



Contents lists available at ScienceDirect

Journal of Computational Science

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



From digital hype to analogue reality: Universal simulation beyond the quantum and exascale eras

Peter V. Coveney^{a,b,*}, Roger R. Highfield^c

^a Centre for Computational Science, University College London, Gordon Street, London, WC1H 0AJ, UK

^b Institute for Informatics, Science Park 904, University of Amsterdam, 1098 XH, Amsterdam, Netherlands

^c Science Museum, Exhibition Road, London, SW7 2DD, UK

ARTICLE INFO

Article history:

Received 10 January 2020

Accepted 3 March 2020

Available online xxx

Keywords:

Computer simulation

Digital computing

Analogue computing

Quantum computing

Exascale computing

Machine learning

Artificial intelligence

ABSTRACT

Many believe that the future of innovation lies in simulation. However, as computers are becoming ever more powerful, so does the hyperbole used to discuss their potential in modelling across a vast range of domains, from subatomic physics to chemistry, climate science, epidemiology, economics and cosmology. As we are about to enter the era of quantum and exascale computing, machine learning and artificial intelligence have entered the field in a significant way. In this article we give a brief history of simulation, discuss how machine learning can be more powerful if underpinned by deeper mechanistic understanding, outline the potential of exascale and quantum computing, highlight the limits of digital computing – classical and quantum – and distinguish rhetoric from reality in assessing the future of modelling and simulation, when we believe analogue computing will play an increasingly important role.

© 2020 Published by Elsevier B.V.

1. Introduction

We first glimpsed the potential of computers to model the world eight decades ago, in 1936, when Alan Turing devised a hypothetical computing machine while studying the foundations of mathematics [1]. Today, we use computers to describe the turbulence [2] and dynamics [3,4] of air, water and more complex fluids [5], to understand the electronic states of molecules and the kinetics of chemical reactions, for the discovery and understanding of advanced materials, to predict weather and future climate change [6], refugee migration [7], drug design and personalised medical treatments, the creation of virtual organs [8] and, we anticipate, virtual humans too.

In many of these applications, the computer is programmed to solve partial differential equations that are bereft of analytical solutions; and, of course, they can also be used to describe discrete systems, such as lattice gases and other models of fluids, gene regulatory networks with small numbers of molecules [9], agent-based simulations in ecology [10], economics [11] and epidemiology, population dynamics [12], and so on.

Extraordinary progress has been made in recent decades towards the next performance barrier which should be transcended in 2021 with the first exascale computers capable of at least one exaflop [13]. Digital computation is also poised to escape the confines of Moore's law [14] with the advent of quantum computing [15], leading to feats of modelling and simulation far beyond those achievable by classical computing, or so we are led to believe. A 1000-qubit device would theoretically handle more simultaneous calculations than there are particles in the known universe. It is important not to get carried away, however; today, the biggest quantum computers boast only a few tens of qubits.

Recently, new approaches predicated on machine learning (ML) and artificial intelligence (AI)—terms which are often used interchangeably and synonymously in association with “big data”—have become prominent in tackling a range of complex problems and are sometimes regarded as unbounded in terms of the scope of their domains of application. All this has created the widespread expectation among the general public that we can effortlessly use computers to create virtual worlds across a range of domains, from cosmic associations of galaxies stretching over one hundred million light-years to the mesoscale that is most directly accessible to our senses, and from the molecular machines in our cells down to structures within the heart of an atom and inside the particles that comprise its nucleus.

* Corresponding author at: Centre for Computational Science, University College London, Gordon Street, London, WC1H 0AJ, UK.

E-mail address: p.v.coveney@uva.nl (P.V. Coveney).

To simulate different levels of reality, oceans of electrons within myriad microchips, organised within vast numbers of cores that reside inside thousands of nodes in supercomputers are manipulated by billions of tiny switches that flicker on and off billions of times every second. Will this burgeoning virtual environment ever become rich and textured enough to create a faithful replica of our own universe? Some go much further, and even speculate that the cosmos itself arises from interactions between energy and information [16]. Could the universe itself be a quantum computer [17]?

From a consideration of the limits of what is computable, which we discuss further below, we conclude that, although we have come a long way, digital computers – whether classical or quantum – are more restricted in their potential than many realise. Their power resides in their ability to produce vast quantities of numerical data, which lends them an aura of invincibility, yet our article will draw attention to cases where such numerical output can be wrong. Indeed, as we move into the exascale and quantum eras and start to discern what lies beyond them for modelling, there is considerable potential for surpassing the limitations of digital descriptions by falling back on a form of computation which dates back millennia: analogue computing.

2. Dirac's dream

All models and computer simulations rely on theory, representing yet another example of “the unreasonable effectiveness of mathematics”, as Eugene Wigner observed in 1959 [18]. But theory alone is not enough. Three decades earlier, Paul Dirac had remarked that with the general theory of quantum mechanics almost complete, “the underlying physical laws necessary for the mathematical theory of a large part of physics and the whole of chemistry are thus completely known” [19]. However, he added the critical qualification that “the difficulty is only that the exact application of these laws leads to equations much too complicated to be soluble.”

Analytical calculations become intractable relatively quickly, as shown by Henri Poincaré in 1889 with the three-body problem: even the motion of the Sun, Moon and Earth is a non-integrable system, where analytical methods cannot yield an exact solution. Dirac's dream of, for example, replacing chemistry experiments with theory is rendered stillborn since, when mixing two substances together to trigger a reaction under typical laboratory conditions, chemists typically blend of the order of 10^{24} molecules, give or take a few.

We can express this difficulty in terms of algorithmic ‘complexity’, the amount of time required by a Turing machine to solve a problem with a given algorithm. This is a useful lens through which to view the challenge of modelling. The algorithms used to describe computable problems can be divided into two classes, based on the length of time it takes a Turing machine to find the solution to a problem, as a function of some quantity N that measures its size. When problems are polynomial in N (that is, an algebraic power of N), they are said to be in class P and tractable – the length of time required to crack them does not become unbounded as the size of the problem increases, although for high values of N , it can still take an age.

But a class of problem rapidly becomes intractable when the time required to solve it increases in an exponential fashion. These problems, which are not solvable in polynomial time, are said to be in the class NP. Probably the most famous example is the ‘travelling salesman problem’ but there are many other optimisation problems for which no efficient solution algorithm has been found, such as satisfiability and graph-covering problems.

Another way to view the problem of complexity is to consider the extent to which behaviour can be algorithmically ‘compressed’. Compression, developed by Andrey Kolmogorov and Gregory

Chaitin, expresses complexity as the size of the smallest program which computes it or provides a complete description of it. By this reasoning, if a property of a system is algorithmically incompressible, the most compact way of simulating its behaviour is to observe the process itself.

3. Big data need big theory

Even though we do not understand intelligence, and we cannot define it in a scientifically reliable manner [20], “artificial intelligence” has recently returned to the public arena in a prominent manner following a forty-year “AI winter”, notably with the resurgence of artificial neural networks and other forms of ML [21]. In this, the age of so-called “big data”, researchers are increasingly relying on ML to make sense of the world, ironically often without recourse to any form of explanatory insight. Today, there is a pervasive and growing belief that problems of arbitrary complexity are tractable using ML, and in any domain.

The enormous interest in applying ML and AI has proved so intoxicating that some researchers claim that we can dispense with conventional approaches to science, relying on big data to “teach us” how the world works [22,23]. Does this mean the end of theory and the scientific method?

