

Journeys in Non-Classical Computation

A Grand Challenge for Computing Research

Susan Stepney, John A. Clark, Andy Tyrrell : University of York

Colin Johnson, Jonathan Timmis : University of Kent

Derek Partridge : University of Exeter

Andrew Adamatzky, Robert E. Smith : University of the West of England

The Challenge

A *gateway event* [26] is a change to a system that leads to the possibility of huge increases in kinds and levels of complexity. It opens up a whole new kind of phase space to the system's dynamics. Gateway events during evolution of life on earth include the appearance of eukaryotes (organisms with a cell nucleus), an oxygen atmosphere, multi-celled organisms, and grass. Gateway events during the development of mathematics include each invention of a new class of numbers (negative, irrational, imaginary, ...), and dropping Euclid's parallel postulate.

A gateway event produces a profound and fundamental change to the system: once through the gateway, life is never the same again. We are currently poised on the threshold of a significant gateway event in computation: that of breaking free from many of our current "classical computational" assumptions. The Grand Challenge for computer science is

**to journey through the gateway event
obtained by breaking our current classical
computational assumptions, and thereby
develop a mature science of Non-Classical
Computation**

Journeys versus Goals

To travel hopefully is a better thing than to arrive.

– Robert Louis Stevenson, "El Dorado", 1878.

Many Grand Challenges are cast in terms of *goals*, of end points: "achieving the goal, before this decade is out, of landing a man on the moon and returning him safely to earth" [37], mapping the human genome, proving whether $P = NP$ or not. We believe that a goal is not the best metaphor to use for this particular Grand Challenge, however, and prefer that of a *journey*.

The metaphor of a journey emphasises the importance of the entire process, rather than emphasising the end point. In the 17th and 18th centuries it was traditional for certain sections of "polite society" to go on "a Grand Tour of Europe", spending several years broadening their horizons: the experience of the entire journey was important. And in the Journey of Life, death is certainly not the goal! Indeed, an open journey, passing through gateway events, exploring new lands with ever expanding horizons, need not have an end point.

A journey of a thousand miles begins with a single step.

– Lao Tzu, *Tao Te Ching*, Chapter 64, ~600 B.C.

Journeys and goals have rather different properties. A goal is a fixed target, and influences the route taken to it. With an open journey of exploration, however, it is not possible to predict what will happen: the purpose of the journey is discovery, and the discoveries along the journey suggest new directions to take. One can suggest starting steps, and some intermediate *way points*, but not the detailed progress, and certainly not the end result.

Thinking of the Non-Classical Computation Challenge in terms of a journey, or rather several journeys, of exploration, we suggest some early way points that appear sensible to aim for. But we emphasise that these *are* early points, that we spy today as we peer through the gateway. As the community's journey progresses, new way points will heave into view, and we can alter our course to encounter these as appropriate.

The Road goes ever on and on.

– J. R. R. Tolkien, *The Lord of the Rings*, 1954.

Six classical paradigms to disbelieve before breakfast

Classical computing is an extraordinary success story. However, there is a growing appreciation that it encompasses an extremely small subset of all computational possibilities.

In many avenues of life, we create unnecessary limitations. Perhaps the most invidious of these are the implicit assumptions we make. We need to distinguish *this has to be the case* from the merely *this has always been the case*. Discoveries may emerge when what was considered an instance of the former is found to be an instance of the latter. For example, dropping Euclid's parallel postulate gave rise to the whole field of non-Euclidean geometry, arguably paving the way for General Relativity. We wish to encourage similar revolts against the assumptions of classical computing. So below we identify several paradigms that seem to define classical computing, but that may not necessarily be true in all computing paradigms, and we encourage the community to drop, invert, or otherwise perturb these paradigms in whatever ways seem interesting. Our brochure of reality-based journeys is a start.

Many computational approaches seek inspiration in reality (mainly biology and physics), or seek to exploit features of reality. These *reality-based computing* approaches hold great promise. Often, nature does it better, or at the very least differently and interestingly. Examining how the real world solves its computational problems provides inspirations for novel algorithms (such as genetic algorithms or artificial immune systems), for novel views of what constitutes a computation (such as complex adaptive systems, and self-organising networks), and for novel computational paradigms (such as quantum computing).

There is a gulf between the maturity of classical computing and that of the emerging non-classical paradigms. For classical computing, intellectual investment over many years is turning craft into science. To fully exploit emerging non-classical computational approaches we must seek for them such rigour and engineering discipline as is possible. What that science will look like is currently unclear, and the Grand Challenge encourages exploration.

Here we outline some assumptions of classical computation, and ways researchers in different fields are challenging them. In later sections we discuss alternatives in more detail. (Some of the categories arguably overlap.)

It ain't necessarily so.

– George Gershwin, *Porgy and Bess*, 1934

1: The Turing paradigm

classical physics: information can be freely copied, information is local, states have particular values. *Rather*, at the quantum level information cannot be cloned, entanglement implies non-locality, and states may exist in superpositions.

atomicity: computation is discrete in time and space; there is a before state, an after state and an operation that transforms the former into the latter. *Rather*, the underlying implementation substrate realises intermediate physical states.

infinite resources: Turing machines have infinite tape state, and zero power consumption. *Rather*, resources are always constrained.

substrate as implementation detail: the machine is logical, not physical. *Rather*, a physical implementation of one form or another is always required, and the particular choice has consequences.

universality is a good thing: one size of digital computer, one size of algorithm, fits all problems. *Rather*, a choice of implementation to match the problem, or hybrid solutions, can give more effective results.

closed and ergodic systems: the state space can be pre-determined. *Rather*, the progress of the computation opens up new regions of state space in a contingent manner.

2: The von Neumann paradigm

sequential program execution. *Rather*, parallel implementations already exist.

fetch-execute-store model of program execution. *Rather*, other architectures already exist, for example, neural nets, FPGAs.

the static program: the program stays put and the data comes to it. *Rather*, the data could stay put and the processing rove over it.

3: The output paradigm

a program is a black box: it is an oracle abstracted away from any internal structure. *Rather*, the trajectory taken by a computation can be as interesting, or more interesting, than the final result.

a program has a single well-defined output channel. *Rather*, we can chose to observe other aspects of the physical system as it executes.

a program is a mathematical function: logically equivalent systems are indistinguishable. *Rather*,

correlations of multiple outputs from different executions, or different systems, may be of interest.

4: The algorithmic paradigm

a program maps the initial input to the final output, ignoring the external world while it executes. Rather, many systems are ongoing adaptive processes, with inputs provided over time, whose values depend on interaction with the open unpredictable environment; identical inputs may provide different outputs, as the system learns and adapts to its history of interactions; there is no prespecified endpoint.

randomness is noise is bad: most computer science is deterministic. Rather, nature-inspired processes, in which randomness or chaos is essential, are known to work well.

the computer can be switched on and off: computations are bounded in time, outside which the computer does not need to be active. Rather, the computer may engage in a continuous interactive dialogue, with users and other computers.

5: The refinement paradigm

incremental transformational steps move a specification to an implementation that realises that specification. Rather, there may be a discontinuity between specification and implementation, for example, bio-inspired recognisers.

binary is good: answers are crisp yes/no, true/false, and provably correct. Rather, probabilistic, approximate, and fuzzy solutions can be just as useful, and more efficient.

a specification exists, either before the development and forms its basis, or at least after the development. Rather, the specification may be an emergent and changing property of the system, as the history of interaction with the environment grows.

emergence is undesired, because the specification captures everything required, and the refinement process is top-down. Rather, as systems grow more complex, this refinement paradigm is infeasible, and emergent properties become an important means of engineering desired behaviour.

6: The “computer as artefact” paradigm

computation is performed by artefacts: computation is not part of the real world. Rather, in some cases, nature “just does it”, for example, optical Fourier transforms.

the hardware exists unchanged throughout the computation. Rather, new hardware can appear as the computation proceeds, for example, by the addition of new resources. Also, hardware can be “consumed”, for example, a chemical computer consuming its initial reagents. In the extreme, nanites will construct the computer as part of the computation, and disassemble it at the end.

the computer must be on to work. Rather, recent quantum computation results [35] suggest that you don’t even need to “run” the computer to get a result!

Doubtless there are other classical paradigms that we accept almost without question. They too can be fruitfully disbelieved.

The Real World : breaking the Turing paradigm

Real World as its own computer

The universe doesn’t need to calculate, it just does it. We can take the *computational stance*, and view many physical, chemical and biological processes *as if* they were computations: the Principle of Least Action “computes” the shortest path for light and bodies in free fall; water “computes” its own level; evolution “computes” fitter organisms; DNA and morphogenesis “computes” phenotypes; the immune system “computes” antigen recognition.

This natural computation can be more effective than a digital simulation. Gravitational stellar clusters do not “slow down” if more stars are added, despite the problem appearing to us to be $O(n^2)$. And as Feynman noted [23], the real world performs quantum mechanical computations exponentially faster than can classical simulations.

