



ELSEVIER

Parallel Computing 000 (2000) 000–000

PARALLEL
COMPUTING

www.elsevier.com/locate/parco

2 High-performance parallel implicit CFD

3 William D. Gropp^{a,1}, Dinesh K. Kaushik^{a,2},
4 David E. Keyes^{b,*,3}, Barry F. Smith^a5 ^a Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439, USA6 ^b Mathematics and Statistics Department, Old Dominion University, Norfolk, VA 23529, USA

7 Received 7 August 2000; received in revised form 4 October 2000

8 Abstract

9 Fluid dynamical simulations based on finite discretizations on (quasi)-static grids scale well
10 in parallel, but execute at a disappointing percentage of per-processor peak floating point
11 operation rates without special attention to layout and access ordering of data. We document
12 both claims from our experience with an unstructured grid CFD code that is typical of the
13 state of the practice at NASA. These basic performance characteristics of PDE-based codes
14 can be understood with surprisingly simple models, for which we quote earlier work, pre-
15 senting primarily experimental results. The performance models and experimental results
16 motivate algorithmic and software practices that lead to improvements in both parallel sca-
17 lability and per node performance. This snapshot of ongoing work updates our 1999 Bell
18 Prize-winning simulation on ASCI computers. © 2000 Elsevier Science B.V. All rights re-
19 served.

20 MSC: 65H20; 65N55; 65Y05; 76G25

* Corresponding author. Tel.: +1-757-683-3906; fax: +1-757-683-3885.

E-mail address: dkeyes@odu.edu (D.E. Keyes).

¹ This work was supported in part by the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Advanced Scientific Computing Research, U.S. Department of Energy, under Contract W-31-109-Eng-38.

² This work was supported by a GAANN Fellowship from the U.S. Department of Education and by Argonne National Laboratory under contract 983572401.

³ This work was supported by the National Science Foundation under grant ECS-9527169, by NASA under contracts NAS1-97046 and NAS1-19480 (while the author was in residence at ICASE), by Argonne National Laboratory under contract 982232402, and by Lawrence Livermore National Laboratory under subcontract B347882.

21 *Keywords:* Parallel implicit solvers; Unstructured grids; Computational fluid dynamics; High-performance
22 computing

23 1. PDE application overview

24 Systems modeled by partial differential equations often possess a wide range of
25 time scales – some (or all, in steady-state problems) much faster than the phenomena
26 of interest – suggesting the need for implicit methods. In addition, many applications
27 are geometrically complex, suggesting the convenience of an unstructured mesh for
28 fast automated grid generation. The best algorithms for solving nonlinear implicit
29 problems are often Newton methods, which in turn require the solution of very large,
30 sparse linear systems. The best algorithms for these sparse linear problems, partic-
31 ularly at very large sizes, are often preconditioned iterative methods, of multilevel
32 type if necessary. This nested hierarchy of tunable algorithms has proved effective in
33 solving complex problems in such areas as aerodynamics, combustion, radiation
34 transport, and global circulation.

35 When well tuned, such codes spend almost all of their time in two phases: flux
36 computations (to evaluate conservation law residuals), where one aims to have such
37 codes spent almost *all* their time, and sparse linear algebraic kernels, which are a fact
38 of life in implicit methods. Altogether, four basic groups of tasks can be identified
39 based on the criteria of arithmetic concurrency, communication patterns, and the
40 ratio of operation complexity to data size within the task. These four distinct phases,
41 present in most implicit codes, are vertex-based loops, edge-based loops, recurrences,
42 and global reductions. Each of these groups of tasks stresses a different subsystem of
43 contemporary high-performance computers. Analysis of our demonstration code
44 shows that, after tuning, the linear algebraic kernels run at close to the aggregate
45 memory-bandwidth limit on performance, the flux computations are bounded either
46 by memory bandwidth or instruction scheduling (depending upon the ratio of load/
47 store units to floating-point units in the CPU), and parallel efficiency is bounded
48 primarily by slight load imbalances at synchronization points.

49 Our demonstration application code, FUN3D, is a tetrahedral, vertex-centered
50 unstructured mesh code originally developed by W.K. Anderson of the NASA
51 Langley Research Center for compressible and incompressible Euler and Navier–
52 Stokes equations [1,2]. FUN3D uses a control volume discretization with a variable-
53 order Roe scheme for approximating the convective fluxes and a Galerkin discreti-
54 zation for the viscous terms. FUN3D has been used for design optimization of
55 airplanes, automobiles, and submarines, with irregular meshes comprising several
56 million mesh points. The optimization involves many analyses, typically sequential.
57 Thus, reaching the steady-state solution in each analysis cycle in a reasonable
58 amount of time is crucial to conducting the design optimization. Our best achieve-
59 ment to date for multimillion meshpoint simulations is about 15 μ s per degree-of-
60 freedom for satisfaction of residuals close to machine precision.

61 We have ported FUN3D into the PETSc [3] framework using the single program
 62 multiple data (SPMD) message-passing programming model, supplemented by
 63 multithreading at the physically shared memory level. Thus far, our large-scale
 64 parallel experience with PETSc-FUN3D is with compressible or incompressible
 65 Euler flows, but nothing in the solution algorithms or software changes when ad-
 66 ditional physical phenomenology present in the original FUN3D is included. Of
 67 course, the convergence rate varies with conditioning, as determined by Mach and
 68 Reynolds numbers and the correspondingly induced mesh adaptivity. Robustness
 69 becomes an issue in problems that admit shocks or employ turbulence models. When
 70 nonlinear robustness is restored in the usual manner, through pseudo-transient
 71 continuation, the conditioning of the linear inner iterations is enhanced, and parallel
 72 scalability may be improved. In some sense, the subsonic Euler examples on which
 73 we concentrate, with their smaller number of flops per point per iteration and their
 74 aggressive pseudotransient buildup toward the steady-state limit, may be a *more*
 75 severe test of parallel performance than more physically complex cases.

76 Achieving high sustained performance, in terms of solutions per second, requires
 77 attention to three factors. The first is a scalable implementation, in the sense that
 78 time per iteration is reduced in inverse proportion to the number of processors, or
 79 that time per iteration is constant as problem size and processor number are scaled
 80 proportionally. The second is good per processor performance on contemporary
 81 cache-based microprocessors. The third is algorithmic scalability, in the sense that
 82 the number of iterations to convergence does not grow with increased numbers of
 83 processors. The third factor arises because the requirement of a scalable imple-
 84 mentation generally forces parameterized changes in the algorithm as the number of
 85 processors grows. If the convergence is allowed to degrade, however, the overall
 86 execution is not scalable, and this must be countered algorithmically. These factors
 87 in the overall performance are considered in Sections 3–5, respectively, which are the
 88 heart of this paper. Section 2 first expands on the algorithmics. Section 6 details our
 89 highest performing runs to date, and Section 7 summarizes our work and looks
 90 ahead.

91 2. Ψ NKS: a family of parallel implicit solution algorithms

92 Our implicit algorithmic framework for advancing toward an assumed steady state
 93 for the system of conservation equations, $\mathbf{f}(\mathbf{u}) = 0$, has the form

$$\left(\frac{1}{\Delta t^\ell}\right)\mathbf{u}^\ell + \mathbf{f}(\mathbf{u}^\ell) = \left(\frac{1}{\Delta t^\ell}\right)\mathbf{u}^{\ell-1},$$

95 where $\Delta t^\ell \rightarrow \infty$ as $\ell \rightarrow \infty$, \mathbf{u} represents the fully coupled vector of unknowns, and
 96 $\mathbf{f}(\mathbf{u})$ is the vector of nonlinear conservation laws.

97 Each member of the sequence of nonlinear problems, $\ell = 1, 2, \dots$, is solved with an
 98 inexact Newton method. The resulting Jacobian systems for the Newton corrections
 99 are solved with a Krylov method, relying directly only on matrix-free Jacobian-
 100 vector product operations. The Krylov method needs to be preconditioned for ac-

101 ceptable inner iteration convergence rates, and the preconditioning can be the
 102 “make-or-break” feature of an implicit code. A good preconditioner saves time and
 103 space by permitting fewer iterations in the Krylov loop and smaller storage for the
 104 Krylov subspace. An additive Schwarz preconditioner [5] accomplishes this in a
 105 concurrent, localized manner, with an approximate solve in each subdomain of a
 106 partitioning of the global PDE domain. The coefficients for the preconditioning
 107 operator are derived from a lower-order, sparser and more diffusive discretization
 108 than that used for $\mathbf{f}(\mathbf{u})$, itself. Applying any preconditioner in an additive Schwarz
 109 manner tends to increase flop rates over the same preconditioner applied globally,
 110 since the smaller subdomain blocks maintain better cache residency, even apart from
 111 concurrency considerations [28]. Combining a Schwarz preconditioner with a Krylov
 112 iteration method inside an inexact Newton method leads to a synergistic, parallel-
 113 izable nonlinear boundary value problem solver with a classical name: Newton–
 114 Krylov–Schwarz (NKS) [12]. We combine NKS with pseudo-timestepping [17] and
 115 use the shorthand Ψ NKS to describe the algorithm.

