

Parallel Performance Studies for a 3-D Elliptic Test Problem on the 2018 Portion of the Taki Cluster

Richard Ebadi¹, Carlos Barajas², and Matthias K. Gobbert² (gobbert@umbc.edu)

¹Institute for Mathematics, University of Kassel, Germany

²Department of Mathematics and Statistics, University of Maryland, Baltimore County, USA

Technical Report HPCF-2018-19, hpcf.umbc.edu > Publications

Abstract

The new 2018 nodes in the CPU cluster taki in the UMBC High Performance Computing Facility contain two 18-core Intel Skylake CPUs and 384 GB of memory per node, connected by an EDR (Enhanced Data Rate) InfiniBand interconnect. The performance studies use the test problem of the Poisson equation in three spatial dimensions, discretized by the finite difference method to give a very large and sparse system of linear equations that is solved by the conjugate gradient method. The algorithm is known to be memory-bound so it is a good test for the architecture of the nodes and the parallel network connecting them. Strong scalability studies varying the number of processes per node as well as the number of compute nodes demonstrate excellent scalability when using multiple nodes as well as very good scalability using multiple cores per node.

1 Introduction

The UMBC High Performance Computing Facility (HPCF) is the community-based, interdisciplinary core facility for scientific computing and research on parallel algorithms at UMBC. Started in 2008 by more than 20 researchers from ten academic departments and research centers from all three colleges, it is supported by faculty contributions, federal grants, and the UMBC administration. The facility is open to UMBC researchers at no charge. Researchers can contribute funding for long-term priority access. System administration is provided by the UMBC Division of Information Technology, and users have access to consulting support provided by dedicated full-time graduate assistants. See hpcf.umbc.edu for more information on HPCF and the projects using its resources.

The tests reported here use the new 2018 portion of the CPU cluster in taki. This portion of the CPU cluster consists of 42 compute nodes with two 18-core Intel Xeon Gold 6140 Skylake CPUs (2.3 GHz clock speed, 24.75 MB L3 cache, 6 memory channels, 140 W power), for a total of 36 cores per node, 384 GB memory (12 × 32 GB DDR4), and a 120 GB SSD drive. These nodes are connected by a network of four 36-port EDR (Enhanced Data Rate) InfiniBand switches (100 Gb/s bandwidth, 90 ns latency) to a central storage of more than 750 TB. See hpcf.umbc.edu for photos, schematics, and more detailed information, also on the other portions of the taki cluster.

This report uses the Poisson equation with homogeneous Dirichlet boundary conditions on a unit cube domain in three spatial dimensions. to test the performance of the 2018 portion of the CPU cluster in taki. Two earlier reports also considered the problem, namely [8] with the same linear solver, but implemented in Matlab, and [9] using a parallel C code using MPI, but with a different linear solver. Many other previous reports considered the two-dimensional analogue of this test problem (see [2] and references therein), most recently using compute nodes with Intel Skylake CPUs [2] on taki 2018 and [1] on Stampede2. Discretizing the spatial derivatives in the Poisson equation by the finite difference method yields a system of linear equations with a large, sparse, highly structured, symmetric positive definite system matrix. The linear system resulting from the two-dimensional version of the test problem is a classical test problem for iterative solvers and contained in several textbooks including [4, 5, 6, 10]. The parallel, matrix-free implementation of the conjugate gradient method as appropriate iterative linear solver for this linear system involves necessarily communications both collectively among all parallel processes and between pairs of processes in every iteration. Therefore, this method provides an excellent test problem for the overall, real-life performance of a parallel computer on a memory-bound algorithm. The results are not just applicable to the conjugate gradient method, which is important in its own right as a representative of the class of Krylov subspace methods, but to all memory-bound algorithms. The implementation uses the C programming language, with MPI (Message Passing Interface) for communications between distributed-memory cluster nodes.

The results demonstrate excellent scalability when using multiple nodes due to the low latency of the high-performance interconnect and good speedup when using all cores of the multi-core CPUs. The studies varying the number of processes per node demonstrate very good scalability using multiple cores per node.

The remainder of this report is organized as follows: Section 2 details the test problem and discusses the parallel implementation in more detail, and Section 3 summarizes the solution and method convergence data. Section 4 contains the strong scalability studies using MPI-only code on taki 2018.

