UNIVERSITY OF AMSTERDAM

Distributed Multiscale Computing
J. Borgdorff

Summary

Multiscale models combine knowledge, data, and hypotheses from different scales. Simulating a multiscale model often requires extensive computation. This thesis evaluates distributing these computations; an approach termed distributed multiscale computing (DMC). First, the process of multiscale modeling is examined, in order to describe it in a general and effective way. Then, multiscale models are described with a scale-aware component-based approach, treating them as a set of coupled single scale models. The computational architecture of multiscale applications is then specified with the multiscale modeling language. Such a specification can be analyzed for its structural and computational characteristics, using a task graph, and it can be used as the basis for an implementation with Multiscale Coupling Library and Environment 2 (MUSCLE 2). MUSCLE 2 executes multiscale applications on local and distributed machines with a low overhead. As a use case, a model of in-stent restenosis (ISR3D) is described as a set of coupled single scale models, specified with the multiscale modeling language, and implemented and executed with MUSCLE 2. During the making of this thesis, the performance of distributed multiscale computing was also tested. The experiments showed that DMC lead to a decrease in resource consumption under specific circumstances. Five other applications from several domains evaluated DMC, and derived different benefits from it: an increase of simulation speed or a decrease in resource consumption by using heterogeneous machines; or an increase in simulation speed by using more resources altogether. Given these results, DMC is deemed viable for heterogeneous multiscale models and for users with limited local computing resources.