Technology in Society 30 (2008) 388–396

Contents lists available at ScienceDirect

Technology in Society

journal homepage: www.elsevier.com/locate/techsoc

Frontiers of research and future directions in information and communication technology

Angel G. Jordan *

Electrical and Computer Engineering and Robotics, Carnegie Mellon University, Wean Hall 4618, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA

ARTICLE INFO

Keywords:
ICT
Communications
Computers
Electronics
Software
Hardware
Microelectronics
Semiconductors
Blades
Computers

ABSTRACT

Information and communication technology (ICT), characterized by continual innovation and rapid technological change, is having a tremendous impact on society. Research and development in ICT are being conducted throughout the technology sector, and fundamental research is being carried out in many university departments of electronics, computer science, and engineering. This paper deals with frontiers of research and trends in selected areas of ICT, including computer hardware, microelectronics, and semiconductor devices and materials—areas that are leading the innovations in ICT. It briefly describes emerging technological developments, and concludes with advances in software engineering.

1. Introduction

The convergence of the fields of information technology (IT) and telecommunications has produced the term "information and communication technology" (ICT), with "communication" including the field of telecommunications. ICT research is conducted under a number of umbrellas and fields including computer science and engineering (computers, computation, communication, and information science), and electronics engineering (robotics and artificial intelligence), all related to ICT. Publications in these broad fields appear in numerous journals, including those produced by the Institute of Electrical and Electronics Engineers (IEEE) and the Association for Computer Machines (ACM).

The remainder of this article is a brief discussion of what I see as the current status of research and future directions in the areas of information and communications, which are driving the advances in ICT.

2. Advances in computer hardware

Research and development in computer hardware are leading to significant reductions in the size of components and subsystems. This is due not only to the ability to incorporate more transistors on chips, but also to the switch from conventional servers to *blade* servers in computer installations. For most IT departments, blades are a huge improvement over conventional rack-mounted units. This compact and slim computer (hence the name "blade") typically uses the same Intel or AMD processors and Windows or Linux operating systems as most other servers, but blades consume much less power and take up much less room. This will have an impact on the size of mainframe computers and supercomputers, as they will require less area.

* Tel.: +1 412 268 2590; fax: +1 412 268 2338.
E-mail address: ajordan@cs.cmu.edu

0160-791X/$ - see front matter © 2008 Published by Elsevier Ltd.
doi:10.1016/j.techsoc.2008.05.002

3. Microelectronics and semiconductor devices and materials

For the past four decades, the productivity and performance of microelectronics have advanced at exponential rates unmatched in technological history. The number of transistors per microchip has skyrocketed by a factor of about 100 million, while the cost of a chip has remained virtually constant. Microelectronics has become the principal driver of the modern information revolution, and the ubiquitous microchip has had a profound and pervasive impact on our daily lives.

Microelectronics is in rapid transition to nanoelectronics. For instance, the physical laws governing metal oxide semiconductor field effect transistors (MOSFET) behavior support the projection that device dimensions that began at 25 μm in 1960 can be scaled down from the 0.25 μm range to the 0.025 μm or 25 nm ranges after 2020.

Since microelectronics are integral to many current and future technology developments, research in design, modeling, fabrication, and testing of solid-state devices is ongoing in many university and industrial laboratories. This research includes semiconductor electronic devices, photonics and optical devices used in data storage systems, and the increasingly complex area of optical networking. Semiconductor research is also being conducted into the use of materials other than silicon, such as epitaxial germanium–silicon junctions and devices, nitride semiconductors, and optoelectronic materials and devices.

Another important area of research is electronic design automation (EDA), which has been active for over 25 years. However, advances in EDA alone cannot maintain the progress of integrated circuit (IC) designs. Frontier research is being conducted to develop silicon components created in tandem with design automation tools. These new technologies will facilitate implementation of a silicon system that can be manufactured. Work in this area includes analog and mixed-signal IC design; low-power architectures and circuits; designing for manufacturability; circuit-level modeling, optimization, and synthesis; verification and system design; computer architecture; and hardware–software co-design.

