European Research Consortium
for Informatics and Mathematics

ERCIM

www.ercim.org

EC Consultation on Future European Union Research Policy: Strategy for ICT in Europe

LEGAL NOTICE

ERCIM

ERCIM - the European Research Consortium for Informatics and Mathematics - aims to foster collaborative work within the European research community and to increase co-operation with European industry. ERCIM is the European Host for the World Wide Web Consortium (W3C).

Leading research institutes from eighteen European countries are members of ERCIM. A member institute must be a leading research establishment in its country with excellent links to both national and international research community. All ERCIM members are national centres of excellence, independent of specific commercial ties, they have a strong involvement in the research programmes of the European Union and joint projects with both small and medium size enterprises and large industrial companies.

The members of ERCIM are: -

Austria	AARIT	http://www.aarit.at/
Belgium	FNRS & FWO	http://www.fnrs.be/
Czech Republic	CRCIM	http://www.utia.cas.cz/CRCIM/home.html
Finland	VTT	http://www.vtt.fi/tte/staff/lis/vttercim.html
France	INRIA	http://www.inria.fr/
Germany	FhG	http://www.fraunhofer.de/fhg/EN/index.jsp
Greece	ICS-FORTH	http://www.ics.forth.gr/
Hungary	SZTAKI	http://www.sztaki.hu/
Ireland	Irish Universities Consortium	
Italy	CNR	http://www.cnr.it/sitocnr/home.html
Luxembourg	FNR	http://www.fnr.lu/
Netherlands	CWI	http://www.cwi.nl/
Norway	NTNU	http://www.ntnu.no/
Slovakia	SRCIM	http://www.science.upjs.sk/sscs/poster/poster.htm
Spain	SpaRCIM	http://www.sparcim.org/
Sweden	SICS	http://www.sics.se/
Switzerland	SARIT	http://www.sarit.ch/
United Kingdom	CCLRC	http://www.cclrc.ac.uk/

ERCIM OFFICE

2004, Route des Lucioles
BP 93
F-06902 Sophia Antipolis
France
Tel: +33 4 92 38 50 10
Fax ++33 4 92 38 50 11
E-mail: office@ercim.org

Additional information on ERCIM is available on the Internet, it may be accessed through the ERCIM server (http://www.ercim.org).

Contents

1. EXECUTIVE SUMMARY

ERCIM (www.ercim.org) represents 12000 researchers in major institutions and consortia of universities across Europe and - via the intercommunication of communities with each ERCIM partner - reaches at least one order of magnitude further. ERCIM is the European arm of W3C (World Wide Web Consortium) hosting offices in European countries and contributing heavily to the development of the web. ERCIM institutions also have more than 100 spin-out companies. At the May 2004 ERCIM meeting it was agreed that ERCIM should produce this document at the suggestion of officials from EC DG INFSO Unit F1. Of course, ERCIM has contacts with many parts of the EC and the document could - and should – have wider circulation. The document was approved at the November 2004 meeting of the ERCIM Board of Directors as our strategic view.

The rapid decrease of ICT-costs versus performance makes ICT appear everywhere. In each decade performance versus cost improves by factors of 100-1000 for the basic technologies. This unremitting advance in capability makes ICT the most important change driver in society. Central to a future European R&D (framework) programme is the determination of how to behave and invest in order to be in the forefront of using ICT to develop and finance the socio-economic systems that we wish to have in Europe.

Thus, it is of the utmost importance that the EC continues to increase investment in ICT R&D. ICT is a scientific discipline in its own right, but uniquely one that brings heavily leveraged benefits to other sciences, technology, economic activity, arts and humanities, healthcare and the environment. These benefits are realised when there is an integrated activity from:

(a) fundamental basic research (providing the underpinning assurance for the technology) through

(b) generally used well-engineered components which interoperate (providing the fabric or computing surface) and on to

(c) specialised products and services for particular application areas (leading to direct benefits in wealth creation through both the products and services and the utilisation of those by the business or organisation to achieve its goals).

Importantly this is NOT a linear chain but there is a continual need for interaction among basic/fundamental R&D, R&D to develop commonly applicable components and R&D for specialised application-aware products and services.

The overriding question in planning the volume and the profile of European ICT research is to find the relevant issues for European business and society, select those where Europe has (or can gain) a leading position and to finance them to such a

degree that a competitive advantage can be achieved relative to other regions of the world.

The key to the successful deployment of ICT in Europe – for wealth creation and improvements in the quality of life – comes through excellent fundamental R&D-led developments based upon developing and emerging technologies. These today are partly-thought-out designs or research prototypes and bringing these to the status where they are proven and ready for take-up and mass utilisation through commercial exploitation or open standards activities we consider a prudent, risk-aware approach to utilising the scarce R&D resources for ICT in Europe. Of equal importance in the long run is selected basic research on risky, bold, over-the-horizon issues which may bring within tractability previously intractable requirements.

There are many ways to structure a discussion on Future and Emerging Technologies for ICT. Here we have chosen to consider three main aspects:

(1) user and system centric components, together with connectedness, for architected application systems;

(2) electronic, storage, computing and communication components to support (1);

(3) systems development methods to construct components of (1) and applications on (1).

Within this framework we propose R&D issues that are crucial components in Europe's future knowledge profile.

For further information or discussion please contact: Keith.G.Jeffery@rl.ac.uk Tel: +44 1235 44 6103	Professor Keith G Jeffery Director, IT CCLRC Rutherford Appleton Laboratory Chilton, Didcot Oxfordshire OX11 0QX UK

2. INTRODUCTION

Importance of ICT

The rapid decrease of ICT-costs versus performance makes ICT appear everywhere. In each decade performance versus cost improves by factors of 100-1000 for the basic technologies. This unremitting advance in capability makes ICT the most important change driver in society. Businesses and public administration are changed and renewed. To continue doing things as they were done before leads to being out-competed by those who utilise the new technological possibilities. Businesses and public administration must continuously reinvent themselves. Those that do not are doomed to lose in the global competition. This goes for individual businesses, whole branches of activities, and for nations – and groupings of nations such as the EU - as well.

Thus, it is of the utmost importance that the EC continues to increase investment in ICT R&D. ICT is a scientific discipline in its own right, but one that brings heavily leveraged benefits to other sciences, technology, economic activity, arts and humanities, healthcare and the environment. These benefits are realised when there is encouraged an integrated activity from fundamental basic research (providing the underpinning assurance for the technology) through generally used well-engineered components which interoperate providing the fabric or computing surface and on to specialised products and services for particular application areas leading to direct benefits in wealth creation through both the products and services and the utilisation of those by the business or organisation to achieve its goals. Importantly this is NOT a linear chain but there is a continual need for interaction among basic/fundamental R&D, R&D to develop commonly applicable components and R&D for specialised application-aware products and services.

There are many examples of successful leveraging of basic ICT R&D, and equally many examples of failed systems where the fundamentals were not researched and developed. Huge cost overruns and abandonment of ICT systems – especially revealed in government projects but present also in the commercial world – commonly are due to poor fundamentals. In contrast Tim Berners-Lee (a British mathematician working as a systems developer at a European fundamental physics research laboratory) took three concepts from earlier fundamental ICT R&D and produced the most successful ICT development of all time.

Central to a future European IST programme is the determination of how to behave and invest in order to be in the forefront of using ICT to develop and finance the socio-economic systems that we wish to have in Europe.

Knowledge is the critical resource in modernising our society, both knowledge utilised and generated in development of new products and services, and knowledge that permits the easy assimilation of the new knowledge in, and for the benefit of, the existing society. High "assimilation" of knowledge and its utilisation for benefit is based on a high education level in society as a whole, covering the whole population. This is a necessary prerequisite for being successful in renewing public services, and to find broad markets for ICT-loaded products and private services, e.g., in e-business.

For the development and operation of new products, services and systems the situation is different, and deep specialist knowledge is the crucial factor. All of this is increasingly being done in a stiff global competition. This is particularly so for mass-produced generic 'black box' products, which can be made in comparative isolation from cultural particularities of the products' users. The more intimate the relationship between the ICT product and the socio-economic particularities of their users, the greater the chance for smaller players (especially SMEs) to develop non-global but innovative and wealth-creating businesses that are well-tuned to improving the quality of life for employees and customers. Independent of this, the prerequisite for success is world-class knowledge of ICT issues relevant to the tasks to be performed.

The overriding question in planning the volume and the profile of European ICT research is to find the relevant issues, and to finance them to such a degree that a competitive advantage can be achieved relative to other regions of the world. We have to consider both our ability to build the systems that are needed to create the good society, and products that have the potential of financing the desired society.

A basic challenge to the planning of future EU research in ICT is to determine those issues in basic research that are crucial components in our future knowledge profile, and to pursue those with vigour. The challenge is to select the most appropriate research and development possibilities and to resource them sufficiently to maximise the chances of success.

Another challenge is to reduce the time gap between knowledge creation (through research) and the uptake of the new knowledge in industry, public administration and society at large. The success of the Lisboa declaration depends to a large degree on the uptake speed. Knowledge comes wrapped in the heads of people. Universities are therefore central to providing staff capable of effective technology transfer. New candidates from universities should be equipped with the newest best knowledge, which in turn will be transferred to their employers.

Furthermore, achieving a high degree of interplay and mutual impact between applications of ICT and basic ICT research seems to be correlated with fast upsurge of new businesses, both within ICT as such and within application areas. Encouraging examples are seen in USA of new business creation around universities that receive generous funding of basic technological research.

A continuation of a current practice of forming research consortia of partners from industry, research institutes and universities is recommended in order to support this rapid interaction of basic, strategic and applied R&D. This provides support for the speed of uptake of research results (new knowledge) in society at large, as compared to the emphasis put on bringing out short-term beneficial effects which are mostly restricted to the partners of the consortia.

2.1 Pervasiveness

ICT is everywhere – packaging, clothes, domestic devices, leisure devices, business systems, environmental systems, health systems etc. We cannot ignore it. It will exist. Therefore we should develop and utilise ICT to the improvement of the human condition (especially in Europe) through wealth creation and improvements to the quality of life.