2 The Elliptic Test Problem

We consider the classical elliptic test problem of the Poisson equation with homogeneous Dirichlet boundary conditions (see, e.g., [10, Chapter 8] for the two-dimensional analogue)

$$\begin{aligned} -\Delta u &= f && \text{in } \Omega, \\ u &= 0 && \text{on } \partial\Omega, \end{aligned} \quad (2.1)$$

on the unit cube domain $\Omega = (0, 1) \times (0, 1) \times (0, 1) \subset \mathbb{R}^3$ in three spatial dimensions. Here, $\partial\Omega$ denotes the boundary of the domain Ω and the Laplace operator in is defined as $\Delta u = \frac{\partial^2 u}{\partial x_1^2} + \frac{\partial^2 u}{\partial x_2^2} + \frac{\partial^2 u}{\partial x_3^2}$. Using $N + 2$ mesh points in each dimension, we construct a mesh with uniform mesh spacing $h = 1/(N + 1)$. Specifically, define the mesh points $(x_{k_1}, x_{k_2}, x_{k_3}) \in \bar{\Omega} \subset \mathbb{R}^3$ with $x_{k_i} = h k_i$, $k_i = 0, 1, \dots, N, N + 1$, in each dimension $i = 1, 2, 3$. Denote the approximations to the solution at the mesh points by $u_{k_1, k_2, k_3} \approx u(x_{k_1}, x_{k_2}, x_{k_3})$. Then approximate the second-order derivatives in the Laplace operator at the N^2 interior mesh points by

$$\frac{\partial^2 u(x_{k_1}, x_{k_2}, x_{k_3})}{\partial x_1^2} \approx \frac{u_{k_1-1, k_2, k_3} - 2u_{k_1, k_2, k_3} + u_{k_1+1, k_2, k_3}}{h^2}, \quad (2.2)$$

$$\frac{\partial^2 u(x_{k_1}, x_{k_2}, x_{k_3})}{\partial x_2^2} \approx \frac{u_{k_1, k_2-1, k_3} - 2u_{k_1, k_2, k_3} + u_{k_1, k_2+1, k_3}}{h^2}, \quad (2.3)$$

$$\frac{\partial^2 u(x_{k_1}, x_{k_2}, x_{k_3})}{\partial x_3^2} \approx \frac{u_{k_1, k_2, k_3-1} - 2u_{k_1, k_2, k_3} + u_{k_1, k_2, k_3+1}}{h^2} \quad (2.4)$$

for $k_i = 1, \dots, N$, $i = 1, 2, 3$, for the approximations at the interior points. Using this approximation together with the homogeneous boundary conditions (2.1) gives a system of N^3 linear equations for the finite difference approximations at the N^3 interior mesh points.

Collecting the N^3 unknown approximations u_{k_1, k_2, k_3} in a vector $u \in \mathbb{R}^{N^3}$ using the natural ordering of the mesh points, we can state the problem as a system of linear equations in standard form $Au = b$ with a system matrix $A \in \mathbb{R}^{N^3 \times N^3}$ and a right-hand side vector $b \in \mathbb{R}^{N^3}$. The components of the right-hand side vector b are given by the product of h^2 multiplied by right-hand side function evaluations $f(x_{k_1}, x_{k_2}, x_{k_3})$ at the interior mesh points using the same ordering as the one used for u_{k_1, k_2, k_3} . The system matrix $A \in \mathbb{R}^{N^3 \times N^3}$ can be defined recursively as block tri-diagonal matrix with $N \times N$ blocks of size $N^2 \times N^2$ each, each of which in turn is a block tri-diagonal matrix with $N \times N$ blocks of size $N \times N$ each. Concretely, we have

$$A = \begin{bmatrix} S & -I_{N^2} & & & & & \\ -I_{N^2} & S & -I_{N^2} & & & & \\ & & \ddots & \ddots & \ddots & & \\ & & & -I_{N^2} & S & -I_{N^2} & \\ & & & & -I_{N^2} & S & \\ & & & & & & S \end{bmatrix} \in \mathbb{R}^{N^3 \times N^3}, \quad S = \begin{bmatrix} T & -I_N & & & & & \\ -I_N & T & -I_N & & & & \\ & & \ddots & \ddots & \ddots & & \\ & & & -I_N & T & -I_N & \\ & & & & -I_N & T & \\ & & & & & -I_N & T \end{bmatrix} \in \mathbb{R}^{N^2 \times N^2}$$

