ICT Energy Strategic Research Agenda

October 2016

1 Executive Summary

i. ICT Energy is a community covering ULSI, microarchitectures, high performance computing (HPC), energy harvesting, thermoelectrics, energy storage, ICT system design, embedded systems, efficient electronics, static analysis and the physics of computation.

ii. The reliance of society on the use of Information Communications Technology (ICT) devices and systems is increasing with over 4% of all electricity consumption and 2% of all CO_2 emissions now the result of ICT use. If entertainment, telephones, TV and media which are now being translated onto ICT devices and systems are added then these consumption numbers approximately double. FP7-ICT and Horizon2020 have highlighted ICT as a key engine of growth, with the use of ICT to improve energy efficiency by managing energy demand and use. The energy consumption and carbon dioxide emission from the expanding ICT use with present drivers, however, is unsustainable and will impact heavily on future climate change.

iii. The goal of the ICT-Energy consortium is to identify methods, techniques and directions for making ICT devices and systems sustainable with regard to energy.

iv. Communications consumes significantly more energy per bit of information than any logical calculation either inside an ICT system or between ICT systems. ICT through cloud computing and mobile consumption of video is significantly increasing the volume of data communicated between devices. Energy consumption could be minimised if research to find the ideal distribution of cloud services for present and future markets was undertaken. Energy consumption could also be minimised if smart communications minimising the amount of data being transported is used over other techniques. Finally the development and fast deployment of new communications system with low energy consumption per bit to replace legacy systems is essential to circumvent the enormous increase in data volume from cloud and especially high definition video use.

v. There is no transparency at higher levels of the system stack of ICT systems such as computers to allow software writers or hardware designers to understand how changes to software, communications or the delivery of services and applications affect energy consumption in ICT throughout the system stack. Research to find traceable and transparent energy usage throughout the system stack of ICT systems is required. Only once such research has been successfully completed can compilers and software be written to minimise energy consumption. It is clear that poorly written software is responsible for wasted energy in ICT systems and significant CO_2 emissions at present. It is also clear that education is required to understand energy consumption and for the writing of efficient and low energy software code that is disciplined in ways that compilers will generate efficient runtime operations on ICT systems. At present many programmers especially for apps on phones and tablets have little knowledge of how to write efficient software algorithms that minimise energy consumption. The volume of such software and ICT devices provides large potential to minimise energy consumption.

vi. Whilst the direct contribution to CO_2 emission from ICT devices might be small, ICT devices have the potential to contribute significantly to the reduction of CO_2 in transportation, heating (through smart building control) and manufacturing (autonomous sensors).

vii. ICT devices in the form of autonomous sensors have the potential to significantly improve the quality of life of society both in Europe and across the whole world. In particular personalised healthcare where a wide range of sensors enable early detection of disease could transform patient care in the future. When ill your doctor or hospital phones you up and tells you that you need to arrange an appointment to get treatment. Such sensors have the potential to improve the quality of life of society through environmental monitoring (monitoring pollution, CO_2 and O_2) and security monitoring.

viii. Transistors are approaching the minimum amount of energy per switch and the access resistance of electrical interconnects is a key issue for future scaling of energy in all ICT hardware. The failure of the International Technology Roadmap for Semiconductors after 2013 provides evidence of the end of scaling and also of Moore's law, the economic law driving future ICT hardware systems. Future scaling of transistors is therefore unlikely to result in any significant reduction in system energy consumption and power. This is further exasperated by the product volume requirements for return on investment at the leading edge CMOS technology nodes. Radical new devices, interconnect solutions and system architectures are required if reductions in power per device are to be achieved in the future. Technologies that can be delivered on legacy foundries would also significantly reduce the financial costs of adoption. As Europe is the leader in low power design and consumer embedded systems, this is an enormous opportunity for the EC to lead as applications with old wired devices move towards portable solutions.

ix. Energy storage at the large scale for high performance computers or cloud datacentres is an enormous problem preventing the large scale deployment of renewable sources to power large scale computing systems. Radical solutions are required to remove the dependency on diesel generators and high carbon sourced electricity to maintain robust and reliable deployment of computing and communication services.

x. For autonomous systems, significant improvements in energy harvesting and energy storage at the small scale would also provide disruptive solutions to the use of smart sensors for a host of applications in personalised healthcare, environmental monitoring, industrial monitoring, security and transportation. Such applications have the potential for significant improvements in the quality of life and reductions in energy consumption and CO_2 emissions.

2 Contents

1 Executive Summary 1
2 Contents 3
3 The ICT Energy Co-ordination Activity 5
4 European Agenda 5
5 Vision, Aims and Objectives 6
6 Challenges for Sustainable Energy ICT 7
6.1 The Energy Issue for ICT 7
7 Relationships to Other Projects, Agendas and Roadmaps 12
8 Drivers 12
8.1 Application Drivers 12
8.2 Computing System Cost Drivers 14
8.3 Communications Drivers 14
9 The System Stack 15
10 Devices 16
10.1 Present and Future Device Performance 16
10.2 Key Device Challenges 17
10.3 Fundamental Device Limits 17
10.3.1 Subthreshold Devices 19
10.3.2 Interconnects and Electrical Wires 21
10.4 Opportunities for Device Improvement 22
11 Circuit Micro-Architecture 23
11.1 Present and Future Circuit Architecture Performance 23
11.1.1 HPC / Data Centres 24
11.1.2 Portable and Embedded Systems 25
11.2 Fundamental Circuit Architecture Limits 26
11.3 Key Challenges 27
11.4 Opportunities for Circuit Architecture Improvement 30
12 System Architecture 32
12.1 Present and Future System Architecture 32
12.2 Key Challenges of System Architecture 32
12.3 Fundamental Limits of System Architecture 33
12.4 Opportunities for System Architecture 33
13 Software 35
13.1 Present and Future Software Performance 35
13.2 Key Software Challenges 35
13.2.1 Energy Modelling of Software 36
13.2.2 Software Energy Analysis 37
13.3 Fundamental Software Limits 38
13.4 Opportunities for Software Improvement 38
13.5 Applied mathematics and numerical simulations using high performance computing (HPC) 39
13.5.1 Present & future performance of applied maths & numerical simulations using HPC 39
13.5.2 Key challenges of future exascale computing 39
13.5.3 Opportunities for improvement of applied maths & numerical simulations using HPC 40
14 Communications 42
14.1 Present and Future Communications Performance 42
14.2 Key Communication Challenges 42
14.3 Fundamental Communications Limits 42
14.3.1 Wireless Communication 42
14.3.2 Photonics Communications 43
14.4 Opportunities for Communications Improvement 44
15 Power and Power Management 45
15.1 Present and Future Power and Management Performance 45
15.1.1 Power Management 45
15.1.2 Renewable and/or Sustainable Energy Sources for HPC and Cloud Datacentres 45

15.1.3 Energy Harvesting for Autonomous Systems 47
15.1.4 Energy Storage for Autonomous Systems 48
15.1.5 Supercapacitors 49
15.2 Key Power and Management Challenges 50
15.3 Fundamental Power and Management Limits 50
15.4 Opportunities for Power and Management Improvement 50
16 Added Value from Optimising at Multiple Levels of the System Stack 51
17 International Comparisons 52
18 Strengths, Weaknesses, Opportunities and Threats to Sustainable Energy ICT in Europe 52
18.1 Strengths 52
18.2 Weaknesses 52
18.3 Opportunities 52
18.4 Threats 53
19 Conclusions 53
20 Appendices 54
20.1 References 54
20.2 Definition of Acronyms 59
20.3 List of Contributors 60

3 The ICT Energy Co-ordination Activity

The goal of the ICT Energy project is to create a coordination activity among consortia involved in the ICT-ENERGY subject with specific reference to bringing together the existing "Toward Zero-Power ICT" community organized within the ZEROPOWER co-ordinated action (2010 to 2013) and the novel "MINECC" (Minimising energy consumption of computing to the limit) community recently funded under the FET Proactive Call 8 (FP7-ICT-2011-8) Objective 9.8 (2012 to 2015).

Seven research projects, carried out by different consortia have been funded under this call and are currently reaching the end of their activity.

ENTRA – Whole System Energy Transparency – http://entraproject.eu

EXA2GREEN – Energy-Aware Sustainable Computing on Future Technology – Paving the Road to Exascale Computing – http://exa2green-project.eu

LANDAUER – Operating ICT Basic Switches Below the Landauer Limit – http://www.landauer-project.eu

PARADIME – Parallel Distributed Infrastructure for Minimization of energy – http://www.paradime-project.eu

PHIDIAS – Ultra-low-power Holistic Design for Smart Bio-signals Computing Platforms – http://www.phidiasproject.eu

SENSATION – Self Energy Supporting Autonomous Computation – http://www.sensation-project.eu

TOLOP – Towards Low Power ICT – http://www.tolop.eu

The ICT Energy co-ordination activity is aimed at assessing the impact of the research efforts developed in the groups involved in the different consortia (overall more than 30 European research groups) and proposing measures to increase the visibility of ICT-Energy related initiatives to the scientific community, targeted industries and to the public at large through exchange of information, dedicated networking events and media campaigns. The activities of our co-ordination activity will inspire more research projects in this emerging area by generating broader acceptance for the developed technology and the benefits of its applications.

ICT-Energy co-ordination activity will facilitate broader interaction and feedback among the consortia members and stakeholders, thereby, consolidating progress in the field. Positive benefits to the European Community are foreseen in all great challenges of energy, security, environment and health by developing a strategic research agenda in low power, energy efficient ICT and making it happen.

4 European Agenda

The Europe 2020 Agenda defines three priorities at the heart of Europe in 2020:

Smart growth – developing an economy based on knowledge and innovation.

Sustainable growth – promoting a more resource efficient, greener and more competitive economy.

Inclusive growth – fostering a high-employment economy delivering economic, social and territorial cohesion.

The statement on energy states:

Reduce greenhouse gas emissions by at least 20% compared to 1990 levels (see Figure 1) or by 30%, if the conditions are right (provided that other developed countries commit themselves to comparable emission reductions and that developing countries contribute adequately according to their responsibilities and respective capabilities); increase the share of renewable energy sources in our final energy consumption to 20%; and a 20% increase in energy efficiency.

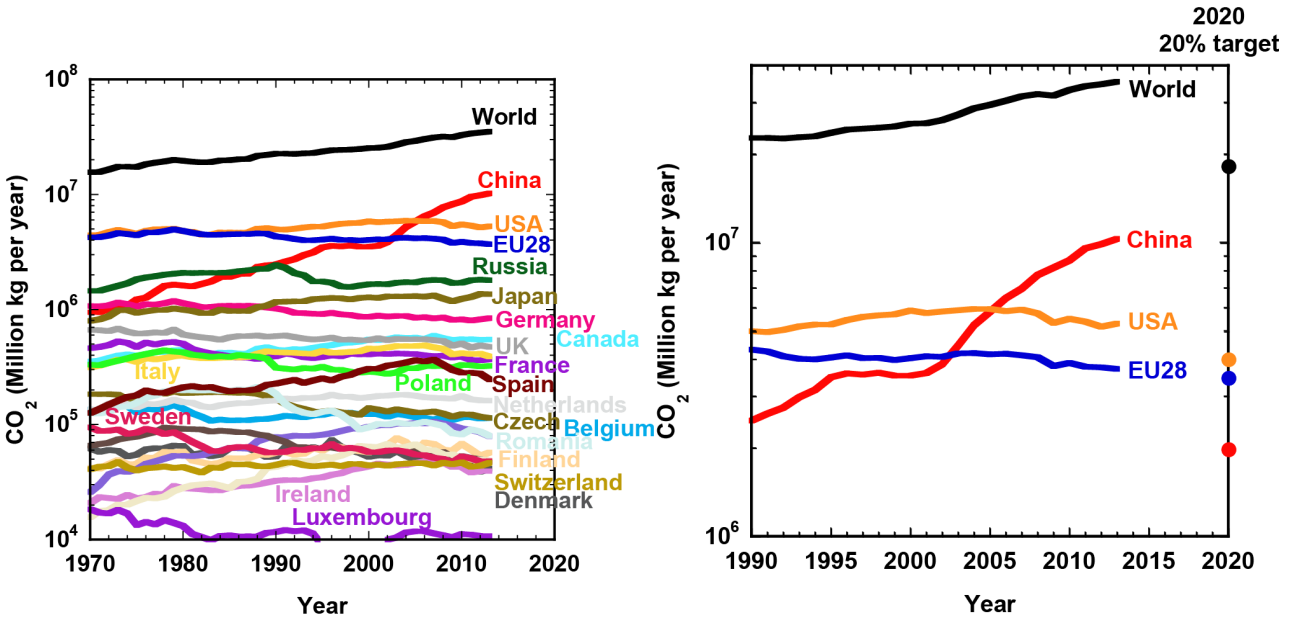

Figure 1: Left – The CO2 emissions per year for EU 28 countries compared to major economies around the world Right – the three major world economies and their proposed 20% reduction of CO2 emissions of 1990 levels for 2020. Source: [EDGAR:2014].

This reduction in greenhouse gas emissions and improved efficiency is whilst undertaking a Flagship Initiative "A Digital Agenda for Europe": The aim is to deliver sustainable economic and social benefits from a Digital Single Market based on fast and ultrafast internet and interoperable applications, with broadband access for all by 2013, access for all to much higher internet speeds (30 Mbps or above) by 2020, and 50% or more of European households subscribing to interconnections above 100 Mbps.

A second Flagship Initiative is "Resource Efficient Europe": The aim is to support the shift towards a resource efficient and low-carbon economy that is efficient in the way it uses all resources. The aim is to decouple our economic growth from resource and energy use, reduce CO_2 emissions, enhance competitiveness and promote greater energy security.

The EU Horizon 2020 Societal Challenges resonate the Europe 2020 Agenda where funding will focus on the following challenges:

Health, demographic change and wellbeing;
Food security, sustainable agriculture and forestry, marine and maritime and inland water research, and the bioeconomy;
Secure, clean and efficient energy;
Smart, green and integrated transport;
Climate action, environment, resource efficiency and raw materials;
Europe in a changing world – inclusive, innovative and reflective societies;
Secure societies – protecting freedom and security of Europe and its citizens.

5 Vision, Aims and Objectives

The Vision of the ICT Energy Co-ordinated Action is to direct research towards sustainable energy solutions for ICT devices and systems.

Sustainable energy was defined by the United Nation's Brundtland Commission "Our Common Future" in 1987 [UN:1987] as requiring fuel or energy sources which have the following criteria:

1. fuel is not significantly depleted by continuous use,
2. no significant pollution or hazards to humans, ecology or climate systems,
3. no significant perpetuation of social injustice.

There are two main aims required to meet this vision. The first is that the consumption of energy by all ICT devices and systems must be reduced. The second is that the use of sustainable energy and in particular renewable energy systems must be increased to power the majority of ICT systems.

In fact, these aims represent also strategic conditions for the future development of ICT itself:

1) If we want to foster the realization of the next generation of High Performance Computing (HPC) systems we need to increase energy efficiency of computing. Exascale computers capable of reaching 10^{18} operations per second require a substantial decrease in the amount of energy dissipated into heat compared to present standards.
2) On the other hand the so-called Internet of Things (IoT) scenario foresees that an ever-increasing number of smart, mobile, sensing and communicating devices will be dispersed into ordinary appliances and tools of common use. In order to meet the performances of foreseeable energy harvester generators, the amount of power required by such devices needs to be significantly reduced.

At present all ICT roadmaps still use cost or performance as the main driver and improved energy efficiency as a secondary issues (i.e. cost or performance can only be achieved if energy issues are considered). If ICT is to become sustainable in terms of energy then energy must be the key driver for all ICT devices and systems. Also a whole ICT energy approach is essential and not just incremental improvements at each level of the system stack.

The key objective of this document is to define the challenges, opportunities and potential mechanisms to strive towards energy sustainable ICT.

6 Challenges for Sustainable Energy ICT

6.1 The Energy Issue for ICT

The use of Information Communication Technology (ICT) systems is now ubiquitous and ranges from small autonomous sensor systems operating at the mW level to high performance computing (HPC) systems requiring tens of MWs for operation. In between these power levels lie a large number of devices including embedded sensors, mobile phones, smart phones, tablets, personal computers, servers and cloud computing storage systems. The annual sales of many of these consumer systems are now at the 100 million to 1 billion level per annum and every device consumes a certain amount of energy which will results in the emission of CO_2.

The ubiquitous use of ICT systems has been driven by the continuous scaling of silicon chips. The original drivers for scaling came from improving the performance of computers but as the size of transistors was reduced, so was the power consumption enabling many portable systems to be developed. Figure 2 present a summary of the scaling of silicon chips and demonstrates how the performance of computers has improved over time along with the number of transistors physically produced on each chip.

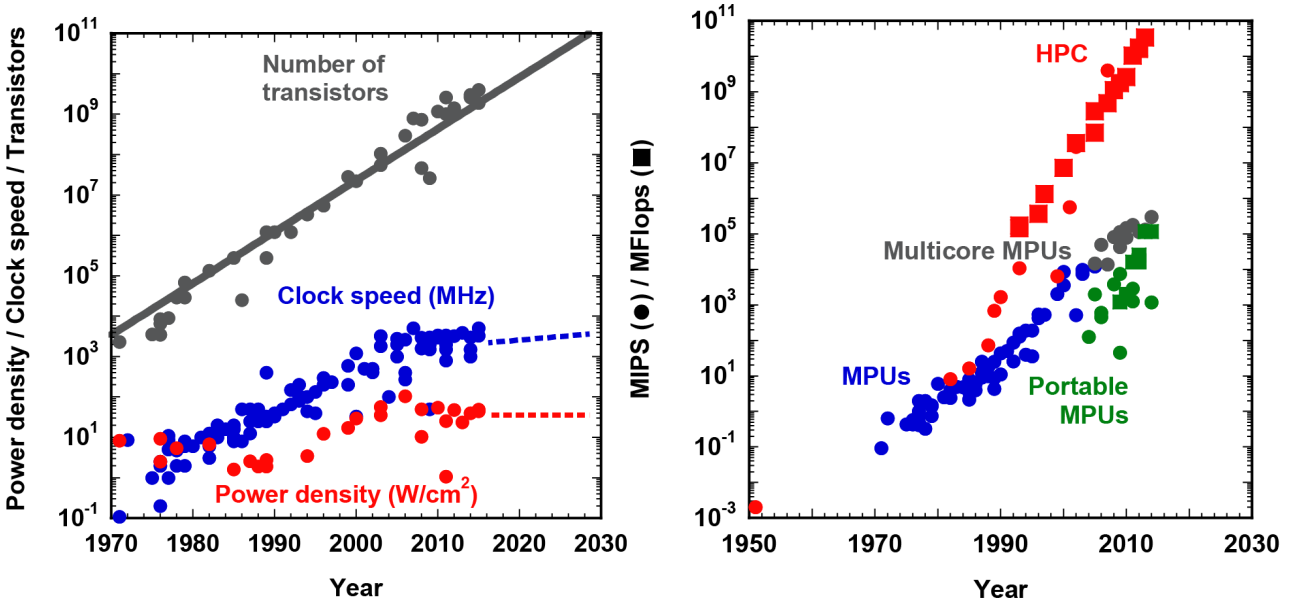

Figure 2: (a) The scaling of the number of transistors, chip clock speed and power density as a function of time. (b) The performance of microprocessor units (MPUs) for computers, portable devices and high performance computing system as a function of time. The circles are data for million instructions per second (MIPs) whilst the squares are million floating point operations per second (MFlops). Sources: Datasheets for processors from Intel, AMD, IBM, Digital, Motorola, Zilog, Samsung, Apple and Top 500 HPC [HPC500:2015].

As the number of transistors increased according to Denhard's scaling rules, the device sizes started to become significant for quantum effects and so leakage currents started to become significant whilst previous CMOS was dominated by dynamic power dissipation (Figure 3). Fred Pollack from Intel first suggested the problems of continuing the scaling of the 1990s without changes to the architecture where the chip power density would scale to that of a nuclear reactor [Pollack:1999]. Indeed in 2006, Intel released a chip with a power density higher than the core of a nuclear pressurised water reactor (PWR). Chip design was changed to reduce this power density but an analysis of the absolute power indicates that with the increasing number of transistors, the total power of MPUs is still increasing over time despite the reduction in power density (Fig. 3(b)). More worrying is the increase in peak power dissipation of low power, portable MPUs which is being driven by applications such as video streaming. This has led to architectures such as the ARM Big.Little which during normal operation can allow low power operation but when computationally intensive tasks must be undertaken, the peak power will increase significantly as required by the video streaming applications.

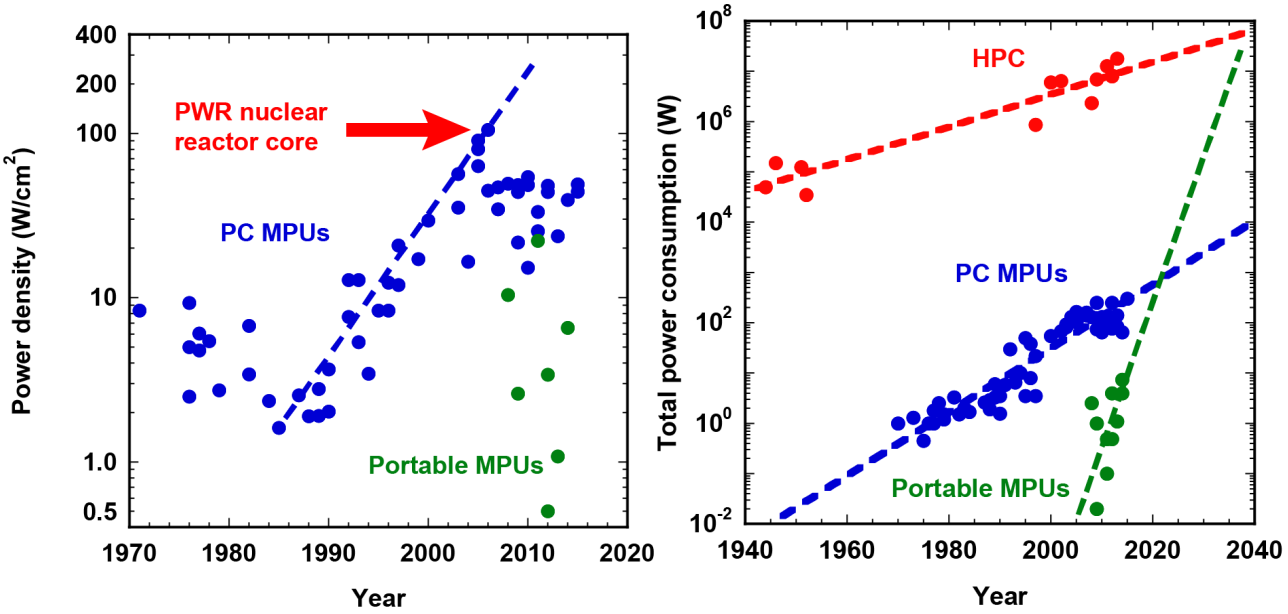

Figure 3: (a) The power density for PC and portable MPUs versus year of first release. The PWR nuclear reactor core power density is 102 W/cm2 [IAEA:2007]. (b) The total power consumption of HPCs, MPUs for PCs (servers, desktop and laptops) and MPUs portable devices (smartphones and tablets). Datasheets for processors from Intel, AMD, IBM, Digital, Motorola, Zilog, Samsung, Apple and Top 500 HPC [HPC500:2015].

The total number of ICT devices being sold each year is also increasing with time (Figure 4(a)). As the number of ICT devices and especially portable devices increases, the amount of data being transmitted by the internet and communication networks is also increasing significantly (Figure 4(b)).

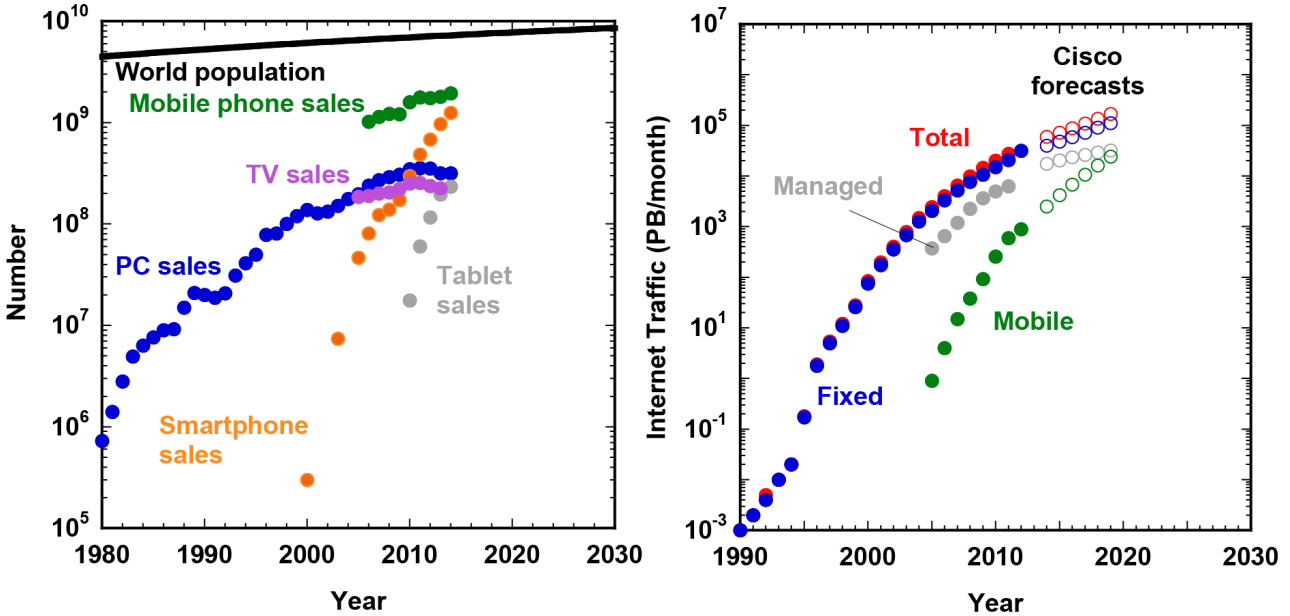

Figure 4: (a) The number of shipped end user device products per annum for personal computers (PCs – both desktops and laptops), mobile phones, smartphones, tablets and TVs (Sources: [Gartner:2014], [UN:2012]). (b) The average internet traffic in terabytes per month for each year. (Source: [Cisco:2015]).

A number of studies have been looking at the energy consumption of ICT devices and systems (Figure 5). While a number of sources especially on web pages provide guesses of the total electricity consumption of ICT devices, there are a number of detailed studies which have tried to accurate estimate the total energy consumption and CO_2 emission [Pickavet:2008][Malmodin:2010][Heddeghem:2014][Smarter:2013]. These studies have used trade data to estimate the number of devices, analysed average use and loading of devices and considered the power scaling to provide estimates of the total electricity consumption (Figure 5) and

CO_2 emission (Figure 6). In particular there are a significant number of studies investigating the energy consumption and scaling of the internet as the present energy consumption of telecommunications is the fastest growing part of ICT energy consumption [Fehske:2011][Lambert:2012][Krug:2014]. The key message is that in 2015 around 4% of the electricity generated worldwide [IEA:2014] is consumed by ICT devices which results in 1 billion tonnes CO_2 equivalent which is about 2.3% of the global emission of CO_2. These studies do not include TV and media uses of ICT systems nor the CO_2 produced from the manufacture of the ICT devices and systems. The suggestion from [Malmodin: 2010] is that TV and media has 82% of the energy consumption of ICT but 131% of the CO_2 emission of ICT. As media is now being transferred to ICT devices and systems, a large part of this consumption may require to be included in the total ICT consumption and emission in the future.

