

Integrative Biology: Modelling Heart Attacks with Supercomputers

Trying to understand what kills one third of the UK population is a challenging task which needs the most powerful supercomputers available.

Gernot Plank (University of Graz), Richard Clayton (University of Sheffield), David Boyd (University of Oxford/CCLRC), Ed Vigmond (University of Calgary)

Heart disease and sudden cardiac death

Although the incidence of cardiovascular disease continues to fall, it remains an important cause of death in the industrialised world. Recent data for England and Wales for 2004 [1] indicate that diseases of the circulatory system accounted for around 190,000 deaths, 37% of the total, with almost half of these being directly attributed to ischaemic heart disease. In many of these cases, the lethal event is likely to have been a catastrophic arrhythmia called ventricular fibrillation (VF). Despite its immense social and economic impact, the mechanisms underlying formation of arrhythmias (arrhythmogenesis) are still debated and remedial therapies are, at best, suboptimal. In the case of life-threatening arrhythmias like VF, electrical defibrillation, the application of an electric shock within minutes of VF onset, is now recognized as the only effective therapeutic option.

The heart is an electromechanical pump. During each beat a wave of electrical excitation originates in the heart's pacemaker, the sinus node, located in the right atria (smaller chambers sitting on top of the ventricles). During a normal beat, the cardiac impulse reaches the ventricle via the only electrically excitable connection, the AV node, and subsequently is carried quickly to the working myocardium, acting as a signal for the muscle to contract (Figure 1). During VF, a state of electrical anarchy exists, with normal electrical activity being replaced by rotating waves which stop the heart beating properly. VF is usually associated with underlying heart disease, but the mechanisms that initiate a particular episode of VF are poorly understood.

Despite the critical role that defibrillation therapy plays in saving human life, our understanding of the mechanisms by which electric shocks halt life-threatening arrhythmias remains incomplete. While recent advances in experimental methodology have provided new characterizations of tissue responses to electrical shocks, inquiry into the success and failure of defibrillation is hampered by the inability of current experimental techniques to resolve, with sufficient accuracy, electrical behaviour confined to the depth of the ventricles during and after the shock. Realistic computational models of the heart, sometimes referred to as "virtual heart simulators", are currently the only viable approach to allowing us to observe all parameters of interest at sufficient spatial and temporal resolution.

The Integrative Biology project

Integrative Biology [2] is a second generation e-Science project, funded by EPSRC, which is developing a Grid-based framework for computational biology driven primarily by the need to understand the two largest killers in the UK population, heart disease and cancer. The project is supporting a globally-distributed team of scientists developing multi-scale models which attempt to link behaviour at the molecular and cellular level to the operation of whole organs, thus enabling us to build realistic virtual heart simulators.

Our aspiration is that, as our models become more detailed, they will enable us to understand these diseases sufficiently well that we can develop therapeutic strategies for treating and possibly in the long term preventing them. The current state of the art in modelling the heart is more advanced than for cancer tumours, although developments there are moving rapidly, and our attempts to develop

more detailed heart models are driving development of the computing infrastructure and will stretch the capability of current generation supercomputers such as HPCx.

The Integrative Biology team which we have assembled to tackle this challenge brings in expertise from many disciplines. Scientists from Oxford, Nottingham, Birmingham, Sheffield and Auckland are working with computational scientists from CCLRC, Oxford, UCL, Leeds, Auckland and IBM. Since the project started, partners have joined from Graz (Austria), Utrecht (Netherlands), Tulane, UCSD and UCLA (USA) and Calgary (Canada), and we expect the team to continue to grow.

The computing services we are developing within the project allow researchers to target simulations at the most appropriate computer system depending on the resources and response needed. They provide data and metadata management facilities, using the Storage Resource Broker [3], for looking after the many datasets created in computational experiments, and visualisation tools for examining results and discussing these collaboratively within the consortium. An associated project also being carried out by the Integrative Biology team, funded by the Joint Information Systems Committee, is developing a Virtual Research Environment which is embedding these services into a portal-based framework so they are easily usable by researchers who are not themselves computing experts.

Modelling the heart

VF can be studied experimentally using voltage sensitive fluorescent dyes or electrodes to map the spread of electrical activation in heart tissue but, as we noted above, this approach only yields information about activity on the heart surface. Modelling has an important role, because models of VF in the heart can be used to simulate and study electrical activation patterns that are consistent with what can be observed on the surface of the heart. This approach is similar in principle to studying the sun; it is only possible to observe the sun's surface, yet an understanding of the physics of stellar nuclear reactions allows the mechanisms that produce these surface observations to be understood.

