THE POWER OF SIMULATION IN THE SCIENCES

Computer modelling has all but conquered science—can we use it to tackle the complexity of the brain?

From the author:
"Modelling and simulation complements the traditional scientific method"

IN BRIEF

- Digital modelling and simulation (M&S) is used to test and predict the behaviour of systems in many scientific and industrial fields, from weather to medicine and space.
- A scientific phenomenon can be unpacked by first abstracting it through mathematical modelling, implementing it through simulation, and viewing the results through visualization.
- M&S is driven by improvements in computing performance predicted by Moore's law. It has become ubiquitous in science and industry after initial military applications in World War II.
- Complex non-linear systems can be described in terms of their emergent properties. Human engagement with M&S through visualization is important for discovery and knowledge creation.
- Simulating the human brain and the whole Earth are two grand challenges that require integration across fields as well as new frontiers of M&S.

By Bob Bishop

Founder and president of the ICES Foundation and the former chairman and CEO of Silicon Graphics, Inc. He is also a member of the HBP Advisory Board.

Modelling and simulation, or M&S, has become indispensable for the modern-day scientist and innovator. The rise of these methods parallels the history of computing itself, and is now being driven by the confluence of internet technology, sensor networks, digital media, social networking, citizen science, and big data—but above all by Moore's law. How did these factors coalesce to make M&S the power tool of 21st-century scientific exploration and knowledge building, and are there any limits to how far simulation can take us?

M&S has become a singular concept and its two component terms are frequently used interchangeably. Unlike older mock-up and analogue techniques, most modelling and simulation is now digital—that is, concepts, systems or devices are recreated on an abstract mathematical level with the help of computers to define their components, interconnections and functions.

AN EXPONENTIAL TRAJECTORY

The key to the power of M&S is Moore's law, a simple rule of thumb that describes the ongoing pace of semiconductor miniaturization. According to the law, the performance of digital computers doubles—at half the price—every two years, a multiplicative effect that generates remarkable price-performance improvements of one thousand every 10 years, one million every 20 years, one billion every 30 years, and so on.

Such an extraordinary pace of productivity improvement must surely come to an end at some point in time—especially as semiconductor linewidths scale down to single-digit nanometres, the size of a single atom. Gordon Moore made his eponymous observation in 1965, correctly anticipating the next 50 years of the semiconductor industry. But he could not have foreseen at that time that the next-generation computers will most likely be built from a mix of carbon nanotubes, graphene and photonic devices and will move photons as well as electrons around, thus extending the "law" for some time to come. Despite a slow-down in the exponential trend, Moore's law may stretch for decades yet thanks to quantum computing built from quantum bits or qubits.

The power of modelling and simulation comes from Moore's law

Even now, there is a way to engineer around the soon-to-be-reached single-digit nanometre semiconductor linewidth barrier, and that is through parallelization. Our most powerful machines are already massively parallel, in that billions of transistors are first fabricated onto semiconductor wafers, and then cut into thousands of processor cores, or elementary processing units. The five fastest machines on the planet today currently have more than 500,000 cores grouped in parallel—first into nodes and then into clusters—that yield performance levels of multiple petaflops (10^{15} or one quadrillion floating point operations per second). This fragmentation, however, imposes an enormous challenge for software to cohesively coordinate arrays of processor cores into a single programming task—at the operating system level, at the application programming level, as well as at the middleware level in between. The problem will surely be exacerbated further when processor core counts in a single system grow to millions, and ultimately to billions by the end of this decade.

Get the specs of the top supercomputer.

New programming languages and methodologies are rapidly emerging to deal with the complexity of large-scale massively parallel systems, although it may be several years before reliable standards evolve. A further complication arises in the programming environment with hybrid hardware architectures that mix standard processor cores with dedicated accelerator cores, as in so-called CPU-GPU configurations. Of course, all such cores are accessing memory at much the same time, moving data in and out, so there are serious memory bandwidth and contention issues to deal with as well—not to speak of the high consumption of electric power in moving data around, which is generally more energy-intensive than computation itself.

