Tony Hoare

From: Robin Milner [Robin.Milner@cl.cam.ac.uk]
Sent: 11 May 2003 15:59
To: Tony Hoare

Tony

In trying to assess the GCs I wrote this series of sketches, one summarising each GC. They probably shouldn't go in our report, but we might circulate them to the PC as an aide memoir, when we ask them to agree to our draft report. If you think they serve purpose then don't hesitate to modify them.

Robin

---

1 In Vivo <=> In Silico

The proposal is to devise and realise fully detailed, accurate and predictive models of some of the most studied life forms in biology, such as the Nematode worm. Phenomena to be modelled include:
development from single cell to adult; cell function and interaction; interaction with environment.

Outcomes would include: showpiece demonstrators (for specific life forms); exprpriments evaluating their expressive and predictive power; conceptual frameworks traded between biology and computer science, influencing further research in both.

If successful, an expected paradigm shift is in the perceived value of computational ideas in life-science modelling.

Many subtasks are proposed, such as: development; object modelling; exemplars; cell virtual machine; cell interaction protocols; concurrent development control. Three possible five-year phases are:
(1) models of exemplars beginning to yield predictions; (2) models with partially accurate sensory responses; (3) the first complete models.

Many laboratories worldwide in computer modelling for biology are cited, and the proposal draws substantially on some eight other GC submissions.

---

2 Science for Global Ubiquitous Computing

Treating the Global Ubiquitous Computer (GUC) as both engineered artifact and natural phenomenon, the Challenge is to develop concepts, calculi, theories and automated tools that are deployed consistently both in the analysis and in the construction of the GUC at each level of abstraction.

This is an ultimate goal, expecting approximate achievement. Success will be measured by the extent to which the generic systems of GC4, and the Exemplar systems mentioned there, are expressed and analysed in terms of the concepts and calculi developed here.

The paradigm shift that is sought is that, as in other engineering disciplines, theories should inform the creation of systems and not merely be retrofitted to analyse them.

Several currently researched concepts and methods will be deployed, including (calculi for): Space and mobility; Security and Privacy; Boundaries, resources and trust; Distibuted data; Game semantics; Hybrid systems; Stochastics; Model-checking. In each of these directions, relevant advances within five years are predicted. New concepts, e.g. for reflectivity, need to be developed.

Three levels are defined for shorter-term projects at the start: (1) experimental applications (Exemplars); (2) experimental generic systems; (3) exploring hierarchy of theories (e.g. the implementation of a logic for trust in terms of more operational concepts). Examples of such projects are suggested.

It is suggested that the first stage in addressing the challenge is to form a GUC network, to define and mount initial projects and to cohere all work on the Challenge as it proceeds.

---

3 Memories for life

The Challenge is to manage individuals' digital memories --all the information in many forms that they keep about themselves-- for the benefit of human life and for a lifetime. It entails extracting useful knowledge from this record, and presenting it effectively to different kinds of user. It extends to using the extracted knowledge to build more intelligent tools, for example to assist medical diagnosis, or to model a person's linguistic ability and thereby customise web pages for him or her.

The challenge will be met when the majority of people can efficiently manage their information stream, and when all of us benefit from our digital memories. A paradigm shift from current thinking is from `computer as personal aid' to `memory as personal aid'.

The Challenge will be addressed by Exemplar applications. All Exemplars will share access to individuals' information streams. Each will be given a time span. For example:

(1) Five years: Search for images or audio by example (not text); Simplified web-pages adjusted to individual linguistic competence.

(2) Ten years: Electronic GP, a model of individual life-style and activities to assist in health advice; Smart electronic `paper' allowing a person's scribblings to be added to their digital memories.

(3) Fifteen years: Analyse subsequences of memories and re-present them as stories; Intelligent maths tutor, using a model of a child's current math knowledge based on attempts hitherto.

Several themes will link the Exemplars: The persistent user models inherent the memories; sensory interaction between user and computer, adjusted to users' skills; extraction of deep structure from memories; adaptation of memory-representations to allow evolving tasks.

---

4 Scalable Ubiquitous Computing Systems

The overall Challenge is to create the novel architectural structures required for scalable ubiquitous computing systems. Many unsolved issues are mentioned: context awareness; trust, security and privacy; seamless communication; low-powered devices; self configuration; information overload; information provenance; support tools for design, analysis, programming; information provenance; human factors; social issues; business model for creating ubiquitous infrastructure.