Models of electrical activity in the heart during VF vary widely in the level of detail that is included. To reduce the computing resources required, pilot simulations often use simplified models of anatomy, which may be a 2D sheet of tissue or a 3D slab representing part of the heart wall, and use simplified representations of the electrical excitability of cardiac cells. However, to properly understand the complexity of a problem like VF, we would ideally like to model the heart as a complete organ.

Mathematically, the virtual heart is based on a set of coupled PDEs which are referred to as bidomain equations. In recent years, most cardiac computer models have used either a computationally significantly less demanding simplification (monodomain) which does not capture electrophysiologically important mechanisms, particularly when defibrillation is under study, or rather coarse discretization to keep simulations computationally tractable.

State of the art models are a trade-off capturing either gross anatomy at limited spatial resolution or representing regions of limited size incorporating microscopical detail. Several questions of major electrophysiological importance, however, require both macro- and microscopic structural details and use of the more complete bidomain formulation. Today, the cardiac modelling community is on the brink of implementing simulators which will allow development of virtual hearts for many different species. The use of high performance computing and state of the art numerical methods to deal with the tremendous computational burden imposed by such simulations is key to success in this endeavour.

Understanding ventricular fibrillation

Current research at Sheffield University supported by the Integrative Biology project is focusing on understanding the mechanisms that initiate and sustain VF. Compute-intensive simulations using whole ventricle detailed anatomy and biophysically detailed models of electrical excitability are run on HPCx, and these build on simulations using simplified models that are run on local HPC resources including the White Rose Grid [4]. A key aspect of this work is to relate the findings to clinical practice. Work on modelling the initiation of VF has already yielded information that could be used to identify patients at risk, and this is the basis of a pilot clinical study about to start in collaboration with clinical colleagues at the Northern General Hospital in Sheffield. Work on the mechanisms that sustain VF is also tightly meshed with clinical and experimental studies, and is one component of a wider project focusing on understanding VF in the human heart that also involves the Universities of Auckland, Utrecht, Oxford, and UCL.

The current version of SCAM (the Sheffield Cardiac Arrhythmia Model) is written in C, uses shared memory parallelism, and runs on a single frame of HPCx. The code has been optimised for the HPCx architecture, and ongoing development aims to further exploit the mixed mode parallelism capability of HPCx, so that simulations can be run across large numbers of frames. Results of simulations using SCAM show initiation and development of fibrillation in the ventricles (Figures 2 and 3).

Virtual experiments on HPCx

In a preliminary study at the University of Graz using the Integrative Biology framework, the feasibility of carrying out a “virtual experiment” was tested using HPCx. A computer model of a ventricle, discretized at an average spatial resolution of 200 μm , was simulated immersed in a conductive bath. At the bath boundaries, two plate electrodes were placed next to the anterior and posterior faces of the ventricle (Figure 4A). To test conditions under which an arrhythmia can be induced, a train of 10 pacing pulses of varying basic cycle length was delivered (Figure 4B). After the last pacing pulse, 2 seconds of activity were simulated to examine whether an induced arrhythmia was sustained or self-terminated (Figure 4C).

Performing this virtual experiment involved the solution of an elliptic PDE (862,515 unknowns), a parabolic PDE (547,680 unknowns) and a set of 21 non-linear ODE's, defined at the same grid as the parabolic PDE. Using a temporal discretization step of 8 μs , the solution scheme had to be repeated 500,000 times to complete the experiment. Preliminary simulations carried out on a Dual Opteron desktop computer suggested that execution times would be around 2 months. Using 128 CPUs of HPCx allowed the execution of a single experiment in only 10 hours. The simulations were carried out using CARP, the Cardiac Arrhythmia Research Package.

A subset of this simulation was repeated using different numbers of CPUs to demonstrate the scalability of the method (Figure 5). The overall computational workload is clearly dominated by the elliptic problem (> 95% of the overall workload). The parabolic PDE, solved by a simple forward Euler integration step, showed super-linear scaling. As expected, the ODE solver scaled linearly since the involved variables do not diffuse and thus no communication is required. The dominating elliptic problem scaled well, although the parallel efficiency decreased slightly when going from 64 to 128 CPUs. Taking into account that this problem size is rather small for the high number of CPUs the scaling efficiency is more than satisfying.

These preliminary results suggest that realistic simulations of a human heart including a torso are feasible on the HPCx platform. In such simulations one has to deal with roughly 20-200 million unknowns (20-200 times larger than in this study). Memory usage and execution times will require the

use of more CPUs. It is expected that parallel efficiency will increase significantly thanks to a more favourable ratio between local computational load and communication.

The future - model-based therapy?

