CARP User Meeting

Edward J. Vigmond (University of Calgary),
Gernot Plank (Medical University of Graz)

September 5–7, 2008

1 Overview of the Field

Heart disease is the number one killer in the industrialized world. Modelling cardiac electrical phenomena is attractive since it allows complete control of the system and one has full knowledge of all components in the system, unlike animal experiments. Cardiac bioelectricity is described by the bidomain equations[1] which links current flow in the extracellular space to current flow within cells through current flowing through the cell membranes:

$$\nabla \cdot (\bar{\sigma_i} + \bar{\sigma_e}) \nabla \phi_e = -\nabla \cdot \bar{\sigma_i} \nabla V_m - I_e \tag{1}$$

$$\nabla \cdot \bar{\sigma_i} \nabla V_m = -\nabla \cdot \bar{\sigma_i} \nabla \phi_e + \beta I_m(V_m, \zeta) \tag{2}$$

which is a parabolic and an elliptical equation which depend on the voltage across the cell membrane, V_m, and the extracellular electric potential, ϕ_e. While the media are linear, a highly nonlinear source source term, I_m, links solutions in the two media.

Simulations of cardiac electrophysiology are becoming increasingly sophisticated and more quantitative. This has arisen as a result of faster computing hardware, better imaging and experimental methodology, and advances in scientific computation. However, simulation of a human heart at a near real time performance is still a challenge as the size of the system is currently prohibitive. Solving the bidomain equations is an inherently expensive procedure since the involved space constants are small (some tens of μm up to 1 mm) and the time scales are very fast. At the same time, domain sizes which are sufficiently large to maintain arrhythmias or to observe wavefront propagation free of boundary effects, is large, at the order of centimetres, and the observation periods are long (some hundreds of milliseconds up to minutes). The dynamics of charge transport across the membrane and between intracellular compartments is described by a set of non-linear ODEs. In general, the set of ODEs is quite stiff. This is particularly true for very recent formulations that rely on Markov state models to describe the cellular dynamics. In these models, the intrinsic time scales vary from the sub-microsecond scale for the fastest processes up to hundreds of milliseconds for the slowest processes. The fast onset of the cellular excitation process, referred to as upstroke of the action potential, leads to potential variations over fairly small spatial domains when electrical wavefronts are traversing the heart (within less than 1 mm the transmembrane voltage covers the full physiological range). That is, discretizations of the bidomain equations using anatomically realistic representations of the whole heart together with recent mechanistically realistic models of the cellular dynamics typically lead to systems of millions of degrees of freedom and hundreds of thousands of time steps. Therefore, even when simulations are executed using the most powerful HPC hardware available today, execution times impede fine-grained explorations of the parameter space of interest.

Finally, the heart is an electrically activated mechanical pump. While electrical pathologies have been recognized as lethal and been the subject of much research, the role of electromechanical coupling and mechano-electric feedback has been receiving less attention. For many clinically relevant scenarios, both aspects, electrical activation as well as mechanical contraction, have to be accounted for in computer modelling studies to contribute additional information. Recent therapeutical strategies aim at improving cardiac output in patients suffering from an impaired synchronization of the ventricles by implanting devices which pace the heart to improve synchrony between the two main pumping chambers, the ventricles. This therapeutical approach is referred to as Cardiac Resynchronization Therapy (CRT). Although quite successful with quantifiable survival benefits for patients, the therapy is expensive and the percentage of non-responders is too high. Computer simulations will be helpful to investigate the mechanisms which determine whether CRT improves cardiac output or not. Such simulations can be used to sweep the parameter spaces (electrode locations, timing and delay for delivery of pacing pulses, etc) to determine optimal configurations.

Our software, CARP, is among the fastest and most versatile for simulating cardiac electrophysiology. With a growing user base of about 50 spread throughout several universities including Johns Hopkins and Oxford, a reexamination of the code was necessary, in order to keep at the forefront in terms of computational efficiency as well as flexibility with particular focus on multiphysics simulations.

