

DISTRIBUTION STATEMENT R
Approved for public release
Distribution Unlimited

Architectural Implications of a Family of Irregular Applications

David O'Hallaron, Jonathan Richard Shewchuk, and Thomas Gross

November 14, 1997

CMU-CS-97-189

DTIC QUALITY INSPECTED 2

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

19980310 125

Abstract

Irregular applications based on sparse matrices are at the core of many important scientific computations. Since the importance of such applications is likely to increase in the future, high-performance parallel and distributed systems must provide adequate support for such applications. We characterize a family of irregular scientific applications and derive the demands they will place on the communication systems of future parallel systems. Running time of these applications is dominated by repeated sparse matrix vector product (SMVP) operations. Using simple performance models of the SMVP, we investigate requirements for bisection bandwidth, sustained bandwidth on each processing element (PE), burst bandwidth during block transfers, and block latencies for PEs under different assumptions about sustained computational throughput. Our model indicates that block latencies are likely to be the most problematic engineering challenge for future communication networks.

Effort sponsored in part by the Advanced Research Projects Agency and Rome Laboratory, Air Force Materiel Command, USAF, under agreement number F30602-96-1-0287, in part by the National Science Foundation under Grant CMS-9318163, and in part by grant from Intel Corporation and Digital Equipment Corporation. The Pittsburgh Supercomputing Center provided time on the Cray T3D and T3E systems. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Advanced Research Projects Agency, Rome Laboratory, or the U.S. Government. Authors' email addresses: {droh,jrs,trg}@cs.cmu.edu

Contents

1	Introduction	1
2	The family of Quake applications	2
2.1	Quake finite element models	2
2.2	Parallel finite element simulations	2
2.3	The Parallel SMVP	4
3	An SMVP performance model	4
3.1	Model of the computation phase	5
3.2	High level model of the communication phase	6
3.3	Low level model of the communication phase	7
3.4	Error bound	8
4	Model validation	10
4.1	Properties of the Quake applications	10
4.2	Communication model validation	13
5	Communication requirements	15
5.1	Bisection bandwidth	15
5.2	Sustained PE bandwidth	16
5.3	Bandwidth and latency tradeoffs	17
6	Concluding remarks	18

About this Report

This report is an extended version of a paper presented at the Fourth International Symposium on High Performance Computer Architecture, Las Vegas, Nevada, February 1998.

The Spark98 suite [14], a collection of 10 portable sequential and parallel SMVP kernels written in C and based on this report's sf10 and sf5 meshes, is available at www.cs.cmu.edu/~quake/spark98.html. The complete set of partitioned San Fernando meshes is available at www.cs.cmu.edu/~quake/meshsuite.html.

1 Introduction

The peak performance of the commodity processors in modern parallel systems is doubling every couple of years. Clearly, communication performance needs to improve as processor performance increases, but the question is how much and in what respects. It is crucial to understand communication requirements because some parts of a high-performance communication system cannot be commodity, and will therefore be expensive.

In general it is important to understand the communication requirements of real applications [5, 10], and these requirements are especially difficult to characterize for the important class of irregular scientific applications that manipulate sparse matrices. Such applications typically model natural phenomena like the flow of air over an airplane wing or the distribution of stresses and strains in a bridge, and tend to be large, complex, and poorly understood.

This report addresses the following question: as microprocessors in parallel systems continue to improve, how much must communication systems improve to run irregular applications efficiently? Specifically, we address this question for the *Quake applications*, a family of three-dimensional unstructured finite element simulations. The Quake applications, described in Section 2, simulate the motion of the ground during strong earthquakes. They were developed as part of an ongoing project at Carnegie Mellon to model earthquakes in the Los Angeles Basin and other alluvial valleys [2].

The running time of the Quake applications is dominated by a sparse matrix-vector product (SMVP) operation that is repeated thousands of times, and the SMVP is the only operation besides I/O that requires the transfer of data between processors. Thus, as modelers we can focus on understanding the behavior of the SMVP operation and abstract away much of the complexity of the application. We derive a simple performance model for operations that consist of separate computation and communication phases—particularly the SMVP—in Section 3. In Section 4, we describe how to estimate the machine-specific parameters in the model, and we validate the model on a real machine.

In Section 5 we apply the model to derive requirements for (1) bisection bandwidth, (2) sustained communication bandwidth per processing element (PE), and (3) block transfer latency and burst bandwidth, under various assumptions about efficiency and local MFLOP rates. (We use *processing element* instead of the common term *node* to avoid confusion with the nodes of finite element meshes.)

Our detailed characterizations of the requirements of the Quake applications reveal that bisection bandwidth is not important for irregular finite element applications; bandwidth at each PE is what matters. For systems with a sustained computational rate of 200 MFLOPS, PEs will need about 300 MBytes/sec of *sustained* bandwidth and 600 MBytes/sec of burst bandwidth to run irregular codes with good efficiency.

Our work is similar in spirit to that of Cypher, Ho, Konstantinidou, and Messina [5], who characterize eight regular and irregular scientific applications in terms of memory, processing, communication, and I/O requirements, and build scalability models for three of the simpler regular applications. However, our approach is different in that our goal is depth rather than breadth. We provide a detailed characterization for a specific family of irregular applications that are real (in the sense that people really care about the results that the Quake applications compute), that we understand completely, that we have complete control over, and that we can make arbitrarily large or small by adjusting the frequency of ground motion.

To show that the generality of our model seems to extend beyond this single family of irregular applications, we observe that there is some evidence that the Quake applications are typical of unstructured finite element simulations. For example, the EXFLOW application from Cypher et al. [5] is a 3D unstructured finite element program that simulates a fluid dynamics problem on 512 PEs. Interestingly, EXFLOW has

nearly identical computational properties as a similarly sized Quake application (called sf2/128, which resolves a wave with a two-second period on 128 PEs). EXFLOW and Quake each require about 2 MBytes of data on each PE. The communication volume/MFLOP is 144 KBytes for EXFLOW, vs. 155 KBytes for Quake. The average number of messages/MFLOP is 66 messages for EXFLOW, vs. 60 messages for Quake. And the average message size is 2.2 KBytes for EXFLOW, vs. 3.6 KBytes for Quake.

These two unstructured finite element applications are from two different scientific domains, yet each has similar computational properties and differs from regular applications in similar ways. Irregular finite element applications like EXFLOW and Quake have an average total communication volume similar to that of the regular applications studied earlier by Cypher et al. [5], but they transfer more messages having a smaller average size than most of those regular applications. Perhaps our most troublesome conclusion is that because the blocks transferred between PEs tend to be small even for large irregular applications, block latency costs cannot be amortized by large messages, and communication latency will need to be a central focus of future efforts to engineer effective communication networks and software.

2 The family of Quake applications

The Quake applications are unstructured finite element codes that were developed to predict ground motion in the San Fernando Valley of Southern California during earthquakes [2]. There are four Quake applications, denoted sf10, sf5, sf2, and sf1. The “sf” is an abbreviation for San Fernando, and the number indicates the period (in seconds) of the highest frequency wave that the simulation is able to resolve. For example, sf10 resolves waves with 10 second periods, sf5 resolves waves with 5 second periods, etc.

Each Quake application consists of (1) a three-dimensional unstructured finite element model of the ground beneath San Fernando, and (2) a parallel finite element program that simulates the propagation of seismic waves through the model for 60 seconds of simulated time.

2.1 Quake finite element models

Each Quake application employs a three-dimensional unstructured mesh, composed of thousands or millions of tetrahedra (i.e., pyramids with triangular bases), and models a volume of earth roughly 50 km x 50 km x 10 km in size (Figure 1). Each tetrahedron is called an *element*, and the vertices of the tetrahedra are called *nodes*. Some finite element simulations use structured meshes constructed from regular grids; however, the Quake applications require unstructured meshes, which can accommodate the wildly varying density of the soils in the valley. To ensure that the simulation is numerically stable, the size of elements in any region of the mesh must be matched to the wavelength of ground motion, which is shorter in softer soils and longer in hard rock. Thus, softer soils need a higher density of smaller elements.