We have demonstrated that it is now feasible to carry out virtual experiments using supercomputers. This leads us to expect that different therapeutic approaches may be studied to investigate possible optimizations. Three avenues of research are particularly interesting:

- 1) How can the dynamic behaviour of the cardiac membrane be modified to terminate or even prevent arrhythmia? This could lead to development of pharmacological treatments to determine which pharmacological targets are the most promising ones.
- 2) Can the most effective therapy, namely implantation of a cardioverter/defibrillator, be further optimized? Is it possible to defibrillate in a smarter way to increase the lifetime of the implanted device, to reduce the time to therapy in the case of failed shocks and to minimize tissue damage? For instance, can it be tested how safety and efficacy of the therapy are affected by geometry and location of the defibrillation electrode, by the instant of shock delivery (relative to the cycle of the arrhythmia) and by the waveform of the delivered shock pulse?
- 3) The only curative treatment to cardiac arrhythmias is catheter ablation, a procedure by which the conduction pathways within the heart are modified. This treatment is particularly effective in the atria, but can we also find applications in the ventricles (the main pumping chamber)?

Future computational models of the heart have to incorporate all these ingredients, the atria, the ventricles, the specialized conduction system and probably a torso model as well to allow direct comparisons with clinically measured parameters like the body surface electrocardiogram (ECG). By combining electrical models with mechanical models we could simulate the mechanical action of the heart and the blood flow driven by contraction allowing comparison with further clinically important parameters such as cardiac output and ejection fractions. However, integrating these parts into one comprehensive model will substantially increase the computational burden. Parallel environments like HPCx are the only viable way of executing such simulations by providing the required address space and keeping execution times tractable.

In conclusion

The Integrative Biology project is building a computational framework which will enable scientists to develop increasingly realistic virtual heart simulators without having to become computer experts themselves. The potential benefit to the population as a whole from using simulation to help us understand, treat, and possibly prevent, heart disease is a strong motivation driving this work forward.

Acknowledgements

We wish to thank other members of the Integrative Biology consortium for their input into, and comments on, the work described in this paper, in particular the role of the scientific collaborators in determining the requirements for the project. We also thank Michael Holden at EPCC for help with installing and optimising our codes on HPCx. We acknowledge the financial support of the EPSRC (ref no: GR/S72023/01) and IBM for the Integrative Biology project, and of the Joint Information Systems Committee for the Integrative Biology Virtual Research Environment project. R. Clayton acknowledges support from the British Heart Foundation (PG03/102/15852). G. Plank was supported by the Austrian Science Fund FWF (R21-N04).

References

- [1] <http://www.statistics.gov.uk/>
- [2] <http://www.integrativebiology.ac.uk>
- [3] <http://www.sdsc.edu/srb/>
- [4] <http://www.wrgrid.org.uk>




Figure 1. Normal Activation Sequence of the Ventricles

(A-B) The cardiac impulse propagates from the atria (smaller chambers sitting on top of the ventricles, not shown) via the specialized conduction system into the ventricles. The specialized conduction system starts as a single strand on top (A), bifurcates into a left and right branch (B) and then further bifurcates into a mesh-like structure. The impulse reaches the left ventricular endocardium (the inner surface forming the left cavity) first (C), then the right ventricular endocardium (D-E). Subsequently, activation occurs within the heart's walls until breakthroughs are observed at the epicardium (F). Seen from the epicardial surface, breakthrough sites act as focal activations which join to form an epicardial wavefront (G-H). Finally, the entire ventricles are in excited state (I-J).

