BEYOND THE HORIZON

June 2006

Beyond the Horizon

Anticipating Future and Emerging Information Society Technologies

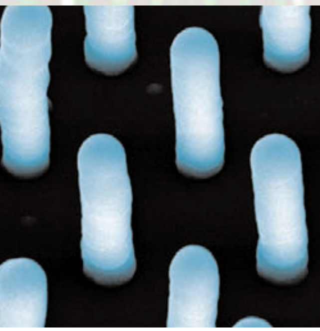


"Foreseeing beyond the horizon — for technology and business opportunities — is any Director's strategic responsibility. Engaging in future emerging technology research will create services supported by skilled staff for deployment within (1) my own organisation performing leading edge R&D and (2) commercially in industry."
Keith G. Jeffery, Director IT, CCLRC, UK

"The ambitious goals for European R&D set in Lisbon enforces identification of key areas with imagination and vision. Roadmapping this research is essential for determining our direction. The 'Beyond-the-Horizon' project is a most valuable contribution to this process."
Jan Karel Lenstra, Director, CWI, The Netherlands

"Rapid technological change demands organizational navigation. Nobody can analyze future technological trends all by himself. Usually we have to rely on the common view. 'Beyond-the-Horizon' provides a rock-solid basis for tracing out one's own course. I enjoyed reading the reports which helped my own analysis very much."
Arne Sølvberg, Norwegian University of Science and Technology, Norway

European Research Consortium for Informatics and Mathematics
ERCIM
www.ercim.org

Information Society Technologies

SIXTH FRAMEWORK PROGRAMME

Introducing 'Beyond-The-Horizon'

Keith G. Jeffery,
Director IT, CCLRC, UK
and ERCIM President.

Information and Communication Technologies (ICT) have made an enormous impact on all aspects of our lives: from commerce and industry through to advances in science, technology and medicine, on to novel techniques and products in arts and humanities and into our homes and everyday life. The personal computer is just over 20 years old and now it is available in a pocket-sized device, integrated with mobile phone, digital camera (including video-clip capability) and media player. How will it (or something else) be realised in 15 years time? How will it impact our lives?

The rapid development of ICT has left in its wake two major problems. The first concerns complexity: how to design and build ICT systems of quality that are reliable, easy to use, adaptive, capable of development and efficient (in both performance and cost). The second is how to foresee changes in the socioeconomic landscape and in the technological advances in all related fields (such as nanotechnology) so that systems can be designed to meet the needs, and advantage of these technological changes can be taken. There is a clear need in Europe for us to grapple with — and come to understand — the future of ICT because of its far-reaching impact. As part of ERCIM's response to this need, the Beyond-The-Horizon Project was launched to identify long-term research needs through a wide ranging foresight activity.

Of course this is not the only activity attacking these problems: ISTAG (the Information Society Technologies Advisory Group) in the last few years has had two ERCIM professionals as members. It produced (among other documents) Grand Challenges as a visionary view of the future of ICT. Expert Groups in various Director-ates of the EC have considered the future of ICT in a specific area and have had many ERCIM participants; one example is the Next Generation GRIDs Expert Group. European Technology Platforms, most of them involving ERCIM members, are also developing their R&D roadmaps for specific domains. ERCIM has the reach to pull together all these threads of activity.

The focus of Beyond-the Horizon is complementary. It aims to identify ideas for research of high-risk but with high potential for radical innovation and long-term payoff. It goes beyond ICT R&D roadmaps that often identify issues where 'no known solutions' are available. It is crucial to avoid tunnel vision in research, to give space to new ideas, to explore radical alternatives for the future and provide options for later developments that will be taken up by industry in the future. At the same time, while research in ICT is largely 'purpose-driven', it is necessary to maintain some distance between research and commercial applications if we want to keep creativity: in ICT, the slow gestation period for research is in sharp contrast with the rapid economic sanctions of the resulting technologies.

ICT has always drawn on wisdom from surrounding disciplines, both for requirements of ICT systems and inspiration for solutions to problems of ICT systems. From mathematics to biology, from materials science to psychology, ICT has obtained important insights and developments. Thus it is no surprise that the future of ICT in the medium-term is seen to be intermixed with (for example) advances in materials science producing billions of miniaturised components, advances in biological systems (from animal colonies to cells to DNA and proteins) to understand their functioning and utilise that in ICT systems, and advances in cognitive science to understand brain function and use the techniques in ICT but also to optimise human-ICT interactions.

Undertaking a process (information-gathering, analysing, predicting) to produce roadmaps for the future of ICT is important for the economic and social well-being of Europe. The results should be utilised not only for the EC Framework Programme but also for influencing national governments' strategies and those of commercial, industrial, medical, social and cultural organisations. These actors must create the conditions for sustained collaboration among the best teams so that the critical mass can be obtained to arrive at substantial results and proofs of concept that drive progress in emerging fields, such as those explored here. Interdisciplinary research is essential because the combination and confrontation of different view points stimulates the exploration of new directions that may lead to significant breakthroughs.

The Beyond-The-Horizon Project

Beyond-The-Horizon (B-T-H) is a European coordination action to identify ICT-related research trends and strategic areas that require support. The action is funded by IST-FET, the Future and Emerging Technologies activity of the EU Information Society Technologies programme. The coordination is handled jointly by ERCIM — the European Research Consortium for Informatics and Mathematics, and FORTH — Institute of Computer Science, Greece.

European Research Consortium for Informatics and Mathematics
ERCIM
www.ercim.org

About ERCIM

ERCIM — the European Research Consortium for Informatics and Mathematics — fosters collaborative work within the European research community: academic, commercial and industrial. The members of ERCIM include leading research establishments from seventeen European countries. Encompassing over 12,000 researchers and engineers, ERCIM is able to undertake consultancy, development and educational projects on any subject related to its field of activity. ERCIM was founded in 1989 and is a European Economic Interest Grouping (EEIG) based in Sophia Antipolis, France.

ERCIM has one national node per country that focuses academic and commercial/industrial expertise on a local scale. These nodes are linked together into a unique European organisation of excellence. ERCIM has gained a major role in Information and Communication Technologies (ICT) in Europe by building this Europe-wide, open network of centres of excellence.

ERCIM is a major representative organisation in its field, and a gateway to relevant ICT research groups in Europe for worldwide international organisations. ERCIM considers it a high priority to develop cooperation with scientists all over the world. ERCIM hosts the European branch of the World Wide Web Consortium, participates in EU activities, and has established cooperations with both the European Science Foundation and the US National Science Foundation.

http://www.ercim.org

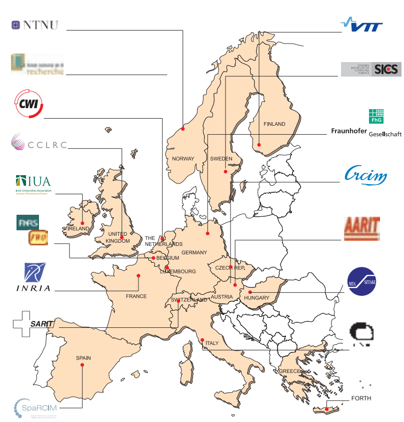

A well-defined, extensive and systematic consultation of the science and technology communities in Europe involved in the problems and challenges mentioned above, has led to roadmaps for six strategic research areas, describing the routes research should take in order to meet the emerging grand challenges.

These areas are: Pervasive computing and communications; Nano-electronics and nanotechnologies; Security, dependability and trust; Bio- ICT synergies; Intelligent and cognitive systems; and Software-intensive systems. In addition, FET itself brought in three strategic areas to be included in the coordination action: Quantum Information Processing and Communication, Complex Systems, and Tera-Device Computing. Similar roadmaps exist for these areas.

