A Heterogeneous Computing Model for a Grand Challenge Problem.

Prof Tim David (University of Canterbury)

It is well known that atherosclerosis occurs at very specific locations throughout the human vasculature, such as arterial bifurcations and bends, all of which are subjected to complex time-dependent fluid dynamics including low wall shear stress. A key player in the pathology of atherosclerosis is the endothelium (a mono-layer of endothelial cells surrounding the lumen), controlling the passage of material to and from the artery wall. It has become increasingly recognised that this layer is not simply a passive barrier, but plays a crucial role in maintaining vascular homeostasis and regulating the passage of materials between the blood and the vessel wall. EC's are coupled to one another in a local sense through junctions and are sensitive to biochemical stimuli such as blood–borne agonists and they are also a source of numerous vasoactive factors, such as nitric oxide (NO), prostaglandin, endothelin and angiotensin.

Any computational model requires that the cellular chemistry within the endothelium and the fluid dynamics of blood flow are closely coupled.

Clearly there are two characteristics length scales to this problem, firstly that of the cell length and secondly the artery dimension. They are orders of magnitude different. The fundmamental problem therefore is how to integrate these two scales in a coupled computational environment. The cellular chemistry is determined on the cell scale whilst the fluid dynamics are determined in a continuum fashion at the scale of arterial radii.

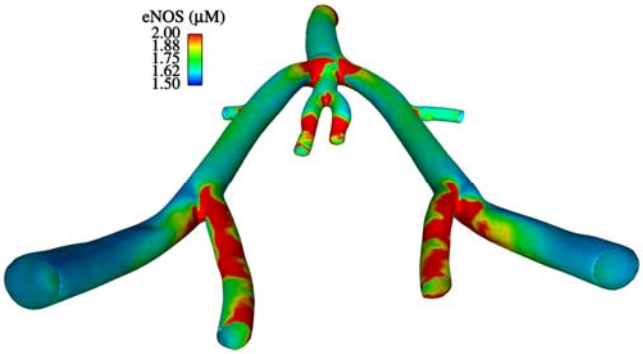

Figure 1

Figure 1 shows the artery length scales with a contour map of eNOS (endoethial derived nitric oxide synthase) , a secondary messenger for the vaso-dilator NO (nitric oxide). This "uncoupled" result was computed on a IBM p-575 using about 8 processors with approximately 10 million control volumes and took about 72 hours of compute time. At this same scale some 1 million endothelial cells would line the artery.

The endothelial cells and their asscoaited set of differential equations governing ion concentrations etc. can be mapped uniquely onto the torus network of the Blue Gene whilst the p-series is ideally suited to the computational fluid dynamic model of a mixed parabolic-elliptic system of equations .

The coupling procedure requires that at each time-step information from the endothelial cellular model be transmitted to the fluid dynamics model and in return information of the fluid motion and chemical concentrations be transmitted to the endothelial model (essentially each is a boundary condition for each other) .

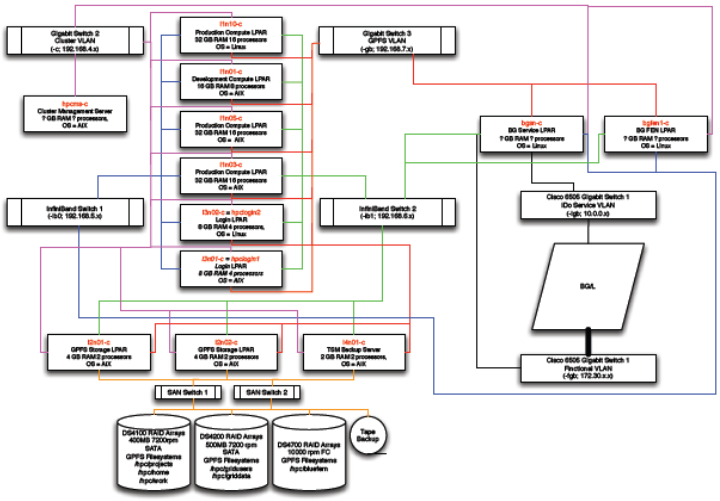

Figure 2

The presentation will show how a Blue Gene and an IBM p-series are connected together via Gigabit ethernet from the i/o nodes of the Blue Gene to the LPARs of the p-sereies as indicated in Figure 2 such that a fully coupled problem can be executed via standard MPI calls. The solution should be such that synchronisation is accomplished with minimum information passing and optimal load balancing. This “concept” may well be used in a generic way to provide for a heterogeneous system capable of supporting a number of grand challenge problems.