2 Recent Developments and Open Problems

Current trends in cardiac modelling are aiming at developing simulation environments that enable researchers to perform "in-silico" experiments. It is expected that the increased level of anatomical and functional detail will contribute to further increase the predictive value of computer simulations which will make simulators an indispensable complementary clinical modality over the next few years. A further important trend is to add mechanical movement to electrical heart models. This is a significant advancement and opens new and exciting perspectives for basic research and clinical applications. Finally, in many recent studies not only the physics of the heart is modelled, but also the experimental techniques employed to observe bioelectrical phenomena to better understand how the physical quantities behave and how they appear to behave when being measured with a particular modality. Such approaches have demonstrated that experimentally measured signals suffer from substantial distortions and are related to the actual physical quantity in a non-trivial manner. Such "secondary" simulations include the simulation of the processes underlying the signal acquisition in mapping techniques such as optical mapping, electrical mapping, the computation of magnetocardiograms, etc.

Great challenges are offered by these current trends that cardiac modelling is becoming *multiphysics*, incorporating electrical and mechanical activity and, at a later stage, also fluid-structure interaction, as well as *multiscale*, where subcellular processes are modeled which seem to have an impact at the tissue and organ level. Besides the constant scientific challenge to improve, develop and validate the quantitative descriptions of cardiac phenomena upon which current simulators are built, dealing with the computational burden and the complex software engineering involved in the development of "in-silico" simulators is probably the most difficult obstacle to overcome on the quest to establish cardiac modeling in new application areas where modeling is not an option today. A major open problem is to increase computational efficiency and speed. Although many problems of practical interest can be tackled today using the fastest HPC facilities available, execution times remain to be a limiting factor that impede quick parameter sweeps. To be of relevance in a clinical context, an increase in computational efficiency of 3-4 orders or magnitude is required. Parallel and scalable implementations that take advantage of upcoming PetaFLOPS computing hardware are sought after to achieve these speedups. Optimizing codes for novel hardware is likely to take cardiac modelling a major step forward, nonetheless, numerical techniques including spatiotemporal adaptivity, linear solver and preconditioning techniques, custom-tailored cache optimized ODE solver techniques for extremely stiff sets of ODEs which are required to better describe the cellular physiology may contribute to further reduce the computational burden and/or increase the efficiency of "in-silico" simulators. A major concern is the scalability of numerical methods. Although it can be expected that the available computational power will increase substantially over the next few years, it is unclear to which degree current simulators will benefit. For instance, with the CARP simulator the entire numerical scheme shows excellent scaling up to 128 processors for problem sizes of 1e6 degrees of freedom or more. Using more processors is inefficient, since

scaling efficiency saturates and, when more processors are added a critical number of processors, N_{crit}, will be eventually hit where execution time starts to increase. However, for many applications execution times are too slow even when using $_{crit}$ processors. Using the anatomically most realistic model of the rabbit ventricles available today [6] with 4.3 million degrees of freedom, solving the monodomain equations with 128 processors is about 4000 times slower than real time. This is a hard limit that cannot be overcome with current computing hardware and currently implemented numerical techniques. When considering human heart, the critical number of processors N_{crit} will be much larger, but the same scalability issue will eventually kick in when the problem sizes per partition are becoming sufficiently small leading to a less favourable ratio between local computational load and communication. No matter how powerful PetaFLOPS computers are, current simulators will not benefit and it appears that near real-time performance is not within reach without a shift in paradigm.

Currently, reliable predictions on which hardware and computing paradigms will prevail, are difficult to make. Regarding classical CPUs, further increases in clock frequency cannot be expected, but rather an increase in the number of available cores. A major bottleneck with multicore processors using the CPU is the limited memory bandwidth which prevents data from being transferred sufficiently quickly from memory to the CPU. Accelerator cards, although unlikely to be suited for all relevant aspects of cardiac modeling, bear significant promise to overcome some of these limitations. To our understanding, the most promising technologies include: i) general purpose graphical processing units (GP-GPU), promoted mainly by NVidia and, to a lesser extent, by AMD, ii) cell processors, promoted by IBM, and iii) the Larrabee processor, promoted by Intel.

Porting an "in-silico" simulator as complex as CARP to a novel non-standardized computing platform entails major coding efforts. Although efforts are under way to develop high-level interfaces, designed to hide differences between different accelerator technologies, it can be expected that the technology will not be in place soon enough, or at least not a the level of maturity required to support software development endeavours as complex as "in-silico" multiphysics simulators. Hence, not all new paradigms can be supported and decisions have to be made which platforms to support. Performance charts over the last few years clearly indicate that GPUs increase their performance at a faster rate than CPUs. The software development kit CUDA, provided by NVidia to develop software for their GPUs, is at a reasonably mature stage and facilitates a fairly smooth migration pathway to run entire software packages such as CARP, or at least specific aspects of it, on a GPU. Based on these observations we decided to invest major efforts into developing GPU based codes to replace or complement specific aspects of CARP. While remarkable speedups with GPU based implementations have been achieved for some classes of problems at a fraction of the cost of a supercomputer, they have yet to be applied to solution of the bidomain problems.