with the tri-diagonal matrix $T = \text{tridiag}(-1, 6, -1) \in \mathbb{R}^{N \times N}$ for the diagonal blocks of S as well as A and identity matrices $I_{N^2} \in \mathbb{R}^{N^2 \times N^2}$ and $I_N \in \mathbb{R}^{N \times N}$ for the off-diagonal blocks of A and S , respectively.

For fine meshes with large N , iterative methods such as the conjugate gradient method are appropriate for solving this linear system. The system matrix A is known to be symmetric positive definite and thus the method is guaranteed to converge for this problem. In a careful implementation, the conjugate gradient method requires in each iteration exactly two inner products between vectors, three vector updates, and one matrix-vector product involving the system matrix A . In fact, this matrix-vector product is the only way, in which A enters into the algorithm. Therefore, a so-called matrix-free implementation of the conjugate gradient method is possible that avoids setting up any matrix, if one provides a function that computes as its output the product vector $q = Ap$ component-wise directly from the components of the input vector p by using the explicit knowledge of the values and positions of the non-zero components of A , but without assembling A as a matrix.

Thus, without storing A , a careful, efficient, matrix-free implementation of the (unpreconditioned) conjugate gradient method only requires the storage of four vectors (commonly denoted as the solution vector x , the residual r , the search direction p , and an auxiliary vector q). In a parallel implementation of the conjugate gradient method, each vector is split into as many blocks as parallel processes are available and one block distributed to each process. That is, each parallel process possesses its own block of each vector, and normally no vector is ever assembled in full on any process. To understand what this means for parallel programming and the performance of the method, note that

an inner product between two vectors distributed in this way is computed by first forming the local inner products between the local blocks of the vectors and second summing all local inner products across all parallel processes to obtain the global inner product. This summation of values from all processes is known as a reduce operation in parallel programming, which requires a communication among all parallel processes. This communication is necessary as part of the numerical method used, and this necessity is responsible for the fact that for fixed problem sizes eventually for very large numbers of processes the time needed for communication — increasing with the number of processes — will unavoidably dominate over the time used for the calculations that are done simultaneously in parallel — decreasing due to shorter local vectors for increasing number of processes. By contrast, the vector updates in each iteration can be executed simultaneously on all processes on their local blocks, because they do not require any parallel communications. However, this requires that the scalar factors that appear in the vector updates are available on all parallel processes. This is accomplished already as part of the computation of these factors by using a so-called Allreduce operation, that is, a reduce operation that also communicates the result to all processes. This is implemented in the MPI function `MPI_Allreduce` [7]. Finally, the matrix-vector product $q = Ap$ also computes only the block of the vector q that is local to each process. But since the matrix A has non-zero off-diagonal elements, each local block needs values of p that are local to the two processes that hold the neighboring blocks of p . The communications between parallel processes thus needed are so-called point-to-point communications, because not all processes participate in each of them, but rather only specific pairs of processes that exchange data needed for their local calculations. Observe now that it is only a few components of q that require data from p that is not local to the process. Therefore, it is possible and potentially very efficient to proceed to calculate those components that can be computed from local data only, while the communications with the neighboring processes are taking place. This technique is known as interleaving calculations and communications and can be implemented using the non-blocking MPI communications commands `MPI_Isend` and `MPI_Irecv` [7].

3 Convergence Study for the Model Problem

To test the numerical method and its implementation, we consider the elliptic problem (2.1) on the unit cube

$$\Omega = (0, 1) \times (0, 1) \times (0, 1)$$

with right-hand side function

$$f(x_1, x_2, x_3) = (-2\pi^2) \left(\cos(2\pi x_1) \sin^2(\pi x_2) \sin^2(\pi x_3) + \sin^2(\pi x_1) \cos(2\pi x_2) \sin^2(\pi x_3) + \sin^2(\pi x_1) \sin^2(\pi x_2) \cos(2\pi x_3) \right),$$

for which the true analytic solution in closed form

$$u(x_1, x_2, x_3) = \sin^2(\pi x_1) \sin^2(\pi x_2) \sin^2(\pi x_3)$$

is known.