Real World as our computer

Taking the computational stance, we may exploit the way the world works to perform “computations” for us. We set up the situation so that the natural behaviour of the real world gives the desired result.

There are various forms of real world sorting and searching, for example. Centrifuges exploit differences in density to separate mixtures of substances, a form of *gravitational sorting*. Vapours of a boiling mixture are richer in the components that have lower boiling points (and the residual mixture is richer in those that have higher boiling points); distillation exploits this to give a form of *thermal sorting*. Chromatography provides chemical means of separation. Ferromagnetic objects can be separated out from other junk by

using industrial-strength magnets. Optics can be exploited to determine Fourier transforms.

Maggots perform the “computation” of eating dead flesh: historically, maggots were used to clean wounds, that is, to perform their computation in a context to benefit us. More recently, bacterial metabolisms have been altered to perform the “computation” of cleaning up pollution.

Access control computations abound. Suitably constructed shape is used to calculate whether the key inserted in a tumbler lock is the correct one. Physical interlocks are exploited for safety and practical reasons across many industries: for example, it is impossible to insert a nozzle from a leaded petrol pump into the fuel tank of a unleaded petrol car.

Real World as analogue computer

We may exploit the real world in more indirect ways. The “computations” of the “real world as our computer” are very direct. Often we are concerned with more abstract questions. Sometimes the physical world can be harnessed to provide results that we need: we may be able to set up the situation so that there is an *analogy* between the computation performed by the real world, and the result we want.

There is an age-old mechanism for finding the longest stick of spaghetti in an unruly pile, exploiting the physics of gravity and rigidity: we can use this to sort by setting up an analogy between spaghetti strand length and the quantity of interest. Mercury and alcohol thermometers use a physical means of computing temperature by fluid expansion: the analogy is between the length of the fluid column and the temperature. Millikan’s calculation of the charge on an electron exploits relationships between velocity of falling oil drops, viscosity of air, the charge on those drops, and the strength of surrounding electric fields.

Classical computing already exploits physics at the level of electron movements. But there are other ways of exploiting nature.

Analogue computing itself exploits the properties of electrical circuits as analogues of differential equations.

DNA computing [4] encodes problems and solution as sequences of bases (strands) and seeks to exploit mechanisms such as strand splitting, recombination and reproduction to perform calculations of interest. This can result in vast parallelism, of the order of 10^{20} strands.

Quantum computing [51] presents one of the most exciting developments for computer science in recent times, breaking out of the classical Turing

paradigm. As its name suggests, it is based on quantum physics, and can perform computations that cannot be *effectively* implemented on a classical Turing machine.¹ It exploits interference, many worlds, entanglement and non-locality. Newer work still is further breaking out of the binary mind-set, with multiple-valued “qudits”, and continuous variables. Research in quantum computing is mushrooming, and it is apparent that we are not yet in position to fully exploit the possibilities it offers. If only small quantum computers were to prove practical then uses could still be found for simulating various quantum phenomena. However, if larger computers prove possible we will find ourselves unprepared.

- Why are there so few distinct quantum algorithms? How can new ones be found?
- How do we discover new a quantum algorithms to solve a given problem? How do we use existing algorithms to solve new problems? How can we find the best algorithms to use given limited computational resources? More generally....
 - What would a discipline of quantum software engineering look like? (See later for more detail.)
- How can quantum computers be harnessed most effectively as part of a hybrid computational approach?

Real World as Inspiration

Many important techniques in computer science have resulted from observing the real world. *Meta-heuristic search* techniques have drawn inspiration from physics (simulated annealing), evolution (genetic algorithms [27] [48], genetic programming [7] [40]), neurology (artificial neural networks [10] [39] [47] [58]), immunology (artificial immune systems [19]), plant growth (L-systems [57]), social networks (ant colony optimisation [11]), and other domains.

These have all proved remarkably successful, or look highly promising, yet the science underpinning their use comes nowhere near matching the science of classical computing. Given a raft of nature-inspired techniques we would like to get from problem to solution

¹ Analogue (as in continuous) computing also breaks the Turing paradigm. But the real world is neither analogue nor classically discrete; it is quantum. So analogue computing might be dismissed as of theoretical interest only. However, the same dismissal might then be made of classically discrete (classical) computation! (The real world is also relativistic, but that paradigm has not been embraced by computation theory, yet.)

efficiently and effectively, and we would like to reason about the performance of the resulting systems. But this falls outside the classical refinement paradigm.

- What would a science of non-classical refinement look like? A science would allow us, for example, to reason confidently about the behaviour of neural networks in critical applications, to derive highly effective systems targeted at highly limited resources.

In the virtual worlds inside the computer, we are no longer constrained by the laws of nature. Our simulations can go beyond the precise way the real world works. For example, we can introduce novel evolutionary operators to our genetic algorithms,

novel kinds of neurons to our neural nets, and even, as we come to understand the embracing concepts, novel kinds of complex adaptive systems themselves. The real world is our inspiration, not a restriction.

- How can we use nature inspired computation to build “better than reality” systems? What are the computational limits to what we can simulate?
- What is the best you can do given many components, each with highly restricted memory and processing ability?

Massive parallelism : breaking the von Neumann paradigm

Parallel processing (Cellular Automata [65], etc) and other non-classical architectures break out of the sequential, von Neumann, paradigm. (The fact that the sequential paradigm is named after von Neumann should not be taken to imply that von Neumann himself was an advocate of purely sequential computation; indeed, he was also one of the early pioneers of CAs [50].)

Under the classical paradigm assumptions, any parallel computation can be serialised, yet parallelism has its advantages.

Real-time response to the environment. The environment evolves at its own speed, and a single processor might not be able to keep pace. (Possibly the ultimate example of this will be the use of vast numbers of nanotechnological assemblers (*nanites*) to build macroscopic artefacts. A single nanite would take too long, by very many orders of magnitude.)

Better mapping of the computation to the problem structure. The real world is intrinsically parallel, and serialisation of its interactions to map the computational structure can be hard. Parallelism also permits collocation of each processor and the part of the environment with which it interacts most. It then permits collocation of the software: software agents can roam around the distributed system looking for the data of interest, and meeting other agents in a context dependent manner.

And once the classical paradigm assumptions are challenged, we can see that serialisation is not necessarily equivalent.

Fault tolerance. Computation requires physical implementation, and that implementation might fail. A parallel implementation can be engineered to continue working even though some subset of its processors have failed. A sequential implementation has only the one processor.

Interference/interaction between devices. Computation requires physical implementation, and those implementations have extra-logical properties, such as power consumption, or electromagnetic emissions, which may be interpreted as computations in their own right (see later). These properties may interfere when the devices are running in parallel, leading to effects not present in a serialised implementation. (Possibly the ultimate example of this is the exponentially large state space provided by the superposed parallel qubits in a quantum computer.)

The use of massive parallelism introduces new problems. The main one is the requirement for *decentralised control*. It is just not possible to have a single centralised source exercising precise control over vast numbers of heterogeneous devices (this is merely a covert attempt to serialise the system). Part of this problem is tackled by the sister *Grand Challenges in Ubiquitous Systems*, and part is addressed in the later section on *open processes*.

In the eye of the beholder : breaking the output paradigm

The classical paradigm of program execution is that an abstract computation processes an input to produce an output. This input-output mapping is a

logical property of the computation, and is all that is important: no intermediate states are of interest, the computation is independent of physical

realisation, and different instances of the computation yield precisely the same results.

Computation, however, is in the eye of the beholder. Algorithms are implemented by physical devices; intermediate states exist, physical changes happen in the world, different devices are distinguishable. Any information that can be observed in this physical world may be used to enrich the perceived computation [16].

Logical Trajectory Observations

An executing algorithm follows a *trajectory* through the logical state space. (Caveat: this is a classical argument: intermediate *quantum* computational states may be in principle unobservable.) Typically, this trajectory is not observed (except possibly during debugging). This is shockingly wasteful: such logical information can be a computational resource in its own right. For example, during certain types of heuristic search the trajectory followed can give more information about a sought solution than the final “result” of the search itself.

- How can logical observations made during execution be used to give useful information?

Physical Trajectory Observations

An executing algorithm is accompanied by physical changes to the world: for example, it consumes trajectory-dependent power as it progresses, and can take trajectory-dependent time to complete. Such physical resource consumption can be observed and exploited as a computational resource, for example, to deduce features of the logical trajectory. (For example, some recent attacks on smart cards have observed the power consumption profile and data-dependent timing of

internal operations to deduce secret key information [14].) Such physical observations provide a very powerful source of information, currently exploited mainly by attackers, but available for more general computational use.

- What physical observations are feasible, and correlated with logical trajectories?
- What new uses can be found for such physical observations?

