interfaces between all layers that I propose are needed in order to exchange information about resource usage and needs in both directions. Application developers can expose information about resource needs that cannot be deduced from run-time measurements and system software can get information about factual resource usage, including energy that also can provide information for better informed decisions about scheduling both hardware resources and software components such as threads and tasks.

3. Target Outcome / Expected Impact

Depending on which area we are focus on, the outcome can be very different. Policy makers can enforce or provide encouragements to promote the use of ICT as support to reduce energy consumption in society. If focus is on the ICT area itself, I believe we can expect a normalized energy reduction of about 50 % with sustained effectiveness.

4. Suitability for ICT and FET / Long-term Vision

TBA

5. Communities

TBA

Laurent Dussopt,
CEA - LETI, Grenoble

1. Challenge / Motivation

What are the characteristics of the challenge? What is the motivation behind it? What is the contribution to transforming ICT and ICT research with respect to energy efficiency? Will it profit from new insights gained through multidisciplinary perspectives?

Current mobile terminals autonomy is insufficient and roughly constant over the years due to moderate improvements of batteries capacity and increased number of functions and processing needs. A similar autonomy issue penalizes the domain of wireless sensor networks.

A major challenge is to bring disruptive improvements in the energy efficiency of communication and signal processing systems through micro/nanomechanical-oriented architectures.

Micro/nanomechanical structures can provide elementary functions (switching, filtering, mixing) with very high quality factor, i.e. high power efficiency. Going further toward communication system architectures fully, or nearly fully, built from mechanical devices would open the way to orders of magnitude in energy savings.

This topic is multidisciplinary since it has to built from micro/nanotechnologies and wireless communication systems, including RF architectures, IC design and antennas.

2. Suggested Approach / Specific Topics

What are the foundational transformative research topics? What approaches can be suggested? List research topics that would be relevant and together bring ICT forward with respect to energy efficiency.

Both top-down or bottom-up approaches should be envisaged in order to develop a long-term vision of the perspectives and bottlenecks of this topic. In a bottom-up approach, micro/nano-devices (cantilevers, resonators, nanowires, carbon nanotubes, etc.) are investigated to implement basic functions (filters, oscillators, antennas, detectors, mixers, etc.) and eventually reach the demonstration of sub-systems. A particular topic of interest is the investigation of arrays of nano-devices operating collectively.

Since these new building blocks will not be traded for today's semiconductor functions on a one-to-one basis, specific micro/nanomechanical-oriented communication systems architectures have to be developed to take the best benefits of the available technologies.

3. Target Outcome / Expected Impact

What would be the specific outcome of research supported in this area? What would be the expected impact of a breakthrough in this area in a time frame of 10 to 20 years? What are the success measures / indicators?

The outcomes targeted at short-middle term is the demonstration of key building blocks at the component, circuit or sub-system level, and the validation of new hardware architectures. In 10-20 years, the perspective is to enable the realization of (quasi-)fully mechanical RF front-ends.

The main impact would be an improvement of the autonomy of mobile communication terminals and wireless sensors by one or several orders of magnitude. Mobile communication terminals will benefit from increased potential for new power-hungry functions such as wireless video streaming for instance. Wireless sensor networks will see a dramatic improvement in miniaturization and energy performance, which in turn will enable numerous industrial and ambient intelligence applications.

4. Suitability for ICT and FET / Long-term Vision

In what way and how far does it go beyond the state of the art? What makes this suitable for ICT-FET as opposed to mainstream ICT? Is it vision-driven and high-risk, embryonic or foundational? Is there already a (pre-)established research road map for this topic(s)? What would be the added value of FET funding?

Although some elementary devices have been already demonstrated (filter, mixer), many others are still to be investigated and do not have an established feasibility. More importantly, state-of-the-art devices are not suitable for a one-to-one replacement into current RF front-ends due to incompatible technologies or different impedance domains. The move from elementary functions to full systems will require a joint effort from several communities to refound the way RF signals are processed, and the actual feasibility of these architectures is not established. In this respect, it is a high-risk research area.

It is expected from a FET funding to bring a critical mass of multidisciplinary teams on this topic with a perspective of outcomes in the industrial domain of at least ten years.

5. Communities

Is this topic addressing an existing or new community? What areas of expertise are to be involved? Is it a multidisciplinary area, and if so is this a new or an existing mix of disciplines? What would be the critical mass of European researchers needed to carry out research in this area? What is the current situation?

The current situation on this topic is isolated efforts of some research groups focused on a specific elementary device, and therefore not addressing the long-term vision at the system level. This topic is addressing several existing communities: micro and nano electromechanical systems, electromagnetism and antennas, IC design and architectures.

Chiara Francalanci - Represented by Eugenio Capra, Politecnico di Milano

1. Challenge / Motivation

Research on green ICT is inherently **interdisciplinary**. A new ICT is green if it allows energy savings. A new green ICT has a practical impact if energy savings are economically feasible, i.e. they translate into lower overall costs.

Costs represent an economic variable that has technical drivers and, therefore, requires both technical and economic skills to be managed, both in research and in practice. As a result, research on green ICT should be justified in terms of costs (as opposed to energy savings only) and should involve **technical** and **economic/managerial** competences.

2. Suggested Approach / Specific Topics

First, three basic research fields should be distinguished:

1) **green ICT**, defined as new ways to design ICTs with lower power consumption, within this topic, three main research areas should be distinguished:

    a) **green hardware**, defined as the design of new hardware with lower power requirements,

    b) **green data centers**, defined as the design of new methodological approaches to design and manage data centers with lower power requirements,

    c) **green distributed ICT**, defined as defined as the design of new methodological approaches to design and manage distributed ICTs (PCs, printers, faxes, etc.), and

    d) **green software**, defined as the design of new methodological approaches to develop code with lower energy requirements (with a further distinction between **embedded** and **enterprise** software).

2) **ICT for a greener business**, defined as new ICTs or new applications of ICTs that reduce the energy consumption of the business processes where they are applied, and

3) **WEEE** (waste electrical and electronic equipment) in ICT.

3. Target Outcome / Expected Impact

Most research in this field is mainstream. Breakthrough long-term innovation can be obtained on topic 1.1 (see above), green hardware and green software.

Mainstream research should have ambitious targets in terms of percent savings, with an emphasis on feasibility to be applicable. Many existing solutions require such investment costs that energy savings do not pay off. This means that the reason behind the investment is not green ICT and related research should not be labelled as research on green.

Green software in particular is a potential field of breakthrough and visionary innovation.

Whereas most of current researches focus on green hardware, which is phisically responsible for the consumption of energy, it is software that “guides” the processor and that is consequently responsible for most of the energy consumend by an IT sysem.

Elaborating information *per se* requires energy, according to recent quantum physics researches. Quantum physics also quantifies the minimum theoretical amount of energy needed to commute a bit of information (Margolus-Levitin theorem), which according to the current state of our knowledge could be optimally represented by the spin of an electron. Of course this minimal amount of energy is much lower than current consumptions. This gap between theoretical and actual consumptions is due to all the inefficiencies introduced by the different architectural layers of a computational system, e.g. because we use transistors rather than atoms to store and elaborate bits.