But stepping back from these underlying technological matters, the question is how to apply such a fount of power to the task of modelling and simulating real-world biogeophysical phenomena.

WAR, WEATHER, THE WORLD?

Digital computing actually first took off through simulation and modelling work done during World War II—encryption and decryption of messages, calculation of missile trajectories and the routing of trans-Atlantic convoys. The single-processor machines of the day had a Von Neumann architecture not much different from the single-processor cores being built today, except they were made of electro-mechanical devices and radio tubes rather than semiconductor materials. Weather calculations also figured prominently in this early work, ultimately leading to numerical weather forecasting techniques and the climate modelling that we know today. Since World War II, M&S has fanned out into myriad applications, tackling issues in engineering, vibration analysis, fluid and aerodynamics, nuclear, particle and astrophysics, oceanography, seismology, and even cosmology. More recently, modelling and simulation has made great inroads in molecular, chemical, and pharmaceutical design, genomics and systems biology, and even surgical simulation.

> Modelling lets you scope out all possible scenarios cheaply and safely

Beyond science, the industrial products that we use today are mostly designed, tested and ultimately fabricated using the output from modelling and simulation done on digital computers. How can such a wide range of research questions and industrial processes be addressed by the same tools? The answer is rather simple. Not only is M&S quick and cheap; the underlying processes and dynamics can be accurately described by simple, reductionist, deterministic, linear mathematical models, whether you are dealing with cars or canned food. Thanks to Moore's law and our ever-increasing processing power, such models have become tractable and computable.

We are quickly learning, however, that much of our planet and its environment, as well as the space beyond, work through non-linear feedback mechanisms, intertwining processes that cannot be accurately captured by simple first principles. Complexity and chaos raise their ugly heads and confound the orderly world of linear modelling and simple simulations. One way out of this muddle is to harness more complex maths—using regression analysis and advanced statistical methods to extract correlations of variables and then sorting such correlations by relevance to deduce a hypothesis about causality (Bayesian inference is a common way to build a probability estimate for a hypothesis as new evidence is acquired). Although we may never know the exact causal chain that leads from action to reaction in this world of multi-layered, multivariate, non-linear complexity, we can, in our ignorance, describe the various macro phenomena that manifest as emergent properties and leave it at that.

Indeed, as science plunges headlong into the nanoscale world, M&S will be one of the only tools we have to comprehend the complex inner workings, and emergent properties will increasingly become a commonly understood reference model to explain macro-level manifestations. Strangely enough, mechanisms of this very nature also hold in social sciences and humanities research, where large numbers of complex human interactions can be said to yield properties that emerge from the crowd. Agent-based modelling, which gives each person in a crowd certain personal characteristics and interactive properties, is the specialized technique used to deal with these matters.

THE HUMAN TOUCH IN THE SIMULATED WORLD

Today scientists, engineers and technologists are pondering better ways to cast the physical processes and phenomena of physics, chemistry and biology into new and streamlined models and algorithms to make the work of digital computing faster and cheaper. Quite often such research activities are globally networked over the internet or through dedicated high-speed communication lines so that researchers can instantly share their results and mutually brainstorm their way around roadblocks and problems as they occur. Industry, government and academia now have at their disposal full-on interactive and immersive digital environments.

We have entered an era of co-creation where hardware, software and algorithms all mesh together, and each can optimize or de-optimize the other, as the case may be. What has become apparent is that large-scale massively parallel computing requires special attention to reliability, accessibility and serviceability; how to deal with inevitable component failures, for example. Gracefully degrading a machine in the case of component failure until a full recovery can be made, retracing prior steps if necessary, is an art form of its own.

How can we apply the power of parallelization to modelling real-world phenomena?