The consultations consisted of:

- brainstorming workshops during 2005, in which eminent researchers from both academia and industry drafted agendas for basic research in the six identified areas
- a plenary workshop in Paris, December 2005, where representatives of the six thematic groups together with prominent invitees like Nobel laureate Claude Cohen-Tannoudji and IST-FET Head of Unit Thierry Van der Pyl, discussed the outcome of the individual workshops, and explored research challenges arising at the intersections of the different areas
- on-line consultation of research communities in Europe through the B-T-H website, where the draft documents describing the six areas were available for comment, after which they were finalized during April 2006.

This process has mobilised a wide and multi-disciplinary community providing input at all levels — individual, research groups, institutions as well as funding agencies — to the benefit of the future European Information Society.

About FET

The Future and Emerging Technologies (FET) action is part of the IST Programme — itself part of the EU 6th and forthcoming 7th Framework Programmes. Its main objective is to promote the blossoming of new research ideas and help them mature, thus founding the technologies for the future information society. This implies support of high-risk collaborative research, specifically in areas lacking an agreed roadmap, where options are unclear or ideas speculative. FET support aims to overcome fragmentation and achieve critical mass, and is implemented in two ways: the open domain (FET-Open) and the proactive scheme.

Research in the open domain is motivated by the fast-moving and unpredictable nature of change and innovation in ICT. Here new and promising ideas are tested and assessed. As FET-Open has no predefined work programme, it is an early indicator of changing research priorities. It acts as a test bed for research to be included later in mainstream activities of the IST Programme, where industrial involvement becomes essential.

The proactive scheme aims to build up critical mass in particularly promising areas with potential for industrial or societal impact. It provides funding to groups of projects that jointly explore their vision, creating critical mass and the appropriate mix of disciplines to build consensus around research agendas in emerging areas that lack a common roadmap and to arrive at substantial results and proof-of-concepts that can drive progress in emerging fields.

Research in FET involves collaboration between different disciplines, including physics, mathematics, chemistry, biology, and the life sciences as a whole. Within this inter-disciplinary landscape, FET also supports new emerging research communities to organise themselves, and the coordination of national research programmes.

http://www.cordis.lu/ist/fet/

Information Society Technologies

Contents

Long-Term Challenges in ICT

Technology is steadily developing in relatively well-known directions such as miniaturisation, greater efficiency, connectivity, embedding, personalisation, richer content and interactivity. While each of these still present great challenges, the road for short- and medium-term research, development and application is largely traced, and visible until the horizon.

Longer-term research in Information and Communications Technology (ICT) is aimed at exploring the world Beyond The Horizon. It is of central importance for future technologies, especially when it aims to overcome current roadblocks, or leads to radical alternatives and new views on technology and what it can do for society. Our best theories and methods are already overwhelmed by the sheer size and complexity of some of today's ICT applications, like the Internet. Moreover, current approaches, for example miniaturisation, will reach their ultimate limitations, and more input from science is needed to overcome these limitations. In addition, fundamental breakthroughs and eye-openers are expected to arise from the cross-fertilisation of ICT with areas such as biology, physics, cognitive science, and social science. These areas not only serve as rich sources of inspiration; the notions of information and computation are at the heart of understanding the interfaces between these different scientific approaches to reality and ICT is increasingly providing the 'glue' between them. Faced with such high expectations, ICT is under pressure to reinvent itself.

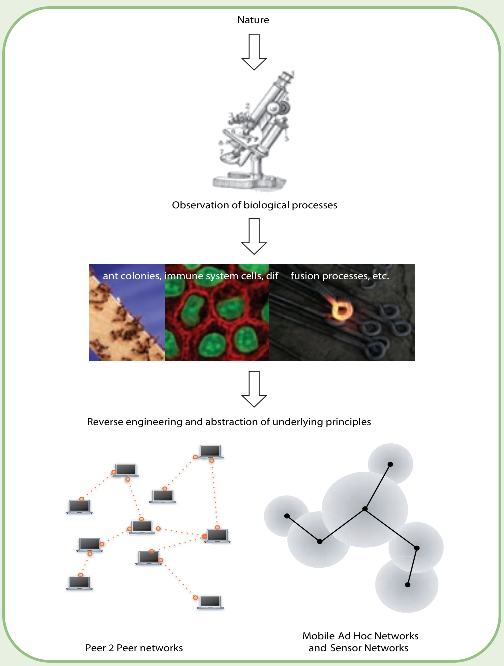

Source: Istituto Dalle Molle di Studi sull'Intelligenza Artificiale (IDSIA), Switzerland

Rethinking the nature of Information and Computation: From biology to dynamic networks.

In this context, a radical interdisciplinary exploration of new and alternative approaches to future and emerging ICT-related technologies is urgently needed. Science and technology evolve, and our society with them, in several ways that require a fundamental reconsideration of our theories, methods, technologies and applications, and provide a frame of reference for long-term research trends in ICT. These ways include the following:

- **Rethinking the Nature of Information and Computation**

  Viewing computers as input/output boxes processing information is increasingly anachronistic. Their increased embedding in everyday objects, their direct coupling with the world through sensors, and the distributed and many-user nature of contemporary applications no longer naturally fit the highly idealised, still prevailing model of computation as formulated by Alan Turing back in the 1930s. Moreover, there are many *natural systems molecules, living cells, the brain, the weather, the economy — that can be viewed as doing some kind of computation as well, but that defy classical models. For instance, now that the limits of today's computational intelligence are becoming clearer, the emerging cognitive science draws heavily on other disciplines (neurosciences and psychology, in particular) to reshape its research agenda. If ICT is to regain its role in this endeavour, it must adapt its own models and dare to adopt new ones.*

Interdisciplinary Research

Once upon a time, chemists didn't talk to physicists, physicists to biologists, biologists to psychologists, and mathematicians to anyone else. These times are over. The frontiers of scientific disciplines are disappearing. For further progress in science reaching out between disciplines it is essential to understand the language of the others, learn from each other's accomplishments, and cooperate with the strangers across the border. ICT act as a kind of 'glue', pervading almost all scientific disciplines, and as a driver for new research in those disciplines. For example, large-scale computational schemes are invented and applied in areas varying from weather prediction to modelling the living cell. In turn, ICT itself is inspired by the scientific fields it is applied to: new computing paradigms are emerging, based on insights from biology, the cognitive and the social sciences, and (quantum) physics. As these developments require close collaboration between birds of different feathers, support of interdisciplinary research carried out in Europe-wide frameworks is a cornerstone of the FET action.

The Beyond-the-Horizon project, funded by IST-FET and coordinated jointly by ERCIM and ICS-FORTH, Greece, precisely addresses this interdisciplinary aspect as inherent in research for the future society.

http://www.beyond-the-horizon.net/

The Physical Realisation of ICT

Miniaturisation has been at the heart of microelectronics, and will remain so for a long time. However, beyond making things smaller, there is a more fundamental trend of reducing computation to a minimal physical process. For instance, in quantum computing fundamental properties of matter, rather than some abstraction (the 'bit') on top of it, are used to encode information. In molecular computing the natural dynamics of molecules is used as the computing process. But also at the macro scale, intelligent materials like specifically tailored polymers utilise the synergies between physical dynamics and computation, e.g., for robotic grasping or locomotion. The capacity to exploit the characteristics of different 'computational materials' would lead to new types of architectures and components, but also to completely new ways of manufacturing them, e.g., by growth or self-assembly. These tendencies are repositioning ICT as an integral part of nature rather than as something alien to it.