A software application executes a certain number of computations on a defined number of bits in order to obtain a result. From a logical perspective, the energy consumption of a software application can be estimated as:

$$EC_{logical}(f) = E(f) \cdot C_c \cdot T_d$$

where $E(f)$ is the energy required by a single bit status commutation at frequency f, C_c is the computational complexityof the application that is executed (i.e., the number of commutations requires) and T_d is the thermodynamic depth of the computation that is performed onto the problem representation (i.e., the number of bits).

In order to approach the theoretical maximum level of energy efficiency, research should focus on all three factors.

Optimization of $E(f)$ is mainly related to green hardware, as it is about storing and commuting a single bit with as little energy as possible.

C_c can be affected through software applications and compilers, which ultimately determine the effective number of elementary operations to be executed.

Eventually, T_d is related to how information is stored in the system and to the efficiency in the use of storage.

It should be noted that the goal green IT research is to minimize the product of the three factors, and not the single factors, which may not independent from each other. For example, data memorization and the specific utilization of storage may affect the complexity of software.

5. Suitability for ICT and FET / Long-term Vision

Breakthrough long-term innovation can be obtained on topic 1.1 (see above), green hardware and green software.

6. Communities

The green grid, but it is not a community, rather an aggregation of large companies that manage their interactions as a community. Many white papers are released, which is positive.

Luca Gammaitoni,
NiPS Laboratory, Università di Perugia

1. Challenge / Motivation

What are the characteristics of the challenge?

The challenge that we propose to address here is the description of **energy transformation processes at the nanoscale** aimed at unveiling new mechanisms for powering next generations of ICT devices.

What is the motivation behind it?

An ICT device is a machine that inputs information and energy (under the form of work), processes both and outputs information and energy (mostly under the form of heat). Energy efficiency is usually defined as the *percentage of energy input to a device that is consumed in useful work and not wasted as useless heat*, however this definition does not apply when we have to deal with processes taking place at nanoscale. The well-known laws of heat and work transformation that lie at the base of the classical thermodynamics are going to need a rethinking. The very basic mechanism behind energy dissipation requires a new definition when non-equilibrium processes involving only few degrees of freedom are considered.

What is the contribution to transforming ICT and ICT research with respect to energy efficiency?

Future ICT will be largely affected by nanoscale devices that process information while transforming work in to heat and heat into work. Pioneering work developed by R. Landauer in the sixties has shown that information processing is intimately related to energy management issues ("information is physical"). The dream of realizing highly efficient computation devices has to deal with both energy and information dissipation as in the paradigmatic example of reversible computers.

The long-term goal of this challenge is to make possible the realization of extremely low power ICT devices with a significant impact on energy efficiency on a much broader scale. A new generation of energy efficient ICT device has to deal with energy transformation processes at nanoscale.

Will it profit from new insights gained through multidisciplinary perspectives?

This challenge requires unavoidably a multidisciplinary approach where competences from fields as diverse as physics, computer science, electronic eng., mechanical eng., chemistry and biology are brought together in a coordinated effort.

2. Suggested Approach / Specific Topics

What are the foundational transformative research topics?

The foundational research topics for this challenge cover at least the following general subjects: "non-equilibrium thermodynamics", "thermal and charge transport" in solid state systems, "statistical fluctuation theory", "nonlinear dynamics", "quantum and classical information theory".

What approaches can be suggested?

To address this challenge we can imagine different approaches that complement each others. On the theoretical physics side we need a better comprehension of the dissipation mechanisms and non-equilibrium relaxation in solid-state systems, together with a new understanding of the charge and heat transport. From computer science side we need a deeper understanding of the intimate relation between information dissipation and energy management. From biophysics point of view (both theoretical and experimental) we need a better comprehension of the mechanism of molecular motors in the presence of nonequilibrium random fluctuations much in the line of “bio-inspired” energy management. From electronic eng. perspective we need new ideas for state dependent devices and nanoscale “switches”. From eng. point of view an interesting approach will be the study of N(O)EMS as basic elements to implement the ICT functions, as a more energy efficient alternative to semiconductor electron devices.

List research topics that would be relevant and together bring ICT forward with respect to energy efficiency.

To address this challenge the following topics are relevant:

- work and heat transformation at nanoscale
- charge and heat transport in solid state devices
- phonon engineering
- reversible computing
- bio-inspired energy management
- noise operated logic gates

To mention a few.

3. Target Outcome / Expected Impact

What would be the specific outcome of research supported in this area?

The specific outcome of research in this area will cover the following results:

- the significant decrease of the power needed by the most common ICT devices;
- an extension of the laws of thermodynamics to nanoscale systems;
- a fundamental physics description of the processes related to heat dissipation;
- the growth of a novel field related to phonon engineering for ICT applications;
- the experimental realization of reversible logic gates;
- the realization of ICT in-body, ICT in-civil system and transport vehicles.

What would be the expected impact of a breakthrough in this area in a time frame of 10 to 20 years?

The expected impact in 10 years would be the growth of the bases for a new technology based on nanoscale devices that harness thermal fluctuations and process information with reduced power.

The expected impact in 20 years would be the production of new very low power ICT devices that are quasi-energetically autonomous with respect to a centralized power system.

What are the success measures / indicators?

In order to evaluate the progress of the research in this area the following indicators will be relevant:

- Scientific production on the subject of energy efficiency at nanoscale. It is important that in the next few years a wide scientific community produces scientific contents on this topic.
- Experimental results on nanoscale systems where energy dissipation and information production are tested.
- Reversible logic gates prototypes.
- Low dissipation charge and heat transport models with experimental tests (see e.g. thermo-electric related research).

4. Suitability for ICT and FET / Long-term Vision

In what way and how far does it go beyond the state of the art?

State of the art ICT basic devices are dominated by power dissipated in heat. Presently the main effort is aimed at cooling down the heat produced during computation with specific attention to the charge transport on one hand and on the other hand on reducing the voltage operating levels up to the point of not compromising the error rate due to voltage fluctuations.

We propose to address the problem at a very fundamental level:

- what are the basic mechanisms behind the heat production?
- How can we take advantage of the fluctuations instead of avoiding them?
- How the physics of the heat and charge transport can be merged with the phonon engineering in order to advance the computing tasks?

These questions should convey the idea that the proposed approach is way beyond the state of the art.

What makes this suitable for ICT-FET as opposed to mainstream ICT? Is it vision-driven and high-risk, embryonic or foundational?

In order to reach the expected goal we need a completely new approach that is extremely risky and still very challenging. It is not simply an incremental progress toward the reduction of heat production in room temperature conductors or new technology beyond CMOS. It is a new, visionary approach that challenges the very basic foundation of thermodynamics. We propose to understand the dissipative mechanisms at nanoscale with the aim at setting the bases for a new thermodynamics of ICT devices.

Is there already a (pre-)established research road map for this topic(s)? What would be the added value of FET funding?

Related to the very challenge proposed here, there are embryonic research efforts, distributed at European level that are today barely visible. These are sparse efforts that need focus in order to become a coordinated initiative. The FET funding would have the role of focussing attention on the topic and bringing to coalescence research branches that are developing independently in different communities and different countries.

5. Communities

Is this topic addressing an existing or new community?

I believe that this challenge is addressing a new community.

