Cray CoE Science Delivery Programme

Computational Support Systems for Guiding Delivery of Cardiac Therapies

G.Plank (Oxford e-Science Research Centre, University of Oxford),
N.A. Trayanova (Department of Biomedical Engineering, John Hopkins University, Baltimore),
Kevin Roy (CoE)

Scientific Background

Ventricular arrhythmias in the setting of myocardial infarction are associated with high risk of sudden cardiac death. While antiarrhythmic medications are used in the treatment of infarct-related ventricular arrhythmias, proarrhythmia can often occur from antiarrhythmic agents, limiting their utility. There is a significant survival benefit of implantable cardioverter-defibrillators (ICD), however, the cost of ICD implantation, in all patients who meet entry criteria, would be extremely expensive to society. Therefore, alternative approaches to terminating infarct-related ventricular arrhythmias have been developed over the last decade. One such approach is catheter based ablation. With catheter ablation, tissue is heated to destroy its ability to generate and conduct electrical signals, and thus terminate the arrhythmia. Catheter ablation in the setting of supraventricular tachycardia is now a safe and effective therapy with high success rates and low levels of complication. Such levels of success have not, however, been achieved in catheter ablation of infarct-related ventricular tachycardia, which is associated, unfortunately, with only a 58% initial success rate and a 71% eventual success following repeated procedures. Additional approaches to improve electrical conduction in partially-viable tissue within and around the infarct scar (and thus terminate arrhythmia) and increase contractility of the region of infarction includes gene therapy approaches and the delivery and engrafting of stem cells. These approaches remain, however, in their infancy, although the expectation for potential therapeutic application in the treatment of patients with myocardial infarction holds high promise.

Catheter ablation of infarct-related ventricular arrhythmias remains a challenging procedure hindered by the complex three-dimensional (3D) nature of the infarction reentrant pathways, which renders them difficult to localize. The low efficacy of the catheter ablation procedure stems from the fact that current clinical voltage and pace mapping techniques to identify the sites for ablation are associated with numerous limitations, including ambiguities in correlating maps with anatomy, as well as possible missed critical sites due to the point-by-point sampling nature of current methods. In addition, the complex 3D nature of the reentrant pathways prevents the correct identification of the optimal sites of ablation on the basis of electrical interrogation of the surfaces by the current mapping systems. Finally, complications such as perforation and emboli can be as high as 8%, likely due to the often prolonged duration of the procedure. There is an urgent need for new methodologies that can result in accurate identification of critical reentrant circuit features and sites for ablation and thereby improve the efficacy

of the therapy. Similar approaches to identify critical reentrant circuit features will also have dramatic effect on the efficacy of the delivery of gene and stem cell therapies for termination of ventricular arrhythmias.

Computational modelling of the heart can address this challenge and thus become an essential tool in the integrated approach to diagnosis and treatment of infarct-related arrhythmias. Today we stand at the threshold of a new era in computational modelling of the heart. Anatomically detailed, tomographically reconstructed models of hearts from various species are being developed, integrating functions from the molecular level to the electromechanical interactions in the intact heart. Such models hold great promise for enhanced interpretation of clinical and physiological measurements, for improving the basic understanding of the mechanisms of cardiac dysfunction in disease conditions, and for prediction of the success of therapy. Our intent is to take advantage of these new developments and to combine them with advanced image processing techniques to bring image-based computational modelling of cardiac (dys)function out of the realm of basic research and into the clinic.

The Application

About half a year ago, the scalability of the CARP code was limited to about 512 cores when considering a realistic test case of simulating a single beat of a human heart. Within a recent DCSE project scalability of the CARP code has been dramatically improved. It could be demonstrated that a single beat of a human heart can be simulated within about 4 minutes (real-time lag factor of ~250), showing quite efficient scalability up to 16k cores (Fig. 1).

