Disruptive Solutions for Energy Efficient ICT

Expert Consultation Workshop
8 & 9 February 2010

Report

Future and Emerging Technologies
Proactive

European Commission
Information Society and Media

LEGAL NOTICE
By the Commission of the European Union, Information Society & Media Directorate-General, Future and Emerging Technologies Proactive unit.

Table of contents

Executive summary

Some eight orders of magnitude separate the energy efficiency of conventional computers from what is theoretically possible. Closing this gap would lead to a significant improvement in the energy efficiency of information and communication technology (ICT).

Experts from a variety of disciplines held wide-ranging discussions about how to achieve this goal at the Disruptive Solutions for Energy Efficient ICT meeting in Brussels on 8-9 February 2010.

Participants divided their ideas into two broad categories. The first focuses on device-related approaches that examine the physical problems and conceptual challenges facing scientists hoping to better understand the limits of computing and exploit them to achieve better efficiency. The second category focuses on computer systems and architectures and how disruptive changes in these areas would lead to transformational improvements in energy efficiency.

An FET Proactive Initiative focused on these areas would engage a wide range of multidisciplinary research communities and could yield significant improvements in the energy efficiency of ICT.

1. Introduction

Future Emerging Technologies (FET) Proactive is a part of the Information and Communication Technology (ICT) Programme of the European Commission. FET supports catalytic ideas expected to have transformative impact on society and technology.

The Disruptive Solutions for Energy Efficient ICT meeting in Brussels brought together a diverse group of experts to brainstorm and elaborate where and in what ways FET research could improve energy efficiency in ICT. The goal was to look for disruptive ideas with a high potential pay-off that cannot be reached with traditional funding instruments.

2. Synthesis of discussions

The experts each presented the research challenges relating to ICT energy efficiency, the approaches that might bring about transformational change, the communities that would need to be involved and how these could be addressed with an FET Proactive initiative.

The presentations covered a wide range of topics. A number touched on the fundamental limits of computing efficiency and how current computers stand in relation to this limit. Participants heard about the approaches of Von Neumann in the 1920s, Landauer in the 1960s and Levitin and Margolis in the last decade [Capra, Gammaitoni, Levine]. These approaches gave vastly different estimates depending on the nature of the computation involved. Participants discussed in detail how these differences arose and a consensus was reached that the energy consumption in modern computers differed from the theoretical limit by some eight orders of magnitude.

How to close this gap was a significant theme throughout the meeting. Participants outlined various exotic computing paradigms that have the potential for greater efficiency than the binary switching used in conventional computers. These included reversible computing, non-Boolean logic and stochastic computing. They agreed that understanding and evaluating these paradigms was an important goal.

Many participants spoke of the need to develop robust models of computing systems that can be used to study the trade offs between power and performance and help to optimise them [Gammaitoni, Glück, Kaxiras, Pérez, Pierson, Torres]. A corollary to this was the issue of metrics and benchmarks for evaluating and comparing the energy consumption of computing systems [Müller, Torres] and this generated some controversy.

There was general agreement that metrics and benchmarks will be important for future energy efficient ICT efforts. However, the view was also held that their development did not constitute transformative research and so fell outside the remit of an FET Proactive initiative.

The way natural systems store and process information was also discussed. Participants heard that the human vision system is particularly efficient at pattern recognition because of the efficient coding, processing and transmission of information [Linares-Barranco] it exploits. They discussed the possibility of exploiting the information processing ability of biomolecules to create living computers that self-replicate and evolve [Jaramillo].

The human vision system is an example of a task-related computing system that is both highly specialized and highly efficient. Participants discussed the potential for task-related computing in other areas such as sensor networks.

The group heard that changes to software, architecture and computer systems could make important contributions to energy efficiency. For example, improving the energy transparency between software layers could help to make energy consumption a first order parameter in software design [Brorsson, Murphy]. In discussing this issue, participants heard that Intel chips already give developers various ways to control energy consumption, most of which are ignored because of the difficulty in using them. This raised the question of how best to design and implement energy transparency measures.