If that were possible, it would eliminate the need for mechanistic, physics based, modelling and simulation, in favour of acquiring and analysing vast amounts of data. This is particularly seductive in areas of biology and medicine, where there remains a substantial lack of fundamental understanding of processes. AI methods should indeed be able to automate diagnostics, notably through pattern recognition [24], and some treatment methods, based on historical data and existing medical treatments [25]. But there is a limit to what can be achieved and, in medicine, regulatory authorities will always require mechanistic explanations before they can approve new drugs and other forms of treatments to make sense of their action and damaging side-effects. On average, each new “once-size-fits-all drug” today costs billions, takes over ten years to bring to market and – at best – works for around 50 % of the population, roughly speaking (see, for example, studies of antidepressants [27], neuraminidase inhibitors [28], and migraine [29]) – most common medications have small-to-medium effects [26] and for reasons that are well known [30]. Yet we live in the post-genomic era, when sequencing an individual's DNA is accurate and affordable.

Ultimately, one would like doctors to be able to make actionable predictions based on data, such as a genotype, computer models and simulations – by this we mean timely predictions which give doctors the opportunity to improve or cure the condition of a patient, whether in picking one antimicrobial drug in preference to another when confronted with a severe infection, or selecting the best approach for risky life-saving surgery.

Currently, the kind of theory used in biology and medicine tends to offer *post hoc*, i.e. retrospective, explanations of the kind seen in dietary advice and financial markets—for which past performance is no guarantee of future results—rather than scientific ones in their most powerful sense, where predictions are made that, when subsequently tested, are shown to be correct or their failures explained. Indeed, no medical treatment will pass muster with regulatory authorities (such as the FDA in USA and the EMA in Europe) if based simply on the predictions of a ML algorithm. It is essential that an underlying mechanistic explanation is provided.

There is a need for more theory in biology and medicine because no matter the ‘depth’ of neural networks, and the sophistication of data-driven methods, they merely fit curves to existing data. They may be capable of interpolation between data in a training domain provided these relationships are smoothly varying – something which cannot be guaranteed. However, effective extrap-

olation beyond this domain is fraught because these relationships do not capture the structural characteristics – the underlying mechanism – of the system [22,23]. Complex models commonly overfit data, such that they do not perform well on new compared with existing data [31]. When ML is applied in ignorance of fundamental laws of nature, it is likely to deliver unreliable answers [23].

Even using a rigorous predictive statistical framework—often noticeably lacking from ML studies—characterizing average behaviour from big data will not deliver “personalized medicine” for two reasons. The first is that it is an inference-based approach, so it depends on assigning common properties to “related” data sets, but ultimately every patient is different. The second shortcoming arises from the fact that the amount of data is, and almost certainly always will remain, tiny relative to the complexity of biological systems, which are heterogeneous in space and time. Indeed, the larger the data set acquired, the more problems will beset black box methods such as ML: the ratio of false to true correlations increases rapidly with the size of the data set, making it difficult to ascertain which ones are meaningful and which are not without mechanistic explanations [22]. Sauro Succi and Peter Coveney have explored the many further problems that could arise if and when we have too much data to play with [23].

Today the practice of science is neither strict empiricism (data without reason) nor strict rationalism (reason without data); rather, it is an integration of both with respect to experimental design, predictive modelling and verification. Viewed this way, many big data approaches are pre-Baconian in that they are unconstrained by experimental design and place undue emphasis on blind data gathering without even *post hoc* explanations. Devoid of structural explanations, the correlations they reveal are not necessarily evidence of dependency. To move beyond this pre-Baconian thinking, we need more reliance on theory to produce more reliable predictive models.

Where ML methods are at their most powerful and reliable is when they are used in conjunction with mechanistic understanding, which focuses their use on acquisition of data which is properly targeted to clarify complex but known or relationships anticipated on sound theoretical grounds. Because of issues such as nonlinearity, non-locality and hyperdimensions – which are encountered frequently in multi-scale modelling of complex systems, for example – big data should and will ultimately be used to complement and enhance theory, rather than render it obsolete. The field of big data is hugely important but it will not realise its full potential to generate useful knowledge without big theory [22].

4. Classical computational chemistry

Some of the greatest challenges to modelling and simulation can be found in chemistry. Despite all the information gathered about hydrogen and oxygen down the years in centuries of experiments, and developments made in the wake of the quantum revolution, we still cannot convincingly predict the boiling point of water from the “bottom up”, that is from the atomic properties of hydrogen and oxygen.

Even so, much important progress has been made over the past half century, for instance harnessing quantum mechanical and Newtonian laws of motion – see below – to examine how drugs bind with target proteins to aid the development of pharmaceuticals. There is an urgent need for accelerated methods of drug development: the traditional model is broken – it costs \$2.6 billion to develop a new drug, with gestation times of a decade or longer [32], and no longer fit for purpose in our post-genomic era, where we all know that individuals respond to drugs in different ways because their genomes are different. We need a new kind of precision medicine which at least stratifies the human population into different groups which all respond in known ways to a given drug.

Computer based modelling offers a powerful way to accelerate drug development. Ligand binding models are used to identify the binding sites of receptors but it is well known that tiny changes in the structure of these small organic molecules can sometimes produce very large changes in the strength of binding – so-called “affinity cliffs” [33]. This is one of the reasons why the use of ML remains largely unsuccessful here. To correctly capture such variations using these approaches would necessitate inordinately large quantities of data, which are unlikely ever to be forthcoming [23].

To compute how candidate drug molecules interact with a target protein, ensemble-based methods can be used, where collections of individual Newtonian molecular dynamics are initialised randomly by assigning velocities to all atoms based on the Maxwell-Boltzmann distribution. Because the behaviour of these systems is chaotic, and thus exquisitely sensitive to initial conditions, such ensemble methods reduce the uncertainty of individual “one-off” simulations, making such modelling much more accurate and reproducible [34].

By using an ensemble approach and classical dynamics, the binding affinity, or free energy of binding of a ligand to a protein, can be computed with a high degree of accuracy and precision [35]. Critically, such an approach helps to render the methods reproducible. The results tally with experimental data, which is as important for drug discovery as it is essential to the scientific method; moreover, it makes the calculations actionable. These binding affinities can be calculated using high-performance computing within a few hours, thereby ranking all available drugs in order of their effectiveness for a given person, a development that heralds the development of truly personalized medicine [36–38]. As a consequence, it is now feasible to use simulations to provide timely information about an individual’s mutations responsible for acquired resistance, for example that interfere with the binding of breast cancer drugs (such as Tamoxifen and Raloxifen).

Ensemble-based methods help boost the reproducibility of computational science. When it comes to molecular dynamics simulations, we can run ensembles of 50 different simulations to quantify the uncertainty, for instance when calculating binding free energies. Although this puts yet more demands on the size of computers needed for our models, one can run them concurrently on many of today’s supercomputers and will certainly be able to do so on exascale machines.

Aside from the use of ensemble methods, we need reassurance in other ways if we are to trust our models – we seek confidence that simulation codes are solving the correct governing equations (validation), are being used correctly (verification), with a comprehensive estimate of the uncertainties in the result (uncertainty quantification), where verification, validation and uncertainty quantification are collectively known as VVUQ. However, though this need is widely acknowledged, it is far from being universally implemented, not least because of difficulty obtaining sufficient computing power. One of us (PVC) has helped to develop EasyVVUQ to make state of the art VVUQ algorithms available to a wide range of computational scientists [39]. However, many research findings in computational science are not reproducible for other reasons, such as lack of access to the same data sets and the expense of independently gathering these data [40], so reproducibility will continue to be a challenge.

5. Quantum chemistry

Although researchers have long sought methods to model bonds between the atoms in molecules to calculate molecular properties and the interplay between molecules, the applications within chemistry had to wait until the 1960s when computers came into general use for solving these equations and quantum chem-

istry emerged, marking an important step towards realising Dirac's dream.