In the real world, advances in productivity and performance do not increase endlessly. Consequently, the paramount questions now facing microelectronics are: How much longer can we expect these advances to continue? What lies beyond the profound and pervasive impact that microelectronics have already delivered to society?

4. The transistor laser

The transistor has spawned a $300 billion per year semiconductor industry. The transistor makes possible a lifestyle that includes cellphones, PCs, digital cameras, MP3 players, medical imaging systems, set-top boxes, supercomputers, and the Internet. Now it promises to emit light, and with that comes the possibility of speedier broadband communications in both telecommunications networks and within and between chips. A team at the University of Illinois at Urbana-Champaign [1] has an extraordinary prototype transistor that can switch on and off more than 700 billion times per second, faster than any other transistor in the world. It is made from indium phosphide and indium–gallium–arsenide, the same type of compounds used in today's light-emitting and laser diodes.

This discovery will be followed by a more powerful kind of device, a transistor laser. It will put out both electrical signals and a laser beam that can be directly modulated to send optical signals at the rate of 10 billion bits per second. With some further modification, the transistor laser will eventually send 100 billion bits per second or more.

5. Silicon chips

The truth of Moore's Law—the number of transistors that can be placed on an integrated circuit is increasing exponentially, doubling approximately every 2 years—has been demonstrated relentlessly. Data rates have soared, Internet traffic continues to swell, and wired and wireless technology expands constantly to cover continents. Users now expect fast, free-flowing bandwidth whenever and wherever they connect with the world. Within the next decade, the circuitry now found in a rack of today's servers—which are able to churn through billions of bits of data per second and handle all the data processing needs of a small company—will soon fit neatly on a single silicon chip. However, as newer, faster microprocessors emerge, the copper connections that feed the processors in computers and servers will prove inadequate to handle such large amounts of data. At data rates approach 10 billion bits per second, microscopic imperfections in the copper or irregularities in a printed circuit board begin to weaken and distort the signals, even when traveling distances as short as 50 cm.

New board materials and new techniques will provide some performance gains, but also will cost more. Researchers are now replacing copper with optical fiber and electrons with photons. This process is called silicon photonics, and it will result in affordable optical communications. Manufacturers will be able to build optical components using the same semiconductor equipment and methods they now use for ordinary integrated circuits, thereby dramatically lowering the cost of photonics.

Meanwhile, performance gains will be significant: integrated onto a silicon chip, an optical transceiver could send and receive data at 10 billion or even 100 billion bits per second. This kind of bandwidth will, in turn, dramatically alter the ways computers are used. With optical interconnects in and around desktop computers and servers, it will be possible to download movies in seconds rather than hours and conduct ultra-fast searches through gigabytes of image, audio, or text data. Multiple simultaneous streams of video arriving on PCs will open up new applications in remote monitoring and surveillance, teleconferencing, and entertainment.

6. Microelectromechanical systems (MEMS)

As information devices increasingly leave fixed locations and appear in our pockets and palms, they are also coming closer to the physical world, creating new opportunities for perceiving and controlling our machines, structures, and environments. To exploit these opportunities, information systems need to sense and act as well as compute. To this end, research is being conducted that will lead to highly integrated microsystems that interact with the environment and push the boundaries of electrical–mechanical integration. Frontier research is being conducted on sensor and actuator systems derived from mechanical features measured in microns and components numbering from a few to millions. Examples include: integrated complementary metal oxide semiconductor (CMOS)–MEMS fabrication; integrated MEMS design and synthesis methodologies and tools; CMOS–MEMS device parameter extraction; MEMS-based probe storage; microfluidic systems synthesis; polymer-based gas microsensors; CMOS wireless sensing platforms; nanomechanical mixers and filters; tunable radio frequency (RF) MEMS devices and circuits; acoustic/ultrasonic MEMS arrays; and inertial sensor systems on a chip.