What areas of expertise are to be involved? Is it a multidisciplinary area, and if so is this a new or an existing mix of disciplines?

In my opinion it is a multidisciplinary area and it involves an existing mix of disciplines from fields as diverse as physics, computer science, electronic eng., mechanical eng., chemistry and biology.

What would be the critical mass of European researchers needed to carry out research in this area?

It is somehow difficult to estimate how many researchers would guarantee a critical mass capable of ensuring the success. It would be interesting to have something like 50 teams from different fields competing/cooperating in this effort. This might be equivalent to 5-7 medium size consortia.

What is the current situation?

Research activity in this area has been developed independently in different disciplines like: non-equilibrium statistical mechanics (theoretical Physics), reversible computing (theoretical computer science), energetic processes at nanoscsale (Theoretical Chemistry), in addition to the communities of NEMS/NOEMS and bio-inspired technologies.

Robert Glueck,
University of Copenhagen

1. Challenge / Motivation

Challenges are the good design of deterministic and time symmetric programming languages for constructing and developing fully working reversible software systems, and a firm grasp of the theoretical foundations, the design techniques and construction principles of such systems. Reversible software systems are a necessary prerequisite for the usage of practical reversible computing systems, much as it is for conventional computing systems. This class of languages shares reversibility properties with physics, whereas conventional programming languages do not.

The challenge of creating reversible computing systems reminds of the challenges of early computer developments in 1950ies where very little working hardware, no programming languages, software techniques, theoretical knowledge of computing was readily available. Thus, a multidisciplinary approach combining expertise in computer science, electronics, physics, and mathematics is essential. Reversible computing models might also be relevant for quantum computing, which is based on reversible principles by using unitary operations.

According to Moore's Law, transistor sizes shrinks slowly, but steadily. Accompanying this familiar development, are two less-known trends: (a) the energy dissipated per elementary logical step steadily decreases, and (b) the energy dissipation density per computational step steadily increases. If these two trends continue, around the year 2030, the energy dissipation will reach the Landauer limit of about 3 zeptojoule per computational step and the silicon temperature will reach its melting point. If the dissipation density has to be kept below reasonable values, the dissipation itself has to be lowered far beyond the Landauer limit. In order to cross this barrier, there is only one strategy possible: doing all computations reversibly.

2. Suggested Approach / Specific Topics

Regarding reversible SW systems: Language design and implementation, compiler construction and optimizations techniques, co-design of hardware/software and abstract machines development, as well as complexity and computability theory, also exploring the relation with quantum computing.

3. Target Outcome / Expected Impact

Examples:

(a) hardware-software computing system (e.g., for an embedded system) with at least some components based on reversible principles.

(b) main components of a minimal, but fully reversible microprocessor prototype and its instruction set architecture (ISA).

4. Suitability for ICT and FET / Long-term Vision

State-of-the-art: small state-full imperative languages and a computability notion of r-Turing completeness and component designs (for example, fast Fourier transform). Reversible computing systems approach requires cooperation of multiple disciplines at every system level from gate and circuit design, abstract machines, software layers and algorithms. Very far beyond current state. This is the suggested research that FET would support research towards, which would not be possible otherwise. The cooperation of multiple disciplines on reversible computing systems is definitely FET, not "normal research".

5. Communities

In some sense it is a known mix of expertise (computer science, electronics, mathematics, physics) which was, for example, closely practiced in the early computing system designs (1950-ies).

Currently, not more than 10 small groups (typically 1 to 3 persons) active, but working on isolated aspects (eg, Bremen, Cork, Copenhagen, Gent, Warszawa). In some cases, there is only one experts in Europe. The European research groups are also not enough networked to make concerted progress. Such small groups are threaded by abolishment due to funding schemes emphasizing short-term progress further amplifying the Mathaeus effect that leads to a narrowing of the research evolution.

In order to compete with research groups in the USA and Canada, the critical mass of European researchers is a network of about 10 research groups, each of 6 to 10 researchers (including faculty, postdocs, and PhD students).

Alfonso Jaramillo,
Genopole®-CNRS UPS3201 & Ecole Polytechnique, Évry

One of the possible solutions towards a zero-energy ICT would be the development of a living computer. This would even provide the design of self-replicating and evolvable computers. Biochemistry-based information technology is starting to move from inorganic materials to biological macromolecules and from in vitro to in vivo systems. Those systems will have very low power consumption; they will be amenable to massive parallelism and to fault tolerance. This will also enlarge our knowledge of the physical mechanisms underlying cell function and it will also allow the development of a new type of therapeutic agents, by incorporating information-processing molecular devices in living systems.

Most of the biochemistry-based information technology has been done using biological macromolecules in vitro. A key property of biological molecular-scale devices is their molecular recognition, which leads to self-assembly. This bottom up strategy requires designing complex interaction networks, which usually it is only possible using automated computational design methods. Recent work has produced methodologies to design such biological computing devices with the final goal of arriving to an integrated multi-physical and multi-scale computational cell. This requires developing the engineering principles and abstractions needed to build general-purpose large-scale molecular circuits, within constraints imposed by the physical interactions, chemical kinetics and interfering side reactions. This could be seen as implementing a "biological machine code" for an abstract language that has been automatically evolved to have a targeted behaviour. In this process, we will gain fundamental quantitative knowledge of the emergent multi-scale properties that occur when individual biomolecular reactions are organized into basic computational networks motifs, when those basic network motifs are coupled into complex scale-free networks of thousands of molecules, or when massively parallel computing could emerge through suitable cell-to-cell communication in a population of millions of cells. We have already reached the stage where molecular devices could be tested experimentally. As a result of those efforts, computer science will be enriched by a new computing paradigm, which experimentally implements the nature-inspired information representing and processing paradigms generated by artificial biochemistries.

We propose a methodology that implements a type of cell computation. We utilise bacterial plasmid conjugation to generate a combinatorial gene pool that will be later used to screen for biological networks with targeted behaviour. The selection would be implemented either by the bacteria themselves or by using a real-time computer-assisted selection in a new sorting-enabled microfluidic device. For the former, we introduce a fitness function by engineering a comparator genetic device that would be able to change the replication rate of the plasmids after comparing the gene expression of a query gene against a template. We also use computer simulation to aid in the engineering of our bacteria. This involves characterisation, at the single-cell level, of bacterial dynamics. This provides a predictive model for bacterial dynamics for our bacterial evolution at the single cell and population levels. The quantitative characterisation of conjugation at the single-cell level will also allow to better understand the spread of antibiotic resistance in bacteria. Our computational development of the plasmid libraries could be seen as a new type of "bacterial software", for which we will target some applications. We will design libraries to evolve several genetic, RNA and metabolic networks with targeted behaviour. In particular we will try to evolve oscillatory networks, which could be tested against available synthetic circuits.

Stefanos Kaxiras,
University of Patras / Uppsala University

1. Challenge / Motivation

(What are the characteristics of the challenge? What is the motivation behind it?)

Challenge: How to architect future computing systems for power-efficiency.

Since the mid 90's power consumption has become an increasing concern to the computer architects. The inherent inefficiency of the modern CPU featuring dynamic execution led not only to many micro-architectural techniques to reduce power but eventually to an abrupt halt in its development. At one point it became impossible to scale the performance of a single core by increasing its frequency and at the same time avoid exponential increases in power and power density.