Walter Kohn simplified the mathematics in descriptions of the bonding between atoms, a prerequisite for many of today's calculations [41]. He showed that it is not necessary to consider the behaviour of individual electrons – it suffices to know the average number located at any one point in space, creating a computationally simpler method called density-functional theory [42]. This approach can, for example, be used to explain how enzymatic reactions occur [43].

In 1998, Kohn shared the Nobel prize with John Pople, who developed the quantum-chemical methodology that facilitated the theoretical study of the behaviour of molecules in chemical reactions, through his GAUSSIAN computer program [44]. At the beginning of the 1990s, Pople was able to incorporate Kohn's density-functional theory into that code's algorithms, opening further possibilities for analysing more complex molecules.

Today, quantum chemistry can predict the equilibrium structures of molecules, transition states, intermediates and reaction paths, along with various molecular properties and intermolecular interactions to study solvent effects, crystal packing and so on [45]. However, it is hobbled by the intractability of most of its algorithms. The fastest and most approximate of these algorithms scale as powers of N of four, or higher, where N means the number of orbitals which are used in the quantum chemical calculation. The greater N , the more electrons in the system, meaning the larger the molecule under study, but also the more accurate the calculations for a molecule of given size. The most accurate of these calculations are so-called coupled cluster methods; in their full glory they scale factorially with N , which is even more demanding of computing power than those that scale exponentially with the size of the problem. For these reasons, despite the enormous acceleration in the speed of computers over the past several decades, it remains infeasible to calculate the behaviour of large molecules and assemblies of molecules of most relevance to real world chemistry, physics, biology and medicine.

However, a glimpse of the solution to this problem came as long ago as 1981, when Feynman delivered his seminal lecture 'Simulating Physics with Computers' [46]. As he put it, somewhat trenchantly: "Nature isn't classical, dammit, and if you want to make a simulation of nature, you'd better make it quantum mechanical." Many researchers now place their faith in quantum computing, discussed below, to render these calculations tractable within the coming decade or more.

6. Multiscale modelling

In the discovery of advanced materials, similar challenges to those seen in drug development present themselves. It can typically take twenty years and many billions of dollars to design a new material for an aerospace component and see it implemented in an aircraft. The quest is still dominated by trial and error, and extensive experimental testing.

When using modelling to predict the behaviour of materials, there is an additional complication to consider. The breaking of a single chemical bond, by electron rearrangement, can cause the fracture of an entire aircraft wing. That is an example of a multiscale process working to devastating effect. The challenge to modelling is how to convolve many scales – from the precise chemical combination of the atoms and their electronic structure to descriptions of matter which are oblivious to the existence of atoms and molecules.

To address a diverse range of problems, from modelling how breaking bonds lead to cracks, and how the movement of ions along molecular pores affect the performance of organs, one has to start at the quantum mechanical level, then lace this description into

tiered levels on longer length and time scales, typically including atomic and molecular dynamics, a mesoscale domain and, finally, the macroscopic scale, where matter is represented as a continuum (using finite element methods). Combining these theories to bridge different levels of description is not easy [47]. One reason is that the theories used in each domain (quantum, classical, continuum) are not necessarily compatible – for example some have an arrow of time while others are indifferent to time's direction [48].

To describe condensed phases of matter rather than isolated molecules in the gas phase, Martin Karplus [49], Michael Levitt and Arieh Warshel in the 1970s developed hybrid methods – density-functional theories – that blend the best of classical and quantum methods [50]. They could draw on the facility of classical physics, where calculations were relatively simple enough to model large molecules, together with quantum insights into chemical reactions. By making Newton's classical physics work alongside quantum physics they could, for example, simulate how a drug couples to its target protein by performing quantum theoretical calculations on those atoms in the target that interact with the drug, simulating the rest of the protein using less demanding classical physics [51]. Because this methodology has been widely used in organic chemistry and biochemistry, along with heterogeneous catalysis and theoretical calculation of the spectrum of dissolved molecules, they shared the Nobel prize in chemistry in 2013.

There are many more recent examples of multiscale modelling. Studies by one of us (Peter Coveney) have focused on simulating one family of composites in which clays have been combined with synthetic polymers (such as nylon) to produce superior properties [52,53]. We also investigated how graphene influences the properties of epoxy resins [54]. Through our work on the latter problem we seek to reliably design 2D nanocomposite materials through multiscale materials modelling. Reassuringly, quantum simulations of material properties performed by different researchers and with different software are both reproducible and able to produce identical results [55].

These efforts complement those at Los Alamos, where another team is using first principles methods to guide, enhance and refine a variant of the materials genome initiative [56]. They do this by applying statistical methods to direct data collection based on both experimental measurements and electronic density functional theory for the discovery of new materials [57,58]. In the long run, these novel approaches will cut the time to roll out new materials from laboratory to the production line from decades to years, months and, one fervently hopes, days.

ML can complement multiscale modelling, if used judiciously. With the help of ML, it is possible to create robust predictive models [59] such that, where ML reveals a correlation, multiscale modelling can test whether this correlation is causal [60]. If other challenges can be overcome (the creation of robust predictive mechanistic models from sparse data, ways to manage poorly-posed problems and so on) the use of ML to produce less compute-intensive surrogate models will help multiscale modelling to become central to our understanding of complex systems, such as those found in advanced materials, biology and medicine [59].

Multiscale computing is now widespread but, although the field is gaining in maturity, there is a need for further developments to make the most of future exascale and quantum computers. Uncertainty quantification is largely an open subject in this context, but one which may be facilitated by the advent of the exascale.

7. Exascale computing

Within the next two years we expect to see the first exascale computers. What sort of challenges will we then be able to address that are currently out of reach?

A glimpse of the bleeding edge of modelling comes from the winners of the ACM Gordon Bell prize [61] for innovations in supercomputing. This year, the prize went to a team from the Swiss Federal Institute of Technology Zurich for “A Data-Centric Approach to Extreme-Scale *Ab Initio* Dissipative Quantum Transport Simulations”. They used Piz Daint at the Swiss National Supercomputing Centre and Summit at Oak Ridge National Laboratory, US, to better understand the thermal properties of transistors which would, appropriately enough, help manage heat generation and dissipation as computer architecture shrinks [62].

Last year, a team affiliated with the Oak Ridge National Laboratory was recognized for their paper “Attacking the Opioid Epidemic: Determining the Epistatic and Pleiotropic Genetic Architectures for Chronic Pain and Opioid Addiction,” along with a second affiliated with the Lawrence Berkeley National Laboratory for “Exascale Deep Learning for Climate Analytics.” Both reported new performance hikes through running their applications on Summit at the Oak Ridge Leadership Computing Facility. In 2017, a simulation carried out on China’s Sunway TaihuLight created 3D visualizations relating to a devastating earthquake that occurred in Tangshan, China in 1976. Although earthquake prediction is an inexact and emerging research area, the use of supercomputers this way may lead to better prediction and preparedness – in this case, to inform engineering standards for buildings in seismically active zones.

Examples of forerunners of exascale machines include Summit, mentioned above, which has an IBM Power 9 architecture with 6 NVIDIA V100 GPUs per node and a peak performance of around 200 petaflops. Another, Fugaku (Japanese for Mount Fuji), the successor to the K supercomputer being developed by RIKEN and Fujitsu in Kobe, Japan, will have ARM AF64FX accelerators and is scheduled to go into service in 2021, while Aurora, an Intel machine with Intel’s own GPUs, should be installed at DoE’s Argonne National Laboratory near Chicago in 2021 and is intended to hit the exascale in the second half of that year.

Exascale machines are unprecedented for reasons other than raw speed. Their architectures will be far more complicated than previous supercomputers, displaying a level of heterogeneity which makes it challenging for us to programme them effectively. That is why the notion of co-design is so important at this scale: programmers need access to the detailed hardware layout of proposed machines in order to work with manufacturers and systemware developers to ensure that exascale machines can be used effectively for scientific research as soon as they roll off the production line.