7. Data storage systems

Information technology includes the transmission, processing, and storage of information. Magnetic disk drives have evolved from using fifty 24 inch platters to store five megabytes in 1956, to storing over 1 terabyte on four 3.5 inch disks today. This represents a density increase of 10^8, from 2000 bits/inch^2 to over 200 Gbits/inch^2. This huge increase has resulted in disk drives also being used in consumer applications, such as personal video recorders, global positioning systems, and MP3 players, games, and cameras. It is expected that future changes in technology will enable the industry to increase density by several orders of magnitude while further enhancing performance.

Researchers are developing a greater understanding of what is required for future information storage systems. Research is being carried out that will advance information storage technology far beyond current methods of magnetic recording, optical data storage, probe-based systems, and holographic and solid-state memory. Efforts to advance data storage systems (DSS) have taken a systems approach that simultaneously focuses on materials, fabrication, devices, servo systems, signal processing, and coding. Nanotechnology and its applications are at the forefront of research in information storage.

Memory chips may one day rely on organic compounds to store information. Several research groups are reporting promising organic memory prototypes, from devices that function as dynamic random-access memories (DRAMs) to high-capacity, non-rewritable storage media similar to CD-ROMs. Organic memory will not make computers more powerful, but because they do not require clean-room fabrication and endless rounds of photolithography to produce, they might make memory chips and other storage devices easier to manufacture and therefore less expensive.

Researchers are not attempting to offer all-organic devices. Instead, to ease adoption of the technology, they plan to mesh their materials with existing silicon-processing technology. The semiconductor industry has shown considerable interest in organic memory because it is becoming prohibitively expensive to continue expanding the performance of memory chips based on current silicon technology. Companies such as Intel, AMD, IBM, and Philips now have in-house groups working on organic memory devices, or they have partnered with university groups and startups in the field.

8. Quantum computers

Frontier research in quantum computing promises the ability to surpass today's supercomputers by conducting encryption schemes or simulating quantum physics in far less time. The key to the power of quantum computers is the quantum bit, or *qubit*, which is not limited to representing 0 or 1. Qubits exist in fuzzy states that are both 0 and 1, and can be combined to represent many numbers at once. As a result, a quantum computer is similar to a massive parallel computer array whose power grows exponentially with each additional qubit.

The most promising technology for constructing an ultra-powerful quantum computer is the ion trap, a nest of electrodes that holds ions in midair. However, if a quantum computer is ever to be manufactured, the ion traps must be readily built from semiconductors the way computer chips are. Researchers have now built the first such ion-trap chips. Linking multiple chips may enable researchers to manipulate much larger numbers of ions and demonstrate rudimentary components of a quantum computer over the next few years.

9. Robotics

Some important advances in robotics are being made through research in what is called cognitive robotics, which draws inspiration from cognitive science to develop robot programming. Frontier research in robotics is proceeding in several new directions:

- sensor-resetting localization, where small robots, particularly in adversarial environments, are exposed to a variety of actions that cannot be fully modeled *a priori*;

- multi-robot, multi-hypothesis estimation where, in addition to its own individual sensors, a robot in a team can exchange information with its teammates;
- detecting environmental changes from sensory data, where robots constantly encounter changing environmental conditions, and sensory data are used to automatically identify and adapt to the environmental changes; and
- human–robot interaction and collaboration, which gives both humans and robots the ability to decide when to cede control to the other.

Significant advances are being made in understanding manipulation. This involves how a robot can rearrange the world around it by focusing on the physical interactions between the robot and objects being manipulated.

Humanoids is an area of robotics that is attracting leading researchers in the US, Europe, and Japan because it is expected to offer new perspectives on vision, learning, planning, and control theory. Humanoid robots have some human-like features that will lead to a new generation of consumer and assistive technology, especially human–robot interactions.