The multi-core (Chip-Multiprocessor/CMP) revolution came because of the inability to scale a single core to further exploit Instruction-Level Parallelism (ILP) in a powerefficient manner and turned computer architecture in exploring alternative kinds of parallelism (Task/Thread parallelism, Data parallelism). However, this road is also faced with the same problems because of the inability to extract sufficient speedup from parallel programs (Amdahl's law for Power). In short, we do not have a good handle on parallelism for today's applications.

In addition, for the remainder life of the silicon CMOS technology, future process generations are faced with an inability to significantly reduce power consumption via scaling as was the case for previous generations. This is a result of increased leakage power when scaling the threshold voltage to accommodate lower supply voltages. The difference between supply voltage and threshold voltage determines the speed of the devices and the supply voltage determines dynamic power consumption. To maintain both a speed increase and lower dynamic power via scaling, the threshold voltage must be lowered and this exponentially increases static (leakage) power consumption. Furthermore, future process technologies are characterized by increased variability, soft errors, etc., and in general by an overall problematic reliability. Their unreliability is exacerbated by either increasing speed or lowering power (via voltage reductions). The faster and/or lower-power we want to go the more unreliable the hardware becomes.

Within this context the approaches we suggest in this document are based on some key observations:

1. We can no longer hide the struggle against the Power Wall in the hardware layers. We need to expose it all the way to the application. This is especially important for two domains: embedded architectures and HPC. Both can benefit tremendously by a tight feedback between HW and SW concerning power (and thermal) management. This brings software and hardware communities together.

2. Parallelism is key to power-efficiency: the more parallel we can make something the slower (frequency-wise) we can run it for the same performance but for significantly lower power. Parallelism is very hard for the general case. But: accelerators for specific functions/services can be made massively parallel. What are these functions/services? What are their characteristics? Does accuracy (reliability/certainty) matter? Multidisciplinary cooperation is required to understand and build massively-parallel application specific accelerators.

3. Unreliability or Uncertainty in hardware is the new reality: deep submicron technologies are faced with significant process variations, vastly increased soft errors, etc. Trying to

lower power in this environment only increases unreliability. Power + reliability must be handled together. Embrace uncertainty rather than fight it: expose it all the way for applications which can handle it. Likewise, this requires multidisciplinary cooperation to understand how applications could behave in unreliable/uncertain systems.

4. Beyond CMOS: new technologies that seek to compute not by moving charge carriers but rather their quantum properties (e.g., electron spin vs. charge) can yield significantly lower power computing systems. But, it's not easy to compute solely in these domains. Uncertainty is also a fact of life when it comes to dealing with quantum properties. Should we adhere to the design principles of the past or architect computing systems that embrace uncertainty?

What is the contribution to transforming ICT and ICT research with respect to energy efficiency? Will it profit from new insights gained through multidisciplinary perspectives?

The key observations above (and the corresponding suggested approaches) have the potential to dramatically shape the research towards power-efficient computing. An important aspect of these approaches is that they share some common traits: a holistic view of hardware and software, application-specific acceleration, massive parallelism, and hardware "uncertainty" visible to the software. Each approach is by necessity multidisciplinary, i.e., requires the involvement of different communities, sometimes from different areas to come together to make it work.

2. Suggested Approach / Specific Topics

Foundational transformative research topics and approaches include:

1. CMOS

   a. Holistic HW/SW view: Expose power consumption all the way from circuits to applications. Architect systems with a tight integration between hardware and software concerning management of power and thermal issues.

   b. Massive Parallelism: Architect massively parallel application-specific accelerators (vision, robotics, AI, brain modeling, physical systems modeling) for power efficiency: trade parallelism for operational frequency (hence, power).

   c. Integrate power and unreliability: Power and unreliability go hand-inhand in future process technologies (cannot change the one without affecting the other). Expose both to software and all layers of a computer system (OS, middleware, compilers, applications). Study application tolerance to uncertainty and exploit possible tolerance for power-efficiency.

2. Beyond CMOS (for power efficiency)

   a. Quantum-property manipulation: Spintronics, quantum computers, new devices based on manipulating quantum behavior rather than physically moving charge carriers (e.g., electrons).

      a. What can we build? Application-specific accelerators (cryptography, brain modeling, …)

      ii. How we can interface with classical computing? CMOSquantum integration. It is unlikely that we will be able to compute solely on quantum computers but likely that quantum approaches can offer excellent accelerators.

b. Manage uncertainty in quantum computation: how well does it fit with applications? The issue of reliability may play a central role in quantum approaches but this is analogous to the issues we are facing with unreliability in CMOS. Work performed in this area will be applicable in both domains.3. Target Outcome / Expected Impact

3. Target Outcome / Expected Impact

The expected outcome depends on the inherent risk in each approach/topic described above. The more speculative the research topic the more uncertain is the outcome but also the higher is the payoff.

- A new level of power-efficiency for embedded and HPC systems is possible using a holistic hardware-software approach. Success indicators are the benefits above and beyond of what the hardware or the software could achieve alone.

- Massive parallelism has the potential to disconnect performance from power for a specific class of applications that require a fixed level of performance to be useful (e.g., vision) and are inherently parallel.

- Integrating power and reliability and exposing this to the software has the potential of offering a new tradeoff for power (power vs. reliability).

- Manipulating quantum properties instead of physically moving charge carriers has the greatest potential for benefits in power-efficiency. Orders of magnitude better power consumption can be expected for comparable primitive computing operations. Other novel approaches to computing also have such potential and should be researched. However, the challenge for the architect is to understand what computing devices can be built using such technology (possibly only accelerators rather than full-scale general-purpose CPUs) and integrate them with the rest of the CMOS technologies. Uncertainty will also play a major role here and needs to be managed possibly at the application level.

4. Suitability for ICT and FET / Long-term Vision

All of the above approaches are beyond the state of the art and progressively more and more aggressive with respect to power consumption. None of them are trivial and require significant research funding support to move forward. Each has a significant research challenge associated with it: complexity in working across many layers from hardware to software, extracting and exploiting (massive) parallelism, managing unreliability (uncertainty), architecting computing systems out of difficult-to-handle and possibly unwieldy quantum devices with inherent unreliability. There is no better instrument than the FETs to fund such approaches critical for developing significantly more power-efficient computing systems in the future.

5. Communities

Finally, all of the topics and approaches suggested here are multidisciplinary. In order to understand how to architect power-efficient systems we need to bring together the hardware and software communities both from the embedded and the HPC world, application domain experts with architects (for the development of new massivelyparallel, application-specific accelerators), different groups within the architecture community that work on power or reliability in isolation, and, finally, bring together all the different communities that work on developing new devices for computing (from physicists to electrical engineers) with architects and software developers who can build computing systems out such novel devices.

Raphaël D. Levine, The Hebrew University of Jerusalem,
Françoise Remacle, Université de Liège,
Sven Rogge, TU-Delft

1. Challenge / Motivation

Energy efficiency is a central requirement for all new technologies. The current CMOS solution was shown to be energy inefficient, both at the production level and when CMOS devices operate. Complementary roads can be devised for the development of more efficient devices in terms of power consumption: use a bottom up approach for their fabrication, design computing devices that can harvest the energy from the environment for their operation, make the operation of the device and the architecture and the addressing more energy efficient (but still providing a source of energy to drive them). All of this can serve to bring the energy cost per operation closer to the information theory limit.