To check the convergence of the finite difference method as well as to analyze the performance of the conjugate gradient method, we solve the problem on a sequence of progressively finer meshes. The conjugate gradient method is started with a zero vector as initial guess and the solution is accepted as converged when the Euclidean vector norm of the residual is reduced to the fraction 10^{-6} of the initial residual. Table 3.1 lists the mesh resolution N of the $N \times N \times N$ mesh, the number of degrees of freedom N^3 (DOF; i.e., the dimension of the linear system), the norm of the finite difference error $\|u - u_h\| \equiv \|u - u_h\|_{L^\infty(\Omega)}$, the ratio of consecutive errors $\|u - u_{2h}\| / \|u - u_h\|$, the number of conjugate gradient iterations `#iter`, the observed wall clock time in HH:MM:SS and in seconds, and the predicted and observed memory usage in GB for studies performed in serial. More precisely, the serial runs use the parallel code run on one process only, on a dedicated node (no other processes running on the node), and with all parallel communication commands disabled by if-statements. The wall clock time is measured using the `MPI_Wtime` command (after synchronizing all processes by an `MPI_Barrier` command). The memory usage of the code is predicted by noting that there are $4N^3$ double-precision numbers needed to store the four vectors of significant length N^3 and that each double-precision number requires 8 bytes; dividing this result by 1024^3 converts its value to units of GB, as quoted in the table. The memory usage is observed in the code by checking the `VmRSS` field in the the special file `/proc/self/status`.

In all cases, the norms of the finite difference errors in Table 3.1 decrease by a factor of about 4 each time that the mesh is refined by a factor 2. This confirms that the finite difference method is second-order convergent, as predicted by the numerical theory for the finite difference method [3, 6]. The fact that this convergence order is attained also

Table 3.1: Convergence study on taki with serial code except where noted.

N	DOF	$\ u - u_h\ $	Ratio	#iter	wall clock time		memory usage (GB)	
					HH:MM:SS	seconds	predicted	observed
32	32,769	3.0060e-03	—	61	00:00:00	0.01	< 1	< 1
64	262,144	7.7765e-04	3.86	120	00:00:00	0.13	< 1	< 1
128	2,097,152	1.9763e-04	3.93	243	00:00:04	3.79	< 1	< 1
256	16,777,216	4.9807e-05	3.97	493	00:01:09	69.39	< 1	< 1
512	134,217,728	1.2503e-05	3.98	999	00:19:21	1,161.23	4	4.02
1024	1,073,741,824	3.1327e-06	3.99	2,023	05:34:36	20,075.61	32	32.02
2048	8,589,934,592	7.8346e-07	4.00	4,095	116:50:39	420,638.69	256	256.02

confirms that the tolerance of the iterative linear solver is tight enough to ensure a sufficiently accurate solution of the linear system. The increasing numbers of iterations needed to achieve the convergence of the linear solver highlights the fundamental computational challenge with methods in the family of Krylov subspace methods, of which the conjugate gradient method is the most important example: Refinements of the mesh imply more mesh points, where the solution approximation needs to be found, and makes the computation of each iteration of the linear solver more expensive. Additionally, more of these more expensive iterations are required to achieve convergence to the desired tolerance for finer meshes. And it is not possible to relax the solver tolerance, because otherwise its solution would not be accurate enough and the norm of the finite difference error would not show a second-order convergence behavior, as required by its theory. The observed memory usage in units of GB rounds to within less than 1 GB of the predicted usage. This good agreement between predicted and observed memory usage in the last two columns of the table indicates that the implementation of the code does not have any unexpected memory usage in the serial case. The wall clock times and the memory usages for these serial runs indicate for which mesh resolutions this elliptic test problem becomes challenging computationally. Notice that the very fine meshes show very significant runtimes and memory usage; parallel computing clearly offers opportunities to decrease runtimes as well as to decrease memory usage per process by spreading the problem over the parallel processes.