Differential Observations

An executing algorithm is realised in a physical device. Physical devices have physical characteristics that can change depending on environmental conditions such as temperature, and that differ subtly across *logically* identical devices. (Indeed, much of the rationale for digitisation is the removal of these differences.) So one can make observations not merely of the output of a single execution, but of set of outputs from a family of executions, of multiple systems, of different but related systems. For example, if repeated executions of a search each get 90% of elements of a sought solution correct then repeated executions might be combined to give an overall solution.

- How can diversity of multiple computations be exploited?
- How should diversity be engineered? By repeated mutation of a source program? By embracing technologically diverse solution paradigms?

Higher-order Observations

These are observations not of the program execution itself, but of the execution of the program used to design (the program used to design...) the program.

Open processes : breaking the algorithmic paradigm

In the classical paradigm, the ultimate goal of a computation is reaching a fixed point: the final output, the “result” of the computation, after which we may switch off the computer. The majority of classical science is also based around the notion of *fixed-point equilibrium* and *ergodicity* (ergodicity is the property that the system has well defined spatial and temporal averages, because any state of the system will recur with non-zero probability).

Modern theories of physics consider systems that lack repetition and stability: they are *far from equilibrium* and *non-ergodic*. Perhaps the most obvious non-ergodic, far from equilibrium system is that of life itself, characterised by perpetual evolution (change). Most human problems are also best described in such terms; since computation is ultimately in service of such problems, the

implications of non-ergodic, far from equilibrium physics must be considered in relationship to computing’s future.

Consider the most basic of chaotic systems: the logistic process, parameterised by R .

$$x_{t+1} = Rx_t(1 - x_t)$$

The behaviours of various logistic processes as a function of R are shown in Figure 1, where each point on the plot is a point on the attractor.

For values of $1 < R < 3$, these logistic processes have a fixed point attractor. For $R = 3$ they an attractor of period two. As we raise R , the attractor becomes period four, period eight, etc. This *period doubling* continues as we raise R , and the values of R where each doubling occurs get closer together.

For $R > 3.569945671\dots$ the logistic process's attractor goes through an infinite number of values (except for a few "islands" or order, of attractors with multiples of odd periods). There is a *phase transition* from order (the region of period doubling) to chaos ("random" behaviour). The phase transition point at $R = 3.569945671\dots$ is the so-called *edge of chaos* [44].

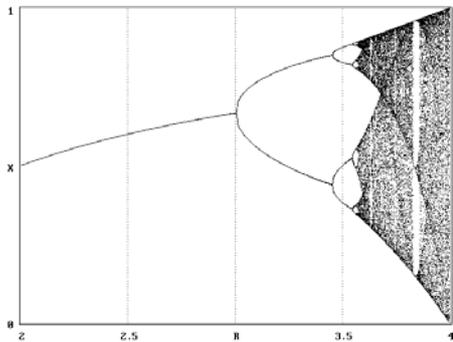


Figure 1: Points on the attractors of various logistic processes, versus the parameter R

Consider a discretised process whose underlying (continuous) dynamics are those of the logistic equation. Imagine taking measurements from this process. Take very coarse measurements: say the process outputs 1 if $x > 0.5$, and 0 otherwise; and take samples of length L bits. For a given L , construct an automaton that represents the process. So now the logistic processes generated by various values of R are being interpreted as a variety of automata: *logistic machines*. It turns out that there is a clear phase transition (a peak in the machine size versus the entropy of the bit sequence) as we move from the period doubling region to the chaotic region.

At the phase transition, the machine size versus the length of the sequence L , *expands without bound*. That is, at the edge of chaos, the logistic machine requires an infinite memory machine for accurate representation. There is a leap in the level of intrinsic computation going on in the logistic machine at the edge of chaos. (In terms of the Chomsky hierarchy, the machine has gone from the

level of regular grammars to the level of context-free grammars.)

At the edge of chaos, we can add new resources (computational or physical) to get results that are neither redundant (as they are in the structured period doubling regime) nor random (as in the chaotic regime). Within the classical paradigm, such conditions would be anathema, indicating unceasing variety that never yields "the solution". But in life-like systems, there is simultaneously sustained order, and useful innovation. In this setting, emergence of the unforeseen is a desirable property, rather disruptive noise.

Some computational approaches attempt to exploit the biological paradigm: cellular automata, evolutionary computation, recurrent networks (autocatalytic, neural, genomic, immune system, ecological webs, ...), social insect and agent-based systems, DNA-computing, and nanite-systems that build themselves. However, in most of these cases, the implementations of such systems have been locked into themselves, *closed*, unable to take on new matter or information, thus unable to truly exploit emergence.

We should consider *open systems*, systems where new resources, and new kinds of resources can be added at any time, either by external agency, or by the actions of the system itself. These new resources can provide gateway events, that fundamentally alter the character of the system dynamics, by opening up new kinds of regions of phase space, and so allowing new possibilities. Computational systems are beginning to open themselves, to unceasing flows of information (if not so much to new matter). The openness arises, for example, through human interactivity as a continuing dialogue between user and machine [64], through unbounded networks, through robotic systems with energy autonomy. As computers become ubiquitous, the importance of *open systems physics* to understanding computation becomes critical. The solutions we expect from people are ongoing processes, and this should be our expectation from computers too.

A coherent revolutionary Challenge, that also respects the past

Classical physics did not disappear when modern physics came along: rather its restrictions and domains of applicability were made explicit.

Similarly, the various forms of non-classical computation will not supersede classical computation: they will augment and enrich it. And when a wide range of tools is available, we can pick the best one, or the best combination, for each job. For example, it might be that using a quantum

algorithm to reduce a search space, and then a meta-heuristic search to explore that, is more effective than using either algorithm alone.

We would like

to create a general flexible conceptual framework that allows effective and efficient exploitation of hybrid approaches, including classical and non-classical components

The journey is the important thing. At various points in journey-space researches will alight to mark their way, leaving behind diary entries to which they may return at a later date. In common parlance these intermediate recordings may be regarded as “achievements”. Opportunities are manifold. We expect journeys relevant to the sub-disciplines to be articulated separately; several have already been prepared. These are

- Non-Classical Physics – Quantum Software Engineering
- Non-Classical Refinement – Approximate Computation
- Computing in non-linear media – reaction-diffusion and excitable processors

- Artificial Immune Systems
- Non-Classical Interactivity – Open Dynamical Networks
- Non-Classical Architectures – Evolving Hardware

See later for details. Also relevant are the sister Ubiquitous Systems challenges.

It is important these separate journeys are not seen as independent explorations. Rather, their results and insights should provide valuable groundwork for the overarching challenge, to produce

<p>a fully mature science of all forms of computation, that unifies the classical and non-classical paradigms</p>
--

The Grand Challenge Criteria

It arises from scientific curiosity about the foundation, the nature or the limits of a scientific discipline. It arises from questioning the assumptions of the classical paradigms, and aims at the creation of a new science.

It gives scope for engineering ambition to build something that has never been seen before. It aims to build a new science; the engineering opportunities will follow.

It will be obvious how far and when the challenge has been met (or not). It will never be met fully: it is an open journey, not a closed goal. The science will continue to mature, until itself overtaken by the next paradigm shift.

It has enthusiastic support from (almost) the entire research community, even those who do not participate and do not benefit from it. No. However, in the best tradition of paradigm shifts, the change will occur.

An important scientific innovation rarely makes its way by gradually winning over and converting its opponents: it rarely happens that Saul becomes Paul. What does happen is that the opponents gradually die out, and that the growing generation is familiarised with the ideas from the beginning.
– Max Planck, *Scientific Autobiography*, 1949

It has international scope: participation would increase the research profile of a nation. This is a new fundamental area of computer science.

It is generally comprehensible, and captures the imagination of the general public, as well as the esteem of scientists in other disciplines. Much popular literature already exists in several of these areas, written by scientists in other disciplines (quantum computing, complexity, nanotech, ...),

and so they and the general public are arguably already ahead of the CS community!

It was formulated long ago, and still stands. Its seeds have been around for a long time, but it has only recently become of obvious importance.

It promises to go beyond what is initially possible, and requires development of understanding, techniques and tools unknown at the start of the project. The structure of the Challenge mirrors the journey suggested by this criterion.

It calls for planned co-operation among identified research teams and communities. It is a multi-disciplinary Challenge, with contributions needed from a range of research specialities.

It encourages and benefits from competition among individuals and teams, with clear criteria on who is winning, or who has won. There need not be a single “winner”. Diversity of solutions should be encouraged to be applicable to a range of application domains. Winners may emerge in particular application domains, as the strengths of the various techniques become clear.

It decomposes into identified intermediate research goals, whose achievement brings scientific or economic benefit, even if the project as a whole fails. There are several components to the Challenge that can be explored in parallel.

It will lead to radical paradigm shift, breaking free from the dead hand of legacy. Non-classical computing is a radical paradigm shift!

It is not likely to be met simply from commercially motivated evolutionary advance. Applications might be supported by industry, but it is unlikely that the development of the underlying science would be.