2. Suggested Approach / Specific Topics

New computer and architecture paradigms for implementing high level logic operations (more complex than one switching operation) in an energy efficient way, including providing inputs and reading outputs and cascading logic devices.

Design of the computing units. Molecules are inexpensive to synthesize in large numbers but the energy cost of controlling, positioning, and addressing at the molecular level is right now not energy efficient.

Benchmarking of the power consumption with respect to CMOS power consumption for the same logic operation.

3. Target Outcome / Expected Impact

Development of low power consumption technology for information and communication applications. The new technology should be judged by comparison of the implementation of a complex logical operation, one that requires a switching network, with CMOS performances. The comparison should not only occur on the device level, but include the energy cost of addressing and reading.

4. Suitability for ICT and FET / Long-term Vision

The departure from conventional device implementations and architectures with an aim at energy efficiency is a break with mainstream ICT and thus state-of-the-art and risky. Realistically we must aim towards low-power computational circuits within existing fabrication technology. Only this way can one hope to achieve a paradigm shift in the billion-dollar ICT fabrication world, that by the very nature of large scale industrial processes, are somewhat conservative.

5. Communities

Theorists for the design and simulation of the operation of the new devices.

Architecture engineers to embed the new device concepts.

Chemists, solid state physicists for synthesis, realization of the circuits, and characterization.

Bernabé Linares-Barranco,
National Microelectronics Center, Sevilla

1. Challenge / Motivation

Energy efficient bio-inspired (visual) sensing and processing.

Vision computation is a high computational problem, requiring high energy. If solved for low energy, findings can be exported to other sensory means.

Present day vision sensing and processing systems are extremely power hungry. A camera acquires sequences of still frames. Each frame has to be processed pixel by pixel, for example performing large number of convolutions, to extract first low level features (contrast and different scales, orientation of segments), which are combined later into higher level features (first simple shapes, later more sophisticated figures), which can be used to perform object recognition. Ideally, this also needs to be performed under uncontrolled lighting and contrast conditions.

However, biology teaches us that it has solved the vision problem with higher performance than human-made technologies, and with much lower energy as well. Biology does not represent visual information as a sequence of frames, but as a continuous flow of visual information represented as sequences of spikes generated asynchronously at the individual cell (pixel) level. Biology is capable of representing dynamic visual information with codes using a highly compressed number of spikes per unit time, while maintaining all the relevant information and reducing or eliminating redundant information. Energy is consumed as a per spike basis, or equivalente, as a per relevant-information quantum basis.

Also, biology has come up with a highly energy-efficient way of interchanging information among its parts (modules): spikes. When a spike travels from the brain to a mussel (about 1m) in about 40ms, since spike duration is only about 2ms, the spike sender only needs to provide charge for a travelling 5cm segment. This contrasts strongly with present day high speed transmission lines, where the sender needs to fully charge and discharge a full line. Even worst, usually high speed links have 50ohms resistive impedance, which implies a permanent energy dissipation, even in the absence ofinformation transmission. Consequently, this interconnectsenergy- dissipation problem poses a severe limitation when scaling down systems. Nanotechnology experts are becoming to be aware of this limitation for the expected new high density systems that future nanotechnology might produce.

Consequently, for improving energy efficiency new (bio-inspired) ways of representing information are required, as well as new ways communicating information, at the level of information theory, information computation, and physical information transmission.

2. Suggested Approach / Specific Topics

A. Sensors:

New energy efficient (visual) sensors are required, not based on capturing sequences of still frames, since this results in an overwhelming amount of (redundant) information that requires processing by later stages. New approaches have to be developed where pixels themselves decide to transmit information above a relevance level. Information has to be represented with frame-free approaches, or at least with frame-constraint approaches where frame time adjusts dynamically to relevant-information content.

B) Processing Modules:

Processing modules for feature extraction and computation in feature space have to be compatible with the highly compressed (frame-free) coding coming from the sensors. Processing architectures need to be scalable and preserve energy efficiency, while concentrating on relevant information.

C) Information Exchange:

Exchange of information from the sensors to the processing modules and between processing modules should be based on spikes (information quantums, events). These information-quantums have to be exchanged between modules with new physical elements that do not dissipate static energy, such as traditional transmission lines or modern LVDS (low voltage differential signalling) links, which dissipate energy even in the absence ofinformation. Novel approaches could consider, for example, soliton technolgy where self regenerating sharp spikes travel through non linear transmission lines, similarly to spikes in biological axons.

D) Architectures:

We know that the brain is structured hierarchically and processes waves of spikes coding relevant information. Furthermore, in what is known as fast recognition (when a human is capable of recognizing some shape in a flashing image in less than 200ms), we know that the neurons involved in the recognition process only have fired one single spike, thus consuming a minimum of energy in the whole process. Biology has achieved this not only because information is represented by highly efficient codes, and because information is transmitted by highly efficient means, but also because the overall architectural hierarchy is highly efficient in processing the available information with a minimum of information and energy loss. Thus, novel (hardware) architectures are required to process with minimum energy the highly compressed information represented by the spiking codes.

E) Learning and Adaptation:

Biology has accomplished all of the above (highly efficient information-relevant codes, optimum parameters for the processing modules, and optimum hierarchical architectural structures) by including learning and adaptation. For example, in the visual system, early stages are plastic (allow learning) only during early stages of development and remain hard-wired afterwards. Later processing stages may remain plastic always. Thus learning and adaptation is tuned specifically depending on the information abstraction level. Also, learning and adaptation is compatible with the information representation code. In this sense, an important progress has been done in recent years in, for example, Spike-Time-Dependent-Plasticity (STDP). This and similar techniques need to be further developed to both understand biological computing structures and develop artificial bio-inspired highly information and energy efficient schemes.

3. Target Outcome / Expected Impact

The target outcome would be energy and information efficient vision systems for recognition applications, to be used in environments with uncontrolled lighting conditions and with low energy consumption requirements, such as portable systems. An interesting application example is automatic car driving, where a lot of continuous information (lanes, other vehicles, pedestrians, traffic signs) has to be sensed, extracted and recognized within sub-second delays, and critical unexpected events (such as a falling or crossing object) have to detected and processed on the fly to guarantee security. In order to have such a sophisticated vision system deployed by the car industry, it has to be performed by a reasonably simple computing system consuming a reasonable energy (we cannot mount a computer cluster on every car).

The expected impact of a breakthrough in this area in 10-20 years could be the availability of compact and low power highly sophisticated vision artefacts capable of robust and confident recognition, to be used in a variety of applications in the car industry, surveillance and security, domestic robotics and domotics, automatic or semi-automatic surgery, and so on.

In the meantime success measures and indicators could be a partial implementations, such as advances at the sensor level, at the processing modules level, at the interconnect levels, at the computational level, and at the architectural level, yielding prototypes or demonstrators that improve in some partial aspects with respect to conventional frame-constraint approaches.

4. Suitability for ICT and FET / Long-term Vision

In what way and how far does it go beyond state of the art?