Figure 2 records the sizes of the San Fernando meshes. When the wave's period is halved, its frequency doubles, and the number of nodes increases by a factor of nearly eight—a factor of two for each of three dimensions. As a general rule, for each node in the mesh, a simulation uses about 1.2 KByte of memory at runtime to accommodate the storage of several vectors and sparse matrices. For example, sf2 requires about 450 MBytes of memory at runtime.

2.2 Parallel finite element simulations

During each of 6000 time steps, a Quake finite element simulation executes a sparse matrix-vector product (SMVP) operation of the form $y = Kx$, where x and y are vectors of length $3n$ (here n is the number of

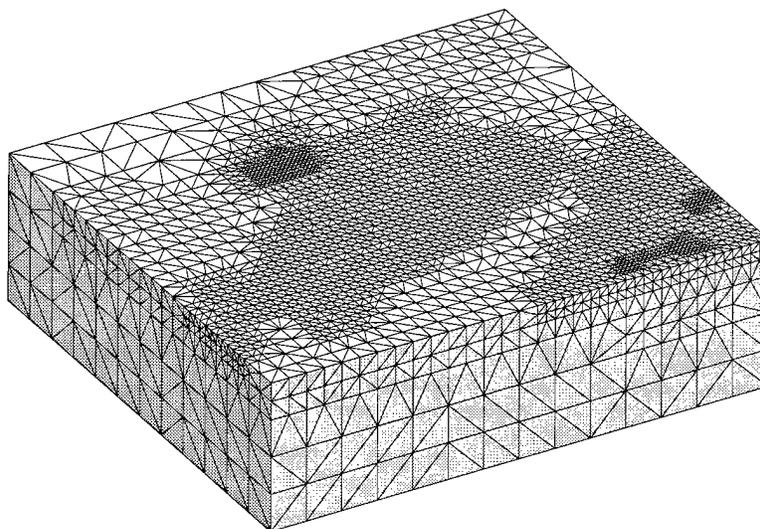


Figure 1: Finite element mesh for the sf10 model of the San Fernando Valley.

Number of ...	sf10	sf5	sf2	sf1
nodes	7,294	30,169	378,747	2,461,694
elements	35,025	151,239	2,067,739	13,980,162
edges	44,922	190,377	2,509,064	16,684,112

Figure 2: Sizes of the Quake meshes.

nodes), and K is a sparse $3n \times 3n$ *stiffness matrix*. Because an explicit time-stepping method is used, there are no other parallel operations (such as dot products or preconditioning). The vectors x and y have length $3n$ because each vector represents three degrees of freedom— x , y , and z displacements—for each node of the mesh. K can be likened to an adjacency matrix of the nodes of the mesh; K contains a 3×3 submatrix for each pair of nodes that are connected by an edge of the mesh (including self-edges). The stiffness matrix is extremely sparse; each node is connected to an average of 13 neighbors (in addition to itself), so each row of K contains an average of $14 \times 3 = 42$ nonzero floating point numbers [15].

The simulations are parallelized using Archimedes, a domain-specific tool chain for finite element problems [2, 17]. To generate a simulation that will run on p PEs, Archimedes partitions the mesh into p disjoint sets of elements. Each set is called a *subdomain*; a one-to-one mapping is established between PEs and subdomains. The program that partitions each mesh into subdomains is based on a recursive geometric bisection algorithm [12] that divides the elements equally among the subdomains while attempting to minimize the total number of nodes that are shared by multiple subdomains, and hence the total communication volume. The geometric partitioning algorithm has provable asymptotic upper bounds on the number of shared nodes, and in practice generates partitions that are competitive with those produced by other modern partitioning algorithms [7, 3, 8].

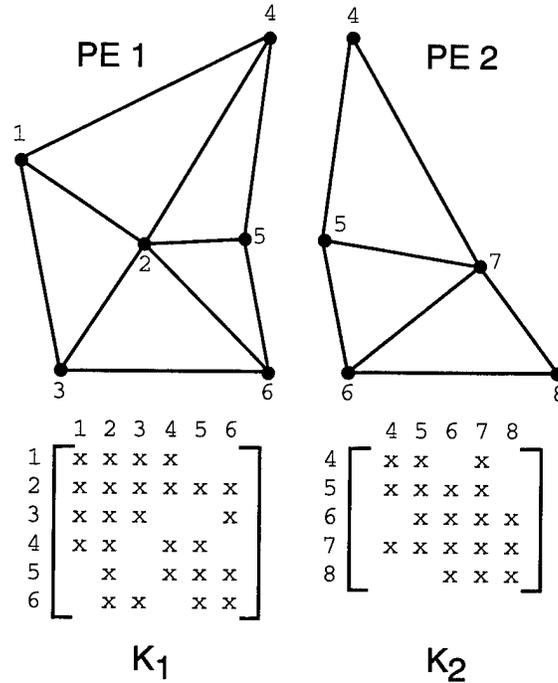


Figure 3: A finite element mesh and corresponding stiffness matrix K , distributed between two PEs. Each X represents a 3×3 submatrix.

2.3 The Parallel SMVP

The running time of the Quake applications is dominated by the execution of SMVP operations, which consume over 80% of the total running time in the sequential case. Furthermore, the SMVP operations are the only operations (besides I/O) that require interprocessor communication. So even though the Quake applications are complicated, we can abstract away most of the complexity and focus on the performance of a simple well-defined parallel SMVP kernel.

To compute the global SMVP operation $y = Kx$ on a set of PEs, we must consider the data distribution by which vectors and matrices are stored. The vectors x and y are stored in a distributed fashion according to the mapping of nodes to PEs induced by the partition of elements among PEs. If a node i resides in several PEs (because i is a vertex of several elements mapped to different PEs), the values x_i and y_i are replicated on those PEs. The matrix K is distributed so that K_{ij} resides on any PE on which nodes i and j both reside. Figure 3 demonstrates this method of distributing data.

With this method of distributing data, the global SMVP $y = Kx$ is performed in two steps. First, each PE computes a local SMVP over the subdomain that resides on that PE. Second, PEs that share nodes communicate and sum their nodal y values into correct global values for each node. In Figure 3, PEs 1 and 2 must communicate with each other to resolve the values of the shared nodes 4, 5, and 6.

3 An SMVP performance model

Each global SMVP consists of a simple sequence, executed simultaneously on each PE: a local SMVP (the computation phase) followed by an operation in which each PE exchanges and sums data with each

Parameters for Equation (1)	
F	flops per PE per SMVP
C_{max}	maximum communication words per PE per SMVP
E	target efficiency
T_f	amortized time per flop (inverse of sustained flops)
Parameters for Equation (2)	
B_{max}	maximum communication blocks per PE per SMVP
C_{max}	maximum communication words per PE per SMVP
T_l	time per communication block (block latency)
T_w	time per additional block word (inverse of burst bandwidth)
Computed quantities	
T_{smvp}	running time for sparse matrix-vector product (SMVP)
T_{comp}	running time for computation phase of SMVP
T_{comm}	running time for communication phase of SMVP
T_c	amortized time per comm word (inverse of sustained bandwidth)
β	upper bound on relative overestimate of communication time
M_{avg}	average message size (words)

Figure 4: Summary of symbols.

PE it shares nodes with (the communication phase). The computation and communication phases are synchronous, i.e., each phase begins with a global barrier synchronization.¹ Further, we assume (1) that during the communication phase, the summing operations are overlapped with the exchanges, and (2) that the time to perform the exchanges determines the running time of the communication phase. These assumptions are reasonable, as the summing operations can be handled by the PE and the exchanges by a network interface.