There is an argument that, because of its pervasiveness, the ICT should be developed where it is applied – separately e.g. in each business sector, in healthcare, in leisure. However, such an approach leads to

(a) massive duplication of effort

(b) no learning from experience

(c) systems which cannot interoperate to allow benefit to be gained from emergent interdisciplinary and cross-disciplinary opportunities

The very pervasiveness provides massive opportunities for business and leisure, for quality of life improvements and wealth creation, through joined-up, interoperating systems at all scales.

2.2 Benefits of ICT

ICT (and therefore R&D to improve the ICT available) is beneficial in three major ways:

(a) It supports the personal information and interaction environment thus liberating the person from tedious chores and enriching their life and interaction with others in social or business contexts – it 'enlarges the personal contribution';

(b) it supports commercial activity: mineral extraction, process plant; environment, agriculture and fisheries, business, healthcare, scientific and technical research and development, transport and travel, leisure activities, culture, media… it 'magnifies the contribution' from each of these areas of human activity;

(c) in its own right computer science and information systems engineering provides a scientific discipline to pursue fundamental questions in natural philosophy and thus increase our understanding of the world, natural and man-made phenomena;

all aimed at wealth creation and improving the quality of life.

2.3 Structure of the Document

There are many ways to structure a discussion on Future and Emerging Technologies for ICT. We consider that applications (and their requirements) can drive the requirements of ICT, but that equally ICT developments can provide technologies that can change dramatically the capability of applications and the way to handle applications – indeed creating whole new business opportunities. Here we have chosen to consider three main aspects:

1. User and system centric components, together with connectedness, for architected application systems
2. electronic, storage, computing and communication components to support (1)
3. system development methods to construct components of (1) and applications on (1)

User centric components include those for interaction, intelligent support, and personalised services.

System centric components include those for ubiquitous, pervasive and ambient computing, and data analysis / reduction, information and knowledge processing, modelling and simulation and decision support leading to GRIDs technologies as defined in the NGG Reports (www.cordis.lu/ist/grids)

Platforms for information and communication services have traditionally been constituted by two types of platforms, i.e. platforms for information services and platforms for communication services. However, there are presently synergies to be efficiently developed by joint efforts from developers from these separated ideological camps to constitute a common integrated platform - expressed as middleware - for information and communication services. We see today that the user and system centric components are incorporating elements from the communication platforms at the same time as the communication service platforms incorporate information service elements. These communication service platforms are also applications that constitute a complex ICT system and the boundary between these service platforms and the user and system centric components are seemly vanishing. We foresee future open integrated platforms for information and communication services

Electronic, Storage, Computing and Communication components include those required for information processing, storage and transfer in a highly heterogeneous distributed environment.

System development methods will need to evolve to produce flexible self-managing systems composed dominantly of components, which interoperate reliably and in a dynamically self-managed way and that meets the requirements of the end-user/ application. Thus system development methods for applications as well as the components of the platform for information and communication services platforms are needed.

2.4 Scope of the Document

The analysis and recommendations in this document are intended to accelerate developments that follow visible technological trends. The proposed R&D is based on developing, proving and encouraging take-up of currently existing ideas, research prototypes, emerging or nascent technologies. We have made no efforts to develop more speculative scenarios except in a few areas where current approaches appear not to be producing results and where such futuristic research should be focussed to produce results. We provide a special short section covering this topic and the 20 year horizon. However, there are more than enough demanding research challenges within a limited scope of realistic predictions of what the next 15-20 years may bring.

We have also provided a section with a few example application area requirements to demonstrate how the approach proposed here will, indeed, provide a sound basis for development of the products, systems and services needed for Europe in the coming years.

Finally in two annexes we provide a brief list of the relevant qualifications of ERCIM and some scenarios used to stimulate the process of defining requirements and thus required lines of R&D.

3. USER AND SYSTEM CENTRIC COMPONENTS, TOGETHER WITH CONNECTEDNESS, FOR ARCHITECTED APPLICATION SYSTEMS

3.1 Introduction

A typical application system in the IT domain spans both user- and system-centric areas. However, since it is possible to utilise the same or a similar user-centric component to access many system-centric services and facilities (e.g. web browser / proforma access to banks, travel agencies, libraries, entertainment for information or control) and conversely it is possible to utilise several different user-centric systems to access one system-centric service or facility (e.g. telephone, mobile PDA, laptop or desktop computer, another intelligent computer system acting for an end-user) the two are considered separately. The key is, of course, the interfaces between these two areas of components and their characteristics of performance, security, and expressivity, representativity (related to semantics). This area will be dominated by technology based on metadata, agents and brokers.

Today the user and system centric components are using the services of some communication service platform. The communication service component offers services for communication session establishment, information transfer, and disconnection, but also will also offer functionality for service creation, deployment, execution and management. Functions like mobility handling and capability and resource handling are also functions to be covered by the communication service component.

3.2 User-Centric Components

The user centric view on ICT is concerned with ICT as a support discipline for achieving better support for satisfying human needs in other activity areas. The development of application-enabling technologies requires a multidisciplinary approach involving both domain-dependent and -independent aspects, and in particular: (i) disciplines related to the application of the technical infrastructure for building ICT-based systems in support of other activities (for example within organisations), and (ii) disciplines related to the development of ICT-applications independently from specific application domains, e.g., software engineering, artificial intelligence, digital libraries, human-computer interaction.

Application-enabling technologies deal with software systems that help users carry out tasks by providing relevant information, as well as providing computation and cooperation support. They generate, manipulate, analyse and store both structured and unstructured data, and they automate, control, monitor, and support the tasks at hand in other ways. The inter-dependence of software, humans and various artefacts is central to the user centric view on ICT. This complex mixture of technological and human resources is today crucial in many domains, and information systems integrate and apply more specialized technologies to meet individual or organizational needs. In such a context, it is particularly important that ICT technologies are developed in such a way as to cater for the needs and requirements of the widest possible range of users and uses.

The user centric view on ICT comprises several research areas, of which the following are of particular importance: Interaction, Intelligent support, Personalized services, Ubiquitous computing, and in the Systems Development area Networked communities, Emerging development processes. Additionally there is an horizontal requirement, concerning the quality and reliability of the resulting applications.

In each of these areas substantial research is needed. Each of them is discussed below.

3.2.1 Interaction – body and environment

During the next 20 years we envisage a rapid development of multi-modal interaction. This comprises substantial development in sensor technology, speech technology, displays, and audio/video communication. Additional modalities of increasing importance are speech, eye tracking, gesture tracking, and biometrics. Also for Virtual Reality, immersive displays audiospace technology and haptic devices are essential. Finally, a very important aspect is the presentation of information to non-experts, and the associated dialogue.

> The requirement is for R&D to develop the hardware sensors and effectors and associated software systems to detect human communication modalities and act intelligently upon them.

Sensor technology moves in the direction of large sensor networks of "smart dust", that is, very small sensors with much processing power and low energy requirements. Battery technology (or more generally, power source technology) is at the heart of this development, and in the development of communication networks where energy is only needed during short spurts of sensor communication activity. This opens a vast area of new applications, ranging from tracking of human condition and behaviour to tracking of the physical environment. At the same time, methods and systems are needed to protect privacy from unauthorized monitoring, especially when such monitoring is performed via "smart dust" that is so small as to go unnoticed.

> The requirement is for R&D to develop the hardware components down to nano-scale and also to develop appropriate efficient software for sensor control and communication

Many of the new input technologies have an inherent uncertainty (speech recognition errors, the interpretation of a medical sensor regarding a specific human condition, noise in transmission, etc.). This calls for context-based interaction (intelligent dialogue), where we allow delayed decisions using probabilities and models from the different modalities as well as dialogue histories and user preferences. To provide for user friendly interactions the sensor semantics must be coordinated with the semantics of the users. The systems must be able to detect and repair misunderstandings.

The requirement is for R&D to develop appropriate efficient software to manage the incompleteness and uncertainty of multimodal multimedia input and maintain intelligently the syntactic and semantic interpretation, modifying with additional information

3.2.2 Intelligent support

As more data becomes available for computers to retrieve and to analyse, there is an increasing market for systems that help users to solve problems. Applications are found in, e.g. learning and teaching, decision support, information retrieval, and many more. The big challenge ahead is how to get sense out of, and make appropriate use of, the vast amount of unstructured data that is increasing by the day.

Contemporary methods are of a syntactic nature, e.g., methods for information retrieval, information extraction, text mining, They make use of statistical approaches like normalized word frequencies and simple morphological information about words to analyze document content, and do not go into the semantic logical part. Semantic approaches are still not reliable enough, and the processing limitations prevent the methods from being employed in many information management and learning systems.

Over the next years semantic analysis methods will have to improve so that we can deal with the semantics of data rather than various statistical and morphological aspects. Methods and approaches from linguistics, statistics, and logic will have to be tightly integrated and used on-the-fly to help users find and interpret the information they need.

Structured domain knowledge will be necessary in order to interpret the data in the right context. The same data may mean different things to different users in different contexts. Different knowledge models must consequently be developed for different contexts, to reflect different users' experiences, interests, and preferences. The knowledge models must make sure that fragmented and ambiguous data become meaningful and relevant to the users. They should allow a user to specify only what he wants to achieve and let the system be responsible for figuring out how to do it by retrieving and combining the necessary data and services. Current research on "semantic web" issues should be strengthened and extended; the concept of layered semantics – to reduce the problem space and provide context aware structuring – is a promising line of R&D.

Semantically-enriched metadata and service descriptions that adhere to domain and context-specific knowledge and process models should be employed for information and service integration. Their combination will lead to the achievement of a further step in interoperability on the Internet, namely that of Semantic Web processes.