Humans are more than just number crunchers in the modelling and simulation scheme—soft human factors like ease-of-use, machine-human interfaces, and the role of intuition and imagination should figure prominently in the acts of exploration and discovery. Humans excel at adaptation, following twists and turns and testing new hypotheses; couple that with machine learning and we are slowly building a compendium of new knowledge.

One of the most exciting aspects of the modelling and simulation world is the human component. When M&S processes can be visualized online in an interactive and immersive manner, human engagement and perception can be exceptionally aroused. Deep engagement arises naturally in a cockpit where the simulation process is online and computationally steered on-the-fly by user interaction and manipulation. This is clearly the reason for the high success of flight simulation training, space mission planning, safety operational training, interactive entertainment and even the world of distributed multi-player video games, many of which are simply online derivatives of off-line simulations from the staid scientific, engineering and industrial space.

M&S—a generic computational tool—can deliver remarkable insight into the design, operation and functionality of the world around us, and can be readily networked for use by scientists, mobile users and citizen scientists worldwide. The tool becomes more powerful year-by-year because Moore's law is its underlying engine; it also benefits economically from the unit volumes of consumerism, since the individual digital processing cores are common to both the professional and consumer markets and decrease in cost with mounting volumes. M&S allows us to pose "what if" questions and to scope out any number of scenarios in a cost-effective manner, in complete safety.

TROUBLE IN PARADISE?

Although it is physically safe to compute, visualize and immerse oneself in a virtual replica of the real world, and to navigate within with great curiosity and abandon, there are some pitfalls to the process that cannot be ignored. Open questions and assumptions have to be confronted for any modelling and simulation exercise to be successful.

First, there many branches of mathematics—which branch best describes

RELATED PROJECTS:

The Illustris Simulation: Realistic Virtual Universe
—The Illustris Project

The ICES Foundation: International Centre for Earth Simulation
—The ICES Foundation

VPH Institute: Building the Virtual Physiological Human
—The VPH Institute

The Human Brain Project: Unifying our understanding of the brain
—The Human Brain Project

NASA's Global Climate Modelling program
—NASA

the physics, chemistry and biology under investigation? Questions about computability can be numerous: what level of data is necessary to commence the calculations, the so-called initial conditions? Do we have these data at hand? What is the quality of the data? Do we have the meta-data that describe how the initial data were collected? How big are the initial data? Are they available locally or remotely? How can we assimilate these initial data into the model itself? Are they in the right units or is conversion required? Can the data be immediately assimilated or is extrapolation to the appropriate grid points of the model necessary? How do we deal with data that are continuously updated, or unstructured, event-driven data, not to mention video and image data?

With an initial data set that has been assimilated, the computation will proceed something like this. With a digital computer (but not an analogue computer), when we compute forward in time the machine does so in discrete incremental time steps. We use numerical methods (a sophisticated branch of mathematics that uses differential equations) to calculate the time evolution of the model, hopefully without incurring great errors in the process. If the phenomenon under investigation varies continuously in time, then this discrete time step computation will only be an approximation of the underlying reality; some form of error will be generated and propagated. Managing this ongoing and growing error within certain bounds is key to the accuracy of the M&S effort, critically determining the ultimate usability of the results. If we are modelling and simulating a truly time-discrete phenomenon, this type of error won't become relevant.

> The exaflood of big data has arrived

Next we deal with any errors in the structure of the model itself, where the fundamental scientific understanding and description of the phenomenon under investigation may be incomplete or mistaken. In this case, averaging over the simulation outcomes from several models will hopefully factor out any shortcomings, a process known as ensemble averaging. If a model has gaps or missing descriptors, parameterizing and substituting certain "tuneable" control variables may help. The act of ensemble averaging can then be applied to these control variables by varying or perturbing their values stochastically and running multiple simulations accordingly.

Modelling and simulation has been applied in many research fields.