References and Further Reading

- [1] Editorial article. *Nature Immunology*, 3 (10), 883, October 2002.
- [2] Andrew Adamatzky. *Computing in Nonlinear Media and Automata Collectives*. IoP Publishing, 2001.
- [3] Andrew Adamatzky, ed. *Collision-Based Computing*. Springer, 2002.
- [4] Leonard M. Adleman. Molecular Computation of Solutions to Combinatorial Problems. *Science*, 266:1021-1024, Nov 1994.
- [5] Thomas Back, David B. Fogel, Zbigniew Michalewicz, eds. *Evolutionary Computation I: basic algorithms and operators*. Institute of Physics, 2000.
- [6] Per Bak. *How Nature Works: the science of self-organized criticality*. OUP, 1997.
- [7] Wolfgang Banzhaf, Peter Nordin, Robert E. Keller, Frank D. Francone. *Genetic Programming, An Introduction: on the automatic evolution of computer programs and its applications*. Morgan Kaufmann, 1998.
- [8] Albert-Laszlo Barabasi. *Linked: the new science of networks*. Perseus, 2002.
- [9] Hugues Bersini, Francisco J. Varela. Hints for Adaptive Problem Solving Gleaned from Immune Networks. In H. P. Schwefel, H. Mühlenbein, eds, *Parallel Problem Solving from Nature*. Springer, 1991.
- [10] Christopher M. Bishop. *Neural Networks for Pattern Recognition*. OUP, 1995.
- [11] Eric W. Bonabeau, Marco Dorigo, Guy Theraulaz. *Swarm Intelligence: from natural to artificial systems*. OUP, 1999.
- [12] Daryl W. Bradley, Andy M. Tyrrell. Hardware Fault Tolerance: an immunological approach. In *Proc IEEE Conf on System, Man, and Cybernetics*. 2000.
- [13] Cristian S. Calude, Gheorghe Paun. *Computing with Cells and Atoms*. Taylor & Francis, 2001
- [14] Suresh Chari, Charanjit S. Jutla, Josyula R. Rao, Pankaj Rohatgi. *Power analysis: attacks and countermeasures*. In Annabelle McIver, Carroll Morgan, eds. *Programming Methodology*. Springer, 2003.
- [15] Bastine Chopard, Michel Droz. *Cellular Automata Modeling of Physical Systems*. CUP, 1998.
- [16] John A. Clark, Susan Stepney, Howard Chivers. Breaking the model: finalisation and a taxonomy of security attacks (submitted)
- [17] David W. Corne, Marco Dorigo, Fred Glover, eds. *New Ideas in Optimization*. McGraw Hill, 1999.
- [18] Dipankar Dasgupta, ed. *Artificial Immune Systems and Their Applications*. Springer, 1999.
- [19] Leandro N. de Castro, Jonathan Timmis. *Artificial Immune Systems: a new computational intelligence approach*. Springer, 2002.
- [20] Leandro N. de Castro, Fernando J. von Zuben. An Evolutionary Immune Network for Data Clustering. *SBRN '00, Brazil*, 84-89. IEEE, 2000.
- [21] Marianne Delorme, Jacques Mazoyer, eds. *Cellular Automata: a parallel model*. Kluwer, 1999.
- [22] J. Doyne Farmer, Norman H. Packard, Alan S. Perelson. The Immune System, Adaptation, and Machine Learning. *Physica D*, 22, 187-204, 1986.
- [23] Richard P. Feynman. Simulating Physics with Computers. *Int. J. Theor. Phys.* 21(6/7). 1982.
- [24] Stephanie Forrest, ed. *Emergent Computation: self-organizing, collective, and cooperative phenomena in natural and computing networks*. MIT Press, 1991.
- [25] Stephanie Forrest, Alan S. Perelson, Lawrence Allen, Rajesh Cherukuri. Self-Nonself Discrimination in a Computer. *Symposium on Research in Security and Privacy*, 202-212. IEEE, 1994
- [26] Murray Gell-Mann. *The Quark and the Jaguar*. Abacus, 1994.
- [27] David E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley, 1989.
- [28] Prabhat Hajela, Jun Sun Yoo. Immune Network Modelling in Design Optimization. In [17].
- [29] Emma Hart, Peter Ross. The Evolution and Analysis of a Potential Antibody Library for Use in Job Shop Scheduling. In [17].
- [30] John H. Holland. *Hidden Order: how adaptation builds complexity*. Addison-Wesley, 1995
- [31] John H. Holland. *Emergence: from chaos to order*. OUP, 1998.
- [32] Yoshitero Ishida. Distributed and Autonomous Sensing Based on Immune Network. In *Proc Artificial Life and Robotics, Beppu*. AAAI Press. 1996.
- [33] Henrik Jeldtoft Jensen. *Self-Organized Criticality: emergent complex behaviour in physical and biological systems*. CUP, 1998.
- [34] Niels K. Jerne. Towards a Network Theory of the Immune System. *Annals of Immunology*, 125C, 373-389, 1974.
- [35] Richard Jozsa: Characterising Classes of Functions Computable by Quantum Parallelism. *Proc. R. Soc. Lond. A* 435. 1991
- [36] Stuart A. Kauffman. *The Origins of Order: self-organization and selection in evolution*. OUP, 1993.
- [37] John F. Kennedy. Announcement to the US Congress. 25 May, 1961.
- [38] Jeffrey O. Kephart. A Biologically Inspired Immune System for Computers. In Rodney A. Brooks, Pattie Maes, eds, *Artificial Life IV*. MIT Press, 1994.
- [39] Teuvo Kohonen. *Self-Organization and Associative Memory*. Springer, 1988.
- [40] John R. Koza. *Genetic Programming: on the programming of computers by means of natural selection*. MIT Press, 1992.
- [41] John R. Koza. *Genetic Programming II: automatic discovery of reusable programs*. MIT Press, 1994.
- [42] John R. Koza, Forrest H. Bennett III, David Andre, Martin A. Keane. *Genetic Programming III: Darwinian invention and problem solving*. Morgan Kaufmann, 1999.

- [43] L. F. Landweber, E. Winfree, eds. *Evolution as Computation*. Springer, 2002.
- [44] Christopher G. Langton. Computation at the Edge of Chaos: phase transitions and emergent computation. In [24].
- [45] Christopher G. Langton, ed. *Artificial Life: an Overview*. MIT Press, 1995.
- [46] Benoit B. Mandelbrot. *The Fractal Geometry of Nature*. Freeman, 1977.
- [47] Marvin L. Minsky, Seymour A. Papert. *Perceptrons*. MIT Press, 1988.
- [48] Melanie Mitchell. *An Introduction to Genetic Algorithms*. MIT Press, 1996.
- [49] Mark Neal, Jonathan Timmis. Timidity: A Useful Emotional Mechanism for Robot Control? *Informatica: Special Issue on Perception and Emotion based Reasoning*, 2003. (in press)
- [50] John von Neumann. *Theory of Self-Reproducing Automata*, edited by A. W. Burks. University of Illinois Press, 1966.
- [51] Michael A. Nielsen, Isaac L. Chuang. *Quantum Computation and Quantum Information*. CUP, 2000.
- [52] Mihaela Oprea, Stephanie Forrest. Simulated evolution of antibody gene libraries under pathogen selection. In *Systems, Man and Cybernetics*. IEEE, 1998.
- [53] Derek Partridge. On the difficulty of really considering a radical novelty. *Minds and Machines*, 5(3) 391-410, 1995.
- [54] Derek Partridge. Non-Programmed Computation. *Comm ACM*, 43, 293-301, 2000.
- [55] Gheorghe Paun. *Membrane Computing: an introduction*. Springer, 2002.
- [56] Heinz-Otto Peitgen, Peter H. Richter. *The Beauty of Fractals: images of complex dynamical systems*. Springer, 1986.
- [57] Przemyslaw Prusinkiewicz, Aristid Lindenmayer. *The Algorithmic Beauty of Plants*. Springer, 1990.
- [58] David E. Rumelhart, James L. McClelland. *Parallel Distributed Processing*. MIT Press, 1986.
- [59] Tanya Sienko, Andrew Adamatzky, Nicholas G. Rambidi, Michael Conrad, eds. *Molecular Computing*. MIT Press, 2003.
- [60] Derek J. Smith, Stephanie Forrest, David H. Ackley, Alan S. Perelson. Modeling the effects of prior infection on vaccine efficacy. In [18].
- [61] Susan Stepney. Critical Critical Systems. In *Formal Aspects of Security, FASEC'02*. LNCS vol 2629. Springer, 2003.
- [62] Tommaso Toffoli, Norman H. Margolus. *Cellular Automata Machines*. MIT Press, 1985.
- [63] Duncan J. Watts. *Small Worlds: the dynamics of networks between order and randomness*. Princeton University Press, 1999.
- [64] Peter Wegner. Why interaction is more powerful than algorithms. *CACM*, 40(5). 1997.
- [65] Stephen Wolfram. *Cellular Automata and Complexity: collected papers*. Addison-Wesley, 1994.
- [66] Andrew Wuensche, Mike Lesser. *The Global Dynamics of Cellular Automata*. Addison Wesley, 1992.