116 To implement Ψ NKS in FUN3D, we employ the PETSc package [3], which fea-
 117 tures distributed data structures – index sets, vectors, and matrices – as fundamental
 118 objects. Iterative linear and nonlinear solvers are implemented within PETSc in a
 119 data structure-neutral manner, providing a uniform application programmer inter-
 120 face. Portability is achieved through MPI, but message-passing detail is not required
 121 in the application. We use MeTiS [16] to partition the unstructured mesh.

122 The basic philosophy of any efficient parallel computation is “owner computes,”
 123 with message merging and overlap of communication with computation where
 124 possible via split transactions. Each processor “ghosts” its stencil dependencies on its
 125 neighbors’ data. Grid functions are mapped from a global (user) ordering into
 126 contiguous local orderings (which, in unstructured cases, are designed to maximize
 127 spatial locality for cache line reuse). Scatter/gather operations are created between
 128 local sequential vectors and global distributed vectors, based on runtime-deduced
 129 connectivity patterns.

130 As mentioned above, there are four groups of tasks in a typical PDE solver, each
 131 with a distinct proportion of work to datasize to communication requirements. In
 132 the language of a vertex-centered code, in which the data are stored at cell vertices,
 133 these tasks are as follows:

- Vertex-based loops
 - state vector and auxiliary vector updates
- Edge-based “stencil op” loops
 - residual evaluation, Jacobian evaluation
 - Jacobian-vector product (often replaced with matrix-free form, involving
 139 residual evaluation)
 - interpolation between grid levels
- Sparse, narrow-band recurrences
 - (approximate) factorization, back substitution, relaxation/smoothing
- Vector inner products and norms
 - orthogonalization/conjugation
 - convergence progress checks and stability heuristics.

146 Vertex-based loops are characterized by work closely proportional to datasize,
147 pointwise concurrency, and no communication.

148 Edge-based “stencil op” loops have a large ratio of work to datasize, since each
149 vertex is used in many discrete stencil operations, and each degree of freedom at a
150 point (momenta, energy, density, species concentration) generally interacts with all
151 others in the conservation laws – through constitutive and state relationships or
152 directly. There is concurrency at the level of the number of edges between vertices
153 (or, at worst, the number of edges of a given “color” when write consistency needs to
154 be protected through mesh coloring). There is local communication between pro-
155 cessors sharing ownership of the vertices in a stencil.

156 Sparse, narrow-band recurrences involve work closely proportional to data size,
157 the matrix being the largest data object and each of its elements typically being used
158 once. Concurrency is at the level of the number of fronts in the recurrence, which
159 may vary with the level of exactness of the recurrence. In a preconditioned iterative
160 method, the recurrences are typically broken to deliver a prescribed process con-
161 currence; only the quality of the preconditioning is thereby affected, not the final
162 result. Depending upon whether one uses a pure decomposed Schwarz-type pre-
163 conditioner, a truncated incomplete solve, or an exact solve, there may be no, local
164 only, or global communication in this task.

165 Vector inner products and norms involve work closely proportional to data size,
166 mostly pointwise concurrency, and global communication. Unfortunately, inner
167 products and norms occur rather frequently in stable, robust linear and nonlinear
168 methods.

169 Based on these characteristics, one anticipates that vertex-based loops, recur-
170 rences, and inner products will be *memory bandwidth limited*, whereas edge-based
171 loops are likely to be only *load/store limited*. However, edge-based loops are vul-
172 nerable to *internode bandwidth* if the latter does not scale. Inner products are vul-
173 nerable to *internode latency* and *network diameter*. Recurrences can resemble some
174 combination of edge-based loops and inner products in their communication char-
175 acteristics if preconditioning fancier than simple Schwarz is employed. For instance,
176 if incomplete factorization is employed globally or a coarse grid is used in a multi-
177 level preconditioner, global recurrences ensue.

178 3. Implementation scalability

179 Domain-decomposed parallelism for PDEs is a natural means of overcoming
180 Amdahl’s law in the limit of fixed problem size per processor. Computational work
181 on each evaluation of the conservation residuals scales as the volume of the (equal-
182 sized) subdomains, whereas communication overhead scales only as the surface. This
183 ratio is fixed when problem size and processors are scaled in proportion, leaving only
184 global reduction operations over all processors as an impediment to perfect per-
185 formance scaling.

186 In [18], it is shown that on contemporary tightly coupled parallel architectures in
187 which the number of connections between processors grows in proportion to the

188 number of processors, such as meshes and tori, aggregate internode bandwidth is
 189 more than sufficient, and limits to scalability may be determined by a balance of
 190 work per node to synchronization frequency. On the other hand, if there is nearest-
 191 neighbor communication contention, in which a fixed resource like an internet switch
 192 is divided among all processors, the number of processors is allowed to grow only as
 193 the one-fourth power of the problem size (in three dimensions). This is a curse of
 194 typical Beowulf-type clusters with inexpensive networks; we do not discuss the
 195 problem here, although it is an important practical limitation in many CFD groups.

196 When the load is perfectly balanced (which is easy to achieve for static meshes)
 197 and local communication is not an issue because the network is scalable, the optimal
 198 number of processors is related to the network diameter. For logarithmic networks,
 199 like a hypercube, the optimal number of processors, P , grows directly in proportion
 200 to the problem size, N . For a d -dimensional torus network, $P \propto N^{d/d+1}$. The pro-
 201 portionality constant is a ratio of work per subdomain to the product of synchrono-
 202 nization frequency and internode communication latency.

203 3.1. Scalability bottlenecks

204 In Table 1, we present a closer look at the relative cost of computation for PETSc-
 205 FUN3D for a fixed-size problem of 2.8 million vertices on the ASCI Red machine,
 206 from 128 to 3072 nodes. The intent here is to identify the factors that retard the
 207 scalability.

208 From Table 1, we observe that the buffer-to-buffer time for global reductions for
 209 these runs is relatively small and does not grow on this excellent network. The
 210 primary factors responsible for the increased overhead of communication are the
 211 implicit synchronizations and the ghost point updates (interprocessor data scatters).

212 Interestingly, the increase in the percentage of time (3–10%) for the scatters results
 213 more from algorithmic issues than from hardware/software limitations. With an
 214 increase in the number of subdomains, the percentage of grid point data that must be
 215 communicated also rises. For example, the total amount of nearest neighbor data
 216 that must be communicated per iteration for 128 subdomains is 2 gigabytes, while
 217 for 3072 subdomains it is 8 gigabytes. Although more network wires are available
 218 when more processors are employed, scatter time increases. If problem size and

Table 1
 Scalability bottlenecks on ASCI Red for a fixed-size 2.8 M-vertex case^a

Number of processors	Percentage of time		
	Global reductions	Implicit synchronizations	Ghost point scatters
128	5	4	3
512	3	7	5
3072	5	14	10

^a The preconditioner used in these results is block Jacobi with ILU(1) in each subdomain. We observe that the principal nonscaling factor is the implicit synchronization.

219 processor count are scaled together, we would expect scatter times to occupy a fixed
 220 percentage of the total and load imbalance to be reduced at high granularity.

221 3.2. Effect of partitioning strategy

222 Mesh partitioning has a dominant effect on parallel scalability for problems
 223 characterized by (almost) constant work per point. As shown above, poor load
 224 balance causes idleness at synchronization points, which are frequent in implicit
 225 methods (e.g., at every conjugation step in a Krylov solver). With NKS methods,
 226 then, it is natural to strive for a very well balanced load. The p-MeTiS algorithm in
 227 the MeTiS package [16], for example, provides almost perfect balancing of the
 228 number of mesh points per processor. However, balancing work alone is not suffi-
 229 cient. Communication must be balanced as well, and these objectives are not entirely
 230 compatible. Fig. 1 shows the effect of data partitioning using p-MeTiS, which tries to
 231 balance the number of nodes and edges on each partition, and k-MeTiS, which tries
 232 to reduce the number of noncontiguous subdomains and connectivity of the sub-
 233 domains. Better overall scalability is observed with k-MeTiS, despite the better load
 234 balance for the p-MeTiS partitions. This is due to the slightly poorer numerical
 235 convergence rate of the iterative NKS algorithm with the p-MeTiS partitions. The
 236 poorer convergence rate can be explained by the fact that the p-MeTiS partitioner
 237 generates disconnected pieces within a single “subdomain,” effectively increasing the
 238 number of blocks in the block Jacobi or additive Schwarz algorithm and increasing
 239 the size of the interface. The convergence rates for block iterative methods degrade
 240 with increasing number of blocks, as discussed in Section 5.