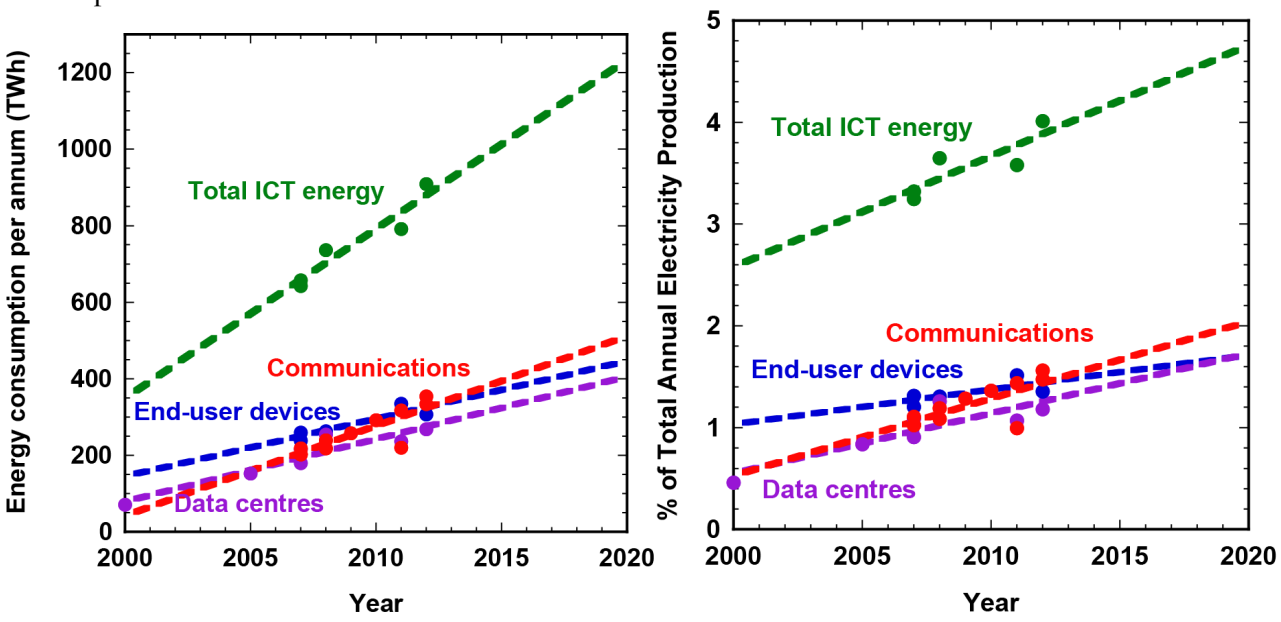

Figure 5: (a) The estimated energy consumption per annum for data centres, PC devices (including desktops, laptops and tablets), communications (internet, networks, mobiles, smartphones) and data centres (servers including cloud computing) plus the total annual ICT energy consumption. Total ICT energy does not include any entertainment and media use such as TV, HiFi, DVD, CD or radio. The ICT energy also excludes all manufacture and disposal of ICT devices Sources: [Pickavet:2008] [Malmodin:2010] [Heddeghem:2014] [SMARTer:2013].

Table 1 is taken from [Krug:2014] which investigates the energy and carbon costs for communications networks which includes an estimate of the embodied carbon from manufacture and also the carbon emissions from the use of the devices (in-use carbon). For the majority of present ICT devices which are connected to the internet through fixed lines, the carbon emission from manufacture (embodied carbon) is far greater than the carbon emission from use. While such modelling always has to estimate the average use of devices and some users it does demonstrate that at present the carbon emission from manufacture are higher than that from usage. This may not be the case in the future if consumers demand high definition video on wireless devices.

System	Embodied carbon (kg CO_2)	In-use carbon (kg CO_2)	% device penetration
Laptop	97 ± 30	24 ± 10	61%
Games machine	102 ± 10.2	59 ± 5.9	52%
PC	108 ± 40	99 ± 40	44%
Smartphone	19 ± 6	4 ± 2	39%
Small games device	13 ± 4	1 ± 0.4	32%
E-reader	13 ± 4	1 ± 0.4	14%
Tablet	35 ± 10	9 ± 4	18%

Table 1: The embodied carbon emissions from manufacture of products (embodied carbon), the in-use consumption of carbon and the % device penetration used for some of the modelling of carbon emission from ICT devices and systems. Source [Krug:2014].

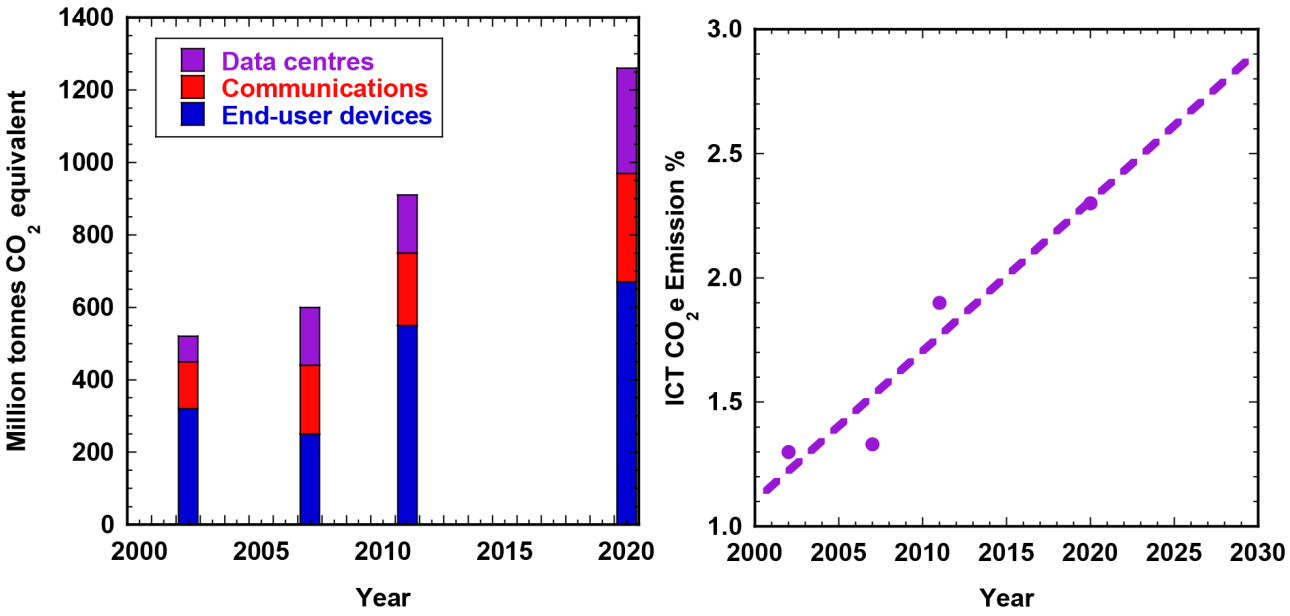

Figure 6: (a) The amount of CO2 equivalent emitted from the manufacture and use of ICT equipment, infrastructure and systems per annum along with predictions for 2020. (b) The percentage of ICT CO2 equivalent emissions as a percentage of total CO2 emissions. Sources: [SMARTer:2020] [Malmodin:2010].

There are a large number of market surveys predicting the future of the ICT market and all of them suggest growth in a significant number of areas. As an example, Figure 7 plots the 2013 share of the Integrated Circuit sales versus the compound annual growth rate (CGAR) for a wide range of ICT devices and systems from [ELC:2014]. Only standard PCs and set top boxes are predicted to be static or decrease whilst all other areas are predicted to grow significantly suggested that the number of ICT devices will continue to grow in the foreseeable future.

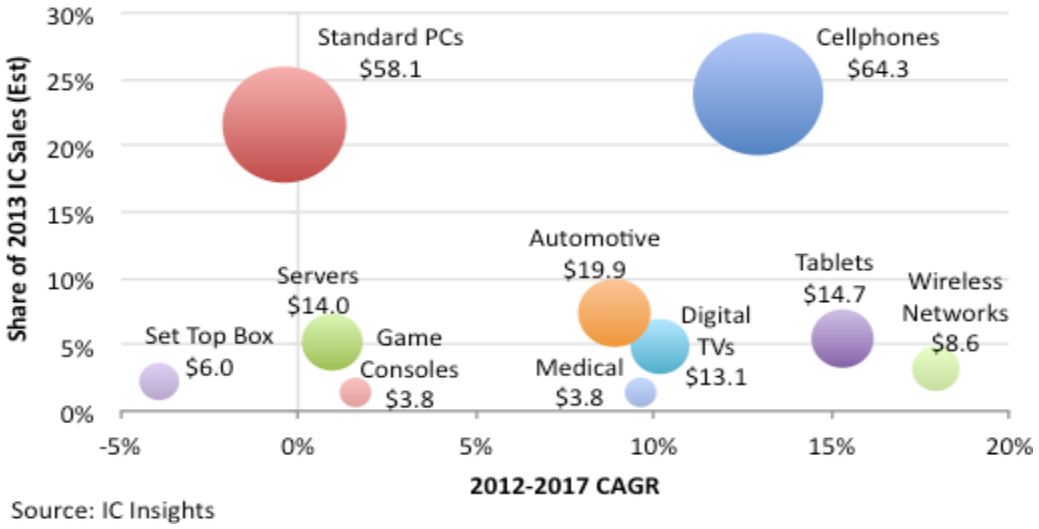

Figure 7: Market analysis with the predicted growth rates of different ICT applications. From [ELC:2014].

An obvious way to reduce CO_2 emissions is to use sustainable and in particular renewable energy generation sources (Figure 8) provides a comparison between the power requirements of ICT devices and systems and available sustainable energy generation technology. The real issue is that most sustainable energy technologies require significantly higher capital outlay for installation (e.g. photovoltaic (PV), hydro) and have long payback periods. Also many of the renewable sources cannot deliver a constant supply of energy and require both storage and/or alternative power supply mechanisms to maintain a constant energy supply. At the small scale batteries and super capacitors are sufficient with power management to deal with such issues. At the large scale for HPC and cloud computing, the storage of large amounts of energy is

problematic and only pump-storage hydro can reach the required volume of energy in a sustainable manner. Such hydro schemes can only be built in suitable environments where large reservoirs with significant height difference can be built and such environments are seldom where the energy is required.

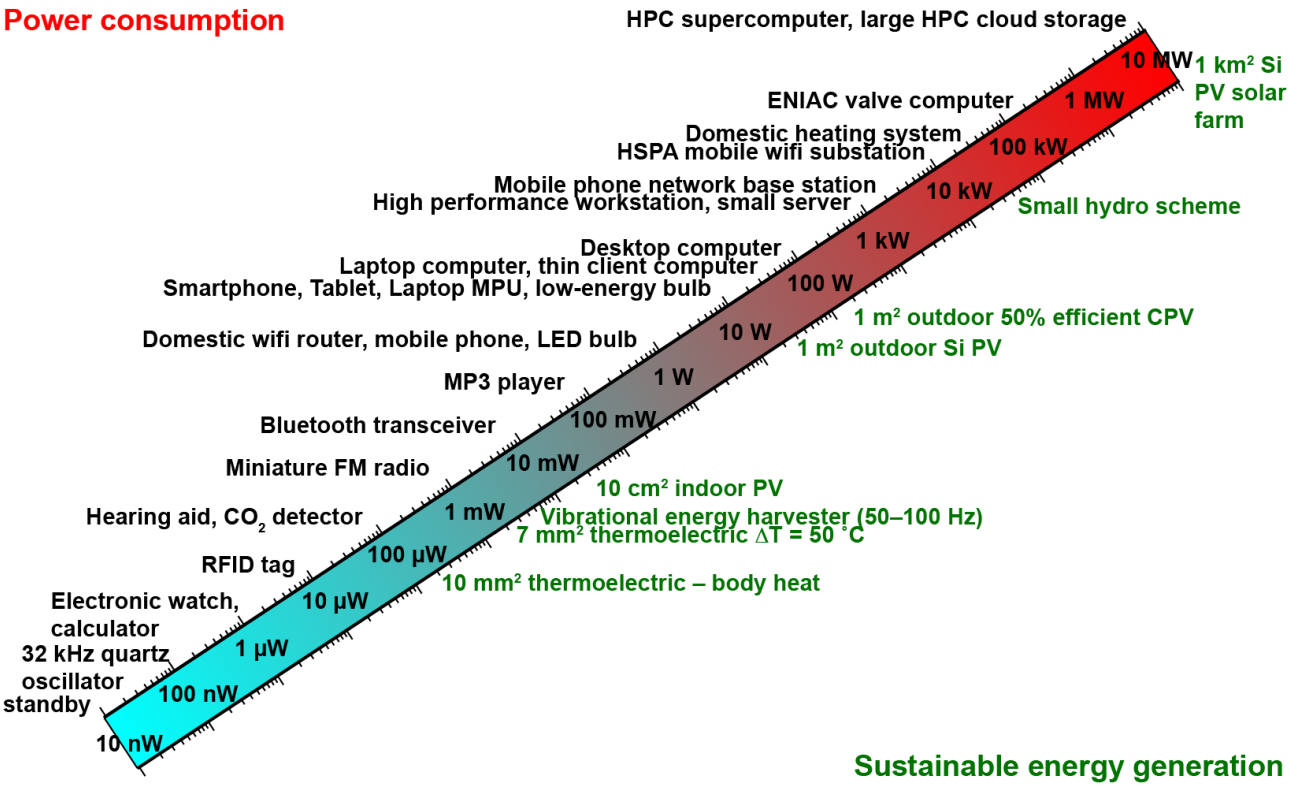

Figure 8: The typical power consumption of different ICT devices and systems versus the power that can be generated from sustainable / renewable energy generation devices and systems.

7 Relationships to Other Projects, Agendas and Roadmaps

There are many research agendas and roadmaps with different visions of lower power ICT devices and systems but none of these documents address mechanisms to achieve sustainable energy ICT. At present virtually all ICT roadmaps still use cost or performance as the main driver and improved energy efficiency as a secondary issues (i.e. cost or performance can only be achieved if energy issues are considered). The single project also investigating the energy issues is the IEEE Green ICT project [GreenICT:2015].

Electronics Enabling Efficient Energy Use, December 2009

High Performance and embedded Architecture and Compilation (HIPEAC) "The HIPEAC Vision for Advanced Computing in Horizon 2020" March 2013

International Technology Roadmap for Semiconductors [ITRS:2013] December 2013

A European Industrial Strategic Roadmap for Micro- and Nano-electronic Components and Systems: Implementation Plan, June 2014

The International Exascale Software Project Roadmap

European Technology Platform for High Performance Computing, ETP 4 HPC Vision Paper

AENEAS and CATRENE "Innovation for the Future of Europe: Nanoelectronics Beyond 2020"

AENEAS, ARTEMIS and EPoSS "2015 Multi Agenda Strategic Research and Innovation Agenda for the ECSEL Joint Undertaking" October 2014

8 Drivers

8.1 Application Drivers

As demonstrated in Figure 7, presently there is an increasing demand for ambient intelligence devices, various kinds of sensor networks for safety and environmental monitoring and for monitoring of health. Such devices are components in the internet of things (IoT). These all require distributed powering systems, batteries or RF-sources for operation. Wiring is expensive, batteries have to be replaced and operation distance for RF-powered devices is short, not to mention the pick-up coils with non-negligible dimensions. The current CMOS-based ubiquitous devices can operate at minimum power levels ranging from 1 μW to 10 μW at around 100 kHz frequencies. Radio transmitters usually need 1 mW of power, but by using burst transmission with low duty cycle the average power level can be substantially reduced. If one can realize an energy harvester with a capacity to deliver power of a few μW into battery or capacitor storage, it would open ways to large number of applications. Also, the availability of this kind of power sources would boost the development of even lower power devices, leading to the vastness of autonomous nanoscale ICT systems for implants and in vivo health monitoring, environmental warning and hazard preventing networks and for other safety measures. This is the present domain of what is sometime called energy harvesting or energy scavenging approaches. In Figure 8 the energy requirements of a number of common electronic devices are shown.

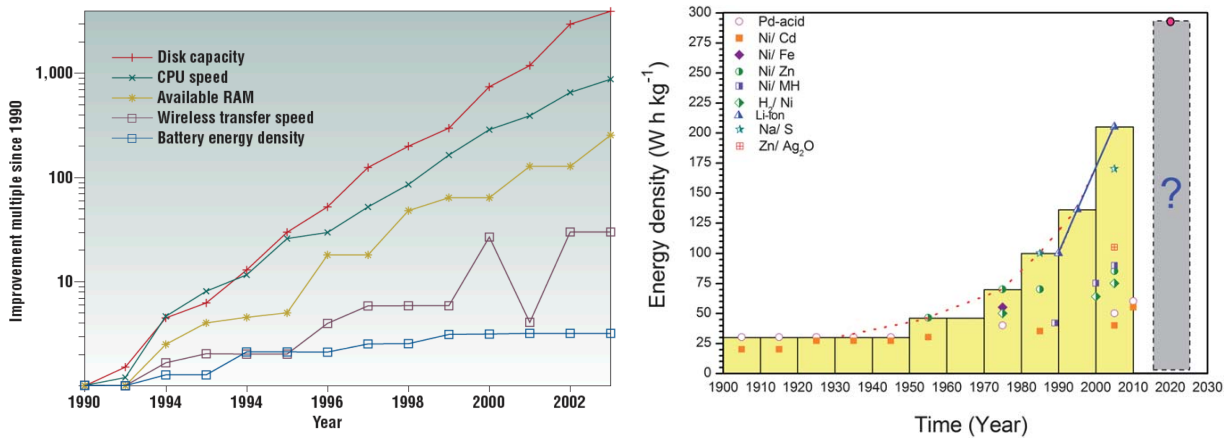

Figure 9: (a) Improvement of the battery energy density compared to other significant figures in the ICT panorama in the last twenty years. (b) The improvements in battery energy density since 1900. Sources: [Paradiso:2005][Zu:2011].

Standard solutions, i.e., disposable batteries, cannot always be employed due to a number of reasons. Traditional batteries were invented by A. Volta in 1799, and were widely used to provide electrical energy to any possible application until the late 1800s – early 1900s, when wiring of cities relegated batteries to mobile applications. In the late 1900s, as electronic components became smaller and required less power, batteries could grow smaller, enabling the wide diffusion of todays wireless and mobile applications. Although economical batteries are a prime agent behind this expansion, they also limit the penetration into new markets. One important factor is the battery energy density (i.e. the amount of energy stored in a given system or region of space per unit volume, or per unit mass, depending on the context). Fig. 9 demonstrates the relative improvement of the battery energy density compared to other significant figures in the ICT panorama in the last twenty years. It is clearly demonstrated that the battery energy density exhibits the slowest trend among those presented here. Although it is true that the research in this field is very active and the situation might improve in the future, to date (and for the foreseeable future) we have no viable solution to the demand of cheap, reliable, high density portable power sources. On top of such difficulties there is a problem related to the battery finite lifetime: indeed the ubiquitous computing dream of "wireless sensors everywhere" is accompanied by the nightmare of battery replacement and disposal when dealing with primary or non-rechargeable batteries which have been utilised for autonomous systems to date. The use of rechargeable batteries which has been critical in enabling mobile ICT development would also benefit autonomous devices or systems. Unlike portable electronic devices however, a system led approach must be implemented where a hybrid energy source combining harvesting and storage is utilised as the wireless sensors will never be connected to the grid for recharging. For wireless sensors in the IoT the batteries can be recharged by the harvesting source to guarantee energy provision during periods of low available energy for harvesting. For such reasons, a new approach based on the exploitation of energy harvested where and when available has attracted considerable attention in recent years [Paradiso:2005][Roundy:2003][O'Mathuna: 2008][Gammaitoni:2012][Rohan:2014].

This new approach somehow implies a significant change of paradigm. In the "battery-only era", the energy was managed in large, centralized compounds where it was transformed from some available source into chemical energy and then stored into batteries in order to be transported and subsequently installed in the mobile devices which were then recharged daily. With this new approach energy is harvested from the environment and transformed locally in real time into electric energy ready to be used when and where necessary. Buffering or temporarily storage of electric energy co-located with the devices is, however, still required for most applications. The energy density per unit weight or volume available in state-of-the-art batteries is orders of magnitude higher than that available from energy harvesting sources further emphasising the need for a hybrid device. Significant improvements must also be made to minimise device size for the hybrid energy elements to match the footprint or volume of the wireless sensors they enable. Long-life hybrid energy systems require new routes to process materials of appropriate dimension to meet energy and power needs while also exhibiting stabilised material interfaces such as provided by solid state materials to enable thousands of recharge cycles for 10 year wireless sensor operation. Following this new paradigm, in the last twenty years, a number of studies on energy harvesting techniques have been performed (see [ENIAC:2015] for reviews). In a wider context, electronic devices currently account for 15% of household electricity consumption, but their share is rising rapidly, mainly due to growing demand in Africa and the developing world. Next to the need for more secure and greener energy supplies at a large scale, immediate action has to be taken to employ alternative energy sources and reduce power consumption in consumer electronics at all levels.

8.2 Computing System Cost Drivers

Rather than companies owning their own computers and systems, there is a financial drive towards cloud computing which potentially allows reduced cost of ownership through the use of a large, centrally managed pool of storage and computer resources [Baliga:2011]. Cloud computing is defined as a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources that can be rapidly provisioned and released with minimal management effort or service provider interactions. In most cases the server is a cloud datacentre. Such cloud computing models provide benefits to the providers through economies of scale since at any one time substantial numbers of end users are inactive. The end users benefit from having data and services available from any location, from data being backed up from a centrally managed facility, increased capacity when required and usage based charging (i.e. no requirement for large one-off investments in hardware every 3 to 5 years). The issue from an energy perspective is whether the increased communication between devices is of higher or lower energy consumption than the reduction in energy by using lower power clients rather than higher power individual PCs and the efficiency of storage and computation on a large cloud datacentre server.

Initial modelling of cloud computing systems has demonstrated a range of energy efficiency gains but only for some applications and only if modern low power clients are used with the systems. Private local area network cloud systems demonstrate better energy efficiency [Baliga:2011] and suggest future opportunities for energy reductions with different approaches to cloud computing.

For HPC whilst performance is still the key driver, the electricity running cost of the system is now being strongly considered in the majority of present and future HPC systems. It is still the cost of the energy that requires HPC people to consider energy rather than a strong driver to improve overall computational efficiency although the first signs that this is changing are starting to appear. Whilst Figure 3 provides a linear scaling which includes the valve systems of the 1940s and 1950s, if only the large scale parallel core systems from circa 2005 are considered the scaling of the total power reaches 100 MW soon after 2015. Clearly new approaches are required.

8.3 Communications Drivers

The area which is dominating the increase in energy consumption (see Figure 5) in ICT systems is communications. Video is presenting moving from standard definition to high definition (HD) and 4K. Table 2 presents the resolution and number of pixels in each frame of a wide range of video systems. The energy issues in expansion of data from standard video formats to HD is clear as PAL to 1080p format corresponds to an increase in the number of pixels by a factor of 7. HDTV to the lowest resolution 4K format is another expansion by nearly a factor of 4 so in total moving from standard definition to 4K represents a 27.6 times

increase in the required data per frame. Today the largest volume of network traffic is video and this is predominantly films and TV for entertainment. There are also other forms of video which are expanding and are expected to impact the bandwidths and volume of communications in the future. This expansion includes cloud-gaming, education, health and person to person communications. This provides new future challenges energy wise as the present techniques to manage the energy impact of video (e.g. multi-cast and caching) may no longer be applicable for these new formats.

Application	Standard	Resolution	Display aspect ratio	Pixels
Analogue TV	PAL, SECAM	520 x 576	4:3	299,520
Analogue TV	NTSC	440 x 486	4:3	213,840
Analogue TV	VHS PAL/SECAM	310 x 576	4:3	178,560
Digital HDTV	720p (FWXGA)	1280 x 720 1366 x 768	16:9 16:9	921,600 1,049,088
Digital HDTV	1080i, 1080p	1920 x 1080	16:9	2,073,600
Digital 4K TV	2160p	3840 x 2160	16:9	8,294,400
Digital 4K TV	4320p	7680 x 4320	16:9	33,177,600
Digital TV	DVD PAL	720 x 576	16:9	414,720
Digital Film	Digital cinema 2x	1998 x 1080	2.39:1	1,757,184
Digital Film	Digital cinema 4x	3996 x 2160	1.85:1	8,631,360
Digital Film	IMAX Digital	5616 x 4096	1.37:1	23,003,136
Video conference	CIF	352 x 288	1.22:1	101,376
Video conference	4CIF	704 x 576	1.22:1	405,504
Video conference	16CIF	1408 x 1152	1.22:1	1,622,016
Video conference	Skype 1	640 x 480	4:3	307,200
Video conference	Skype 5.11	1920 x 1080	16:9	2,073,600

Table 2: The pixel sizes for a range of video formats including analogue TV, digital TV, digital film and video conference [WikipediaTV:2015][Wright:2012].

Communications systems are also changing towards a rental model similar to parts of the server market moving to renting storage and services on the cloud. This may go well beyond the simple mobile contract of today which aims to recover the large capital expenses especially of smartphones in more affordable monthly instalments and start to look at resource sustainability of the phones.

9 The System Stack

The system stack provides a complete view of a whole ICT system. Whilst one might expect cloud high performance computers and autonomous sensors to have completely different systems, the system stacks are very similar in many areas and provide an opportunity to learn how to improve each from the experience of the other. Figure 10 provides a general overview of the system stacks for both high performance computers and an autonomous sensor. It also suggests that there is an opportunity to try to use common solutions at different levels of the system stack to reduce power requirements.