We note that the results in Table 3.1 agree with past results for this problem, namely with Table 7.3 (first subtable) in [8], which used the software Matlab, but the same method as linear solver. This ensures that the parallel performance studies in the next section are practically relevant, since a correct solution of the test problem is computed.

4 Performance Studies on taki 2018 Using MPI-Only Code

This section describes the parallel performance studies for the solution of the test problem on the 2018 portion of the CPU cluster in taki. This portion of the CPU cluster consists of 42 compute nodes with two 18-core Intel Xeon Gold 6140 Skylake CPUs (2.3 GHz clock speed, 24.75 MB L3 cache, 6 memory channels, 140 W power), for a total of 36 cores per node, 384 GB memory (12×32 GB DDR4), and a 120 GB SSD drive. Figure 4.1 shows a schematic of one of the compute nodes, showing also the two Intel UPI connections between the CPUs and indicating that each CPU has 6 memory channels to a DDR4 memory of 32 GB. These compute nodes are connected by a network of four 36-port EDR (Enhanced Data Rate) InfiniBand switches (100 Gb/s bandwidth, 90 ns latency).

The results in this section use the default Intel compiler and Intel MPI, currently version 18.0.3, which are accessed on taki through the wrapper `mpiicc`. We use the compiler options `-O3 -std=c99 -Wall`.

The SLURM submission script uses the `srun` command to start the job. The number of nodes are controlled by the `--nodes` option and the number of MPI processes per node by the `--ntasks-per-node` option. For a performance study, each node that is used is dedicated to the job with the remaining cores idling, if not all of them are used, using the `--exclusive` flag. Correspondingly, we request all memory of the node for the job by `--mem=MaxMemPerNode`.

We include the `OMP_PLACES` and `OMP_PROC_BIND` environment variables¹ in the SLURM script. The environment variable `OMP_PLACES=cores` is used to list the cores of the CPUs on the node as the places that OpenMP threads are pinned on, while `OMP_PROC_BIND` chooses the order of places in the pinning. The value `OMP_PROC_BIND=close` means that the assignment goes successively through the available places, while `OMP_PROC_BIND=spread` spreads the threads over the places. Studies in [1] with `OMP_PLACES=cores` using the choices of `OMP_PROC_BIND=close` and `OMP_PROC_BIND=spread` resulted in run time that were almost the same for corresponding values of nodes and processes

¹<http://pages.tacc.utexas.edu/~eijkhout/pcse/html/omp-affinity.html>

per node used. This should have been expected, since the environment variables tested are supposed to influence the placement of OpenMP threads, yet the code used here is an MPI-only parallel code.

We conduct complete performance studies of the test problem for four progressively finer meshes of $N = 256, 512, 1024, 2048$. These studies result in progressively larger systems of linear equations with system dimensions ranging from over 16 millions for $N = 256$ to over 8 billions for $N = 2048$.

Table 4.1 collects the results of the performance studies on the 2018 portion of the CPU cluster in taki. For each mesh resolution of the four meshes with $N = 256, 512, 1024, 2048$, the parallel implementation of the test problem is run on all possible combinations of nodes from 1 to 32 by powers of 2 and processes per node from 1 to 32 by powers of 2. The table summarizes the observed wall clock time (total time to execute the code) in HH:MM:SS (hours:minutes:seconds) format. The upper-left entry of each subtable contains the runtime for the 1-process run, i.e., the serial run, of the code for that particular mesh. The lower-right entry of each subtable lists the runtime using 32 cores of both 18-core processors, i.e., not using all available cores, on 32 nodes for a total of 1024 parallel processes working together to solve the problem. Some of the values for $N = 256$ and $N = 512$ are not available if the number of total processes used is too high compared to N . This is because the domain is divided along the x_3 direction and distributed among all processes. Meaning that N has to be higher or equal than the total number of processes used.