Given these assumptions, the running time for each global SMVP is

$$T_{smvp} = T_{comp} + T_{comm},$$

where T_{comp} is the time required for all PEs to finish their computation phases, and T_{comm} is the time required for all PEs to finish their communication phases. For reference, Figure 4 summarizes all symbols introduced in this section.

3.1 Model of the computation phase

The unit of work for the computation phase is a floating-point operation (*flop*), which is either a scalar add or multiply. Although the flop is not an ideal way to measure work in every application, it is reasonable for the SMVP operation because the (minimum) number of flops each PE must perform can be counted exactly and is independent of the compiler and its optimization level. If a PE's local matrix has m nonzeros, then the local SMVP requires at least $F = 2m$ flops; one add and one multiply for each nonzero. If each flop requires an average time of T_f , then the running time for a local SMVP is FT_f . Modern mesh partitioners

¹It is in principle possible, with difficult modifications to the applications, to improve performance by overlapping the computation and communication phases, but the implementations of the Quake application do not do this, so the models in this section follow suit. (This report undertakes to discuss how architects can accommodate scientific users, and not vice versa.) By not modeling any overlap, we obtain conservative bandwidth and latency estimates, and avoid possibly slowing the program by complicating the runtime system.

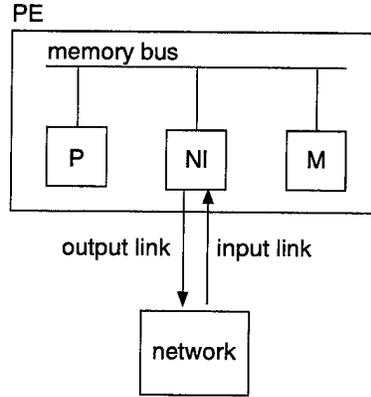


Figure 5: Processing element (PE) model (P: processor, M: memory system, NI: network interface).

do an excellent job of distributing computation evenly across PEs [15], so we will assume that each PE performs $F = 2^m$ flops and thus the running time for the computation phase is

$$T_{comp} = FT_f.$$

Since T_f includes all hardware and software overheads (e.g., loads, stores, various miss penalties, pipeline stalls, etc.) it is difficult to predict. However, T_f can be precisely measured for a particular application running on a particular system. For example, measurements of the local SMVPs from the Quake applications show a steady $T_f = 30$ ns for the Cray T3D (150 MHz Alpha 21064, cc -O3) and $T_f = 14$ ns for the Cray T3E (300 MHz Alpha 21164, cc -O3).

3.2 High level model of the communication phase

The hardware part of the communication system of a PE is modeled as a network interface (NI) with an input link and output link, as shown in Figure 5.

The running time of the communication phase is given by

$$T_{comm} = C_{max}T_c,$$

where C_{max} is the maximum number of words communicated by any PE during the communication phase and T_c is the average time per communication word. The rate T_c^{-1} is the *sustained bandwidth* across the links of the network interface during the communication phase. T_c includes the time to transfer data in and out of a PE, as well as all software and hardware overheads.

We define the *efficiency* E to be the proportion of SMVP time devoted to computation; that is, $E = T_{comp}/T_{smvp}$. Given this definition, $T_{comm}/T_{comp} = (1 - E)/E$. We have seen that $T_{comm}/T_{comp} = (C_{max}T_c)/(FT_f)$, and hence

$$T_c = \frac{F}{C_{max}} \left(\frac{1 - E}{E} \right) T_f. \quad (1)$$

Equation (1) describes at a high level the relation between sustained computation bandwidth (T_f^{-1}) and sustained communication bandwidth (T_c^{-1}). The model is interesting because it cleanly separates the various factors that relate computation performance to communication performance: The *computation/communication ratio* F/C_{max} is a property of the application and the partitioner; T_f is a property of the processor

architecture and compiler; E is a target efficiency imposed by the user. This separation of factors is similar in spirit to the familiar CPI model for uniprocessor performance [9].

In Section 5 we use Equation (1) to investigate the following question: given a target efficiency E for an SMVP running on PEs with local computational rates of T_f^{-1} , what is the required communication performance of the system? The answers to this question provide insight, for this class of applications, into the sustained performance we need from our communication systems as processor speeds continue to increase.

3.3 Low level model of the communication phase

Equation (1) provides a high level model of sustained performance in terms of streams of words being transferred between memory and the network interface. However, in a lower level and more realistic view of communication, the unit of work during the communication phase is the transfer of a *block* between the network and the local memory system of the PE. The time for this transfer is called the *block transfer time*. A block is a transfer unit, a group of words that move together from one PE to another. Blocks may be fixed-sized or variable-sized. For example, a block might be a cache line in a CC-NUMA system [11], a message in a message passing system [13], a bulk asynchronous data transfer between two PEs' memories [16], or a page in a software distributed shared memory system [1].

The transfer time for a block i of l_i words is $T_l + l_i T_w$, where T_l is the constant *block latency*, T_w is the constant marginal cost of transferring each additional word, and T_w^{-1} is the *burst bandwidth*. Notice that the block latency does not include any delay through an interconnection network. We only model the overhead of transferring data between the network interface and local memory, in view of previous findings that most of the cost of communication on modern systems is incurred at the individual PEs [18]. In essence, we assume that the interconnection network has infinite capacity and constant latency. In Section 4.2, we offer empirical evidence that this is a reasonable assumption for the SMVP running on tightly coupled systems.

Modern mesh partitioners do an excellent job of balancing computation (F) and a good job of minimizing the global volume of communication. Unfortunately they do less well in balancing the total number of blocks B_i and the total volume in words C_i sent and received by each PE i [15]. As a simplifying assumption, we pessimistically assume that the PE that transfers the maximum number of words (C_{max}) is the same PE that transfers the maximum number of blocks (B_{max}). In this case, the same PE has the longest communication phase, and thus

$$T_{comm} = B_{max}T_l + C_{max}T_w.$$

Finally, since $T_c = T_{comm}/C_{max}$, we have

$$T_c = \frac{B_{max}}{C_{max}}T_l + T_w. \quad (2)$$

T_l and T_w are properties of the communication system, including the architecture, hardware implementation, and software libraries. In Section 4.2, we describe a simple method of estimating these parameters on real systems. For the Cray T3E, our measurements indicate that $T_l \approx 22 \mu s$ and $T_w \approx 55 ns$.

For message-passing systems, B_{max} (the maximum number of blocks transferred by one PE) is purely a property of the application and the partitioner, but on shared-memory systems it is also a property of the architecture, as the data sent from one PE to another may have to be broken up into multiple blocks (i.e., cache lines). Thus, Equation (2) characterizes sustained communication bandwidth during the communication phase of the SMVP in terms of basic properties of the communication system and the application. In

subdomains	sf10	sf5	sf2	sf1
4	1.00	1.00	1.00	1.00
8	1.00	1.00	1.00	1.00
16	1.09	1.10	1.07	1.00
32	1.01	1.01	1.15	1.00
64	1.03	1.08	1.11	1.05
128	1.03	1.04	1.04	1.11

Figure 6: Computed upper bounds on β_{max} , the maximum factor by which T_{comm} and T_c are overestimated for each Quake application.

Section 5, we use this model to explore bandwidth and latency tradeoffs in communication systems as the base microprocessors continue to improve.

The models in Equations (1) and (2) are similar in some ways and different in others to the LogP model [4]. LogP is a general performance model for bulk-synchronous parallel (BSP) computations [19], where a program is viewed as a series of *supersteps* separated by barrier synchronizations. During a superstep, each PE performs local computation and transfers a limited number of messages. The models we have developed are for a restricted class of BSP programs where each superstep (SMVP) consists of a distinct computation phase followed by a distinct communication phase. In general, our models view performance at a lower level of abstraction than LogP. The parameters T_f , T_w , F , B_{max} , and C_{max} have no counterparts in LogP. On the other hand, our T_l parameter is similar to the overhead parameter o in LogP.