Journey: Non-Classical Physics – Quantum Software Engineering

This journey of Non-Classical Computation covers

**the development of a mature discipline of
Quantum Software Engineering**

We wish to be ready to exploit the full potential of commercial quantum computer hardware, once it arrives, projected to be around 2020 (or, less optimistically, “20 years from now”, no matter when “now” is).

We might have to wait a while for commercial quantum computers, but when they arrive, Moore’s law suggests they will grow in power *very quickly*. Doubling a classical computer’s register length (roughly) doubles classical computing power, but adding just *one bit* to a quantum computer’s register doubles quantum computing power. We need to be ready to exploit these devices once they appear. However, the majority of today’s theory of computation, algorithms, programming languages, specification models, refinement calculi, and so on, is purely classical. The challenge is to build the corresponding languages, tools and techniques for quantum software engineering.

We need to raise the level of thinking about quantum programs. Today we reason about

quantum programs predominantly at the level of quantum *gates*: imagine how far classical computing would have progressed if the only language we had to describe programs was that of AND and OR gates! Most importantly, we need a new paradigm (or paradigms) for thinking about quantum computations, to augment the existing classical declarative, functional, and imperative paradigms.

The whole of classical software engineering needs to be reworked and extended into the quantum domain.

Foundations

Much foundational work is still needed. We need further developments of the fundamentals of quantum computability: the **Universal Quantum Turing Machine**. We need to investigate **quantum algorithmic complexity**: time, space, “parallel universe space”, and any other parameters of interest.

We have models of classical computation – von Neumann machines with fetch-execute-store, imperative, functional and logic languages, etc – that let us write and reason about classical

programs without worrying about logic levels, transistors, gates, etc. In much the same way we need **metaphors and models of quantum computation**, that enable us design and reason about quantum algorithms without recourse to QM, unitary matrices, etc. Does Deutsch's many-worlds description provide the best programming metaphor, or are there better ones? Whatever the actual metaphors chosen, they must be formalised into new computational models.

We need theories and models of that weirdest quantum process of all: that of **quantum entanglement**. Two qubit entanglement is relatively well understood – but multi qubit entanglement, and qudit entanglement, are barely understood.

Languages and Compilers

We need to determine the **fundamental building blocks** of quantum programming: is there a simple extension of GCL? of classical logic languages? of classical assembly languages? is an entirely new paradigm needed?

We need to design suitable **assembly level** and **high level Q-languages** (analogues of classical imperative, declarative, and functional languages, at 3rd, 4th, 5th generation, and beyond). We need to design and build the corresponding **Q-compilers** for these languages.

We need to design and implement (initially, simulate) **new Q-algorithms** (beyond the current ones of Min Max, Shor's period finding algorithm used for factorization, and Grover's algorithm for DB searching). What classes of algorithms may be quantised? How may certain well-known classical algorithms be quantised?

We need to develop suitable **reasoning systems** and **refinement calculi** for these languages. (Even sequential composition is different in the quantum regime, due to the fundamental unobservability of the intermediate state.) Although higher level specifications may well abstract away from details of any underlying classical *versus* quantum implementation, there may be certain application-specific **quantum specification languages**, for example, for **quantum protocols**.

Methods and Tools

Before commercial quantum computers are available, we have to make do with simulations on classical machines. We need to implement powerful **quantum computer simulators**, in order

to perform computational experiments and validate language and algorithm designs. (Computational resources for simulating quantum algorithms can be exponentially large. Something like a simulation engine of multiple FPGAs might be appropriate, to get the required massive parallelism.)

We need to discover what **higher level structuring techniques and architectures** are suitable for quantum software. In particular, can classical structuring (such as object-orientation, or component based software), be extended to incorporate Q-software? How can classical and quantum paradigms co-exist? (It seems likely that, at least to start with, most software will remain classical, with a "call-out" to quantum power as needed. But the development process needs to be able to handle such hybrid developments seamlessly.)

Given that quantum execution is *in principle* unobservable, we need to discover new **debugging and testing techniques** for these Q-languages.

We need to design ways of **visualising** Q-algorithm execution, as an aid to understanding, design, and implementation.

Novel Quantum possibilities

Quantum computing can do some things that cannot even be simulated by discrete deterministic classical computers. We need to extend quantum software engineering to encompass these new domains.

Quantum devices can produce **genuine random numbers**; classical digital simulations can produce only *pseudo*-random numbers. We need to investigate the differences this causes, if any. In the short term, will a quantum computer simulator need to be hooked up to a genuinely random number source? In the longer term, what new power, what new difficulties, might emerge as a result of genuine randomness?

Quantum **entanglement** offers many new possibilities, such as information teleportation. We need to understand how entanglement can be applied to produce genuinely new algorithms, and new kinds of protocols.

Journey: Non-Classical Refinement – Approximate Computation

This journey of Non-Classical Computation is

to develop a science of approximate computation, and to derive from it a well-founded discipline for engineering approximate software

A radical departure from discrete correct/incorrect computation is required, a shift away from logics towards statistical foundations, such that meaningful estimates of ‘confidence’ emerge with each approximate result. This implies that probabilities play an integral part in computation throughout the process. The component probabilities and the eventual confidence estimates, if secured by large numbers (e.g. repeated sampling from a proposed distribution), imply a computational effort that is becoming increasingly feasible as a result of hardware advances as well as innovative developments in statistical modelling theory (e.g. reversible-jump Markov Chain Monte Carlo methods).

Classical computation versus approximations

The classical, discrete, view of computation has each step as either correct or incorrect, and the middle ground excluded. This naturally leads to formal logics as the dominant underpinning framework. The programmer devises the “formula”, which is intended to be an exact solution to the problem; this symbol structure is translated into a machine executable form and the manipulations that the programmer envisaged are performed automatically, at high speed and with complete accuracy.

Consider the symbol structures being manipulated by a trained a multilayer perceptron (MLP), for example. These are not formulae composed of operators and variables that admit a ready mapping to the operations and parameters of the human conception of the problem. One consequence is that any adjustment to the function to be computed by an MLP involves complete retraining, because code-fixing is not an option. The “formulae” cannot reasonably be devised by a programmer; they must be automatically generated from data samples.

Typically, the output of an MLP classifier, a real-value, is arbitrarily thresholded to obtain a class label. This and other inherent weaknesses of an approximate classifier constructed with empirically determined (suboptimal) values for its parameters are widely acknowledged. *Ad hoc*-ery is rife in neural computing, but work on error-bars already points the way towards a well-founded science.

These innovative developments to move beyond the constraint of correct/incorrect results from hand-crafted formulae are but piecemeal strategies; they need to be woven into the basic fabric of a comprehensive model for approximate computation, not stitched-on to classical computation as useful extras, or mere curiosities.

How would the classical paradigm be shifted?

Taking the viewpoint that the computational task is an unknown (or intractable, see later) function, the computational goal is to approximate it in a way that holds the promise of reasonable optimality, but crucially associates a *meaningful estimate of confidence* with every output computed. In general terms, data-driven software development supplants specification-driven; computational tasks are viewed as data-defined rather than (abstract) specification-defined.

In detail, the approach might be through a survey, sampling by, say, Markov Chain Monte Carlo methods across a continuum of potentially viable models. By doing this within a Bayesian framework, rationalisable probabilities are attached to various elements throughout the computational process. The outcome is a weighted average across a range of modelling possibilities. It is a well-founded approximation whose validity emerges as a secure estimate from the computational processes employed. The infrastructure of the new paradigm seeks to avoid searching, comparing and selecting from amongst a discrete set of alternative models (and hence commitment to a specific model, or even discrete set of alternative models) by maintaining the range of potential models as a set of continuous parameters; probability theories, secured by large-number sampling, provide the over-arching framework.

A fundamental basis of continuity avoids the brittleness inherent in discrete, classical computation. Notice, for example, that the necessary discretisation of the real numbers that plagues classical computation is not similarly problematic for MLPs, despite their fundamental dependence upon the real continuum.

Initially at least, classical computation will provide the virtual machine upon which the approximate computations will run, but hardware innovations coupled with the establishment of generally applicable approximation algorithms could change that dramatically. However, building the required confidence in a classically programmed virtual machine is not the same scale of problem as doing

it individually for every piece of application software.

The initial challenge is to begin to establish the limits and the potential infrastructure of such a science of approximate computation. This includes major subdomains, such as a discipline of engineering approximate software. It also involves the identification and integration into a coherent framework of many activities that are currently pursued under a variety of labels, for example, statistical pattern recognition, some varieties of data mining, statistical data modelling, some technologies of inductive generalization or data-driven computation.

A science of approximate computation: when and where?

The new science of approximate computation will not oust the classical one; it will sit alongside it as a new weapon in an armoury of well-founded alternative computational techniques to be used when appropriate.