The increase in speed achieved by these machines will come from massive parallelism, where performance is boosted through the aggregation of the speed of all the cores, accelerators, and nodes in the machine via ultrafast proprietary networks. Data management is an increasing challenge, along with power demand [64]. To curb the incessant increase in power required to run these vast machines, individual cores are not generally increasing in performance; if anything, they are decreasing in clock speeds. To put it colloquially, they are becoming “fatter rather than faster” [63]. While they can accelerate many tasks, others are harder to speed up significantly, not least because transistors are already tightly packed on computing chips. One can liken the performance limitations of modern supercomputers to the limitation on the gestation of a human foetus; it takes nine months and cannot be sped up, though multiple babies can be born concurrently on that time scale across a human population.

Significant thought has been put into determining what forms of computing will be best suited to running on exascale machines. With Alfons Hoekstra, Simon Portegies Zwart, and Bastien Chopard, David Coster, Peter Coveney has laid out a description of the kinds of computational algorithms and “patterns” best suited for deployment on exascale architectures [47]. While there will certainly be

scope for a limited number of monolithic applications that can run across the full production partition, there will be relatively few such cases; more generally, we expect to see the further rise of multiscale, multi-physics codes that stitch together numerous single scale models to effectively bridge the time and space scales of the most important scientific and engineering challenges of the day.

One of the more extraordinary efforts set to come into its own in exascale machines is dedicated to the creation of virtual humans, also called digital twins [65]. There are multiscale models of the skeletal and musculoskeletal systems [66]. Efforts are also under way to develop virtual livers [67], lungs [68] and kidneys [69]. In particular, great progress has been made since the 1960s in modelling the heart [70], flows in the aorta [71] and to integrate knowledge at the tissue, cellular, and molecular levels. A team led by one of us (Peter Coveney) is currently simulating the blood circulation from the hips down in Yoon-sun, based on high-resolution ($0.1 \times 0.1 \times 0.2$ mm) colour cross-sections from a frozen cadaver of a 26-year-old South Korean woman. The simulation runs on up to one half of SuperMUC-NG in Garching (due to its then systemware, so far it has only been possible to run on half of its 311,000 cores, although the first benchmarks across the entire production partition were acquired in late November 2019). The simulation of 3D macroscopic blood flow on a full human scale is now becoming possible [72].

In the exascale era, it will be possible to combine multiple models, such as the human heart and blood circulation, to describe the cardiovascular system at the scale of a body. There are even greater ambitions. Functioning computer models of a brain or part of a brain, from rodent to human, are now being created c.f. EU Human Brain Project [73], Blue Brain [74] and equivalents elsewhere, e.g. in the USA. Here, however, the hype has sometimes overtaken reality. Some dream of creating computers that can realise artificial forms of consciousness [75], a concept that is even more elusive than that of intelligence. Others speculate about transferring the contents their own brains onto a computer [76], or even that natural intelligence inhabits a computer simulation [77]. Given the limitations of digital computers, discussed later, these ambitions seem far-fetched.

8. Quantum computing

New horizons beckon with the rise of quantum computing. It purports to offer the means, at some ill-defined but supposedly not-too-distant future, to perform calculations in milliseconds that would take classical supercomputers the age of the universe to do.

In the early 1980s, extending the work of Turing half a century earlier, Richard Feynman [46], Paul Benioff and Yuri Manin [78] provided the groundwork for quantum computing. In 1985, David Deutsch proposed a universal computer based on quantum physics, which would have calculating powers that Turing’s classical computer (even in theory) could not simulate.

Quantum computing remains a relatively immature discipline and progress is being achieved despite the ‘measurement problem’, that is, the lack of deep understanding of the measurement process: there is no generally accepted theory of how (or whether) wave function collapse occurs and what causes it (it lies outside the agreed theory of quantum mechanics) to convert a probabilistic quantum quantity into a classical certainty. Paradoxically, a great deal of effort is being expended to maintain quantum coherence without knowing how to prevent collapse.

Though it is unclear exactly what problems quantum computers will be most effective at solving most agree, as suggested earlier, that quantum chemistry should be one of its main applications [79]. The central expectation of many in quantum computing is that we will “crack” the computational intractability of electronic

structure calculations which we discussed earlier using quantum computers – “the quantum simulation of quantum chemistry”. Various attempts have already been made to use quantum computers (or quantum computer simulators) to solve chemistry problems, such as the calculation of lithium hydride’s and other molecular ground-state energies [80], which suggest that accurate ground-state energy calculations should scale linearly with molecular size rather than exponentially, as in classical computing.

Quantum computers have thus far been used to calculate the ground-state energy for hydrogen [81] and molecules of increasing size, up to water and BeH₂ [82], and there is reason to think that quantum simulations of classically intractable molecules will be viable. [83]. The dream is that one day researchers will be able to use quantum simulations to help in the design entirely new molecules for use as medical drugs. But for quantum computation to become widely used, algorithmic improvements will be needed [84]. Currently, much hope is being placed on the use of variational quantum eigensolvers, which are hybrid schemes combining quantum and classical computational steps, but the field is still in flux and most effort is of a purely theoretical nature. Work by one of us (Peter Coveney), in collaboration with colleagues at Imperial College London and Tufts University, focuses on how quantum algorithms for quantum chemistry can be optimised so as to maximise the use of available quantum computational resources, which are likely to be limited in the near future [85].

9. Coping with chaos

We have already alluded to the care that must be taken with modelling and simulating complex systems that are sensitive to chaos, the subtle mathematical idea that simple equations can generate apparently random behaviour. The field of weather forecasting was one of the first to appreciate the need for caution because of the ‘butterfly effect’, where a butterfly flapping its wings in London can, in principle, cause a subsequent hurricane in the Philippines because the system – in this case the Earth’s atmosphere – is so delicate [48]. Models have to be kept on track by correcting them with satellite and other sensor-based data, else they would drift away from the true behaviour of the weather.

There is, however, a more fundamental problem that may stymie progress in simulation because of the digital nature of modern computers, whether classical or quantum. Digital computers often handle four billion rational numbers that range from plus to minus infinity, the so-called ‘single-precision IEEE floating-point numbers’, which refers to a technical standard for floating-point arithmetic established by the Institute of Electrical and Electronics Engineers in the 1950s; they also frequently use double precision floating-point numbers, while half precision has become commonplace of late in the running of machine learning algorithms.

However, it is seldom realised by computer and computational scientists, let alone day to day users of these digital machines, that the way these numbers are distributed is highly nonuniform. There are as many IEEE floating point numbers between 0 and 1 (where there are one billion of them) as there are from 1 to infinity, and as many between 0.25 and 0.5 as between 0.5 and 1.0.

Plainly, digital computers only use a very small subset of the rational numbers; and there are infinitely more irrational than rational numbers, which are ignored by all digital computers. The IEEE floating point numbers are all dyadic fractions, that is numbers whose denominators are powers of two. That is a poor representation even of the rational numbers. Recent work by one of us (Peter Coveney) in collaboration with Bruce Boghosian and Hongyan Wang at Tufts University, demonstrates that there are major errors in the computer based prediction of the behaviour of arguably the simplest of chaotic dynamical systems, the generalised

Bernoulli map, regardless of the precision of the floating point numbers used [86]. As such, it is a newly discovered pathology of the IEEE floating point number system. Earlier work also points in the same direction [87]. It is remarkable that the simulation community remains ignorant of these problems in the use of digital computers. These findings strongly suggest that the digital simulation of all chaotic systems, found in models used to predict weather, climate, molecular dynamics, chemical reactions, fusion energy and much more, contain sizeable errors [87]. By the same token, the use of data from these chaotic simulations to train ML algorithms will be subject to learning artefacts. For example, Pathak et al. claim that machine learning can be used for “model-free” prediction of the time evolution of chaotic dynamical systems from data alone [88]. This is illustrated using data generated from the solution of a known chaotic model called the Kuramoto-Sivashinsky equation, suggesting that a low prediction error is obtained for about 8 Lyapunov times. Since this model and the data generated from it is likely to contain significant errors arising from the aforementioned pathology, reproducing such errors with an ML algorithm is not per se compelling evidence that model-free prediction works.

