Predictive Analysis of Pipelined Wavefront Codes Using a Re-usable Analytic Model*

Pipelined Synchronous Wavefront Computations form a major portion of the high performance scientific computing workload at organisations such as the Los Alamos National Laboratories in the U.S and the Atomic Weapons Establishment in the U.K. To aid the design, procurement and optimization decisions of these applications on high-end computers there has been and continues to be considerable research interests in performance engineering these codes.

In this work we present several speculative performance projections for pipelined wavefront computations, using a recently developed reusable analytic model. Particularly we show the quantitative and qualitative results for (1) assessing application and hardware configurations (2) application bottleneck analysis including the impact of changing the compute performance of a processor and the contribution of network latency/bandwidth (3) developing new metrics for procurement and configuration and (4) assessing new application design modifications.

We validate the analytic model's quantitative predictions against (1) a large Cray XT4 (2) a commodity CMP cluster based on dual-Intel Xeon processors interconnected by an InfiniBand network and (3) a predictive discrete event simulator developed by the High Performance Systems Group at Warwick. The high correlation of the results from these validations gives high confidence in the accuracy of our predictions. This analysis further demonstrate the versatility of our methods in answering wavefront code design questions in a low cost and highly efficient manner.

*Results adapted from previous publications at IPDPS 2008, UKPEW 2008 and work in progress.