Embracing Change

Everything flows, said Heraclitus 2500 years ago, but he didn't know current ICT systems: they are extremely rigid, requiring the users to adapt to them, rather than the other way around. They are not evolving with changing needs, neither do they adapt to different contexts of use. Moreover, building them requires time consuming, expensive and error-prone engineering processes. And, once put into practice, running software systems are not easily replaced. Changing or adapting them can be a nightmare, leading more often to failure than to success. In contrast, biological systems have extraordinary capability to change over time. They grow, adapt, self-assemble, replicate, heal, self-organise, evolve and recycle themselves. If we can not master these features for ICT, then our technological future will not be sustainable.

Physical-Virtual Confluence — example of an augmented reality application: Analysis of a goal keeper's movements. The goalkeeper faces a virtual 3D adversary. with 'autonomous' capabilities.

In other words, a capacity for change, no matter at what time scales, must be embraced as a fundamental characteristic of any ICT system.

- **Physical-Virtual Confluence**
  With advanced media and interface technologies (e.g., for multi-player games in augmented reality), further enhanced by insights on human perception and action, it is clear that the borderline between the real and the virtual is rapidly blurring away. This will have tremendous consequences for the way in which we communicate, do science and business, the way we learn and work, and live our daily lives. In parallel, ICTs will increasingly blend with the biological world. Advances in molecular biology and the neurosciences, but also in psychology and the social sciences are completely transforming our ability to both understand and program biological, cognitive and behavioural functions. This has deep implications for the way in which we relate to and interact with machines. It is no longer helpful to think of a 'system' as ending at the technological interfaces, but it will include the user(s) and the broader context and process within which it operates. Instead of the prevailing technology-centred approach, a new systems approach is needed.

- **ICT Transforming Society**
  The real power of technology comes from its embedding in, and interaction with human systems and organisations. Effects that are often unanticipated by the designers emerge from the intertwining of complex networks of increasingly sophisticated devices and services with large-scale societal activities. Some of these effects can be positive (like the market dynamics generated in online marketplaces) and others negative (cascading or escalating failures in communication networks). Yet others open up new possibilities that businesses and citizens readily pick up and exploit, sometimes with malicious intents (computer viruses). In other words, applying ICT is no longer a neutral act. The ways in which ICT changes, moulds and becomes part of the things to which it is applied, needs to be much better understood if we want to integrate future technologies like personal robotics, massive commercial services or immersive collaborative environments, in a secure, trustworthy and sustainable way.

Given these trends, FET research addresses a wide range of core ICT communities, that cover at least the areas listed below, and are complemented where needed by other disciplines:

- Pervasive Computing and Communications
- Nano-electronics and Nanotechnologies
- Security, Dependability and Trust
- Bio-ICT Synergies
- Intelligent and Cognitive Systems
- Software Intensive Systems
- Quantum Information Processing and Communication
- Complex Systems Research
- Tera Device Computing.

Pervasive Computing and Communications

A Pervasive Computing scenario.

Computing devices are already pervading into everyday objects, in such a way that users no longer notice them as separate entities. Appliances, tools, clothing, accessories, furniture, rooms, machinery, cars, buildings, roads, cities, even whole agricultural landscapes increasingly embody miniaturised and wireless -thus invisible- information and communication systems, establishing information technology rich socio economic systems with the potential to radically change the style of how we perceive, create, think, interact, behave and socialise as human beings, but also how we learn, work, cultivate, live, cure, age as individuals or in societal settings.

Prospective advances in microprocessor-, communication- and sensor/actuator-technologies envision a whole new era of computing systems, seamlessly and invisibly woven into the 'fabric of everyday life', and hence referred to as Pervasive Computing. Their services will be tailored to the person and the context of their use. After the era of keyboard and screen interaction, a computer will be understood as secondary an artefact, embedded and operating in the background, with its complete physical environment acting as an interface (primary artefact). Pervasive Computing aims at interaction with digital information by manipulating physical real-world artefacts as 'graspable interfaces', by simultaneously involving all human senses, and by considering interaction related to the semantics of the situation in which it occurs.

The future pervasive computing senses and controls the physical world via many sensors and actuators, respectively. Applications and services will therefore have to be greatly based on the notions of context and knowledge, will have to cope with highly dynamic environments and changing resources, and will need to evolve towards a more implicit and proactive interaction with humans. Communication must go beyond sending information from one fixed point to another, considering situations where devices cooperate and adapt spontaneously and autonomously in the absence of any centralized control.

A vast manifold of small, embedded and mobile artefacts characterize the scenarios envisaged by Pervasive Computing and Communications. The challenges are related to their ubiquity, their self-organisation and interoperation, their ability to perceive and interprete their situation and consequently adapt the services they offer, and to the different modes of user interaction with those services. Pervasive communications call for new architectures based on device autonomy, fragmented connectivity, spatial awareness and data harnessing inside each network node.

Examples of everyday objects with computing devices that users do not notice anymore as separate entities. The characteristics of pervasive computing devices are that they are aware of the user's presence: sensitive, adaptive and responsive to the users' capabilities, needs, habits and emotions, and ubiquitously, safely and securely accessible via natural interaction.

The key technical problems and milestones towards achieving this vision are:

Networked Societies of Artefacts

Going beyond their individual capability to localize and recognize other artefacts as well as humans and their intentions, groups of artefacts organized as 'societies' (characterized by huge size and enormous complexity) will have to share competencies, to collectively act in a sensitive, proactive, and responsive way according to the perceived and anticipated needs, habits, and emotions of the users. Coordinated goal-oriented artefact communities are supposed to be the 'interface', via which humans will ultimately be served. The research can benefit from models derived from the social sciences, for example the reaction to and the prediction of user needs, thereby bridging the gap between ICT and social science.

Evolveable Systems

Designing Pervasive Computing and Communications environments might not be feasible due to their complexity. Instead, such systems could grow from their origin driven by predefined or emerging goals. This ability to evolve is a key feature: in order to cope with the continuously changing contexts, conditions, and purpose of their use, systems must become self-configuring, self-healing, self-optimizing and self-protecting, from a hardware and software point of view. Research must go beyond the current state of the art in context-awareness and become future-aware in the sense that the system will have a certain anticipation of future contexts of its use.

Human Computer Confluence

Human computer confluence refers to an invisible, implicit, embodied or even implanted interaction between humans and system components. New classes of user interfaces may evolve that make use of several sensors and are able to adapt their physical properties to the current situational context of users. In the near future visible displays will be available in all sizes and will compete for the limited attention of users. Examples include body worn displays, smart apparel, interactive rooms, large display walls, roads and architecture annotated with digital information — or displays delivering information to the periphery of the observers perception. Recent advances have also brought input and output technology closer to the human, even connecting it directly with the human sensory and neural system in terms of in-body interaction and intelligent prosthetics, such as ocular video implants. Research in that area has to cover both technological and qualitative aspects, such as user experience and usability.

Nano-Electronics and Nanotechnologies

Experience over the last half-century shows that computing power doubles every 18 months (Moore's law). This is primarily due to the ongoing miniaturization of computing elements, which form the basis for the recent explosion of information technologies and electronic products. Even though the semiconductor industry is already producing transistors with critical dimensions well below 100 nm (1 nanometre = one millionth of a millimetre), it is expected that the downsizing achievable with existing technologies will reach a plateau in 10-15 years time, with storage capacities of Terabytes and peak performances of TeraFLOPS for a standard chip (Tera = one million times one million, the present 'Giga-era' being a factor thousand less). These predictions underline the need to explore technological alternatives, which would extend Information Technology capabilities beyond the expected limitations of current semi-conductor (CMOS, or Complementary Metal Oxide Semi-conductor) technology and support the current trend towards diversification and increased complexity of technologies.