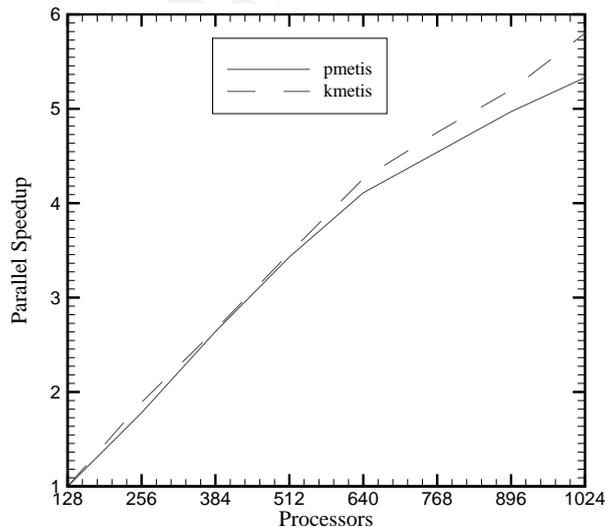


Fig. 1. Parallel speedup relative to 128 processors on a 600 MHz Cray T3E for a 2.8 M-vertex case, showing the effect of partitioning algorithms *k*-MeTiS, and *p*-MeTiS.

241 *3.3. Domain-based and/or instruction-level parallelism*

242 The performance results above are based on subdomain parallelism using the
 243 message passing interface (MPI) [13]. With the availability of large scale SMP
 244 clusters, different software models for parallel programming require a fresh assess-
 245 ment. For machines with physically distributed memory, MPI has been a natural and
 246 successful software model. For machines with distributed shared memory and
 247 nonuniform memory access, both MPI and OpenMP have been used with respect-
 248 able parallel scalability. For clusters with two or more SMPs on a single node, the
 249 mixed software model of threads within a node (OpenMP being a special case of
 250 threads because of the potential for highly efficient handling of the threads and
 251 memory by the compiler) and MPI between the nodes appears natural. Several re-
 252 searchers (e.g., [4,20]) have used this mixed model with reasonable success.

253 We investigate the mixed model by employing OpenMP only in the flux calcula-
 254 tion phase. This phase takes over 60% of the execution time on ASCI Red and is an
 255 ideal candidate for shared-memory parallelism because it does not suffer from the
 256 memory bandwidth bottleneck (see Section 4). In Table 2, we compare the perfor-
 257 mance of this phase when the work is divided by using two OpenMP threads per
 258 node with the performance when the work is divided using two independent MPI
 259 processes per node. There is no communication in this phase. Both processors work
 260 with the same amount of memory available on a node; in the OpenMP case, it is
 261 shared between the two threads, while in the case of MPI it is divided into two
 262 address spaces.

263 The hybrid MPI/OpenMP programming model appears to be a more efficient way
 264 to employ shared memory than are the heavyweight subdomain-based processes
 265 (MPI alone), especially when the number of nodes is large. The MPI model works
 266 with larger number of subdomains (equal to the number of MPI processors), re-
 267 sulting in slower rate of convergence. The hybrid model works with fewer chunkier
 268 subdomains (equal to the number of nodes) that result in faster convergence rate and
 269 shorter execution time, despite the fact that there is some redundant work when the
 270 data from the two threads are combined due to the lack of a vector-reduce operation
 271 in the OpenMP standard (version 1) itself. Specifically, some redundant work arrays
 272 must be allocated that are not present in the MPI code. The subsequent gather

Table 2
 Execution time on the 333 MHz Pentium Pro ASCI Red machine for function evaluations only for a 2.8 M-vertex case, showing differences in exploiting the second processor sharing the same memory with either OpenMP instruction-level parallelism (number of subdomains equals the number of nodes) or MPI domain-level parallelism (number of subdomains is equal to the number of processes per node)

Nodes	MPI/OpenMP threads per node (s)		MPI processes per node (s)	
	1	2	1	2
256	483	261	456	258
2560	76	39	72	45
3072	66	33	62	40

273 operations (which tend to be memory bandwidth bound) can easily offset the ad-
 274 vantages accruing from the low-latency shared-memory communication. One way to
 275 get around this problem is to use coloring strategies to create the disjoint work sets,
 276 but this takes away the ease and simplicity of the parallelization step promised by the
 277 OpenMP model.

278 **4. Single-processor performance modeling and tuning**

279 In this section, we describe the details of per processor performance and tuning.
 280 Since the gap between memory and CPU speeds is ever widening [14] and algo-
 281 rithmically optimized PDE codes do relatively little work per data item, it is crucial
 282 to efficiently utilize the data brought into the levels of memory hierarchy that are
 283 close to the CPU. To achieve this goal, the data structure storage patterns for pri-
 284 mary (e.g., momenta and pressure) and auxiliary (e.g., geometry and constitutive
 285 parameter) fields should adapt to hierarchical memory. Three simple techniques have
 286 proved very useful in improving the performance of the FUN3D code, which was
 287 originally tuned for vector machines. These techniques are interlacing, blocking, and
 288 edge reordering. They are within the scope of automated compiler transformations
 289 in structured grid codes but, so far must be implemented manually in unstructured
 290 codes.

291 *4.1. Interlacing, blocking, and edge reordering*

292 Table 3 shows the effectiveness of interlacing, block, and edge reordering (de-
 293 scribed below) on one processor of the SGI Origin2000. The combination of the
 294 three effects can enhance overall execution time by a factor of 5.7. To further un-
 295 derstand the dramatic effect of reordering the edges, we carried out hardware counter
 296 profiling on the R10000 processor. Fig. 2 shows that edge reordering reduces the

Table 3
 Execution times for Euler flow over M6 wing for a fixed-size grid of 22,677 vertices (90,708 DOFs in-
 compressible; 113,385 DOFs compressible)^a

Enhancements			Results			
Field interlacing	Structural blocking	Edge reordering	Incompressible		Compressible	
			Time/step (s)	Ratio	Time/step (s)	Ratio
			83.6	–	140.0	–
×			36.1	2.31	57.5	2.44
×	×		29.0	2.88	43.1	3.25
		×	29.2	2.86	59.1	2.37
		×	23.4	3.57	35.7	3.92
×	×	×	16.9	4.96	24.5	5.71

^aThe processor is a 250 MHz MIPS R10000. Activation of a layout enhancement is indicated by “×” in the corresponding column.

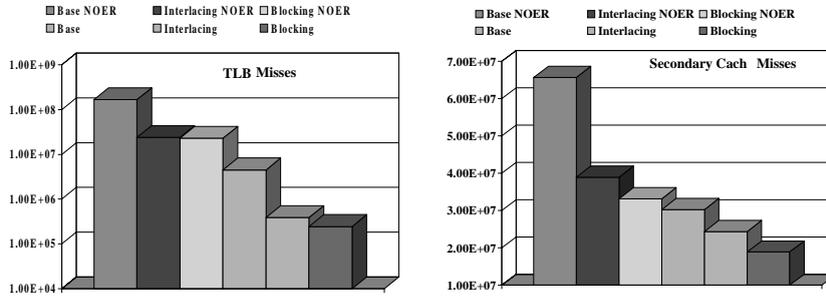


Fig. 2. TLB misses (log scale) and secondary cache misses (linear scale) for a 22,677-vertex case on a 250 MHz R10000 processor, showing dramatic improvements in data locality due to data ordering and blocking techniques. (“NOER” denotes no edge ordering; otherwise edges are reordered by default.)

Table 4

Comparison of optimized to original performance, absolute and as percentage of peak, for PETS-FUN3D on many processor families

Processor	Clock	Peak	Opt. MF/s	Orig. MF/s	Opt.% peak	Orig.% peak	Ratio
R10000	250	500	127	26	25.4	5.2	4.9
RS6000/P3	200	800	163	32	20.3	4.0	5.1
RS6000/P2	120	480	117	15	24.3	3.1	7.8
RS6000/604e	333	666	66	15	9.9	2.3	4.4
Pentium Pro	333	333	60	21	18.8	6.3	3.0
Alpha 21164	600	1200	91	16	7.6	1.3	5.7
Ultra II	400	800	71	20	8.9	2.5	3.6

297 misses in the translation lookaside buffer (TLB) cache by two orders of magnitude,
 298 while secondary cache misses (which are very expensive) are reduced by a factor of
 299 3.5. (The TLB cache is used in virtual memory address translation.)

300 Table 4 compares the original and optimized per processor performance for sev-
 301 eral other architectures. The ratio of improvement in the last column varies from 2.6
 302 to 7.8. Improvement ratios are averages over the entire code; different subroutines
 303 benefit to different degrees.