10. Computer graphics

Important advances in computer graphics are continuing, with strong emphasis on realism. Progress occurs as significant quantities of data are collected, whether by motion-capture equipment, special-purpose optical rigs, or pre-computed results from simulation. These data will be used to develop interface technologies for applications, such as synthesizing motion for humanoid robots and human characters; interactive haptic simulations (haptic technology refers to technology that interfaces with the user via the sense of touch by applying forces, vibrations, and/or motions to the user); perceptual metrics for human animation; texture synthesis and analysis; tracking human motion; and reconstructing information in foggy environments.

11. Computer programming

Research in this area focuses on programming that includes not just language design and implementation but also specification, verification, implementation, evaluation, and validation of programs. This requires formal languages with precise semantics that provide a direct link to the running code. Tools are needed to build software systems and evaluate their utility both analytically and empirically. This work covers a broad spectrum of research, ranging from abstract theories of programming concepts to large-scale implementations of those ideas in working software systems.

In the future, languages will have full formal definitions with checked proofs of their safety and security. State-of-the-art design and implementation of certifying compilers suggests that in the future all compilers will be certifying compilers. Certifying compilers generate object code that carries with it a machine-checkable proof of safety and security properties, eliminating or greatly reducing the need to rely on the correctness of a compiler to ensure the satisfactory behavior of a binary executable.

12. Software architecture

Research in software architecture seeks to improve existing computer systems and to develop new ones that will increase performance, improve reliability, or adapt to new computing environments. A major challenge for software architecture is in finding ways to exploit the largely unbounded hardware resources that can be fabricated, especially when planning for systems that will be built 10–20 years from now. Advances are being made in:

- performance improvements in application-specific memory, to close the ever-growing gap between memory system and processor performance;
- novel memory systems that extract streams of correlated data and move them between processors and memory to hide memory latency (the time between initiating a request for a byte or word in memory until it is retrieved);
- reliable, scalable server architecture that enables both cost/performance scalability and will tolerate failure within and across chips; and
- fast, accurate, flexible simulation and prototyping techniques that target fast, accurate and flexible computer system evaluation.

Radically new systems are being pursued in an effort to make a processing element that can be shrunk to a sub-millimeter scale. This could evolve into the ability to create programmable matter, consisting of millions of *catoms* (tiny robots that can morph into any shape), each with its own processor, wireless network interface, power source, actuators, and display.

13. Networking

In recent decades, the Internet has grown from a small experimental network that facilitated communication among researchers, to a global infrastructure that connects millions of people. Today, the Internet is the largest and the most complex computer system ever created, and its complexity is reflected in the design, engineering, and management of computer networks and protocols. Even though the Internet and its underlying protocols are engineered artifacts, it is increasingly difficult to understand how the Internet works and how various protocols interact. Complexity also contributes directly to the fragility of the Internet. Even minor disturbances or changes in one part of the Internet can create large-scale cascading effects that cause widespread failure and performance degradation in applications and critical network services.

These problems are exacerbated by the fact that there is a growing need to develop and deploy new capabilities, such as security, multicasting, and mobility. The challenge is to support these new capabilities without increasing the complexity of the overall system, and to understand how new solutions will perform and interact with other network components. Network researchers are engaged in fundamental research to understand, contain, and reduce network complexity. These efforts are proceeding along the lines of Internet measurement, measurement-driven protocol design, next-generation network and protocol architecture, and self-managing wireless networks.

Other ambitious projects seek answers to the following question: "Given the benefit of hindsight and our current understanding of network requirements and technologies, if we were not bound by existing design decisions and could design the network and protocols from first principles (i.e., a clean-slate design), what should be the design?" The "100×100 Project" [2] involves multiple institutions (Carnegie Mellon, UC-Berkeley, Stanford, Rice, and Fraser Research Internet2) and an interdisciplinary team of network researchers, economists, and security researchers. The project focuses on

- network control, an area that needs radical improvement;
- security, where the Internet today is plagued by a variety of malicious attacks such as e-mail viruses, worms;
- denial of service (DoS), an attack on the network that denies a user or organization the services of a resource; and distributed denial of service (DDoS);
- self-managing wireless networks that often suffer from serious *contention* (a condition in telecommunications where two or more stations attempt to transmit at the same time over a shared channel), poor performance, and security problems;
- protocol design**,** recognizing that the Internet Protocol (IP) service model has remained largely unchanged since it was invented 30 years ago.