State of the art is frame-constraint vision sensing and processing. A new energy-efficient frame-free vision sensing and processing paradigm requires development of new sensors, new processing modules, new information coding principles, new system level architectures, new learning paradigms compatible with the new information representation codings. Some of these aspects are presently being developed independently by groups from various disciplines. Efforts have to be invested to integrate all these aspects into a common goal.

This is certainly high risk research, which still requires a strong academic load. Many aspects have to be coordinated properly, each of which requires independent success. It is not ready for mainstream ICT.

5. Communities

This topic would address different communities, such as neuroscientists, computational neuroscientists and (vision) model developers, hardware engineers and electrical engineers, specific chip developers, machine learning engineers, and system integration engineers. However, despite the strong academic flavour, marketing could also be of interest both in the short and long term, since partial solutions may find new markets.

Matthias S. Müller,
TU Dresden

1. Challenge / Motivation

What are the characteristics of the challenge? What is the motivation behind it? What is the contribution to transforming ICT and ICT research with respect to energy efficiency? Will it profit from new insights gained through multidisciplinary perspectives?

In the last two decades of High Performance Computing a rapid grow of performance was made possible by the driving force of Moore's law and the increase of parallelism. The performance reported in the TOP500 list almost doubled every year, resulting in a performance increase of three orders of magnitude in just 11 years. The continuation of that growth would result in 1.000.000 times the performance in 22 years. Many technological break-through are required to achieve this performance increase. Most of the current studies and initiatives focus at the Exascale level, a 22 years timeframe has a focus on Zettascale systems.

- Already today, the energy consumption of large scale HPC systems is a dominating factor of the TCO. If the current trend continues there will only a few sites capable to install, power and finance the operating costs of HPC systems. Current installations have power consumptions in the range of a few MW. A rough estimate for a power consumption that would allow operating a reasonable number of HPC systems is about 20 MW or 50 Gigaflops/Watt compared to 50-500 Megaflops/Watt of today's systems.

- The International Exascale Software Project (IESP) focuses on software for exascale systems and has a timescale of 11 years. Power Management was identified as a crosscutting consideration affecting all software components.

- The International Technology Roadmap for Semiconductors (ITRS) has a focus on semiconductor technology and maintains a roadmap for the next 15 years. The ITRS statement "computation performance, in some suitable metric, must be increased by one-to-two orders of magnitude by 2020" clearly demonstrates that HPC is in need for much more significant energy reductions, since a reduction between two and three orders of magnitude is required.

2. Suggested Approach / Specific Topics

What are the foundational transformative research topics? What approaches can be suggested? List research topics that would be relevant and together bring ICT forward with respect to energy efficiency.

- New memory technologies are required to improve performance, maintain a reasonably Byte/Flop ratio and to reduce the substantial energy dissipation inherent in current memory technologies

- New methods of computing should be investigated approaching the physical limits of computation. Research should be conducted to analyze these limits for potential implementation, i.e. a thermodynamic analysis of limits in which the physics of an emerging research implementation in the Boltzmann-Heisenberg limit is studied. Examples include carbon-based nano-electronics, spin-based devices, ferromagnetic logic, atomic switches, and nano-electro-mechanical-system (NEMS) switches [see ITRS 2009], but may also be extended to cover quantum computing.

- A new class of algorithms is required for at least a subset of the approaches.

- In general suitable metrics taking the energy consumption into account should be comprehensively introduced to drive the further development in all areas of HPC. This metric needs to be at least partly application specific, going beyond the simple Flops/Watt that is sometimes used today.

3. Target Outcome / Expected Impact

What would be the specific outcome of research supported in this area? What would be the expected impact of a breakthrough in this area in a time frame of 10 to 20 years? What are the success measures / indicators?

- The need for higher performance introduced parallelism into computing. The HPC community developed corresponding methods that are now mainstream and in wide use. Similarly, the need for reducing energy consumption could drive technology developments in hard- and software required to achieve the aggressive goal of 20MW for even the largest systems. The HPC community demonstrated that it can implement even radical changes if it is required to achieve the goal.

- Maintaining the exponential growth of performance at the same rate with a constant power budget is the ultimate goal and also the success indicator.

4. Suitability for ICT and FET / Long-term Vision

In what way and how far does it go beyond the state of the art? What makes this suitable for ICT-FET as opposed to mainstream ICT? Is it vision-driven and high-risk, embryonic or foundational? Is there already a (pre-)established research road map for this topic(s)? What would be the added value of FET funding?

- Radical changes in the way HPC is performed need to be implemented to achieve the goal of exponential performance growth within a constant power budget. Comparisons with the more conservative targets set by ITRS clearly show that it is a high risk endeavor.

- FET funding could provide the incentives for research and industry to analyze high risk approaches, rather than just maintaining a competitive edge.

5. Communities

Is this topic addressing an existing or new community? What areas of expertise are to be involved? Is it a multidisciplinary area, and if so is this a new or an existing mix of disciplines? What would be the critical mass of European researchers needed to carry out research in this area? What is the current situation?

- The required research is multidisciplinary and needs to involve computer-scientists, physicists, engineers, mathematicians and application domain experts. The HPC community in Europe is already multi-disciplinary but the contribution of other research communities is clearly required, e.g. from fields like embedded/mobile devices or new approaches of computing.

Amy Murphy,
Bruno-Kessler-Foundation, Trento

Motivation

While some computing domains are traditionally identified as large consumers of energy, e.g., data centers, others are explicitly focused on energy savings, e.g., wireless sensor networks and programming for other resource constrained devices. Although it is unlikely that the techniques from the resource-constrained world can be applied directly in the unconstrained world, some of the principles can be applied, yielding solutions with lower energy consumption.

Target Outcome

Proposals can focus on any aspects of resource constrained computing from novel hardware to resource-aware system architectures, where the computing costs are not measured in monetary terms but rather in environmental impact caused by the use of the computing infrastructure.

One interesting open issue is if these novel, energy-aware solutions can co-exist with existing systems, and under what conditions. Interestingly, solutions in closed environments, such as data centers, would not be constrained to such co-existence, and may prove to be fruitful test scenarios.

Expected Impact

- The results are likely to lead to distributed computation systems with reduced energy demand with respect to traditional systems.
- This, of course, will likely impact performance, but the trade-offs must be carefully studied and exposed to the system designers in the process.

David Pérez,
Innaxis, Madrid

1. Challenge / Motivation

Challenge 1 (CREATE.NET)

Build a carbon-neutral Internet.

Motivation: while ICT today is not contributing widely to the global GHG figure, its impact is increasing. The question to be asked is the following one: if we had to re-design the Internet from scratch, and if our objective was to make it carbon-neutral, how would we design it? This involves re-thinking the architectural principles of the Internet as we know it, relying on contributions from networking, software systems, distributed computing, complex systems etc.

Challenge 2 (INX)

Green P2P Networking Architectures

Motivation: P2P is the largest contributor to the internet traffic as of 2009. P2P architectures, initially designed for sharing resources (storage, network bandwidth optimisation but also computing power) does not include any green-driven design criteria.