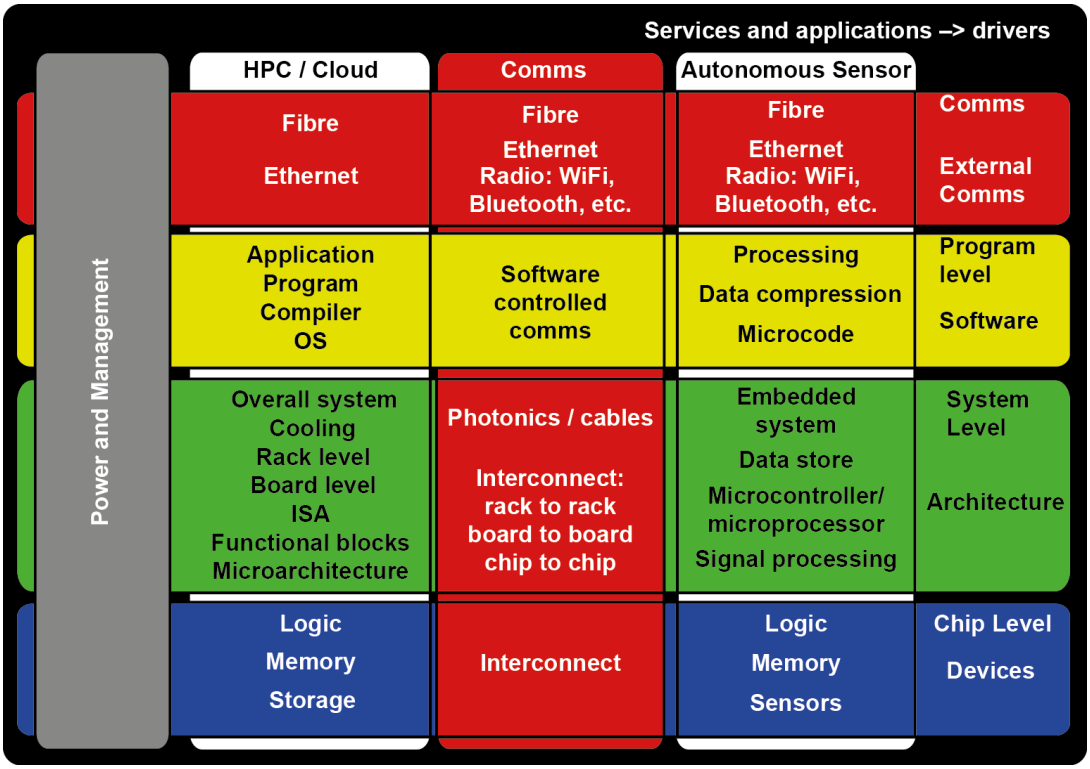

Figure 10: A schematic diagram of the system stacks for a high performance computing (HPC), cloud computing and autonomous sensors. The figure indicates the common elements in the system stacks between these different ICT systems. Whilst HPC and cloud have very similar system stacks, the applications require significantly different performance and costs at most levels of the stack.

Science and Technology Issues for Sustainable Energy ICT

10 Devices

10.1 Present and Future Device Performance

In the last forty years the semiconductor industry has made impressive progresses in reducing the size of the Complementary Metal-Oxide-Semiconductor (CMOS) Field Effect Transistor (FET) components and thus increasing the computational density of standard microprocessors. The continuous scaling to smaller dimensions is now facing a foreseeable end due to increasing heat production as a side effect of the computation process.

Figure 2 on the left reproduces the amount of energy dissipated as heat per unitary surface, per second, in different CPUs, as a function of year of production [Sources: Datasheets for processors from Intel, AMD, IBM, Digital, Motorola, Zilog, Samsung, Apple and Top 500 HPC [HPC500:2015] [Pop:2010]]. The steep increase in CPU power density is mirrored by the increasing fraction of energy spent in cooling activities in the functioning of US data centres as reproduced in the figure above on the right. In order to address this situation in the US, the Semiconductor Industry Association (SIA) launched in 2000 a strategic initiative aimed at demonstrating novel computing devices capable of replacing the CMOS FET as a logic switch in the 2020 timeframe. The 2001 International technology Roadmap for Semiconductors (ITRS) Emerging Research Device Technical Working Group focused attention on the logic switch dissipation and proposed [Zhirnov:2003] a simplified analysis of a generic electronic switch at thermal equilibrium. The switch was modelled as a potential barrier separating two quantum wells (an idealized FET channel between source and drain contacts). The analysis showed that the channel could conceivably be scaled down to ~1.5 nm and the transistor could have a minimum switching speed of ~40 fs – significantly smaller and faster than the FETs of today. In order to avoid leakage at room temperature due to a finite height of the potential barrier, however, the voltage could not be scaled as rapidly as the physical dimensions, and the resulting power density for these switches at maximum packing density would be on the order of 1 MW/cm^2. Energy

efficiency in operating ICT devices is presently considered an objective of extremely high relevance. As demonstrated in Figure 5, the world wide electricity consumption today is around 4%. 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 increasingly important to consider and improve the energy efficiency of ICT. In the short term, it will be an obvious and practical solution to better exploit the potential of technologies that already exist or are currently in development. On the long term, new and disruptive ideas will be required [FET:2010].

10.2 Key Device Challenges

1. The cost of the lowest energy devices requires the latest CMOS technology node which now requires enormous volumes due to the cost of the foundries and technology.
2. The scaling of transistors to smaller dimensions is now not expected to decrease the switching energy significantly.
3. If there is a change of the basic switching device to move to a significantly lower energy technology then circuits architectures, design tools, verification, operating systems and software may require rewriting or complete changes for basic operation or optimal performance.
4. Driving interconnects with multiple fan-out or an antenna have fundamental energy and noise limits that makes ultra-low energy consumption difficult.

10.3 Fundamental Device Limits

As we have briefly discussed here the ICT-Energy relation has two main aspects that are presently under the focus of research. These two aspects have some important common basic roots that become more evident if we consider a generic ICT device like a machine that processes information while transforming work into heat and heat into work. Pioneering research developed by J. Von Neumann and by R. Landauer in the last century has shown that information processing is intimately related to energy management ("information is physical" [Landauer:1961]). As a matter of fact, an ICT device (see Figure 11) is a machine that inputs information and energy (under the form of work), processes both and outputs information and energy (under the form of heat).

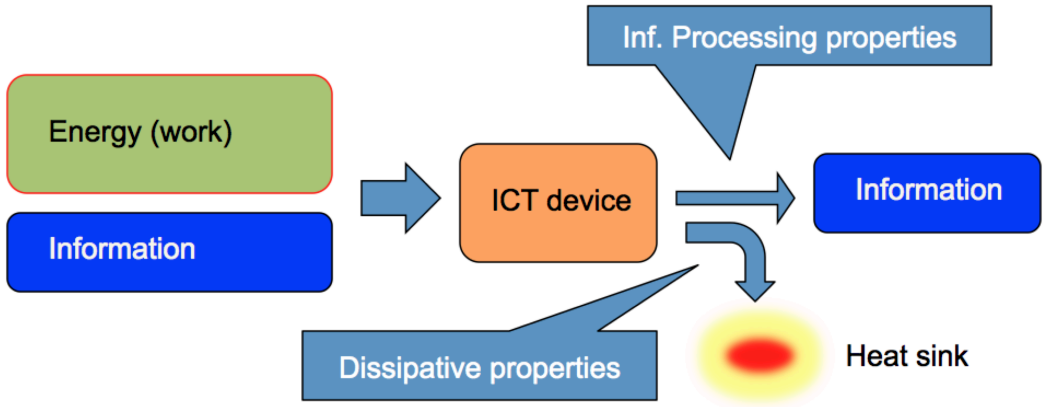

Figure 11: An ICT device is a machine that inputs information and energy (under the form of work), processes both and outputs information and energy (mostly under the form of heat).

In this perspective energy dissipation via heat production and energy transformation processes are two aspects of the same topic: energy management at micro and nano scales. 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. This definition, however, may not apply when we have to deal with processes taking place at the nanoscale. The well-known laws of heat and work transformation that lie at the base of the classical thermodynamics require to be redefined. The very basic mechanisms behind energy dissipation requires a new definition when non-equilibrium processes involving only a few degrees of freedom are considered. These aspects are presently under consideration in a number of laboratories in Europe, US and Japan. One of the main challenges is the timely identification and substantiation of new directions for the powering of small scale

devices, i.e. the management of microscale energies, that have a high potential for significant breakthrough and that may become the foundations of the ICT and innovations of tomorrow[Gammaitoni:2015].

The major driver for many decades for the scaling of CMOS transistors and logic has been the high speed, high performance metrics of high performance computing. With such drivers, power dissipation was a secondary issue up until 2000. At this stage, the critical dimensions in CMOS devices became sufficiently small that the static power dissipation through leakage and parasitic currents started to became larger than the dynamic power dissipation from switching (Figure 2). Simulations at the time quickly demonstrated that the continued constant electric field scaling would quickly produce dissipated power resulting in temperatures well above the melting points of the metals and even the semiconductors being used for the systems. These problems were the results of a number of fixed parameters in the transistor which did not scale in a standard linear, quadratic or cubic way with the gate length of the transistor including the bandgap and voltage threshold (the voltage at which the transistor is switched on into a digital "1" state). New types of scaling rules as well as new designs and materials were introduced to reduce the power dissipation.

FETs, however, have fundamental power and speed limits and even switching to new materials, these cannot be overcome.

The thermodynamic limit was defined by Landauer [Landauer:1961] which gives as the minimum energy required for a binary resetting operation:

$$E_{min,TD} \geq k_B T \ln 2 = 2.87 \times 10^{-21} \text{ J / operation at 300 K}$$

where k_B is Boltzmann's constant and T is the temperature.

Heisenberg's uncertainty principle also provides a limit in that the uncertainty in a switching time Δt can only be determined with an uncertainty in the energy

$$E_{min,QM} \equiv \Delta E \geq \frac{h}{\Delta t^2}$$

where h is Planck's constant and Δt is the uncertainty in the switching or delay time.

The switching energy of a MOSFET transistor is given by

$$E_{MOSFET} = \alpha C_L f V_{DD} \tau + I_{leakage} V_{DD} \tau + P_{SC} \tau$$

where α is the activity factor, C_L is the load capacitance, f is the circuit frequency (clock frequency), V_{DD} is the supply voltage, τ is the delay or switching time (= 1/f), $I_{leakage}$ is the leakage current and P_{sc} is the short circuit power dissipation.

Whilst in principle the switching energy of a MOSFET could be reduced by reducing the supply voltage V_{DD} (indeed this is what scaling over the last 40 years has been doing to reduce the energy and power dissipation), the finite bandgap of a semiconductor provides a lower limit below which the transistor will not switch on. Moving to another semiconductor with a smaller bandgap allows some improvement to reduce energy but the fundamental physics of the p-n junction for the contacts provides another limit for switching MOSFETs one and off. The minimum voltage is a function of the sub-threshold slope, the slope of the device between the on-current, I_{on} (the digital "1") and the off-current, I_{off} (the digital "0"). This slope is defined as

$$SS = (Log10)\left(\frac{k_B T}{q}\right)\left(\frac{C_{total}}{C_{ox}}\right)$$

where q is the electronic charge, C_{ox} is the gate oxide capacitance, C_{total} is the sum of the oxide, depletion and interface state capacitances. Once the capacitances have been optimised and defects minimised, then they tend to 1 leaving the minimum sub-threshold slope only dependent on fundamental constants and temperature. For a device at 300 K, the minimum slope is about 60 mV/decade which fixes a minimum threshold voltage for a MOS transistor that cannot be circumvented. Only by moving to a radically new technology can lower power dissipation potentially be achieved.

The ITRS roadmap [ITRS:2013] along with a number of other review articles [Pop:2010][Nikonov:2012] provides details of potential technologies being reviewed to reduce the subthreshold slope by moving away from p-n junctions and using quantum mechanical tunnelling to allow lower V_{DD}.

Figure 12 provides a comparison of present and future high performance and low power CMOS against future proposed device technologies in the ITRS roadmap [ITRS:2013]. Present low power CMOS is already at 3 aJ per operation (excluding large fan-out and interconnect impedance) and only predicted to reduce by a factor of 3 over the next 13 years. The majority of improved future technologies do not produce any significant reduction in energy consumption per operation. Indeed a 100 times the Landauer thermodynamic limit appears to be difficult to improve on when all the new proposed device technologies are considered.

This 100 times the Landauer limit at 300 K is related to the energy barriers in all the switching devices that provide a significant I_{on} / I_{off} ratio as required for the circuit architectures. A number of new types of device suggest that this barrier can be circumvented. The EC ICT Energy Landauer FET project has demonstrated a switch with $8x10^{-20}$ W of switching power at room temperature, albeit at the slow switching speed of ms. More importantly this switch demonstrates that the 100 x Landauer limit can be broken.

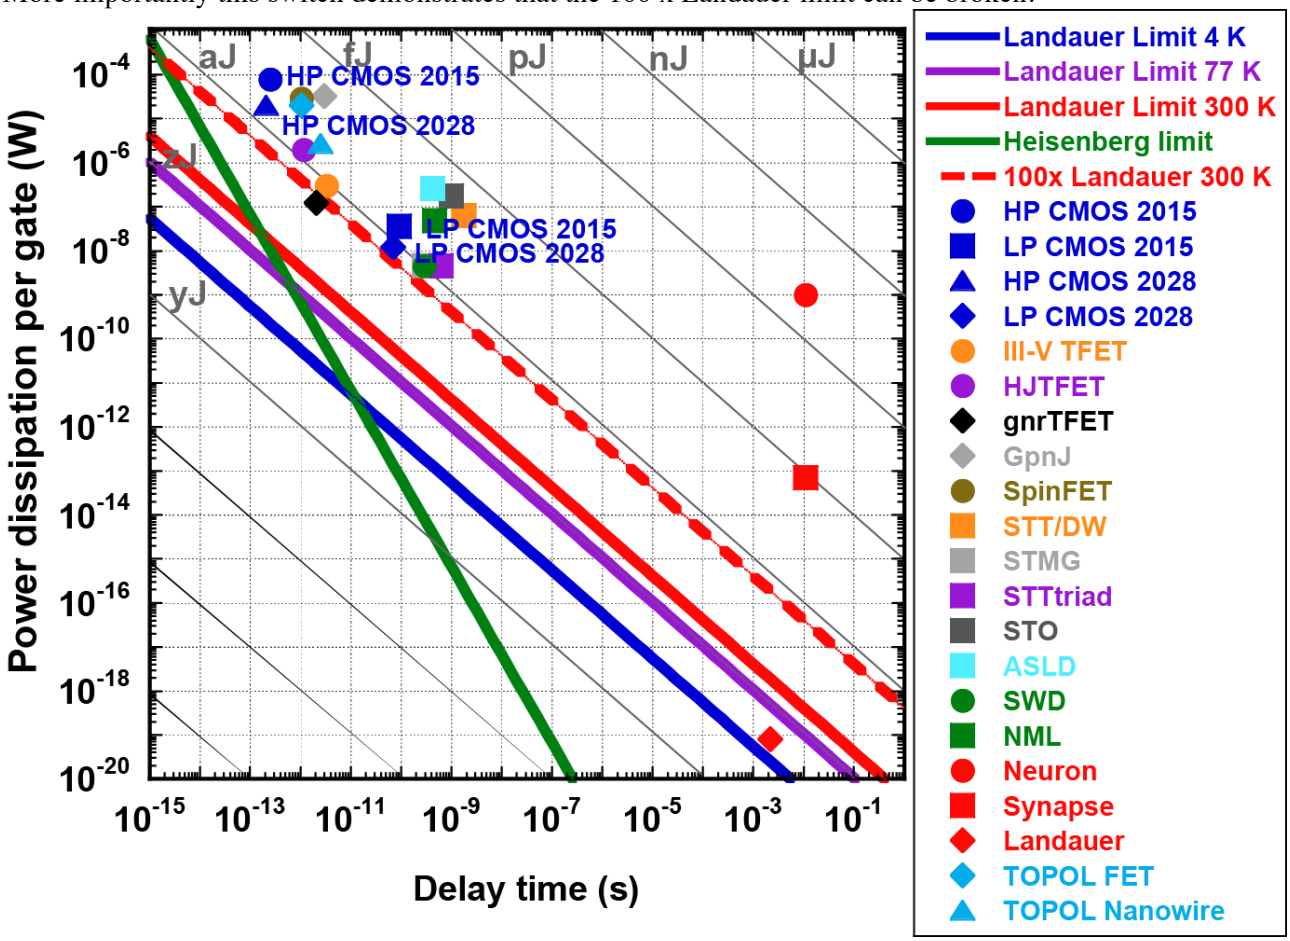

Figure 12: The intrinsic power dissipation per device switch versus delay time for CMOS and future devices as proposed in the ITRS future emerging technology roadmap [ITRS:2013][Nikonov:2012]. The diagonal grey lines are lines of constant energy. The figure indicates a limit of about 1 aJ per switch for CMOS devices and 250 zJ per device for the best proposed future switching device. Key: CMOS HP (high performance CMOS), CMOS LP (low power CMOS), iii-vTFET (III-V tunnel FET), HJFET (heterojunction FET), gnrFET (graphene nanoribbon TFET), GpnJ (graphene pn-junction), spinFET (Sughara-Tanaka spin FET), SST/DW (spin torque domain wall), STMG (spin torque majority device), STTtriad (spin torque triad), STO (spin torque oscillator), ASLD (all spin logic device) and NML (nanomagnetic logic). Results from the MINECC projects are the Landauer MEMS device (Landauer) [López-Suárez:2015], the Si CMOS FET from TOPOL (TOPOL FET) and the Si nanowire from TOPOL (TOPOL Nanowire).

10.3.1 Subthreshold Devices

Figure 12 indicates that all present transistors operate in the regime where the switching power is > 100 x the Landauer limit at 300 K. This is predominantly related to the energy barrier between the source and the drain of the transistors. In MOSFETs, this barrier is dependent on the p-n junctions between the channel and the contacts and is therefore also dependent on the bandgap of the semiconductor material. The subthreshold slope which determines how quickly a transistor can be switched on or off with a fixed voltage is limited to a minimum value of 60 mV/dec at 300 K if a pn junction is used. Ideally a steeper subthreshold whiuch requires less voltage is required as then the switching power can be reduced compared to MOSFETs or any other device with an energy barrier than must be overcome to switch the device on or off. Steep subthreshold devices are designed to operate without this energy barrier. A number of technologies have been proposed but two are the main ones being pursued. The first approach is tunnel devices where electrons or holes

quantum mechanically tunnel through the energy barrier rather than a voltage being applied to lower (or raise) this energy barrier to switch the device on (or off). The second is negative capacitance devices where the dielectric constant of the gate insulator is chosen to provide a negative capacitance and therefore allow subthreshold slopes below 60 mV/dec.

Tunnel field effect transistors (TFETs) is now a large research area as part of the subthreshold device research field. Figure 13 provides a comparison of the on-current to off-current (I_{on}/I_{off} ratio) comparing silicon CMOS to a range of TFETs and transistors in group IV and III-V materials. High performance CMOS has an I_{on}/I_{off} between 10^4 and 10^5 whilst low power CMOS operations with $\geq 10^7$. The best results from III-V TFETs is $5x10^6$ with a minimum subthreshold slope of 21 mV/dec [Tomioka:2012] which looks very promising. The major problem to date is that the I_{on} is only 5 nA/μm at V_{DD} = 1 V which is far too low for present circuit architectures. Figure 14 presents the I_{on} versus minimum subthreshold slope for a range of V_{DD} values. All the III-V technologies close to the I_{on} values of CMOS have worst subthreshold slopes so more work is required before this is a viable route to new devices for circuit architectures. The work does indicate that there are approaches where the 100 x Landauer limit at 300 K could be broken and so finding devices with subtheshold slopes below 30 mV/dec with high I_{on} could significantly reduce the power dissipation and switching energy for low power electronics.

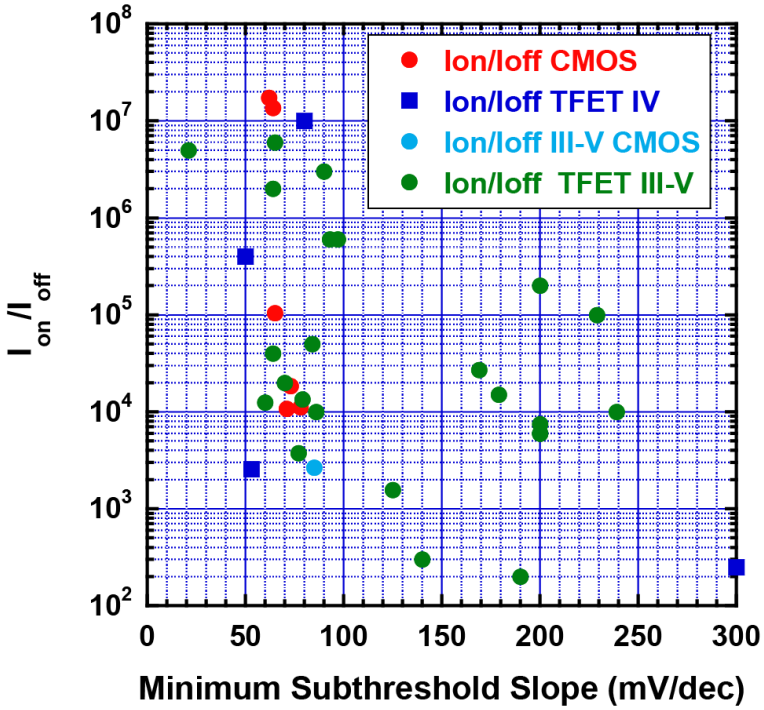

Figure 13: The I_{on}/I_{off} ratio versus the minimum subthreshold slope for a range of different device technologies. Sources: [Chang:2013][Choi:2007][Colinge:2010][Deleonibus:2014][Dewey:2011][Ford:2011][Fujimatsu:2012][Iida:2011][Jan:2012][Kim:2009][Krishnamohan:2008][Leonelli:2010][Li:2012][Liu:2014][Mayer:2008][Mohata:2011][Mohata:2011A][Mohata:2012][Mookerjea:2009][Natarajan:2014][Noguchi:2015][Rajamohanan:2015][Tomioka:2012][Wang:2015][Wu:2013][Yu:2013][Zhao:2010][Zhao:2011][Zhou:2012][Zhou:2012A][Zhao:2014].

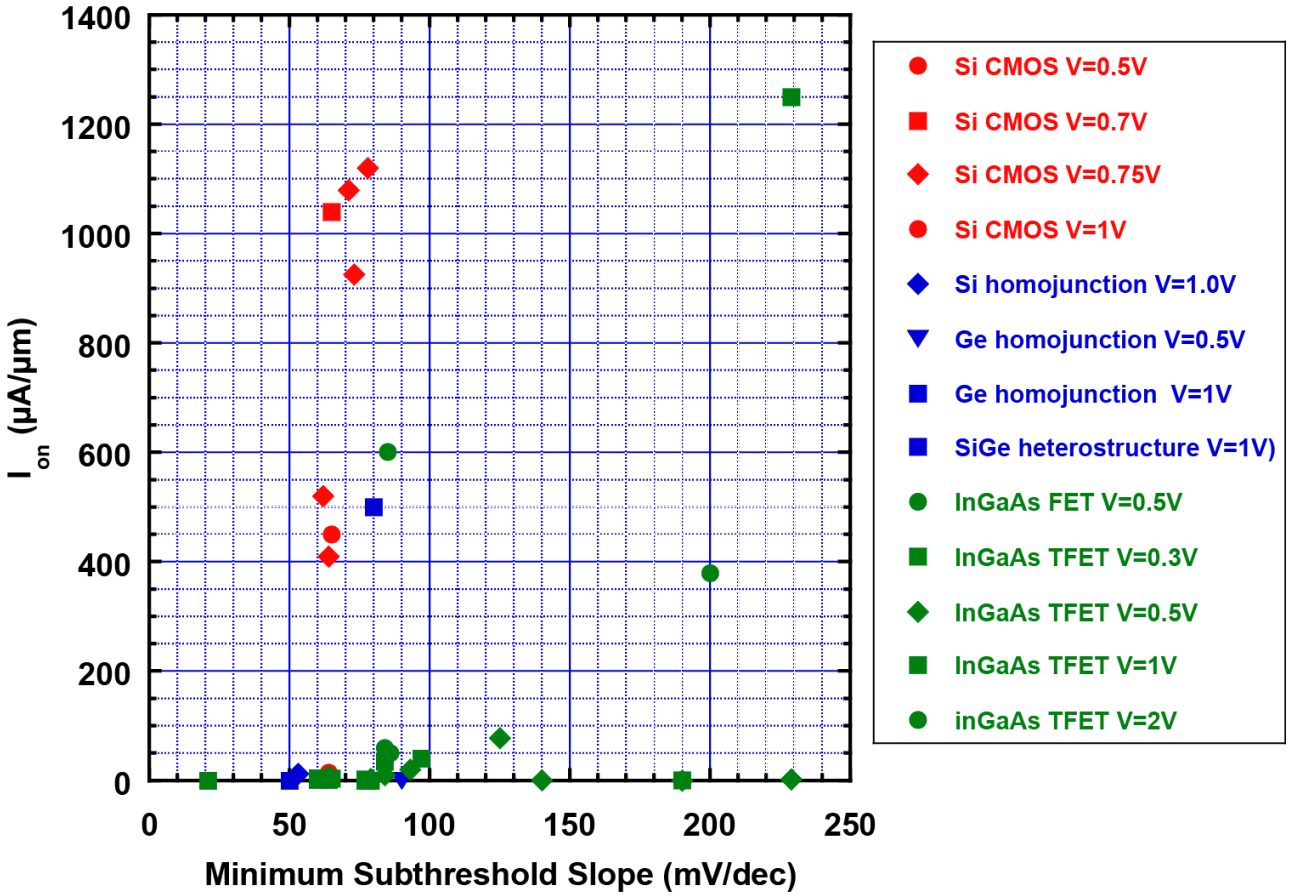

Figure 14: The I_{on} versus the minimum subthreshold slope for a range of steep sub-threshold device technologies. Sources: [Chang:2013][Choi:2007][Colinge:2010][Deleonibus:2014][Dewey:2011][Ford:2011][Fujimatsu:2012][Iida:2011][Jan:2012][Kim:2009][Krishnamohan:2008][Leonelli:2010][Li:2012][Liu:2014][Mayer:2008][Mohata:2011][Mohata:2011A][Mohata:2012][Mookerjea:2009][Natarajan:2014][Noguchi:2015][Rajamohanan:2015][Tomioka:2012][Wang:2015][Wu:2013][Yu:2013][Zhao:2010][Zhao:2011][Zhou:2012][Zhou:2012A][Zhao:2014]

10.3.2 Interconnects and Electrical Wires

For any electrical wire, the energy consumed to transmit a bit of information is related to the capacitance, C and the voltage, V. If the capacitance is approximated by $C \sim \varepsilon_0 L$ assuming an isolated wire where ε_0 is the permittivity of a vacuum and L is the length then the energy for transmitting a bit of information in a metal interconnect or wire can be approximated as

$$E_{wire} = \frac{CV^2}{2} \sim \frac{\varepsilon_0 L}{2}\left(\frac{k_B T}{e}\right)^2$$

This gives a fundamental limit of 3 x 10^{-15} J per bit.m for isolated interconnects or wires [Zhirnov:2014]. The other fundamental limit is shot noise limit relating to the quantum of electric charge in electrons being transported along interconnects. At 300 K this corresponds to an energy of 2.9 x 10^{-21} J/bit.

In real systems the actual values are significantly larger than the fundamental limits. In particular the signal to noise and fan-out require significantly higher energy consumption especially when circuit speeds are high. Examples of the energy per bit for different interconnects as a function of year (ITRS technology nodes) has been estimated by [Miller:2009] and presented in Figure 15.