It will be appropriate to use whenever a computational task is defined more by samples of desired or observed behaviour than by an abstract specification. It will also be appropriate to use whenever the problem is well defined but computationally intractable, where the particular task is appropriate for approximate solutions, albeit with a 'confidence' measure attached; there is no prohibition on certainty emerging as an extreme of approximation.

Consider an illuminating extreme --- safety-critical software. Such systems would seem to absolutely require the classical strategy: they must be correct. However, the practical impossibility of this requirement leads to a slight relaxation: it is typically couched in terms of a very low failure/error rate, and the major component of the required assurances is extensive testing. The bulwark of statistical reasoning, as an integral part of the testing, is thus dragged in by the back door (as it were) – how much better to integrate it into the fabric of the computation from beginning to end, instead of tagging in on the end as a stopgap for verification failure?

Will 'programming' an approximation computer be more difficult than conventional programming? All we can say is it will be fundamentally different – for example, data analysis, selecting sampling strategies, rather than formula derivation. The 'programming' difficulties that confront the user of this new paradigm will be directly determined by how successful we are in formulating the fundamental model(s) of approximate computation.

Nothing in the above *requires* a new paradigm: any of the innovations envisaged could be realised within the scope of classical computation, as some already are. However, although a screwdriver can be used to open a tin, it is quicker, neater and generally preferable to use a well-designed tin opener for the task.

Journey: Computing in non-linear media – reaction-diffusion and excitable processors

Nature diffuses, reacts and excites. Does it compute? Everyday life gives us instances of propagating structures: dynamics of excitation in heart and neural tissue, calcium waves in cell cytoplasm, the spreading of genes in population dynamics, forest fires. Could we use the travelling structures — emerging in reaction-diffusion and excitable systems — to do computation? This journey of Non-Classical Computation is

to develop a science of computation using spatio-temporal dynamics and propagating phenomena, in many-dimensional amorphous non-linear media

What is so particular about reaction-diffusion systems? A non-linear chemical medium is *bistable*: each micro-volume of the medium has at least two steady stable states, and the micro-volume switches between these states. In the chemical medium, fronts of diffusing reactants

propagate with constant velocity and wave-form; the reagents of the wave front convert reagents ahead of the front into products left behind. In an excitable chemical medium the wave propagation occurs because of coupling between diffusion and auto-catalytic reactions. Auto-catalytic species produced in one micro-volume of the medium diffuse to neighbouring micro-volumes, and thus trigger an auto-catalytic reaction there. So an excitable medium responds to perturbations that exceed the excitation threshold, by generating excitation waves.

Why are excitation waves so good for computing? Unlike mechanical waves, excitation waves do not conserve energy, rather, they conserve waveform and amplitude, do not interfere, and generally do not reflect. So excitation waves can play an essential role of information transmission in active non-linear media processors.

Specialised non-linear media processors

A problem's *spatial representation* is a key feature of reaction-diffusion processors. Data and results are represented through concentration profiles of the reagents, or spatial configurations of activity patterns. A computation is also defined in a physical space. The computation is realised by spreading and interacting waves of the reagents, or excitation patterns. A computational code, or program, is interpreted in a list of possible reactions between the diffusing components, and in a form of diffusive or excitation coupling between micro-volumes of the computing medium. Usually, such properties could not be changed online. However they can be determined and adjusted to work towards the solution of a particular problem. Therefore most reaction-diffusion processors are intentionally designed to solve a few particular problems — they are specialised. Examples of working laboratory prototypes of specialised reaction-diffusion computing devices include Belousov-Zhabotinsky chemical medium image processors implemented (Kuhnert-Agladze-Krinsky and Rambidi designs), chemical processors for computation of a skeleton of planar shape, plane sub-division (Voronoi diagram), shortest path (Adamatzky-Tolmachev-De Lacy Costello designs), chemical processors for robot navigation and taxis (De Lacy Costello-Adamatzky implementations).

Experimental chemical computers are very rare species — there are just a handful of chemical processors in the world — why? Because the design of every chemical processor requires at least a chemist and a computer scientist. So one of the actual tasks deals with the make-up of non-classical scientists is

to dissolve boundaries between chemical-physical sciences and theoretical computer science, training a new generation of scientists who tackle theory and experiments with ease

Computational (logical) universality of non-linear media: dynamical vs. static architectures

A natural process is called *computationally universal*, by analogy with mathematical machines, if it potentially can calculate any computable logical function, so realises a functionally complete set of logical gates in its spatio-temporal dynamics. So in the state of the given system, we need to represent information quanta (usually TRUE and FALSE values of a Boolean variable), logic gates (where information quanta are processes), and routes of information transmission or connections between the gates. There are two ways to build a

logical circuit in a non-linear system: structural, or *stationary*, and structureless, or *dynamic*, designs.

In a *stationary* design, which underlines an architecture-based universality, a logical circuit is embedded into a system in such a manner that all elements of the circuit are represented by the system's stationary states; the architecture is static and its topology is essential for a computation. The architecture-based universality allows for applying conventional solutions to unconventional materials: we could fabricate varieties of traditional computers made of non-standard materials (glass tubes filled with chemicals, molecular arrays, excitable tissues). Examples of stationary architectures of non-classical computers include Hjelmfelt-Weinberger-Ross mass-transfer or kinetic-based logical gates (a logical gate is constructed by adjusting flow rates of chemical species between several reactors), Tóth-Showalter circuits (a logical circuit comprises several narrow tubes filled with Belousov-Zhabotinsky excitable chemical medium, the tubes are connected via expansion chambers where logical functions are implemented by interacting wave fronts), and Adamatzky-de Lacy Costello palladium gate (implemented in simple non-excitable reaction-diffusion chemical system).

Most biological systems are “architecture-based computationally universal”. Could they compute better if they lose this compartmentalisation? If all neurons in our brain were to dissolve their membranes and fuse into an amorphous mass, could this “gooware” be computationally more powerful?

Dynamic, or collision-based, computers employ mobile self-localisations, which travel in space and execute computation when they collide with each other. Truth-values of logical variables are represented by absence or presence of the travelling information quanta. There are no pre-determined wires: patterns can travel anywhere in the medium, a trajectory of a pattern motion is analogous to a momentarily wire. A typical interaction gate has two input ‘wires’ (trajectories of the colliding mobile localisations) and, typically, three output ‘wires’ (two representing the localisations’ trajectories when they continue their motion undisturbed, the third giving the trajectory of a new localisation formed as the result of the collision of two incoming localisations). The travelling is analogue to information transfer, while collision is an act of computation, thus we call the set up *collision-based computing*. There are three sources of collision-based computing: Berlekamp-Conway-Guy proof of universality of Conway's Game of Life via collisions of glider streams, Fredkin-Toffoli conservative logic and cellular automaton implementation of the billiard ball

model (Margolus block cellular automata), and the Steiglitz-Kamal-Watson particle machine (a concept of computation in cellular automata with soliton-like patterns). A wide range of physical, chemical, and biological media are capable of collision-based computing. Thus, for example, this type of computation can be implemented as localised excitation in two- and three-dimensional excitable lattices, as breathers and solitons travelling along polymer chains (and DNA molecules), as excitons in mono-molecular arrays (like Scheibe aggregates), and even as quasi-particles in gas-discharge systems and two-component reaction-diffusion systems.

So far we can implement a gate or two in collision of propagating localisations — what about a collision-based chip?

Complexity and Computation

Non-linear sciences are a centrepiece of contemporary sciences, from physics to biology. The dynamics of non-linear media are becoming a crucial tool in understanding complex behaviour in natural and artificial systems, emergent behaviour, complexity, and self-organized criticality. Non-linear dynamics of large-scale massive systems, described in terms of physical systems or their continuous and discrete mathematical and computational models, are typically recruited at the present time to explain the nature of complexity, to predict the behaviour of biological and social systems, and to discover the novel properties of multi-component systems. To develop a solid framework of computing in non-linear media we must answer a number of questions residing at the edge between complexity and computation.

What families of computational tasks are solved in chemical and biological non-linear media?

How are the computing abilities of non-linear media related to the behavioural complexity of the media itself? Is it necessarily true that a system with a more complex behaviour can solve more computationally complex problems than a system with less complex behaviour?

What is a relation between complexity of space-time dynamics of a non-linear medium, and computational complexity of the medium as a computer?

How do we exert some control over the dynamics of non-linear media? What engineering

approaches are required to form interfaces between conventional computers and experimental prototypes of non-linear media based computers?

How do we program non-linear medium computers? What is a trade-off between the medium's complexity and the medium's programmability? Does complexity reduce programmability?

Cellular automata: Non-linear medium mathematical machines

The field of cellular automata — massive-parallel locally-connected mathematical machines — flourishes and occupies a significant part of computational sciences. A cellular automaton is a lattice of uniform finite automata; the automata evolve in discrete time and take their states from a finite set. All automata of the lattice update their states simultaneously. Every automaton calculates its next state depending on the states of its closest neighbours.