3.4 Error bound

In deriving Equation (2), we assume that the PE with the most words to communicate also has the most blocks to communicate. In practice this assumption may not be true. The question is, by how much do we overestimate T_{comm} (and thus T_c) as a result? For a given partitioned mesh, this question is easy to answer if we know the machine parameters T_l (block latency) and T_w (inverse burst bandwidth). However, we would like to establish a worst-case bound that applies to all machines, regardless of their parameters.

The length of the communication phase is determined by the PE that takes the longest; that is,

$$T_{comm} = \max_{i \text{ is a PE}} (B_i T_l + C_i T_w).$$

Our simplifying assumption is to calculate T_{comm} as if there is some PE i such that $B_i = B_{max}$ and $C_i = C_{max}$. Let β be the factor by which this assumption causes our model to overestimate T_{comm} . In other words,

$$\beta = \frac{B_{max} T_l + C_{max} T_w}{\max_i (B_i T_l + C_i T_w)}.$$

Let β_{max} be the largest possible value of β over all values of T_l and T_w . Figure 6 shows that β_{max} is not much larger than one for any of the Quake applications. Hence, Equation (2) is a reasonably good model even when our assumption is not satisfied. The rest of this section describes how the numbers in Figure 6 are obtained, and may safely be skipped by the reader not interested in minute technical detail.

Let $\alpha = T_w/T_l$. Our assumption will cause us to overestimate T_{comm} (and hence T_c) by

$$\begin{aligned}\beta &= \frac{B_{max}T_l + C_{max}T_w}{\max_i(B_iT_l + C_iT_w)} \\ &= \frac{B_{max} + \alpha C_{max}}{\max_i(B_i + \alpha C_i)} \\ &= \min_i \frac{B_{max} + \alpha C_{max}}{B_i + \alpha C_i}.\end{aligned}\tag{3}$$

Clearly, α is a property of the communication system, and reflects the balance between block latency and burst bandwidth. To ensure that our model is reasonably accurate for any machine, we wish to find the value of α that yields the most pessimistic (largest) value of β :

$$\beta_{max} = \max_{\alpha} \min_i \frac{B_{max} + \alpha C_{max}}{B_i + \alpha C_i}.$$

To simplify the problem a bit, first consider the behavior of this expression for a single PE j , which happens to be the processor that yields β_{max} when substituted for i in the formula above. Specifically, consider the function

$$\beta_j(\alpha) = \frac{B_{max} + \alpha C_{max}}{B_j + \alpha C_j}.$$

What value of α maximizes β_j ? It is apparent on inspection that the derivative $d\beta_j/d\alpha = (C_{max}B_j - B_{max}C_j)/(B_j + \alpha C_j)^2$ is either everywhere positive, everywhere negative, or uniformly zero over $\alpha \geq 0$; hence, there are no critical points, and β_j is maximized at an extreme value of α .

However, the notion of "extreme values of α " is not straightforward, and does not necessarily correspond to the values $\alpha = 0$ and $\alpha = \infty$. Instead, the extreme values of α are the smallest and largest values for which PE j minimizes Expression (3). However, the PE that minimizes Expression (3) depends on α . If α is swept from zero to infinity, the PE that minimizes Expression (3) (called the *active PE* in the terminology of optimization) may change several times; each α value where such a change occurs determines a candidate for β_{max} .

Hence, β_{max} may be computed as follows. Draw a two-dimensional plot in which each processor i is represented by a point whose position is (B_i, C_i) . Compute the convex hull of these points, and consider the edges that face up and to the right. (If there are no such edges, then one of the points is (B_{max}, C_{max}) , and thus $\beta = 1$.) Each edge ij of the upper right convex hull represents a change in the active PE from one of its endpoints (say, i) to the other (say, j) as α is swept from zero to infinity. The transition occurs when the value of α is such that the edge ij is orthogonal to a line having slope α . For this value of α , $\beta_i(\alpha) = \beta_j(\alpha)$. β_{max} is found by taking the maximum such β value from among those determined by all the edges of the upper right convex hull.

Being too lazy to compute convex hulls, we implemented a simpler computation that determines an upper bound on β_{max} . Suppose there are only three PEs, having (B, C) parameters of (B_j, C_j) , $(B_{max}, 0)$, and $(0, C_{max})$, respectively. Because of the last two PEs, the denominator $\max_i(B_i + \alpha C_i)$ cannot be less than B_{max} , nor can it be less than αC_{max} ; hence, β cannot be greater than two.

When $\alpha = 0$, the active PE is $(B_{max}, 0)$ (which clearly maximizes $B_i + \alpha C_i$), yielding $\beta = 1$, and when $\alpha \rightarrow \infty$, the active PE becomes $(0, C_{max})$, also yielding $\beta = 1$. The value of α that maximizes β occurs between these extremes. As α increases from zero, the active PE changes from $(B_{max}, 0)$ to (B_j, C_j) when

$B_{max} = B_j + \alpha C_j$, and changes from (B_j, C_j) to $(0, C_{max})$ when $B_j + \alpha C_j = \alpha C_{max}$. In each of these two cases, we solve for α and substitute the result into β_j , yielding

$$\beta_j \left(\alpha = \frac{B_{max} - B_j}{C_j} \right) = 1 + \frac{C_{max}(B_{max} - B_j)}{C_j B_{max}}$$

and

$$\beta_j \left(\alpha = \frac{B_j}{C_{max} - C_j} \right) = 1 + \frac{B_{max}(C_{max} - C_j)}{B_j C_{max}}.$$

The larger of these two values is an upper bound on β_{max} (over all values of α), even if there are other processors that we have not taken into account.

If there are more than three PEs, such a bound may be computed for each PE. Each such bound is valid, so we use the smallest. Hence,

$$1 + \min_{i \text{ is a PE}} \left[\max \left\{ \frac{C_{max}(B_{max} - B_i)}{C_i B_{max}}, \frac{B_{max}(C_{max} - C_i)}{B_i C_{max}} \right\} \right].$$

is an upper bound on β_{max} . This bound is not generally tight, but we have used this expression to compute the values in Figure 6.

4 Model validation

The parameters in Equation (1) are straightforward to determine: F and C_{max} are static application properties, T_f is a constant property of the compiler and the processor, and the efficiency E is imposed by the user. Similarly for Equation (2), the parameters B_{max} and C_{max} are static application properties. However, the block latency T_l and inverse burst bandwidth T_w are system properties that require some care to estimate. In this section we summarize the properties of the Quake SMVP instances (including the parameters F , B_{max} , and C_{max}), describe a simple method for estimating T_l and T_w , and then plug these parameters into our models to predict running time on a Cray T3E.

4.1 Properties of the Quake applications

Figure 7 summarizes the properties of the various Quake SMVP instances. (See O'Hallaron and Shewchuk [15] for a complete characterization of the applications.) The values of B_{max} and C_{max} in this table are always even, because each message from PE i to PE j is matched by a message from j to i of equal length. (The values of C_{max} are also divisible by three, because there are three degrees of freedom per node.) The values of B_{max} indicate the degree of subdomain adjacency; for instance, if $B_{max} = 46$, then each subdomain shares mesh nodes with at most 23 other subdomains, and some PE must send a message to and receive a message from each of 23 other PEs. These values assume that blocks may be arbitrarily large, and hence each PE sends at most one block to each other PE. On a fine-grained shared memory machine, where the unit of interchange between PEs may be as small as a cache line, B_{max} may be much larger.