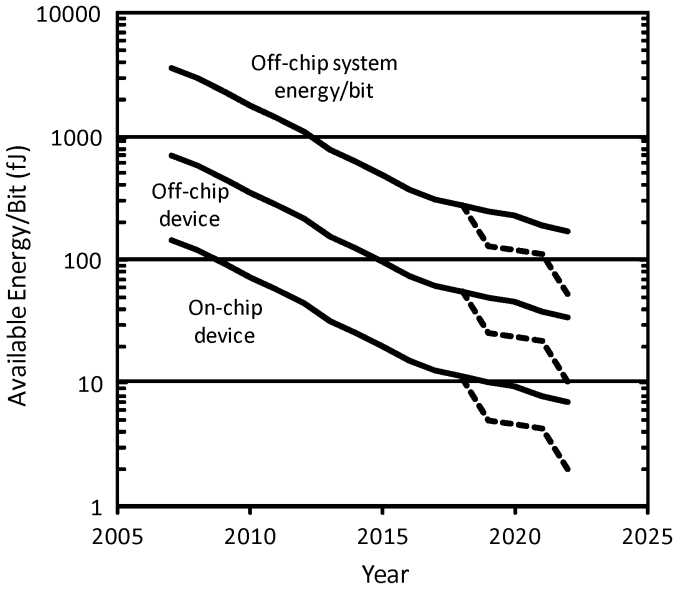

Figure 15: The available energy per bit for wired on-chip, off-chip and off-chip system level interconnects using data from the ITRS to calculate the values. Source [Miller:2009].

10.4 Opportunities for Device Improvement

The ITRS Roadmap [ITRS:2013] provides significant details of the challenges for devices and the opportunities for new device concepts. In particular the opportunities can be divided into continued scaling of conventional transistor devices (termed More Moore) and adding new device concepts for computation (termed Beyond Moore) or new functionality onto the base CMOS technology (termed More than Moore).

More Moore is looking at the challenges of scaling CMOS devices to dimensions below 10 nm and a significant portion of this work is now investigating new channel materials that will allow higher performance with lower voltages for reduced power operation. The materials include Ge, III-Vs, carbon based materials, magnetic materials and phase change materials. The major issue now being recognised by the Semiconductor Research Association and others is that scaling the transistor to smaller sizes may no longer result in lower energy devices. The key concept in More Moore is that the circuit architectures will be similar to present CMOS architectures and only the devices size, voltages or currents are changed.

More than Moore is now a very diverse field of study but the opportunities for research can be described as:

1. Devices with learning capability e.g. devices that learn logic configurations and devices that learn by example.
2. Sub-threshold slope transistors allowing voltages for operation below the normal p-n junction limits are also being investigated.
3. Integration of RF devices and components.
4. Optical devices e.g. Si photonics for communications and optical sensing modalities.
5. Micro Electro Mechanical Systems (MEMS) and Nano Electro Mechanical Systems (NEMS).
6. Integration of sensors.

Beyond Moore is the opportunity to completely change how information is processed. This field is investigating the fundamental limits of information theory and if any of the known limits can be circumvented through innovative techniques. One example includes investigating if the Landauer limit requiring heat to be dissipated through the storage and erasure of information can be circumvented to allow zero energy switching. Spin wave devices are another example for a radical new technology that circumvent many of the limits of conventional transistor logic. Investigating methods of integrating device operation with energy harvesting has also been suggested where the heat dissipated and normally lost could be harvested to improve the overall thermodynamic system efficiency. The Beyond Moore research area has the largest potential to reduce energy consumption of ICT but is also the most difficult to implement into

systems as solutions may be radically different from conventional CMOS technology, architecture and systems.

11 Circuit Micro-Architecture

The micro-architecture level considers devices that are the integration of many of the underlying technologies. Transistor technologies and silicon processes are combined into useful blocks such as memory, control and arithmetic that together form a computational device, often a processor or an accelerator. The scope for that device is very broad, including ultralow power embedded devices, through general purpose processors, up to high performance network-on-chip components.

Micro-architecture exploits what is physically possible with contemporary technology and presents an interface through which other hardware and software can use the device. In hardware terms, this interface is of course physical, and will typically obey a specified protocol. In software terms, the processing device presents a set of possible operations through an Instruction Set Architecture (ISA). If the ISA is a description of the behaviours of the device, then the micro-architecture is the implementation of those behaviours.

This section explores the implications that micro-architectures have in the energy consumption of ICT devices. It covers contemporary trends and their relation to a number of long-standing 'laws' of computing, as well as current or impending challenges that pose significant research challenges if improvements in features, performance and energy are to continue.

At present it is clear that micro-architectures have significantly different drivers for HPC, data centres, general purpose (e.g. PCs) and embedded systems / portable systems.

Semiconductor device technology has hit a fundamental efficiency limit, called the “energy wall” that prevents the reduction of energy consumption when the transistor size is scaled down in the current manufacturing technologies for the forthcoming technology nodes. Based on the current projections, a ten-fold improvement in chip energy-efficiency is needed to maintain the information technology (IT) energy scalability in the next decade.

11.1 Present and Future Circuit Architecture Performance

The International Technology Roadmap for Semiconductors [ITRS], collects data on and monitors the progress made across a broad spectrum of semiconductor devices. The ITRS's Design technology Working Group (TWG) provides insights into trends and emerging solutions, many of which are relevant to micro-architecture. This determines, or at least predicts, what systems will look like in the next few years. The work of the TWG sets the tone for challenges that must be overcome by researchers and industry alike, as well as identifying potential new challenges on the horizon, that may soon become barriers to technological progression.

In the TWG's most recent system drivers roadmap report, summarised in [ITRS2013], a number of continuing trends are observed, as well as the emergence of a number of new challenges.

Frequency limits

First and foremost, the frequencies at which modern devices operate is almost fixed, in that year-on-year increases in clock frequency are now very small (Figure 2).

Both the 2011 and 2013 editions of the ITRS system drivers roadmaps observe a 1.04x frequency change per year. This is a stark contrast to the 1.41x change seen a decade before (Table 3). This operating ceiling has resulted in new technology drivers, particularly multi-core technology, becoming prevalent. The up-take of multi-core has happened at different times in different sectors, as will be discussed in a later section.

Year of projection	Projected yearly clock frequency increase
Pre-2001	41%
2001	17%
2007	8%
2011	4%
2013	4%

Table 3: The ITRS per-year clock frequency increase projections.

The activity factor of a device is the proportion of its transistors that switch, on average, in order to carry out operations. In micro-architectural terms, two adders could be compared in terms of how many transistors must switch in order for them to add two numbers. The adder with a smaller activity factor could be considered more efficient. However, this must be reconciled with the delay in the critical path in the component (determining the maximum clock speed), as well as the overall number of transistors (contributing to static power through leakage current), and is therefore a multi-dimensional consideration.

The ITRS observes in 2013 that activity factors are changing by 0.95x per year, as was the case in 2011. This, combined with the previously described frequency limit, are indicative of technology operating at the edge of practical power dissipation.

There are various other factors and metrics established by the ITRS. However, many of these have a more direct affect on layers beneath the micro-architecture, such as silicon layout and other physical traits. As such, these are not considered in this section.

It is clear that progress in operating frequencies and activity factors are not yielding significant changes that will provide anything beyond the status quo, particularly in energy consumption. Therefore, future research must strive either to disrupt these trends, or deliver dramatic improvements through other means not covered by these metrics.

11.1.1 HPC / Data Centres

Both at a small scale (embedded systems) and at a large scale (datacenters), the so-called “economic meltdown trend of Moore’s law” [Brill:2012] transcends in a dramatic increase in the computation and cooling energy costs. If no major paradigm shift in the design and operation of chips is introduced, the EU estimates that, only datacenters, will likely account for an astronomic 124 TWh by 2020. This trend can only be drastically curbed by supporting research and innovation that promise orders of magnitude gains in energy efficiency, not incremental improvements.

Realizing this vision requires inter-disciplinary research at the boundaries of multiple scientific domains, as well as developing and integrating innovations in several research areas, namely, computer engineering and cooling design, large-scale computing system simulation, software generation and optimization, statistical network modeling and model predictive control theory.

In the field of HPC, one of the most promising lines is the development of energy-efficient processing architectures building blocks that would significantly enhance the energy-proportionality of server processing power at the deep submicron (i.e., beyond 28 nm) era. The development of these building blocks will be achieved by using emerging technologies such as Fully Depleted Silicon On Insulator (FDSOI) [FDSOI:2012], and integrated microfluidic cooling and power delivery [Sridhar:2014]. This development would require modelling the power and thermal dissipations involved in the processing units, memory hierarchy, and the cooling and power delivery networks at the server level.

During the last year the EC FET MINECC project TOLOP has contributed in two ways:

1. We completed the characterization (energy, performance and temperature) of full HW IPs synthetized and manufactured in 28 nm FDSOI technology tailored to emulate parts of parallel accelerators for HPC and many-core server systems.
2. We developed PowerCool, the first-of-its-kind mathematical model that can be used to explore architectures of micro-scale on-chip microfluidic fuel cell networks for joint cooling and localized power generation and delivery for memories and logic of many-core servers. As part of the PowerCool development, we extracted silicon-calibrated power and performance models for server components and implemented microfluidic-based on-chip power-cooling delivery network.

We have explored for the first time the potential of FDSOI technology featuring body-biasing and voltage scaling for ultra-aggressive power and thermal management.

Figure 16 shows the impact of temperature variation on maximum frequency of the prototype when operating at different voltage levels. The device is always in thermal inversion as the maximum frequency always increases with the silicon temperature. Moreover, when powered with low voltage (low power operating mode) the device can see up to 8x performance variation when operating at high temperature.

Body bias is an effective knobs to dynamically tune the system performance and leakage power consumption. This is of primary interest for the YINS project [YINS:2015] when we want to explore innovative solutions to thermally constrained data-center architectures. Indeed body bias increases its

effectiveness as the voltage supply decreases which is a region of interest for massively parallel and energy-efficient architectures. In this region of interest this body-bias technology knob can increase the performance of the 140% and decreasing the leakage percentage of almost the 100%.

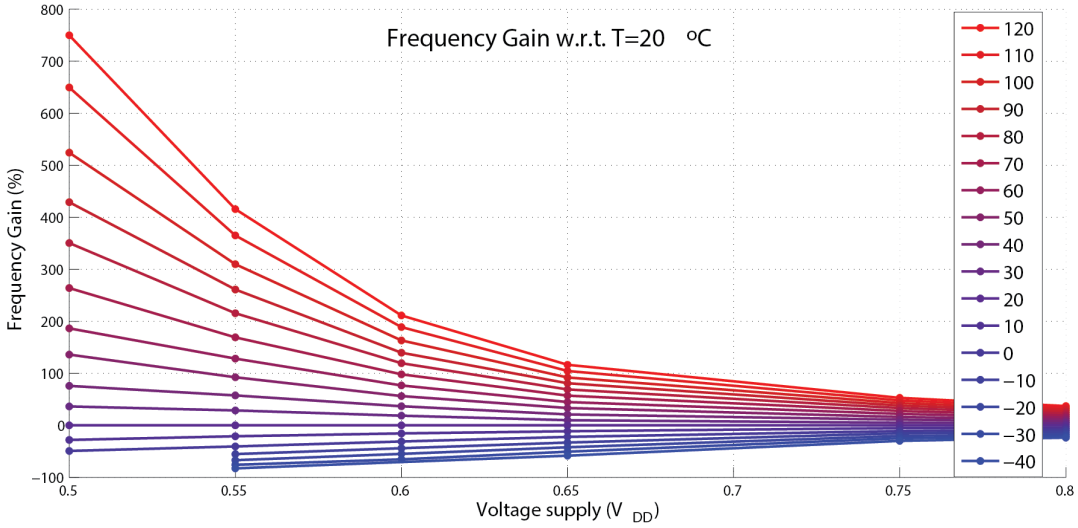

Figure 16: Effect of temperature on maximum frequency achieved by the FDSOI silicon prototype within the voltage range of 0.5V - 0.8V and temperature range of -40°C - 120°C.

Brand new technologies, like FDSOI are good candidates to be coupled with liquid cooling solutions. To this end, we have developed PowerCool, the first-of-its-kind mathematical model that performs electro-chemical modeling and simulation of integrated microfluidic power generation in full MPSoCs architectures. The accuracy of the model has been validated against fine-grained multi-physics simulations of flow cells in the COMSOL software that is unsuitable for EDA because of large simulation times. PowerCool model is demonstrated to be up to 425x times faster than COMSOL, and have a worst-case error of only 5%.
Micro-channel fluid flows usage with on-chip circulation (i.e. liquid cooling) has been proposed in the recent years as an alternative cooling mechanism to air cooling. Recent work has proven that it can be used effectively as a primary cooling technique in current server and prospective 3D architectures. In this prospective cooling technology for next-generation servers, a micro-channel-based cold plate is mounted on the corresponding computing unit, and single-phase water is injected to remove the heat generated by this computing unit. Fluids are injected in this design through micro-channels and, through forced convection heat transfer; the fluid absorbs the dissipated heat from the server processing units. This cooling technology reduces dramatically the cooling costs with respect to conventional air-cooling techniques [Sabri:2011].

11.1.2 Portable and Embedded Systems

"Cyber Physical System" (CPSs) describes a new generation of embedded devices, with the ability to seamlessly integrate with the environment. CPS platforms provide heterogeneous components dedicated to signal acquisition (e.g.: analogue-to-digital converters), processing (digital signal processors) and environment manipulation (actuators). They typically include wireless communication and run on batteries or energy harvesters.

To meet design goals within a small energy budget envelope, the design of each of the components of CPSs is highly tailored to the targeted application. In this direction, during its second year, Phidias has progressed towards this goal by exploring both domain-specific memory organizations and specific circuitry to optimize Compressive Sensing applications [Mamaghanian:2014].

We investigated the benefits of a domain-specific memory organization. The proposed hybrid solutions feature both 8-transistor and 6-transistor SRAM implementation, allowing ultra-low power modes within a single supply domain. The scheme is capable of normal operation at a high voltage supply level. At a lower level, it presents a fully reliable memory partition (implemented as 8T) while the rest of the memory (6T) is state-retentive.

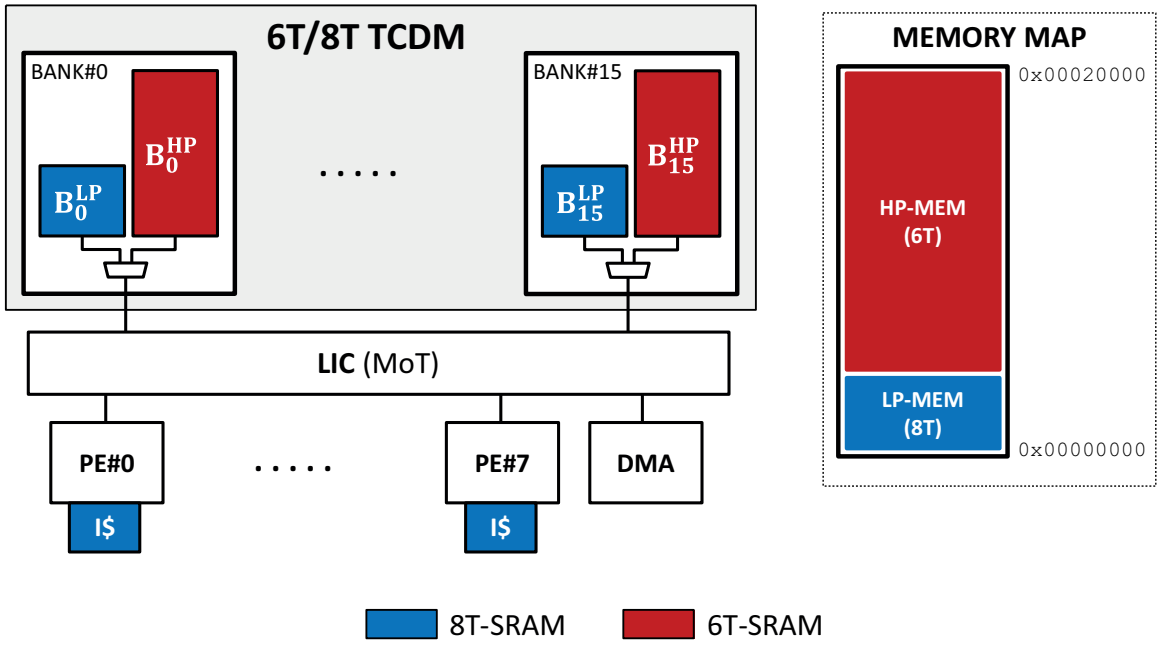

Figure 17: Hybrid 6T/8T memory architecture and the memory map.

The platform (see Figure 17) leverages the characteristic of CS-based applications, which are usually composed of two phases: data collection and computation. The first phase is characterized by low-workload/low-memory requirements and a long duration (LP Phase), while the latter consists of short bursts of high-workload computation (HP Phase). The proposed memory organization supports the two phases at two different supply voltage regimes, while minimizing the hardware overhead implicit in 8T-SRAM designs. The proposed solution improves energy efficiency of the system for the range 0-90% of HP/LP ratio. Considering a typical scenario with a 5% ratio between HP and LP phases, the proposed solution proves to be ≈ 25% more efficient than the baseline architecture with a homogeneous 6T-SRAM memory sub-system.

Regarding the design of specific circuitry to optimize operations, PHIDIAS has explored solutions beyond the state of the art leveraging the compressed sensing paradigm. More in the detail, a novel ultra-low-power Analogue Front End (AFE) has been prototyped. The AFE is an energy-critical circuit in both digital and analogue CS implementations that acquires the bio-signals. The AFE forms a significant percentage of the overall system power consumption. The digital computation and the memory power can be reduced significantly just by lowering the supply voltage (V_{DD}) to ≤ 0.5V. Compressive sampling can lower this power even further as it will reduce the amount of data. However, the AFE power will then dominate the system power consumption and needs to be reduced simultaneously.

The AFE comprises of an instrumentation amplifier (IA), filters and an analogue-to-digital converter (ADC). Although, in literature, low voltage, power-efficient ADCs and filters exist, a low voltage power-efficient IA does not. Therefore, the AFE power consumption cannot be reduced by lowering the supply voltage without trading off on IA precision. This constraint leads to the design of complicated power management circuits, which further reduce the power efficiency and are hence undesirable.

To overcome the challenges posed by the design of IA at lower supply voltages, an alternate paradigm of time-domain based amplification is implemented in the current work (details in [Benatti:2014]). The new ultra-low-power analog-front-end achieves comparable noise performance with the state-of-the-art solutions, yet being operational at an ultra-low voltage of 0.35V while consuming only 210 nW of power.

11.2 Fundamental Circuit Architecture Limits

This section looks at a set of 'laws', or long-holding observations with respect to computing devices. These laws have a significant effect on micro-architectural design decisions, which will be discussed here.

Moore's Law is perhaps the most famous law in ICT, having predicted a 2x increase in chip density every two years. However, properties previous attributed to Moore's Law, such as aforementioned operating frequencies, have not necessarily held. These pose potentially more interesting challenges that will run alongside the so-called “Beyond Moore” and “More than Moore” efforts discussed in Chapter 10.

In 1974, Robert Dennard [Dennard:1974] observed a set of scaling characteristics for circuit performance. Dennard's scaling factors, much like Moore's Law, have a significant impact on how we implement chips. The most problematic factors currently, are threshold voltages and interconnect. Threshold voltages no longer scale with new technology nodes, leading to the search for novel near- [Zha:2009] and sub-threshold [Kau:2012] device implementations examined in other chapters of this research agenda. Interconnect scaling has a significant impact on micro-architecture, as it governs the speed at which data can be moved around a chip. In a micro-architecture with many cores, each with multiple functional blocks, the interconnect can become a bottleneck to performance.

Whilst transistors have benefited from performance increases as they have shrunk, interconnects have not, their delay product remaining close to constant [Bohr:2007]. The interconnect now plays a large part in the delays present in memory hierarchies. Thus, the structure of cache hierarchies – something that is very much a micro-architectural choice – is governed not just by the speed of the storage cells, but the delays in moving data between those caches and the processor.

Fast, first-level caches cannot be large, because it becomes impossible to transport data across them and still close timing constraints. Thus, modern high performance processors have multiple cache layers, increasing in capacity and latency as they get further from the computational part of the processor.

The amount of energy consumption from a circuit architecture design for a given CMOS technology node is heavily dependent on how specific (i.e. optimised for a single or a few tasks) or how general (i.e. undertake many different computations) a design has to deliver. Applications Specific Integrated Circuits (ASICs) designed for a single task can be optimised proving the lowest energy consumption but such designs have no flexibility and cannot be reprogrammed. For microprocessors or microcontrollers that must be able to undertake a wide range of tasks, optimisation to reduce energy consumption is significantly more difficult.

Historically the reduction of transistor dimensions from the scaling through Moore’s law resulted in the power dissipation reducing as each new technology node became available. As static power dissipation increased and became larger than dynamic power dissipation for the recent technology nodes, scaling no longer offers a reduction in power if the architecture remains constant. The ultimate limits from architecture designs are almost impossible to derive but there is general agreement by academia and industry that new architectures are more promising to reduce power consumption than improving the energy consumption of the basic switching device in the circuit.

11.3 Key Challenges

Advances in physics, transistor design and device manufacturing techniques can benefit micro-electronic devices of all kinds. However, micro-architectural design decisions are heavily influenced by the target market of the resultant product. Whilst all devices strive to achieve good efficiency, in terms of both performance and power, the application area will dictate design constraints such as size and maximum power.

For example, in HPC, power dissipation per-server must be considered in order to adequately provision the electricity supply as well as extract the waste heat. To that end, HPC computer architectures such as server processors and MPUs, work within power constraints of in the order of tens or occasionally hundreds of watts. However, an embedded device, that must be kept in a self-contained environment and run on a battery for ten years, will have a significantly smaller constraint imposed upon it.

This section will provide an overview of micro-architectural characteristics grouped into four areas: deeply-embedded, embedded / mobile, general purpose and servers / high-performance. These are not necessarily strict boundaries, and properties often transfer between areas over time, as technology or commercial pressures permit.

Deeply-embedded

These devices may operate under one or more of the following constraints:

Physical size must be small, from millimetres to a few centimetres, depending on application.

Power may be intermittent or limited, often sub-watt or sub-milliwatt.

Operating temperatures may vary significantly and the removal of excess heat quickly may not be possible.

Reliability is essential, because servicing may be difficult, expensive, or impossible.

Predictable behaviour may be required to guarantee safety criteria or always-correct device functionality.

An example of a deeply-embedded device is the processor in a smart card. It must fit within a credit card form factor, cannot be modified once sent to the customer, be powered by a battery-backed device and obey strict security protocols. All of the above points are relevant in an energy context. The energy consumption of the device dictates the temperature it runs at, the amount of cooling apparatus surrounding it, how long it will live. And of course, predictable energy consumption is required to guarantee a particular battery life.

The micro-architectures of deeply embedded systems take various forms, but the following traits are common:

They provide predictable execution times for many or all of the ISA instructions they support.

Their functional blocks, such as arithmetic and memory units, are often simpler than larger counterparts, to reduce power, improve predictability, and keep the device small.

Their memory hierarchy is flat, avoiding caches that would impact predictability and increase device complexity.

Programs may execute directly out of integrated flash storage, with RAM only used for read-write data.

Prolific architectures in this area include AVR, ARM's Cortex M-series and PIC. They feature compact instruction sets (often 8- or 16-bit instructions), with short execution pipelines and in-order execution. This means opportunities for performance enhancement are limited, but their implementation is simple. The relative simplicity of such devices aids activities such as modelling their behaviour in order to predict energy consumption. However, properties such as the memory layout limit the types of application that can feasibly be run on such devices.

Looking forward, we can expect the safety and real-time requirements of deeply-embedded systems to persist. However, research must strive to shrink these devices into sub-millimeter and beyond, with a desire for micro- or nano-watt power envelope devices with comparable performance today, whilst milli-watt devices deliver improved performance and capabilities. Multi-core is not yet common in deeply-embedded systems, but in order to deliver the above improvements in the light of limits to Moore scaling, we can expect it to become essential.

A particularly lucrative opportunity in deeply-embedded research is a “zero power” idle or sleep mode, with close to instant response time. When not responding to an event, the device should, ideally, consumer no energy at all. However, following an event (such as inbound sensor data), it must be able to return to a responsive state.

Embedded / mobile

The line between embedded and deeply-embedded is often blurry, but for the purposes of this report, we group *regular* embedded devices with mobile devices, in order to set the grouping in terms of energy requirements. Such devices may still have real-time constraints, small size requirements and sub-watt power envelopes.

An example of an embedded device is the controller chip on a hard disk or solid state disk. They cannot be replaced, so their failure effectively renders the entire disk useless. The software running on them is difficult or undesirable to update in the interests of data security. They interact with components that have strict timing protocols, and missed deadlines will potentially result in data loss or corruption. They must fit within the form-factor of the device, for example the whole device may be as small as 12x16mm (the smallest of the possible M.2 format of expansion devices used in laptops at the time of writing). However, the performance requirement and available power are higher than deeply embedded devices, hence this may be considered embedded rather than deeply-embedded.

Mobile and embedded systems typically have the following micro-architectural properties:

Heavily integrated into System-on-Chip (SoC), providing various peripherals and multi-core computational capabilities in a single chip.

Larger storage and memory capacity than deeply-embedded, in the order of Gigabytes in current devices.

Cache hierarchies for better memory performance.

Sub-watt power constraints.

Must fit within relatively small form factors, such as a mobile phone.

More complex application sets and usage patterns.

Mobile and modern embedded devices feature performance that is competitive with general purpose hardware from less than a decade previous. A contemporary smart-phone has more compute power than a ten year old desktop PC and consumes significantly less energy. Much of this is thanks to lower-level improvements, as described by trends such as Moore's Law and Dennard Scaling.

If one is to expect the same to be true in another decade, then we must counter abstraction inefficiencies as systems become more complex. For example, May's Law states that software efficiency reduces to counteract any improvement in hardware efficiency. At the same time, as devices become even more integrated into lifestyles of consumers, both users and app developers must be given better visibility of how and why energy is consumed by their apps. This transparency will encourage accountability for embedded software energy efficiency, and narrow the gap between the best efficiency a hardware platform is capable of, and the efficiency achieved when running a particular set of applications.

General purpose

Global use of computing devices is shifting to a more mobile-centric approach. As such, the features of mobile devices often compete with those in more traditional general purpose devices such as the desktop PC. However, general purpose computing is less constrained than mobile and embedded computing, and this is reflected in the micro-architecture:

Power envelope of tens of watts

Multiple layers in the cache hierarchy (three-layer caches or more).

Less tightly integrated, with separate physical components for RAM, peripheral devices, etc.

The processor architecture may be similar to mobile devices. For example, both x86 and ARM instruction sets are used in desktop and mobile computing products. However, general purpose micro-architectures that use these instruction sets may be more complex, with increased pipeline length, more functional units, multi-threading capabilities and higher operating frequencies.

The trends in device usage suggest that general purpose computing will merge with mobile computing, backed by servers such as those provide cloud services. To that end, the general purpose computing category may eventually disappear, rather embedded/mobile will become the de-facto general purpose computing platform. Thus, research into low-energy micro-architecture may create more benefits if it focuses on embedded and server related areas.

Servers / high performance

The properties of high performance and server devices are similar to general purpose, but their form factor and tolerances are different.

Multiple servers may occupy dense racks.

Power envelopes may be higher than general purpose, due to more aggressive cooling.

Underlying architectures are similar to general purpose, typically at leading-edge, with additional resiliance features such as error correction.

The biggest difference in the server class is that amassing the cost per device is somewhat easier – a company may have many servers, and so that company has an interest in managing the total cost of ownership (TCO), including energy and cooling bills. Of course, service platforms such as the cloud simplify this for many companies, but ultimately this shifts the management responsibility to the cloud provider. The servers, whoever manages them, are densely packed, with processors each consuming tens of watts, if not more.