304 4.1.1. Field interlacing

305 Field interlacing creates the spatial locality for the data items needed successively
 306 in time. This is achieved by choosing

$$u1, v1, w1, p1, u2, v2, w2, p2, \dots$$

308 in place of

$$u_1, u_2, \dots, v_1, v_2, \dots, w_1, w_2, \dots, p_1, p_2, \dots$$

310 for a calculation that uses u, v, w, p together. We denote the first ordering “inter-
311 laced” and the second “noninterlaced.” The noninterlaced storage pattern is good
312 for vector machines. For cache-based architectures, the interlaced storage pattern
313 has many advantages: (1) it provides high reuse of data brought into the cache, (2) it
314 makes the memory references closely spaced, which in turn reduces the TLB misses,
315 and (3) it decreases the size of the working set of the data cache(s), which reduces the
316 number of conflict misses.

317 4.1.2. Structural blocking

318 Once the field data are interlaced, it is natural to use a block storage format for the
319 Jacobian matrix of a multicomponent system of PDEs. The block size is the number
320 of components (unknowns) per mesh point. As shown for the sparse matrix–vector
321 case in [10], this structural blocking significantly reduces the number of integer loads
322 and enhances the reuse of the data items in registers. It also reduces the memory
323 bandwidth required for optimal performance.

324 4.1.3. Edge and node reorderings

325 In the original FUN3D code, the edges are colored for good vector performance.
326 No pair of nodes in the same discretization stencil share a color. This strategy results
327 in a very low cache line reuse. In addition, since consecutive memory references may
328 be far apart, the TLB misses are a grave concern. About 70% of the execution time in
329 the original vector code is spent serving TLB misses. As shown in Fig. 2, this
330 problem is effectively addressed by reordering the edges.

331 The edge reordering we have used sorts the edges in increasing order by the node
332 number at the one end of each edge. In effect, this converts an edge-based loop into a
333 vertex-based loop that reuses vertex-based data items in most or all of the stencils
334 that reference them several times before discarding it. Since a loop over edges goes
335 over a node’s neighbors first, edge reordering (in conjunction with a bandwidth
336 reducing ordering for nodes) results in memory references that are closely spaced.
337 Hence, the number of TLB misses is reduced significantly. For vertex ordering, we
338 have used the Reverse Cuthill McKee (RCM) [7], which is known in the linear al-
339 gebra literature to reduce cache misses by enhancing spatial locality.

340 4.2. Performance analysis of the sparse matrix–vector product

341 The sparse matrix–vector product (or “matvec”) is an important part of many
342 iterative solvers in its own right, and also representative of the data access patterns of
343 explicit grid-based stencil operations and recurrences. While detailed performance
344 modeling of this operation can be complex, particularly when data reference patterns
345 are included [26,27,29], a simplified analysis can still yield upper bounds on the
346 achievable performance of this operation.

347 In [10], we estimate the memory bandwidth required by sparse matvecs in un-
 348 structured grid codes, after making some simplifying assumptions that idealize the
 349 rest of the memory system. We assume that there are no conflict misses, meaning that
 350 each matrix and vector element is loaded into cache only once until flushed by ca-
 351 pacity misses. We also assume that the processor never waits on a memory reference;
 352 that is, that any number of loads and stores are satisfied in a single cycle.

353 The matrix is stored in compressed rows (equivalent to PETSc's AIJ format) or
 354 block AIJ (BAIJ format) [3]. For each nonzero in the matrix, we transfer one integer
 355 (giving the column incidence) and two doubles (the matrix element and the corre-
 356 sponding row vector element), and we do one floating-point multiply-add (fmadd)
 357 operation (which is two flops). Finally, we store the output vector element. Including
 358 loop control and addressing overheads, this leads (see [10]) to a data volume estimate
 359 of 12.36 bytes per fmadd operation for a sample PETSc-FUN3D sparse Jacobian.
 360 This gives us an estimate of the bandwidth required in order for the processor to do
 361 all $2 * n_{nz}$ flops at its peak speed, where n_{nz} is the number of nonzeros in the Jacobian.
 362 Unfortunately, bandwidth as measured by the STREAM [21] benchmark is typically
 363 an order of magnitude less. Alternatively, given a measured memory bandwidth
 364 rating, we can predict the maximum achievable rate of floating-point operations.
 365 Finally, we can measure the achieved floating-point operations. The last four col-
 366 umns of Table 5 summarize the results of this combined theoretical/experimental
 367 study for a matrix with 90,708 rows and 5,047,120 nonzero entries from a PETSc-
 368 FUN3D application (incompressible) with four unknowns per vertex. For this ma-
 369 trix, with a block size of four, the column incidence array is smaller by a factor of the
 370 block size. We observe that the blocking helps significantly by reducing the memory
 371 bandwidth requirement. In [10], we also describe how multiplying more than one
 372 vector at a time requires less memory bandwidth per matvec because of reuse of
 373 matrix elements. We can multiply four vectors in about 1.5 times the time needed to
 374 multiply a single vector. If the three additional vectors can be employed in a block
 375 Krylov method, they are almost free, so algorithmic work on block-Krylov methods
 376 is highly recommended.

377 To further incriminate memory bandwidth as the bottleneck to the execution time
 378 of sparse linear solvers, we have performed an experiment that effectively doubles the
 379 available memory bandwidth. The linear solver execution time is dominated by the

Table 5
 Effect of memory bandwidth on the performance of sparse matrix–vector products on a 250 MHz R10000 processor^a

Format	Bytes/fmadd	Bandwidth (MB/s)		Mflop/s	
		Required	Achieved	Ideal	Achieved
AIJ	12.36	3090	276	58	45
BAIJ	9.31	2327	280	84	55

^a The STREAM benchmark memory bandwidth [21] is 358 MB/s; this value of memory bandwidth is used to calculate the ideal Mflop/s. The achieved values of memory bandwidth and Mflop/s are measured using hardware counters.

380 cost of preconditioning when (as in our production PETSc-FUN3D code) the Ja-
 381 cobian-vector products required in the Krylov methods are performed in a matrix-
 382 free manner by finite-differencing a pair of flux evaluations. Since the preconditioning
 383 is already very approximate, we have implemented the data structures storing
 384 PETSc's preconditioners in single precision while preserving double-precision in all
 385 other parts of the code. Once an element of the preconditioner is in the CPU, it is
 386 padded to 64 bits with trailing zeros, and all arithmetic is done with this (consistent
 387 but inaccurate) double precision value. The consistency is required to suppress the
 388 contamination of the Krylov space with roundoff errors. The loss of accuracy in the
 389 preconditioner is irrelevant to the final result, which satisfies the true linearized
 390 Newton correction equation to required precision, and it is nearly irrelevant to the
 391 convergence rate of the preconditioned iteration. However, it is very relevant to the
 392 execution time, as shown in Table 6. Asymptotically, as the preconditioner matrix
 393 becomes the dominant noncacheable object in the workingset, the running time of
 394 the linear solution is halved, as evidenced by a comparison of columns 2 and 3 in
 395 Table 6.

396 The importance of memory bandwidth to the overall performance is suggested by
 397 the single-processor performance of PETSc-FUN3D shown in Fig. 3. The performance
 398 of PETSc-FUN3D is compared with the peak performance and the result of
 399 the STREAM benchmark [21], which measures achievable performance for memory
 400 bandwidth limited computations. These comparisons show that the STREAM results
 401 are much better indicators of realized performance than the peak numbers. The
 402 parts of the code that are memory bandwidth-limited (like the sparse triangular
 403 preconditioner solution phase, which is responsible for about 25% of the overall
 404 execution time) are bound to show poor performance, as compared with dense
 405 matrix–matrix operations, which often achieve 80–90% of peak.

406 The importance of reducing the memory bandwidth requirements of algorithms is
 407 emphasized by reference to the hardware profiles of the ASCI machines, which are
 408 scheduled to reach a peak of 100 Tflop/s by 2004. Table 7 shows the peak processing
 409 and memory bandwidth capacities of the first four of these machines. The “white”
 410 machine is being delivered to the US Department of Energy at the time of this
 411 writing. The “blue” and “red” machines rank in the top three spots of the Top 500
 412 installed computers as of June 2000 [9]. The last column shows that memory

Table 6

Execution times on a 250 MHz R10000 processor for the linear algebra phase of a 357,900-vertex case with single- or double-precision storage of the preconditioner matrix

Number of processors	Computational phase			
	Linear solve (s)		Overall (s)	
	Double	Single	Double	Single
16	223	136	746	657
64	60	34	205	181
120	31	16	122	106

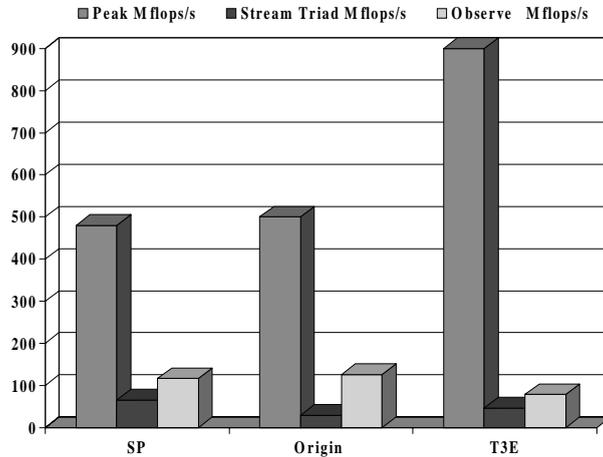


Fig. 3. Sequential performance of PETSc-FUN3D for a 22,677-vertex case.