There are some interesting points to make about the numbers in Figure 7. First, conventional wisdom holds that sparse codes like the SMVP are communication intensive. However, this is not always the case. As we see for sf2, which is a reasonably large problem, the computation/communication ratios F/C_{max} vary from large (~ 450 for sf2/4) to moderate (~ 50 for sf2/128). (We use the notation sfx/y to denote an

Subdomains		sf10	sf5	sf2	sf1
4	F	453,924	1,899,396	24,640,110	162,372,024
	C_{max}	2,352	7,746	55,338	186,162
	B_{max}	6	6	6	6
	M_{avg}	369	1,290	8,682	27,540
	F/C_{max}	193	245	445	872
8	F	235,566	970,740	12,414,006	81,602,442
	C_{max}	2,550	7,080	35,148	151,764
	B_{max}	12	12	10	14
	M_{avg}	237	699	4,152	13,761
	F/C_{max}	92	137	353	538
16	F	122,742	496,872	6,278,076	41,116,374
	C_{max}	2,208	5,292	28,482	119,280
	B_{max}	18	20	16	18
	M_{avg}	159	342	1,920	7,434
	F/C_{max}	56	94	220	345
32	F	64,980	257,004	3,191,436	20,740,734
	C_{max}	2,172	4,476	24,018	87,228
	B_{max}	30	30	26	26
	M_{avg}	87	213	1,239	4,044
	F/C_{max}	30	57	133	238
64	F	34,956	134,424	1,632,708	10,511,586
	C_{max}	1,764	4,296	20,520	73,062
	B_{max}	38	40	36	38
	M_{avg}	57	135	765	2,712
	F/C_{max}	20	31	80	144
128	F	18,954	70,956	838,224	5,332,806
	C_{max}	1,740	3,360	16,260	51,048
	B_{max}	62	52	50	46
	M_{avg}	36	135	459	1,515
	F/C_{max}	11	21	52	104

Figure 7: Quake SMVP properties. F : flops per PE. C_{max} : maximum communication words on any one PE. B_{max} : maximum communication blocks on any one PE. M_{avg} : Average message size (words). F/C_{max} : computation/communication ratio.

instance of application sf_x partitioned into y subdomains.) This common misconception about sparse codes is probably due to the fact that researchers have not had the opportunity to run sufficiently large problems.

In fairness, though, there are related problems for which the computation/communication ratio is smaller. For instance, a heat conduction simulation would use only one degree of freedom per node (the temperature at each node), and each matrix nonzero would be a scalar, rather than a 3×3 submatrix. Hence, the computation would be nine times less, but the communication volume would only be three times less. Even so, the computation/communication ratio would still be large for large problems.

Second, while the computation/communication ratios are reasonably high for large problems, as the problem sizes grow by a factor of ten, we see that the computation/communication ratios grow only by a factor of two. This is not surprising; consider that a good partition of an n -node 3D mesh will produce $\mathcal{O}(n^{2/3})$ shared nodes (for the same reason that an $\mathcal{O}(n)$ -node cube has a surface area of $\mathcal{O}(n^{2/3})$ nodes). Hence, the computation/communication ratio is $\mathcal{O}(n^{1/3})$, and a factor-of-ten increase in n yields

sf10 msg size	subdomains					
	4	8	16	32	64	128
3		2		18	72	266
6		0	2	10	48	148
9-12		0	8	32	66	228
15-24		2	2	20	74	228
27-48		2	6	18	64	272
51-96		2	4	46	144	394
99-192	2	8	30	78	150	90
195-384	2	6	30	28		
387-768	6	10				
771-1536						

(a) sf10

sf5 msg size	subdomains					
	4	8	16	32	64	128
3				6	36	96
6				6	18	60
9-12		2	6	16	24	94
15-24		0	6	12	38	128
27-48		0	2	14	80	184
51-96		0	8	24	74	258
99-192	2	8	42	116	366	
195-384		0	22	72	168	154
387-768		10	38	38	10	
771-1536	8	14				

(b) sf5

sf2 msg size	subdomains					
	4	8	16	32	64	128
3				4	14	32
6				0	4	12
9-12				4	8	28
15-24			4	2	16	42
27-48			4	8	22	50
51-96			4	8	22	108
99-192			6	12	48	142
195-384			6	28	62	230
387-768			2	22	94	338
771-1536		2	10	40	142	256
1539-3072		6	30	78	84	8
3075-6144		14	22	4		
6147-12,288	8	4				
12,291-24,576						
24,579-49,152						

(c) sf2

sf1 msg size	subdomains						
	4	8	16	32	64	128	
3				2		12	
6				6	2	12	
9-12				2	8	24	
15-24				4	12	30	
27-48			2	10	6	44	
51-96				2	14	52	
99-192				2	12	24	96
195-384			4	14	30	110	
387-768		2	2	18	38	150	
771-1536			6	12	94	250	
1539-3072			8	46	102	290	
3075-6144		4	12	34	128	226	
6147-12,288		4	30	66	64		
12,291-24,576	6	16	20	4			
24,579-49,152	2	2					

(d) sf1

Figure 8: Histograms of message sizes (in words) for the Quake applications partitioned among 4, 8, 16, 32, 64, or 128 PEs. Message sizes are multiples of three because the applications store three degrees of freedom per node.

roughly a factor-of-two increase in that ratio. Our point is that while large SMVPs indeed have reasonable computation/communication ratios, these ratios do not increase quickly with increasing problem size, as they do for cubic problems like dense matrix multiply. Thus, we cannot rely on simply increasing the problem size to guarantee good efficiency.

Third, even when blocks are as large as possible (i.e., each PE sends at most one block to any other PE), the mean block size M_{avg} is surprisingly small. For example, the largest mesh (sf1) running on 128 PEs has an average message size of only about 1,500 words. Thus, we cannot rely on large blocks to amortize high block latencies. The large number of small messages that characterize finite element applications can be seen clearly in Figure 8, whose histograms show how message sizes are distributed for each of the Quake applications.

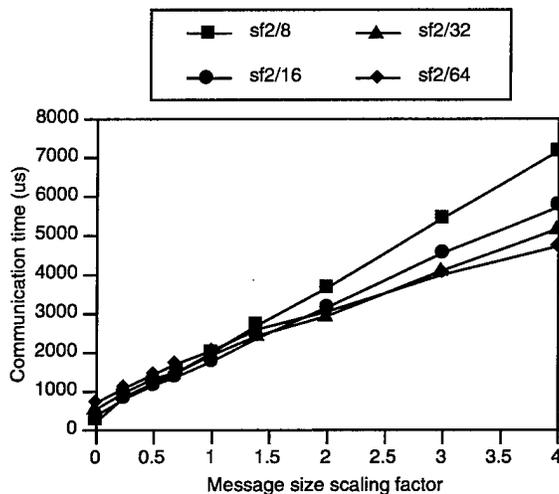


Figure 9: Scaling of the sf2 SMVP on the Cray T3E.

Finally, for sf1/128 each PE communicates with up to 18% of the other PEs. Thus, the Quake SMVP occupies an interesting middle ground between difficult applications that require all-to-all communication (like two-dimensional fast Fourier transforms), and simple applications wherein PEs communicate with a few neighbors (like finite differences on regular grids).

4.2 Communication model validation

Here we describe the method we use to estimate the parameters T_l and T_w , and offer some evidence that Equation (2) is a reasonable model.

We begin with our implementation of the parallel SMVP from the sf2 Quake application, and produce several modified versions in which we scale the size of each message by some factor c . Changing the message sizes in this manner does not necessarily yield a working application; however, by measuring the change in communication time as a function of changing message length, we can distinguish between bandwidth-related and latency-related overheads. We measure and plot the elapsed time $y(c)$ of the communication phase as a function of the scaling factor c . Performing this step for a range of scaling factors yields a curve, which (happily) resembles a line, as Figure 9 illustrates. Four curves are plotted, for instances of the sf2 application partitioned among 8, 16, 32, or 64 PEs of a Cray T3E.