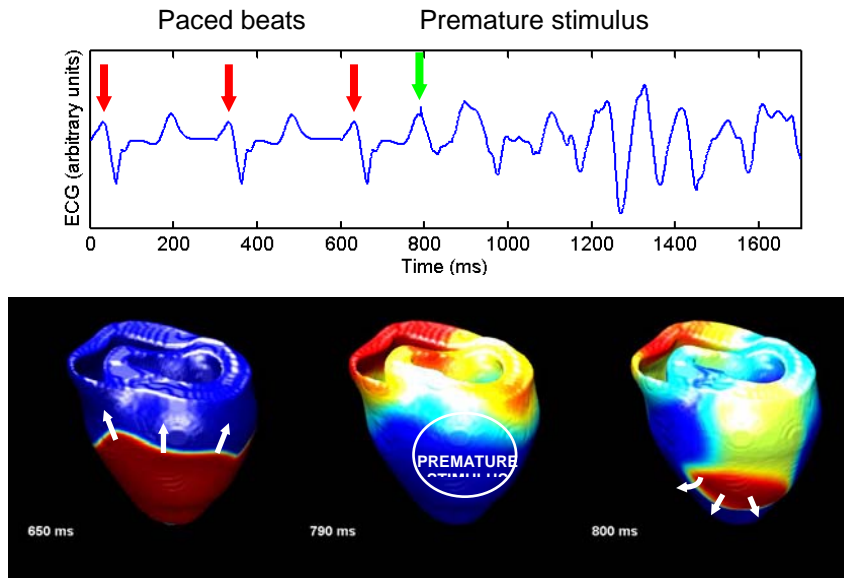


Figure 2. Simulation showing the initiation of fibrillation in the ventricles

Top panel shows simulated electrocardiogram (ECG). Red arrows indicate three stimuli to the apex (bottom) of the heart that result in normal paced beats at intervals of 300ms, and green arrow indicates premature stimulus delivered to the heart wall. The premature stimulus results in the onset of fibrillation, shown by rapid and self-sustained activity in the simulated ECG. Bottom panel shows snapshots of the simulation, where electrical activation is colour coded with blue indicating resting tissue and red indicating active tissue. The first frame (650 ms) shows the propagation of the third paced beat shortly after the stimulus has been applied to the apex (bottom) of the heart. The second frame (790 ms) shows the state of the heart just before the premature stimulus is applied over the region shown. Part of this region has yet to recover, shown by the yellow and light blue regions. The third frame shows the effect of the premature stimulus. The activity resulting from the stimulus can only propagate outwards and downwards because it cannot propagate into areas that are still recovering.

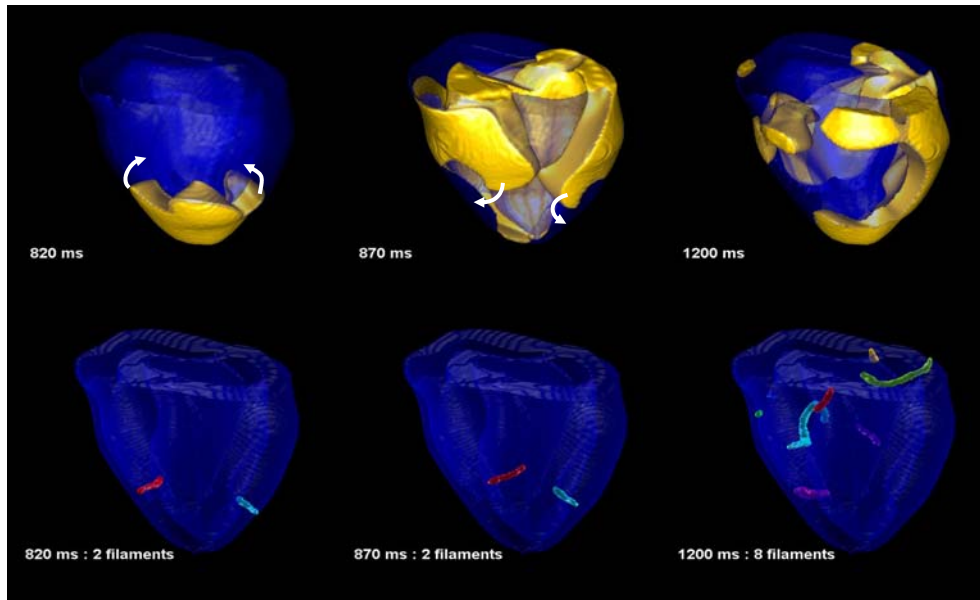


Figure 3. Simulation showing the development of fibrillation in the ventricles

Top panel shows isosurface views, where electrically active regions are enclosed by yellow surfaces. The first frame (820 ms) shows how the downward propagating activity shown in Figure 2 is beginning to curl around as the tissue recovers from the paced beat, and forms a pair of counter-rotating scroll waves. In the second frame (870 ms) these scroll waves have rotated by about 180 degrees, but the activation pattern is unstable, and by the third frame (1200 ms) the initial scroll wave pair has broken up into multiple interacting waves. Bottom panel shows scroll wave filaments – the lines around which the scroll waves rotate. In the first two frames there are two filaments, one for each of the scroll waves. In the third panel (1000 ms) there are 8 filaments, reflecting the more complex activation pattern resulting from the initial instability.