Although the so-called “next generation arithmetic” of John Gustafson, based on entities he calls unums and posits [89], offers more flexibility than the IEEE floating point numbers, it cannot mitigate this problem. Research is now underway by Boghosian and Coveney to find alternative approaches that might render such problems computable on digital computers.

Among possible solutions, one that seems guaranteed to succeed is analogue computing, an older and more powerful idea, able to handle the numerical continuum of reality in a way that digital computers can only approximate. While digital computation has seemingly rendered nearly all forms of analogue computation obsolete since the 1950s, various novel approaches are now under development. For example, an analogue approach to the simulation of quantum chemistry problems that relies on the combination of ultracold atoms in optical lattices and cavity quantum electrodynamics, opens up the possibility of analogue quantum simulations of the electronic structure of molecules [90].

10. Universe in a computer

Finally, let’s consider the ultimate simulation – could we, one day, simulate our entire universe in a computer? Many hurdles face this extraordinary and hubristic idea: the probabilistic and non-deterministic nature of quantum theory; our incomplete knowledge of the cosmos; our inability to capture the richness of the universe using digital computers; along with deeper philosophical concerns, such as to how such a computer, being part of the universe, must also simulate itself simulating itself in an infinite recursion. The anthropocentric idea that the universe itself is a computer, which rests on the assumption that it operates in the same way as we solve problems, runs into difficulty because of quantum phenomena and the measurement problem [91].

In the wake of the first claim [15] of ‘quantum supremacy’ one can also expect quantum computing to gain increasing attention. Even so, as we gaze toward the long-term prospects for modelling, many of the basic premises of current computing rest on somewhat shaky foundations. In the medium term, less emphasis should be placed on the world of ones and zeros and more on analogue computers, both quantum and classical.

Important progress will be made by complementing increasingly powerful digital computers with classical and quantum analogue machines. In 1941, Claude Shannon provided a theoretical model, the general-purpose analogue computer (GPAC) [92]. Four decades later, Pour-El & Richards [93] and Lipshitz and Rubel

[94] showed that we can solve any ordinary differential equation with this type of device, given appropriate boundary conditions as part of the input. There is also a fully quantum version of analogue computing, known as continuous-variable quantum computing (CVQC), first described by Lloyd and Braunstein [95]. Carver Mead's concept of neuromorphic computing is alive today [96], particularly within the Human Brain project, albeit only one of the two approaches being pursued along these lines is truly analogue in nature [97,98,99].

We may encounter new limitations: the results returned from calculations performed on analogue computers could, at the limit, turn out to be dependent on the nature of the matter (polymers, molecules and atoms) of which they are comprised, just as digital computers are limited by the nature of the digits that they manipulate. Even so, it should be pointed out that the most advanced computational device that we know of happens to be analogue: the human brain. Humans are able to correctly assign the truth of Gödelian statements which cannot be divined by digital computers, underpinning Roger Penrose's famous argument against the ability of digital computers to produce true artificial intelligence [100]. Moreover, as supercomputers get more energy hungry (the peak energy consumption of Summit is around 13 MW) it is humbling to remember that the human brain requires around 20 W to operate.

Beyond the quantum and exascale computer eras, the future of modelling and simulation may well come to rely more on analogue computing than has hitherto been expected.

Authorship conformation form

All authors have participated in (a) conception and design, or analysis and interpretation of the data; (b) drafting the article or revising it critically for important intellectual content; and (c) approval of the final version.

This manuscript has not been submitted to, nor is under review at, another journal or other publishing venue.

The authors have no affiliation with any organization with a direct or indirect financial interest in the subject matter discussed in the manuscript.

Declaration of Competing Interest

None.

Acknowledgments

P.V.C. is grateful for funding from the UK EPSRC for the UKCOMES High-End Computing Consortium (EP/R029598/1), from MRC for a Medical Bioinformatics grant (MR/L016311/1), the European Commission for the CompBioMed, CompBioMed2 and VECMA grants (numbers 675451, 823712 and 800925 respectively) and special funding from the UCL Provost. He is grateful for many stimulating conversations with Bruce Boghossian, Ed Dougherty, John Gustafson, Alfons Hoekstra, Peter Sloot and Sauro Succi.