Although these curves are not perfectly linear, they are close enough that we have chosen to estimate block latency and burst bandwidth in the most straightforward way. The block latency is determined from the y -intercept of each “line” by computing

$$T_l = \frac{y(0)}{B_{max}}, \quad (4)$$

and the burst bandwidth is determined from each line’s slope by computing

$$T_w = \frac{y(1) - y(0)}{C_{max}}. \quad (5)$$

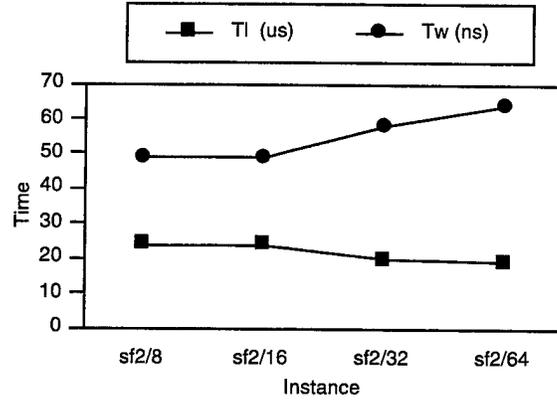


Figure 10: Measured values of T_l and T_w on the Cray T3E. Note that T_l is measured in different units (a thousand times larger) than T_w .

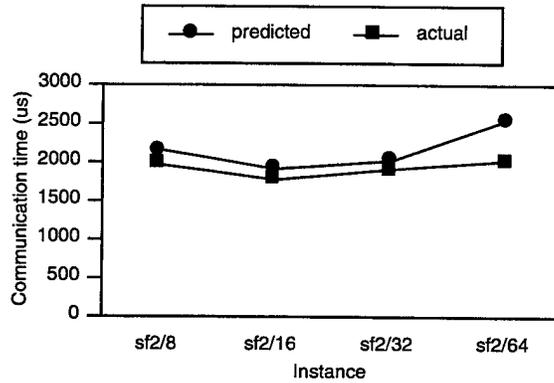


Figure 11: Predicted vs. actual time of the sf2 SMVP communication phase on the Cray T3E ($T_l = 22 \mu\text{s}$, $T_w = 55 \text{ ns}$).

The fact that communication times are nearly linearly related to the size of the messages indicates that network congestion is not an issue. If it were, then we would see concave curves rather than nearly straight lines.

As the number of processors increases, the communication time devoted to block latency (as revealed by the y -intercepts in Figure 9) increases because the number of messages per PE increases. However, the corresponding slope decreases, because the total communication volume per PE decreases. Hence, block latency becomes increasingly influential as the number of processors rises, because there is a larger number of smaller messages.

By applying Equations (4) and (5) to each of the curves in Figure 9, we find the estimates for T_l and T_w shown in Figure 10. The fact that T_w (or T_l) isn't identical in all four cases shows that our model isn't perfect, but the values vary little enough to give us some confidence. We use the averages of these values over these four instances as our final estimates for block latency and inverse burst bandwidth for the Cray T3E: $T_l \approx 22 \mu\text{s}$ and $T_w \approx 55 \text{ ns}$.

Plugging these estimates and the application properties from Figure 7 into Equation (2) yields the predicted communication phase times shown in Figure 11. While not exact, they provide reasonable predictions of actual running time.

Although we have shown in this section that it is possible to obtain reasonable estimates of T_l and T_w for the Cray T3E, these parameters are properties of complex hardware and software systems, and thus are not really constants, as we see in Figure 10. There may be communication systems for which Equation (2) is not a tenable model. For our purposes, however, the constant estimates for T_l and T_w seem to be reasonable.

5 Communication requirements

In this section, we use our performance models to address the following question: As the base commodity microprocessors in parallel systems get faster, what kind of performance will be needed from communication systems in order to run the Quake SMVPs efficiently?

We will investigate this question for two hypothetical machines: A “current” machine that runs the local SMVP at a sustained rate of 100 MFLOPS ($T_f = 10$ ns), and a “future” machine that runs the local SMVP at a rate of 200 MFLOPS ($T_f = 5$ ns). (Specifications are for 64-bit floating-point values.) Computational rates of 100 to 200 MFLOPS may seem low, but they reflect the reality that the actual performance of irregular applications is far less than peak rated performance, largely because of irregular memory reference patterns and because the data structures are too large to fit in cache. For example, a single PE of a Cray T3E (300 MHz Alpha 21164, cc -O3) runs the local SMVP from the Quake applications at approximately 70 MFLOPS ($T_f = 14$ ns), which is only 12% of the peak rated performance of 600 MFLOPS.

The sf2 SMVP will serve as a running example. Sf2 is a good example for a number of reasons. First, it is clearly large enough to qualify as a real problem by any standard (the sparse coefficient matrix is $1.14M \times 1.14M$ in size). Second, even though it is quite large, it has a wide range of computation/communication ratios (F/C_{max}), from a high of 450:1 when partitioned into four subdomains, down to 50:1 for 128 subdomains.

5.1 Bisection bandwidth

Bisection bandwidth is a popular measure of communication system performance, but it proves to be unimportant for Quake SMVPs. To compute the bisection bandwidth requirements for a Quake SMVP running on p PEs, we create a symmetric $p \times p$ matrix m such that m_{ij} is the number of 64-bit words transferred from PE i to PE j . If we assume that PEs $0, \dots, p/2 - 1$ are on one side of the bisection and PEs $p/2, \dots, p - 1$ are on the other side, then

$$V = 2 \sum_{i=0}^{p/2-1} \sum_{j=p/2}^{p-1} m_{ij}$$

words cross the bisection during the communication phase, and the sustained bisection bandwidth requirement is $V/(C_{max}T_c)$. Figure 12 shows the required bisection bandwidths given different assumptions of PE performance and overall efficiency. Here, Equation (1) is used to calculate $C_{max}T_c$. The worst case of 700 MBytes/sec ($E = 0.9$ and $T_f^{-1} = 200$ MFLOPS) is quite modest, on the order of the bandwidth of a couple of links in a modern system. Bisection bandwidth is unlikely to be an issue for SMVPs as local PE performance increases.

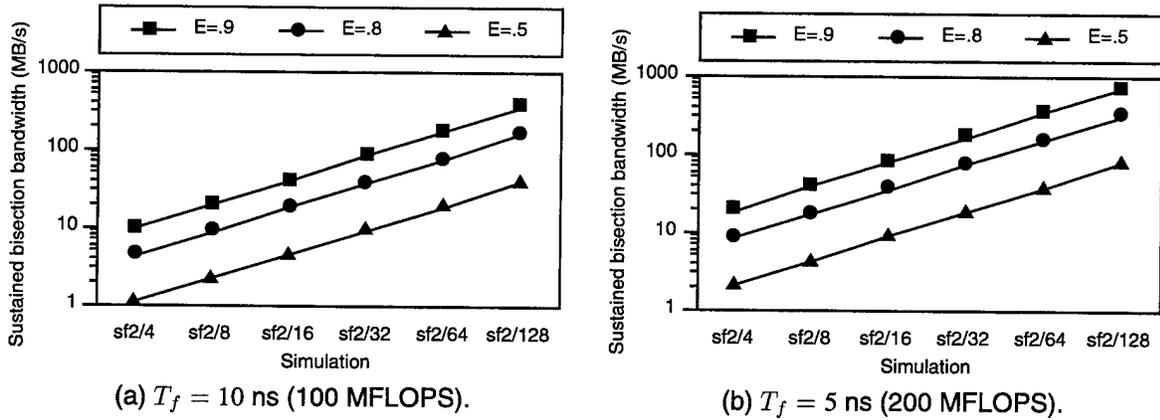
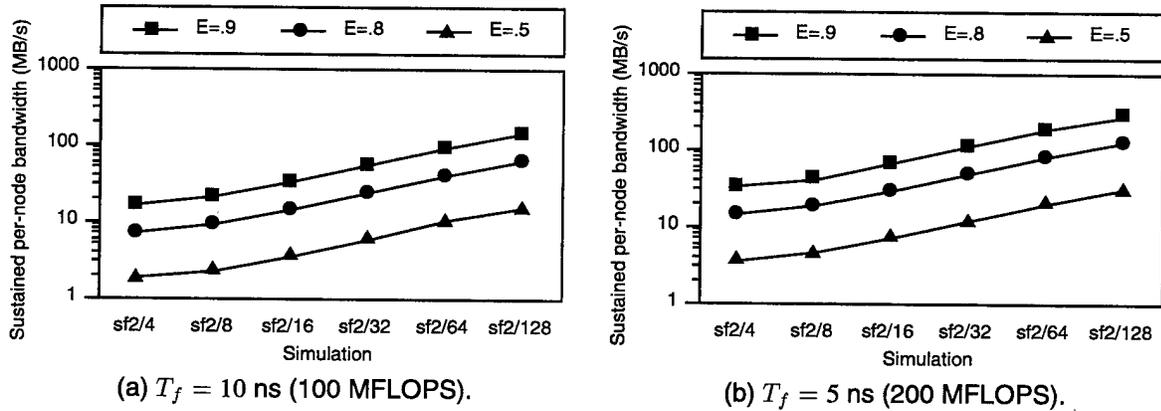