The requirement is for R&D to develop user interface software to interact intelligently with the end-user, using multimodal and multimedia techniques, to ascertain not only the request but also the intent of the request and to match this with communications to the system-centric components to ascertain contextual, historical and domain knowledge and utilise this in improving the interaction

3.2.3 Universal personalised services

Users of ICT systems leave traces that can be collected and used to assess their past behaviour. By collecting and analysing these traces automated systems (intelligent artificial agents) will help to find relevant information, carry out the appropriate transactions, and provide a computer-supported environment that is tailored to the individual user. E.g., information overload through digital television and an increasingly voluminous internet makes filtering and personalization of the available information important.

First-generation personalization systems are based on techniques like collaborative filtering, clustering, and classification algorithms. Contemporary decision support systems are capable of automating relatively simple decision making in limited domains, e.g., screening of medical images, classification of spam-mail.

Emerging technologies related to data warehousing and data mining will in the future be more integrated in everyday life. Agent technology is central to the area.

Protection of privacy becomes very important. Also for this reason there is an urgent need for intensified research in technology associated with universal personalised services.

The requirement is for R&D to provide appropriate intelligent agent technology, configurable by metadata, representing the end-user (or a participating intelligent system). The agent together with the metadata should recall behaviours in context, should be able to link to the system-centric environment for contextual, historical and domain knowledge and utilise this in improving the interaction

3.2.4 Ubiquitous Computing

The underlying goal is to move computers away from the user's attention, where they are applied subconsciously and embedded in many small and highly specialized devices within our continuously changing environment. This vision of ubiquitous computing has been around for a while, but it is only now starting to take shape. We envision that in the next decade ubiquitous applications will get out of research laboratories and into everyday life of people.

However we expect this adoption to be problematic for some years to come because of limitations in the technology, hardware and running costs, but also and foremost for the worries of the general public around these technologies. This will force the developers of these applications to take seriously into consideration the social and ethical implications that the developed technology brings along. Additionally, it will bring about the need of ensuring that new technologies are developed from the start in such a way as to proactively address the requirements of diverse target user groups with diverse individual characteristics and abilities. This implies the availability of appropriate development methodologies, techniques and tools.

Ubiquitous computing has implications for research fields like software architectures, computer-supported cooperative work, interaction design, and multimodality

The requirement is for R&D to provide better resilience and availability of systems for the end-user; the communication should be seamless despite mobility and the system should manage online/offline modes, intermittent connectivity, synchronisation of files and processes, backup and recovery. It must manage (together with the system-centric components) security and privacy issues including those concerned with identification, authentication and authorisation. It is particularly important that the new environment is proactively envisaged as being accessible and usable by the largest possible target user population.

3.2.5 *Games, Gaming and Edutainment*

Games provide a powerful mathematical framework with a well-developed theory as well as very general and versatile models of computation. In particular they capture in a natural way the aspect of interaction. Many synthesis and validation tasks, formula checking, or query evaluation problems can be formulated in game-theoretic terms. Accordingly, specifying a module amounts to formally describing a game, synthesizing a module amounts to computing a winning strategy and verifying a module against a specification amounts to checking that a strategy is winning. The growing expectation on the quality (as close as possible to reality) of games can be met only by interactive simulation based on physical models.

Games enhance user-system-user communication. With advanced graphics – including virtual and augmented reality - and formalised interaction structures they provide ideas for business applications. Virtual intelligent human-behaving characters can act roles. Games allow joint adventures in a virtual world, and these techniques can be used for education, training and situational analysis. With mobile user devices and merging mobile phone technology real and virtual worlds can interact under user control.

The requirement is for R&D on interactive simulation based on partial differential equations, particle methods, intelligent characters and more realistic avatars, graphics closer to reality and more intuitive user-interfaces.

3.2.6 The quality of the resulting applications

One key point in that the resulting applications need to be of enough quality, This raises the issue of software reliability and quality, probably the most important problem of software development nowadays. The recent history is full of software errors that cause considerable damage (for instance, software failure of the Ariane 501 space shuttle, the lost of the NASA Mars Climate Orbiter, the use of a non reliable software for a device used in treating cancer patients resulting that most of them died because of the overexposure to radiation in Panama, etc.) A recent article in *The Economist*[1] (19 June 2003) emphasized the growing interest in software tools that can analyse code as it is being written, and automate the testing and quality-assurance procedures: "To achieve predictable quality in software-making, just as in car making; the more you automate the process, the more reliable it is". "The Panama incident causes some industry experts to consider the possibility that more stringent regulation of software development is necessary" was reported by media, putting software quality and reliability (i.e. the fragility of the software infrastructure and the cost of bugs) as a fundamental societal problem. But it is also an economical problem, as reported several times (a 2002 study of America's National Institute of Standards and Technology (NIST) [2]estimated that software bugs cost to the American economy more than $60 billion a year or about 0.6% of gross domestic product).

With the increasing reliance of organisations on the ICT systems then there is the probability that system failures will lead to litigation against system suppliers. There is therefore a need for increased R&D into the safety aspects of systems in the widest sense, as well as the security of systems.

The requirement is for R&D on theories, methods, techniques and supporting platforms is needed to allow the future generation of software systems and services to be validated, measured and certified for quality attributes, including reliability, interoperability, reusability, safety, security and non-functional aspects, based on a given specification, requirements, or constraint, in order to keep the methods of quality assessment in pace with the growing complexity of software, and increase the predictability of software quality.

3.3 System Centric Components

3.3.1 Introduction

The next 10 years will see progressive device interconnection such that the systems environment will consist of millions of interconnected nodes. The nodes will range from RFID components through embedded components to mobile PDAs to laptop or desktop user computers (on the user-centric side) to servers with computation and / or data storage capabilities and on to supercomputers and massive data warehouses

[1] http://www.economist.com/displaystory.cfm?story_id=1841081

[2] *NIST Study: Software Bugs Take Bite Out of Nation's Economy, www.nist.gov/director/prog-ofc/report02-3.pdf*

(on the system-centric side). Novel and exotic devices will be developed and require network interconnection.

There is thus a massive requirement to develop component software that :

(a) by its very existence permits the rapid, inexpensive and reliable construction of applications by utilising these preconstructed components for all but the application-specific system components

(b) permits those applications to intercommunicate freely subject to security and privacy constraints

(c) provides security and privacy features that can be configured appropriately for the application

(d) provides self-tuning, self-managing and self-healing capabilities to ensure continuity for the applications

(e) provides adequate semantics interfaces more than merely syntactical ones.

The component software modules will provide functionalities as described below.

3.3.2 *Identification and authentication (of the user or system)*

Current password-based systems are failing; most passwords can be cracked by normal computer power within a few days at most. The addition of personal information challenges – especially if not obvious ones like 'email address', 'postal address' or 'mother's maiden name' improves security dramatically. However, the responses to these challenges could be deduced with intelligently connected systems and sufficient computer power.

Cryptology provides mathematical techniques. Research in this area is motivated by security threats, by the erosion of the computational difficulty of the mathematical cryptology base and by the requirements of new applications. It concerns strong foundations, and cryptographic implementations which appear in many different industrial contexts such as cell phones, pay TV, banking cards, and of course Internet protocols.

The requirement is for R&D into better cryptography, better digital certificate processes and (for human users) biometrics.

3.3.3 *Authorisation (to use resources)*

Once the user or process is identified and authenticated, their use of resources is controlled by permissions. Many different systems require permissions in different formats and styles. The requirement is for each organisation to maintain a single permissions resource, suitably protected and replicated, such that a change of circumstance for a user is recorded immediately in one place only and then cascaded as needed. Thus a person who, during absence, has to act for another should have changed permissions for the appropriate time period only. The permissions update process is tedious and error-prone; permissions should be deduced from other information automatically.

The requirement is for R&D into systems providing permissions (authorisations) automatically based on deduction from known and trusted facts

3.3.4 Resource scheduling and synchronisation including distribution of program code and / or data

Given a distributed environment of nodes with varying capabilities in computation, data / information / knowledge provision, detection of facts from the real world then there is a need to schedule optimally the request from the end-user utilising these resources. The scheduling has to take account both of the priority of a particular end-user request and its optimal scheduling among those of other end-users. Such scheduling may involve simple synchronisation of reserved time across multiple processing nodes or it may involve some preparatory work in shipping code to particular processing nodes or data to particular processing nodes.

The requirement is for R&D into systems providing intelligent scheduling utilising a vast amount of information in order to optimise and allowing for preparatory actions before the actual execution of the end-user request

3.3.5 Backup, recovery, checkpointing, restart

Reliability in a distributed system requires that data and software is backed up, can be recovered in the event of some malfunction and that long-running software codes can be restarted from checkpoints –all in order to optimise both response to the end-user and to optimise resources.

The requirement is for R&D into systems providing intelligent backup, recovery, checkpointing and restarting taking into account the user request and the condition of the distributed system

3.3.6 Archiving and curation

Digitally stored information and knowledge is overtaking rapidly that stored on paper. There is a need to archive it (preservation over time) and to curate it (so that it is comprehensible and useful to others than the originators over infinite time). The former requires a system for detecting old media and formats and translating to new (migration). The latter requires metadata that is both machine readable and machine understandable (i.e. treatable by logic) – in other words far beyond the current Dublin Core standard. Furthermore, restrictive metadata (rights, security, privacy, charges) is required.

> The requirement is for R&D into systems providing intelligent media migration and into advanced metadata systems (including multilingual domain ontologies) to allow intelligent autonomic operation of systems

3.3.7 Filestore access and update

Existing local and distributed filestore systems have been in use for many years. There is a need for better optimised systems – integrated with restrictive metadata controlling access, distribution and archiving – to meet the challenges of advanced multimedia formats and more complex access patterns. In particular, there is a need for efficient handling of replication, distribution of access requests, partitioning of datasets, management of RAID-based datasets and network file storage systems.