References

- [1] A.M. Turing, On Computable Numbers, With an application to the Entscheidungsproblem: corrected version, Proc. Lond. Math. Soc. 2 (43) (1937) 544–546, <http://dx.doi.org/10.1112/plms/s2-43.6.544>.
- [2] U.U. Frisch, Turbulence: The Legacy of A.N. Kolmogorov, Cambridge University Press, Cambridge, 1995.
- [3] S. Succi, The Lattice Boltzmann Equation for Complex States of Flowing Matter, Oxford University Press, Oxford, 2018.
- [4] Succi, The Lattice Boltzmann Equation for Fluid Dynamics and Beyond, Oxford University Press, Oxford, 2001.
- [5] E. Stein, R. de Borst, T.J.R. Hughes, Encyclopedia of Computational Mechanics, Volume 3: Fluids, Wiley, Hoboken, 2004.
- [6] The UK Earth System Modelling Project – Development and Community Release, 2019 (Accessed 03 December 2019) <https://ukesm.ac.uk/>.
- [7] D. Suleimenova, D. Bell, D. Groen, A generalized simulation development approach for predicting refugee destinations, Sci. Rep. 7 (2017) 13377, <http://dx.doi.org/10.1038/s41598-017-13828-9>.
- [8] D. Noble, From the Hodgkin-Huxley axon to the virtual heart, J. Physiol. 580 (Pt 1) (2007) 58015–58022, <http://dx.doi.org/10.1113/jphysiol.2006.119370>.
- [9] R. Dehghannasiri, B.J. Yoon, E. Dougherty, Optimal experimental design for gene regulatory networks in the presence of uncertainty, IEEE/ACM Trans. Comput. Biol. Bioinform. 14 (2015) 938–950, <http://dx.doi.org/10.1109/TCBB.2014.2377733>.
- [10] V. Grimm, S.F. Railsback, Individual-Based Modeling and Ecology, Princeton University Press, Princeton, 2005.
- [11] F.C. Billari, T. Fent, A. Prskawetz, J. Scheffran (Eds.), Agent-Based Computational Modelling: Applications in Demography, Social, Economic and Environmental Sciences (Contributions to Economics), Physica-Verlag, Heidelberg, 2006.
- [12] J. Wu, X. Wang, B. Pan, Agent-based simulations of China inbound tourism network, Sci. Rep. 9 (2019) 12325, <http://dx.doi.org/10.1038/s41598-019-48668-2>.
- [13] US Department of Energy and Intel to Deliver First Exascale Supercomputer, 2019 (Accessed 03 December 2019) <https://www.eurekaalert.org/pub-releases/2019-03/dnl-udo031819.php>.
- [14] G.E. Moore, Cramming more components onto integrated circuits, Electronics 38 (8) (1965) 114–117, <http://dx.doi.org/10.1109/N-SSC.2006.4785860>.
- [15] F. Arute, K. Arya, R. Babbush, et al., Quantum supremacy using a programmable superconducting processor, Nature 574 (2019) 505–510, <http://dx.doi.org/10.1038/s41586-019-1666-5>.
- [16] S. Lloyd, Programming the Universe: A Quantum Computer Scientist Takes on the Cosmos, Alfred A Knopf, New York, 2006.
- [17] S. Lloyd, The Universe as Quantum Computer, arXiv:1312.4455v1 [quant-ph], 2013.
- [18] E. Wigner, The unreasonable effectiveness of mathematics in the natural sciences, Commun. Pure Appl. Math. 13 (1) (1960) 1–14.
- [19] P.A.M. Dirac, Quantum mechanics of many-electron systems, Proc. R. Soc. Lond. A 123 (792) (1929) 714–733, <http://dx.doi.org/10.1098/rspa.1929.0094>.
- [20] A. Hampshire, R.R. Highfield, B.L. Parkin, A. Owen, Fractionating human intelligence, Neuron 76 (6) (2012) 1225–1237, <http://dx.doi.org/10.1016/j.neuron.2012.06.022>.
- [21] A. Krizhevsky, I. Sutskever, G.E. Hinton, ImageNet classification with deep convolutional neural networks, in: F. Pereira, C.J.C. Burges, L. Bottou, K.Q. Weinberger (Eds.), Advances in Neural Information Processing Systems, Vol. 25, Curran Associates Inc., Red Hook, 2012, pp. 1097–1105.
- [22] P.V. Coveney, E. Dougherty, R. Highfield, Big data need big theory too, Philos. Trans. R. Soc. A 374 (2016) 2080, <http://dx.doi.org/10.1098/rsta.2016.0153>.
- [23] S. Succi, P.V. Coveney, Big data: the end of the scientific method? Philos. Trans. R. Soc. A 377 (2019), <http://dx.doi.org/10.1098/rsta.2018.0145>.
- [24] S.M. McKinney, M. Sieniek, V. Godbole, et al., International evaluation of an AI system for breast cancer screening, Nature 577 (7788) (2020) 89–94, <http://dx.doi.org/10.1038/s41586-019-1799-6>.
- [25] S. Jamal, S. Goyal, A. Shanker, A. Grover, Predicting neurological adverse drug reactions based on biological, chemical and phenotypic properties of drugs using machine learning models, Sci. Rep. 7 (1) (2017) 872, <http://dx.doi.org/10.1038/s41598-017-00908-z>.
- [26] S. Leucht, B. Helfer, G. Gartlehner, J.M. Davis, How effective are common medications: a perspective based on meta-analyses of major drugs, BMC Med. 13 (2015) 253, <http://dx.doi.org/10.1186/s12916-015-0494-1>.
- [27] E. Turner, A. Matthews, E. Linardatos, R. Tell, R. Rosenthal, Selective publication of antidepressant trials and its influence on apparent efficacy, N. Engl. J. Med. 358 (2008) 252–260, <http://dx.doi.org/10.1056/NEJMsa065779>.
- [28] T. Jefferson, M.A. Jones, P. Doshi, C.B. Del Mar, R. Hama, M.J. Thompson, E.A. Spencer, I. Onakpoya, K.R. Mahtani, D. Nunan, J. Howick, C.J. Heneghan, Neuraminidase inhibitors for preventing and treating influenza in adults and children, Cochrane Database Syst. Rev. 4 (2014) CD008965, <http://dx.doi.org/10.1002/14651858.CD008965.pub4>.
- [29] S. Kam-Hansen, M. Jakubowski, J. Kelley, I. Kirsch, D. Hoaglin, T. Kaptchuk, R. Burstein, Altered placebo and drug labeling changes the outcome of episodic migraine attacks, Sci. Transl. Med. 6 (218) (2014) 218ra5, <http://dx.doi.org/10.1126/scitranslmed.3006175>.
- [30] B. Goldacre, Bad Pharma: How Drug Companies Mislead Doctors and Harm Patients, Fourth Estate, London, 2012.
- [31] D.M. Hawkins, The problem of overfitting, J. Chem. Inf. Comput. Sci. 44 (1) (2004) 1–12, <http://dx.doi.org/10.1021/ci0342472>.
- [32] J.A. DiMasi, H.G. Grabowski, R.A. Hansen, Innovation in the pharmaceutical industry: new estimates of R&D costs, J. Health Econ. 47 (2016) 20–33, <http://dx.doi.org/10.1016/j.jhealeco.2016.01.012>.
- [33] L. Pérez-Benito, N. Casajuana-Martin, M. Jimenez-Roses, H. van Vlijmen, G. Tresadern, J. Chem. Theory Comput. 15 (3) (2019) 1884–1895, <http://dx.doi.org/10.1021/acs.jctc.8b01290>.
- [34] S. Wan, G. Tresadern, L. Perez-Benito, H. van Vlijmen, P.V. Coveney, Accuracy and precision of alchemical relative free energy predictions with and without replica-exchange, Adv. Theory Simul. (2019) 1900195, <http://dx.doi.org/10.1002/adts.201900195>.