14. Security

Computer security is concerned with protecting computers and networks from misuse and interference. Computer attacks include actions by an attacker to manipulate the functional behavior of the system (attacks on integrity); actions that induce the disclosure of information to unintended parties (attacks on confidentiality); and actions that infringe on the system's intended use (attacks on availability). Computer security also focuses on other types of misuse, such as e-mail that is commercial and unsolicited (spam), or intended to trick a user into disclosing information that can be used for identity theft (phishing).

Computer security worldwide has grown dramatically in the past 20 years, due in large part to the widespread use of e-commerce. A number of companies and academic institutions have introduced technologies to facilitate e-commerce, such as electronic payment systems and rigorous approaches to analyzing security protocols that are used to exchange private information over public networks like the Internet. The attacks of September 11, 2001, in the US, although not targeted at information technology, underlined the vulnerability of critical infrastructures of all sorts to determined adversaries. Within the US, this led to the creation of the Department of Homeland Security, which partnered with Carnegie Mellon's CERT/CC to create the United States Computer Emergency Readiness Team [3].

Research in cyber security includes: advances in cryptography, application of formal methods (e.g., model checking, theorem proving) to detect vulnerabilities in security protocols, detecting vulnerabilities in the compositions of systems, and modeling and implementing rich access control mechanisms. Innovative techniques are being developed to support the deployment of software that provides strong guarantees of behavior. One such development is proof-carrying code, i.e., software that allows consumers to easily validate certain software properties because the software producer has included a machine-checkable proof of these properties.

Important advances are also being made in intrusion detection. To respond to attacks launched anonymously and/or a long distance from secure data networks such as the Internet, some research efforts are exploring techniques to detect or filter various types of attack traffic. Others are exploring techniques for tracking the source of certain types of potent attacks, such as self-propagating worms.

15. Large-scale distributed systems

Two trends drive the growth of distributed systems. First, many applications utilize a distributed computing paradigm in which users access a distributed service infrastructure. Second, embedded stand-alone systems are becoming part of the global infrastructure. Distributed computing has been around since the early days of the ARPANET, but the scale of today's service infrastructure is unprecedented. The rapid deployment of sensors and sensor networks in many environments further blurs the line between embedded systems and open distributed systems.

Self-management is becoming more apparent. Self-management means that the system automatically optimizes one or more system features in response to changes in the load, system status, and environment. This trend has arisen due to costs and complexity. Research in this field borrows from organizational behavior in corporations, and the technologies in artificial intelligence and control systems. Other areas of research in large-scale distributed systems deal with information access, peer-to-peer service, and sensor networks.

16. Mobile computing

Mobile computing was born in the early 1990s with the advent of full-function laptop computers and wireless LANs. Although many basic principles of distributed system design apply, four key constraints of mobility forced the development of specialized techniques. These constraints are: (a) unpredictable variations in network quality; (b) lower trust and robustness of mobile elements; (c) limits imposed by weight and size; and (d) concern for battery-power consumption.

Mobile computing is an active and evolving field that includes:

- mobile networking, including mobile IP, ad hoc protocols, and techniques for improving performance in wireless networks;
- mobile information access, including disconnected operation, bandwidth-adaptive file access, and selective control of data consistency;
- support for adaptive applications, including transcoding (digital-to-digital conversion from one codec to another) by proxies and adaptive resource management;
- system-level energy-saving techniques, such as energy-aware adaptation, variable-speed processor scheduling, energy-fair ad hoc networking, and energy-sensitive task and memory management; and
- location sensitivity, including location sensing and location-aware system behavior.