Mobile architecutres such as ARM are beginning to encroach into the server space, with examples such as Cavium's ThunderX processors and APM's HeliX devices, both of which use the 64-bit variant of ARM for server-grate compute and infrastructure roles. The continued proliferation of multi-core supports this movement, and it is likely we will see lower power per device, but packed at a higher density, so that overall the power per rack is still a challenge [Rasmussen:2005]. The interconnect between these systems also becomes of particular research interest, as the cost of data movement between these devices does not scale with the reduction in energy per processing unit.

Summary

The following Table 4 summarises the important properties of this section, capturing both the current state of

devices in the four defined classes, as well as the expected progression in the future. Our research goals over the next decade must seek to facilitate these progressions.

Feature	Deeply-embedded	Embedded / mobile	General purpose	Server / HPC
Physical size	Millimeter chips on very small systems.	Millimeter scale chips, centimeter scale systems.	Centimeter chips fitting desktop form-factors.	Centimeter chips fitting rackmount form-factors.
Future	*Sub-millimeter or less.*	*Similar scales.*	*Similar.*	*Smaller devices to increase density.*
Power availability	Milliwatts or less, from battery or energy harvesting.	Sub-watt	Tens of watts	Low hundreds of watts.
Future	*Micro- and nano-watt power, intermittent from harvesters or low capacity batteries.*	*Milli-watt*	*Tens of watts or less.*	*Similar, but less per device density increases.*
Heat	Potentially limited methods for removing excess heat.	*Pocket-friendly* heat removal constraints.	Airflow-based heat removal, low noise for office environments.	High-airflow or water-cooling.
Future	*Similar.*	*Similar.*	*More limited airflow as devices become increasingly mobile-like.*	*Increased use of water-cooling, higher ambient operating temperatures.*
Performance	MIPS	GIPS	10s of GIPS	TIPS
Future	*10s of MIPS or more.*	*10s of GIPS*	*Similar.*	*Many TIPS*
Memory	Kilobytes. Flat hierarchy.	Megabytes or gigabytes. Caches for performance.	Tens of gigabytes. More cache layers and virtual memory.	Tens or hundreds of gigabytes, NUMA.
Future	*Megabytes or more. Access must remain predictable.*	*Many gigabytes.*	*Potentially hundreds of gigabytes. Virtual memory complexity increases.*	*Hundreds of gigabytes or more. NUMA complexity increases.*
Integration	Limited components in a single chip or die.	SoC with multi-core and many peripherals.	Less than mobile; more modular as a result.	Similar to general purpose.
Future	*Increasing amount of integration, similar to SoC.*	*NoC, vast heterogeneity.*	*Will increasingly move towards mobile-like approaches.*	*Tighter integration as density increases and to enable higher performance.*

Table 4: A summary of the architectural properties and considerations by device class.

11.4 Opportunities for Circuit Architecture Improvement

At the micro-architectural level we foresee three main directions of progress:

Increased heterogeneity, with specialised functional blocks for particular tasks.

Many-core systems, with various heterogeneous blocks as described above, combined with groups of homogeneous clusters.

Network-based interconnects or complex, multi-layered buses, to connect these many components, along with caches, memory and peripherals, together.

All of these pose scalability and programmability problems, some of which must be addressed at the micro-architecture level, whilst others can be dealt with in other levels in the stack. A number of major challenges arising from the above directions will now be posed. These will form at least part of the necessary research agenda for micro-architecture in the coming decade.

Heterogeneity

Specialisation of compute resources is emerging as a requirement for very low energy computing and so micro-architectures must adapt to this [Borkar:2011]. The outcome is heterogeneous SoCs with many blocks, each dedicated to doing a particular task well.

These types of platforms pose several research questions:

How do we choose which blocks to make available in a given micro-architecture?

In a system with a strict power budget, how do we control how many, and which blocks, can be active simultaneously?

How should these components be presented for programming? For example, what is a good ISA for each block?

Can workloads be easily migrated between blocks?

Under what access model do we enable the movement of data between blocks, and at what granularity?

Many-core

In some respects this is an extension of heterogeneity's research challenges, but with some considerations for grouped homogeneous resources.

What are efficient methods of sharing a variety of workload types in parallel on a group of functional blocks?

In a very large system, subject to challenges such as process variability, how do we guarantee consistent or predictable behaviour across many cores?

At what granularity do we apply groupings, and manipulation of their behaviour, such as voltage and frequency scaling?

Communication

Underpinning micro-architectures, the ability to move data around the system is essential, but its limited scaling poses significant challenges:

What topologies provide a good balance between connectivity, predictability, low energy and high performance?

How do we trade-off the cost of moving data to a specialised functional block, compared to the cost of operating on it with a more local, but less efficient block?

How do we make micro-architectures expose the energy cost of communication to higher abstraction layers, to aid scheduling, allocation, or even programming strategies?

In shared interconnect systems with security or mixed-criticality concerns, how do we avoid low-priority or unprivileged communication flows from adversely affecting critical-priority or secure data?

These are largely unanswered research questions, into which further effort must be invested.

Huge amount of Industrial R&D efforts have been invested in new technologies to guarantee further silicon CMOS process scaling, beyond the limits of traditional MOS devices [ITRS:2013]. These efforts have produced very successful breakthroughs, such as FDSOI transistors [FDSOI:2012]. FDSOI is giving fast payoffs, and it is proving itself the most likely winner for the next couple of technology nodes [FMOSF:2012]. Thanks to the insertion of the ultra-thin buried oxide, FDSOI brings a number of fundamental improvements in the transistor electrical characteristics, while it still leverages the very mature planar technology for chip fabrication, such as: up to 35% faster operating frequency for the same power consumption at high voltage, and up to 100% faster at low voltage. FDSOI features lower power consumption, and FDSOI process is simple (does not require stressors or similarly complex techniques used in other processes). New memories and logic components will benefit from the use of this new technology in the upcoming years, which will enable higher-level architectural solutions.

At the micro-architectural level, energy efficiency advancements must take the form of a synthesis of improvements at lower levels, and respond to the evolving needs at higher levels. We must strive to provide predictable, transparent behaviour, not just functionally, but for properties such as energy consumption. The cost of data movement must be more readily exposed, and significant effort must be put into efficiently

moving (or avoiding moving) data around a system, as its contribution to energy consumption will continue to increase. This must be complemented by novel, intuitive methods for presenting the increasingly heterogeneous resources that form emerging multi-core systems across the various device classes, so that higher levels of the system stack can fully understand and exploit the underlying hardware.

12 System Architecture

The main challenge for the future evolution of chips is clearly the energy scalability problem. Semiconductor technology has indeed hit a fundamental efficiency limit, called "energy wall" that prevents reduction of energy consumption when transistor size scales down in current manufacturing technologies for forthcoming technology nodes. Therefore, disrupting architectural innovations will be necessary to push the evolution of chip design and keep the improvement pace and compensate for this technology shortage.

12.1 Present and Future System Architecture

Similarly to the scenario of circuit architecture, system architecture will require from multidisciplinary research efforts to achieve holistic energy-efficient design and management. Advances in computer architecture must be carried out in two information propagation directions: The first direction is bottom-up where technology-related parameters (e.g.: FDSOI, Stacked-DRAMs, etc) are propagated from new technologies and chip level to the large-scale datacenter level. In particular, circuit-level technological parameters will be used in the server and datacenter architecture explorations, while the chip level cooling parameters will be exploited by the large-scale cooling and energy reuse technologies. The second direction is top-down where the target large-scale datacenter operation characteristics (in terms of run-time workload variations at software level and infrastructure operating conditions), will be exploited to tune further the lower level architectural, cooling and technological design aspects.

12.2 Key Challenges of System Architecture

In both mobile and HPC microprocessors there are conflicting demands for high performance and energy efficiency. Circuit architectures to deliver these conflicting performance requirements are being addressed through the heterogeneous integration of a range of embedded cores. One example is the ARM big.LITTLE architecture [BIGLITTLE:2013] where smaller cores are employed to process simple, less demanding tasks to save energy whilst larger cores are optimised for the high performance tasks which are more energy demanding. This is a general trend and, in parallel to it, there are a wide range of complementary techniques that are either being used in commercial processors or are being researched in academia to reduce the power consumption through architecture design. These include:

- Dynamic voltage and frequency scaling: high voltages and frequencies are used for high performance task delivery whilst low voltages and frequencies deliver reduced energy consumption [Hunumaiah:2012].
- Clock gating and clock distribution: this is a dynamic power management technique where additional logic switches are placed between the clock and the clock input of the processors logic to disconnect the logic preventing clock cycling when it is not required [Bassett:2012].
- Power domains: power domains are sections of a core in a processor that can be completely powered down to reduce energy consumption without removing the supply to the system [ARM:2013].
- Pipeline balancing: is the dynamic adjustment of the resources of the pipeline of a processor such that it retains performance whilst reducing the power consumption [Bahar:2001].
- Caches and interconnects: it has been demonstrated that too large caches waste energy whilst too small caches limit bandwidth and performance. Careful optimisation is therefore required to both minimise energy consumption and maintain performance [Zeng:2008].
- Dynamic partial reconfiguration: A number of floating point gate array (FPGA) manufacturers now offer dynamic partial reconfiguration which enables reconfiguration of parts of the FPGA whilst the other regions are still active [Jooya:2009].
- Composable and partitionable architectures: this is where a set of low power small cores can be aggregated together dynamically to form a larger single-threaded processor when required for higher performance [Changkyu:2007].

- Coarse grained reconfigurable array architecture: are designed to be reconfigurable at the module or block level rather than at the gate level in order to trade off flexibility for reduced reconfiguration time [Rakossy:2012].

At a larger scale, we present some of the new architectural solutions that are being investigated in the context of datacenters and HPC in European projects:

- Advanced cooling infrastructures and energy reuse:
  - A passive thermosyphon (gravity driven) cooling system for servers and racks, using energy only to remove the waste heat out of the datacenter room.
  - Innovative systems able to re-convert generated heat into electricity, such as, the pressure reverse osmosis (PRO) process. The objective is to provide datacenters a solution to absorb a part of their waste heat and re-produce electricity.
- Datacenter thermal control and management
  - Resolution of complex run-time multi-objective optimization (performance, power and temperature) in high-level IT workload management (allocation and scheduling) and physical resource management (processors, memories, storage, network and cooling infrastructure configuration).
  - Exploration of several optimization schemes and control approaches such as adaptive dynamic programming, networked control, and predictive control at the large-scale of datacenter level.
- System-level energy-efficient datacenter architecture design:
  - Definition of more efficient system-level exploration approaches to develop energy-efficient thermal-aware architectures for servers, racks and datacenter rooms by using a holistic and tighter integration of computing and cooling power costs.
  - Development of scalable simulation methods of large-scale computing systems that integrate power and thermal modeling to help datacenter designers make robust predictions to anticipate potential failures of components, and to accurately estimate the necessary provision in global energy during datacenter conception and throughout its lifetime.

12.3 Fundamental Limits of System Architecture

In MPSoC design, joint architecture optimization and integration of low-power components in novel architectures is the only way to keep increasing the performance while staying under the power budget. Solutions like application specific and heterogeneous memory hierarchy, for instance, heterogeneous 3D architectures (stacked-DRAM) or efficient hardware implementation of components (analog and digital) for ultra-low-power sensing. On top of this, all the on-chip components (both logic and memories) are increasingly affected by process variability, which means that they can no longer operate always under the best conditions, requiring self-adaptive architectures. Therefore, this scenario makes more important than ever to develop scalable functional simulation frameworks to explore the limits of parallelization and global power reductions in MPSoCs.

In the field of HPC, as volume server costs drop, electricity costs will emerge as a substantial fraction of the overall cost of ownership in datacenters. The fundamental question beyond the state-of-the-art is how to bridge the efficiency gap between emerging data-centric and scale-out workloads and modern server and network platforms. There is a large efficiency gap between existing server architectures and what the emerging data-centric scale-out workloads need.

12.4 Opportunities for System Architecture

In order to evaluate new solutions for ultra-low-power heterogeneous architectures, we finalized the design of our own multi-core cycle-accurate virtual platform simulator. It enables cycle-accurate simulations and precise power estimations of the target digital architecture using back-annotated power consumption values from low-level syntheses and post-layout analysis.

While post-layout simulations can indeed offer precise energy consumption estimation, they also require a lengthy synthesis and simulation process, making them impractical for early evaluation and/or design space exploration. Cycle-accurate simulations retain a precise evaluation of run-time characteristics while allowing

faster simulations. This aspect is of great relevance for embedded WBSN applications, as the input bio-signals have slow dynamics (e.g., in the example of heart monitoring, the normal heart rate ranges from 60 to 100 beats-per-minute), thus requiring long simulations to capture the energy efficiency and workload balance of different architectural configurations.

The virtual platform supports multiple computing cores, whose architecture and compiler have been developed with the Lisa Tool-suite from Coware [LISA].

The availability of this tool allowed us to explore domain-specific designs of architectural elements like, for example, investigating the benefits derived from application-specific memory subsystems.

At a higher level, we have also developed a framework for rapid evaluation of CS architectures, complementing the cycle-accurate virtual platform. This architectural abstracted model has been developed in matlab and aims at providing high-level insights of the benefits of different task-allocation schemes, in terms of energy and performance, in digital and analog CS nodes. With this framework we studied the problem of allocating tasks on the PHIDIAS target architecture subjected to process variability, focused on the study of tasks allocation and shutdown policies for Compressive Sensing application in parallel architectures and to comparatively assess the performance of different CS acquisition strategies under energy constraints.

As an example of the potential of this type or architectural simulators, Figure 18 summarizes a small study that experiments with architectural variations based on the Ultra-Low-Power (ULP) TamaRISC cores. The performance of a single-core general purpose TamaRISC system is compared with the one of a multi-core implementation employing up to eight cores, presenting advanced interconnects and synchronization strategies to allow energy savings through Single Instruction Multiple Data (SIMD) operations. In addition an Application Specific Instruction-set Processor (ASIP), derived from TamaRISC, and dedicated to the CS application, can effectively increase the energy efficiency, both in single- and multi- core configurations.

The two approaches are orthogonal, and they both result in substantial energy savings. In fact, even in the case when ASIP cores are employed, a multi-core solution is more efficient than a single-core one when multiple signal windows are processed concurrently (see Figure 18).

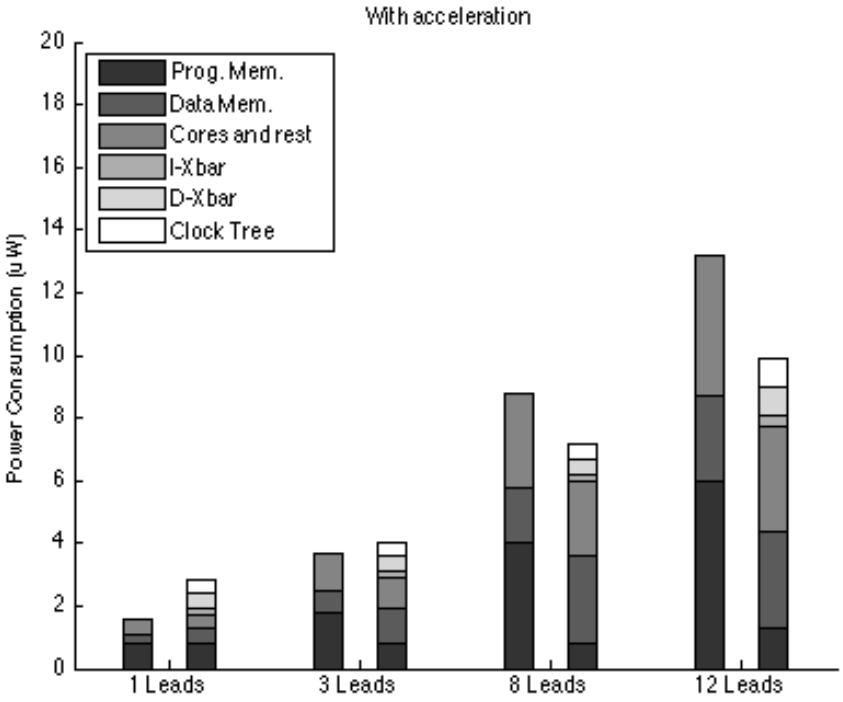

Figure 18: The power consumption of single-core and 8-cores TamaRISC-CS systems, for different numbers of concurrently acquired signals.

In the field of many-core servers (HPC), we have also been exploring new architectural solutions, like die-stacked DRAM caches, to improve latency, bandwidth and power characteristics for in-memory computing. As a result we introduced a novel stacked-DRAM cache design, the Unison Cache, that

incorporates the tag metadata directly into the stacked DRAM to enable scalability to arbitrary stacked-DRAM capacities. Unison Cache employs large, page-sized cache allocation units to achieve high hit rates and reduction in tag overheads, while predicting and fetching only the useful blocks within each page to minimize the off-chip traffic. Our evaluation using server workloads and caches of up to 8GB reveals that Unison cache improves performance by 14% compared to Alloy Cache due to its high hit rate, while outperforming the state-of-the art page-based designs that require impractical SRAM-based tags of around 50MB.

13 Software

13.1 Present and Future Software Performance

Software controls the amount of energy consumption in a system in a variety of ways. Possibly the most significant is that many facilities in a system operate at the request of software, meaning that processor cores and system level consumers such as radios or displays can consume more or less energy depending on software algorithms. Accepted wisdom states that because of this, the best energy efficiency is achieved by having all software operate as quickly as possible, meaning that facilities can be disabled (minimizing their consumption) for longer periods. This is known as reducing static power.

In addition, software also causes energy consumption through dynamic power, the cost of circuits charging and discharging as they signal digital information. Examples include the activation of different banks of memory according to access patterns, driving of data buses, and switching activity as gates in arithmetic units converge on an output. While static power tends to be controlled by the algorithm behind a piece of software, dynamic power is governed by its particular implementation. It's relation to higher level programming constructs tend to be poorly understood.

As documented elsewhere (section 10) silicon manufacturers have been unable to deliver speed improvements in the last decade that follow the trends of the past. In the realm of software, this is known as "The free lunch being over" [Sutter:2005]: speed improvements will demand effort from the software developer in the future. Correspondingly, we cannot expect that hardware will deliver better energy performance in the future, and so the developer must contribute to energy reduction too.

An obvious barrier at present is that very few software developers have much idea of how much power their programs dissipate or which parts of a program are energy hotspots. This might even be different for the same program from one platform to another. There are good software engineering reasons for the programmer's ignorance of energy, namely, to allow programs to be ported to different platforms and to allow program design as higher levels of abstraction. The large conceptual gap from hardware, where energy is consumed, to high-level languages and programming abstractions has been created by decades of computer science research and compiler advances. Somehow, the developer using a high-level language has to understand the energy induced by the software at the hardware level, without having to measure it on a machine.

A clear theme emerges, that the developers of the future will rely less on the performance of processors, but instead be energy-aware. This means that they tune their software to work optimally on the available hardware, or, if feasible, choose hardware specific to the software application. The growth in popularity of heterogeneous systems HIPEAC Workshop on Energy Efficiency with Heterogeneous Computing IEEE International Conference on Application-specific Systems, Architectures and Processors may provide many options in the future. There is a lower bound on how much energy software can consume, since it is not always possible or feasible to choose hardware specific to a particular problem, and thus “energy-optimality” is defined relative to available hardware.

13.2 Key Software Challenges

Given a program to be implemented in some programming language and a hardware platform on which it is to be executed, we may ask whether it is energy-optimal. The energy limit for software is a relative and pragmatic one; what is the least energy a given program should consume on a given hardware platform? We

assume that in answering this question the program could be redesigned to use a more energy-efficient algorithm (for that hardware platform).

Software designers and developers typically first target functionality (getting the program to do what it is supposed to do), then performance (doing it as fast as possible) and thirdly minimisation of code development costs (productivity). Optimisation of performance is highly important in some fields, especially in HPC. This usually means optimising for minimising the time of a computation (i.e. optimising performance and high speed). However, there has not been significant work on optimising the energy consumption or understanding the ultimate energy consumption limits. This provides significant opportunities for optimising energy consumption. LSI Logic suggest that dedicated low power hardware circuitry may save 20% whilst changes to software to better control the power states could provide power savings of a factor of 3 to 5. In short, more energy is wasted by software than by hardware.

In a general-purpose piece of software, reducing energy consumption will mean reducing the amount of static and dynamic power the application draws. This process requires two distinct steps: that of identifying the sources of each kind of power, and then reducing that amount of consumption. Identifying static power consumption corresponds to evaluating the amount of time taken for the program to execute, while dynamic power requires either analysis of the machine instructions that the program compiles to, or some (more or less approximate) model of the energy usage of higher-level software. Both of these techniques are complex, and not immediately available to modern software developers, limiting their ability to make decisions to reduce energy.

The key research challenge is to bridge the conceptual gap and make energy consumption transparent through the layers. When programmers are energy-aware, there are a number of measures that can be taken to optimise for energy efficiency of software [Roy: 1997]:

1. Choose the best algorithm for the problem at hand and make sure it fits well with the computational hardware. Failure to do this can lead to costs far exceeding the benefit of more localised power optimisations.
2. Minimise memory size and expensive memory accesses through algorithm transformations, efficient mapping of data into memory and optimal use of memory bandwidth, registers and cache.
3. Optimise the performance of the application, making maximum use of the available parallelism.
4. Take advantage of hardware support power management.
5. Select instructions, sequence them and order operations in a way that minimizes switching in the CPU and datapath.

To achieve energy transparency, and hence enable energy-aware software development, advances are required in two major areas: energy modelling of software, and software energy analysis methods.

13.2.1 Energy Modelling of Software

Tiwari and Malik [Tiwari:1996] have suggested a method for analysing the energy consumption and energy optimisation of software using instruction level power analysis where they define the energy cost of a program, E_p as

$$E_p = \sum_i (B_i \times N_i) + \sum_{i,j} \left(O_{i,j} \times N_{i,j}\right) + \sum_{ik} E_k$$

where the first term is the instruction base cost, B_i of each instruction, i occurring Ni times; the second term is the circuit state overhead, $O_{i,j}$ for each instruction pair i and j; and the final term is the other instruction effects (stalls, cache misses, etc.). The model is not accurate for all computational systems and software and in particular there are many examples where input/output and protocols govern activity and energy consumption. For example, in sensor and network systems, the software can spend a significant proportion of time waiting for a sensor or server to send appropriate information. There is therefore significant opportunity for developing accurate energy consumption software, which will allow better predictions so that software can be optimised to minimise energy consumption.

One of the biggest problems and therefore a significant opportunity if addressed is that the hardware has to guess when it can use a low power state since at present the software frequently takes too long. This is due to the fact that software designers do not understanding hardware and the fact that hardware states are not easily accessed at the software level. There is therefore significant opportunity if the software designers understand how the software influences the energy consumption of the energy-saving features of the hardware. This is key to accurate software energy modelling and optimisation on modern processors.

Modelling of the other instruction effects can be challenging too: matters such as cache behaviour can depend on uncontrolled memory layout, and branch prediction may be affected by input data patterns. The effects of these features depend on attributes of the software and it's environment (such as memory locality), and the manner in which hardware reacts to these attributes. Modelling these attributes is an open problem.

13.2.2 Software Energy Analysis

The software analysis problem is essentially to determine how many times each basic unit is executed. An energy model (such as Tiwari et al.'s model outlined above) associates energy costs with each basic computational unit of the program (for instance, the expressions Bi and Oi,j in Tiwari et al.'s model). The task of static analysis thus is to determined the values of Ni and Ni,j in Tiwari et al.'s model in order to compute the energy. These values are dependent on the input of the program and possibly other context of the execution.

One approach to analysis is dynamic: the program execution is simulated (with a range of input data) and the number of times each unit is executed is logged, giving an approximate average value for the counts. This approach is limited solution in coverage, scalability and accuracy. A more comprehensive approach is given by static analysis based on abstract interpretation.

The overall goal of static program analysis is to derive information about a program's execution without executing it, relying only on the semantics of the programming language. Modern compilers already incorporate static analysis in their optimisation stages, for example to detect live variables and dead code, and to perform optimisations such as those based on common sub-expressions, available expressions and constant propagation.

In recent years, resource analysis has been developed as a branch of static analysis using abstract interpretation, focussing on non-functional aspects of a program's execution such as time and memory consumption. Resource consumption can be derived as a function of the size of the input data, in much the same way as the complexity of a program is usually expressed using "big-O" notation as a function of input data size. Static analysis using abstract interpretation inherently involves approximation; thus resource analysis yields upper and lower bounds on resource usage for a given input data size.

Software energy analysis faces many challenges; static analysis inherently involves a trade-off of precision against complexity of the analysis. Obtaining tight bounds of energy usage depends on several factors, including accurate propagation of data size measures and extraction and solution of the relations expressing energy consumption in terms of data sizes. Both of these problems are solvable for a large class of useful programs, but if program structure departs from standard patterns, precision may be rapidly lost.

Furthermore, although upper and lower bounds on energy consumption are useful, information about the distribution of consumption within those bounds is even more so. For example, it may be that most execution cases of a program result in consumption close to the lower bound, while the upper bound is reached only in a few outlying cases, or vice versa. From the distribution, estimates of average energy consumption can be derived. One approach to obtaining this kind of information is to perform probabilistic static analysis of a program with respect to its energy consumption. This is a special case of probabilistic output analysis, whose aim is to derive a probability distribution of possible output values for a program from a probability distribution of its input. The output in this case is energy consumption.

Complex software is extremely difficult to analyze, especially due to non-determinism arising from threads, caches, pipelines, races and operating system effects. Static analysis of multi-threaded code is especially difficult since precision is easily lost due to thread interleaving. Accurate analysis of the timing

and synchronisation behaviour of threads is a prerequisite for energy analysis in which the energy of an instruction depends on how many threads are active simultaneously. Finding good trade-offs of complexity and precision in analysis of complex software systems.

13.3 Fundamental Software Limits

It seems unlikely that the energy consumption of software is subject to fundamental limits in the same sense as hardware. Software is written using programming languages and translated into codes that are executed by hardware. The translation and execution process is separate from the software and depends on tools such as compilers, interpreters and schedulers and the hardware platform. All of these may influence its energy consumption and can be varied independently of the software itself. Thus one cannot expect absolute energy limits associated with software. In addition, there could be many different programs that have the same functional behaviour – which of these is the most energy-efficient? The answer to this question also depends on the execution context outlined above as well as the algorithmic complexity of the different programs. As mentioned above, the energy limit for software is thus a relative and pragmatic one; what is the least energy a given program should consume on a given hardware platform? We assume that in answering this question the program could be redesigned to use a more energy-efficient algorithm (for that hardware platform). One may also ask whether there is a more suitable hardware platform available.

The lower energy bound limitation stems from the generality of hardware. The main benefit of software is that it is reconfigurable: one can take a general-purpose processor and deploy many different applications to it, perform field updates, and otherwise alter behaviour without altering hardware. This necessitates that the processor is substantially more generalised than the application, to allow for reconfiguration. It must have a large enough set of behaviours (i.e., be Turing complete) to be programmed, as well as having sufficient performance and resources to meet application budgets for time and space.

This leads to a significant capability gap between hardware and software. On one hand, the hardware must have a high capability to allow its reconfiguration for different applications, but on the other hand any particular application will only require a subset of those capabilities. DVFS and gating techniques allow disabling un-needed capabilities to some extent, while heterogeneous systems can provide application specific acceleration facilities.