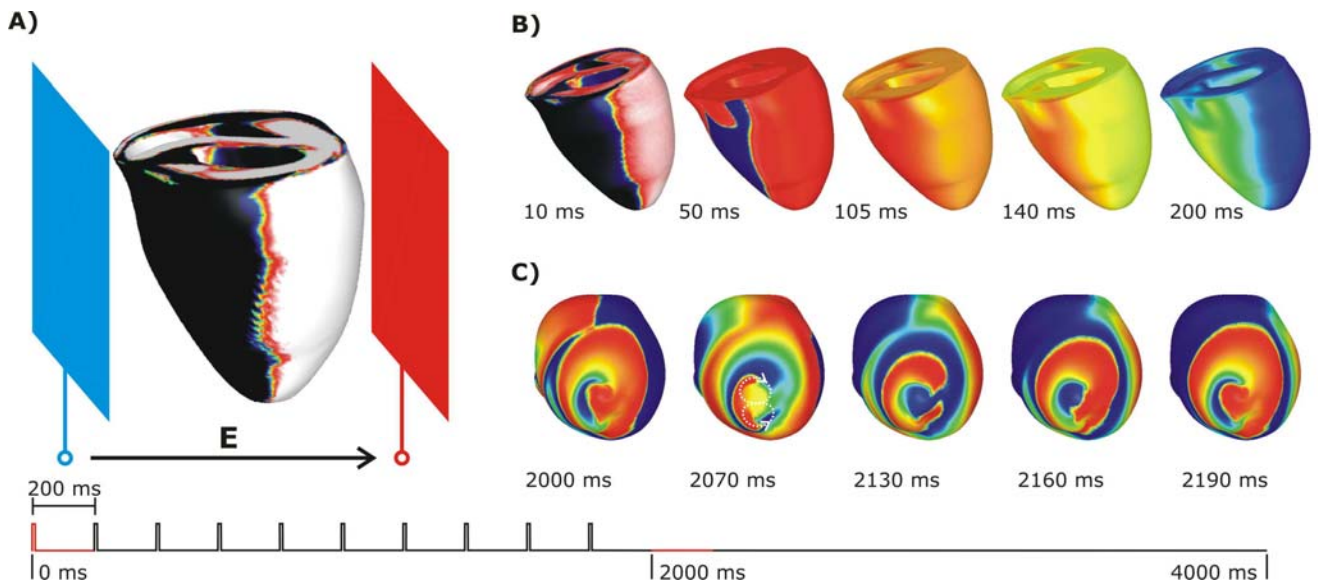


Figure 4. A virtual experiment. Setup for a “virtual experiment” to induce an arrhythmia in the ventricles by applying an electrical pacing protocol: A) The ventricles are immersed in a conductive fluid and placed between two plate electrodes. B) Electrical activation of the ventricles during the stimulation period. C) After the last pacing pulse, a so-called figure-of-eight re-entry (named after the movement of the tips of the wavefronts) ensued and was sustained until the end of the simulation run at 4000 ms.

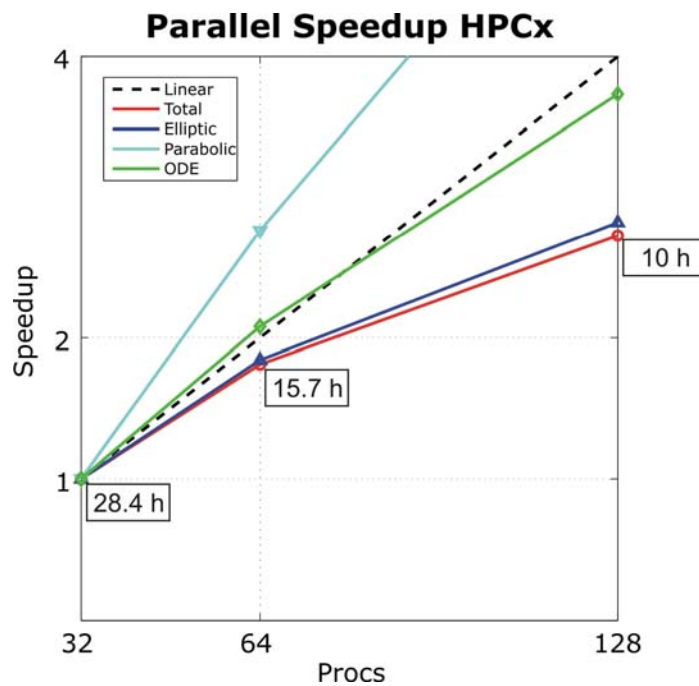


Figure 5. Benchmark results: scaling of different portions of the bidomain computation. As expected, the ODE part scaled linearly (no communication required). The parabolic problem, solved by a simple forward Euler step, basically involved only a matrix-vector product which showed super-linear scaling. Computations were dominated by the elliptic problem which scaled reasonably well, particularly if one takes into account that the problem size is small for the number of CPU's employed in these simulations. Overall execution time in hours is shown as a function of the number of CPU's.