Mobile computing shares many research themes with pervasive computing. The term "pervasive computing," also known as "ubiquitous computing," was introduced in a seminal 1991 paper entitled "The Computer for the 21st Century" by Mark Weiser of Xerox PARC [4]. He observed that "The most profound technologies are those that disappear. They weave themselves into the fabric of everyday life until they are indistinguishable from it." The key to pervasive computing is the creation of environments that are saturated with computing and communication yet are gracefully integrated with human users. When articulated in the early 1990s, this was a vision too far ahead of its time since the hardware technology needed to achieve it simply did not exist. It is only now, 16 years later, that the computing and wireless communications technologies needed for its realization are becoming available.

17. Artificial intelligence

With the explosion of information on the Internet, in data warehouses, and from physical sensory artifacts, artificial intelligence (AI) faces large-scale problems. Pioneering contributions from AI are making possible innovations in intelligent digital libraries, data representations and algorithms, market clearing technologies, and complex task-embedded robotic applications. Advances in AI contribute to multi-agent systems, and to progress toward building personalized intelligent computer assistants that can learn with experience.

The broad field of AI also includes vision and speech. Computer vision is about acquiring and interpreting the visual world. Today, cameras are everywhere, the number of images and videos is overwhelming, and automatic visual information processing has never been more important. Although vision systems have enjoyed great success in controlled and structured environments, several challenges must be overcome for these systems to be successful in unstructured, real-world situations.

Fundamental contributions also are being made in cartography and photo interpretation, such as automated analysis and interpretation of aerial or satellite imagery, and applications of a larger Image Understanding Program continue to grow. Fundamental contributions are being made in digital mapping using knowledge-intensive techniques to conduct detailed analyses of remotely sensed imagery. Research ranges from low-level vision (e.g., stereo matching and scene registration), to systems for cartographic feature analysis (e.g., buildings, roads, airports), to information fusion using multiple cooperative methods, as well as multi-spectral imagery and large-scale databases for advanced distributed simulation. All these areas are now part of frontier research.

Innovative research is being undertaken in the detection and recognition of facial expression analysis, gaze tracking, marker-less human tracking, and daily human activity recognition. The convergence of computer vision and graphics has offered new opportunities for frontier research on image-based rendering, and physics-based and data-driven modeling and animation.

Research at the intersection between AI and cognitive science is called *computational neuroscience*. This involves research on computational models of human cognition. Efforts are underway that go beyond applying computational techniques to biology, and now focus on the study of the computational principles of the brain itself. Researchers in AI are in close contact with researchers in neuroscience in their quest for a better understanding of the computational principles of neural computation and natural intelligence.

18. Machine learning

The broad goal of machine learning is to automate the steps in the decision-making process used by humans to carry out many tasks, so that computer-automated predictions can make a task more efficient, accurate, and/or cost-effective than it would be using only human decision making. There are three challenges with machine learning:

1. *generalization*, because as the dimensions of the data increase, an exponentially increasing amount of data is required to make good predictions;
2. *computation*, because as data become more abundant and high-dimensional, they become increasingly difficult to manipulate, identify matching patterns, and otherwise process in order to construct and apply models efficiently; and
3. *representation*, because data and other information relevant to a task come from many different sources, including expert human knowledge.

An unprecedented investment is being made in the collection of data, with enormous archives being formed. Financial transactions and computer traces are being logged, biological data are being collected using increasingly fast machines to scan genomes, hyper-spectral satellite imagery is being stored on a massive scale, and web documents are appearing at an explosive rate. Developing effective ways to extract useful information from these data stores is the challenge, with the goal of providing a greater return on the investment in data. This goal drives much of machine learning research.

One research trend in machine learning addresses learning in games, where there are multiple learners with different interests. Other areas include:

- on-line learning, with connections between problems in machine learning, on-line algorithms, and optimization;
- astrostatistics, where the explosion of sky survey data has made non-parametric statistics and machine-learning methods of great interest; and
- intrusion detection, which involves monitoring networks to discover unauthorized use or other suspect behavior.

Considerable research has been undertaken recently on automatically mapping brain images onto cognitive models of the thought process. Research in machine learning intersects with computational biology, and with computer graphics and computer vision, using data-driven techniques for problems that are difficult to model parametrically. New research is applying machine learning methods to wireless sensor networks to develop efficient distributed algorithms for inference, learning, and control. A notable trend in machine-learning research is the increasing use of statistical methodology.