Participants also heard that the ICT industry well focused on improving the energy efficiency of their products. One danger is that FET research could end up having little impact on these efforts if it is not carefully targeted [Müller].

Following extensive discussion about these presentations, participants agreed to divide their ideas into two broad categories. The first focuses on the physical problems and conceptual challenges facing scientists hoping to better understand and exploit the limits of computing. These come under the heading “Device-related challenges” in section 3 where the discussions that motivated each challenge are outlined.

The second category focuses on approaches related to computer systems and architecture. These are outlined and discussed in section 4.

Sections 5 and 6 outline the potential impact of these ideas and the communities that will have to be engaged to achieve them.

3. Device-related challenges

3.1 Losses of energy in elementary devices such as transistors, biomolecules or atoms are due to computing, storage and input and output functions. Another source of losses comes from the communication between elementary devices of the same kind and interfaces between devices of different kinds. New devices and communication mechanisms are needed with improved function per energy unit.

The elementary building blocks of computers are transistors: switches which can process and store data. The workhorse device in today's computers is the field effect transistor made using CMOS technology. This switches on and off extremely quickly but generates large amounts of waste heat. Significant energy losses also occur when information is sent from one transistor to another, when it is sent between parts of a chip with different functions such as from computing to storage, and when the data is addressed and read out to other devices, such as displays [Van Poucke].

The group heard that the ICT industry is already devoting considerable resources to improving the energy efficiency of conventional computing devices which means that FET funding has to be carefully targeted [Müller].

The group discussed a number of other computing mechanisms, such as exotic switches consisting of simple atoms or large biomolecules [Remacle], which have the potential to be significantly more efficient than field effect transistors. Post-Boolean logic in which data is represented in higher bases also has the potential to be more energy efficient [Remacle, Levine].

Another energy saving possibility is to send data in the form of spike-like solitons which mimic the way neurons carry information. The advantage here is that only a small section of the transmission line is charged rather than its entire length as in conventional devices [Linares-Barranco].

In addition, nano and microelectromechanical components have the potential to reduce the energy consumption of wireless radio devices by an order of magnitude [Dussopt].

3.2 New computational means beyond switching should be considered, such as noise driven, stochastic and mechanical computation.

While modern computers are built from deterministic switches, other computing paradigms beyond this are also possible, such as noise driven logic gates in which noise could be the information carrier and the power driver [Gammaitoni]. Efficient mechanical computation may also soon be possible thanks to the recent advances in nano and microelectromechanical device technologies. The improvements in energy efficiency that are possible with these and other approaches should be investigated.

3.3 Conceptual challenges: understanding energy dissipation, thermodynamic/quantum physics limits, computing at finite power or at finite rate, reversible computing and information dissipation.

The pioneering work of Von Neumann, Landauer and Margolis and Levitin place fundamental limits on the energy efficiency that is possible in computing devices [Capra, Gammaitoni, Levine]. But there are still many fundamental areas where our conceptual understanding needs to be improved. For example, it may never be possible to reach the theoretical limits of computing efficiency but an important but unanswered question is how close is it possible to get.

The nature of energy dissipation is another area where our understanding is incomplete. Thermodynamic laws place important limits on the energy efficiency of macroscale ICT devices but it is not clear how, or even whether, these limits apply to nanoscale systems where statistical thermodynamic laws may not be appropriate [Gammaitoni]. Neither is it clear what the limits are in the non-adiabatic, non-static regime [Levine].

One type of energy loss occurs when information is discarded, in which case it dissipates as heat. This process is endemic in modern computing: data is routinely overwritten and discarded. But the destruction of data is not intrinsic to computation and can be avoided or at least minimised [Glück]. One such technique is reversible computing, in which every computation can run in reverse, preserving the information it handles. Avoiding the unnecessary destruction of information, and the waste heat this generates, could lead to substantial increases in efficiency.

3.4 Ways to recover lost energy or harvest ambient energy are possibilities.

Since the dissipation of some energy is unavoidable, steps should be taken to recover it. Such techniques are closely linked to the ability to harvest ambient energy. Examples include nanoelectromechanical devices which have the potential to harvest ambient energy and make ICT devices self sufficient [Dussopt]. The group noted that the Towards Zero-Power ICT initiative has similar goals.