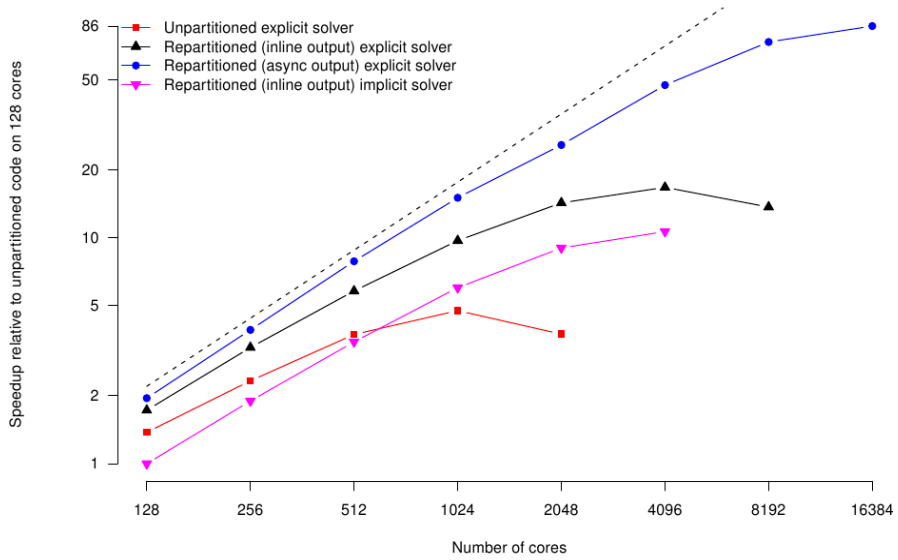

Figure 1: Parallel efficiency and speedup of CARP at different stages during the DCSE project: The original code relied on linear partitioning (referred to as "unpartitioned"). Subsequently, parMetis partitioning was included which provided both better performance and scalability. Finally, asynchronous IO was implemented to further improve performance and scalability particularly for higher number of cores where IO turned out to be a relevant contributor to overall execution time.

Although the improvements clearly open new and exciting perspectives, enabling applications which were not before, some limitations remain. The high parallel efficiency has been achieved by resorting to a cheap explicit solver. Although this reduced the maximum stable time step, since a single explicit solver step was more than one order of magnitude cheaper than one implicit solver step, overall the explicit method was quicker. For most setups this is likely to be the case, but in those cases where even higher spatial resolutions are required the method becomes impractical quickly since stability constraints enforce even smaller time steps. With implicit methods scalability levelled off earlier at around 4k cores (Fig.1).

Experiments Targeted

An essential component in bringing these approaches to the clinic is the development of heart-torso models of each patient to investigate the signature of cardiac therapy delivery on body surface ECG (Fig 2). Addressing the computational challenges in this process constitutes the goal of this project. Successful execution of the proposed studies will constitute a major step towards the integration of computational modelling in the diagnosis and treatment of cardiac disease, and in patient care in general.

The work proposed here will allow the code to be able to tackle the large scale systems required to simulate the whole heart-torso models and we plan to show the feasibility with one real-world application and will eventually become part of the main production workload for this code.