Table 7
Peak processing and memory bandwidth profiles of the ASCI machines

Platform	Number procs.	Sys. peak (TF/s)	Proc. peak (MF/s)	BW/proc. (MB/s)	BW/proc. (MW/s)	(MF/s)/(MW/s)
White	8192	12.3	1500	1000	125.0	12.0
BlueMtn	6144	3.1	500	390	48.8	10.2
BluePac	5808	3.9	666	360	45.0	14.8
Red	9632	3.2	333	266	33.3	10.0

413 bandwidth, in double precision words per second, is off by an order of magnitude
 414 from what is required if each cached word is used only once. As the raw speed of the
 415 machines is increased, this ratio does not improve. Therefore, algorithms must im-
 416 prove to emphasize locality. Several proposals for discretization and solution
 417 methods that improve spatial or temporary locality are made in [19]. Many of these
 418 require special features in memory control hardware and software that exist today
 419 but are not commonly exploited by computational modelers in high-level scientific
 420 languages.

421 *4.3. Performance analysis of the flux calculation*

422 Even parts of the code that are not memory intensive often achieve much less than
 423 peak performance because of the limits on the number of basic operations that can
 424 be performed in a single clock cycle [10]. This is true for the flux calculation routine
 425 in PETSc-FUN3D, which consumes approximately 60% of the overall execution
 426 time.

427 While looping over each edge, the flow variables from the vertex-based arrays are
 428 read, many floating-point operations are done, and residual values at each node of

429 the edge are updated. Because of the large number of floating-point operations in
430 this phase, memory bandwidth is not (yet) a limiting factor on machines at the high
431 end. Measurements on our Origin2000 support this; only 57 MB/s are needed to keep
432 the flux calculation phase at full throttle [10]. However, the measured floating-point
433 performance is still just 209 Mflop/s out of a theoretical peak of 500 Mflop/s. This is
434 substantially less than the performance that can be achieved with dense matrix–
435 matrix operations.

436 To understand where the limit on the performance of this part of the code comes
437 from, we take a close look at the assembly code for the flux calculation function.
438 This examination yields the following workload mix for the average iteration of the
439 loop over edges: 519 total instructions, 111 integer operations, 250 floating-point
440 instructions of which there are 55 are `fmadd` instructions (for $195 + 2 \times 55 = 305$
441 flops), and 155 memory references. Most contemporary processors can issue only
442 one load or store in one cycle. Since the number of floating-point instructions is less
443 than the number of memory references, the code is bound to take at least as many
444 cycles as the number of loads and stores.

445 If all operations could be scheduled optimally for this hardware – say, one
446 floating-point instruction, one integer instruction, and one memory reference per
447 cycle – this code would take 250 instructions and achieve 305 Mflop/s. However,
448 dependencies between these instructions, as well as complexities in scheduling the
449 instructions [22,24], make it very difficult for the programmer to determine the
450 number of cycles that this code would take to execute. Fortunately, many compilers
451 provide this information as comments in the assembly code. For example, on the
452 Origin2000, when the code is compiled with cache optimizations turned off (con-
453 sistent with our assumption that data items are in primary cache for the purpose of
454 estimating this bound), the compiler estimates that the above work can be completed
455 in about 325 cycles. This leads to a theoretical performance bound of 235 Mflop/s
456 (47% of the peak on the 250 MHz dual-issue processor). We actually measure 209
457 Mflop/s using hardware counters. This shows that the performance in this phase of
458 the computation is restricted by the instruction scheduling limitation. A detailed
459 analytical model for this phase of computation is under way.

460 4.4. Performance comparison

461 In Fig. 4, we compare three performance bounds: the peak performance (based on
462 the clock frequency and the maximum number of floating-point operations per cy-
463 cle), the performance predicted from the memory bandwidth limitation, and the
464 performance based on operation issue limitation. For the sparse matrix–vector
465 multiply, it is clear that the memory-bandwidth limit on performance is a good
466 approximation. The greatest differences between the performance observed and
467 predicted by memory bandwidth are on the systems with the smallest caches (IBM
468 SP and T3E), where our assumption that there are no conflict misses is least likely to
469 be valid.

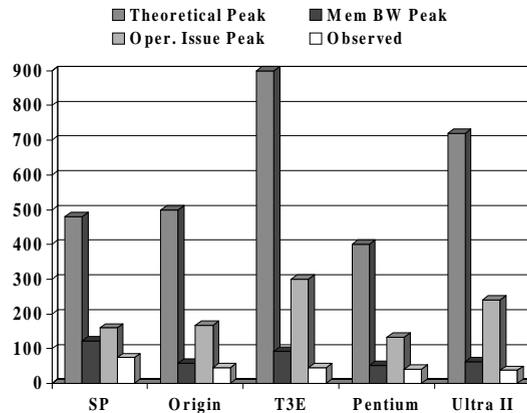


Fig. 4. Three performance bounds for sparse matrix–vector product; the bounds based on memory bandwidth and instruction scheduling are much more closer to the observed performance than the theoretical peak of the processor. Memory bandwidth values are taken from the STREAM benchmark Website.

470 5. Convergence scalability

471 The convergence rates and, therefore, the overall parallel efficiencies of additive
 472 Schwarz methods are often dependent on subdomain granularity. Except when ef-
 473 fective coarse-grid operators and intergrid transfer operators are known, so that
 474 optimal multilevel preconditioners can be constructed, the number of iterations to
 475 convergence tends to increase with granularity for elliptically controlled problems,
 476 for either fixed or memory-scaled problem sizes. In practical large-scale applications,
 477 however, the convergence rate degradation of single-level additive Schwarz is
 478 sometimes not as serious as the scalar, linear elliptic theory would suggest. Its effects
 479 are mitigated by several factors, including pseudo-transient nonlinear continuation
 480 and dominant intercomponent coupling. The former parabolizes the operator, en-
 481 dowing diagonal dominance. The latter renders the off-diagonal coupling less critical
 482 and, therefore, less painful to sever by domain decomposition. The block diagonal
 483 coupling can be captured fully in a point-block ILU preconditioner.

484 5.1. Convergence of Schwarz methods

485 For a general exposition of Schwarz methods for linear problems, see [25]. Assume
 486 a d -dimensional isotropic problem. Consider a unit aspect ratio domain with quasi-
 487 uniform mesh parameter h and quasi-uniform subdomain diameter H . Then problem
 488 size $N = h^{-d}$, and, under the one-subdomain-per-processor assumption, processor
 489 number $P = H^{-d}$. Consider four preconditioners: point Jacobi, subdomain Jacobi, 1-
 490 level additive Schwarz (subdomain Jacobi with overlapped subdomains), and 2-level
 491 additive Schwarz (overlapped subdomains with a global coarse problem with ap-
 492 proximately one degree-of-freedom per subdomain). The first two can be thought of

493 as degenerate Schwarz methods (with zero overlap, and possibly point-sized sub-
 494 domains). Consider acceleration of the Schwarz method by a Krylov method such as
 495 conjugate gradients or one of its many generalizations to nonsymmetric problems
 496 (e.g., GMRES). Krylov–Schwarz iterative methods typically converge in a number
 497 of iterations that scales as the square-root of the condition number of the Schwarz-
 498 preconditioned system. Table 8 lists the expected number of iterations to achieve a
 499 given reduction ratio in the residual norm. The first line of this table pertains to
 500 diagonally scaled CG, a common default parallel implicit method, but one that is *not*
 501 very algorithmically scalable, since the iteration count degrades with a power of N .
 502 The results in this table were first derived for symmetric definite operators with exact
 503 solves on each subdomain, but they have been extended to operators with non-
 504 symmetric and indefinite components and inexact solves on each subdomain.

505 The intuition behind this table is the following: errors propagate from the interior
 506 to the boundary in steps that are proportional to the largest implicit aggregate in the
 507 preconditioner, whether pointwise (in N) or subdomainwise (in P). The use of
 508 overlap avoids the introduction of high-energy-norm solution near discontinuities at
 509 subdomain boundaries. The 2-level method projects out low-wave number errors at
 510 the price of solving a global problem.

511 Only the 2-level method scales perfectly in convergence rate (constant, indepen-
 512 dent of N and P), like a traditional multilevel iterative method. However, the 2-level
 513 method shares with multilevel methods a non-scalable cost-per-iteration from the
 514 necessity of solving a coarse-grid system of size $O(P)$. Unlike recursive multilevel
 515 methods, a 2-level Schwarz method may have a rather fine coarse grid, for example,
 516 $H = O(h^{1/2})$, which makes it less scalable overall. Parallelizing the coarse grid solve
 517 is necessary. Neither extreme of a fully distributed or a fully redundant coarse solve
 518 is optimal, but rather something in between.