On the one hand, developments of new material structures at the nanometre scale, such as nanotubes and nanowires, bring opportunities for multidisciplinary contributions from electronics, material science, physics, chemistry, and biology. On the other hand, nano-electronics and nanotechnologies are expected to have a very strong impact on future economic growth in information technology, as well as in medicine, energy and materials.

The following research challenges are identified:

'System-ability' of Emerging ICT Technologies and Devices

Currently-available deterministic design tools for CMOS will not be useable for the design of systems with over 100 billion devices. Radically-new design methods and tools must be developed to abstract from the designer the new characteristics and variability of nano-devices, while reaching the system-level objectives of low power consumption, reliability, evolvability, and affordable design effort. System-level tools will also need to be complemented by atomistic simulation tools at the device level.

Interfacing Cell-Level Biology with Nano-Electronics

To experiment with and develop new ways to implement bidirectional communication between man-made electronic systems and living entities, at the cellular scale and below, and to develop new information processing or communication elements and interfaces based on biological structures or growth methods. This would open the way to replicate biological properties, such as growing, shrinking and reconfiguring information processing systems. Applications could include new sensors for quick, clean and cheap medical diagnostics, and future health products that counter rising costs in our ageing society.

Future Interconnects for System Integration

To research the basis for radically new on-chip interconnects, both at the architecture and the materials level, and to achieve lower delays and power consumption, while providing required signal integrity and ease of design and manufacturing. Whereas critical dimensions of transistors are now below 100 nanometres, the widths of interconnects are still about ten times larger, and most of the power is dissipated there rather than in active devices, leading to short battery life in laptops and handheld devices. The proposed research will aim at overcoming these limitations.

Post CMOS Devices and Storage

In view of the increasing obstacles in the pursuit of shrinking of silicon transistors according to Moore's law, new concepts must be developed for switches, memories, interconnects or other functions, to provide levels of performance, power consumption, size and/or costs substantially better than the expected limits of current industrial technologies such as silicon CMOS devices. This requires exploring information carriers such as electrons, spins, photons, or phonons, new material structures at the nanoscale, and to devise new ways to represent and process information. For example, devices based on photonics principles would be highly desirable for their integration with interconnects, parallel processing and communication networks.

Nano-Electromechanical Systems (NEMS)

To explore the contributions of NEMS to information processing and communication devices, leading to multifunctional integrated circuits, pushing further down the physical limits; and to explore the application potential of massively parallel arrays of NEMS. Developments in nanoscale mechanics could revolutionize modern Information Technology by bringing back mechanics as an efficient high-speed low-power technology to store and process information.

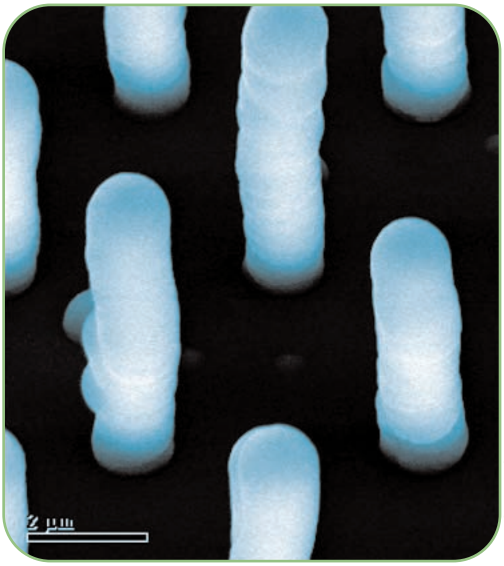

Source: NASA

Array of vertically aligned carbon nanotubes grown using plasma enhanced chemical vapor deposition is intercalated with SiO_2 to provide a bottom-up process for developing interconnects. These vertical interconnects are ideal for DRAM applications and three-dimensional architectures.

Atomic-Scale Information Technology

To explore new computing paradigms based on devices at the atomic scale, where the laws of quantum mechanics take over from classical mechanics, to achieve new information representation, processing, communication and sensing capabilities. Computing capabilities could also be developed inside a single molecule or within specially-crafted structures at the atomic-scale.

Security, Dependability and Trust

There is a need for security, dependability and privacy for communication infrastructures.

The information society is becoming ever more complex, but at the same time also more fragile. Information and communication systems are being interconnected through an increasing use of open communication infrastructures such as the Internet, the 3G Telecom infrastructures and the future Galileo infrastructure, leading to a vast web of interdependencies between these systems. In many areas, such as industry, administration, banking, energy, transportation, telecommunication, public health and defence, this is leading to new vulnerabilities and threats, for example espionage, cyber-terrorism and privacy issues.

Moreover, the widespread use of standard applications and network interfaces facilitates the rapid propagation of accidental faults or deliberate acts of aggression. While these would have been contained in one network or application environment in the past, they can now propagate freely to affect all areas of human activity, with a cascading effect anywhere in the world.

Further security issues arise in the context of new technology deployments, such as small networked objects with little resources (e.g. RFIDs, sensors and future nano-systems), smart artefacts, and emerging technical developments, such as quantum computing and quantum cryptography.

In the face of these technical and societal changes, there is an obvious need to provide the users of new and emerging information and communication systems with better security, dependability and privacy protection.

The following research challenges have been identified:

Ambient Security, Dependability and Privacy

The mass diffusion of networked devices, and their increased embedding in the environments that surround us, must be supported by mechanisms for enhancing confidence in their usage. Traditional mechanisms based on setting boundaries and creating firewalls cannot cope with the large numbers of devices and their mobility. Moreover, security and cryptographic mechanisms are needed for very small networked objects, such as RFIDs and sensors, which lack the necessary resources (such as energy) for using existing techniques and protocols. Citizens also need to be assured of the security of the systems on which they depend.

The challenge is to provide a guaranteed level of security, dependability and privacy in

Source: NASA

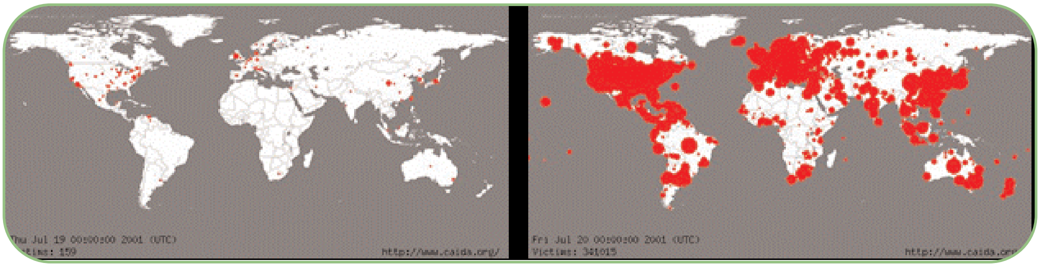

Modern worms have demonstrated that they can infect tens of thousands of computers worldwide in a few hours..

such a large-scale, heterogeneous and dynamic context with, in some cases, very limited resources. This may be achieved through new mechanisms such as bio-inspired, adaptive and evolvable mechanisms, cryptographic algorithms and security mechanisms. This will need to be supported by technologies for provable security and dependability, for secure programming, as well as mechanisms and tools for assessing and proving the trustworthiness of an open and complex information and communication infrastructure.