19. Theoretical computer science

Other areas of computer science having an impact in ICT are scientific computing and computational molecular biology.

Scientific computing tackles problems associated with activities such as: (a) building numerical models, (b) building computer models, and (c) querying and analyzing experimental and synthetic data. By nature, scientific computing is broad and multi-disciplinary, drawing on physics, applied mathematics, algorithms, computational geometry, parallel and distributed computing, database systems, signal processing, graphics, and visualization. Some examples of advanced research in scientific computing are:

- new methods for building special, data-driven, numerical models that quickly and accurately simulate complex moving and colliding objects;
- developing software tools that simplify the construction of large-scale parallel finite element simulations; and
- new techniques for representing, indexing, and querying meshes and data that are stored in databases.

Scientific computing is often interdisciplinary. For example, earth scientists and civil engineers advancing pioneering techniques for modeling ground motion during earthquakes; mathematicians, engineers and physicians working jointly on methods for modeling blood flow; astronomers and astrophysicists working on database techniques for the World-Wide Telescope; engineers working on the interactive simulation parts of an aircraft.

Computational molecular biology is a flourishing area of research carried out in collaborations between computational and biological scientists. Areas of focus include applying machine learning and data-mining techniques to large biological knowledge bases, biological modeling, computation of high-throughput laboratory methods for large-scale, systematic studies of protein structure and function, analysis of biological image data, and computational genomics.

20. Human–computer interaction

Human–computer interaction (HCI) is a broad and dynamic discipline that combines scientific and engineering knowledge derived from computing with that of human and social sciences. Many of these research projects involve researchers from several disciplines, including computer science, interaction design, and cognitive science.

Research on HCI is proceeding in the following areas:

- learning technologies;
- tools and technologies to enhance the work of many researchers and extend it for solutions to problems within a domain;
- software tools and toolkits for interactive software development and intelligent tutoring systems;
- producing a psychologically valid computational model of users' behavior;
- evaluating software architectures;
- creating information visualizations;
- developing technology-based assistive technologies that support social, physical and intellectual tasks;
- people and robots; and
- arts and entertainment.

21. Software engineering

Software engineering is the branch of computer science and engineering that develops practical, cost-effective solutions to computing and information-processing problems by applying scientific knowledge and developing software systems that serve mankind. The distinctive character of software raises special issues about its engineering. Software is design intensive, symbolic, abstract, and constrained by intellectual complexity more than by physical laws. Software engineering rests on a body of core computer science concepts related to data structures, algorithms, programming languages and their semantics, analysis, computability, and computational models. It also rests on a body of engineering knowledge related to engineering processes, tradeoffs, costs, conventionalization and standards, quality and assurance.

These are complemented by the social and economic context of engineering, which includes the process of creating and evolving artifacts, as well as issues related to policy, markets, usability, and socioeconomic impacts, which provide a basis for shaping the engineered artifacts to be fit for their intended use. Software engineering results may be established qualitatively through empirical study of software systems, through empirical study of the software development process, or through formal analysis.

Important advances in software engineering are being made along two lines: (1) developing improved models and techniques for the design and implementation of large complex software systems that must satisfy rich, often under-specified, requirements in their operating environments, and (2) developing improved models and techniques for formal analysis of software components to assure or determine specific, formalizable properties. Software engineering research overlaps somewhat with research described elsewhere in this paper, namely software architecture, security, and programming.

Until recently, architectural design has been largely based on ad hoc choice, informal experience, and local expertise. Significant advances now focus on ways to make this knowledge precise, codified, and available to engineers as a matter of routine engineering. This has resulted in the development of several architecture description languages, and tools to support the analysis of architectural designs and generation of systems from architectural blueprints. Today self-monitoring and self-adaptation of systems (if they exist at all) are handled in an ad hoc manner and embedded in the application code itself. Systems of the future must take more responsibility for their own health and welfare. Important advances are moving toward a new approach in which monitoring and adaptation mechanisms are engineered outside the application. Application of economic theories is also being used to help aid in reconfiguration decisions in situations where resources are constrained.