519 *5.2. Algorithmic tuning for Ψ NKS solver*

520 The following is an incomplete list of parameters that need to be tuned in various
 521 phases of a pseudo-transient Newton–Krylov–Schwarz algorithm.

- Nonlinear robustness continuation parameters: discretization order, initial time-
 523 step, exponent of timestep evolution law.

Table 8

Iteration count scaling of Schwarz-preconditioned Krylov methods, translated from the theory into problem size N and processor number P , assuming quasi-uniform grid, quasi-unit aspect ratio grid and decomposition, and quasi-isotropic operator

Preconditioning	Iteration count	
	In 2D	In 3D
Point Jacobi	$O(N^{1/2})$	$O(N^{1/3})$
Subdomain Jacobi	$O((NP)^{1/4})$	$O((NP)^{1/6})$
1-level Additive Schwarz	$O(P^{1/2})$	$O(P^{1/3})$
2-level Additive Schwarz	$O(1)$	$O(1)$

- 525 • Newton parameters: convergence tolerance on each timestep, globalization strat-
526 egy (line search or trust region parameters), refresh frequency for Jacobian pre-
conditioner.
- 528 • Krylov parameters: convergence tolerance for each Newton correction, restart di-
529 mension of Krylov subspace, overall Krylov iteration limit, orthogonalization
mechanism.
- 531 • Schwarz parameters: subdomain number, quality of subdomain solver (fill level,
number of sweeps), amount of subdomain overlap, coarse grid usage.
- Subproblem parameters: fill level, number of sweeps.

533 *5.2.1. Parameters for pseudo-transient continuation*

534 Although asymptotically superlinear, solution strategies based on Newton’s
535 method must often be approached through pseudo-timestepping. For robustness,
536 pseudo-timestepping is often initiated with very small timesteps and accelerated
537 subsequently. However, this conventional approach can lead to long “induction”
538 periods that may be bypassed by a more aggressive strategy, especially for the
539 smooth flow fields.

540 The timestep is advanced toward infinity by a power-law variation of the switched
541 evolution/relaxation (SER) heuristic of Van Leer and Mulder [23]. To be specific,
542 within each residual reduction phase of computation, we adjust the timestep ac-
543 cording to

$$N_{\text{CFL}}^\ell = N_{\text{CFL}}^0 \left(\frac{\|f(u^0)\|}{\|f(u^{\ell-1})\|} \right)^p,$$

545 where p is a tunable exponent close to unity. Fig. 5 shows the effect of initial CFL
546 number (the Courant–Friedrich–Levy number, a dimensionless measure of the
547 timestep size), N_{CFL}^0 , on the convergence rate. In general, the best choice of initial
548 CFL number is dependent on the grid size and Mach number. A small CFL adds
549 nonlinear stability far from the solution but retards the approach to the domain of
550 superlinear convergence of the steady state. For flows with near discontinuities, it is
551 safer to start with small CFL numbers.

552 In flows with shocks, high-order (second or higher) discretization for the con-
553 vection terms should be activated only after the shock position has settled down. We
554 begin such simulations with a first-order upwind scheme and switch to second-order
555 after a certain residual reduction. The exponent (p) in the power law above is
556 damped to 0.75 for robustness when shocks are expected to appear in second-order
557 discretizations. For first-order discretizations, this exponent may be as large as 1.5. A
558 reasonable switchover point of the residual norm between first-order and second-
559 order discretization phases has been determined empirically. In shock-free simula-
560 tions we use second-order accuracy throughout. Otherwise, we normally reduce the
561 first two to four orders of residual norm with the first-order discretization, then
562 switch to second. This order of accuracy applies to the flux calculation. The pre-
563 conditioner matrix is always built out of a first-order analytical Jacobian matrix.

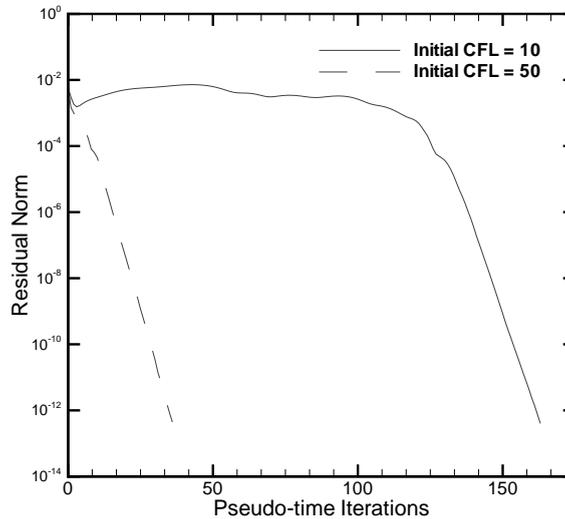


Fig. 5. Residual norm versus iteration count for a 2.8 M-vertex case, showing the effect of initial CFL number on convergence rate. The convergence tuning of nonlinear problems is notoriously case specific.

564 5.2.2. Parameters for Krylov solver

565 We use an inexact Newton method on each timestep [8]; that is, the linear system
 566 within each Newton iteration is solved only approximately. Especially in the be-
 567 ginning of the solution process, this saves a significant amount of execution time. We
 568 have considered the following three parameters in this phase of computation: con-
 569 vergence tolerance, the number of simultaneously storable Krylov vectors, and the
 570 total number of Krylov iterations. The typical range of variation for the inner
 571 convergence tolerance is 0.001–0.01. We have experimented with progressively
 572 tighter tolerances near convergence, and saved Newton iterations thereby, but did
 573 not save time relative to cases with loose and constant tolerance. The Krylov sub-
 574 space dimension depends largely on the problem size and the available memory. We
 575 have used values in the range of 10–30 for most of the problems. The total number of
 576 linear iterations (within each nonlinear solve) has been varied from 10 for the
 577 smallest problem to 80 for the largest one. A typical number of fine-grid flux eval-
 578 uations for achieving 10^{-10} residual reduction on a million-vertex Euler problem is a
 579 couple of thousand.

580 5.2.3. Additive Schwarz preconditioner

581 Table 9 explores two quality parameters for the additive Schwarz preconditioner:
 582 subdomain overlap and quality of the subdomain solve using incomplete factor-
 583 ization. We exhibit execution time and iteration count data from runs of PETSc-
 584 FUN3D on the ASCI Red machine for a fixed-size problem with 357,900 grid points
 585 and 1,789,500 degrees-of-freedom. These calculations were performed using
 586 GMRES(20), one subdomain per processor (without overlap for block Jacobi and

Table 9

Execution times and linear iteration counts on the 333 MHz Pentium Pro ASCII Red machine for a 357,900-vertex case, showing the effect of subdomain overlap and incomplete factorization fill level in the additive Schwarz preconditioner^a

Number of processors	<i>ILU(0) in each subdomain</i>					
	Overlap					
	0		1		2	
	Time (s)	Linear Its	Time (s)	Linear Its	Time (s)	Linear Its
32	688	930	661	816	696	813
64	371	993	374	876	418	887
128	210	1052	230	988	222	872
	<i>ILU(1) in each subdomain</i>					
32	598	674	564	549	617	532
64	334	746	335	617	359	551
128	177	807	178	630	200	555
	<i>ILU(2) in each subdomain</i>					
32	688	527	786	441	–	–
64	386	608	441	488	531	448
128	193	631	272	540	313	472

^aThe best execution times for each ILU fill level and number of processors are in boldface in each row.

587 with overlap for ASM), and $ILU(k)$ where k varies from 0 to 2, and with the natural
 588 ordering in each subdomain block. The use of $ILU(0)$ with natural ordering on the
 589 first-order Jacobian, while applying a second-order operator, allows the factorization
 590 to be done in place, with or without overlap. However, the overlap case does require
 591 forming an additional data structure on each processor to store matrix elements
 592 corresponding to the overlapped regions.

593 From Table 9 we see that the larger overlap and more fill help in reducing the total
 594 number of linear iterations as the number of processors increases, as theory and
 595 intuition predict. However, both increases consume more memory, and both result in
 596 more work per iteration, ultimately driving up execution times in spite of faster
 597 convergence. Best execution times are obtained for any given number of processors
 598 for $ILU(1)$, as the number of processors becomes large (subdomain size small), for
 599 zero overlap.

600 The additional computation/communication costs for additive Schwarz (as com-
 601 pared with block Jacobi) are the following.

- 602 1. Calculation of the matrix couplings among processors. For block Jacobi, these
 603 need not be calculated.
 - 604 2. Communication of the “overlapped” matrix elements to the relevant processors.
 - 605 3. Factorization of the larger local submatrices.
 - 606 4. Communication of the ghost points in the application of the ASM preconditioner.
- 607 We use restricted additive Schwarz method (RASM) [6], which communicates on-
 608 ly when setting up the overlapped subdomain problems and ignores the updates

609 coming from the overlapped regions. This saves a factor of two in local commu-
610 nication relative to standard ASM.