Challenge 3 (INX)

Energy efficiency-driven Clouds / MetaClouds

Motivation: Cloud Computing is an emerging trend that is expected to deliver solutions to some Green concerns (e.g. power limitations and power optimisations in data centers).

The technical definition of the interoperability of clouds and meta-clouds maybe done through the main criteria identified *(1)* but there is no green driver in the agenda. Cloud interoperability does not take carbon footprint as an input.

Main reasons why data centers are limited in growth is due to the restrictions bringing enough power to grow - cloud computing technology is to make data center operations greener, but not clear if the requirement is there when designing the cloud or implementing interoperability criteria.

(1) Current priorities in CC:

Legal issues, Economical aspects, systems management, data handling, programming models, scale and elasticity.

2. Suggested Approach / Specific Topics

Challenge 1 (CREATE.NET)

Build a carbon-neutral Internet.

Topics:

- New architectural approaches to communications and networking

- Protocols able to self-adapt their behaviour depending on the operational conditions to minimize energy consumption
- Low-power switches and routers
- Energy-optimized congestion control
- Energy-optimized data centers
- Distributed real-time monitoring of network power consumption Approaches:
- Beyond end-to-end architectures: contentcentric networking
- Biologically-inspired content dissemination strategies
- Optimal design of protocols with energy as an objective and not as a constraint

Challenge 2 (INX)

Green P2P Networking Architectures

New architectural approaches to green optimisation of P2P networks, which will include understanding of green impact of all clients providing resources to the distributed network. Green Peer Governance would provide governing mechanism to the network allowing a minimal consumption of resources without relaying in devices that are permanently connected (e.g. enabling wake on P2P traffic) looking into new complex power-governing paradigms for all devices that need network resources.

Challenge 3 (INX)

Energy efficiency-driven Clouds / MetaClouds

Cloud Computing and specifically meta-clouds are in the early stages as a research topic. Few initiatives are offering cloud computing as a service.

Pushing ICT research to incorporate green design from the early stages of development, understanding what it takes to adopt green requirements at the same level that other functional requirements (e.g. scale or elasticity).

3. Target Outcome / Expected Impact

Challenge 1 (CREATE.NET)

Build a carbon-neutral Internet.

A novel architecture and a suite of protocols able to set the foundations for a future carbonneutral internet by reducing by 1000 times the power consumption of the current Internet.

Challenge 2 (INX)

Green P2P Networking Architectures

Mechanisms to reduce the impact of those connected devices (mainly computers, but different devices will be connected to P2P networks in the future) that generate 50% of the traffic of the internet.

Challenge 3 (INX)

Energy efficiency-driven Clouds / MetaClouds

Impact on cloud computing research and development (hard to predict?).

4. Suitability for ICT and FET / Long-term Vision

Challenge 1 (CREATE.NET)

Build a carbon-neutral Internet.

Current approaches are mostly related to:

- Optimize the power consumption of data centers
- Use of energy harvesting techniques for feeding embedded ICT devices
- Devise suitable techniques for reducing power consumption at the device level.

Our approach takes a system-level, holistic viewpoint, and aims at optimizing the energy consumption of the Internet as a whole. The architectural aspect in particular makes it suitable for a FET initiative: a change in the basic Internet architecture can take place only over long time horizons (>5 years), making it unsuitable for mainstream ICT.

The GreenTouch initiative (see below) constitutes a first approach towards this problem, but it focuses on information transmission rather than on the network as a single system

Challenge 2 (INX)

Green P2P Networking Architectures

Current P2P networks do not take into account green drivers. Due to the expectations of a high number of devices connected to the internet and increased requirements in bandwidth use, it is possible that P2P traffic keeps increasing.

While some initiatives may exist at application level, the vision for a distributed P2P architecture that operates under green drivers goes beyond current ICT initiatives and thus, could be suitable for a FET vision.

Challenge 3 (INX)

Energy efficiency-driven Clouds / MetaClouds

As said before, Cloud

Computing is starting as a research topic (future ICT Challenge?), however, green clouds or meta-clouds will not be in the agenda until a first generation of clouds have proven their business models, which might be too late to impose green criteria on the design.

Today requirements for cloud interoperability and metaclouds are on the performance side (not enough standards develop and bandwidth to transfer clouds) and therefore, providing a vision for green cloud approaches could be foundational for the future.

5. Communities

Challenge 1 (CREATE.NET)

Build a carbon-neutral Internet.

Areas of expertise:

- design of low-power electronics
- distributed network architecture
- content-centric networking
- distributed network monitoring and management
- optimization and control
- complex systems/networks

Challenge 2 (INX)

Green P2P Networking Architectures

Distributed network architecture design Complex Systems

Challenge 3 (INX)

Energy efficiency-driven Clouds / MetaClouds

- Cloud Computing expertise.
- Network architecture.

Some interesting references

- SusteIT project .
- GreenTouch

Jean-Marc Pierson,
University of Toulouse

1. Challenge / Motivation

- Multi-level mechanisms exist or coexist to reduce energy consumption of ICT while guaranteeing performances:
    - Design and production
    - Operational machines: mainly hardware possibilities and consolidation approaches for software stacks, optimization problems
    - Operational environment: air-conditioning, building

→ Lack of modelling and integration, isolated views

2. Suggested Approach / Specific Topics

- Enforce the modelling of energy consumption
    - Develop finer energy meters
    - Develop mathematical models of energy consumption (different levels, and aggregated to have holistic view)
- Develop standards for benchmarking hardware and applications w.r.t energy consumption
- Develop integrated performance-guaranteed solutions easy-to-use (transparent, hidden)
- Integrate energy cost and energy footprint

3. Target Outcome / Expected Impact

- Energy-aware usage of ICT
- Energy-controlled usage of ICT
- A not-increasing ICT impact: do more with the same energy footprint
- Success indicated by users' satisfaction (QoS) w.r.t wasted energy

4. Suitability for ICT and FET / Long-term Vision

- FET could help to integrate diverse views:
    - Hardware
    - Software
    - Operational environment (a/c, building, …)
    - Economic
    - Environmental concerns (quality of energy)

- Energy efficiency must not be limited to one side of the problem

5. Communities

- Multidisciplinary teams:
    - HW, SW, MW (architectures, systems, optimization, large scale systems)
    - Mathematicians, Electrical/Electronical Engineers
    - Energy providers
    - Economists

Existing efforts

- COST IC0804 : Energy efficiency in Large Scale Distributed Research

Jordi Torres,
Universitat Politecnica de Catalunya - Barcelona Supercomputing Center

1. Challenge / Motivation

1.1 Challenge

Transform ICT energy efficiency into a more formal and scientific endeavor, establishing a holistic theoretical model to a better understanding, and represent power consumption of the ICT components and the interaction of them.

ICT are challenged to rethink the resource/service management strategies, adding energy efficiency to a list of critical operating parameters that already include availability, reliability and performance.

1.2 Motivation

Existing models are very restricted, normally concerned with a single component, and embody many simplifications. Clearly it is needed the development of a better understanding of the interactions between heterogeneous components of the current complex ICT systems and to formalize the nature and behaviour of these components.

2. Suggested Approach / Specific Topics

2.1. Suggested Approach

1. **Try any see out of the ICT world in itself and think globally**, considering smartgrid or the economic aspects and economic management of ICT in a global way (e.g. move Data Centers next to Power Plants, and move the data through the optical fiber).