Producing a tractable heart simulator will only result from the collaboration of experts at every level, from computer hardware, to computer scientists, to mathematicians, software engineers, experts in mesh generation, physiologists, and clinicians. To succeed in the quest of developing an efficient, flexible general purpose simulator that can be employed to execute "virtual experiments" without any or very little trade-offs in terms of accounting for anatomical or functional details of the heart, a major shift in the academic approach to software engineering is required. Unlike in the past, where fairly simple simulators were developed in a single laboratory within the course of a single PhD thesis that were quite often custom-tailored for a single specific problem, endeavours such as the development of the CARP simulator require a distributed network based approach to reflect the multifaceted challenges where several groups make longterm commitments to further advance the code to ensure the long-term sustainability of the code development process and to steer the development effort into directions that support those applications that are of particular interest to the research community. The development of the CARP simulator started in 2002. The development process has been set up as a distributed effort.

This workshop invited specialists from all disciplines relevant to the development and application of "in-silico" modeling tools with a focus on application to building a virtual human heart. It was a first step towards building a larger community that Further, despite major advancements, reflected in an increase in the number of users by an order of magnitude over the last 3 years, Not all aspects of "in-silico" modelling can be addressed easily within an academic environment and may be in an non-academic setting. Interestingly, in the field of mechanics, most research groups rely on commercial modeling tools which seem to be well suited. This may be a viable option in the field of cardiac modeling as well as soon as commercial modeling packages become available. Industrial partners were invited to discuss and explore future directions.

3 Presentation Highlights

The workshop was organized to cover three main topics which were *Geometric Modelling*, *Numerical Techniques and Novel Computing Paradigms for Large-scale Bidomain Simulations* and *Modelling Applications within an Experimental and Clinical Context*.

3.1 Geometric Modelling

3.2 Numerical Techniques and Novel Computing Paradigms for Large-scale Bidomain Simulations

To address the challenge of reducing execution times two major themes were discussed. First, what are the bottlenecks in current codes and which algorithms are suited to overcome current limitations. Secondly, since it is unlikely that major increases in computational efficiency can be achieved by novel numerical techniques alone, particular attention has to be paid to explore in which way novel computing paradigms can be exploited, which gains in efficiency can be expected and how well these paradigms are suited for cardiac modeling applications.

Both themes cannot be seen in isolation. One of the major themes presented was that algorithms need to be developed which are parallel from the beginning and not as an afterthought. Computers are not getting faster, but we are getting many more cores. Novel numerical approaches that are not well suited to execute efficiently on upcoming computing platforms, are of limited interest. Highly sophisticated spatio-temporal modeling approaches, although of great interest from a numerical point of view, are a candidate to fall into this category and, consequently, were not considered at the workshop.

The computationally most expensive portion of a numerical scheme to solve the bidomain equations is the solution of the elliptic PDE. The fastest known solver for this problem relies on using a multilevel preconditioner for the conjugate gradient (CG) iterative solver. Our group demonstrate the efficiency of this technique for both a geometric variant [4] as well as an algebraic (AMG) variant [3] which is better suited for "in-silico" modeling where complex grids are considered. Dr. Gundolf Haase presented a novel AMG implementation which was implemented by his group as part of their Parallel Toolbox (PT) solver library. Relative to the Boomer-AMG method used in previous studies, the PT implementation is more lightweight and the multilevel strategy is simpler. First benchmarks suggest that the PT implementation, although requiring about twice the number of iterations over Boomer AMG to converge, results in a three-fold increase in terms of execution time. A further advantage of the PT implementation of AMG is the option to compile for execution on a GPU. Dr. Liebmann presented impressive performance results for the AMG-PT solver running on a GPU. He presented results measured on a Quad GPU server costing $4000 to perform 2.5 times faster than 32 traditional CPUs in a cluster interconnected by a low-latency Infiniband network costing $200.000. The other major contributor to the overall workload is the solution of the system of ODEs. The numerical technique implemented in CARP, essentially, a non-standard finite difference technique based on the work of Rush and Larsen[], seems to be sufficiently simple and thus very well suited for being executed efficiently on a CPU. Dr. Leon presented first results measured with a GPU implementation of a very recent rabbit ventricular ionic model [5]. The reported speedups on the GPU were equally impressive. Preliminary results suggest that GPUs are very well suited for solving the two computationally most burdensome components of bidomain simulations, the solution of the elliptic PDE and the set of ODEs.