> The requirement is for R&D into systems providing intelligent, self-managing and self-optimising distributed filestore systems handling multiple media types

3.3.8 Database access, query, update, fusion

Since the hierarchic databases of the fifties, through to object-relational DBMSs there has been improvement in performance and ease-of-use. However, great challenges remain in the areas of homogeneous access to heterogeneous distributed databases (involving schema reconciliation, ontologies and semantic management), query over hyperlinked multimedia, proactive databases which on change of state trigger events and processes, logic databases which can assure self-consistency (particularly over incomplete and inconsistent data) and can interact intelligently with processes, transactions across large numbers of distributed nodes and concomitant recovery or compensation strategies.

> The requirement is for R&D into systems providing intelligent, self-managing and self-optimising systems handling multiple media types and with mechanisms for homogeneous access to (including fusion of responses from) heterogeneous distributed databases. There are particular requirements in ensuring consistency by transaction and recovery / compensation mechanisms.

3.3.9 Data cleaning, analysis (including statistics), visualisation (including virtual reality)

As business, control room and academic decision-making depends more on accurate information derived from data then the quality of the data (in terms of representing the real world accurately) becomes increasingly important. The subsequent analysis and visualisation depends on quality data for the techniques themselves to have integrity. The techniques for data analysis and visualisation need to be developed to deal with incomplete and inconsistent data. Improved price-performance for visualisation is required, probably requiring R&D into fundamentals.

The requirement is for R&D into systems providing intelligent, self-managing and self-optimising data analysis and visualisation systems front-ended by data quality assurance and repair systems. There is a particular need to improve the price-performance of visualisation systems, especially augmented and virtual reality where the wearable aspect is also a key issue.

3.3.10 Data mining, correlation, summarisation and extraction

Data mining has demonstrated both promise and utility. However it is somewhat of a 'brute force' approach. There is a need for R&D into guided data mining using knowledge from earlier work and collected user intuition, assisted by statistical techniques and aided by 'steering' based on visualisations of the data in n-space. The results (hypotheses or correlations) should be superimposed graphically on the visualised data for human validation or assurance allowing further interaction and steering. Analogous techniques for data summarisation (including intelligent summarisation of text and other media) are required to assist the end-user overcome information overload. Extraction of important concepts from data or information (including full text, audio and images) requires intelligent systems with knowledge.

Integration and mining of distributed and heterogeneous (D&H) information and data sources should be also highlighted. In particular in the era of post-genomics the need to relate both clinical and genomic data is raised as a major demand, and the realisation of respective advanced Biomedical Informatics solutions will support the raising individualised clinico-genomics environment. Data mining in such a context concerns the "classical" information as currently maintained in the health record as well as new tissue and molecule-based information as soon as this is being collected. The change from late stage diagnosis (expensive cure) towards early detection or even prediction of disease (cheap cure or intervention) will not only improve the quality of the health of the individual, but also likely at the same time reduce the overall costs of the healthcare systems. The use of genetic and proteomic data in addition to clinical symptoms for medical decision-making will contribute to the expected, continued shift towards evidence-based medicine. Similar advantages are envisaged in commerce, industrial processes, control room systems, environmental management systems, agricultural advisor systems and many more.

The requirement is for R&D into systems providing extraction and summarisation over distributed & heterogeneous data including full text and images or audio using intelligent knowledge-assisted techniques optimised through graphically represented interactive steering for the end-user request

3.3.11 Simulation and modelling including numerical analysis

Much of modern decision-making relies not just on information to date (historical) but also on 'what if' information generated synthetically (modelling). Furthermore, understanding of processes (physical, biological, medical, sociological, economic…) is assisted by simulation using computer systems – so-called 'in silico' experimentation - and role-plays. These two techniques can be used together or separately, but share many underlying technological components especially advanced mathematical algorithms (numerical analysis) and advanced algorithms representing real-world processes. Statistical data reduction techniques are also relevant.

The requirement is for R&D to provide algorithms that are dynamically optimisable to any current configuration of distributed computation environment and to improve the fidelity of the algorithmic representation of the real world. There is a need for R&D into self-modifying (genetic) algorithms.

The future of more world-representative simulation will need more complex algorithms able to cope with the increasing non linearity of the phenomena under scrutiny. Non linearity and, more dramatically, singularities will be the hot topics. This implies an increasing investment in the fields of optimization, differentiable and non, i.e. combinatorial and integer programming, and in the field of non linear partial differential equations. We will see also flourishing the sector of numerical non linear algebra, i.e. the solution of polynomial systems of equations, and related topics of commutative algebra.

In the context of simulation the concept of simulated reality is characterized by the combination of numerical simulation (virtual engineering), optimization, decision support (reverse engineering) and high-level visualization (e.g. virtual reality).

Simulated reality opens the market for complete novel software products, including planning and development tools for products and processes, tools for virtual material design and interactive tools for decision support with built-in online-optimization. Simulated reality spans computational engineering for simulation of production processes (from single aspects to process chain), rapid prototyping and product lifecycle management (PLM). Furthermore, a methodical integration of optimization (reverse engineering) into simulated reality leads to novel instruments for planning and designing production processes, for optimizing product features and for virtual material design. Tools for multiscale simulation especially address the integration of material and component simulation, because local features of the material can be integrated into the product design. Finally, molecular dynamics (MD) allows mapping at the atomistic scale e.g. allow simulation based optimization of diverse steps to be taken in the field of surface processing (coating, splicing etc.).

In the near future - by integration of modern hardware and software concepts (HPC, Grid computing, parallelisation, storage and manipulation of data) - we can get over the historical separation between technical simulation, optimization and high-level visualization.

The requirement is for R&D in computational engineering (multiphysics, simulator coupling, embedding technical-scientific simulation in system theoretical analysis), in reverse engineering for complex applications of simulations (online-optimization / offline-simulation, successive optimization in hierarchical models, multicriteriel optimization, tools for interactive decision support) and in multiscale simulation (scale coupling of scales in space and time, efficient numerics (hierarchical methods, multiscale integration), modular software-architecture, parallel computers, Grid computing)

3.3.12 Decision support including control room functions

Human operators find themselves responsible for increasingly complex systems, from share trading to air traffic control and from chemical production plant to power stations. There are similar decisions to be made in agriculture (when to spray pesticide and how much), in the environment (how to clean up an oil spill) and in everyday life (which car should I purchase). All these decision-making activities require the support of systems built using the components described previously working together harmoniously. However, in addition the human operator will almost certainly wish to understand the reasoning behind the proposed course of action recommended by the system (explanation) and to query individual steps of that inference. The human operator may well wish to know the predicted result of 'what if we do this'. The system needs to learn which strategies were correct and which were not or were less appropriate.

The requirement is for R&D to provide systems with user interfaces appropriate to supporting a human operator in decision-making including explanation and assurance and with learning capability to store empirical knowledge for comparison with inference. This implies the use of metadata and logic-based systems. There is an underlying linkage to data retrieval and data mining over heterogeneous distributed data.

3.3.13 Workflow

Many processes in business, industry or normal life have a natural sequence of steps, some of which are preconditions for others, some of which can be executed in parallel. The assembly of a workflow representing the end-user request and consisting of sub-processes (components) each with its own requirements for access to data, information, knowledge, computation or human input requires advanced technology. This technology has to construct and modify dynamically the workflow intelligently taking account of all known constraints, changes in availability of resources and priorities. Contemporary ERP systems tend to be quite inflexible, hardly adaptable at runtime, and they primarily support organisationally agreed processes. Contemporary workflow management systems have been focused to deal with exceptions, and have thus offered some support for adaptive processes. These types of systems, however, have not dealt with emergent processes, which seem to encompass an increasing part of organisational activity.

The requirement is for R&D to provide systems which construct workflows - utilising sub-process components - which are flexible and dynamically reconfigurable depending on changing constraints, resource availability, user requests

3.3.14 Semantic Interoperability

Software semantic interoperability is focused on the possibility to integrate software pieces developed by (possible) different teams, in different moments, and (possible) using different methods and platforms. In order to obtain a safe integration semantic information on composed software (more than current merely syntactical one) is needed.

One of the important problems when dealing with software components is how to ``ask'' a component about the services it offers. Current systems just offer the possibility to consult the name and the signature of such these services. The ``semantic'' behaviour of a component can be identified by the programmer (but not by the program itself) using the documentation provided by the component designer. This documentation is usually very informal and can not be consulted automatically by another component.

The requirement is for innovative theories, languages, methods, and tools for supporting the interactive and automatic management of software component specifications, including appropriate metadata standards for those specifications.

In parallel, there is the need for data, information and knowledge interoperability. As already hinted in several sections, a great problem concerns homogenous access to, and use of, distributed heterogeneous data, information and knowledge. Techniques for matching schemas and metadata, based on ontologies and logic programming are promising but there is a crucial need for languages with better-structured syntax and deeper, layered and structured semantics so that the complex real world may be represented accurately and its representation can be manipulated with confidence.

The requirement is for innovative theories, languages, methods, and tools for supporting access to, and management of, heterogeneous, distributed sources of data, information and knowledge. This requires appropriate metadata standards to allow appropriate syntactic and semantic representation which in turn permits semi-autonomic operation.

3.3.15 GRIDs

It should be noted that the European GRIDs concept encompasses all these areas whereas the original North American GRID concept was limited to metacomputing (linked supercomputers). Thus we have a definition of GRIDs:

GRIDs – the GRID concept of metacomputing (linked supercomputers) extended to include datastores, detectors and instrumentation all interconnected to provide a surface for information systems in the widest sense which interact intelligently with the mobile end-user in a pervasive environment to determine the requirement and then execute it utilising the computation, data, information, knowledge and data collection capabilities available.

GRIDs can be seen as a resource sharing (and therefore resource saving) initiative; it can be seen as provision of resources hitherto unimaginable; it can be seen as a universal surface for any application. GRIDs requires ICT R&D of the highest calibre to assure top-to-bottom integrity through the architecture stack and lateral interoperability between application domains to allow interdisciplinary interworking and a consistent user landscape whatever the task or tasks being performed. Since GRIDs is the technology that may coordinate and encompass the others and can act as the orchestrator of the various R&D strands:

> The requirement is for R&D to produce standards for the interfaces and a widely-agreed architectural stack.