Central to the Challenge is to identify a collection of engineering design principles and techniques that apply at every scale. These will be tested against several system desiderata, e.g. means of negotiation for resources, self-stabilisation and suitable annotation of data. New programming models are required, e.g. to manage faults at every level of structure.

These generic system attributes will be tested in the field in Exemplar applications such as: Always-on health care; Gathering evidence, privacy and surveillance: Zero road fatalities.

Possible sub-challenges would focus on sensor networks, zero-power systems and structures to manage performance problems like failure and latency. Three five year phases may focus on (1) privacy and programming models; (2) scale management; (3) single systems architecture.

A significant shift in engineering practice will be that failures are not eliminated, but reported, attributed and managed at every level.

Several interests are complementary to those of GC2, which should develop theories for the concepts (e.g. trust) that underlie the engineering principles.

---

5 Architecture of Brain and Mind

The Challenge is to create a computational architecture of brain and mind inspired by both neuroscience and cognitive function, that describes how neuronal activity corresponds to cognitive abilities (e.g. adaptability), contributes to control and elimination of human mental disorders, and leads to artefacts incorporating some human cognitive functions.

The aim is to abstract and formalise principles of operation, rather than directly mimic chemical and physical mechanisms.

The essential paradigm shift will be the new-found experimentally juistified explanation of how brain gives rise to mind; previous explanations of mind and brain are disjoint. Consequences of the explanation are likely to be new computing paradigms (both hardware and software), and for AI cognitive science the integration of symbolic and neural-level theories --previously pursued separately.

The first steps of the project will create a set of functional requirements that an architecture might meet. The Challenge will be tackled by a succession of sub-challenges --increasingly sophisticated architectures, each meeting a subset of the identified functional requirements. These sub-challenge projects define a natural strategy for attacking the Challenge, since each such project will inform its successors. The success of the entire enterprise will be measured by how far this sequences of projects proceeds towards the human brain and mind.

The Challenge incorporates key aspects of some fifteen submissions to the Grand Challenge Exercise.

---

6 Dependable Systems Evolution

The proposal aims to put the dependability of computer systems on a scientific footing, both as originally designed and as they are evolved in the field. It embodies a science-based definition of dependability; a system is dependable if reliance can be justifiably upon it, in terms such as functionality, availability, safety and security. A condition for justifiability is the rigour of the evidence supplied.

Further, this justification must not deteriorate as the system is evolved to meet changing requirements.

A sub-challenge is to elevate legacy systems also to the same standard of dependability.

Emphasis is placed upon tools (e.g. for verification) equally based on underlying science. Within fifteen years the project will produce prototype tools and example of their use, persuasive enough that tools vendors will develop and market them.

The main shift envisaged in existing thought and practice is in the the science underpinning dependability, and the design methodology to attain it.

Three subtasks are envisaged: study and enhancemant of legacy and COTS systems; calculi, tools and techniques for dependable systems development; study of systems with evolving requirements, with the goal of maintaining dependability.

Support has already been successfully canvassed in many countries.

The goals overlap with those of GC2: Science for Global Ubiquitous computing, though there is significant difference of emhpasis.

---

7 Journeys in Non-Classical Computation

The challenge is to journey through the `gateway events' obtained by breaking our current classical computational assumptions, and thereby develop a mature science of non-classical computation. The essence of the challenge is methodical examination of the consequence of breaking our traditional paradigms. Six are discussed:

(a) The Turing paradigm: computation is logical. Rather, examine all ways in which the real world computes for us; in particular, quantum computing.

(b) The von Neumann paradigm: conmputation is sequential. Rather, examine massive parallelism, decentralised control, fault tolerance.

(c) The output paradigm: what matters is the result of a computation. Rather: observe trajectories, side-effects -- even power-consumption.

(d) The algorithm paradigm: the computation is a closed system. Rather: consider open systems, that may acquire resource on fly.

(e) The refinement paradigm: meeting a specification exactly. Rather: admit approximate solutions, with measures of confidence.

(f) The `computer as artefact' paradigm: computation is not part of the real world. Rather: nature computes, e.g. optical Fourier transforms.

For this Challenge, paradigm shift is not just a consequence but the object of study. Two sub-challenges are identified: (1) To develop a mature discipline of quantum software engineering; (2) develop a science and engineering discipline of approximate computation.

There is a link with GC2, in the sense that some models of ubiquitous computing entail (or allow) the breaking of classical paradigms.