- [35] P.V. Coveney, S. Wan, On the calculation of equilibrium thermodynamic properties from molecular dynamics, *Phys. Chem. Chem. Phys.* 18 (2016) 30236–302340, <http://dx.doi.org/10.1039/C6CP02349E>.
- [36] D. Wright, B. Hall, O. Kenway, S. Jha, P.V. Coveney, Computing clinically relevant binding free energies of HIV-1 protease inhibitors, *J. Chem. Theory Comput.* 10 (3) (2014) 1228–1241, <http://dx.doi.org/10.1021/ct4007037>.
- [37] S.K. Sadiq, M.D. Mazzeo, S.J. Zasada, S. Manos, I. Stoica, C.V. Gale, S.J. Watson, P. Kellam, S. Brew, P.V. Coveney, Patient-specific simulation as a basis for clinical decision-making, *Philos. Trans. R. Soc. A* 366 (1878) (2008) 3199–3219, <http://dx.doi.org/10.1098/rsta.2008.0100>.
- [38] S. Wan, D. Kumar, V. Ilyin, U. Al Homsri, G. Sher, K.R.A. Knuth, P.V. Coveney, From Genome to Personalised Medicine: Cancer Treatment and Discovery of Novel Variants in Qatar, Preprint, 2019.
- [39] D.W. Wright, R.A. Richardson, W. Edeling, J. Lakhili, R.C. Sinclair, V. Jancauskas, D. Suleimenova, B. Bosak, M. Kulczewski, T. Piontek, P. Kopta, I. Circa, H. Arabnejad, O.O. Luk, O. Hoenen, J. Weglarz, D. Crommelin, D. Chroan, P.V. Coveney, Building Confidence in Simulation: Applications of EasyVUUQ, Preprint, 2019.
- [40] R.D. Peng, Reproducible research in computational science, *Science* 334 (6060) (2011) 1226–1227, <http://dx.doi.org/10.1126/science.1213847>.
- [41] W. Kohn, Nobel Lecture, Electronic structure of matter—wave functions and density functionals, *Rev. Mod. Phys.* 71 (5) (1999) 1253–1266, <http://dx.doi.org/10.1103/RevModPhys.71.1253>.
- [42] P. Hohenberg, W. Kohn, Inhomogeneous Electron gas, *Phys. Rev.* 136 (3B) (1964) 864–871, <http://dx.doi.org/10.1103/PhysRev.136B864>.
- [43] G. Roos, P. Geerlings, J. Messens, Enzymatic catalysis: the emerging role of conceptual density functional theory, *J. Phys. Chem. B* 113 (41) (2009) 13465–13475, <http://dx.doi.org/10.1021/jp9034584>.
- [44] J.A. Pople, Theoretical models for chemistry, in: D.W. Smith (Ed.), Proceedings of the Summer Research Conference on Theoretical Chemistry, Energy Structure and Reactivity, John Wiley & Sons, New York, 1973.
- [45] P. von R. Schleyer (Ed.), *Encyclopedia of Computational Chemistry*, John Wiley & Sons, New York, 1998.
- [46] R.P. Feynman, Simulating physics with computers, *Int. J. Theor. Phys.* 21 (6) (1982) 467–488, <http://dx.doi.org/10.1007/BF02650179>.
- [47] A.G. Hoekstra, B. Chopard, D. Coster, S.P. Zwart, P.V. Coveney, Multiscale computing for science and engineering in the era of exascale performance, *Philos. Trans. R. Soc. A* 377 (2019), <http://dx.doi.org/10.1098/rsta.2018.0144>.
- [48] P.V. Coveney, R. Highfield, *The Arrow of Time: The Quest to Solve Time's Greatest Mystery*, Fawcett Columbine, New York, 1992.
- [49] A. Warshel, M. Karplus, Calculation of ground and excited state potential surfaces of conjugated molecules. I. Formulation and Parameterization, *J. Am. Chem. Soc.* 94 (16) (1972) 5612–5625, <http://dx.doi.org/10.1021/ja00771a014>.
- [50] M. Levitt, A. Warshel, Computer simulation of protein folding, *Nature* 253 (1975) 694–698, <http://dx.doi.org/10.1038/253694a0>.
- [51] H.M. Senn, W. Thiel, QM/MM methods for biomolecular systems, *Angew. Chem. Int. Ed. Engl.* 48 (7) (2009) 1198–1229, <http://dx.doi.org/10.1002/anie.200802019>.
- [52] J. Suter, D. Groen, P.V. Coveney, Chemically specific multiscale modeling of clay-polymer nanocomposites reveals intercalation dynamics, tactoid self-assembly and emergent materials properties, *Adv. Mater.* 27 (6) (2015) 966–984, <http://dx.doi.org/10.1002/adma.201403361>.
- [53] J. Suter, D. Groen, P.V. Coveney, Mechanism of exfoliation and prediction of materials properties of clay-polymer nanocomposites from multiscale modeling, *Nano Lett.* 15 (2015) 8108–8113, <http://dx.doi.org/10.1021/acs.nanolett.5b03547>.
- [54] M. Vassaux, R.C. Sinclair, R.A. Richardson, J.L. Suter, P.V. Coveney, The role of graphene in enhancing the material properties of thermosetting polymers, *Adv. Theory Simul.* 2 (5) (2019), 1800168, <http://dx.doi.org/10.1002/adts.201800168>.
- [55] M. Vassaux, R.C. Sinclair, R.A. Richardson, J.L. Suter, P.V. Coveney, Toward high fidelity materials property prediction from multiscale modelling and simulation, *Adv. Theory Simul.* 3 (1) (2019), 21900122, <http://dx.doi.org/10.1002/adts.201900122>.
- [56] NIST, Materials Genome Initiative at NIST, 2019 (Accessed 06 December 2019) <https://mgi.nist.gov/>.
- [57] D. Xue, P.V. Balachandran, J. Hogden, J. Theiler, D. Xue, T. Lookman, Accelerated search for materials with targeted properties by adaptive design, *Nat. Commun.* 7 (2016) 1124, <http://dx.doi.org/10.1038/ncomms11241>.
- [58] T. Lookman, P.V. Balachandran, D. Xue, G. Pilania, T. Shearman, J. Theiler, J.E. Tubermatis, J. Hogden, K. Barros, E. BenNaim, F.J. Alexander, A perspective on materials informatics: state-of-the-art and challenges, in: T. Lookman, F.J. Alexander, K. Rajan (Eds.), *Information Science for Materials Discovery and Design*, Springer Series in Materials Science, vol. 255, Springer, Cham, 2016.
- [59] M. Alber, A.B. Tepole, W.R. Cannon, D. Suvranu, S. Dura-Bernal, K. Garikipati, G. Karniadakis, W.W. Lytton, P. Perdikaris, L. Petzold, E. Kuhl, Integrating machine learning and multiscale modeling—perspectives, challenges, and opportunities in the biological, biomedical, and behavioral sciences, *NPJ Digit. Med.* 2 (2019) 155, <http://dx.doi.org/10.1038/s41746-019-0193-y>.
- [60] W.W. Lytton, J. Arle, G. Bobashev, S. Ji, T.L. Klassen, V.Z. Marmarelis, J. Schwaber, M.A. Sherif, T.D. Sanger, Multiscale modeling in the clinic: diseases of the brain and nervous system, *Brain Inform.* 4 (4) (2017) 219–230, <http://dx.doi.org/10.1007/s40708-017-0067-5>.
- [61] Association for Computing Machinery, ACM Gordon Bell Prize, 2019 (Accessed 08 December 2019) <https://awards.acm.org/bell>.
- [62] A.N. Ziogas, T. Ben-Nun, G.I. Fernández, T. Schneider, M. Luisier, T. Hoefler, A data-centric approach to extreme-scale *ab initio* dissipative quantum transport simulations, SC'19: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, vol. 1 (2019), <http://dx.doi.org/10.1145/3295500.3357156>.
- [63] S. Alowayyed, D. Groen, P.V. Coveney, A.G. Hoekstra, Multiscale computing in the exascale era, *J. Comput. Sci.* 22 (2017) 15–25, <http://dx.doi.org/10.1016/j.jocs.2017.07.004>.
- [64] Defense Advanced Research Projects Agency, Exascale Computing Study: Technology Challenges in Achieving Exascale Systems, in: P. Kogge (Ed.), 2008, Retrieved from <http://www.cse.nd.edu/reports/2008/TR-2008-13.pdf>.
- [65] P.V. Coveney, J.P. Boon, S. Succi, Multiscale modelling at the physics-chemistry-biology interface, *Philos. Trans. R. Soc. A* 374 (2080) (2016) 20160335, <http://dx.doi.org/10.1098/rsta.2016.0335>.
- [66] M. Viceconti, Multiscale Modeling of the Skeletal System, Cambridge University Press, Cambridge, 2012, <http://dx.doi.org/10.1017/CBO9781139049627>.
- [67] Heidelberg Institute for Theoretical Studies, Virtual Liver, 2014 (Accessed 10 December 2019) <https://www.h-its.org/projects/virtual-liver/>.
- [68] K.S. Burrows, A.J. Swan, N.J. Warren, M.H. Tawhai, Towards a virtual lung: multi-scale, multi-physics modelling of the pulmonary system, *Philos. Trans. R. Soc. A* 366 (1879) (2008) 3247–3263, <http://dx.doi.org/10.1098/rsta.2008.0073>.
- [69] P.J. Harris, R. Buyya, X. Chu, T. Kobialka, E. Kazmierczak, R. Moss, W. Appelbe, P.J. Hunter, Thomas S. Randall, The virtual kidney: an eScience interface and grid portal, *Philos. Trans. R. Soc. A* 367 (1896) (2009) 2141–2159, <http://dx.doi.org/10.1098/rsta.2008.0291>.
- [70] D. Noble, Modeling the heart, *Physiology* (2004) 191–197, <http://dx.doi.org/10.1152/physiol.00004.2004>.
- [71] M. Artoli, A.G. Hoekstra, P.M.A. Sloot, Mesoscopic simulations of systolic flow in the human abdominal aorta, *J. Biomech.* 39 (5) (2006) 873–884, <http://dx.doi.org/10.1016/j.jbiomech.2005.01.003>.
- [72] J.W.S. McCullough, R.A. Richardson, A. Patronis, R. Halver, R. Marshall, M. Ruefenacht, B.J.N. Wylie, T. Odaker, M. Wiedmann, B. Lloyd, E. Neufeld, G. Sutmann, A. Skjellum, D. Kranzlmüller, P.V. Coveney, Towards Blood Flow in the Virtual Human: Efficient Self-Coupling of HemeLB, Preprint, 2019.
- [73] Human Brain Project, 2019 (Accessed 12 December 2019) <https://www.humanbrainproject.eu/en/>.
- [74] Blue Brain Project, 2019 (Accessed 12 December 2019) <https://www.epfl.ch/research/domains/bluebrain/>.
- [75] J. Reggia, The rise of machine consciousness: studying consciousness with computational models, *Neural Netw.* 44 (2013) 112–131, <http://dx.doi.org/10.1016/j.neunet.2013.03.011>.
- [76] G.M. Martin, Brief proposal on immortality: an interim solution, *Perspect. Biol. Med.* 14 (2) (1971) 339–340, <http://dx.doi.org/10.1353/pbm.1971.0015>.
- [77] N. Bostrom, Are you living in a computer simulation, *Philos. Q.* 53 (211) (2013) 243–255, <http://dx.doi.org/10.1111/1467-9213.00309>.
- [78] Y.I. Manin, Vychislimoe i ne vychislimoe (Computable and Noncomputable), *Soviet Radio*, 1980 (In Russian).
- [79] C. Hempel, C. Maier, J. Romero, J. McClean, T. Monz, H. Shen, P. Jurcevic, B.P. Lanyon, P. Love, R. Babbush, A. Aspuru-Guzik, R. Blatt, C.F. Roos, Quantum chemistry calculations on a trapped-ion quantum simulator, *Phys. Rev. X* 8 (3) (2018), 031022, <http://dx.doi.org/10.1103/PhysRevX.8.031022>.
- [80] A. Aspuru-Guzik, A.D. Dutoi, P.J. Love, M. Head-Gordon, Simulated quantum computation of molecular energies, *Science* 309 (5741) (2005) 1704–1707, <http://dx.doi.org/10.1126/science.1113479>.
- [81] B.P. Lanyon, J.D. Whitfield, G.G. Gillett, M.E. Goggin, M.P. Almeida, I. Kassal, J.D. Biamonte, M. Mohseni, B.J. Powell, M. Barbieri, A. Aspuru-Guzik, A.G. White, Towards quantum chemistry on a quantum computer, *Nat. Chem.* 2 (2010) 106–111, <http://dx.doi.org/10.1038/nchem.483>.
- [82] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J.M. Chow, J.M. Gambetta, Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets, *Nature* 549 (2017) 242–246, <http://dx.doi.org/10.1038/nature23879>.
- [83] P.J.J. O'Malley, R. Babbush, I.D. Kivlichan, J. Romero, J.R. McClean, R. Barends, J. Kelly, P. Roushan, A. Tranter, N. Ding, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsford, A.G. Fowler, E. Jeffrey, A. Megrant, J.Y. Mutus, C. Neill, C. Quintana, D. Sank, A. Vainsencher, J. Wenner, T.C. White, P.V. Coveney, P.J. Love, H. Neven, A. Aspuru-Guzik, J.M. Martinis, Scalable quantum simulation of molecular energies, *Phys. Rev. X* 6 (3) (2016), 031007, <http://dx.doi.org/10.1103/PhysRevX.6.031007>.
- [84] D. Wecker, B. Bauer, B.K. Clark, M.B. Hastings, M. Troyer, Gate-count estimates for performing quantum chemistry on small quantum computers, *Phys. Rev. A* 90 (2) (2014), 022305, <http://dx.doi.org/10.1103/PhysRevA.90.022305>.
- [85] A. Tranter, P.J. Love, F. Mintert, P.V. Coveney, A comparison of the Bravyi–Kitaev and Jordan–Wigner transformations for the quantum simulation of quantum chemistry, *J. Chem. Theory Comput.* 14 (11) (2018) 5617–5630, <http://dx.doi.org/10.1021/acs.jctc.8b00450>.
- [86] B.M. Boghosian, P.V. Coveney, H. Wang, A new pathology in the simulation of chaotic dynamical systems on digital computers, *Adv. Theory Simul.* 2 (12) (2019), 1900125, <http://dx.doi.org/10.1002/adts.201900125>.