GRIDs middleware (supporting the applications by the provision of commonly required services) and the GRIDs foundationware (providing the operating system level facilities required for GRIDs by interfacing to GRIDs middleware (above) and supplementing the capabilities of existing operating systems (below) require wide-ranging R&D.

> The requirement is for R&D in fundamentals of operating systems, security, trust, resource sharing and scheduling, self healing, self managing and self tuning systems and many other aspects of computer science and information systems engineering.

3.4 Communication service components

Many of the information services that today are handled by end-user devices may have to be shifted towards the network. Key drivers are lack of trust among users, and insufficient technical knowledge on the user's part. Candidates for new solutions are among others: service creation, service discovery, authentication, authorization, charging and payment, and auto-configuration. Service quality will be the key issue, comprising service platforms, and issues of dependability, security and safe engineering practices.

The requirement is for R&D to understand the parameters controlling optimal location of a service or task in varying network topologies and with multi-network configurations involving different network capabilities.

The future service platforms will have to provide for third party service developers and service providers. Platforms for information services and telecom services have been different in the past. Convergent platforms will integrate peer-to-peer technology with client-server technology. Quality-of-Service (QoS), including information security, and negotiated Service Level Agreements (SLA) will become key issues. Architectures and execution frameworks must allow services and applications to be incrementally developed and dynamically deployed while maintaining high levels of QoS with emphasis on availability and security. This also calls for the timely development of compositional engineering methods, and improved methods and tools for dynamic and quantitative verification and validation. Self configuring systems may hold the potential to reduce the perceived user interface complexity and establish adequate trust, and thereby lowering the end-user acceptance threshold for new services.

Dependability is a key issue and comprises the means to ensure that we may trust, i.e., depend on, the systems and their services, as well as the threats to them and their attributes. Dependability encompasses availability, reliability, safety and has issues in common with security. Dependability must be engineered into the systems from their conception. Dependability can not (with success) be added afterwards. When new successful technologies move from a toy phase, where dependability is a minor issue, to a phase where they penetrates society, then dependability becomes of extreme importance, cf. telephony, cars and planes as historical examples. This knowledge should guide the research into the new "ambient intelligent" vision.

The requirement is for R&D to develop self-managing, self-configuring, self-healing and self-optimising networks such that – across multiple networks with different characteristics of performance, reliability, security and cost the end-user sees a homogeneous optimal network path

4. ELECTRONIC, STORAGE, COMPUTING AND COMMUNICATION COMPONENTS

4.1 Introduction

Ever-increasing end-user expectations of ICT systems rely to a considerable degree on the underlying hardware components. The dramatic increases in use of mobile phones, PDAs and embedded systems are driven heavily by improved hardware both in price/performance and in power consumption and battery / energy source technology. The ubiquity of internet access and the widespread deployment of inexpensive sensors (including webcams and mobile phones themselves) makes it conceivable that everything, everywhere will be recorded by ICT systems.

4.2 Technology for Interconnected Computer Systems

There is an end-user / consumer expectation that Moore's law (power per unit price doubles every 18 months) will continue and that the power/price ratio will continue to increase at the same rate as historically. Industry predictions contradict this. The major problems concern development of materials used in chips (CMOS and beyond), circuit design down to the nanoscale level and associated problems concerning power consumption and heat dissipation.

> The requirement is for R&D on suitable semiconducting materials and their properties, at nanoscale.

> The requirement is for R&D on techniques of chip design which reduce power consumption and at least contain heat generation or provide appropriate heat dissipation techniques.

There is an associated latency problem between processors and memory with increasing use of complex caching and pipelining techniques being adopted. This latter trend is optimal for a particular chip but makes software portability (and therefore dynamically reconfigurable systems) difficult.

> The requirement is for R&D to improve latency tolerance and latency hiding while making their exploitation transparent to software.

Thus there is a high expectation that Moore's law will not continue to operate with current technology (indeed the curve is already flattening), and that new approaches are required. There are two main alternative paths:

(a) provide additional throughput by parallelisation / distribution techniques

(b) discover or develop fundamentally different models of computing

Under (a) the idea would be to partition any problem and distribute it in a close or widespread distributed set of resources, including processors, memory, storage devices and communications. The optimisation of a particular task over a particular set of resources has been researched for years in the high performance computing (supercomputing) field. The real problem is to develop intelligent systems that can take any task and configure it optimally for the particular set of resources available at the time (based on their performance, cost, security etc).

The requirement is for R&D to provide intelligent systems to reconfigure dynamically a given task to take maximum advantage of the available parallel, distributed resources of processors, memory, storage and communications

Interconnection Technology will play a pivotal role. The information society relies on an increasing number of compute engines and high-performance servers, and it relies on them being interconnected to each other. Servers are made as clusters of interconnected "blade" computers. In the field of closely-coupled parallelism, there is always a requirement for improved performance. As the demand for throughput keeps climbing, buses are replaced by switches and switching fabrics in interconnects of all ranges – from networks-on-chip to inter-chip and system-area, through local and wide area networks. Current switch and router architectures are immature, leaving room for significant advances. Commodity switches are predicted to emerge as a basic building block, pretty much like the microprocessor is a building block today. Commodity switches are likely to affect the WAN/LAN router market, too, in the same way that commodity PC's pushed supercomputers to evolve into clusters of workstations.

The requirement is for R&D on improved switch and interconnects architectures for intra-chip, inter-chip, and inter-system communication, so as to reduce latency and increase performance. Commodity switches are needed that will become the building block for networks of all ranges..

Under (b) there are already research teams working on bio-inspired computing and quantum computing. While many industry experts predict that these developments will not be adopted by mainstream IT suppliers, nonetheless they may yield techniques and technology that can be used in specialised applications ranging from security systems to self-adjusting computations and from self-guided information searches to novel technologies for pattern recognition in applications ranging from speech to video to DNA sequencing and matching.

The requirement is for speculative R&D on new models of computing particularly the bio-inspired and quantum techniques.

4.3 Technology for Storage Systems

Disk-based storage systems currently outperform Moore's law in that the volume of storage per unit cost doubles every 9 months. With increased use of multimedia, storage demands will continue to increase extremely rapidly and with developing socio-legal framework constraints there is a tendency to store everything – partly because it is economic to do so and partly because it is now possible. However, while capacity increases dramatically access speeds improvements are less impressive.

The requirement is for R&D on improved access speeds to storage systems

Current media are developing with increased storage density. However, for large systems storage requirements it is uneconomic at present to keep everything online on disk storage. Furthermore, the use of multiple copies of data on disk storage (through replication or RAID techniques) reduces the useful storage volume available. Nearline storage provided by cassette tape robots and CD jukeboxes will continue to be used but may suffer problems of electromechanical reliability. These systems also usually have relatively slow access and transfer rates and require available disk storage for use of the data by software. Commonly they just store files and provide the capability to restore those archived files to disk. There is a need for more intelligent archiving. Work at the San Diego Supercomputer Centre on a system with appropriate metadata to improve the relevance and recall of appropriate files indicates a way ahead.

The requirement is for R&D on improved access speeds and transfer speeds for nearline storage systems, and a need for improved metadata catalogues describing nearline archived files to improve relevance and recall. There is an associated need for intelligent staging (caching) of data from nearline storage to ensure timely availability and minimal transfer dependent on update state and last nearline copy

4.4 Technology for interface devices

Current trends towards ubiquitous computing and ambient intelligence imply a rapidly increasing market for novel interface techniques. A key to success is in the minimisation of the energy needs of the devices. A rule of thumb is that the energy needed for executing a piece of logic in software requires 1000 times more energy as is needed for hardware execution. Progress over a wide range of hardware techniques is needed, e.g., electronic, acoustic, photonic and micro/nano-electromechanical devices, circuits, components, sensors, actuators, systems and their interfacing with the rest of the world. Typically, transmitters and receivers are employed for such interfacing. The research will have to embrace methods for modelling, simulating, prototyping and manufacturing such devices and systems. The research should span a variety of applications, e.g. remote sensing and surveillance, localization and positioning, micromechanical operations, medical and automotive electronics.

> The requirement is for R&D to embrace methods for modelling, simulating, prototyping and manufacturing such devices and systems. The research should span a variety of applications, e.g. remote sensing and surveillance, localization and positioning, micromechanical operations, medical and automotive electronics.

Technological progress in the fields of nanotechnology, MEMS (and future NEMS), transmitters and receivers, are key drivers for the development of sensor technology and interface devices.

Transmitters are of the most energy-consuming parts in communication devices. The drive for lighter and smaller transmitters calls for improving properties like linearity, efficiency, bandwidth and power consumption. The power amplifier is the main source of distortion and energy consumption in a transmitter chain.

> The requirement is for increased R&D on energy sources, power amplifiers and transmitters – particularly to optimise power consumption.

4.5 Technology for Access and Core Networks

To supply broadband for all and access everywhere are two key challenges for the next 15-20 years. Those are two different goals. The combination of the two, to supply broadband access everywhere is not a realistic goal.

4.5.1 Access

The relative cost of wired, long distance bulk data transport will continue to be small, relative to other cost factors. Access, services and content will still be the cost drivers. Because of the low cost of wired data transport it is quite probable that broadband access will mainly be wired except for the last few meters; hence the connection may still be perceived as wireless by the end user.

The other goal of Access Everywhere focuses on making basic telecommunications services available in an (almost) ubiquitous and seamless fashion, e.g., voice phone, e-mail, SMS/MMS, location and navigation services, emergency calls, banking transactions and information search via the Internet. These may not include all types of broadband services. Even when taking into account all anticipated technological advances and price reductions, there will still exist technological limitations, for instance due to cost and available frequency spectrum. There are fundamental trade-offs between accessible capacity, distance, degree of mobility, and cost per bit, which imply that full broadband access in sparsely populated areas with little or no infrastructure will always be expensive. A variety of wireless solutions, from satellite links to fixed radio access, will be important for implementing Access Everywhere.