Dynamicity of Trust

The lack of trust is one of the main barriers for the establishment of a secure and dependable Information Society. Future ICT systems will involve thousands of millions of devices, and will no longer be able to depend on setting boundaries and firewalls for its security. Instead, it will require a capability for managing and negotiating trust relationships, adapted to the level of security required in a given situation. The challenge is then to obtain greater understanding of partial trust, security-based trust (where trust follows from security), and trust-based security (where security is achieved through a trusted partnership), and to use this understanding to realise a high level of trust of the citizen in the deployment, economic viability and social acceptance of systems and services. This will require expertise and joint research in several fields outside ICT, such as economy and sociology.

Quantum Technology and Cryptology for Information Security

Quantum technology rests on the use of photons (light particles) rather than electrons, as in the silicon industry (transistors). On the one hand, it offers a range of opportunities for future cryptology, for example through the possibility to generate truly random values or to distribute bits of information with absolute security. On the other hand, it is well known that quantum technology is a threat to current cryptographic methods, since a quantum computer may break current asymmetric cryptographic techniques in a few seconds 10-15 years from now. A deeper understanding of the impact of these technologies on cryptology is mandatory, and one should therefore study the assumptions on which quantum computing may act and their consequences for cryptographic methods. At the same time we should lay the foundations for new algorithms to effectively resist code-breaking attempts by quantum computers. This challenge obviously needs to be addressed in close collaboration with the Quantum Information Processing and Communication (QIPC) community.

Bio-ICT Synergies

Convergence of the biological, information, nano-, and cognitive sciences is opening new prospects for science and technology. A major limiting roadblock is sharing and understanding knowledge from different disciplines that have developed their own terminologies and concepts. However, over the last decade successful inter-disciplinary collaborations have emerged. In neuro-informatics, ICT learns from the ways information is processed in our neural system by studying the associative capabilities of the human brain, or from the behaviour of ant colonies and bee swarms and how these self-organise. Engineers and biologists join forces in 'biomimetics', where mechanisms found in nature are used to inspire technological design, of, for example, 'lifelike' computers capable to growing, self-repairing and adapting. ICT does not only draw benefit from biological insights, it also furthers them, e.g. in computational biology to model the living cell. Interdisciplinary collaboration seeking bio-ICT synergies will lead to new and different perspectives, both in biology and in ICT.

A joint focus on major challenges will help to bridge the gap between the bio- and techno-disciplines, and to mobilize their synergies. The gap is threefold: conceptual, temporal, and physical.

Conceptual Gap

A common conceptual basis for Bio-ICT research is needed. This can be achieved in joint modelling projects, leading to better biological system modelling, new types of computational architectures and to new ways of computing that reflect the salient 'information processing' type of activities in living organisms (cells, neurons, genes, metabolism, etc.).

Temporal Gap

The dynamic-static dichotomy between biological and ICT systems is one of the key hurdles in the natural combination of living and artificial systems. The static nature of the latter is incompatible with the processes of change that govern the former. This leads to biological systems having to take all the burden of adaptation in the synergistic combination, and to ICT systems prone to becoming obsolete, malfunctioning and disturbing.

Physical Gap

The physical incompatibility between biological and ICT systems ('wet-dry dichotomy') requires work on interfacing, embedding and on hybrid combinations of biological and ICT systems. Though this is usually viewed as prosthesis work, there is a much wider potential. For instance, brain-machine and neural interfaces could cover a large field of applications ranging from motor prostheses to any kind of sensory prostheses. One could even imagine ICT-based interactions directly within metabolic processes in organisms (smart drugs, for instance).

Three broad challenges have been identified as major steps towards addressing these gaps.

New Computational Modelling Paradigms

The aim is to find new ways of computing capable of capturing, relating and integrating the different levels of complexity in biological systems — from the molecule to the cell, tis-

Connecting the natural, digital and artificial worlds: cell modelling, synthesis and hybrid systems.

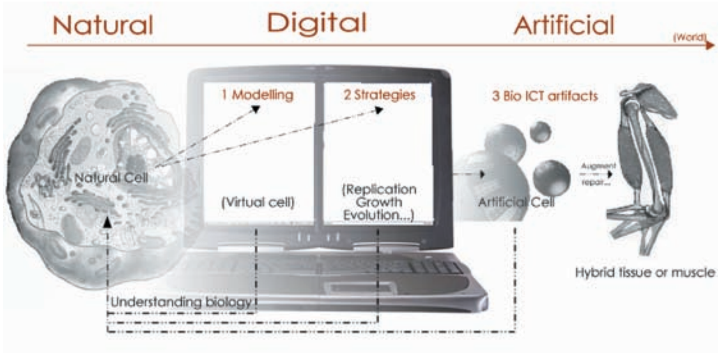

sue, organ, system, individual, population and ecosystem. A thorough understanding of information processing in biological systems could lead to life-like hardware consisting of large numbers of simple devices that operate in a highly parallel fashion at comparatively low speed and with very low power dissipation. New bio-inspired computation, capable of autonomy and self-organisation, could be crucial in many other areas, for instance in the management of renewable energy sources, such as large, distributed parks of wind turbines, or of extended sensor networks surveying airports.

Bio-Inspired Strategies of Growth, Development and Evolution

Technological systems should mimic the extraordinary capacity of biological systems to grow, adapt, self-assemble, replicate, heal, self-organize, and evolve. These different strategies of change are not independent but operate at different time scales, either at the individual or population level. The idea of using these strategies for problem solving is, of course, not new – think of genetic algorithms – but recent advances in understanding of biological processes have not been exploited yet (for example, for more complex genotype-phenotype mappings). This could lead to potential breakthroughs, not only in software but also for new types of growing, self-assembling or evolving hardware at the nano- and micro-scale or intelligent materials applicable in ambient interfaces (displays, active surfaces).

Bio-ICT Artefacts

Bio-ICT artefacts, i.e. artificial entities seamlessly integrated into biological systems, such as artificial retinas or physiologically coupled artificial limbs, can have a significant social and industrial impact. Examples include diagnostic technologies or controlled drugs release. The main research challenges are: (a) to develop new information theories and modelling techniques to capture how living systems achieve their remarkable robustness and performance such as in sensing, action, memory, adaptation, homeostasis or others, (b) to validate the results with respect to real biological systems; (c) to apply these theories to develop new perspectives in domains other than biology; (d) to apply these theories to the design of ICT technologies, that replicate, complement or substitute for these basic capabilities of living systems or are interfaced to them; (e) to study how such technologies can sustainably adapt and evolve to match, over long periods of time, with evolving needs, and be compatible with various natural processes of change (e.g., growth, learning, aging).

Intelligent and Cognitive Systems

Intelligent and cognitive agents that can perceive, understand, and interact with their environment, but also evolve and learn in order to achieve human-like performance in activities requiring context-specific knowledge, are far beyond the current state of the art and will remain so for many years to come. A promising approach to tackle this grand challenge is to study and develop complete agents that are: embodied and self-sufficient (sustain themselves for a long time); situated (acquire information about the environment through their own sensory system); and autonomous (function independently of external control). Intelligence requires a body, hence a core topic is how intelligent behaviour emerges from the interaction of brain, body and environment. Motivated by recent research in robotics, smart materials, neuroscience and cognitive sciences, the focus is on the following strategic areas: (i) the role of growth and development in cognitive systems, (ii) conceptual understanding of the action/perception coupling of embodied agents — with the environment, with humans and with their peers, (iii) research into how materials and morphology (body shape, position and nature of sensors) contribute to the agent's successful operation, and (iv) development of a theoretical framework. The latter should not be seen as a separate research theme, but is to be developed alongside each of the research themes that are detailed below. These research themes address the challenge of embodied intelligence from different, partially overlapping angles.