Traditional engineering design discipline calls for designs to be evaluated long before they are implemented and for designs to be evaluated in terms of their overall value to the client, including costs and contextual considerations as well as capability. In what is called value-based design, the focus is on bringing these aspects of engineering practice to software engineering by developing a framework for evaluating early-stage designs that predict the value a user will see in an implementation. Research in value-based design analysis draws on techniques from economics as well as computer science. Techniques are being developed that predict the properties of a finished software system before the expense of development is incurred.

Other research being undertaken in the software engineering field are: dynamic systems, where software engineers focus on how to specify and predict the performance, dependability, and interoperability of software-intensive systems; product-line systems to enable widespread product-line practice through architecture-based development; and software engineering process management to allow software-dependent organizations to use proven, process-focused methods for improving product costs, schedules, and quality.

22. Conclusion

This paper has briefly discussed a number of areas within the ICT field that are having and will continue to have a profound impact on society. Computer programming, software architecture, and theoretical computer science provide some of the major underpinnings of ICT. Computer hardware, microelectronics, semiconductor devices and materials, and data storage systems, are integral parts of the broad field of ICT. Robotics and computer graphics, along with the transistor laser, future silicon chips, microelectromechanical systems, and quantum computers represent additional stages of the ICT revolution. Networking, security, large-scale distributed systems, mobile and pervasive computing, and software engineering represent well-established areas within ICT that are in constant evolution. Artificial intelligence, machine learning, and human–computer interactions are enablers of current and future ICT. Missing in this discussion is nanotechnology, which is emerging as a full-fledge technology that is attracting the attention of multiple research efforts, and is already having a major impact in ICT.

The contents of this paper reflect the views of the author, who is to a great extent inspired and influenced by his colleagues at Carnegie Mellon, an institution recognized worldwide for its strength in information and communication technology.

References

[1] PhysOrg website. Retrieved from ⟨http://www.physorg.com/news10584.html⟩.
[2] For further information on the "100 × 100" project, see ⟨http://www.100 × 100network.org/⟩.
[3] ⟨http://www.us-cert.gov⟩.
[4] Scientific American. The computer for the twenty-first century 1991:94–110 Available at ⟨http://www.ubiq.com/hypertext/weiser/SciAmDraft3.html⟩.

Angel Jordan is University Professor Emeritus of Electrical & Computer Engineering and Robotics, and Provost Emeritus at Carnegie Mellon where he was Provost from 1983 to 1991. Thereafter he was the Keithley University Professor of Electrical & Computer Engineering and Robotics until 2000. While Provost, he was instrumental in founding the Software Engineering Institute and forming Carnegie Mellon's School of Computer Science. Prior to becoming Provost in 1983, he was Dean of the Carnegie Institute of Technology (Carnegie Mellon's Engineering School) from 1979 to 1983 and Professor of Electrical and Computer Engineering. While Dean, he was instrumental in forming Carnegie Mellon's Robotics Institute. He was Head of the Department of Electrical and Computer Engineering from 1969 to 1979 and the Whitaker Professor of Electrical Engineering from 1972 to 1980. Dr. Jordan was successively an Assistant Professor, Associate Professor, or Professor of Electrical Engineering in 1959–1965. He earned his M.S. and Ph.D. from Carnegie Mellon in 1959 and the degree *Licenciado en Ciencias Físicas* from the Universidad de Zaragoza, Spain. He is a Member of the National Academy of Engineering, USA; Corresponding Member of the National Academy of Engineering, Spain; Fellow of the Institute of Electrical and Electronics Engineers; and Fellow of the American Association for the Advancement of Science. He is *Doctor Honoris Causa* from Universidad Politécnica de Madrid, Spain, the Universidad Pública de Navarra, Spain, and Universidad Carlos III de Madrid.