- [87] T.D. Sauer, Shadowing breakdown and large errors in dynamical simulations of physical systems, *Phys. Rev. E* 65 (2002), 036220, <http://dx.doi.org/10.1103/PhysRevE.65.036220>.
- [88] J. Pathak, B. Hunt, M. Girvan, Z. Lu, E. Ott, Model-free prediction of large spatiotemporally chaotic systems from data: a reservoir computing approach, *Phys. Rev. Lett.* 120 (2) (2018), 024102, <http://dx.doi.org/10.1103/PhysRevLett.120.024102>.
- [89] J.L. Gustafson, *The End of Error: Unum Computing*, Chapman & Hall/CRC Computational Science Series, London, 2015.
- [90] J. Argüello-Luengo, A. González-Tudela, T. Shi, P. Zoller, J.I. Cirac, Analogue quantum chemistry simulation, *Nature* 574 (2019) 215–218, <http://dx.doi.org/10.1038/s41586-019-1614-4>.
- [91] K. Wharton, *The Universe is Not a Computer*, arXiv:1211.7081 [quant-ph], 2012.
- [92] C.E. Shannon, Mathematical theory of the differential analyzer, *J. Math. Phys.* 20 (1–4) (1941) 337–354, <http://dx.doi.org/10.1002/sapm1941201337>.
- [93] M.B. Pour-El, J.J. Richards, *Computability in Analysis and Physics*, 1st ed., Springer, Berlin, 1989.
- [94] L. Lipshitz, L.A. Rubel, A differentially algebraic replacement theorem, and analog computability, *Proc. Am. Math. Soc.* 99 (2) (1987) 367–372, <http://dx.doi.org/10.1090/S0002-9939-1987-0870803-1>.
- [95] S. Lloyd, S.L. Braunstein, Quantum computation over continuous variables, *Phys. Rev. Lett.* 82 (1999) 1784–1787, <http://dx.doi.org/10.1103/PhysRevLett.82.1784>.
- [96] K. Roy, A. Jaiswal, P. Panda, Towards spike-based machine intelligence with neuromorphic computing, *Nature* 575 (2019) 607–617, <http://dx.doi.org/10.1038/s41586-019-1677-2>.
- [97] N. Engheta, Circuits with light at nanoscales: optical nanocircuits inspired by metamaterials, *Science* 317 (2009) 1698–1702, <http://dx.doi.org/10.1126/science.1133268>.
- [98] A. Silva, F. Monticone, G. Castaldi, V. Galdi, A. Alu, N. Engheta, Performing mathematical operations with metamaterials, *Science* 343 (2014) 160–163, <http://dx.doi.org/10.1126/science.1242818>.
- [99] N.M. Estakhri, B. Edwards, N. Engheta, Inverse-designed metastructures that solve equations, *Science* 363 (2019) 1333–1338, <http://dx.doi.org/10.1126/science.aaw2498>.
- [100] R. Penrose, *The Emperor's New Mind: Concerning Computers, Minds and The Laws of Physics*, Oxford University Press, Oxford, 1989.

Peter V. Coveney holds a chair in Physical Chemistry and Professor of Applied High Performance Computing in the Institute of Informatics at the University of Amsterdam, and Professor Adjunct at Yale University School of Medicine (USA). He is Director of the Centre for Computational Science and an Honorary Professor in Computer (both at UCL). He is a Member of the Academia Europaea.

Roger R. Highfield is the Science Director of the Science Museum Group. He is also Visiting Professor of Public Engagement at the University of Oxford and UCL, and a member of the UK Medical Research Council.