Figure 12: Sustained bisection bandwidths required for sf2.

Figure 13: Sustained PE bandwidth (T_c^{-1}) required for sf2.

5.2 Sustained PE bandwidth

Figure 13 plots the sustained bandwidth for each PE (T_c^{-1}) that is required for the sf2 SMVP under different assumptions of PE performance and overall efficiency. The required bandwidths are computed using Equation (1) and the application properties tabulated in Figure 7. On a system with 100-MFLOP PEs, maintaining a sustained rate of about 150 MBytes/sec per PE during the communication phase is sufficient to run all instances of the sf2 SMVP at 90% efficiency. On systems with 200-MFLOP PEs, a sustained PE bandwidth of about 300 MBytes/sec will be required to run all of the sf2 SMVPs at 90% efficiency.

A sustained PE bandwidth of 300 MBytes/sec seems like an easy performance target until one remembers that it includes all overheads, including software overheads, local strided read and write copies, and other latencies; we are not concerned here with peak link bandwidth ratings. For example, even though the optimal throughput of strided copies on the Cray T3D is 30–40 MBytes/sec [18], current implementations of sf2 achieve at best a measured and sustained bandwidth of 10 MBytes/sec using the C interface to the vendor-supplied MPI library [2].

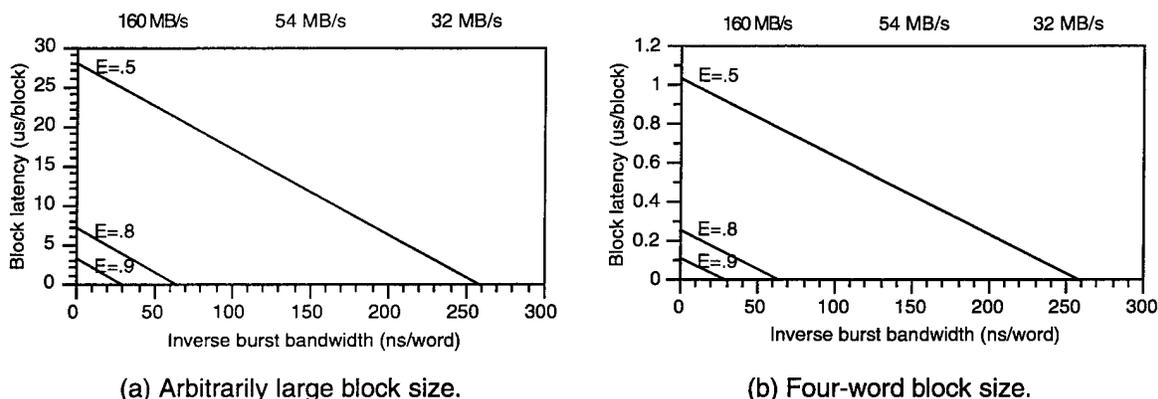


Figure 14: burst bandwidth and latency tradeoffs for sf2/128 (200 MFLOP PE).

Even more daunting is the goal of running the Quake applications at roughly 80% efficiency on high-speed networks of workstations. This goal is attractive because it would make it possible to obtain enough memory to run much larger simulations, but demands sustained per-PE bandwidths of 100 MBytes/sec.

5.3 Bandwidth and latency tradeoffs

There is an interesting tradeoff between the burst bandwidth and latency that are necessary to achieve some target sustained bandwidth on each PE. If block latency is made smaller, then burst bandwidth can be larger, and vice versa. Equation (2) quantifies this tradeoff, which is shown in Figure 14 for the sf2 SMVP running on 128 200-MFLOP PEs. Any point along a diagonal line represents a pairing of burst bandwidth (T_w^{-1}) and latency (T_l) that meets the sustained PE bandwidth requirement (T_c^{-1}) for sf2/128 given in Figure 13(b). The curves are generated by first using Equation (1) to compute the required inverse sustained bandwidth T_c , and then using Equation (2) to define the relationship between T_l and T_w .

Figure 14(a) shows the burst bandwidth and latency tradeoffs under the assumption that blocks are as large as possible, and thus each PE sends at most one block to any other PE during a single communication phase. This is the norm in message passing systems or DSMs that aggregate blocks [6]. The interesting point about this graph is that latency matters for the SMVP. Even if burst bandwidth is driven to infinity, observed block latency must not exceed $3 \mu\text{s}$ if the code is to run at 90% efficiency.

We can also use Equation (2) to get a sense of the latency and bandwidth tradeoffs if blocks are small, fixed-sized objects like cache lines. To model the tradeoff for fixed-sized blocks consisting of four 64-bit words, we set $B_{max} = C_{max}/4$ and again plot block latency as a function of burst bandwidth, in Figure 14(b). Here, block latency is even more crucial, as expected. For example, if burst bandwidth is infinite, then observed block latency must not exceed 100 ns if the SMVP is to run at 90% efficiency.

For a given application, a reasonable network design goal is to achieve values for T_l and T_w^{-1} such that half of the time for the communication phase is due to block latency and the other half to burst bandwidth. We refer to a burst bandwidth chosen thusly as the *half-bandwidth*, and its corresponding latency as the *half-bandwidth latency*. These figures are useful targets for architects of communication systems because overengineering a network's bandwidth cannot relieve the latency constraint by more than a factor of two, and vice versa. Figure 15 plots the half-bandwidths and half-bandwidth latencies for the entire space of sf2 SMVPs, for both the maximal block size and the four-word block size.

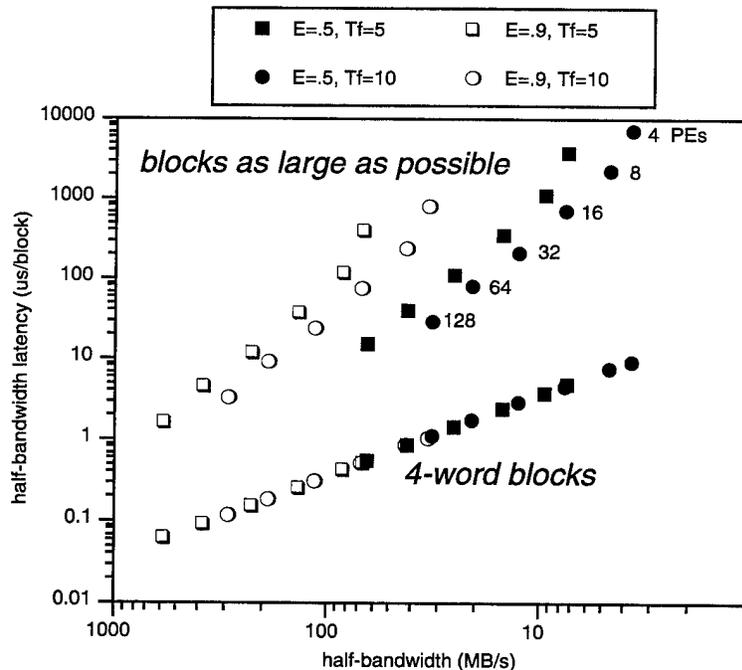


Figure 15: Half-bandwidths and latencies for the sf2 SMVP.