Cellular automata models of reaction-diffusion and excitable media capture essential aspects of natural media in a computationally tractable form, and thus could be adopted as a tool for automatic design of non-linear media computers, development of reaction-diffusion algorithms, and pre-experiment verifications.

Discovering Computation

How do we find reaction-diffusion or excitable media to fulfil our computational goals in real wet-ware? There is not much choice at the moment. There are dozens of oscillating chemical reactions, yet most look quite similar, and so almost everybody experiments mainly with Belousov-Zhabotinsky media. The advantage of such ubiquity is the chance to verify each other's experiments. At the molecular level the situation is not as good: we can fabricate molecular arrays, but there are almost no reports on any feasible computing experiments, either with classical waves, or with mobile-self localisations. Which problems can be solved in what types of non-linear media? Should we fabricate these media from scratch or could we instead search for already existing species in nature?

What are the principal characteristics of spatially-extended non-linear media (chemical, physical or biological) that enable them to implement useful computation?

Journey: Artificial Immune Systems

The inspiration and the analogy

There is a growing interest in the use of the biological immune system as a source of inspiration to the development of computational systems [19]. The natural immune system protects our bodies from infection with a complex interaction of white blood cells, called B Cells and T Cells. Upon encountering an antigen (an infecting item), B Cells are stimulated by interacting with the antigen, and, with the help of T Cells, undergo rapid cloning mutation. This is an attempt by the immune system to kill off the invading antigen and prepare the immune system for another infection from that antigen (or similar antigen). The immune system maintains a *memory* of the infection, so that if ever exposed to the same antigen again, a quicker response can be elicited against the infection.

There are many facets of the immune system that can be considered useful for computation, including pattern recognition, feature extraction, learning, noise tolerance, memory, and inherent distributed parallel processing. For these and other reasons, the immune system has received a significant amount of interest as a metaphor within computing. This emerging field of research is known as Artificial Immune Systems (AIS).

Essentially, AIS is concerned with the use of immune system components and processes as inspiration to construct computational systems. AIS is very much an emerging area of biologically inspired computation. This insight into the natural immune system has led to an increasing body of work in a wide variety of domains. Much of this work emerged from early work in theoretical immunology [34] [22] [9], where mathematical models of immune system process were developed in an attempt to better understand the function of the immune system. This acted as a mini-catalyst for computer scientists, with some of the early AIS work being on fault diagnosis [32], computer security [25], and virus detection [38]. Researchers realised that, although the computer security metaphor was a natural first choice for AIS, there are many other potential application areas that could be explored, such as machine learning [20], scheduling [29], immunised fault tolerance [12], and optimisation [28]. In addition, AIS has been offering better understanding of the immune system [60] [52], whose mechanisms are hugely complex and poorly understood, even by immunologists. The field of AIS is both a powerful computing paradigm and a prominent

apparatus for improving understanding of complex biological systems.

Questions can also be asked such as: How do we construct truly autonomous evolving systems that are capable of adapting to an ever-changing environment? How do we construct systems that can implement complex control mechanisms that are beyond the capabilities of current approaches? How do we cope with the massive increase in complexity of systems that are being given to the information technology society as a whole?

AIS algorithms have the possibility of breaking the *algorithmic paradigm* in two ways. First, they capture the immune system's mechanism of exploiting randomness, therefore removing the idea that "randomness is bad". Secondly, the immune system is inherently a continually learning system with no end point, with no "final output". Clearly, current solutions to such problems have made some progress. However, with the increases in scale and complexity come new and ill-understood demands on computational systems. This has resulted in many systems being inflexible, *ad hoc*, difficult to configure, and impenetrably arcane to maintain. Therefore, alternative ways to construct a new generation of more autonomous and self-organising computational systems are being sought.

The models

There are two main models for AIS: the population based models (or selection models), and the network model (see [19] for details), which have impacts on different areas of the main Grand Challenge.

The Selection model

The immune selection model is computationally inspired by the processes during early maturation of immune cells, before they are released into the lymphatic system. It uses some particular algorithm (positive, negative, clonal, ...) to select a set of *recognisers* (supervised learning) or *classifiers* (unsupervised learning), of self or non-self (details depending on the precise algorithm).

This model fits well with the other bio-inspired soft learning systems, such as neural nets and genetic algorithms. The major contributions to the Grand Challenge are in the area of *breaking the refinement paradigm*.

In all these soft learning approaches, there is a discontinuity between the problem statement and the bio-inspired solution. With both NNs and AISs, the solution is distributed over the entire

system. Each artificial antibody may recognise several different antigens: the specific response to a particular antigen is a global property of all the antibodies. The complex response emerges from the simpler behaviour of individual parts.

The way point questions specific to AIS include:

- What are the effects on the selection algorithm of parameter selection, with regards to the outcome and applicability of these algorithms?
- Can we observe the computational trajectory taken during selection and recognition to get useful information?

The immune system selection model forms an excellent exemplar for breaking the refinement paradigm. The challenge is to develop a *science of non-classical refinement*, that permits quantitative reasoning about *all* bio-inspired algorithms, including AISs, in both a bottom up and top down manner:

- *understanding* and *predicting* the global recognisers and classifiers that emerge from a collection of local non-specific agents
- a means to *design* and *implement* appropriate sets of recognisers or classifiers for particular applications, in a rigorous (but possibly non-incremental) way
- *quantitative description methods* that enable rigorous reasoning about the behaviour of AISs, such that they can be used reliably in critical applications

Taking inspiration and input from all the bio-inspired learning algorithms, major way points on the Non-Classical Computation journey are

- a *general theory of learning systems* that includes neural, evolutionary, and immune systems as special cases
- use of the general theory to develop *more effective kinds* of learning systems, inspired by, but not based upon, any known biological processes

The Network model

The immune system network model is computationally inspired by the biological processes used to maintain a dynamic “memory” of immune responses, in a system where the lifetime of individual immune memory cells is on the order of weeks, yet the memory itself persists on the order of years or decades.

Investigations and a deeper understanding of the nature of scale free networks and their relation to complex systems is required, to allow a greater understanding of a variety of network type structures. In particular, the formalisation of IS molecular-cellular network by means of modern

graph theory (small-world models, scale-free networks theorisation) should be pursued, to depict the topological features and attributes affecting the functionality of the network. Graph theory is one of the most effective and advantageous instruments for understanding the evolution of network systems and a comprehension of the basic principles of their structural organisation and evolution. As such, it is needed to find the best solutions to the problems of real world networks. This approach, proceeding from the formalisation of elements of the network and their interactions as nodes and links, allows structuring a topology whose characterising features can be derived from analytical and numerical solutions. Modern graph theory has already been successfully exploited for studies of the topological and connective features of existing real world networks like, for example, citations of scientific papers and networks of collaborations, WWW and Internet, biological networks as neural networks, metabolic reactions network, genome and protein network, ecological and food webs, world web of human contacts and languages, telephone call graphs, power grids, nets of small world components. As far as we know, a similar approach has never been applied to the study of realistic (beyond the “undistinguishable clones” hypothesis) IS network peculiarities (see also [1]). By exploring the nature of scale-free networks, immune system offer insight into *breaking the von Neumann paradigm* in terms of allowing for massive parallelism at sensible computational costs.

The biological models

Like many other biologically inspired computational ideas, the computer science and biology of immune systems are developing in parallel. The natural immune system, in particular, is exceedingly complicated, and not understood at all well. Additionally, the immune system does not act in isolation. Indeed, there are many interactions with other biological systems, such as the nervous systems and endocrine (or hormonal) systems. The interactions between these systems lead, in part, to the biological concept of *homeostasis*: a steady internal state. By exploring these ideas, there is the potential to break many of the well-held paradigms outlined in this challenge, such as the *algorithmic* and *refinement* paradigms. Limited work has already begun on this approach, with the development of a small controller for robots [49].

Currently, the field of AIS is limited to the development of algorithms *in silico*; questions are still to be asked similar to that of DNA and quantum computing, such as: it is possible to construct computational devices based on the chemical process inherent in the immune system? The immune system has a wealth of complexity for

computation: rather than just extract metaphors from it, is it possible to exploit the biological mechanisms? The current discipline of AIS may have been inspired by biology, but it is painfully clear that AISs are but a pale shadow of the vast complexity of subtlety of the natural immune system. Computer scientists, mathematicians and immunologists working together can ask, and answer, some deep and interesting questions. For example:

- How might we use the real immune system, and other real physical and biological systems, for computation?
- To what extent is the working of the immune system, and other biological systems, dictated by the physical substrate? Can all putative “immune” responses be realised on all substrates? Do some diseases exploit *computational* constraints of the immune system to defeat it?
- How can we use models to decide which parts of the biology are necessary for correct robust

functioning, which parts are necessary only because of the particular physical realisation, and which parts merely contingent evolutionary aspects?

- How can we use nature inspired computation to build “better than reality” systems? What are the computational limits to what we can simulate?