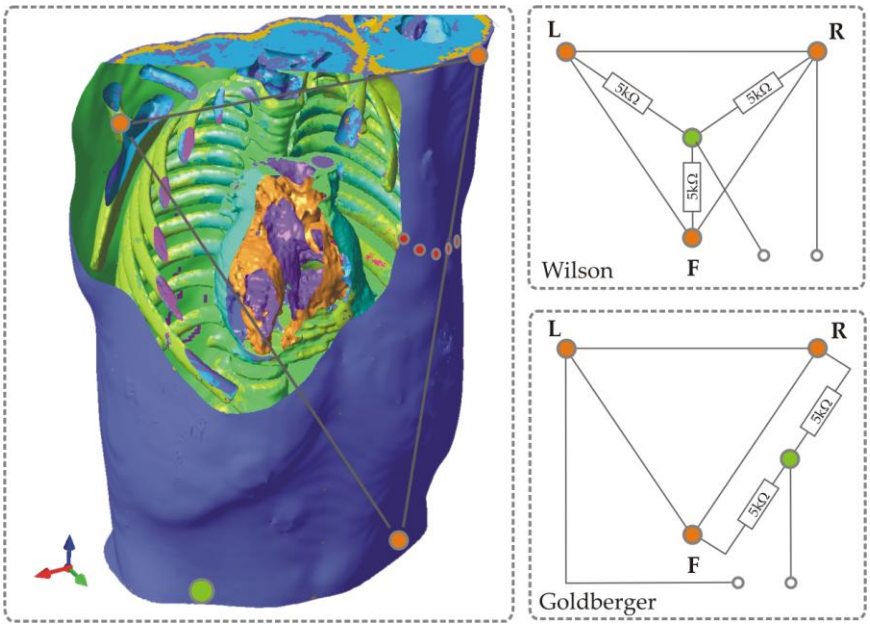

Figure 2: Left panel: Finite element model of a torso including the heart. Various tissue types were identified during the segmentation process to be able to assign tissue-specific conductivities during simulation runs. Right panels: Basic setups for computing the clinical standard 12-lead electrocardiograms (ECG). including Einthoven, Goldberger and the precordial Wilson leads (red circles in right panel)

Computational Challenges

Further improvements in HPC hardware and simulator technology is required to enable the desired in-silico experiment with near real-time performance. Required improvements are summarized here:

1. *Reduce execution times by further improving scalability which requires algorithmic improvements.*
2. *Reduce execution times by increasing the compute power of a single core. With existing CPU designs this is impossible, GPUs seem to be the more promising solution.*
3. *Improve memory scalability of codes to enable the use of high-resolution models with para-cellular resolution, or models which explicitly account for current flow in a surrounding medium such as a torso in a clinical study.*

Only task 3 is suited for being worked on within this proposal. As discussed during a previous meeting, to tackle 1 specific algorithmic expertise is required which is best addressed through other avenues. Task 2 remains a very important goal of the group which would allow a speedup of 10; the group would welcome the opportunity to work with Cray on this in a future proposal.

For the objectives of the proposed study task 3, memory scalability, needs to be addressed to accommodate the large model sizes required to accommodate a human heart immersed in a discretized torso model. As preliminary tests with high-resolution whole heart models indicated it seems as if there is a sequential memory footprint in the code, in particular the non-parallel data structures for holding the nodal information of the FE mesh, limits model size to roughly 40 million degrees of freedom

Resources

The following resources (engineering and compute time) are required:

1. *Memory profiling: Memory usage of the code needs to be profiled to better identify which parts of the code need to be revised. Based on preliminary measurements within a DCSE project we have a rough understanding which code sections are problematic and how to solve this problem. The CoE effort is estimated to be 2 weeks.*
2. *Parallelizing sequential memory structures: Based on insights gained in #1, some data structures and code sections have to be modified to cut down memory requirements. The CoE effort here is around 2 weeks.*
3. *Computational resources: For carrying out benchmarks and execute the proposed preliminary study of studying the signature of intra-cardiac therapy delivery upon body surface ECG requires roughly 1,000,000 AUs.*

We also propose a closer study of the other key application factors:

4. *Discover the limiting factor on achieving higher scalability and performance with the slower, but more stable, implicit steps. Also asses the scalability of the explicit solver at the higher core counts that will become possible. The CoE effort is estimated to be 2 weeks.*

The project requires some preparation from the scientific side, mainly some improvements in the segmentation of the heart and torso models as well as the assignment of structural properties to the heart model. This preliminary work may take a few weeks; August 1st is a realistic start date, and to complete the work by October 1st.

The total contribution requested for this project would amount to 6 weeks CoE effort and 1 Million AUs.

Deliverables

The improvement of the application is the primary deliverable and the completion of the successful modelling of the heart-torso model. This is a very significant result in this area and will generate enough material for two significant papers and also this high profile work is suitable for further publication within the HECToR community and further afield.

We plan to submit a further proposal to continue to develop this application towards real-time modelling.