Finally, chaos itself is a phenomenon of dynamic, deterministic non-linear systems, and will occur naturally as any such system approaches certain well-identified domains of its operational phase-space (the phase-space charts the trajectory of the system over time and can be visualized).

M&S can thus be subject to serious compounding errors—initialization error, time step error, model structural error, parameterization error, and chaos itself, each one of which can do damage to the credibility of the modelling outcome unless treated with great mathematical integrity. Fortunately, advanced statistical methods can help to manage the error in most cases, but definitely not in all cases. For example, the accuracy of numerical weather forecasting techniques fades beyond the 10-day mark for precisely these reasons.

BRIDGING THE GAPS

Conceptually, there is no limit as to how and where modelling and simulation can be applied. It is one of the few tools that can deal with the exaflood of big data that is now upon us. Computational models can clarify data sets whose sheer size would otherwise place them outside of human understanding. We could go as far to say that nobody really understands a scientific phenomenon until it is modelled, simulated and visualized.

There is still much room for progress, particularly in some areas of multi-disciplinary science where several niche models, each with their specialised units of analysis, must be brought together and integrated into a full holistic picture. Individual models can be coupled together with bridging algorithms, taking their intrinsic timing differences into account, or a fundamentally new integrated monolithic model can be built from the ground up, which may require significant work on the programming front.

Selecting one of these paths will be one of two main challenges for the Human Brain Project in its quest for a unified understanding of the elements that contribute to the functioning and malfunctioning of the human brain. The second challenge strays into the cognitive space of human consciousness and the domain of qualia, or subjective experiential phenomena of colour, sound, taste, smell, touch, feelings, emotions, and memory. Addressing these fundamental mysteries will stretch our modelling and simulation tools for years to come.

The same can be said for modelling the whole Earth system, another grand challenge that is similarly beset by the enormity of integrating hundreds of specialised sciences into unified whole. This task is especially timely as its outcome aligns with the safety, health and well-being of the entire planet. Nature herself has done a wonderful job at making the totality work seamlessly; we scientists can only claim success at splintering her into a thousand silos.

RELATED NEWS:

K Computer runs largest ever ensemble simulation of global weather
—Riken

Models: How computers make predictions
—Student Science

Processing, Visualizing, and Understanding Data from High-Performance Computing
—IEEE

An Evolution Toward a Programmable Universe
—New York Times

In Silico: First Steps towards a computer simulation of the human body
—The Independent