4. Computing architecture and systems-related challenges

4.1 Energy savings are possible by developing and exploiting systems that are inherently massively parallel.

In 2009, IBM's BlueGene/P supercomputer simulated the function of a cat's brain consisting of 1.6 billion neurons, albeit at a rate 600 times slower than real time. BlueGene/P has 146,456 processors with 147,000GB of memory and requires a 1.4MW power supply. By contrast, the human brain with 25 billion neurons, runs on a mere 20W [Kaxiras, Linares-Barranco]. Among the keys to such energy efficiency is massive parallelism, which allows computations to proceed more slowly and therefore more efficiently. The group agreed that understanding and exploiting systems that are massively parallel will yield substantial energy savings.

4.2 Uncertainty, randomness and unreliability are a result of the device properties, for example when running CMOS at low voltage. New computational models are needed that take this into account.

One of the consequences of computing at low power is that noise plays a greater role. The noise comes from sources such as thermal fluctuations and causes errors and uncertainty in calculations. Noise will also be an issue with future nanoelectronic CMOS technologies in which device properties will show signification variation. New computational models that deal with noise efficiently will be needed if low-power computing is to be exploited to its fullest extent [Benini].

4.3 Novel computing paradigms are needed and their energy efficiency evaluated. Examples include efficient information coding, biorelated computing, algorithm and system complexity management, reversible computing and post-Boolean logic.

Almost without exception, conventional computers run on binary digital code using a small set of architectures and programming languages. And yet a number of other, more efficient computing paradigms are known.

For example, the information coding mechanisms employed by the human vision system allow the brain to recognise objects and react to them with much greater speed than conventional computers and using less power [Linares-Barranco]. Biological systems that process information encoded in biomolecules may lead to self-replicating and evolvable computers [Jaramillo]. And computers that exploit non-Boolean logic [Remacle] or reversible software [Glück] have the potential to operate highly efficiently. It is important to determine exactly how efficient these systems can be and in what circumstances.

4.4 Software models and programming methodologies are needed that focus on the trade-off between energy and performance/precision, taking into account efficiency metrics beyond number of operations per Watt.

One way to save energy in computing systems and sensor networks is to turn off parts of systems that are not being used, such as radio communications during idle periods. But this inevitably leads to a trade off against the performance of the system. For example, the time it takes to switch on devices when they are needed can reduce the latency of a system [Murphy]. Maximising the benefits of such trade offs requires new ways of modelling a system's performance and activity [Pierson, Torres]. Smart software will also be needed that understands the application and can determine when these trade offs can be made. This will require new ways of measuring efficiency that go beyond a simple measure of the number of operations per Watt.

This kind of approach will be important for the creation of a carbon-neutral internet, green P2P networking architectures and energy efficient Clouds [Pérez].

The group discussed the need for a new generation of metrics and benchmarks for measuring and comparing energy efficiency [Müller, Pierson, Torres]. It agreed that such benchmarks will be an important part of future efforts to study efficiency. However, the view was also held that the development of these would not constitute transformative research and so fell outside FET Proactive's remit.

4.5 Whenever possible, the research should consider multi-level approaches to the design and modelling of systems for low power consumption.

One important way to achieve better energy efficiency will be to optimise computing systems at all levels in their hierarchy, ie by considering the hardware, middleware and software together. Such an approach would benefit from advances such as power aware interfaces between these levels for better monitoring and control of power [Brorsson].

4.6 Energy-efficiency gains may be possible with new networking architectures and communication schemes.

The group believes that new networking architectures and communication schemes should lead to substantial energy savings but decided that it lacked the expertise to make specific recommendations.

4.7 Significant gains can be made by working in both areas together.

While there are important gains to be made from looking at either hardware or software systems, the group felt that it would be possible to make important contributions by tackling both areas together. An example is task-related computing in which hardware and software are optimised for a specific task such as pattern recognition and machine vision.