> The requirement is for increased R&D on mobile telecommunications component devices, particularly antennae and power circuits, to achieve better price/performance. As an element of this, an overall more efficient use of the radio medium.

To deal with mobile end-user-devices a seamless handover between access networks based on different technologies, potentially operated by different commercial entities are needed in addition to hand-over between different geographical areas supported by the same technology. The user may prefer to manage his/her access for cost, security and safety reasons.

> The requirement is for technology for seamless provision of access across technological and operational boundaries, especially with improved speed over current handover.

> The requirement is for R&D on intelligent end-user system agents controlled by metadata is required to provide the end-user with self-management and self-healing capabilities on the connected device. Similarly the agent technology should be able to choose the optimal network (price, performance, reliability, security) for the services required.

4.5.2 Provision of transport services

It is unlikely that there will be only one network. More likely there will be a multitude of networks. A person or device will typically be connected to several of these concurrently. It is not given that the various networks are operating under the same administrative domain with a common operating policy. This has implications for the technologies, types of networks, and business model. There will both be a general competition between telco-operators operating in the same geographic area, and a new set of operators offering services.

Other traffic has surpassed voice traffic. The technology for handling voice as just another data communication flow is maturing. However, the income of the large telecommunication corporations is still dominated by voice traffic revenue. Such an imbalance between traffic and revenue cannot continue for long. The resulting shift in revenue flow may change both consumer behaviour and the industry itself. The same network(s) will carry a traffic with highly differing value to the end-user, required security, dependability and real time requirements. Similarly, the potential income and cost of providing the various services for the operator will differ.

The requirement is for R&D on (optimal) combinations of business models and technological solution to encourage take-up of services using a mix of technologies, and co-operation between entities to provide transport services with managed QoS.

The fundamental problems of network realisation will remain as they are today. Radio frequencies resources will be scarce resources; efficient use of them will remain a problem. Improvements in the capacity on the different mediums, e.g. optical, will continue. The fundamental problems of resource allocation, robustness, reliability, service discovery, and security will remain. Current technology addresses these problems only to a limited degree and do not scale to the required dimensions. A solution suited for experimental networks will not be acceptable for networks that are a fundamental part of the society's infrastructure.

The requirement is for R&D on intelligent, autonomous network management, under a co-operation between end-user-agents and the management entities of competing and co-operating network operators.

Although the relative cost (€/bit) of "bulk bit transport" is low, it is not insignificant, and it triggers services with very high data transport requirements like (3D) high quality video. Optical technology provides low cost transmission. However, packet/burst based optical switching poses a number of unresolved issues. Electronic routing and switching technologies meet scaling challenges. Technological advances with respect to switching/routing may also have implications on the overall network architecture.

The requirement is for R&D on optical and electronic switching and routing principles and related architectures to meet the foreseen extreme capacity requirements in the future fixed network.

4.5.3 Ad-hoc and sensor networks

Ad-hoc network where devices configure themselves into networks, which may connect to the "permanent" communication infrastructure, has been a hot topic in resent years.

Traffic patterns will change towards machine to machine communication. The emergence of small and cheap sensors is a formidable challenge. A wide range of sensors will be integral parts of the future networks. The sensors not only gather data, they are themselves also part of the forwarding network and the collection and fusion of data. Sensors therefore form a new class of networks that are fundamentally different from today's network.

The requirement is for R&D to develop network topologies, configuration systems and management systems appropriate to large numbers of intermittently-connected sensors with varying capabilities as network store/forward devices, as processors and as configurable instruments for detection across a range of environments and domains.

5. SYSTEMS DEVELOPMENT TECHNOLOGY AND SUPPORTING THEORY

5.1 Introduction

At present the business of building ICT systems is perceived as slow, expensive and error-prone. There are many 'war stories' of large failed IT projects in government, defence, industry and commerce. Equally, but less frequently publicised, there are 'war stories' of such projects that are exemplary and delivered great business benefit. The latter involved experienced systems engineers, well-qualified in fundamental computer science know-how, who used selected and relevant ideas, techniques and prototype products generated from the R&D activities.

Nonetheless, there is a clear need for novel techniques to improve the effectiveness (relationship to the real requirement) and efficiency (cost to provide the system) of ICT development. This is particularly emphasised in the new world of massively distributed systems with self-managing, self-optimising and self-healing properties and in the increased user-expectations of ICT systems. A related aspect is the safety of such systems in their performance; an area where litigation threats should drive research and where research on techniques and standards for validation and certification should provide assurance.

The systems development dream has for a long time been 'executable specifications' i.e. the end-user specifies that which is required and the system is somehow built automatically to meet that requirement.

Taking the example scenarios of 'life in 2020' in Annex 2, it is worth examining what is obstructing realisation of this vision now. The obstacles appear to be:

(a) Economics: building systems appropriate to the application domain economically (in development and utilisation costs)

(b) Reliable and self-adjusting: building systems appropriate to the application domain that behave sufficiently reliably and safely and with sufficient self-adjusting flexibility and self-management to reduce utilisation costs

(c) Interoperation: building systems that can interconnect / intercommunicate to permit convenient e-business (whether the business is scientific research, travel, production process, logistics or supermarket)

(d) Task partitioning: building systems that have components spread geographically (home, office, handheld…) which partition the task appropriately for the application, environment and end-user role and location

(e) Socio-legal: building systems that are compatible with sociological and legal standards including security, privacy, freedom of information, ethics

(f) Easy-to-use: building systems where the effort threshold to use is sufficiently low and the immediate gratification of use is sufficiently high

The systems development process to construct the self-managing systems described in earlier sections will be based on certain core or base technologies in order to ensure interoperability, self-management, self-healing, self-configuring and other desirable properties including reliability and safety.

5.2 The New Environment of Systems Development

5.2.1 Networked communities and social implications

We may expect that pervasive usage of ICT will lead to different social communities, probably more ad-hoc and volatile. New forms of aggregation, bypassing traditional forms of governance and control, may promote new and richer forms of participation of individuals in community life. Everybody will be online most of the time. People are offline only when they choose so. Everybody become members of numerous networked communities. At the beginning of the millennium we witness the emergence of new forms of social aggregation, called by someone "social mobs", supported by widespread usage of mobile technologies.

If on one side the widespread use of ICT can lead to a more democratic society, on the other it also brings the risk to more control on citizens and lack of privacy. We expect that the next fifteen years will be largely characterized by the struggle between these opposites and the consequent definition and continuous revision of different policies.

A new set of requirements will emerge, e.g., it must become easier to influence on changing system functionality ("design-for-all"), and ICT systems must be better able to evolve when the desires of their users change, both automatically and in interaction with the user. From a process point of view both of these require more efficient methods for performing interdisciplinary work and more awareness in the design teams of the social, legal, and ethical implications of their systems.

The requirement is for R&D providing systems and techniques to ascertain and verify user requirements (what users want, not what they believe they want) as they change and evolve dynamically and flexibly. The specifications of requirement should be mappable to existing systems components (user-centric, system-centric including hardware, software and communications facilities) to allow ready fabrication of a tailored system to meet the requirement.

5.2.2 *Emerging development processes*

Evolvable ICT systems will have to be based on open standards, not just for infrastructure and middleware, but also for vertical components established by specific application segments. Component-based development will have to mature to the point that systems development implies less programming and more configuration of standard and generic components. Although we already see this trend in contemporary systems development, much more need to be done.

There are several unresolved basic issues in systems development: 1) Understanding the users, their environment and what they desire, and 2) controlling the process of gaining this understanding, and (3) translating this understanding into reliable ICT systems using the relevant technology. These are neither trivial nor well understood.

To be able to continuously evolve and migrate, systems will have to be self-contained and self-describing (no difference between system and documentation) using executable models of function and design. It will become very important that system developers become able to predict the consequences of the proposed systems modifications. New methods for model integration and migration are also needed.

The requirement is for R&D providing tools and techniques to build self-describing, self-modifying systems which can be integrated, for model driven systems development methods, and for methods for predicting consequences of design proposals

The preferred systems development technique will involve close end-user interaction with the developers, rapid prototype building and testing and user appreciation and training throughout the process, and a possibility for a collapse of the design-time/run-time dichotomy for important parts of the systems.

The requirement is for R&D providing tools and techniques for better supporting cooperative systems development with end-users and developers

5.2.3 *Challenges in education and research*

The anticipated development of ICT applications requires strong multi-disciplinary approaches combined with deep ICT specialist knowledge. There is a need to change the educational paradigm to reflect better that practitioners have to cooperate with people of different backgrounds while contributing with their particular competence.

The need for switching roles between viewing technology as a system and seeing it through the eyes of its users, requires a new generation of ICT professionals. Strong analytical and theoretical skills are needed to push the technology forward and take full advantage of the technological basis established. Education becomes a life-long pursuit.

A cooperative and open approach to the development and use of information systems poses new challenges to education and research. On the one hand, we need to adhere more to standards and be more efficient in applying reusable or generic software rather than developing it ourselves. On the other hand, projects will be more loosely organized, and we need to learn how to organize our work on an individual basis and get accustomed to more ad-hoc ways of cooperating with other people.

Finally, we need to encourage creativity. Whereas we in the past have improved the technology by imitating and automating manual approaches, we will soon be in a position to do things that we could never imagine before. We need to stimulate innovation and creativity both in education and research.

The requirement is for R&D to provide tools and techniques to support the cooperative information systems development process encompassing end-users and IT experts. Support will include knowledge-based systems with past experience and domain knowledge, catalogues of available system components, an information system of standards especially for interfaces, system fabrication facilities, appropriate diagramming and documenting tools

5.3 Base Technologies

5.3.1 metadata, agents, brokers

The flexibility and resilience required can be achieved by the use of metadata, agents and brokers. The separation of the description of the local systems environment (metadata) from the execution of tasks within that environment (agents) is crucial, is as the mediation between agents (themselves configured and controlled by metadata) by brokers.