Dr. Kickenger summarized his meshing mesh method. Using an Octree approach and a dual mesh, he is able to rapidly mesh segmented medical images into computationally usable grids. His software is among the best available commercially, in terms of the size of the meshes he can generate, the time in which he can do it, and the accurate surface representations.

3.3 Modelling Applications within an Experimental and Clinical Context

.

Dr. Albert Kim gave an overview of clinical approaches for the treatment of arrhythmias. He outlined the shortcomings of current approaches and advocated the use of computer simulations to explore more targeted and less destructive therapies.

Dr. Trayanova showcased work performed in her lab which is cutting edge in terms of application to clinical issues. One member of her lab, Dr. Gurev, presented his work on electromechanics. He showed his simulations, which are a first, involving defibrillation and a mechanically contracting heart, based on electrical propagation with mechanoelectric feedback. His work further emphasized the need for including mechanical activity, since it is ultimately a major clinical measurand as well as has a profound influence on cardiac electrical activity through feedback of stretch-sensitive channels.

4 Scientific Progress Made

For the first time, many people working on specific aspects of the cardiac electrophysiology simulator were able to see how their contribution fit into the whole scheme.

5 Outcomes of the Meeting

Several potential collaborations also became apparent. Dr. DiMartino work had concentrated on mechanical aspects only. She will now work with Drs. Vigmond and Comtois to integrate electrical activity with mechanical.

Representatives from the medical device companies were impressed with the level of sophistication that cardiac modelling had reached. There was interest in pursuing licensing of the CARP software so that it could be used for device design.

Drs. Vigmond and Plank were able to have discussions with researchers involved at all levels of the software. Several aspects became apparent after the meeting:

- the software needs to remain flexible and not be tied into any one computing paradigm or numerical technique
- ease of use is paramount as the software becomes more wide spread, with more automatic setting of parameters needed
- in the immediate future, GPU computation will be very attractive. As such, we will need to rethink algorithms to be parallel from the beginning, and match algorithms to the hardware.
- geometrical models will continue to become more detailed as imaging improves, necessitating greater heterogeneity that our software must be able to capture

References

[1] R. Plonsey, “Bioelectric sources arising in excitable fibres (ALZA lecture),” *Ann Biomed Eng.*, vol. 16, no. 5, pp. 519–546, 1988.

[2] G. Pank G, L. Zhou, J.L. Greenstein, S. Cortassa, R.L. Winslow, B. O’Rourke, N.A. Trayanova. “From mitochondrial ion channels to arrhythmias in the heart: computational techniques to bridge the spatio-temporal scales,” *Philos Transact A Math Phys Eng Sci.* vol. 366, no. 1879, pp. 3381–3409, 2008.

[3] G. Plank, M. Liebmann, R. Weber dos Santos, E.J. Vigmond, G. Haase, “Algebraic multigrid preconditioner for the cardiac bidomain model,” *IEEE Trans Biomed Eng.* vol. 54, no. 4, pp. 585-596, 2007.

[4] R. Weber dos Santos, G. Plank, S. Bauer, E.J. Vigmond, “Parallel multigrid preconditioner for the cardiac bidomain model”, *IEEE Trans Biomed Eng.* vol. 51, no. 11, pp. 1960–1968, 2004.

[5] Mahajan A, Shiferaw Y, Sato D, Baher A, Olcese R, Xie LH, Yang MJ, Chen PS, Restrepo JG, Karma A, Garfinkel A, Qu Z, Weiss JN. “A rabbit ventricular action potential model replicating cardiac dynamics at rapid heart rates,” *Biophys J.* vol. 94, no. 2, pp. 392–410, 2007.

[6] G. Plank, R.A.B. Burton, P. Hales, M. Bishop, T. Mansoori, M. Bernabeu, A. Garny, A.J. Prassl, C. Bollensdorff, F. Mason, F. Mahmood, B. Rodriguez, V. Grau, J.E. Schneider, D. Gavaghan, P. Kohl, "Generation of histo-anatomically representative models of the individual heart: tools and application," *Philos Transact A Math Phys Eng Sci.*, under review.