The Great Climate Model
—Forbes

Simulated earth illustrations by Alexander Vostriakov and Melissa Cochrane

MODELLING AND SIMULATION IS USED TO STUDY PHENOMENA ACROSS MANY RESEARCH FIELDS

TOUCH ICONS TO READ MORE

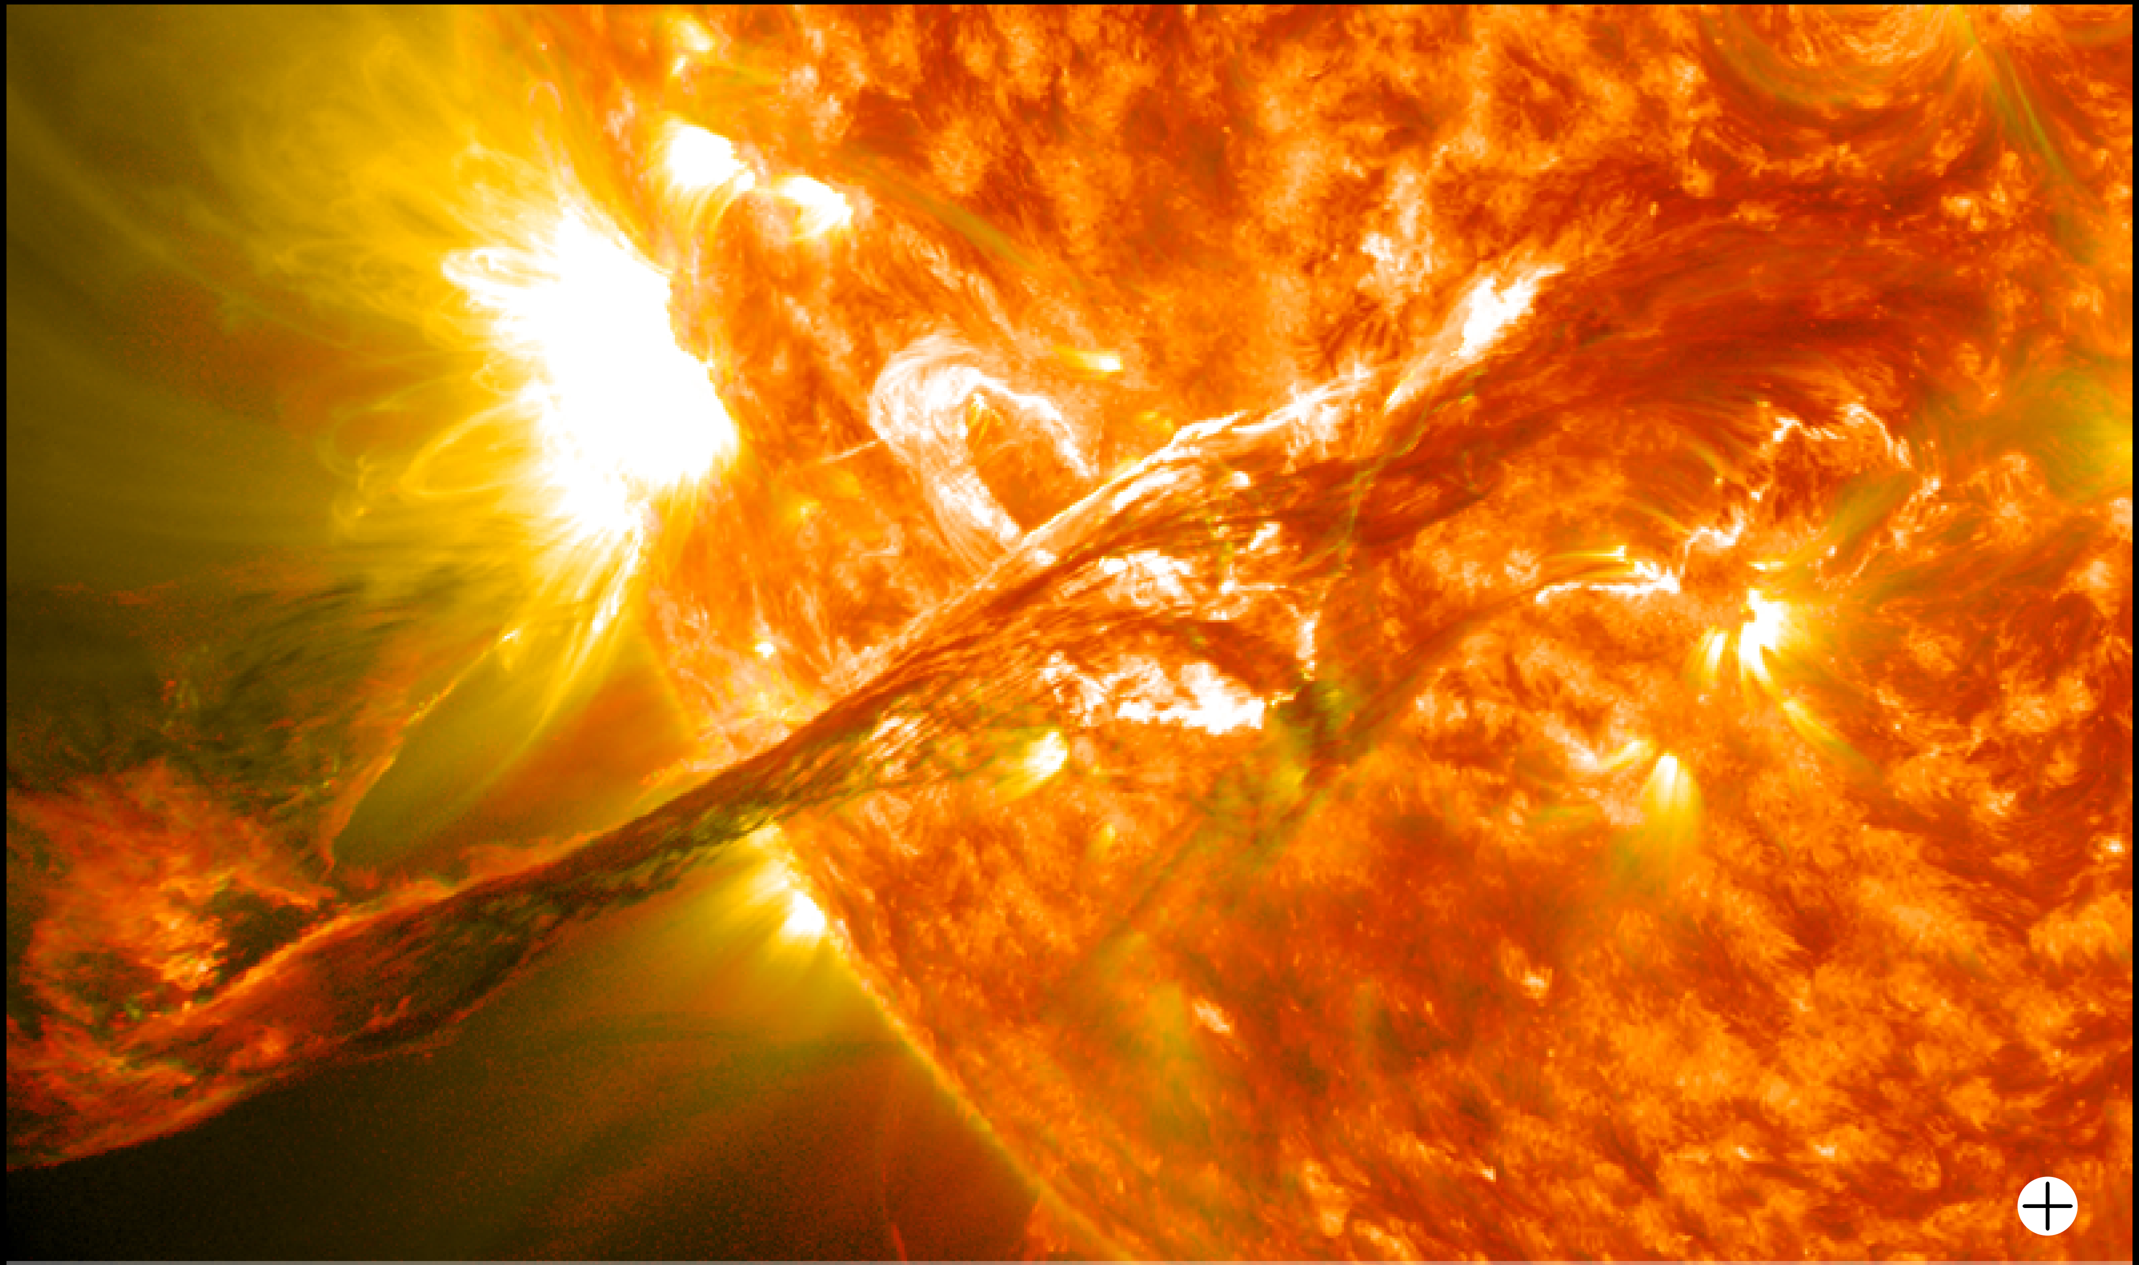

Heliophysics and the impact of the solar wind and coronal mass ejections on Earth