Ultimately, the most energy-efficient implementation of an application is one that has dedicated hardware suited for the application and context, a solution that is often economically infeasible. More generalized hardware leads to a lower cost, but typically greater energy consumption. Developers can adjust their software to be more efficient on their chosen platform, however this requires detailed understanding and creativity [XXX xref to challenges].

13.4 Opportunities for Software Improvement

The most significant factor in improving the energy efficiency of software is the developer. It is currently impossible for automated processes to fully customise an algorithm to take full advantage of available hardware features, and thus it is up to the creativity of the developer to do so. To create the best opportunities for efficiency the developer should be provided with as much information about their software's energy consumption as possible: at all levels, from the switching cost to static power and system-level energy consumption.

Enabling the developer in such a way will allow for energy-aware exploration of the design space, allowing informed decisions to be made so that the developer can pursue minimal energy. Such benefits are not limited to individual applications either: developers of software libraries, compiler optimisations, code schedulers and any other software facility will be able to make energy-orientated design decisions.

Better software design for hardware should also enable better selection of hardware. Understanding of precisely how a piece of software consumes energy would allow more informed selection of heterogeneous systems, particularly if the heterogeneous hardware's features could be characterised in a way that can be matched against software features. Such a process would greatly simplify the benchmarking and evaluation required when selecting a platform to develop on.

Continuing on this theme, improving the developers ability to target well suited heterogeneous systems will create more opportunities for energy consumption reduction. The most obvious barrier on this front is the challenge of concurrency: numerous computational hardware accelerators are available [Buchty:2012] which deliver performance through parallel execution of tasks. Developing concurrent parallel, however, is notoriously difficult. Improved techniques, language support and tooling to ease its development will make such accelerators more accessible.

In summary, new paradigms and tools for the design of software, based on energy transparency, are required if the energy consumption is to be reduced. Currently the main drivers for software production are high performance, minimising time for operation and minimising production cost. These drivers must include minimising energy. For this to succeed designers should be able to:

Allocate compute resources in units of energy, not just time.
Capture both execution duration and power efficiency.
Force application developers to think about energy and understand and energy model of the software.
Access energy-aware tools for developing energy-efficient software and code.

13.5 Applied mathematics and numerical simulations using high performance computing (HPC)

13.5.1 Present and future performance of applied mathematics and numerical simulations using HPC

In the second half of the 20th century, numerical simulation has been established as the third pillar of scientific discovery, complementing theoretical analysis and experiment. The performance of HPC hardware has since then increased rapidly, steadily driven by the ever-increasing demands for computing power in many areas of science. Today's challenges in numerical simulation are characterized by two main issues, the extreme range of temporal and spatial scales and the complex, non-linear interaction of multiple physical processes.

Future exascale supercomputers will be capable of performing a billion floating point operations per second – a work rate that equates the computing power of about one hundred of today's supercomputers. The availability of such immense computing power will allow for fundamental progress in answering complex scientific research questions and future opportunities, as they arise in environmental sciences, energy, economy, health or security. Exascale computing and high fidelity modelling and simulation are an enabling technology of strategic importance for Europe. However, the path towards exascale computing involves numerous challenges, with a key one being power consumption.

Important publications and reports include [Ang:2012], [Kogge:2008], [Keyes:2013], [EMWG:2014] and the many references therein.

13.5.2 Key challenges of future exascale computing

Power consumption. Reduction of the power dissipation is meanwhile a major driver of changes in HPC infrastructure. This includes the compute nodes themselves, which may be equipped with heterogeneous architectures such as multicore CPUs, many-core co-processors, graphics processing units (GPUs), FPGAs, or others, and local memory, furthermore it includes all network, storages, power supplies, cooling and building infrastructure. Simply scaling current technology to exascale would result in a prohibitive power demand by far exceeding the envisaged goal of a 20 MWatt limit for an exascale machine. Clearly, hardware and software need to be transformed to achieve a tolerable power demand. Complementing the traditional budgeting by means of CPU time on supercomputers, metrics reflecting the energy consumption of simulations need to be developed and implemented for proper assessment and monitoring to operators and feedback to users.

Extreme concurrency. Since processor clock speeds are not increasing significantly any more due to power density limitations, increased performance is achieved by placing more processing cores on a chip and using multi-threading. Supercomputers typically comprise several levels of parallelism ranging from the shared memory multi-threading parallelism within a multicore processor, a many-core device or a GPU, over non-uniform memory access or distributed memory platforms using several processors, many-core devices or GPUs within a single compute node, to interconnected clusters of hundreds or thousands of compute nodes and even geographically distributed computing centres. It is assumed that exascale machines will have orders of magnitude more parallelism than today's petascale systems. New algorithms need to be developed which are able to exploit such extreme concurrency, avoiding typical limiting factors such as synchronization and communication. Hereby the reproducibility of results and code correctness must be treated with special care, and new metrics of accuracy need to be developed, since dynamic scheduling and execution environments may lead to non-deterministic behaviour.

Limited memory bandwidth. As for processors, power is also a limiting factor for memory. The memory density and bandwidth is with current technology not increasing proportionally to the processor performance. While still the memory per compute node increases, the available bandwidth per core will decrease. Thus, algorithms need to cope with memory bandwidth limitations.

Data locality. Analogously, while the cumulated memory bandwidth per compute node increases, the bandwidth per core will decrease. Also the network throughput is not increasing at the same rate as the performance per compute node does. Emerging memory techniques like non-volatile memory, 3D stacking, in-memory computing, deep caches or others, might be used in future HPC machines to improve memory performance and reduce energy consumption. Algorithms therefore need to become aware of the data locality and minimize data transfers, since these consume a significant amount of energy.

Resilience. With increasing system size, the rate of hardware failures might rise to levels where the probability that a failure occurs during a simulation is not any more negligible. Resilience techniques need to be developed and adopted in algorithms, so that results can be recovered and simulations can be continued after a hardware failure. Traditional check-pointing strategies might not be suitable when requiring synchronization and file I/O, or might even take longer than the mean time to failure. In addition, power reduction techniques such as near threshold voltage computing might introduce soft failures. Algorithms are required which are tolerant to certain types of failure through their mathematical nature, and local recovery techniques need to be developed. Again, reproducibility and code correctness need to be addressed carefully due to the non-deterministic nature of failures.

13.5.3 Opportunities for improvement of applied mathematics and numerical simulations using HPC

Problem formulation, modelling and discretization. Mathematical optimization and uncertainty quantification (UQ) will be increasingly used with the advent of more powerful HPC resources. Instead of implementing them as an outer loop around the forward simulation, techniques which include optimization and UQ in the problem formulation from the very beginning may allow for sophisticated analysis and use of different numerical algorithms, and may provide more accurate results, although such formulations may lead to more complex systems of equations. The algorithmic implications and trade-offs between time, energy consumption and quality of results need to be explored.

Numerical solvers. After the mathematical model has been formulated and discretised, the problem is typically stated in the form of a finite-dimensional non-linear or linear system of algebraic equations. These need to be solved numerically by means of non-linear and linear solvers, which consume the biggest portion of time and energy to solution in the vast majority of simulations. Thus, improving the numerical solvers is crucial and several aspects, outlined below, need attention. In high performance computing, time and energy to solution or performance and power are naturally correlated in many cases. Performance optimization yields not always, but often also a reduction of the energy consumption as a consequence. Therefore not only energy-related aspects are discussed, but also performance issues.

- **Energy-aware algorithms.** Already today, and increasingly important in the future, algorithms need to be able to respect power and energy constraints. Power and energy need to be added to the conventional design goals of performance and correctness, and metrics for assessment need to be established. Standard interfaces and APIs for collecting power and energy information have to be developed, supported by accurate measurements through built-in hardware counters and sensors or external measurement devices.

- **Avoiding data transfer.** Data transfer at all levels including local memory to cache, within local memory, read/write to storage, and transfer through the network, is expensive both in terms of time and energy compared to floating point operations. Algorithms need to be investigated which reduce the need for data transfer, and which exploit and improve data locality.

- **Data compression.** The volume of transferred data can be reduced if the data is compressed. However, the compression itself is an additional computation which consumes time and energy. Algorithms need to be investigated for their feasibility to use data compression, and the trade-off with time and energy needs to be taken into account.

- **Multiple precision algorithms.** Not all applications require the full IEEE-754 double precision accuracy, which is however often used by default. Algorithms need to be investigated for their feasibility to use multiple, lower precision data formats. This might speed up the computation, reduce the memory requirements and reduce the data transfer, which all can contribute to a reduced time and energy consumption. However, the numerical properties must carefully be considered since multiple precision algorithms might experience different numerical stability.

- **Reducing synchronization.** In most algorithms, at some point the computation must be synchronized across the machine which usually imposes waiting and idle times. One example of a global synchronization which appears in myriads of algorithms is the computation of the dot product, where all processors need to provide a local contribution and the aggregate result is distributed. Research is needed for restructuring of existing and development of new algorithms which reduce the synchronizations.

- **Randomization and sampling algorithms.** Another approach to reduce synchronization and data transfer is the use of randomized or sampling-based algorithms. Such algorithms work on decoupled, independent sub-parts or instances of the problem and synchronize only locally or occasionally. However, such algorithms may show non-deterministic behaviour, or even sometimes fail to return a result. Research should address both the modification of existing deterministic algorithms towards randomization and sampling, and the developments of new algorithms. The numerical properties and the suitability for specific applications need to be considered.

- **Adaption to load imbalance.** Ill-balanced codes can infer substantial penalties on performance and energy consumption. Load imbalance may occur even for initially well-balanced simulations due to different numerical properties of the problem evolving in different temporal or spatial domains, after recovery from failure, or imposed by energy management.

- **Scheduling and memory management.** In order to efficiently use the massively parallel and heterogeneous platforms, power and energy-aware run-times, scheduling and memory techniques are required.

- **Auto-tuning algorithms.** Complementing the scheduling and memory management on the runtime level, algorithms need to be able to detect and tune themselves to the architecture. Numerical software libraries need to become able to choose the most efficient variant of an algorithm for a particular hardware and a particular application in an automatic way.

14 Communications

14.1 Present and Future Communications Performance

The ability to communicate information between devices, memory, storage and systems is fundamental to ICT systems. Whilst many concentrate on the energy from switching logic and memory devices, communication between devices and especially circuits or systems can be many orders of magnitude greater than logic operations. In this section the communications for the higher levels of the system stack in Figure 10 will be considered (predominantly now photonics and wifi but legacy copper cables still have significant deployment) whilst the electrical device interconnects are considered in section 10.

14.2 Key Communication Challenges

1. As the amount of energy is proportional to the amount of information transmitted and the distance information is transported, the increase exponential increase in information being transported through communication systems has the largest increase in contribution to ICT energy consumption and this is expected to increase considerably over time (Fig. 5).
2. Maintaining the system quality of service and reliability.
3. Understanding the various optimisations across the network stack and across the end to end path
4. Migration paths

14.3 Fundamental Communications Limits

The fundamental limits to communication are the following:

1. Shannon's law: The required energy per bit to detect a signal at a receiver never decreases below a minimum. There is an absolute minimum energy per bit beyond which communication is not possible. Using a temperature of 300 K, this is 2.9 x 10^{-21} J/bit.
2. As the fundamental limit to Shannon's law is approached, the computational complexity required to understand the signal approaches infinity.
3. As the available channel bandwidth is decreased, the minimum energy per bit increases.
4. For a given channel bandwidth, if the required capacity is increased then the minimum energy required per bit must also increase.
5. Photonic system fundamental limits: whilst it has been postulated that fundamentally it should be possible to design a photonic communication system that is thermally noise limited, in practice the dominant noise is the shot noise of the receiver. Since light is quantised in discrete bundles called photons, the counting of photons by a photodetector has a statistical Poisson distribution corresponding to the shot noise limit of the photodetector with the distribution indicating the level of noise in the photodetection process. For photons of 1550 nm wavelength for telecoms systems, the shot noise is of order 2 x 10^{-19} J/bit, significantly higher than the Shannon limit for electrical signals of 2.9 x 10^{-21} J/bit.

14.3.1 Wireless Communication

For wireless communications the energy to transmit a bit of information is given by

$$E_{wireless} = N_{photons} E_{photons}$$

where $N_{photons}$ is the number of photons given for uniformly radiating wireless as $N_{photons} \sim \frac{4\pi r^2}{\lambda^2}$ and E_{photon} is the photon energy given by $E_{photons} = h\nu = \frac{hc}{\lambda}$. The energy to transmit a bit of information [Zhirnov:2014] is therefore given by

$$E_{wireless} \sim \frac{4\pi r^2 hc}{\lambda}$$

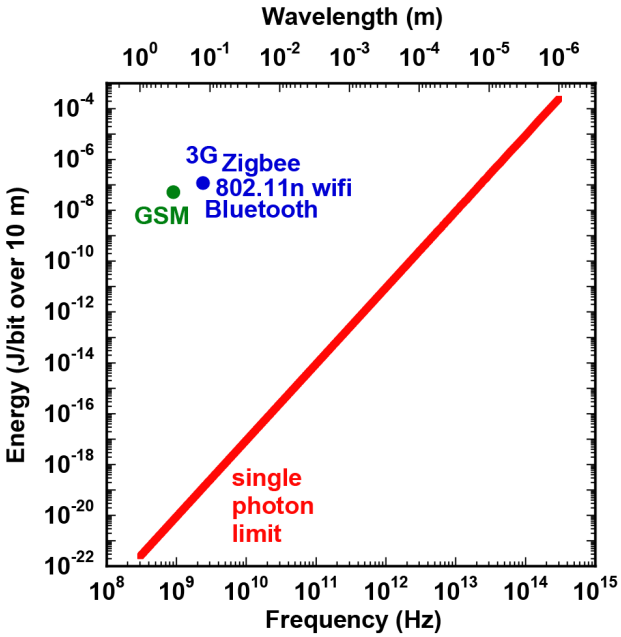

Figure 19: The energy per bit to transmit wireless data 10 m using a number of technologies at a range of transmission frequencies as compared to the fundamental single photon limit for omni-directional wireless transmission.

Figure 19 present the ultimate energy per bit that can be transmitted over 10 m distance by wireless using the single photon limit calculation and compares this to 6 different wireless technologies. All the wireless technologies are around the 60 to 200 nJ range per bit which is many orders of magnitude above the fundamental limits indicating that there is significant potential for improvements in the energy consumption of wireless communications.

14.3.2 Photonics Communications

A number of authors have estimated the energy cost of sending information over the internet [Baliga:2011]. For example in 2009, Baliga et al. undertook a modelling study which included the energy consumption of the core, metro and edge, access and video distribution networks. This included the energy consumption from switching and transmission equipment. They found an energy consumption per bit of 75 μJ at low access rates which decreased to 2 to 4 μJ at an access rate of 100 MB/s. This study estimated that the internet communication components alone accounted for 0.4% of electricity consumption in broadband enabled countries at that point with this percentage predicted to increase significantly as the bandwidth increases. They modelled how the energy per bit will reduce depending on how fast photonic technology improves and with a 10% rate of improved energy efficiency provides 700 nJ/bit by 2023 whilst a 20% rate of improvement results in 120 nJ/bit. A key finding of the study was that the predicted rate of improvement with the future photonic technology in reducing the energy consumption per bit was not sufficient to reduce the energy consumption in the future as demand increases.

Miller [Miller:2009] has investigated the requirements for chip-to-chip and on-chip photonic communications which provides an idea of the ultimate limits for photonic communications. The limits on individual photonic components is likely to be around 10 fJ/bit for the most power hungry suggesting that the limit for short distance photonic communications (< 10 m distance) is around 100 fJ/bit.

For the rack-to-rack optical interconnects for exascale computing for 2019, the US Department of Energy calculated that every pJ/bit of optical power results in a total contribution of 0.8 MW for the complete system power [HOTI:2015]. As of 2014, a typical rack-to-rack optical interconnect at 40 Gbits/s is operating at around 40 pJ/bit [AVAGOTECH:2014].

Table 5 and Table 6 provide data from BT [Krug:2014] demonstrating the typical energy per bit for wired network systems including broadband, access, backhaul, metro and core. Whilst the latest photonic systems can reduce the energy per bit well below the values in tables 5 and 6, as many telecoms network systems are only replaced every 25 years there are a lot of legacy systems that run at far higher energy than the latest technology. Many of these legacy networks are still running at mJ per bit.

	PSTN (μJ/bit)	DSL (μJ/bit)	ADSL (μJ/bit)	FTTC (μJ/bit)
Home	0	183	56	35
Access	1716	176	29	20
Backhaul, metro & core	1200	70	8	9

Table 5: Energy-per-bit for various access types from [Krug:2014].

	Energy (μJ/bit)	Power (W/line)
Home devices	174	20
Legacy networks	2916	4
Broadband access (inc. network termination in home)	116	11
Broadband core (inc. access to data centre)	17	1.5
Data centre	1 – 50	0.1 - 5

Table 6: Typical figures for network energy use from [Krug:2014].

14.4 Opportunities for Communications Improvement

From the photonic device side, experimentally squeezed photon systems have already demonstrated sub-shot noise limited photodetection. The Heisenberg uncertainty principle links the uncertainty in the number of photons detected by the photodetector with the uncertainty in the phase of photons being measured having to be 0.5 or larger. By squeezing the uncertainty in the number of photons the uncertainty in the phase must increase but the statistics become binomial and therefore sub-Poisson. This allows sub-shot noise detection if the photons are placed into Fock or number states [Loudon:2000]. Whilst this has been demonstrated in laboratories the practical demonstration on telecoms systems has not been achieved providing an opportunity for improved energy efficiency.

One of the key issues for photonic networks is that they still require conversion from electrical to optical signals at many points which consumes significant amounts of energy. Also it is still far cheaper to do electronic processing than optical processing again requiring additional energy expensive electrical to optical conversions. The energy consumption of current communications systems is still dominated by the electrical processing. The energy consumption could be reduced if an understanding of the energy consumption associated with the resource management taking into account the control, switching and transmission energy could be achieved. This is an opportunity for a near term energy improvement.

At the moment optical cables between boards and racks is available for datacentres and HPC but there is still significant copper interconnects before you get to the microprocessor, on-chip memory or disk storage. The development of integrated Si photonics where the enormous yields of silicon foundries can be used to produce far cheaper and larger bandwidth (through parallel channels) has the potential to reduce energy per bit far faster than older technologies. Also such technology is also being developed for chip-to-chip photonic communications but also on-chip communications for the higher level and longer interconnects. If the volumes can be reached then the costs should allow large scale deployment.

There is already substantial drivers and programmes aimed at reducing the energy consumption of communication systems at all levels of ICT and the internet: wireless, chip, board, rack, system, metro and core. The biggest issue is whether these reductions are aggressive enough to counteract the increase in the bandwidth of internet and microwave wireless traffic. The IoT will require increased communication bandwidth but as many of the systems are battery powered and require wireless transmission, there are significant incentives to minimise communication bandwidth, distance and time to save battery power. High-definition video is another issue. As use of high-definition video on demand increases, the systems architecture of the internet to reduce long-haul delivery requires investigation with an aim to reduce energy consumption which is proportional to the bandwidth. Studies on cloud computing architectures suggest that more local distributed servers do reduce energy consumption but only if the clients of the users are low power [Baliga:2011]. A key result from analysing cloud computing [Baliga:2011] is that energy

consumption for communications can be significantly reduced if direct, high-bandwidth links could be installed between major servers that required communication in a distributed cloud computation scheme. This is an area of opportunity to understand the optimum distribution network for future ICT. Indeed it may require regulation to direct companies toward energy efficient approaches rather than the cheapest solution for customers.

Baliga et al. [Baliga:2009] have studied the energy consumption for a range of cloud and HPC distributed computing paradigms. For both public (wide area network) and private (local area network) cloud storage of office based work then communications is the largest energy consumption for more than 0.02 downloads per hour for public cloud storage and 0.1 downloads per hour for private cloud storage. The energy consumption overall is, however, lower than each user having a medium performance computer at home undertaking tasks. Their analysis suggests that in an optimum public network, 1 download per hour corresponds to 0.1 W per user whilst 100 downloads per hour corresponds to 18 W/ hour which must be compared to a thin client of 8 W or a medium performance home computer of 80 W. Their analysis concludes that energy savings are achieved by using low-end laptops for routine tasks and cloud processing services for computationally intensive tasks, instead of mid-range or high-end PCs, provided the number of computationally intensive tasks is small. The communication energy consumption in private (local area network) cloud systems is negligibly small compared to the other energy consuming parts of the overall system.

Baliga [Baliga:2009] concludes by stating that energy consumption of cloud computing requires to be considered as an integrated supply chain logistics problem, in which processing, storage and communication are all considered together. They demonstrated that this approach provides more energy efficient use of computer power provided the computational tasks are small on the distributed machines but also indicate that this can fail in many systems.

There is a significant opportunity to determine ideal distributed and cloud computing approaches through undertaking research in this area.

15 Power and Power Management

15.1 Present and Future Power and Management Performance

15.1.1 Power Management

Whilst national grids that transmit electricity to industry and domestic properties for use all transmit using high voltage ac to minimise losses, all ICT systems operate with dc power supplies requiring power transformation and management. At present switch-mode power supplies dominate ac to dc conversion for most ICT devices such as PCs, laptops, smart phones and mobile phones. Such converters are used since they have greater efficiency than other technologies such as linear power supplies because the switching transistor dissipates little power when acting as a switch and spends very little time in any high dissipation energy transition. Silicon power switches including power MOSFETs and insulated gate bipolar transistors (IGBT) depending on the power and regulation requirements.

To improve efficiencies in power conversion, most research is concentrating on developing new materials for power switches which reduce the Ohmic losses and the on-resistance. SiC and GaN are the main materials being developed as the wider bandgap and higher electrical conductivities should improve the conversion efficiencies. As an example GaN is predicted to have 50% higher conversion efficiency than Si for power switches and just converting all electrical drives to GaN is predicted to save 9% of the electricity consumption in the UK which corresponds to removing the equivalent generation capacity of 5 advanced gas cooler nuclear reactors [BIS:2011]. The potential for energy savings across Europe and the world are enormous and have been predicted to be of order of €1,400 Bn per annum [BIS:2011] if GaN technology fully replaces present Si power switches.

15.1.2 Renewable and/or Sustainable Energy Sources for HPC and Cloud Datacentres

Already many HPC cloud servers are being located so that renewable energy can be used to power at least part of the systems. Some of the best examples of low carbon renewables that can be used are photovoltaic, hydro and wind (Table 7). In every case the correct environment is required for each of the renewable

technologies. A good example of this is photovoltaics which is being heavily used for a significant number of datacentres for cloud computing. Figure 20 provides the available energy averaged over 24 hours and 365 days of the year for a range of cities around the planet earth. The actual available energy over 24 hours will be these values times the conversion efficiency for the PV technology being used. The record PV efficiencies are presented in Figure 21 but the majority of deployed PV solar farms use crystalline Si PV cells typically with starting efficiencies of 20 to 22%. Therefore a solar farm in Athens will produce about 41 W/m^2 and to power a 10 MW HPC or cloud datacentre will require at least 250,000 m^2 of PV area. Due to dust and dirt along with the harsh environment that PV operates in, the efficiency drops off with time so a significantly larger area is required for long term sustainable energy generation. Also substantial energy storage is required to capture the energy which is only available during the day so it can be deployed at night.

No renewable technology at present continuously delivers the volume of energy required for HPC or datacentres. Large scale energy storage of renewables is also a significant challenge. Pump storage hydro schemes are the best for large scale storage but requires the correct environment with large reservoirs and a significant height differential. 10 MW levels are certainly achievable even with height differentials of only tens of metres but for low heads of water the volume of water requires to be increased significantly. There is also significant research in batteries, super capacitors and phase change materials with the aim to produce large scale energy storage for renewable technologies. MW scale battery energy storage systems have been deployed [http://www.energystorageexchange.org/], however lower cost, proven long-term and reliable energy storage solutions at the levels requires further research and development.

Technology	**Energy Density**	**CO_2 emission / kWh**	**CO_2 emission / kg fuel**
Coal	29 MJ/kg	960 g	2.685 kg
Diesel	45.4 MJ/kg	778 g	2.772 kg
Natural gas	53.6 MJ/kg	443 g	2.750 kg
Nuclear	20,000 MJ/kg	66 g	Lifecycle CO_2
Solar photovoltaic (c-Si)	1000 W/m^2 peak <200 W/m^2 / 24 h / year	32 g	Lifecycle CO_2
Hydroelectric	0.001 MJ/kg, <0.25 W/m^2	10 g	Lifecycle CO_2
Wind onshore	0.00006 MJ/kg, 1.2 W/m^2	10 g	Lifecycle CO_2

Table 7: The energy density, CO2 emissions per kWh energy generated and CO2 emissions per kg of fuel consumed for different fuels and renewable energy sources. Source [Sovacool:2008] Lifecycle CO2 emissions are from the manufacture, installation, operation and maintenance of the system.

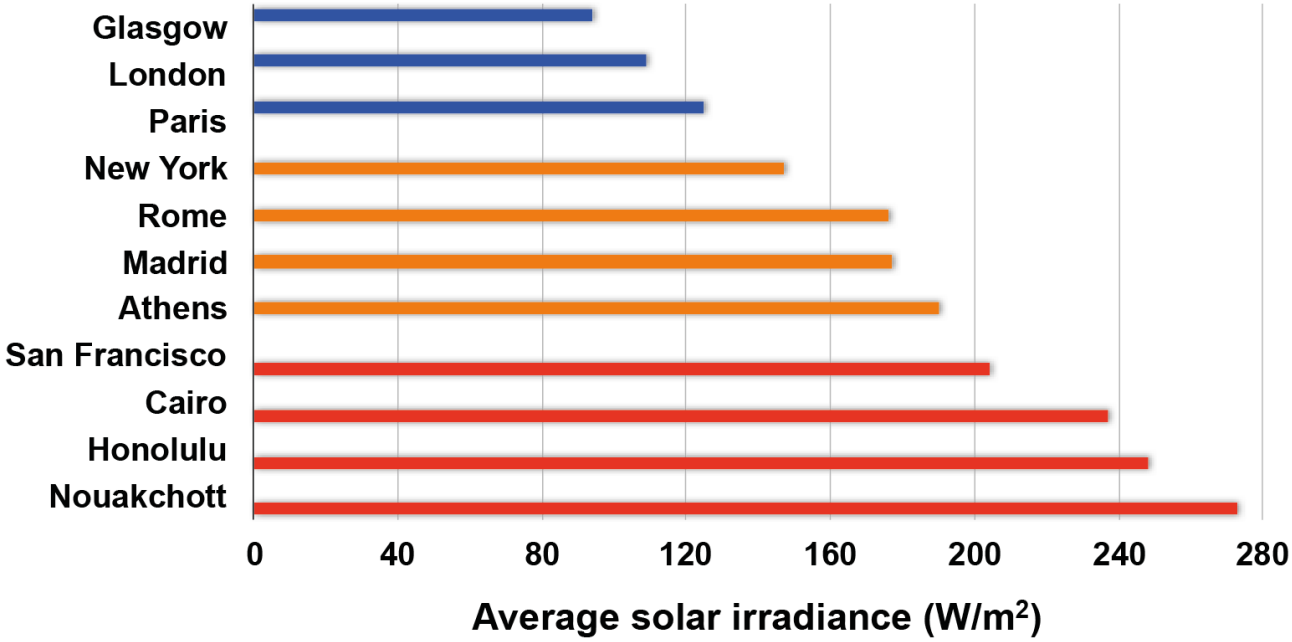

Figure 20: The average solar power available at different points on the earth integrated over 24 hours and 365 days of the year. The peak power is significantly higher than these values but these are the ones available 24/7. Source [Boyle: 2012].