611 5. Inversion of larger triangular factors in each iteration.

612 The execution times reported in Table 9 are highly dependent on the machine
613 used, since each of the additional computation/communication costs listed above
614 may shift the computation past a knee in the performance curve for memory
615 bandwidth, communication network, and so on.

616 *5.2.4. Other algorithmic tuning parameters*

617 In [11] we highlight some additional tunings that have yielded good results in our
618 context. Some subsets of these parameters are not orthogonal but interact strongly
619 with each other. In addition, optimal values of some of these parameters depend on
620 the grid resolution. We are currently using derivative-free asynchronous parallel
621 direct search algorithms [15] to more systematically explore this large parameter
622 space.

623 We emphasize that the discussion in this section does not pertain to discretization
624 parameters, which constitute another area of investigation – one that ultimately
625 impacts performance at a higher level. The algorithmic parameters discussed in this
626 section do not affect the accuracy of the discrete solution, but only the rate at which
627 the solution is attained. In all of our experiments, the goal has been to minimize the
628 overall execution time, not to maximize the floating-point operations per second.
629 There are many tradeoffs that enhance Mflop/s rates but retard execution comple-
630 tion.

631 6. Large-scale demonstration runs

632 We use PETSc's profiling and logging features to measure the parallel perfor-
633 mance. PETSc logs many different types of events and provides valuable information
634 about time spent, communications, load balance, and so forth for each logged event.
635 PETSc uses manual counting of flops, which are afterwards aggregated over all the
636 processors for parallel performance statistics. We have observed that the flops re-
637 ported by PETSc are close to (within 10% of) the values statistically measured by
638 hardware counters on the R10000 processor.

639 PETSc uses the best timers available at the user level in each processing envi-
640 ronment. In our rate computations, we exclude the initialization time devoted to I/O
641 and data partitioning. To suppress timing variations caused by paging in the exe-
642 cutable from disk, we preload the code into memory with one nonlinear iteration,
643 then flush, reload the initial iterate, and begin performance measurements.

644 Since we are solving large fixed-size problems on distributed-memory machines, it
645 is not reasonable to base parallel scalability on a uniprocessor run, which would
646 thrash the paging system. Our base processor number is such that the problem has
647 just fit into the local memory.

648 The same fixed-size problem is run on large ASCI Red configurations with sample
649 scaling results shown in Fig. 6. The implementation efficiency is 91% in going from

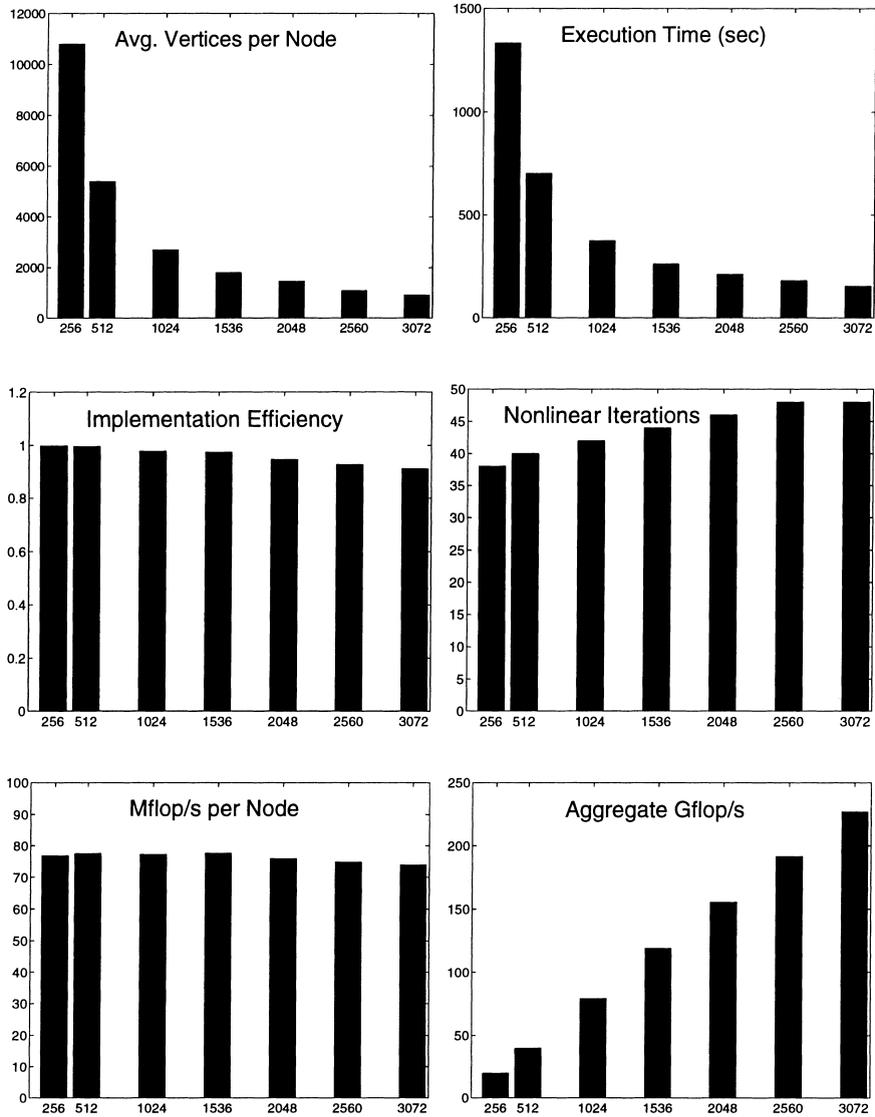


Fig. 6. Parallel performance for a fixed size mesh of 2.8 M vertices run on up to 3072 ASCII Red 333 MHz Pentium Pro processors.

650 256 to 3072 nodes. For the data in Fig. 6, we employed the `-procs 2` runtime
 651 option on ASCII Red. This option enables 2-processor-per-node multithreading
 652 during threadsafe, communication-free portions of the code. We have activated this
 653 feature for the floating-point-intensive flux computation subroutine alone. On 3072
 654 nodes, the largest run we have been able to make on the unclassified side of the

655 machine to date, the resulting Gflop/s rate is 227 (when the preconditioner is stored
 656 in double precision). Undoubtedly, further improvements to the algebraic solver
 657 portion of the code are also possible through multithreading, but the additional
 658 coding work does not seem justified at present.

659 Fig. 7 shows aggregate flop/s performance and a log–log plot showing execution
 660 time for our largest case on the three most capable machines to which we have thus
 661 far had access. In both plots of this figure, the dashed lines indicate ideal behavior.
 662 Note that although the ASCI Red flop/s rate scales nearly linearly, a higher fraction
 663 of the work is redundant at higher parallel granularities, so the execution time does
 664 not drop in exact proportion to the increase in flop/s. The number of vertices per
 665 processor ranges from about 22,000 to fewer than 1000 over the range shown. We
 666 point out that for just 1000 vertices in a three-dimensional domain, about half are on
 667 the interface (e.g., 488 interface vertices on a $10 \times 10 \times 10$ cube).

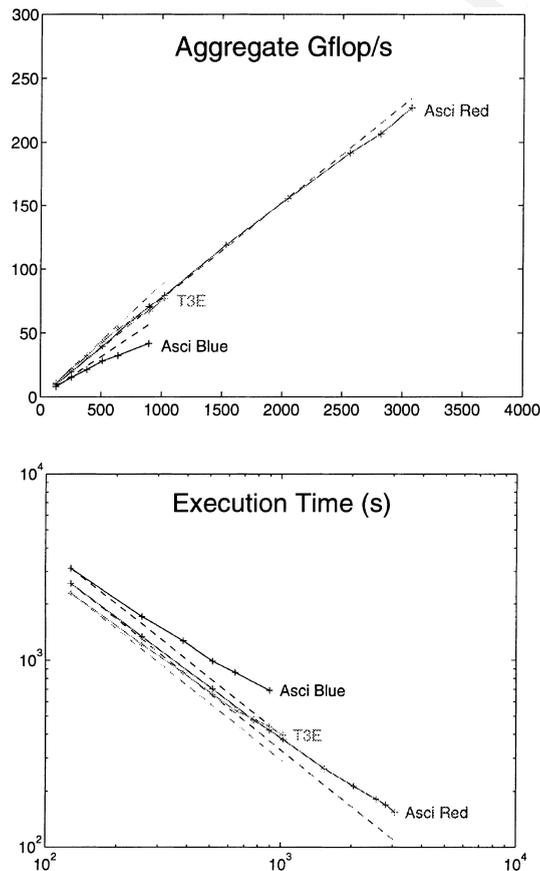


Fig. 7. Gigaflop/s ratings and execution times on ASCI Red (up to 3072 dual processor nodes), ASCI Pacific Blue (up to 768 processors), and a Cray T3E (up to 1024 processors) for a 2.8 M-vertex case, along with dashed lines indicating “perfect” scalings.