The requirement is for R&D providing metadata (with appropriate syntax and semantics) that is machine-understandable (unlike current machine-readable metadata) that is processable by inference to allow intelligent self-management of systems. The metadata should describe the entity of interest, should proscribe its use (through restrictions on security, privacy, usage, charges) and should provide supportive domain information (multilingual ontologies, thesauri)

The requirement is for R&D providing agents which, configured by metadata, can act autonomously and intelligently to perform their tasks simultaneously cooperating (through brokers) with other agents to maintain an awareness of the overall state of the system and modifying their behaviour accordingly

The requirement is for R&D providing brokers which, mediating agents, can ensure end-to-end task requirements are met and can also 'police' the overall system assuring performance and security.

5.3.2 Information systems engineering, knowledge engineering, software engineering

Current techniques rely heavily on human interpretation of requirements and human knowledge of available techniques and components. The use of advanced intelligent systems to assist in requirements capture, to propose new possible business benefits and configurations based on new technological opportunities and to assist in building dynamic and flexible systems is a key area of R&D. This is likely to include novel languages with formal syntax and dynamically interpretable semantics to express end-user requirements and component to component messages and to describe the entities (objects) of interest to the application.

The requirement is for R&D to develop tools and techniques to build flexible, self-managing systems utilising components. This in turn is likely to require the development of novel expressive languages.

6. CHALLENGES AND NOVEL SOLUTIONS

6.1 Introduction

Although the thrust of this document is the minimal risk approach of accelerating the development of emerging appropriate systems, components and techniques by the use of fundamentally-based R&D, there are some requirements that appear not to be soluble using this approach. These concern problems that might be tractable given some breakthrough with a new architecture or method. In general they are problems that previously received less attention because they were considered intractable, even with the projected increases in compute power, data storage, intelligence of systems, network speeds and so on.

6.2 Challenges

The ICT research community has to cope with new large-scale infrastructures connecting together millions of computers. This change in dimension makes obsolete a lot of previous discovery or development, in particular in the field of distributed systems and applications as well as protocols. New visions are necessary in order to accommodate these new architectures. Autonomic computing is such a global approach which may lead to "invisible software" systems in the same way as we aim at providing "invisible hardware". Self-* properties (self-diagnosis, self-healing, self-optimisation,, …) are key features which must be thought of from the very initial design phases.

Designing and building large and complex software systems is a real challenge for computer scientists and software engineers. It is a key issue, as software must reach a much higher level of quality than at present. New concepts, methods and tools have to be invented in order to build the systems of the future, including such properties as mobility, ubiquity, globality of resources.

It is probable that the Von-Neumann model of computing is reaching its limits and that new models of computing have to be invented and experimented. Several avenues are presently investigated including: amorphous computing, bio-inspired computing, chemical computing, quantum computing. It is important for the future that those models attract sufficient research investments in order to become alternatives to classical models. As just one example, let us consider briefly quantum computing:

After the development of Shor's algorithm for getting the prime factorization of a given number in polynomial time in 1994, quantum computing is becoming a fast growing, new subject leading to a new computational paradigm. The principles of quantum computing based on the laws of quantum physics. In contrast to a classical system, the number of states of a quantum system has exponential size which is one of the sources of computational power. Up to now there is no realization of hardware for a quantum computer. Thus we have to simulate a quantum computer but for that we need a parallel computer with exponential memory size having exponential computing power. Currently there are two known algorithms: the Grover and the

Shor algorithm. The Grover algorithm is the search in a data base and produce a quadratic speed up in contrast to the classical algorithm. The Shor algorithm is the problem to find the prime factorization for a fixed number using in the RSA encryption algorithm. This algorithm has an exponential speed up in contrast to the classical algorithm. There is the hope that quantum algorithms can be solved any NP problem in polynomial time. Clearly for this class of problems R&D is required in such technologies.

6.3 Way Forward

It is quite possible that the 'extending by developing promising emerging technologies' approach recommended in this paper will solve the problems – for example a systems development method based on the production of systems using metadata, agents and brokers may well address the problem of millions of nodes, especially when linked with the next generation GRIDs ideas outlined above.

However, as indicated, it is likely that there remain some problems – or some intractable problems may be brought within the scope of tractability – by novel architectures and techniques.

7. EXAMPLE APPLICATION DOMAINS

7.1 Introduction

Three representative application domains - biomedical informatics, environmental modelling and automotive systems engineering – are considered because they provide clear examples of the merits of the approach proposed in this document, namely that ICT can be built that is applicable across a range of domains with specialisation only in the particular applications software to a particular domain. The following examples illustrate the extraction and abstraction of the requirements into the general, most cost-effective ICT R&D prospectus outlined in previous sections.

7.2 Biomedical Informatics

Recent landmark achievements in genomics and the increased importance of genetics in healthcare are changing the clinical landscape. The new knowledge coming out of life science projects is expected to (a) change the world as much as or more than the Internet, (b) transform the pharmaceutical and health care industries and (c) profoundly improve the practice of medicine. Classical epidemiological and clinical research on the one hand, and genomic research on the other, separately considered, are no longer enough for advancing in the so-called genomic medicine.

Genomic medicine integrates Molecular Medicine, which aims to explain life and disease in terms of the presence and regulation of molecular entities; and Personalised Medicine, which applies genotypic knowledge to identify predisposition to disease and develops therapies adapted to genotype of patients. The former is driven towards gaining knowledge about the disease, while the latter tries to know and clinically use individual genetic information. It should be stressed that individual genotypic information, essential to such approaches, must be the subject to stringent security.

The integration and exploitation of all the data and information generated at all levels by the disciplines of bioinformatics, medical informatics, medical imaging and clinical epidemiology requires a new synergetic approach that enables a two-way dialogue between them. The primary objectives of such dialogue are the development of new methods, technologies, tools and applications enabling the discovery and creation of novel diagnostic and therapeutic approaches.

As an example, Pharmacogenetics being the study of drugs in relationship to genetics, seeks to determine the genetic basis of variability in drug responses, and deals with such issues as how an individual's genetic characteristics influence the efficacy of drug treatments (efficacy pharmacogenetics) and the propensity for adverse drug reactions (safety pharmacogenetics).

Pharmacogenetics can be used to develop drugs that are targeted to individuals with specific genetic characteristics (polymorphisms or haplotype). In the scientific literature the importance of genetically mediated drug response has now been demonstrated in a significant number of studies and pharmacogenetic analysis or provision for pharmacogenetic analysis is now increasingly being incorporated into clinical trials. This will eventually herald a new era of what has become known as 'personalised medicine' whereby the drugs prescribed to individuals will depend on their personal genetic makeup, especially where there are significant cost or risk implications.

A new breed of systems and software tools are necessary to convert the enormous amount of data that geneticists and molecular biologists can obtain at their labs into information that physicians and other health care providers can use (and the converse – to codify and anonymise clinical phenotypic data for analysis by researchers).

To enable the vision of ***Individualised Healthcare*** of the future, significant R&D challenges need to be addressed including:

- Ontologies and ontology-based integration of heterogeneous biological and clinical databases;
- Data and text mining algorithms from the exponentially growing medical records, research literature and biological databases;
- Methods and tools for knowledge discovery, representation and visualization;
- Decision support systems for personalized health care;
- Informatics and computational methods in support of drug discovery, rational drug design, clinical trials and pharmacogenomics;
- Simulations and modelling of molecular interactions, metabolic pathways, cells and tissues;
- Molecular and metabolic imaging methodologies in medicine;
- Grid-based approaches for demanding molecular-biomedical computational applications;
- Security technologies for genetic data in e-Health Records.

All the application-based research challenges presented are discussed in the sections above, thus demonstrating the strong link between application requirements and generalised IT systems with application specialisation.

7.3 Environmental Modelling

Modelling of the biosphere with more and more biotic and abiotic components will be a great challenge. Climate research (space weather included) with models in different scales and different resolution will require new architectures with access to distributed resources. Branch oriented simulation systems will be the adequate software tools which can be flexibly adapted to the special structure and data of complex environmental systems.

There are special demands of environmental applications: Complexity (dimension, structure with abiotic and biotic subsystems), Scale (amount of data, distribution, heterogeneous), Modelling for different purposes (scenario analysis, emergency response, risk management, etc), Need for adaptability (coupling of models, parameter adjustment, etc.), Longevity of data, applicability to different purposes.

This leads to research themes : Parallel, distributed and Grid computing, Knowledge from data, decision support, Intelligent/adaptive user interfaces and visualisation, Standardisation of metadata and system interfaces, Workflows for automatic access to distributed resources, Genericity of information and simulation systems.

Clearly these research themes, supported by application requirements in this domain, are all proposed in the earlier sections of this document.

7.4 Automotive Systems Engineering: An example for safety-critical dependable embedded applications

Already now, in cars of the upper class, there are about 80 ECUs (Embedded Computing Units) and five bus systems, controlling comfort as well as safety critical functions. Most of the innovations (80-90%) in cars are ICT-driven, especially product individualization and differentiation are based on ICT. The cost of electronics and software in such a car will rise from 25% to more than 40%. On the other hand, according to reports at the 2003 informatics conference in Germany, 55% of car failures are caused by electronics and software, and the "X-by-wire" implementation plans had to be delayed by major players in the field by years. Diagnosis and maintenance in the field are again a challenge – because of complex electronic systems.

In the "Dependable Embedded Systems"-Roadmap of the European AMSD-project IST-37553 (http://www.am-sd.org) was developed a vision on future systems design, challenges and how to meet them, for critical applications in the areas: Automotive, Aerospace, Railways, Medical Devices/Systems, Industrial Automation and Process Control.