We choose the mesh resolution of $1024 \times 1024 \times 1024$ in Table 4.1 to discuss in detail as example. Reading along the first column of this mesh subtable, we observe that by doubling the number of processes from 1 to 2 we approximately halve the runtime from each column to the next. We observe the same improvement from 2 to 4 processes as well as from 4 to 8 processes. We also observe that by doubling the number of processes from 8 to 16 processes, there is still a significant improvement in runtime, although not the halving we observed previously. Finally, while the decrease in run time from 16 to 32 processes is small, the run times still do decrease, making the use of all available cores advisable. We observe that the behavior is analogous also in all other columns for this subtable. This behavior is a typical characteristic of memory-bound code such as this. The limiting factor in performance of memory-bound code is memory access, so we would expect a bottleneck when more processes on each CPU attempt to access the memory simultaneously than the available 6 memory channels per CPU; see Figure 4.1.

Reading along each row of the $1024 \times 1024 \times 1024$ mesh subtable, we observe that by doubling the number of nodes used, and thus also doubling the number of parallel processes, we approximately halve the runtime all the way up to 32 nodes. This behavior observed for increasing the number of nodes confirms the quality of the high-performance InfiniBand interconnect. Also, we can see that the timings for anti-diagonals in Table 4.1 are about equal, that is, the run time for nodes=2 and and processes per node=1 is almost same as for nodes=1 and processes per node=2, for instance. Thus, it is advisable to use the smallest number of nodes with the largest number of processes per node.

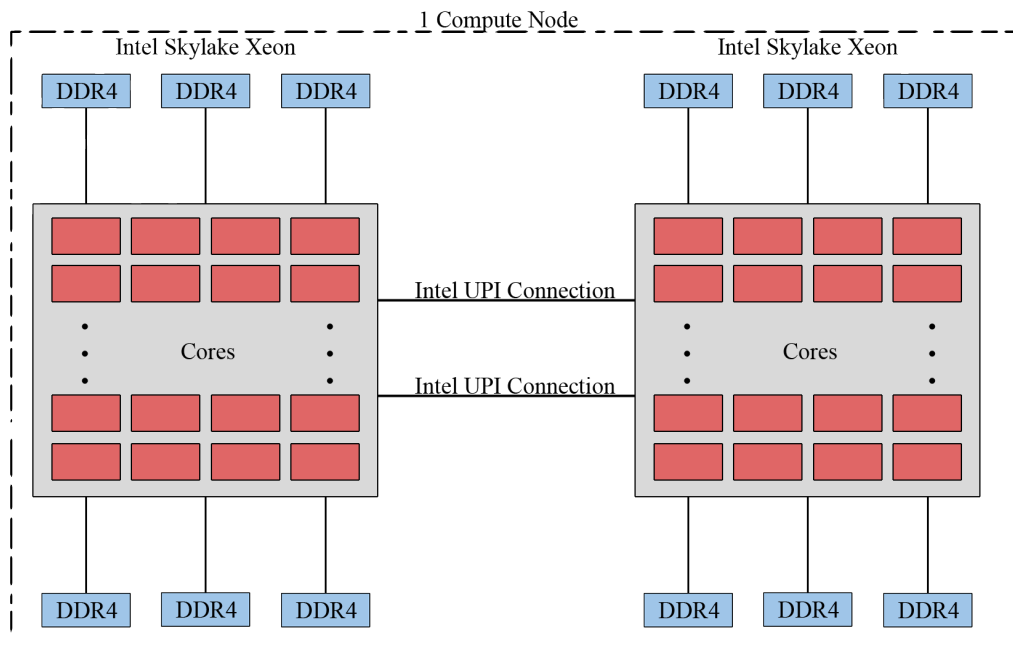


Figure 4.1: Schematic of a compute node with two Intel Skylake Xeon CPUs.

Table 4.1: Wall clock time in HH:MM:SS on taki 2018 using the default Intel compiler and Intel MPI, version 18.0.3.