2. **Identify global metrics** that allow a global evaluation base on our partial metrical existing system. The metrics themselves should be measurable. Remember that without measures we can't evaluate any advance. Also, let's remember that isolated metrics can produce a crash. To create appropriate standards should be done.

3. **To build abstract and strong models** that will survive the changes in technology. That means that models should be valid when computer systems evolve over time. This models should capture the most important factors of the systems while allow abstract reasoning. The models can be composed by different levels in an abstraction hierarchy that includes power-performance tradeoffs at multiple levels of granularity.

4. Create more powerful **formal optimization techniques**. The models proposed will allow formalizing behaviors and interactions that help the use of more sophisticated (current approaches to power management rely on simple algorithms based on heuristics) formal optimization techniques that can be inspired in the ones used in other areas.

5. A new **Autonomic Computing cores** are required. Today the optimizations at different system levels interfere with each other. This makes the behavior of the current ICT systems unmanageable in execution time. This requires novel optimization techniques that implements self-* properties at run time that include vertical and horizontal dialog between the different components.

6. **Encourage integrative efforts** between different research areas to solve the challenged expressed in the previous points. Often researchers in each of these areas are often isolated from each other which make collaboration difficult. Further collaboration among researchers from different disciplines will permit the development of the challenges outlined above.

2.2. Specific Topics

The Challenge proposed can be addressed at all levels of ICT systems design such as:

- System architecture: A proper architecture design can have a significant impact on the energy requirements.
- Algorithm design: For many software engineers it is new space to consider the energy footprint of an algorithm as a relevant quality parameter when they program an algorithm.
- Runtime and Compiler: The object code can have a decisive influence on the power consumption during execution.
- Middleware design: Middleware are challenged to rethink the service management strategies.
- Operating System design: Energy-aware integrated resource management has to be supported by new operating systems.
- Hardware design: Some new evolving trends in hardware design can further improve the energy efficiency.
- Network design: the same issues pointed in the previous hardware and OS paragraphs should be considered.
- Storage: The volume of stored data continues to increase. The I/O has a considerable influence on the power consumption.
- Data Centers: This is related with all the previous issues in a comprehensive coverage. Also could be considered because of their impact, issues such as cooling systems, server density, operational issues, etc.
- …

And also out of the ICT area such as it is the supply of energy or business models of the ICT infrastructures (e.g. nowadays, management oriented to performance objectives is progressively replaced by management oriented to Business Level Objectives (BLOs), that brings the providers the opportunity to improve the economic performance of the sales of their resources).

3. Target Outcome / Expected Impact

3.1. Target Outcome

New models that cross today research areas boundaries to understand and optimize the properties of the current systems in terms of power, energy and thermal. This will lead to new ways of designing hardware, middleware, protocols, management models, etc. that support and enable energy efficiency.

3.2. Expected Impact

It is expected to have a bigger adoption of this kind of technologies. The expected impact of a breakthrough in this area in a time frame of 10 to 20 years is the reduction of a use of energy and to avoid an energy shock, in consequence the reduction of the emission of CO_2. This can make easier the transition of the research into products and services that will improve the economic competitiveness of EU. It is expected that research performed on this area will lead to capabilities beyond the energy consumption of ICTs as we know them today.

4. Suitability for ICT and FET / Long-term Vision

The approach proposed goes beyond current state of the art of power management models that currently rely on simple algorithms based on heuristics and are very restricted. But the goal that we are seeking for is not easy to describe, because of the current and evolving complexity of ICT systems.

A very important effort has to be done to define the methodology to make this process possible, and it is not evident right now how and we will get there. Furthermore, the proposed holistic modeling for sure will require prototyping components at different levels (software, middleware, hardware, communications, ….) beyond current existing components. This exploratory research fits with the FET funding program.

5. Communities

Teams composed by researchers coming from different disciplines whose expertise are referred in section 2.2.

Annex I: Terms of Reference

Disruptive Solutions for Energy Efficient ICT
Expert Consultation Meeting

8-9 February 2010, Brussels

1. Context

Future Emerging Technologies (FET) Proactive is a part of the Information and Communication Technology (ICT) Programme of the European Commission. FET acts as a pathfinder for the ICTs of the future supporting foundational long-term research and technological innovation and by fostering the emergence of new European research communities in ICT.

FET addresses evolutionary and revolutionary approaches through multidisciplinary cooperation. It explores novel future technology options, identifies new drivers for research, and brings new science into ICTs to expand their scientific foundations. To that aim, it supports radically new and catalytic ideas expected to have transformative impact on society and technology.

FET Proactive structures research in a number of proactive initiatives, which typically consist of a group of projects funded around a common theme. The themes are shaped through interaction with the research community and focus on the mission outlined above. Useful pointers to further information about FET proactive are provided in the Annex.

2. Objectives of the Consultation Workshop

It will be a clear future priority world wide to decrease mankind's carbon footprint. According to the SMART2020 study, the share of ICT on the world wide energy consumption today is in the range of 2-5%. Given that the use of ICT will further increase and the overall energy consumption will hopefully decrease due to the help of ICT and other measures, it is expected that the share of ICT on the world wide energy consumption will grow in the future. Hence, it becomes more and more important to consider and improve the energy efficiency of ICT. On the short term, it will be an obvious and practical solution to exploit better the potential of technologies that already exist or are currently in the making. On the long term, new and disruptive ideas will be needed, and we must start to search for those ideas already now.

In this context, the meeting "New and Disruptive Solutions for Energy Efficient ICT" will bring together visionary scientists with diverse expertise to brainstorm and elaborate where and in what ways FET research could best contribute to an increased energy efficiency of ICT, how FET like research could unveil these required new and disruptive ideas.
The intention of this workshop is not to focus on short term solutions, but to concentrate on solutions that cannot be reached with traditional technological approaches and which may be pursued in a possible FET proactive initiative within the ICT Work Programme 2011–12 and later. The meeting will take place in Brussel on the 8th and 9th of February 2010. Participation is by invitation only.

Participants are invited to share their views and work together to identify key research challenges addressing the energy efficiency of ICT. The discussions should reflect on different aspects of ICT (software, hardware, communications,…) at different levels (devices, components, system architectures …) individually and from an integrated point of view.
The brainstorming and discussions will be open and unconstrained and is expected to lead to the emergence of radically new ideas inspired by the cross-fertilisation of different aspects and backgrounds. A rapporteur will write down draft conclusions, which will be subsequently consolidated using participants' feedback.

Participants are invited to submit by **January 28th, 2010,** using the template provided, their ideas on key scientific challenges addressing the energy efficiency of ICT and to present them at the meeting within 5 minutes maximum. Please note that we are looking for candidate topics and challenges at the research programme level rather than at the level of single project ideas.

Please look at the FET website http://cordis.europa.eu/fp7/ict/fet-proactive/ to learn more about us. An overview of FET Proactive Initiatives is given on the website http://cordis.europa.eu/fp7/ict/fet-proactive/areas_en.html.

You also invited to start the discussions already in our CaFETeria beforehand. For this please follow this link: http://cafeteria.ning.com/forum/topic/listForContributor?user=1hzopwgkceh44