Conclusions

AIS do not break all the classic computational paradigms: for example, they do not (yet?) use concepts from quantum physics. However, they do challenge some of the major paradigms. The selection model is a good exemplar for examining alternatives to the refinement paradigm, and the network model is an excellent exemplar for examining open network dynamics and emergence, necessary for a full science of complex adaptive systems.

Journey: Non-Classical Interactivity – Open Dynamical Networks

Dynamic reaction networks can have complex non-linear interactions, and feedback where reaction products may themselves catalyse other reactions in the network. They exhibit the emergent complexity, complex dynamics, and self-organising properties [6] [36] of many far-from-equilibrium systems. These systems, and others, can self-organise into regions “at the edge of chaos”, neither too ordered nor too random, where they can perform interesting computations (or computation analogues). There are many dynamic network models that occur in biological and social systems, from Kauffman’s autocatalytic networks [36], and genomic control networks, through dynamical models of neural networks and cytokine immune networks, to ecological food webs, and social and technological networks.

All these subject areas could benefit from better networks models. Much of the existing mathematical network theory is restricted to static, homogeneous, structured, closed networks, since these are the simplest, most tractable models to work with. However, these are not realistic models of biological networks: for example, antibodies rove around the body (network, system, ...) looking for the anomalies, and new kinds of attacks call for new kinds of defence. The journey is

to develop a pragmatic theory of dynamic, heterogeneous, unstructured, open networks

Dynamic: the network is not in steady state or equilibrium, but is far from equilibrium, governed

by attractors and trajectories. (*Swarm networks* may offer insights to this kind of dynamics [11])

Heterogeneous: the nodes, the connections, and the communications can be of many different types, including higher order types.

Unstructured: the network connectivity has no particular regularity: it is not fully regular, or fully connected, or even fully random. Clearly there need to be *some* kinds of regularity present, but these are likely to be of kinds that cannot be reasoned about in terms of simple averages or mean field notions; they are more likely have fractal structure. Some recent advances in *Small World* networks offer intriguing new insights [8] [63].

Open (metadynamic): the structures are unbounded, and the components are not fixed: nodes and connections may come and go; new *kinds* of nodes and connections may appear.

A general theory of such networks would have wide applicability. Such a theory is a basic requirement of complex systems development in general, one application of which is pervasive, or ubiquitous, computing (the subject of another Grand Challenge). Such a theory a necessary way point for answering many challenging questions.

Computation as a dynamical process. What are the various attractors of a dynamical computation? How can we encourage the system to move to a

“better” attractor? How can we map the route through intermediate attractors that it should take?

Computation at the edge of chaos. What are its capabilities? How can we hold a system at the edge, far from equilibrium, to perform useful computations? How can we make it self-organise to the edge?

Designed emergence. How can we design (refine) open systems that have desired emergent properties? And do not have undesired emergent properties?

Open systems science. What are the fundamental properties of open systems? How can we predict the effect of *interventions* (adding new things, or removing things) to the system? How can we understand the effect of a gateway event that opens up new kinds of regions of phase space to a computation? How can we design a system such that gateway events, natural changes to phase space, can occur endogenously?

Journey: Non-Classical Architectures – Evolving Hardware

This journey of Non-Classical Computation is

to develop (biologically-inspired) computing hardware that can adapt, evolve, grow, heal, replicate, and learn

Computation Models

Biological inspiration in the design of computing machines finds its source in essentially three biological models:

- **phylogenesis** (P), the history of the evolution of the species
- **ontogenesis** (O), the development of an individual as orchestrated by his genetic code
- **epigenesis** (E), the development of an individual through learning processes (nervous and immune systems) influenced both by the genetic code (the innate) and by the environment (the acquired).

These three models share a common basis: the genome.

Phylogenesis: evolution

The process of evolution is based on alterations to the genetic information of a species through two basic mechanisms: *selective reproduction* and *variation*. These mechanisms are non-deterministic, fundamentally different from classical algorithms, and potentially capable of providing astonishingly good solutions to problems that are formally intractable by deterministic approaches. Existing analytical and experimental tools are not designed for tackling such stochastic search algorithms, however. We need new tools and methodologies for generating novel results.

Phylogenesis already provides considerable inspiration for algorithm design, in the discipline of *evolutionary computation* [5] [43], which includes genetic algorithms [27] [48] and genetic programming [7] [40]. It has yet to have such an impact on the conception of digital hardware,

however. Koza *et al.* pioneered the attempt to apply evolutionary strategies to the synthesis of electronic circuits when they applied genetic algorithms to the evolution of a three-variable multiplexer and of a two-bit adder. Evolutionary strategies have been applied to the development of the control circuits for autonomous robots, and other research groups are active in this domain. Although technical issues pose severe obstacles to the development of evolvable electronic hardware, there is still much to be gained from evolutionary design given the appropriate hardware and software mechanisms.

Ontogenesis: growth

Ontogenesis concerns the development of a single multi-cellular biological organism. A set of specific mechanisms define the *growth* of the organism: its development from a single mother cell (*zygote*) to the adult phase. The zygote divides, each offspring containing a copy of the genome (*cellular division*). This continues (each new cell divides, creating new offspring, and so on), and each newly formed cell acquires a functionality (liver cell, epidermal cell, ...) depending on its surroundings, its position in relation to its neighbours (*cellular differentiation*).

Cellular division is therefore a key mechanism in the growth of multi-cellular organisms, impressive examples of massively parallel systems: the $\sim 6 \times 10^{13}$ cells of a human body, each one a relatively simple element, work in parallel to accomplish extremely complex tasks. Development processes inspired by biological growth should provide relevant insights on how to handle massive parallelism in silicon. There are also great gains to be achieved by using ontogenetic mechanisms with regard to fault tolerance and reliability.

Epigenesis: learning

The human genome contains $\sim 3 \times 10^9$ bases, yet an adult human body contains $\sim 6 \times 10^{13}$ cells, of which

$\sim 10^{10}$ are neurons, with $\sim 10^{14}$ connections. The genome cannot contain enough information to completely describe all the cells and synaptic connections of an adult organism. There must be a process that allows the organism to increase in complexity as it develops. This process, *epigenesis*, includes the development of the nervous, immune, and endocrine systems.

Epigenetic, or *learning*, mechanisms have already had considerable impact on computer science, and particularly on software design. The parallel between a computer and a human brain dates to the very earliest days of the development of computing machines, and has led to the development of the related fields of artificial intelligence and *artificial neural networks*.

Living organisms interact with their environment and respond to sensory inputs. In many cases this behaviour is learnt over a period of time, after which a specific stimulus will trigger the same, possibly context dependent, response. Such behaviour is mainly controlled by *spiking neurons* and their interactions. Novel hardware developments are being inspired by these observations.

A more recent addition in the general area of hardware systems and epigenetic processes are *artificial immune systems*. Here the sophisticated mechanisms associated with “fault tolerance” in nature have been adapted for electronic hardware system designs [12].

Complexity and Reliability

As systems become more complex it becomes increasingly difficult to provide comprehensive fault testing to determine the validity of the system. Hence faults can remain in a system, and manifest themselves as errors. Furthermore, faults may be introduced into hardware from external sources, such as electromagnetic interference. Components within a system can die; no transistor will function forever. These faults can ultimately cause a system to fail. The ability of a system to function in the presence of such faults, to become *fault tolerant*, is a continually increasing area of research.

Through millions of years of refinement, biology has produced living creatures that are remarkably

fault tolerant. They can survive injury, damage, wear and tear, and continual attack from other living entities in the form of infectious pathogens. Biology manages to take huge amounts of potentially unreliable matter and use self-checking, self-repair, self-reconfiguration, multiple levels of redundancy, multiple levels of defence, even removing suspected cells, to develop complex biological organisms that continue to work in an extremely hostile environment.

While we consider our systems to be complex, how might one compare a 747 jet with the complexity of an ant, of a 2 year old child, let alone the human nervous system, the human immune system? As technology moves towards nano- and quantum-devices the current issues relating to complexity will appear trivial. How might we design systems with such parallelism, such complexity? How will we ensure that they continue to function correctly over long periods of time, and in unpredictable environments?

The Journey

How can we “evolve” systems of the complexity we will be dealing with produced by technology in 10-20 years? How can we “grow” systems high-reliability designs? How can we build systems that can learn from, and adapt to, their environment in a way that improves their performance, that can become immune to attacks, both internal and external, that can learn to use all of the resources available to them.

What is the effect of “design by evolution” on silicon systems whose microscopic computing paradigm is itself biologically-inspired? What is the interaction between evolutionary processes and the natural imperfections in non-digital chips? How can evolutionary processes capture functionality from such an imperfect computing substrate that conventional design cannot? In particular, when the silicon system is itself adaptive and can “learn”, what is the optimal interaction between “design by evolution” and subsequent adaptation for specific purpose? Natural systems use both methods: how can silicon computation or its successors benefit?