Mind-Body Co-Development and Co-Evolution

The standard approach in evolutionary robotics is to use an existing robot as it is, and evolve its controller. However, in order to maximally exploit the design power of evolution and development, controllers and robot morphologies have to evolve simultaneously. Permanent interaction of an agent's body with the environment during growth enables its 'mind' to develop. However, we lack insight into how the development of organisms is influenced by the interaction with the real world. We also lack growable materials for mimicking developmental processes (although modular growth and self-assembling materials have been demonstrated) and materials like artificial skin or muscles, and deformable tissues. These are necessary ingredients to build embodied artificial systems inspired by developmental processes in nature.

Developmental Robotics

The main goal of this theme is to mimic aspects of ontogenetic development (the growth of an organism from birth to maturity) using robots, and learn how cognitive capabilities emerge from this development. During this process the agent continuously interacts over extended periods of time with its physical and social environment, thereby developing its skills. The mechanisms underlying ontogenetic development are still largely unknown, for example, how sensory-motor and social interaction 'cooperate' to produce cognition. Similarly for 'embodiment', neither the mutual influence of sensor-motor and information

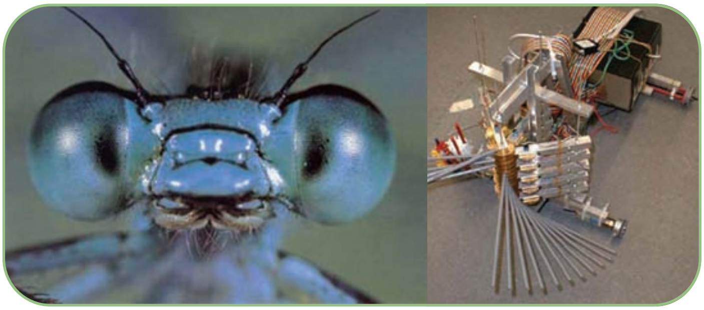

Morphological computation. Facet distribution of insect eye and robot system 'Eyebot' with an artificial compound eye, which can autonomously modify the angular positions of the individual facets. The physical arrangement of the facets — the morphology of the facet eye — performs a kind of 'computation', thus facilitating neural processing.

processes, nor why an agent should engage in ever more complex interactions are well understood. Proper materials and sensors are also lacking (haptic sensors, for example, are still very poor compared to human skin on the finger-tips), as is the knowledge how to integrate very complex systems.

Systems and Materials that Can Grow

Recent research strongly suggests that physical properties of agents, like their shape and their material, play an essential role in shaping their behaviour. Through growth, biological organisms can form highly complex morphological structures that perform important functions for an agent in real time. An example from nature is the growth of the eye, which is regulated through neural and biochemical circuits based on sensory input while interacting with the environment. Although it is not expected that growable materials will become available any time soon, some aspects of physical growth can be studied through advanced principles of autonomous modular systems or self-assembling materials. The challenge is to explicitly apply concepts of morphology and growth in exploring embodied artificial systems. This could lead to robots that can be more easily deployed on the factory floor and in private homes, as well as micro-scale machinery with integrated sensing and information processing abilities for medical use.

Morphological Computation

As complex dynamical systems, embodied agents are fundamentally different from computers: they perform processes by 'morphological computation', in which control not only resides in the brain or a microprocessor, but is also distributed throughout the organism and deeply embedded in it. Examples are facet distributions in insect eyes, stiffness properties of muscle-tendon systems, and deformability of tissue on finger tips. In the real world the time factor frequently forces agents to exploit the materials and the morphology, including that of the brain, for sensory-motor control. This different way of organizing a computation is still poorly understood and requires theoretical work as well as case studies. The research must address aspects such as: shape, materials, sensors, brain structure, signalling mechanisms, and morphological change.

Design for Emergence

Behaviour always results from interaction of an agent with the environment, i.e., behaviour is emergent: it cannot be understood (or designed) on the basis of the internal control program (or 'brain') only. The fundamental research question then is how purposive (goal-directed) behaviour of an agent can be achieved without destroying the emergent nature of this behaviour, i.e., without designing the goal into it.

Software-Intensive Systems

Software has become a central part of a rapidly growing range of applications, products and services in all sectors of economic and social activity. Systems in which software interacts with other software, systems, devices, sensors and with people are called software-intensive systems. Examples include embedded systems in cars and airplanes, (wireless) telecom systems, web services for business applications, control systems for large infrastructures (electricity, transport), and medical support systems (diagnosis, surgery, patient care, e-health for the elderly, etc.). Everyday life increasingly depends on such systems and they are becoming ever more distributed, heterogeneous, decentralized and interdependent.

In the next ten to fifteen years systems will have to operate in large, open and non-deterministic environments, while demands for quality will increase dramatically. Instead of reacting to this by developing computer-oriented systems that people have to adapt to, one has to develop human-oriented systems into which computers integrate seamlessly. Future software-intensive systems should be adaptive to unforeseen changes and situations, while remaining secure and reliable.

Today's grand challenge is then to achieve reliable, secure and trustworthy software-intensive systems that can operate under these conditions. This will require developing practically useful and theoretically well-founded principles, methods, algorithms and tools to model, program and maintain these systems efficiently throughout their lifetime, and possibly forever.

The following three research challenges have been identified.

Engineering Adaptive Software-Intensive Systems

In the prevalent approach to software construction the full set of program behaviours and features is decided when software is designed. This is no longer adequate for tomorrow's pervasive, open and highly dynamic systems. Adaptive systems are required that can adjust to and respond to unforeseen changes in their environment, to evolving requirements, to technologies becoming obsolete or new ones being introduced, and to newly gained knowledge. Such systems will no longer be produced *ab initio*, but more and more as evolving compositions and/or modifications of other, existing systems and components, often created, configured and adapted at runtime, while still secure and reliable. The challenge is to develop algorithms, methods, tools and theoretical foundations that enable effective design of adaptive systems which harness and control properties that emerge through the interaction of systems and their environment.

Managing Diversity in Knowledge

An unforeseen growth of the volume and diversity of the data, content and knowledge is being generated all over the globe. The traditional approach of knowledge management and engineering is top-down and centralised, and what can be expressed and how is fixed at design time. As applications become increasingly open, complex and distributed, the

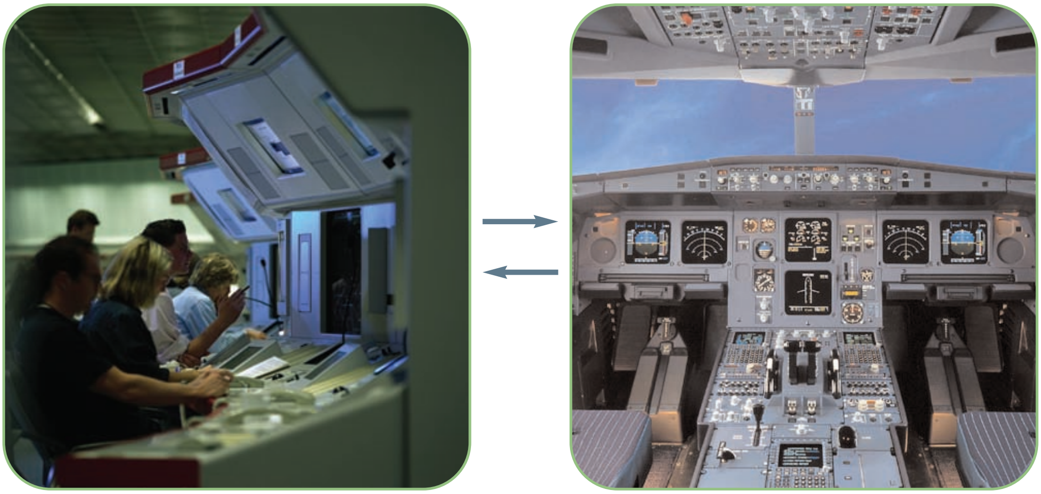

Images: DFS Deutsche Flugsicherung GmbH (left), Airbus (right).