Automotive is clearly the "driver" in this area:

- cost-driven (clear goal: "from supply chain to design chain", to build consistent dependable systems from components of many different suppliers) and
- requiring mass-deployment capabilities for critical (sub-) systems composition ("aerospace safety at automotive cost"), and
- with a clear overall vision: "accident-free driving", and as
- ultimate goal: "Autonomous Driving" (at least on specialized highways and lanes)

Automotive is an exquisite example for a long term vision in the area of critical applications, because the potential for networking to build adaptive systems on several levels of connectivity, which requires the full set of dependability attributes to be fulfilled (safety, security, reliability, availability, maintainability) in a holistic manner, is already there, although the gaps are at the moment bridged by human interaction only:

a. on-board embedded systems, vehicle-bound connectivity
    a. advanced comfort functions (car body electronics, noise suppression, configurable cockpit, navigation, communication, information)
    b. Safety Enhancement
        i. Vehicle dynamics (ABS, ASR, ABC, ESP (Electronic stability program), AAS (Active additional steering), adaptive cruise control road tire friction control),
        ii. Advanced Warning and Control (pedestrian protection, crash avoidance, lane support, track control)
        iii. Driver Monitoring, predictive driver assistance, emergency call system
b. extending autonomous on-board functions with interactive and co-operative systems (European study Final report of the eSafety working group on Road Safety", EC-IST DG, Nov. 2002, http://www.europa.eu.int/information_society/programmes/esafety/index_en.htm)
    a. roadside embedded systems and interaction (e.g. intersection control, speed control, automatic advanced emergency call systems)
    b. Vehicle-to-vehicle communication (advanced adaptive cruise control, traffic throughput optimization)
c. Global connectivity: Integration into regional traffic navigation and control, satellite bound global connectivity

Cost-effectiveness and mass deployment of critical systems in combination with non-critical systems is a trend, where many other application areas will benefit, so that there is a clear need not only for application-specific ICT-technology, but for generic dependable ICT-technology (hardware, software, SoC (Systems-on-Chip), building blocks, communication) which fulfils the requirements of generic functional safety and security standards as well as of sector-specific ones (certification to create trust in these systems !).

There is a need for a generic dependable, holistic systems architecture and design to build safe and secure critical applications in many application areas from component to system-of-systems level in a cost-effective manner, based on safe and secure ICT technology, which allows component based system evaluation, composition and certification as prerequisite for trust in mass deployment of these systems.

These aspects of ICT are all covered in earlier sections of this document.

8. CONCLUSIONS

The key threads running through the document concern overcoming threshold barriers to exploitation and take-up through addressing the concerns of the general public, entrepreneurs and IT specialists. These include provision of more powerful and usable component hardware devices continuing Moore's Law. There are real concerns to be overcome in the areas of security, trust and privacy. Complex systems become incomprehensible and uncontrollable: provision of self-managing, self-healing and self-tuning systems which are dependable and safe is paramount. The need for improved systems development techniques includes component software reuse based on standard interfaces to reduce costs and increase interoperability. It also includes end-user specification capture and involvement through spiral development methods. The use of metadata, agents and brokers appears attractive. Some requirements will only become tractable by the development of novel paradigms and technologies to break free from the limitations of von-Neumann architectures and to reach beyond the apparent slowing of Moore's Law.

It is a truism that we shall have 'Martini Computing' (anytime, anyplace, anywhere). However, the provision of this requires excellent technology which is reliable and self-managing. It also requires intelligent user devices and equally intelligent system components providing an end-to-end service meeting end-user expectations.

Behind all these technologies are the disciplines of mathematics, computer science and information systems engineering dealing with incomplete and inconsistent data, expressive languages, syntax analysis and the intelligent processing of semantics utilising agents, metadata and domain ontologies. It is the continuous interaction between the basic or fundamental R&D in these areas, the generally-applicable system components allowing fast and inexpensive system deployment and the application-specific products and services that will – building on European particular expertise in this style of R&D - give Europe a leading edge in ICT.

The key desirables in an ever more complex world of powerful ICT capability are ease-of-use, interoperability, and trust / security / privacy. It is no surprise that these correspond largely with the 'semantic web' and 'web of trust' concepts of Tim Berners-Lee which we may claim are of European origin. As we move towards GRIDs with ambient intelligent computing as the pervasive architecture, the development of these concepts into implementable engineering artefacts with appropriate performance, reliability, flexibility and security should be the major target of European ICT R&D.

9. ANNEX1: WHY ERCIM?

1. During discussions at the ERCIM 15th Anniversary Meeting in Crete, May 2004, the EC officials present suggested that such a document should be written by ERCIM.

2. ERCIM represents much of the European R&D ICT community through national institutes or consortia in each country but reaching out to industry and academia; ERCIM institutes have ~ 12000 employees; our associates have at least a 10-fold multiplying effect.

3. ERCIM has a long history of innovative and successful R&D with successes registered by each institution at national level (with or without academic or commercial partners) and similarly at European level both individually in consortia with associates and within ERCIM coordinated and managed consortia. ERCIM institutions have > 100 spin-out companies.

4. Because of the rapid advances in ICT the time span between research and application of research results is shorter than in "normal" non-technological research. ERCIM members span the continuum from basic academic research in ICT, mathematics, natural- and social sciences, via technology development in ICT, to consulting on advanced ICT issues in industry and public administration. ERCIM is consequently well placed to voice an opinion as to how R&D activities in the various realms relate and interact to each other.

5. ERCIM has, in the past, produced strategic documents at the request of the EC.

10. ANNEX 2 : LIFE IN 2020

John is old. He's sitting by the window looking out into the garden. He ponders about how to handle those new arrivals in the neighbourhood, the "killer snails" that some guy brought in from a vacation in South East Asia, when he illegally brought back some live plants. Disgusting creatures they are, and poisonous too. And they thrive in the local summer climate. Better search for help in the worldwide digital library. Yep! Relief there is! Somebody in Port Arthur developed a snail sensor! Delivered as smart dust. Better order a handful, to come in a couple of days by air transport. Pretty expensive stuff, 300 Euro a kilogram, but 250 grams will do for the whole garden, dispersed wisely, and then the monitoring will be taken care of by the home information coordination equipment. In a few years some wise guy will probably equip the dust with bio-engineered snail-killer to be delivered with the smart dust, in a "clean-garden" package for old folks who have difficulties bending deep enough for efficient snail-killing. Disgusting creatures!

There are few sensible people around to talk to. Many of the interesting ones are dead, and many of the others are bedridden and mostly interested in their own illnesses and other trivialities. John's eyesight is not what it used to be. And the hearing! Lousy it has become! What a comfort it is to take on the headset so that one can hear properly, and engage in an intellectually stimulating conversation with CoRo - the Conversation Robot, about this and that, and nothing and everything. It is a treat to be able to discuss the birth of Rock'n roll with someone as knowledgeable as CoRo. Almost human-like, and CoRo does not interrupt when told to keep quiet, not like some other folks that we know. And CoRo's facial expression is irrelevant, because John can't see him clearly anyway.

John and Mary's children are worried about their father and mother. Old and fragile they are becoming. Dangerous when they walk the stairs to the basement, in particular when being alone. And the children lead lives of their own, some in other parts of the world, several hours of air-travel away. How easy life is now when the OldieHelper is around to keep track of how they are doing, where they are located, whether they are taking their medicine, whether they eat regularly, and get enough to drink. If one of them should fall and break a leg in the icy roadway leading down from the house, their ICT-filled clothes will soon initiate action by the OldieHelper, which in turn will communicate with the relatives (text reports, images, video with audio) and with the healthcare workers…

John's grandchildren like to walk in the mountains. However, the Norwegian mountains are dangerous; the weather changes fast; there are steep cliffs and deep gullies. There are glaciers which can behave unpredictably and have crevasses. GPS handheld locators have already been in use for many years; but it would be nice if John could follow the progress of his grandchildren across the Folgefonn glacier, with immediate video and audio as well as location which draw a path on the interactive map John has in one of the windows on his system. Of course, should some unfortunate incident should occur, the systems know immediately and the local air-rescue is scrambled to the exact location with already appropriate paramedics and their treatment plan computed. John does not have to worry but can still enjoy their progress.

Of course, John and Mary like to have a good dinner. The intelligent fridge-freezer keeps track of its stocks of food, their sell-by date and state of freshness (and re-orders automatically for home delivery the commonly-used items from the local supermarket). It can compute calories, vitamin quantity and can even suggest appropriate menus and recipes, which if accepted trigger further orders for home delivery. Both John and his wife like to cook, so they choose together the menu and one or other actually does the cooking. Of course, it could all be done with pre-packaged meals and the microwave, but they still enjoy the craftsmanship of real cooking yet are comforted that a fall-back position exists should they need it. Meantime, the wine-store inventory is suggesting that a particularly nice 1998 claret is now at optimal condition and would go nicely with the chosen Norwegian lamb. On this occasion they are dining just 'a deux' but when they entertain friends the kitchen system keeps track of what meals were served last to those friends and can interrogate the systems of the friends to discover what they have been eating recently or on the last visit by John and his wife; subject to privacy constraints it can also interrogate for meal preferences, allergies etc.

John's son works in the oil business. He has at work an integrated system so that the project planning system reminds him of milestones and deliverables, of meetings and travels – without him having to transfer manually from one system to the other. The integration is also intelligent, for example before a meeting with another organisation the system searches and provides appropriate profile information; similarly in an internal meeting with another department the appropriate information is collected and presented before the meeting. The work system interacts with the home system to allow appropriate scheduling if John's son is visiting the theatre or a concert, dining with friends or packing for a ski-trip.

It is worth examining what is obstructing realisation of this vision now. The obstacles appear to be:

(g) Economics: building systems appropriate to the application domain economically (in development and utilisation costs)

(h) Reliable and self-adjusting: building systems appropriate to the application domain that behave sufficiently reliably and with sufficient self-adjusting flexibility and self-management to reduce utilisation costs

(i) Interoperation: building systems that can interconnect / intercommunicate to permit convenient e-business (whether the business is scientific research, travel, production process, logistics or supermarket)

(j) Task partitioning: building systems that have components spread geographically (home, office, handheld…) which partition the task appropriately for the application, environment and end-user role and location

(k) Socio-legal: building systems that are compatible with sociological and legal standards including security, privacy, freedom of information, ethics

(l) Easy-to-use: building systems where the effort threshold to use is sufficiently low and the immediate gratification of use is sufficiently high