5. Impacts

The initiative is expected to yield significant improvements in the energy efficiency of ICT. Estimates of the reduction in energy use range from 50 per cent [Brorsson] to several orders of magnitude [Kaxiras, Murphy]. One important impact will be the exponential increase in the performance of supercomputers with a constant power budget [Müller]. Greater efficiencies should also make possible novel low-power devices with new applications, such as more capable ICT able to operate inside the body [Gammaitoni], energy and information efficient vision systems [Linares-Barranco] and energy harvesting devices that require zero power input.

6. Communities

The initiative will address a multidisciplinary community including:

- computer scientists
- hardware, software and network engineers
- physicists
- neuroscientists
- biochemists
- mathematicians
- economists

To compete with teams in the US and Canada, the critical mass of European researchers would need to be a network of about 10 teams, each consisting of 6-10 researchers [Glück].

7. References (in alphabetical order)

"Benini": from the written contribution of Luca Benini
"Brorsson": from the written contribution and presentation of Mats Brorsson
"Capra": from the written contribution and presentation of Eugenio Capra
"Dussopt": from the written contribution and presentation of Laurent Dussopt
"Gammaitoni": from the written contribution and presentation of Luca Gammaitoni
"Glück": from the written contribution and presentation of Robert Glück
"Jaramillo": from the written contribution and presentation of Alfonso Jaramillo
"Kaxiras": from the written contribution and presentation of Stefanos Kaxiras
"Levine": from the written contribution and presentation of Raphaël D. Levine
"Linares-Barranco": from the written contribution and presentation of Bernabé Linares-Barranco
"Müller": from the written contribution and presentation of Matthias S. Müller
"Murphy": from the written contribution and presentation of Amy Murphy,
"Pérez": from the written contribution and presentation of David Pérez
"Pierson": from the written contribution and presentation of Jean-Marc Pierson
"Remacle": from the written contribution and presentation of Françoise Remacle
"Torres": from the written contribution and presentation of Jordi Torres
"Van Poucke": from the presentation of Bart Van Poucke

Annex I: List of Contributors

Mats Brorsson	KTH School of Information & Communication Technology, Stockholm (SE)
Eugenio Capra	Politecnico di Milano (IT)
Laurent Dussopt	CEA - LETI, Grenoble (FR)
Luca Gammaitoni	NiPS Laboratory, Università di Perugia (IT)
Robert Glück	University of Copenhagen (DK)
Christina Herzog	University of Toulouse (FR)
Alfonso Jaramillo	Genopole®-CNRS UPS3201 & Ecole Polytechnique, Evry (FR)
Stefanos Kaxiras	University of Patras (GR) / Uppsala University (SE)
Raphaël D. Levine	The Hebrew University of Jerusalem (IL)
Bernabé Linares-Barranco	National Microelectronics Center, Sevilla (ES)
Matthias S. Müller	TU Dresden (DE)
Amy Murphy	Bruno-Kessler-Foundation, Trento (IT)
David Pérez	Innaxis, Madrid (ES)
Jean-Marc Pierson	University of Toulouse (FR)
Bart Van Poucke	IMEC (BE)
Françoise Remacle	Université de Liège (BE)
Jordi Torres	Universitat Politecnica de Catalunya - Barcelona Supercomputing Center (ES)

Additional Contributors

Luca Benini	Università di Bologna (IT)
Chiara Francalinci	Politecnico di Milano (IT)
Sven Rogge	TU-Delft (NL)

European Commission Participants:

Mario Campolargo
Wolfgang Boch
Pekka Karp
Ralph Stübner
Rüdiger Martin
Werner Steinhoegl

Rapporteur:

Justin Mullins	Science writer (UK)

Annex II: Terms of Reference

Disruptive Solutions for Energy Efficient ICT
Expert Consultation Workshop

8 & 9 February 2010

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.
Further information about FET proactive initiatives: http://cordis.europa.eu/fet-proactive

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.

European Commission
Directorate-General Information Society and Media
Unit F1 – Future and Emerging Technologies – Proactive

E-mail: infso-ictfet@ec.europa.eu

FET Proactive in the 7th Framework Programme:
http://cordis.europa.eu/fet-proactive

European Commission
Information Society and Media