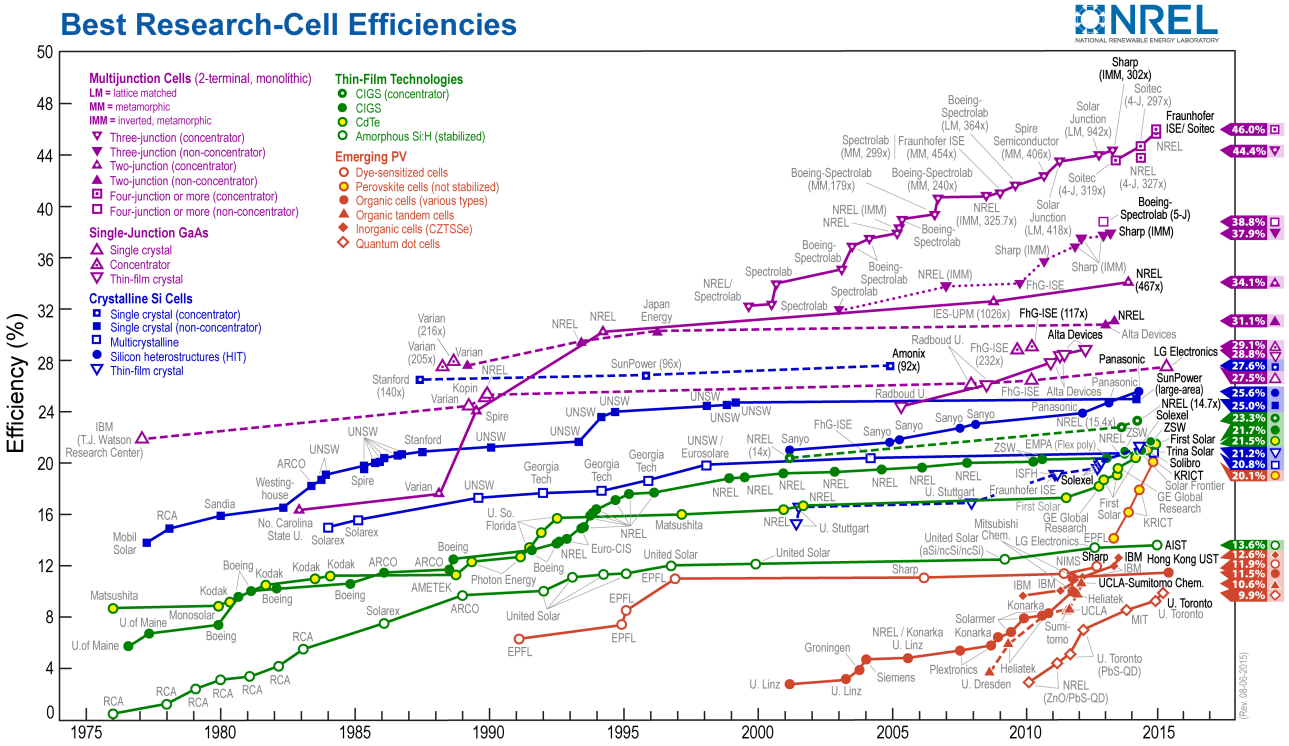

Figure 21: The best reported PV cell efficiencies from different technologies and from single sun to concentrators as a function of the year of report. Source [NREL:2015].

15.1.3 Energy Harvesting for Autonomous Systems

Similar to renewable energy, the energy sources for autonomous sensors must be chosen dependent on the operation environment to maximise the available extraction and conversion of ambient or waste energy into useful electricity. Also all systems require battery or super capacitor storage and appropriate power management to allow the continuous operation since appropriate levels of electrical power may not always be available from energy harvested sources.

Photovoltaics

The available photonic energy outdoors in direct sunlight is provided in Figure 16. For autonomous sensors which can use PV, it is a reliable source during the day and batteries or super capacitors can be used to store up for night time use. Even in northern EU countries the generation levels are around 20 W/m^2 per day providing significant energy for most autonomous sensors.

Indoor PV has significantly lower energy available typically 100 times lower than outdoor direct sunlight. Also most indoor available light is diffuse, scattered off different surfaces and so capturing this light is far more difficult than capturing direct sunlight. The efficiency to capture diffuse sunlight results in PV cells with efficiencies only up to about 10% at present. This is still useful for many autonomous sensors.

Vibrational Energy Harvesters

Vibrational energy harvesters around the size of a drinks can are able to extract up to 5 mWs of power from pumps, trains or vehicles. Such system are now deployed in many industrial and transport applications. Small scale prototypes within few cubic centimetres are able to provide order of µWs of power for sensing and short range communication. For comparison, a lithium ion battery of the same size (330 mL) could store 100 – 300 Wh or at least 20,000 hours of the vibrational energy harvester.

Thermoelectrics

Thermoelectrics convert waste heat into electricity. A temperature gradient is required to generate electricity with the voltage and power output dependent on the size of this temperature gradient. Such

systems are used for smart thermostat controls of heaters and heating systems and also are being develop for automotive vehicles to improve fuel consumption. As 75% of fuel in cars is thrown away as waste heat, harvesting some of this to reduce the alternator dynamic load on the engine to reduce fuel consumption.

Other Energy Harvesting Technologies

A number of other energy harvesting technologies are being researched including rf, miniature heat engines and organic composting systems but none have yet reached products on the market place. The key issues for most other technologies are the problems of generating sufficient power to be useful and practical. Most promising prototypes are those oriented toward hybrid solutions where PV is used together with vibrations and thermoelectrics.

15.1.4 Energy Storage for Autonomous Systems

Microbattery energy storage.

This ICT-Energy Strategic Research Agenda emphasizes the need to bridge the gap between energy requirement for ICT device operation and energy supply from harvesting sources and storage to enable truly autonomous wireless devices. This requires an increase in both the storage capacity of batteries and the efficiency of energy harvesting devices coupled with a decreased demand from the electronics. The ideal scenario would be device integrated energy storage that interfaces with energy harvesting components to provide power on demand to the electronic sensing, communication and display components and operates over the anticipated device lifetime of years.

Ultimately the energy supply components, harvesting and storage should occupy a footprint on chip no larger than the electronics it drives, with 1 mm^2 as an attractive long term target for both. No such energy storage component exists today. Batteries are the most common energy storage option and since their introduction in the 1990's lithium ion batteries have exhibited the highest energy density which has been gradually improved in the intervening period from ~200 Wh/l to ~700 Wh/l through the use of improved materials and processing.

Solid state microbatteries which can be processed/integrated on silicon substrate exhibit similar volumetric energy densities with micron scale thin film materials which necessitates a large area format consistent with the large areas or volumes required by solar or vibrational harvesting, respectively. Solid state devices do offer capacity retention over thousands of cycles [Dudney 2005] which matches the device lifetime requirements. In the 2D, thin-film geometry, current deposition techniques and lithium ion diffusion characteristics in the solid state limits the electrode thickness to several micrometers resulting in a battery dominated by the substrate and other inactive cell components.

As thin films these 2D formats typically exhibit energy densities of ~6 $\mu Wh/mm^2$ or 0.2 J/mm^2. An energy budget of 1 mWh/day can support a wireless sensor node (WSN) used in building energy management with sensing and transmission every 20 minutes. [O'Mathuna 2008] Clearly significant advances are required for energy storage devices to meet the demand in a reasonable footprint. The key challenges are to realise improved energy storage in a significantly decreased footprint for ICT integration and high rate (power) capability during device interrogation and to decrease recharge time. These challenges require;

- Higher energy density materials, particularly at the cathode, where the current material, $LiCoO_2$, is 25 times less energy dense than the lithium metal anode.
- 3D or 1D active materials structuring with increased aspect ratio providing additional material (stored energy) with respect to planar commercial thin film microbatteries.
- Nanoscale active materials with improved electronic conductivity or core/shell structures [Hasan 2010] to facilitate high rate solid state lithium ion transport.

Furthermore, current solid state microbatteries cannot meet the needs of the ICT systems at peak power during measurement and transceiver operation and require a hybrid energy source with a supercapacitor.

This is due to the lithium diffusion limitation in the solid state electrolyte and cathode. On the other hand the solid state construction does facilitate the use of lithium metal anodes which have a large energy capacity (3,600 mAh/g) by comparison with the typical carbon anodes (372 mAh/g) of most lithium ion batteries. If non solid-state electrolytes are to be utilised for higher power outputs then alternative high energy intercalation anodes such as Sn (990 mAh/g) [Whitehead 1999], Ge (1,600 mAh/g) [Laforge 2008] or Si (4,200 mAh/g) [Chan 2008] will be required to prevent dendritic short circuits on cycling. Core/shell [Cui 2009] versions of these anodes may be required to alleviate mechanical stresses leading to poor cycling behaviour and improve the electronic conductivity to access all of the high aspect ratio structures.

Disruptive battery technologies such as a Li/sulphur or Li/air [Abraham 1996; Bruce 2011] can achieve the energy storage requirements of the ICT community. A theoretical energy density of 2.8 mWh/mm^3 (10 J/mm^3) has been estimated for a Li-air battery [Zheng 2008] with non-aqueous electrolytes. It is recognised that the Li/air and other high energy systems have to overcome many obstacles before they will be in widespread deployment but many researchers predict that this could be over the next ten to fifteen years. A very challenging scaling requirement is shown in Figure 22 and that must be coupled with decreased power requirements in the ultimate device. It is clear from the values in Figure 22 that both new materials and 3D structuring of the materials are required to enable the decreased footprint desired for autonomous systems. An increasing focus on improved nanoarchitectures, electrode materials and integrated current collectors is required to surmount these obstacles and deliver high energy density solutions to meet the needs of the electronics industry.

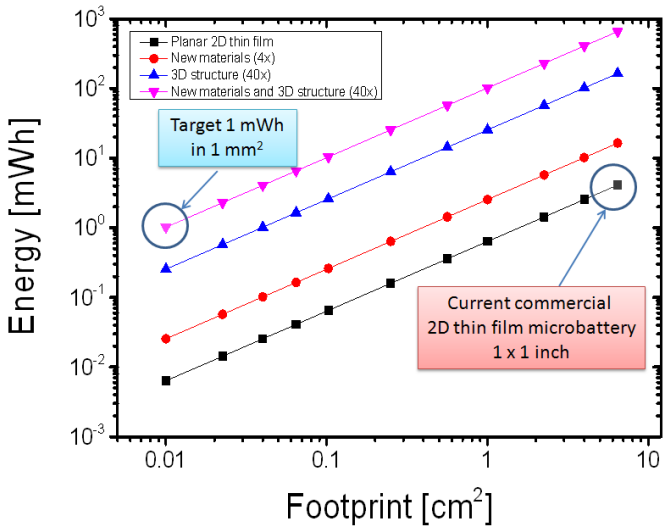

Figure 22: A roadmap for microbattery energy storage requiring the development and integration of new materials which could yield up to 4 times improvement in stored energy and micro or nanoarchitectures to increase the material quantity and surface area to deliver 1 mWh of energy in a 1 mm^2 footprint. Source [Wang 2014].

15.1.5 Supercapacitors

One of the key enabling technologies for IOT applications is increased energy density (Uckelmann 2011).This would ensure further device miniaturisation while also providing greater autonomy of operation for remote devices with improved lifetime capabilities. From a device standpoint, integrated high energy density "smart objects" for IOT applications should be small, rugged, lightweight, and capable of functioning over a wide range of temperatures and under harsh environmental conditions (Miorandi 2012, Kortuem 2010) .

Of the proposed energy storage technologies for "smart objects" supercapacitors are emerging as an attractive candidate to complement advanced lithium batteries, owing to the increasing power density demands being placed on lithium batteries from multifunctional portable devices (Harrop 2014). The key

advantages of supercapacitors for such applications include: high power density, faster charge and discharge rates, and improved cycle-life.

While traditional electrochemical supercapacitors can provide very high specific energy, preventing leakage of the supercapacitors liquid electrolyte solution in a portable or implantable device may be challenging. Also, since exposure of the liquid electrolyte solution to moisture in air can adversely affect its performance, special manufacturing needs are required which to date have prevented integration of supercapacitors into standard manufacturing processes. Thus, there is a market pull for a mechanically durable high performance solid state supercapacitor which can be fabricated using standard semiconductor manufacturing processes (Dunn 2012). While some solid state supercapacitors (SSC's) with highly desirable power and energy density attributes have been demonstrated (Banerjee 2009), significant practical challenges still remain in order to deliver high performance products to meet the future demands of ICT applications.

Desired targets for future solid state supercapacitors to meet these demands would include:

- High energy density > 10 Wh/kg*.
- High power density > 1 X 10^6 W/kg
- Wide operating temperature window (-20 °C to +70 °C)
- Low equivalent series resistance (ESR) < 10mΩ.
- Long cycle life > 1 X 10^6 cyles.

*The energy density of current commercial off-the-shelf supercapacitors is relatively low at less than ~ 10 Wh/kg.(Burke 2014) with current solid state supercapacitor equivalents at ~ 1 Wh/kg.(Banerjee 2009).

15.2 Key Power and Management Challenges

1. Efficiencies are already high in most power supplies and limited by fundamental Ohmic losses. Therefore improvements require lower resistive transistors such as GaN power HEMTs or MOSFETs. As copper is the main inductor metal whilst silver might provide less losses, the increased costs are prohibitive.
2. Cost and reliability are the main drivers in the power management area rather than energy efficiency.
3. Energy harvesting devices including vibrational, thermoelectric and indoor photovoltaic are predominantly at 10s and 100s of μW continuous output power for 1 cm^2 devices rather than mW or higher required for many autonomous sensors.
4. The challenge is to achieve sufficient harvested and stored energy to match the power density requirements of small autonomous sensors.
5. Renewable energy systems such as solar, hydro or wind to run HPC clusters and cloud datacentres require the appropriate local environment to produce sufficient power for the systems. This suggest locating HPCs and cloud datacentres should be located where there is abundant renewable energy.

15.3 Fundamental Power and Management Limits

The transistors in the switched mode power supply are switched fully on or fully off to minimise the amount of time the power supply spends in any high power dissipation state. Therefore the majority of the losses are Ohmic resistance, quiescent current and skin effect losses in either the switching transistor or the copper inductors and wires.

15.4 Opportunities for Power and Management Improvement

1. The power industry is already driving GaN power switches towards markets to reduce losses in power supplies.
2. The [E4U:2009] highlighted how the use of advanced power electronics techniques, like new dc distribution networks could reduce ICT energy consumption by 10%. Further integration of ICT technologies with power electronics could save an additional 20% and the implementation of best practise with smart systems could lead to a 50% reduction in energy consumption of ICT systems.

The technologies required to gain these improvements are only starting to appear on the market and the scale of implementation is still small.

3. As mobile telecoms is one of the biggest areas of growth in ICT, radio base stations for mobile communications could have significant improvements. Conventional base stations in 2009 had an efficiency of 1.2% using 10 kW to transmit 120 W of power. At least 30% saving could be made in such systems [E4U:2009].
4. The discussion above highlights the importance of system level methodologies in optimising the overall integrated solution rather than optimising sub-elements in isolation.
5. Collaboration is required in specification definition and testing, combining expertise in circuit design, fabrication and energy systems for next generation highly miniaturised efficiency, low leakage hybrid or multi-source energy harvesting and storage with novel power management solutions. [Wang:2014]

16 Added Value from Optimising at Multiple Levels of the System Stack

The key message that has become clear from investigating each of the levels of the system stack is that overall energy efficiency can be optimised more aggressively when the design at one level understands the energy issues at another level. This is especially true for levels that control the behaviour of another level e.g. circuit architecture on devices or software on circuit architecture. At the start of ICT systems in the 1960s and 1970s it was possible for engineers to understand all levels of the system stack and design accordingly. As each level has become more complex and the number of transistors and lines of code have moved from thousands into many billions, it has become more difficult for people to understand all levels of the system. A clear message is that only through the joint understanding of how different levels of the stack must be designed to reduce energy consumption will optimum ICT solutions that minimise the use of energy can be found. Therefore suitable education where an expert in e.g. software is provided with sufficient knowledge to understand the energy impact of code at the other levels of the stack is essential if significant energy reductions are to be achieved.

A key challenge is how to design software that enables optimal use of the underlying hardware energy saving capabilities. There is an opportunity through the development of energy transparent systems or models to develop novel compilers that minimise energy consumption well beyond present systems. Also in the software section it was pointed out that education is key so that coders write efficient software that optimised compilers can then produce the most efficient runtime operations otherwise significant amounts of energy are wasted.

As communication is the biggest energy consumption and sustainable energy is environmental dependent, what is the optimum distribution and location of servers for the cloud if it is driven by sustainable energy ICT? Significant work has already been achieved in this space for internet and cloud systems. Similarly as the energy per bit of wireless communication scales with distance, what is the optimum distribution of mobile base stations to minimise energy consumption for wireless communications? Microcells have already been suggested and modelling achieved but could heuristic views of the complete system to deliver services provide new solutions? It should be pointed out that in some circumstances minimising communication is not the lowest energy solution – for example the processing required for heavy compression can cost more energy than you save under most transmission scenarios.

In networks there are many opportunities to reduce energy consumption but also challenges with them. For example how to cope with network systems that are becoming more dynamic – how to design routing protocols that keep the routes available when the routers are asleep. How to get the advances in adaptation into communications systems – reliability, resilience and quality of service being the main issues. Finally migration to new, lower energy technologies is an enormous challenge facing the communications industry.

How would the architecture and software of a computational system look if minimising communication of information was used to optimise a minimum energy consumption system?

17 International Comparisons

A list of the EU and International research groups and programmes is in the ICT Energy Field Survey to look at how the EU position is with respect to the rest of the world.

18 Strengths, Weaknesses, Opportunities and Threats to Sustainable Energy ICT in Europe

18.1 Strengths

1. Europe is a leader in low power circuit design with companies that include ARM Holdings, CSR, Imagination, Wolfson and many others.
2. Europe was the first through the regulatory framework for mobile phones to develop a standard that became the world standard (GSM). Careful choice of new standards that benefit European companies should be a strength.
3. Europe has a much stronger "More than Moore" and embedded systems sectors compared to "More Moore" with strengths in imaging (e.g. medical and mobile phone cameras), navigation (e.g. MEMS gyroscopes) and low power systems.
4. A number of EC networks have produced strong agendas for driving their communities forward. These include ICT Energy, HiPEAC, ARTEMIS ITEA and the European Technology Platform for HPC.
5. Europe has industrial strengths in developing new processes for silicon foundries (e.g. IMEC, Leti) and with industrial tool manufacturers.
6. Europe has significant expertise in energy storage materials research and development which for new applications such as micro or nanosystems and hybrid energy device fabrication can also be implemented in the 'More than Moore' scenario.

18.2 Weaknesses

1. The leading edge silicon foundries with the lowest power transistor sizes are not in Europe and are dominated by non-European industry. The few that are in Europe are related to foreign multinationals circumventing EC importation rules through the use of tax breaks in certain EC countries. Building new foundries benefits from government subsidies and many countries outside of Europe especially in Asia are happy to subsidies the semiconductor industry for long term economic benefit.
2. The main large scale industrial programme to develop beyond CMOS devices is run by the Semiconductor Research Association in the U.S.. Whilst there have been EC equivalents (e.g. SiNANO, MINECC), the scale has been significantly lower and the industrial led programmes have tended to concentrate on more conservative approaches.
3. Whilst Europe is strong in developing new processes and with tool manufacturers, this strength does not translate into large, leading technology node manufacturing foundries in Europe especially for microprocessors.

18.3 Opportunities

1. CMOS is starting to reach the limits of performance improvements from transistor scaling. Any change to a new device technology with lower switching energy is an opportunity for Europe to take a lead in devices, foundries, architectures, design, software and systems with significantly lower energy consumption than any present technology in any industrial roadmap. The move to a different technology to CMOS would require a step change in performance to enable any investment plan to succeed in taking a new technology to market.
2. There are new markets developing for autonomous sensors using embedded technology that can provide significant benefit to society. Such areas include personalised healthcare, smart buildings, smart transport and many others where the use of the sensors can make significant energy and cost savings in each of these markets. Sustainable energy systems in these markets are extremely attractive as there is a strong cost driver to reduce all replacement of power sources by development

of integrated energy harvesting and hybrid energy storage with rechargeable battery materials and supercapacitors that could be used to develop and drive sustainable energy ICT in other sectors.

3. Development of tools and techniques to provide energy transparency throughout the ICT stack would provide a business advantage to companies involved in all aspects of ICT systems. In particular the portable and low power ICT systems could be fully optimised for minimum energy consumption. Such tools are also essential for HPC to being able to achieve the next generation of supercomputers.

18.4 Threats

1. Consumer behaviour for streaming data for smartphones and cloud computing is providing a continuous exponential increase in the volume of data being communicated which is continuously increasing the energy consumed by ICT faster than the present incremental reductions in energy consumption communicated per bit. If the consumers cannot understand why sustainable energy ICT is important then any additional costs or performance changes to achieve sustainable ICT will be difficult to take to market. This suggests that education is key to success.
2. The lack of the leading edge silicon foundries in Europe is a significant risk to future ICT hardware manufacture.
3. Since there are still links between GDP and CO_2 emissions for many economies, the poor European economy may significantly aid the Europe 2020 Agenda of 20% reduction in CO_2 emissions from 1990 levels by 2020 thereby reducing political agendas for reducing CO_2 emissions by other means.

19 Conclusions

It is clear that more sustainable ICT devices and systems is required if a number of the EU Social Agendas are to be met. Communications appears the most power hungry from the initial analysis but further analysis is required to fully understand how to minimise communications energy consumption. A key issue is being to find tools or models to allow energy transparency throughout the system stack. Such tools would allow significant optimisation to reduce energy consumption. The scaling of devices no longer reduces power consumption so can new devices be used which do allow the energy per switch to be reduced.

There is still a lot of work to fully understand energy in ICT devices and systems and especially how to reduce the energy consumption and carbon emissions.

20 Appendices

20.1 References

[Abraham 1996] Abraham, K. M.; Jiang, Z. Journal of the Electrochemical Society **1996**, *143*, 1

[Ang:2012] J. Ang, K. Evans, A. Geist, M. Heroux, P. Hovland, O. Marques, L. Curfman McInnes, E. Ng, and S. M. Wild. Report on the workshop on extreme-scale solvers: Transition to future architectures. Report, U.S. Department of Energy, ASCR (2012).

[ARM:2013] ARM Cortex-A15 MPCore Processor Technical Reference Manual, ARM, pp. 53 - 63 (2013).

[ASAP] IEEE International Conference on Application-specific Systems, Architectures and Processors

[Atzori:2010] L. Atzori, A. Iera, G. Morabito, "The Internet of Things: A survey, Computer Networks", 54(15) pp. 2787-280 (2010).

[AVAGOTECH:2014] http://www.avagotech.com/docs/AV02-3542EN - accessed August 2014

[Bahar:2001] R. Bahar and S. Manne, "Power and energy reduction via pipeline balancing," Proceedings 28th Annual International Symposium on Computer Architecture, 2001, pp. 218–229 (2001).

[Baliga:2009] J. Balig et al., "Energy Consumption in Optical IP Networks" J. Lightwave Technol. 27(13), pp. 2391 – 2403 (2009).

[Baliga:2011] J. Baliga et al., "Green cloud computing: Balancing energy in processing, storage and transport" Proc. IEEE 99(1), pp. 149 – 167 (2011).

[Banerjee 2009] Banerjee, P., Perez, I., Henn-Lecordier, L., Lee, S.B. and Rublof, G.W., 2009. Nanotubular metal-insulator-metal capacitor arrays for energy storage. Nature Nanotechnology, 4(5), pp. 292-296.

[Bassett:2012] P. Bassett and M. Saint-Laurent, "Energy efficient design techniques for a digital signal processor," 2012 IEEE International Conference on IC Design Technology (ICICDT), pp. 1–4 (2012).

[Benatti:2014] S.Benatti, B.Milosevic, F.Casamassima, P.Schonle, P.Bunjaku, S.Fateh, Q.Huang, L.Benini : "EMG-based Hand Gesture Recognition With Flexible Analog Front End" in proceedings of international IEEE conference on BIOmedical Circuit And Systems (BIOCAS 2014), Lausanne, Oct 2014.

[BIGLITTLE:2013] big.LITTLE Technology: The Future of Mobile Making very high performance available in a mobile envelope without sacrificing energy efficiency. White Paper. ARM, 2013.

[BIS:2011] UK Department for Business, Innovation and Skills, "Power Electronics: A Strategy for Success", October (2011).

[Bohr:2007] M. Bohr, "A 30 Year Retrospective on Dennard's MOSFET Scaling Paper" IEEE Solid-State Circuits Society Newletter 12(1), pp. 11 – 13 (2007).

[Borkar:2011] S. Borkar and A.A. Chien, "The future of microprocessors" Communications of the ACM, 54(5), pp. 67 - 77 (2011).

[Boyle:2012] G. Boyle, "Renewable Energy: Power for a Sustainable Future" Oxford University Press (2012).

[Brill:2012] B. K. G. Brill, "The Invisible Crisis in the Data Center: The Economic Meltdown of Moore' s Law," Uptime Institute, Rev. 2, 2007-12.

[Bruce 2011] Bruce, P. G.; Hardwick, L. J.; Abraham, K. M. MRS Bulletin **2011**, *36*, 506

[Buchty:2012] R. Buchty, V. Heuveline, W. Karl and J.-P. Weiss, "A Survey on Hardware-aware and Heterogeneous Computing on Multicore Processors and Accelerators", Concurrency and Computation: Practice & Experience, 24(7) John Wiley and Sons Ltd. (2012).

[Burke 2014] Burke, A. et al. (2014) Review of the Present and Future Applications of Supercapacitors in Electric and Hybrid Vehicles. IEEE International Electric Vehicle Conference, also Institute of Transportation Studies, University of California at Davis, Research Report UCD-ITS-RR-14-23.

[Byrne:2006] R. Byrne and D. Diamond, "Universal tools for biomolecular attachment to surfaces" Nature Mater. 5, 421 (2006).

[Chan 2008] Chan, C. K.; Peng, H. L.; Liu, G.; McIlwrath, K.; Zhang, X. F.; Huggins, R. A.; Cui, Y. Nature Nanotechnology **2008**, *3*, 31.

[Chang:2013] S.W. Chang et al. "InAs N-MOSFETs with record performance of I_{on}=600 μA/μm at I_{off}=100 nA/μm (Vd=0.5 V)" Proc. IEDM 2013, 417 (2013).

[Changkyu:2007] K. Changkyu, S. Sethumadhavan, M. S. Govindan, N. Ranganathan, D. Gulati, D. Burger, and S. Keckler, "Composable Lightweight Processors," 40th Annual IEEE/ACM International Symposium on Microarchitecture, MICRO 2007 pp. 381–394 (2007).

[Choi:2007] W.Y. Choi et al. "Tunneling Field-Effect Transistors (TFETs) With Subthreshold Swing (SS) Less Than 60 mV/dec" IEEE Elec. Dev. Lett. 28, 743 (2007).