Efficiency and safety of air traffic crucially depend on the proper operation of software-intensive systems in the airplane and at ground control, including their intercommunication.

knowledge they contain can no longer be managed in this way. It must be combined with a new, bottom-up approach in which the different knowledge parts are designed and kept 'locally' and independently, and new knowledge is obtained by adaptation and combination of such items. The challenge is to develop theories, methods, algorithms and tools for harnessing, controlling and using the emergent properties of large, distributed and heterogeneous collections of knowledge, as well as knowledge parts that are created through combination of others. The ability to manage diversity in knowledge will allow the creation of adaptive and, when necessary, self-adaptive knowledge systems.

Eternal Software-Intensive Systems

Information represents one of society's most important assets. Continuous and up-to-date access to long-lived and trustworthy information systems that do not age and break, is essential for the economy and for preserving our cultural assets. The traditional approach to the software life-cycle, which presumes that after some time a system becomes obsolete and its content and functionality is lost, needs to be replaced by one that is based on 'eternal' systems, allowing content and functionality to be passed from one generation to the next. Those systems should be able to preserve and update their original functionality and properties in a machine independent way by making it easy to re-program them — or even enable the systems to re-program themselves — in order to take into account changes in the environment in which the system functions. The challenge is to organise software-intensive systems so that they can survive and evolve in a constantly changing world.

Quantum Information Processing and Communication

QIPC is a relatively young multidisciplinary scientific field. It emerged in the 1980s as an alternative to traditional computing, which will be faced with its physical limitations before long. Based on the laws of quantum mechanics, QIPC exploits fundamentally new modes of computation and communication. It holds the promise of immense computing power beyond the capabilities of any classical computer, it guarantees absolute secure communication, and it is directly linked to emerging quantum technologies like new types of sensors. It therefore has the potential to revolutionise many areas of science and technology.

European scientists have been at the forefront of QIPC research since the beginning and the FET programme has given them substantial support and guidance. This has led, among other things, to the elaboration of a common European strategy. The recent roadmap document 'QIPC — Strategic report on current status, visions and goals for research in Europe' was written by the most prominent scientists. It is an impressive joint effort of the European scientific community and a good example of how European research can be adequately structured with the help of the EU. The roadmap identifies the major challenges for the next decade in QIPC's three interconnected sub-fields.

Quantum Computing

At present, a number of physical systems are under investigation for their suitability to implement a scalable quantum processor and remarkable progress towards demonstrating the basic building blocks has been reported. Within the next five years is expected to realise a few-qubit general purpose quantum processor including error correction as a test bed for demonstrating quantum algorithms. Within the ten years from now a first goal is the development of hybrid technologies and architectures for quantum computation, including interfaces between them. A central goal here is to construct laboratory models of quantum computers outperforming their classical counterparts. In parallel, special effort must be made to further develop few-qubit applications. Another goal is to develop special purpose quantum computers with a large number of qubits. These are quantum simulators: programmable quantum systems whose dynamics can be engineered such that it reproduces the dynamics of other many-body quantum systems of interest.

Quantum Communication

The objective of quantum communication is to transfer quantum states from one place to another using entanglement and other non-local characteristics of quantum physics. At present long-distance quantum communication can only be achieved by light signals (photons), but these are dampened with distance. Classical telecommunication solves this problem by 'repeaters' which amplify and reshape the transmitted signal. However, a classical repeater is not suitable for quantum states. Therefore a major scientific and technological challenge is to develop quantum repeaters. This is absolutely necessary for quantum communication both in optical fibre and in free space in order to reach the following two objectives: to demonstrate the feasibility of 'real world' long-distance quantum communication, and to increase by several orders of magnitude the qubit transfer rate. In parallel, building quantum networks has to be a major priority and has to complement the development of computer architectures in order to

In order to grasp the potential of QIPC one has to overcome everyday intuition which is implicitly based on classical physics. Conventional computers perform calculations on 'bits' which can only take the values 0 and 1. Quantum computing is based on 'qubits'. These are two-state quantum systems being in a so-called 'coherent superposition' of the two states '0' and '1'. Qubits can be prepared using, for example, an atom in the ground state (0) and an excited state (1), the spin of an atomic nucleus (up or down), etc. Similarly, registers of several qubits can simultaneously represent many numbers in quantum superpositions. Quantum processors can then evolve initial superpositions of encoded numbers into different superpositions. During this evolution the encoded numbers are affected, and the result is a massively parallel computation performed in a single component of quantum hardware. The drawback is that during such a computation the smallest disturbance from the environment may ruin the delicate superposition (decoherence). Qubits can also become entangled. This property can be used for error correction during computation, or for more efficient transmission of information.

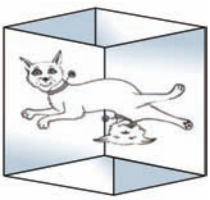

Tobias Baanders, CWI

Quantum Superposition
According to the laws of quantum mechanics, the cat in the box (named after founding father Erwin Schrödinger) can be dead and alive at the same time. The animal is in a superposition of both states with a certain probability. Only if we look inside the box, the observed cat is either dead or alive: observation causes the quantum superposition to collapse to one definite state.

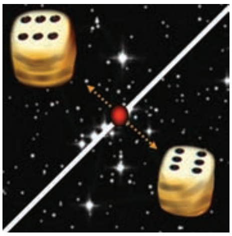

Quantum Entanglement
A pair of futuristic quantum dice is generated in a common source. If one dice is thrown and shows on observation '6', the other one, no matter how far apart, will also definitely show '6'. The two dice are 'entangled': born from one source, they remain correlated forever.

build an entire working quantum computer. Another challenge is to develop quantum cryptography into a commercial product. Quantum cryptography is an outstanding result of the scientific efforts in quantum communication that allows absolute secure transmission of key codes. When transmitting a message encoded with a secret key, eavesdropping is immediately noticed. The reason is that the laws of quantum mechanics imply that the mere act of observing a quantum bit modifies it so that it 'collapses' into a different state.

Quantum Information Theory

The aim is to develop quantum information theory, the quantum analogue of classical information theory (basic to all contemporary computation and communication systems). This very broad research priority pursues several challenging goals. The development of new quantum algorithms remains one of the cornerstones of research. Another one is the development of new computational models and architectures, because there are many different ways how to make quantum systems compute. Actually one does not exclude another and each can lead to the formulation of different quantum algorithms. Despite its amazing power, a quantum computer will be a rather fragile device, susceptible to disturbances and errors. A number of quantum error correction and purification schemes are being developed that can make computation or communication stable. As all communication channels are subject to some level of noise, transmission through noisy communication channels is also a priority. The theory of entanglement is at the heart of QIPC and can lead both to a deeper physical understanding and to novel applications. In particular multi-partite entanglement is expected to have an impact on novel protocols for quantum communication.

More information on QIPC activities in FET: http://www.cordis.lu/ist/fet/qipc.htm

Complex Systems

The nature of science is changing; its focus is shifting from parts and pieces to coherent wholes. For our technological age, confronted by problems of staggering complexity and seemingly impenetrable webs of cause and effect that defy straightforward analysis, this new focus of science comes not a moment too soon.