The bottom portion of the graph shows the requirements when blocks are fixed-sized four-word transfer units, as might occur in a shared-memory architecture. For the easiest case, running at 50% efficiency on four 100-MFLOP PEs, we need a burst bandwidth of about 3 MBytes/sec with a block latency of about 10 μ s. In the most difficult case, running at 90% efficiency on 128 200-MFLOP PEs, we would need a burst bandwidth of about 600 MBytes/sec and a block latency of about 70 ns.

The upper portion of the graph shows the requirements when blocks are made as large as possible (i.e., each PE sends at most one block to any other PE), as on a message-passing multiprocessor or network. For the simplest case, running at 50% efficiency on four 100-MFLOP PEs, we would need a burst bandwidth of about 3 MBytes/sec and a half-bandwidth latency of about 8 ms. However, the most demanding case, where we are running at 90% efficiency on 128 200-MFLOP PEs, requires a burst bandwidth of 600 MBytes/sec and a block latency of about 2 μ s.

6 Concluding remarks

Irregular applications have been poorly understood in the past, in large part because researchers and system builders have not had access to large realistic example applications. This report represents a step toward understanding the implications of large-scale irregular codes on the architecture of high-end systems. We have chosen to study unstructured finite element simulations of strong earthquake-induced ground motion because they constitute real and significant engineering applications currently in daily use, and because their performance is characterized by the performance of a simple SMVP kernel. By examining the behavior of the SMVP, we can diagnose the requirements placed on communication systems in future high-end computers and networks as the performance of commodity PEs continues to increase.

Our main conclusions are: (1) Bisection bandwidth is not an issue. (2) Block transfers tend to be

small, even for large applications. Thus we cannot expect to amortize block latency costs with large messages; other latency hiding techniques must be used, or latency must be reduced. (3) Systems with sustained computational throughput of 200 MFLOPS and maximally aggregated blocks will need about 300 MBytes/sec of sustained bandwidth, 600 MBytes/sec of burst bandwidth, and a block latency under 2 μ s to run unstructured finite element applications with 90% efficiency.

The bandwidth requirements per PE for future systems are aggressive, especially since they must include all hardware and software overheads. Yet they seem achievable. More worrisome are the block latency requirements (2 μ s for large blocks, 70 ns for small blocks), which appear much more difficult to achieve. The numbers in this report provide a strong quantitative argument for reducing latency in future systems, either with more efficient interfaces, latency hiding techniques, or both.

Acknowledgments

Many thanks to our Quake project partners from the CMU Department of Civil Engineering: Jacobo Bielak, Omar Ghattas, Hesheng Bao, Loukas Kallivokas, and Jifeng Xu, who developed the San Fernando simulations and generously shared their extensive knowledge. Gary Miller provided insightful comments. Joe Brandenburg and Mike Barton at Intel sparked our interest in characterizing large finite element simulations.

References

- [1] C. Amza, A. Cox, S. Dwarkadas, C. Hyams, Z. Li, and W. Zwaenepoel, *Treadmarks: Shared memory computing on networks of workstations*, IEEE Computer **29** (1996), no. 2, 18–28.
- [2] H. Bao, J. Bielak, O. Ghattas, L. Kallivokas, D. O'Hallaron, J. Shewchuk, and J. Xu, *Large-scale Simulation of Elastic Wave Propagation in Heterogeneous Media on Parallel Computers*, Computer Methods in Applied Mechanics and Engineering **152** (1998), no. 1–2, 85–102.
- [3] S. Barnard and H. Simon, *A fast multilevel implementation of recursive spectral bisection for partitioning unstructured problems*, Tech. Report RNR-92-033, NASA Ames Research Center, November 1992.
- [4] D. Culler, R. Karp, D. Patterson, A. Sahay, E. Santos, K. Schauer, R. Subramonian, and T. von Eicken, *LogP: A practical model of parallel computation*, Communications of the ACM **39** (1996), no. 11, 79–85.
- [5] R. Cypher, A. Ho, S. Konstantinidou, and P. Messina, *Architetural requirements of parallel scientific applications with explicit communication*, Proc. 20th Intl. Symp. Computer Arch., ACM/IEEE, May 1993, pp. 2–13.
- [6] S. Dwarkadas, A. Cox, and W. Zwaenepoel, *An integrated compile-time/run-time software distributed shared memory system*, Proc. Sixth Intl. Conf. on Architectural Support for Prog. Languages and Operating Systems (ASPLOS VI) (Boston, MA), ACM, October 1996, pp. 186–197.
- [7] J. Gilbert, G. Miller, and S. Teng, *Geometric mesh partitioning: Implementation and experiments*, 9th International Parallel Processing Symposium (Santa Barbara, CA), IEEE, April 1995, pp. 418–427.

-
- [8] B. Hendrickson and R. Leland, *The Chaco user's guide Version 2.0*, Tech. Report SAND95-2344, Sandia National Laboratories, July 1995.
- [9] J. L. Hennessy and D. A. Patterson, *Computer architecture: A quantitative approach (2nd edition)*, Morgan Kaufman, 1995.
- [10] C. Holt, J. Singh, and J. Hennessy, *Application and architectural bottlenecks in large scale distributed shared memory machines*, Proc. 23rd Intl. Symp. Comp. Arch. (Philadelphia, PA), ACM, May 1996, pp. 134–145.
- [11] D. Lenoski, J. Laudon, K. Gharachorloo, W. Weber, A. Gupta, J. Hennessy, M. Horowitz, and M. Lam, *The Stanford DASH multiprocessor*, IEEE Computer **25** (1992), no. 3, 63–79.
- [12] G. Miller, S. Teng, W. Thurston, and S. Vavasis, *Automatic mesh partitioning*, Graph Theory and Sparse Matrix Computation (New York) (Alan George, John R. Gilbert, and Joseph W. H. Liu, eds.), Springer-Verlag, New York, 1993.
- [13] MPI Forum, *MPI: A Message Passing Interface*, Proc. Supercomputing '93 (Portland, OR), ACM/IEEE, November 1993, pp. 878–883.
- [14] D. O'Hallaron, *Spark98: Sparse matrix kernels for shared memory and message passing systems*, Tech. Report CMU-CS-97-178, School of Computer Science, Carnegie Mellon University, October 1997.
- [15] D. O'Hallaron and J. Shewchuk, *Properties of a family of parallel finite element simulations*, Tech. Report CMU-CS-96-141, School of Computer Science, Carnegie Mellon University, December 1996.
- [16] S. Scott, *Synchronization and communication in the T3E multiprocessor*, Proc. 7th. Intl. Conf. on Arch. Support for Prog. Lang. and Oper. Systems (Boston, MA), ACM, October 1996, pp. 26–36.
- [17] J. Shewchuk, *Delaunay refinement mesh generation*, Ph.D. thesis, School of Computer Science, Carnegie Mellon University, May 1997, Available as CMU Tech Report CMU-CS-97-137.
- [18] T. Stricker and T. Gross, *Optimizing memory system performance for communication in parallel computers*, Proc. 22nd Intl. Symp. Comp. Arch. (Santa Margherita di Ligure, Italy), ACM, June 1995, pp. 308–319.
- [19] L. Valiant, *A bridging model for parallel computation*, Communications of the ACM **33** (1990), no. 8, 103–111.