668 7. Conclusions

669 Large-scale implicit computations have matured to a point of practical use on
670 distributed/shared memory architectures for static-grid problems. More sophisti-
671 cated algorithms, including solution adaptivity, inherit the same features *within*
672 static-grid phases, of course, but require extensive additional infrastructure for dy-
673 namic parallel adaptivity, rebalancing, and maintenance of efficient, consistent dis-
674 tributed data structures.

675 Unstructured implicit CFD solvers are amenable to scalable implementation, but
676 careful tuning is needed to obtain the best product of per-processor efficiency and
677 parallel efficiency. The number of cache misses and the achievable memory band-
678 width are two important parameters that should be considered in determining an
679 optimal data storage pattern. The impact of data reorganizing strategies (interlacing,
680 blocking, and edge/vertex reorderings) is demonstrated through the sparse matrix–
681 vector product model and hardware counter profiling.

682 Given contemporary high-end architecture, critical research directions for solution
683 algorithms for systems modeled by PDEs are: (1) multivector algorithms and less
684 synchronous algorithms, and (2) hybrid programming models. To influence future
685 architectures while adapting to current ones, we recommend adoption of new
686 benchmarks featuring implicit methods on unstructured grids, such as the applica-
687 tion featured herein.

688 Acknowledgements

689 We are indebted to Lois C. McInnes and Satish Balay of Argonne National
690 Laboratory, to W. Kyle Anderson, formerly of the NASA Langley Research Center,
691 and to Dimitri Mavriplis of ICASE for collaborations leading up to the work pre-
692 sented here. Debbie Swider of Argonne National Laboratory was of considerable
693 assistance in performing ASCI platform runs. Computer time was supplied by Ar-
694 gonne National Laboratory, Lawrence Livermore National Laboratory, NERSC,
695 Sandia National Laboratories, and SGI-Cray.

696 References

- 697 [1] W.K. Anderson, D.L. Bonhaus, An implicit upwind algorithm for computing turbulent flows on
698 unstructured grids, *Comput. Fluids* 23 (1994) 1–21.
- 699 [2] W.K. Anderson, R.D. Rausch, D.L. Bonhaus, Implicit/multigrid algorithms for incompressible
700 turbulent flows on unstructured grids, *J. Comput. Phys.* 128 (1996) 391–408.
- 701 [3] S. Balay, W.D. Gropp, L.C. McInnes, B.F. Smith, The Portable Extensible Toolkit for Scientific
702 Computing (PETSc) version 2.8, <http://www.mcs.anl.gov/petsc/petsc.html>, 2000.
- 703 [4] S.W. Bova, C.P. Breshears, C.E. Cuicchi, Z. Demirbilek, H.A. Gabb, Dual-level parallel analysis of
704 harbor wave response using MPI and OpenMP, *Int. J. High Performance Comput. Appl.* 14 (2000)
705 49–64.

- 706 [5] X.-C. Cai, Some domain decomposition algorithms for nonselfadjoint elliptic and parabolic partial
707 differential equations, Technical Report 461, Courant Institute, New York, 1989.
- 708 [6] X.-C. Cai, M. Sarkis, A restricted additive Schwarz preconditioner for general sparse linear systems,
709 SIAM J. Scientific Comput. 21 (1999) 792–797.
- 710 [7] E. Cuthill, J. McKee, Reducing the bandwidth of sparse symmetric matrices, in: Proceedings of the
711 24th National Conference of the ACM, 1969.
- 712 [8] R.S. Dembo, S.C. Eisenstat, T. Steihaug, Inexact Newton methods, SIAM J. Numer. Anal. 19 (1982)
713 400–408.
- 714 [9] J. Dongarra, H.-W. Meuer, E. Strohmaier, The TOP 500 List, [http://www.netlib.org/benchmark/
715 top500.html](http://www.netlib.org/benchmark/top500.html), 2000.
- 716 [10] W.D. Gropp, D.K. Kaushik, D.E. Keyes, B.F. Smith, Toward realistic performance bounds for
717 implicit CFD codes, in: D. Keyes, A. Ecer, J. Periaux, N. Satofuka, P. Fox (Eds.), Proceedings of the
718 Parallel CFD'99, Elsevier, Berlin, 1999, pp. 233–240.
- 719 [11] W.D. Gropp, D.K. Kaushik, D.E. Keyes, B.F. Smith, Performance modeling and tuning of an
720 unstructured mesh CFD application, in: Proceedings of the SC2000, IEEE Computer Society, 2000.
- 721 [12] W.D. Gropp, L.C. McInnes, M.D. Tidriri, D.E. Keyes, Globalized Newton–Krylov–Schwarz
722 algorithms and software for parallel implicit CFD, Int. J. High Performance Comput. Appl. 14 (2000)
723 102–136.
- 724 [13] W. Gropp, E. Lusk, A. Skjellum, Using MPI: Portable Parallel Programming with the Message
725 Passing Interface, second ed., MIT Press, Cambridge, MA, 1999.
- 726 [14] J.L. Hennessy, D.A. Patterson, Computer Architecture: A Quantitative Approach, Morgan
727 Kaufmann, Los Altos, CA, 1996.
- 728 [15] P.D. Hough, T.G. Kolda, V.J. Torczon, Asynchronous parallel pattern search for nonlinear
729 optimization, Technical Report SAND2000-8213, Sandia National Laboratories, Livermore, January
730 2000.
- 731 [16] G. Karypis, V. Kumar, A fast and high quality scheme for partitioning irregular graphs, SIAM J.
732 Scientific Comput. 20 (1999) 359–392.
- 733 [17] C.T. Kelley, D.E. Keyes, Convergence analysis of pseudo-transient continuation, SIAM J. Numer.
734 Anal. 35 (1998) 508–523.
- 735 [18] D.E. Keyes, How scalable is domain decomposition in practice? in: C.-H. Lai et al. (Eds.),
736 Proceedings of the 11th International Conference on Domain Decomposition Methods, Domain
737 Decomposition Press, Bergen, 1999.
- 738 [19] D.E. Keyes, Four horizons for enhancing the performance of parallel simulations based on partial
739 differential equations, in: Proceedings of the EuroPar 2000, Lecture Notes in Computer Science,
740 Springer, Berlin, 2000.
- 741 [20] D.J. Mavriplis, Parallel unstructured mesh analysis of high-lift configurations, Technical Report
742 2000-0923, AIAA, 2000.
- 743 [21] J.D. McCalpin, STREAM: Sustainable memory bandwidth in high performance computers,
744 Technical report, University of Virginia, 1995, <http://www.cs.virginia.edu/stream>.
- 745 [22] MIPS Technologies, Inc., [http://techpubs.sgi.com/library/manuals/
746 2000/007-2490-001/pdf/007-2490-001.pdf](http://techpubs.sgi.com/library/manuals/2000/007-2490-001/pdf/007-2490-001.pdf). MIPS R10000 Microprocessor User's Manual, January 1997.
- 747 [23] W. Mulder, B. Van Leer, Experiments with implicit upwind methods for the Euler equations, J.
748 Comput. Phys. 59 (1985) 232–246.
- 749 [24] Silicon Graphics, Inc, [http://techpubs.sgi.com/library/manuals/
750 3000/007-3430-002/pdf/007-3430-002.pdf](http://techpubs.sgi.com/library/manuals/3000/007-3430-002/pdf/007-3430-002.pdf). Origin 2000 and Onyx2 Performance and Tuning Optimization Guide, 1998, Document
751 Number 007-3430-002.
- 752 [25] B.F. Smith, P. Bjørstad, W. Gropp, Domain Decomposition: Parallel Multilevel Algorithms for
753 Elliptic Partial Differential Equations, Cambridge University Press, Cambridge, 1996.
- 754 [26] O. Temam, W. Jalby, Characterizing the behavior of sparse algorithms on caches, in: Proceedings of
755 the Supercomputing 92, IEEE Computer Society, 1992, pp. 578–587.
- 756 [27] S. Toledo, Improving the memory-system performance of sparse-matrix vector multiplication, IBM J.
757 Res. Dev. 41 (1997) 711–725.

- 758 [28] G. Wang, D.K. Tafti, Performance enhancements on microprocessors with hierarchical memory
759 systems for solving large sparse linear systems, *Int. J. High Performance Comput. Appl.* 13 (1999) 63–
760 79.
- 761 [29] J. White, P. Sadayappan, On improving the performance of sparse matrix–vector multiplication, in:
762 *Proceedings of the Fourth International Conference on High Performance Computing (HiPC'97)*,
763 IEEE Computer Society, 1997, pp. 578–587.

UNCORRECTED PROOF