Challenges in an Interconnected World

Rapidly increasing global interdependence, through communication networks like the Internet, has opened unprecedented opportunities for innovation but also brought with it bewildering new risks. The Internet has revolutionized everything from international banking and business management to library science and air traffic control. Yet we seem powerless to prevent serious network inefficiency — today more than two thirds of global email is wasted 'spam'.

Complexity is also showing in areas like the global spread of epidemics (SARS, Avian flu), with worldwide repercussions, or finding for planet Earth a sustainable path into the future, learning to manage her climate and ecosystems. Doing so means understanding a highly complex system of which mankind is an important part, and where the origins of system complexity are as often social and political as they are technological, physical or biological. The demands of complexity will not go away, and traditional approaches are reaching limits of technological and computational feasibility. A science of complex systems that embraces a holistic system view might be an answer.

More is Different

In 1972, the (future) Nobel Prize-winning physicist Philip Anderson published an article in the journal Science titled 'More Is Different.' Anderson was exploring what happens when a number of elements — atoms or molecules, perhaps ants or even people — interact with one another. Interactions lead to messy interdependence; they ratchet up the difficulty of understanding what goes on and why. But Anderson's point was that interactions also lead to 'emergence' — to the spontaneous appearance of features that cannot be traced to the character of the individual parts. Three decades later the deep importance of emergence to every science and every branch of engineering has become undeniable. The present massive paradigm shift in science will have profound effects on every aspect of human life.

Complex systems occurring in the real world generally share certain features: a large number of components (chips, computers, bacteria, people) with strong mutual interactions, rich dynamics with patterns and fluctuations on many scales of space and time, and perpetual evolution. An ant colony is similarly complex, in this sense, as is crowded traffic, the Internet, or a living cell. Many real-world systems share the capacity to organise into highly versatile multi-level structures in order to react and adapt to external conditions. In this respect, complex systems are 'more than the sum of their parts', exhibiting features that cannot be traced to the character of the individual parts.

The traditional way to understand and design such systems is by 'divide and conquer', that is, by breaking problems into smaller component parts. In contrast, complexity science recognizes that system behaviour is often irreducible. Taking its inspiration from the real world, complex systems science is promoting a shift from purely logic-based rational design to a distributed design approach. It harnesses the

Images: bigstockphoto.com, Straßen.NRW, INRIA, MNHN

Complex systems can differ widely in appearance, like an ant colony, crowded traffic, the Internet, and a living cell, but they all share several common features.

capacity of self-organisation that is adapted to the natural complexity and changeability of the real world, both natural and man-made.

The Internet as Organically Growing Artefact

Communications traffic will increase enormously in tomorrow's world, as pervasive sensors share information to do the 'housekeeping' chores of an information-centric world. Billions of intelligent devices are being integrated into functioning networks, and, in the future, not only people, but also innumerable semi-intelligent software agents will use a vastly more powerful Internet. This Internet will behave more like an organism growing and evolving in the absence of any central control.

The sheer volume and diversity of information presents extraordinary challenges to gathering, storing, sharing and finding of information. Complexity science points to a decentralized approach, in which such networks function smoothly by organic growth and where a distributed approach to information search will allow a 'distributed Google' with far greater reach and power.

Exploring the Petabyte Range

Molecular biology has raced past the milestone of the human genome project, and has fully sequenced the genomes of a number of organisms. Now it is recognized that its greatest challenge lies in understanding how thousands of genes and other bio-molecules work together in feedback loops of immense complexity, in order to create the robust and adaptable forms we know as life.

The radical increase in the amount of IT-generated data on living and social systems brings about new challenges related to the sheer size of data and will transform present-day science and engineering. Real-time data from sensors in ecosystems, the human body or in business networks, if combined with more powerful means of exploiting this wealth of data, will radically advance our knowledge of multi-level dynamics of living and social systems, and thereby our capacity to conceive effective therapeutics and effective policies (e.g., in case of epidemics).

The above offers a taste of some of the potentially revolutionary approaches now being explored, ranging from efforts to engineer the 'next generation' search engines that will vastly eclipse the capabilities of today's Google to studies discovering deep links between the structure and growth patterns of artificial systems and living organisms. The research described here clearly illustrates two principal truths of the complex systems viewpoint — first, that complex systems of all kinds reveal deep similarities that make it possible to learn about all by studying any one, and second, that the most profitable inquiries are not abstract but tied to specific real-world problems.

More information on Complex Systems activities in FET: http://www.cordis.lu/ist/fet/co.htm
A 'living roadmap' on complex systems: http://www.once-cs.net.

Tera-Device Computing

Today's most advanced integrated circuits already include more than a billion (10^9 = Giga) transistors. With semiconductor technology still downsizing its components, and with the introduction of even smaller alternative devices, integrated systems with one Tera (10^{12}) devices are likely to be manufacturable within 10 to 15 years. Current approaches to develop the Giga-scale circuits and the associated software already require excessive time and effort, and will not be scalable to the Tera range. (A system is called scalable if it can adapt to increased demands.) New processor architectures must be researched, and a major research effort is needed on the associated system software and design methodologies, to enable the development of Tera-device computers within reasonable time and effort, while guaranteeing scalability, reusability, limited power dissipation, fault tolerance and adaptability to a wide range of applications.

These computing engines will provide the performance required for the mass-market and specialised applications of 2020, including smart phones with cognitive abilities, medical devices for each patient, networked interactive entertainment, embedded electronics in cars and airplanes, up to large computing centres for engineering and scientific applications. The required radically new developments are an opportunity for Europe to reinforce its position in computing, based in part on its firm foothold in embedded computing engines.

A drastic solution for Tera-device computing consists of promoting parallel processing with massive numbers of less complex computing units than those of today's processors. This solution would mitigate design efforts while guaranteeing scalability and would take account of fault tolerance and energy awareness.

Illustration: Tobias Baanders, CWI

Striving for the Tera era.

Hardware Architectures

Research is required into the architectures of simple processor cores and complex interconnection circuits that will combine into networked multi-core systems that are scalable, heterogeneous, and customisable. It should address different potential solutions such as shared memory, message passing, fine-grain parallelism, and unconventional architectures.

System Software

New approaches to programming are needed to express and manage the complexity of highly parallel programs, in particular to allow portability of code onto different parallel architectures. Other research concerns operating systems, run-time environments and automatic program parallelisation, in order to extract parallelism from existing serial programs.

Design Methodologies and Development Tools

The goal is to efficiently develop the system most suited to a given application, and verify and validate it. Factors to be taken into account include power efficiency and control, and reliability/fault tolerance at manufacturing time and run-time.

Published by:
ERCIM EEIG
2004, Route des Lucioles
BP 93
F-06902 Sophia Antipolis
http://www.ercim.org

Text editor: Henk Nieland

Scientific coordination:

Dimitris Plexousakis, ICS-FORTH, Greece

Alois Ferscha, University of Linz, Austria
(Pervasive Computing and Communications)

Colin Lambert, University of Lancaster, UK
(Nano-electronics and Nanotechnologies)

Michel Riguidel, ENST-Paris, France
(Security, Dependability and Trust)

Fernando Martin-Sanchez, Institute of Health Carlos III, Madrid Spain
(Bio-ICT Synergies)

Rolf Pfeifer, University of Zurich, Switzerland
and Alois Knoll, Technische Universität München, Germany
(Intelligent and Cognitive Systems)

Martin Wirsing, Universität München, Germany
(Software Intensive Systems)

Links:

Beyond-The-Horizon project: http://www.beyond-the-horizon.net/
Full report: http://www.ercim.org/publication/policy/policy.html