[Cisco:2015] Cisco White Paper "The Zettabyte Era: Trends and Analysis" (2015) – available http://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/VNI_Hyperconnectivity_WP.html .

[Colinge:2010] J.P. Colinge et al. "Nanowire transistors without junctions" Nature Nano. 5, 225 (2010).

[Constantin: 2012] J.H.-F. Constantin, A. Y. Dogan, O. Andersson, P. A. Meinerzhagen and J. N. Rodrigues et al. "TamaRISC-CS: An Ultra-Low-Power Application-Specific Processor for Compressed Sensing" IFIP/IEEE 20th International Conference on Very Large Scale Integration (VLSI-SoC), Santa Cruz, USA, (2012).

[Cui 2009] Cui, L. F.; Ruffo, R.; Chan, C. K.; Peng, H. L.; Cui, Y. Nano Letters **2009**, *9*, 491.

[Deleonibus:2014] S. Deleonibus et al. "Future Challenges and Opportunities for Heterogeneous Process Technology. Towards the Thin Films, Zero Intrinsic Variability Devices, Zero Power Era." Proc. IEDM 14, 223 (2014).

[Dennard:1974] R.H. Dennard et. al, "Design of Ion Implanted MOSFET's with very small physical dimensions" IEEE J. Solid-State Circuits, 9(5), pp. 256 - 268 (1974).

[Dewey:2011] G. Dewey et al. "Fabrication, characterization, and physics of III-V heterojunction tunneling Field Effect Transistors (H-TFET) for steep sub-threshold swing" Proc. IEDM 11, 785 (2011).

[Dudney 2005] Dudney, N. J. Materials Science and Engineering B-Solid State Materials for Advanced Technology **2005**, 116, 245

[Dunn 2012]Dunn, Bruce S. 2012. SOLID-STATE SUPERCAPACITOR. U.S. Patent 20140035098A1, filed May 14, 2012.

[E4U:2009] Electronics Enabling Efficient Energy Use – https://setis.ec.europa.eu/energy-research/sites/default/files/static-projects/files/D3.1.4%20Final%20E4U%20publication%20V12-Final-Screen.pdf

[EEHCO:2015] HIPEAC Workshop on Energy Efficiency with Heterogeneous Computing (2015)

[EDGAR:2014] Emission Database for Global Atmospheric Research – http://edgar.jrc.ec.europa.eu

[ELC:2014] A European Industrial Strategic Roadmap for Micro- and Nano-electronic Components and Systems: Implementation Plan, June 2014

[EMWG:2014] Report of the Exascale Mathematics Working Group on „Applied Mathematics Research for Exascale Computing", U.S. Department of Energy (2014).

[ENIAC:2006] See e.g. ENIAC "Strategic Research Agenda" (SRA2006): http://www.eniac.eu/web/divers/SRA_2006.pdf and the "i2010 – A European Information Society for growth and employment" initiative, launched by the Commission on 1 June 2005 (http://ec.europa.eu/information_society/eeurope/i2010/what_is_i2010/index_en.htm) - one of the main i2010 three priorities being to strengthen investment in "innovation and research in ICT".

[ENIAC:2015] http://cordis.europa.eu/ist/eniac

[Fehske:2011] A. Fehske, G. Fettweis, J. Malmodin and G. Biczok, "The global footprint of mobile communications: The ecological and economic perspective" IEEE Communications Magazine 49(8), pp. 55-62 (2011)

[FDSOI:2012] N. Planes et. Al, "28nm FDSOI technology platform for high-speed low-voltage digital applications", Proc. Symp. VLSI, 2012.

[FET:2010] "Disruptive Solutions for Energy Efficient ICT", Draft Terms of Reference, FET Proactive, (2010).

[Ford:2011] A.C. Ford et al. "Ultrathin body InAs tunneling field-effect transistors on Si substrates" Appl. Phys. Lett. 98, 113105 (2011).

[Fujimatsu:2012] M. Fujimatsu et al., "71 mV/dec of sub-threshold slope in vertical tunnel field-effect transistors with GaAsSb/InGaAs heterostructure" IEEE Int. Conf. InP Relat. Material 2012, 25 (2012).

[Gammaitoni:2012] Luca Gammaitoni (2012): There's plenty of energy at the bottom (micro and nano scale nonlinear noise harvesting), Contemporary Physics, 53:2, 119-135. http://dx.doi.org/10.1080/00107514.2011.647793

[Gammaitoni:2015] L. Gammaitoni, D. Chiuchiú, M. Madami, G. Carlotti, Towards zero-power ICT, Nanotechnology, IOP Publishing, Volume 26, Number 22, p.222001 (2015)

[Gartner:2015] http://www.gartner.com/technology/research.jsp – Accessed July 2015

[Greenhalgh:2011] P. Greenhalgh, "big.LITTLE processing with ARM Cortex-A15 and Cortex A7" ARM White Paper, 2011.

[GreenICT:2015] http://greenict.ieee.org

[Harrop 2014] Mr Franco Gonzalez and Dr Peter Harrop. "Batteries & Supercapacitors in Consumer Electronics 2013-2023: Forecasts, Opportunities, Innovation." (February 2014). IDTechEx.

[Hasan 2010] Hasan, M.; Chowdhury, T.; Rohan, J. F. Journal of the Electrochemical Society **2010**, 157, A682

[Heddeghem:2014] W. Van Heddeghem, S. Lambert, B. Lannoo, D. Colle, M. Pickavet and P. Demeester, "Trends in worldwide ICT electricity consumption from 2007 to 2012" Computer Commsunications 50, pp 64 – 76 (2014).

[HOTI:2015] http://www.hoti.org/hoti20/slides/Fuad_Doany_IBM.pdf

[HPC500:2015] http://www.top500.org - accessed July 2015

[Hunumaiah:2012] V. Hanumaiah and S. Vrudhula, "Energy-efficient Operation of Multicore Processors by DVFS, Task Migration and Active Cooling," Computers, IEEE Transactions on, vol. 63, no. 2, pp. 349–360, (2012).

[IAEA:2007] International Atomic Energy Agency, "Nuclear Power Plant Design Characteristics" (2007) Available - http://www-pub.iaea.org/mtcd/publications/pdf/te_1544_web.pdf

[IEA:2014] International Energy Authority, "Key World Energy Statistics" (2014) – available http://www.iea.org/statistics/

[Iida:2011] R. Iida et al. Planar-type $In_{0.53}Ga_{0.47}As$ channel band-to-band tunneling metal-oxide-semiconductor field-effect transistors" J. Appl. Phys. 110, 124505 (2011).

[ITRS:2013] International Technology Roadmap for Semiconductors: available – http://www.itrs.net/ .

[Jan:2012] C.-H. Jan et al. "A 22nm SoC Platform Technology Featuring 3-D Tri-Gate and High-k/Metal Gate, Optimized for Ultra Low Power, High Performance and High Density SoC Applications" Proc. IEDM 12, 44 (2012).

[Jooya:2009] A. Z. Jooya and M. Analoui, "Program phase detection in heterogeneous multi-core processors," 14th International CSI Computer Conference, CSICC 2009 pp. 219–224 (2009).

[Kau:2012] H. Kaul et. al, "Near-threshold voltage (NTV) design: opportunities and challenges", ACM/EDAC/IEEE Design Automation Conference pp. 1149 – 1154 (2012).

[Keyes:2013] D. E. Keyes, L. C. McInnes, C. S. Woodward, W. Gropp, E. Myra, M. Pernice, J. Bell, J. Brown, A. Clo, J. Connors, E. Constantinescu, D. Estep, K. Evans, C. Farhat, A. Hakim, G. Hammond, G. Hansen, J. Hill, T. Isaac, X. Jiao, K. Jordan, D. Kaushik, E. Kaxiras, A. Koniges, K. Lee, A. Lott, Q. Lu, J. Magerlein, R. Maxwell, M. McCourt, M. Mehl, R. Pawlowski, A. P. Randles, D. Reynolds, B. Riviere, U. Ruede, T. Scheibe, J. Shadid, B. Sheehan, M. Shephard, A. Siegel, B. Smith, X. Tang, C. Wilson, and B. Wohlmuth. Multiphysics simulations: Challenges and opportunities. Int. J. High Perf. Comput. App., 27:4–83 (2013).

[Kim:2009] S.H. Kim et al. "Germanium-source tunnel field effect transistors with record high ION/IOFF" VLSI Technol. symp. 2009, 178 (2009).

[Kogge:2008] P. Kogge, K. Bergman, S. Borkar, D. Campbell, W. Carlson, W. Dally, M. Denneau, P. Franzon, W. Harrod, K. Hill, J. Hiller, S. Karp, S. Keckler, D. Klein, R. Lucas, M. Richards, A. Scarpelli, S. Scott, A. Snavely, T. Sterling, R. S. Williams, and K. Yelick. Exascale computing study: Technology challenges in achieving exascale systems (2008).

[Kortuem 2010] Kortuem, G., Kawsar, F., Sundramoorthy, V. & Fitton, D. 2010, "Smart objects as building blocks for the internet of things", IEEE Internet Computing, vol. 14, no. 1, pp. 44-51.

[Krishnamohan:2008] T. Krishnamohan et al. "Double-Gate Strained-Ge Heterostructure Tunneling FET (TFET) With record high drive currents and <60mV/dec subthreshold slope" Proc. IEDM 2008, 947 (2008).

[Krug:2014] L. Krug, M. Shackleton and F. Saffre, "Understanding the environmental costs of fixed line networking" e-Energy 14, June 11 – 13 Cambridge, UK (2014) – DOI: 10.1145/2602044.2602057 .

[LaForge 2008] Laforge, B.; Levan-Jodin, L.; Salot, R.; Billard, A. Journal of the Electrochemical Society **2008**, *155*, A181

[Lambert:2012] S. Lambert, W. Van Heddeghem, W. Vereecken, B. Lannoo, D. Colle and M. Pickavet, "Worldwide electricity consumption of communication networks" Opics Express 20(26), pp. B513-B524 (2012).

[Landauer:1961] R. Landauer, IBM J. Res. Dev. 5, 183-191 (1961).

[Leech:2014] C. Leech and T. Kazmierski, "Energy efficient multi-core processing" Electronics 18(1), pp. 3-10 (2014).

[Leonelli:2010] D. Leonelli et al. "Performance Enhancement in Multi Gate Tunneling Field Effect Transistors by Scaling the Fin-Width" Jap. J. Appl. Phys. 49, 04DC10 (2010).

[Li:2012] R. Li et al. "AlGaSb/InAs Tunnel Field-Effect Transistor With On-Current of 78 μA/μm at 0.5 V" IEEE Elec. Dev. Lett. 33, 363 (2012).

[LISA:2006] P. Ienne and R. Leupers. Customizable Embedded Processors: Design Technologies and Applications (Systems on Silicon). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2006.

[Liu:2014] Q. Liu et al. "FDSOI CMOS Devices Featuring Dual Strained Channel and Thin BOX Extendable to the 10nm Node" Proc IEDM 14, 219 (2014).

[Lopez:2014] R. Braojos Lopez, I. Beretta, G. Ansaloni and D. Atienza Alonso. "Hardware/Software Approach for Code Synchronization in Low-Power Multi-Core Sensor Nodes" IEEE/ACM 2014 Design Automation and Test in Europe (DATE) Conference, Dresden, Germany, (2014).

[López-Suárez:2015] M. López-Suárez, I. Neri and L. Gammaitoni, "Sub kBT micro electromechanical irreversible logic gate" arXiv.org Conde-Mat. arXiv:1508.02617

[Loudon:2000] R. Loudon, "The quantum theory of light" Oxford University Press (2000).

[Mamaghanian:2012] H. Mamaghanian, N. Khaled, D. Atienza Alonso and P. Vandergheynst, "Design and Exploration of Low-Power Analog to Information Conversion Based on Compressed Sensing" IEEE Journal of Emerging and Selected Topics in Circuits and Systems, 2(3), pp. 493-501, (2012).

[Mamaghanian:2014] H. Mamaghanian, D. Atienza Alonso and P. Vandergheynst (Dirs.). Compressed sensing: a universal energy-efficient compression scheme for biosignals on wireless body sensor nodes. EPFL, Lausanne, 2014.

[Malmodin:2010] J. Malmodin, A. Moberg, D. Lundén, G. Finnveden and N. Lövehagen, "Greenhouse Gas Emissions and Operational Electricity Use in the ICT and Entertainment & Media Sectors" Journal of Industrial Ecology 14, pp. 770 – 790 (2010).

[Mayer:2008] F. Mayer et al. "Impact of SOI, Si1-xGexOI and GeOI substrates on CMOS compatible Tunnel FET performance" Proc. IEDM 2008, 163 (2008).

[MEDIA:2015] MEDEA+ is the industry-initiated pan-European Programme for advanced cooperative Research and Development in Microelectronics: http://www.medeaplus.org

[Miller:2009] D.A.B. Miller, "Device requirements for optical interconnects to silicon chip" Proc. IEEE 97(7), pp. 1166 – 1185 (2009).

[Miorandi 2012] Miorandi, D., Sicari, S., De Pellegrini, F. & Chlamtac, I. 2012, "Internet of things: Vision, applications and research challenges", Ad Hoc Networks, vol. 10, no. 7, pp. 1497-1516.

[Mohata:2011] D.K. Mohata et al "Self-aligned gate nanopillar $In_{0.53}Ga_{0.47}As$ vertical tunnel transistor" IEEE Dev. Res. Conf. 2011, 203 (2011).

[Mohata:2011A] D.K. Mohata et al. "Demonstration of improved heteroepitaxy, scaled gate stack and reduced interface states enabling heterojunction tunnel FETs with high drive current and high on-off ratio" VLSI Technol. Symp. 2012, 53 (2012).

[Mohata:2012] D.K. Mohata et al. "Demonstration of improved heteroepitaxy, scaled gate stack and reduced interface states enabling heterojunction tunnel FETs with high drive current and high on-off ratio" VLSI Technol. Symp. 2012, 53 (2012).

[Mookerjea:2009] S. Mookerjea et al. "Experimental demonstration of 100nm channel length $In_{0.53}Ga_{0.47}As$-based vertical inter-band tunnel field effect transistors (TFETs) for ultra low-power logic and SRAM applications" Proc. IEDM 09, 949 (2009).

[MOSF:2012] N. Xu et al., "Carrier mobility enhancement via strain engineering in future thin-body MOSFETs", IEEE Electron Device Letters, Vol. 33, pp. 318-320, 2012.

[Natarajan:2014] S. Natarajan et al. "A 14 nm logic technology featuring 2nd-generation FinFET, air-gapped interconnects, self-aligned double patterning and a 0.0588 μm^2 SRAM cell size" Proc. IEDM 14, 71 (2014).

[Nikonov:2012] D.E. Nikonov and I.A. Young, "Uniform methodology for benchmarking beyond-CMOS logic devices" Proceedings IEDM 12, pp. 672-576 (2012).

[Noguchi:2015] M. Noguchi et al. "High I_{on}/I_{off} and low subthreshold slope planar-type InGaAs tunnel field effect transistors with Zn-diffused source junctions" J. Appl. Phys. 118, 045712 (2015).

[NREL:2015] http://www.nrel.gov/ncpv/images/efficiency_chart.jpg

[O'Mathuna] C. Ó.Mathúna, T. O'Donnell, R. Martinez, J.F. Rohan and B. O'Flynn, Energy Scavenging for Long-Term Deployable Mote Networks Talanta, 75, 613-623 (2008).

[Pallister:2015] J. Pallister, K. Eder and S.J. Hollis, "Optimizing the flash-RAM Energy Trade-off in Deeply Embedded Systems", Proc.13th Annual IEEE/ACM International Symposium on Code Generation and Optimization, IEEE Computer Society (2015).

[Paradiso:2005] J. A. Paradiso and T. Starner, "Energy Scavenging for Mobile and Wireless Electronics", IEEE Pervasive Computing 4, 18 (2005).

[Pickavet:2008] M. Pickavet, W. Vereecken, S. Demeyer, P. Audenaert, B. Vermeulen, C. Develder, D. Colle, B. Dhoedt and P. Demeester, "Worldwide energy needs for ICT: The rise of power-aware networking" 2nd International Conference on Advanced Networks and Telecommunication Systems, 2008, pp. 1 – 3 (2008).

[Pollack:1999] F. Pollack "New Microachitecture Challenges in the Coming Generations of CMOS Process Technology, 32nd International Symposium on Microarchitecture (Micro32) (1999) – available http://research.ac.upc.edu/HPCseminar/SEM9900/Pollack1.pdf

[Pop:2010] E. Pop, "Energy Dissipation and Transport in Nanoscale Devices", Nano Research 3, pp. 147–169 (2010).

[Putnam:2014] A. Putnam, A.M. Caulfield, E.S. Chung, D. Chiou, K. Constantinides, J. Demme, H. Esmaeilzadeh, J. Fowers, G.P. Gopal, J. Gray, M. Haselman, S. Hauck, S. Heil, A. Hormati, J.-Y. Kim, S. Lanka, J. Larus, E. Peterson, S. Pope, A. Smith, J. Thong, P.Y. Xiao, and D. Burger, "A Reconfigurable Fabric for Accelerating Large-scale Datacenter Services", Proc. of the 41st Annual International Symposium on Computer Architecuture, IEEE Press, (2014).

[Rajkumar:2010] R. Rajkumar, I. Lee, L. Sha, and J. Stankovic. "Cyber-physical systems: the next computing revolution" Proceedings of the 47th Design Automation Conference (DAC '10) (2010).

[Rakossy:2012] Z. Rakossy, T. Naphade, and A. Chattopadhyay, "Design and analysis of layered coarse-grained reconfigurable architecture," 2012 International Conference on Reconfigurable Computing and FPGAs (ReConFig) 2012 pp. 1–6 (2012).

[Rajamohanan:2015] B. Rajamohanan et al. "0.5 V Supply Voltage Operation of $In_{0.65}Ga_{0.35}As/GaAs_{0.4}Sb_{0.6}$ Tunnel FET" IEEE. Elec. Dev. Lett. 36, 20 (2015).

[Rasmussen:2005] N. Rasmussen, "Guidelines for Specification of Data Center Power Density", (2005) – available http://www.apcdistributors.com/white-papers/Cooling/WP-120%20Guidelines%20for%20Specification%20of%20Data%20Center%20Power%20Density.pdf

[Rohan 2014] J.F. Rohan, M. Hasan, S. Patil, D. Casey and T. Clancy, Energy storage materials and architectures at the nanoscale, in *ICT-Energy - Nanoscale energy management concepts towards Zero-Power Information and Communication Technology* / G. Fagas, L. Gammaitoni, D. Paul and G. Abadal Berini (Eds.), Intech Publications. pp 107-138, (2014)[Roundy:2003] S. Roundy, P.K. Wright, and J. M. Rabaey, "Energy Scavenging for Wireless Sensor Networks" Kluwer Academic Publishers, Boston (2003).

[Roy: 1997] K. Roy and M.C. Johnson," Software design for low power. In *Low power design in deep submicron electronics"*, Wolfgang Nebel and Jean Mermet (Eds.). Kluwer Nato Advanced Science Institutes Series, Vol. 337. Kluwer Academic Publishers, Norwell, MA, USA, pp 433-460 (1997).

[Sabri:2011] M. M. Sabry, et al. "Energy-Efficient Multiobjective Thermal Control for Liquid-Cooled 3-D Stacked Architectures", TCAD, 2011.

[SMARTer:2013] Global e-Sustainability Initiative "SMARTer2020: The Role of ICT in Driving a Sustainable Future" (2013) – available http://gesi.org/portfolio/report/72

[Sovacool:2008] "Valuing the greenhouse gas emissions from nuclear power: A critical survey" Energy Policy, 36, pp. 2940 - 2953 (2008).

[Sridhar:2014] A. Sridhar, M. M. Sabry, P. Ruch, D. Atienza, B. Michel, "PowerCool: Simulation of Integrated Microfluidic Power Generation in Bright Silicon MPSoCs", oral, ICCAD, November 2014, San Jose, CA, USA.

[Sutter:2005] HS05: H. Sutter, "The Free Lunch Is Over: A Fundamental Turn Toward Concurrency in Software", Dr Dobb's Journal, 30(3) http://www.drdobbs.com/ (2005).

[Tiwari:1996] V. Tiwari, S. Malik and A. Wolfe, "Instruction Level Power Analysis and Optimization of Software", Journal of VLSI Signal Processing Systems, 13, pp 223-238, (1996).

[Tomioka:2012] K. Tomioka et al., "Steep-slope tunnel field-effect transistors using III-V nanowire/Si heterojunction" VLSI Technol. Symp. 2012, 47 (2012).

[Uckelmann 2011] Uckelmann, Dieter, Mark Harrison, and Florian Michahelles. "Architecting the Internet ofThings". Berlin: Springer, 2011, pages 212-213.

[UN:1987] United Nation's Brundtland Commission "Our Common Future" (1987) – available http://www.un-documents.net/our-common-future.pdf

[UN:2012] United Nations Department of Economic and Social Affairs "World Population Prospects: The 2012 Revision" (2012) – available at http://esa.un.org/wpp/unpp/panel_population.htm

[Wang] W. Wang, J.F. Rohan, N. Wang, M. Hayes, A. Romani, E. Macrelli, M. Dini, M. Filippi, M. Tartagni and D. Flandre, **Chapter 9 - Smart Energy Management and Conversion** in Beyond CMOS Nanodevices 1 (2014) Editor F. Balestra, Wiley, P 249-276, ISBN: 978-1-84821-654-9.

[Wang:2015] P.Y. Wang and B-Y. Tsui "Experimental Demonstration of p-Channel Germanium Epitaxial Tunnel Layer (ETL) Tunnel FET With High Tunneling Current and High ON/OFF Ratio" IEEE Elec. Dev. Lett. 36, 1264 (2015).

[Whitehead 1999] Whitehead, A. H.; Elliott, J. M.; Owen, J. R. Journal of Power Sources **1999**, 81, 33.

[WikipediaTV:2015] https://en.wikipedia.org/wiki/List_of_common_resolutions#cite_note-21 - accessed November 2015.

[Wright:2015] S. Wright, "Digital compositing for film and video, volume 10" Taylor and Francis (2012).

[Wu:2013] S.-Y. Wu et al. "A 16 nm FinFET CMOS Technology for Mobile SoC and Computing Applications" Proc. IEDM 13, 224 (2013).

[YINS:2015] Developing a radically new thermal-aware design approach for next generation energy-efficient datacenters. http://www.nano-tera.ch/projects/450.php

[Yu:2013] T. Yu et al. "In0.53Ga0.47As/GaAs0.5Sb0.5 Quantum-Well Tunnel-FETs With Tunable Backward Diode Characteristics" IEEE Elec. Dev. Lett. 34, 1503 (2013).

[Zeng:2008] H. Zeng, J. Wang, G. Zhang, and W. Hu, "An interconnect-aware power efficient cache coherence protocol for CMPs," IEEE International Symposium on Parallel and Distributed Processing, IPDPS 2008 pp. 1–11 (2008).

[Zha:2009] Z. Bo et. al, "Energy-efficient subthreshold processor design" IEEE Trans. VLSI 17(8), pp. 1127 - 1137 (2009).

[Zhao:2010] H. Zhao et al. "$In_{0.7}Ga_{0.3}As$ Tunneling Field-Effect Transistors With an I_{on} of 50 μA/μm and a Subthreshold Swing of 86 mV/dec Using HfO_2 Gate Oxide" IEEE Elec. Dev. Lett. 31, 1392 (2010)

[Zhao:2011] H. Zhao et al. "Improving the on-current of $In_{0.7}Ga_{0.3}As$ tunneling field-effect-transistors by p++/n+ tunneling junction" Appl. Phys. Lett. 98, 093501 (2011).

[Zhou:2012] G. Zhou et al. "InGaAs/InP Tunnel FETs With a Subthreshold Swing of 93 mV/dec and ION/IOFF Ratio Near 10^6" IEEE Elec. Dev. Lett. 33, 782 (2012).

[Zhou:2012A] G. Zhou et al. "Novel gate-recessed vertical InAs/GaSb TFETs with record high ION of 180 μA/μm at VDS ~ 0.5 V" Proc. IEDM 2012, 777 (2012).

[Zhao:2014] X. Zhao et al., "InGaAs/InAs Heterojunction Vertical Nanowire Tunnel FETs Fabricated by a Top-down Approach" Proc. IEDM 2014, 1503 (2014)

[Zheng 2008] Zheng, J. P.; Liang, R. Y.; Hendrickson, M.; Plichta, E. J. Journal of the Electrochemical Society **2008**, *155*, A432

[Zhirnov:2003] V.V. Zhirnov, R.K. Cavin III, J.A. Hutchby and G.I. Bourianoff, "Limits to binary logic switch scaling - a gedanken model" Proceedings of the IEEE 91, pp. 1934 – 1939 (2003).

[Zhirnov:2014] V.V. Zhirnov, "Fundamentals of Energy Consumption in ICT Devices" NIPS / ICT Energy Summer School (2014) – available http://www.nipslab.org/sites/nipslab.org/files/NiPS2014_Zhirnov_en%20cons%20ICT.pdf

[Zu:2011] C.-X. Zu and H. Li, "Thermodynamic analysis on energy densities of batteries" Energy Environment Science 4, pp. 2614 – 2624 (2011).

20.2 Definition of Acronyms

CPS – Cyber Physical Systems
CMOS – Complementary Metal-Oxide-Semiconductor
FET – Field Effect Transistor
FPGA – Floating Point Gate Array
HPC – High Performance Computing
ICT – Information Communications Technology
IoT – Internet of Things
ISA – Instruction Set Architecture
ITRS – International Technology Roadmap of Semiconductors
MEMS – Micro Electro Mechanical Systems
MIPS – Million Instructions Per Second
MOSFET – Metal Oxide Semiconductor Field Effect Transistor
MPU – Micro-Processing Unit
NEMS – Nano Electro Mechanical Systems

20.3 List of Contributors

(Alphabetical by surname)

Giovanni Ansaloni, EPFL, Switzerland
David Atienza, EPFL, Switzerland
Zbigniew Chamski, MpicoSys, Poland
Adrian Cristal, Barcelona Supercomputer Centre, Spain
Georgios Fagas, Tyndall National Institute, Ireland
Kerstin Eder, University of Bristol, U.K.
John Gallacher, Roskilde University, Denmark
Luca Gammiatoni, University of Perugia, Italy
Pablo Garcia, EPFL, Switzerland
Vincent Heuveline, Heidelberg University, Germany
Steve Kerrison, University of Bristol, U.K.
Louise Krug, BT, U.K.
Andrew Lord, BT, U.K.
Jeremy Morse, University of Bristol, U.K.
Douglas Paul, University of Glasgow, U.K.
James Rohan, Tyndall National Institute, Ireland
Martin Wlotzka, Heidelberg University, Germany
M. Fernando Gonzalez Zalba, Hitachi Cambridge Laboratory, U.K.
Olivier Zendra, Inria, France
Victor Zhirnov, Semiconductor Research Corporation, U.S.A.

Contributions were also taken from the ICT Energy workshop "Future Energy" in the ICT Research Agenda on the 15th September 2015 in Bristol where this SRA was presented and a number of working groups provided feedback that has been included. Not all the names of the participants have been included in the above list. The participants included a wide range of academics through the system stack and companies including ARM and Intel.