A Cost Model for Communication on a Symmetric MultiProcessor*

Nancy M. Amato† Andrea Pietracaprina‡ Geppino Pucci§
Lucia K. Dale† Jack Perdue†

Abstract

The widespread use of parallel computers has been hampered by the difficulty of exploiting their massive computational potential to an extent that warrants their large cost. Indeed, it has often been noted that theoretically efficient algorithms exhibit poor performance when implemented on real machines. This is due not only to the intrinsic limitations of asymptotic analysis but also to the inadequacies of the cost model employed to predict running times. Although much progress has been made, the development of adequate tools for predicting *actual* algorithm performance on real machines remains one of the most challenging problems in parallel processing. We believe that the achievement of this goal requires a tighter coupling of the cost model to the architecture, and in particular to the memory hierarchy, than has been previously employed.

We investigated the issue of predictivity for a specific architecture, the SGI PowerChallenge, belonging to the family of Symmetric Multiprocessors. In such a system, the cost of an access may vary dramatically: from 1-2 cycles for L1 cache, to tens of cycles for L2 cache, to hundreds of cycles for main memory. Other factors affecting access costs are invalidations, false-sharing, and differing costs for read and write accesses. Our results are summarized below:

Impact of the memory hierarchy We provide evidence that interaction with the memory hierarchy affects communication costs in such a substantial way that none of the existing models (e.g., BSP-like models, LogP, CG), which do not take this interaction into account, can guarantee a reasonable level of accuracy. In particular, we identify computations that perform the same number of reads/writes but whose execution costs differ by up to two orders of magnitude.

Modeling best and worst case accesses We define two families of access patterns that make best-case and worst-case use of the memory hierarchy by maximizing and minimizing, respectively, the locality of reference. Within each family, memory hierarchy effects are under control, and communication costs can be predicted with a high-level of accuracy through a linear cost function of the 'right' parameters: number of reads/writes on a processor, and total number of accesses by all processors.

Prediction Interval On a number of application programs, we show that the cost functions synthesized for the above families provide a prediction interval that can be used to: (i) lower and upper bound the application's actual communication time, and (ii) assess the degree of locality of the application's communication patterns. Thus, the prediction interval may be employed as a profiling tool to gain valuable insights into the application's interaction with the memory hierarchy,

*This research was supported in part by NATO CRG 961243 "Bulk Synchronous Computational Geometry," and by the National Center for Supercomputing Applications under CCR970010N (utilizing the SGI Power ChallengeArray at NCSA, University of Illinois at Urbana-Champaign). The work at Texas A&M was also supported by the NSF by CAREER award CCR-9624315 and grant IRI-9619850.

†Email: {amato,dalel,jkp2866}@cs.tamu.edu. Department of Computer Science, Texas A&M University, College Station, TX, USA.

‡Email: andrea@artemide.dei.unipd.it. Dipartimento di Matematica Pura e Applicata, Università di Padova, Italy.

§Email: geppo@artemide.dei.unipd.it. Dipartimento di Elettronica e Informatica, Università di Padova, Italy.

which may in turn lead to better algorithm design. Moreover, the prediction interval is simple to apply, requiring similar programming effort as the BSP-like models (i.e., counting reads and writes per processor).