(a) Mesh resolution $N \times N \times N = 256 \times 256 \times 256$, system dimension 16777216						
	1 node	2 nodes	4 nodes	8 nodes	16 nodes	32 nodes
1 process per node	00:01:09	00:00:34	00:00:17	00:00:08	00:00:03	00:00:01
2 processes per node	00:00:35	00:00:17	00:00:10	00:00:05	00:00:02	00:00:01
4 processes per node	00:00:18	00:00:10	00:00:05	00:00:02	00:00:01	00:00:01
8 processes per node	00:00:10	00:00:06	00:00:03	00:00:01	00:00:01	00:00:00
16 processes per node	00:00:08	00:00:05	00:00:02	00:00:01	00:00:01	N/A
32 processes per node	00:00:07	00:00:03	00:00:02	00:00:01	N/A	N/A
(b) Mesh resolution $N \times N \times N = 512 \times 512 \times 512$, system dimension 134217728						
	1 node	2 nodes	4 nodes	8 nodes	16 nodes	32 nodes
1 process per node	00:19:21	00:09:47	00:04:52	00:02:28	00:01:13	00:00:36
2 processes per node	00:09:55	00:04:56	00:02:45	00:01:34	00:00:51	00:00:25
4 processes per node	00:05:00	00:02:32	00:01:23	00:00:49	00:00:26	00:00:14
8 processes per node	00:02:38	00:01:21	00:00:46	00:00:25	00:00:15	00:00:08
16 processes per node	00:01:52	00:01:11	00:00:37	00:00:20	00:00:11	00:00:07
32 processes per node	00:01:34	00:00:51	00:00:27	00:00:16	00:00:09	N/A
(c) Mesh resolution $N \times N \times N = 1024 \times 1024 \times 1024$, system dimension 1073741824						
	1 node	2 nodes	4 nodes	8 nodes	16 nodes	32 nodes
1 process per node	05:34:36	02:48:47	01:24:23	00:42:15	00:21:24	00:10:42
2 processes per node	02:50:40	01:25:54	00:47:36	00:22:10	00:12:42	00:07:14
4 processes per node	01:27:40	00:44:27	00:22:14	00:11:49	00:07:05	00:03:52
8 processes per node	00:43:58	00:22:11	00:11:14	00:05:57	00:03:38	00:02:07
16 processes per node	00:29:57	00:15:09	00:07:42	00:04:01	00:02:18	00:01:28
32 processes per node	00:25:04	00:12:53	00:06:44	00:03:56	00:02:05	00:01:13
(d) Mesh resolution $N \times N \times N = 2048 \times 2048 \times 2048$, system dimension 8589934592						
	1 node	2 nodes	4 nodes	8 nodes	16 nodes	32 nodes
1 process per node	116:50:39	46:29:18	23:26:36	11:47:12	05:53:08	02:56:56
2 processes per node	47:04:26	23:39:35	11:54:49	06:00:32	03:06:33	01:35:45
4 processes per node	23:32:45	11:57:10	06:02:15	03:06:01	01:40:09	00:53:11
8 processes per node	11:53:27	06:05:10	03:06:58	01:40:53	00:51:56	00:26:35
16 processes per node	08:04:29	04:03:26	02:06:49	01:04:19	00:33:08	00:18:55
32 processes per node	06:44:52	03:24:50	01:44:28	00:54:07	00:29:04	00:16:29

Parallel scalability is often visually represented by plots of observed speedup and efficiency. The ideal behavior of code for a fixed problem size N using p parallel processes is that it be p times as fast as serial code. If $T_p(N)$ denotes the wall clock time for a problem of a fixed size parameterized by N using p processes, then the quantity $S_p = T_1(N)/T_p(N)$ measures the speedup of the code from 1 to p processes, whose optimal value is $S_p = p$. The efficiency $E_p = S_p/p$ characterizes in relative terms how close a run with p parallel processes is to this optimal value, for which $E_p = 1$. The behavior described here for speedup for a fixed problem size is known as strong scalability of parallel code.

Table 4.2 organizes the results of Table 4.1 in the form of a strong scalability study, that is, there is one row for each problem size, with columns for increasing number of parallel processes p . Table 4.2 (a) lists the raw timing data, like Table 4.1, but organized by numbers of parallel processes p . Tables 4.2 (b) and (c) show the numbers for speedup and efficiency, respectively, that will be visualized in Figures 4.2 (a) and (b), respectively. There are several choices for most values of p , such as for instance for $p = 4$, one could use 1 node with 4 processes per node, 2 nodes with 2 processes per node, or 4 nodes with 1 process per node. In all cases, we use the smallest number of nodes possible, with 32 processes per node for $p \geq 32$; for $p < 32$, only one node is used, with the remaining cores idle. Comparing adjacent columns in the raw timing data in Table 4.2 (a) confirms our previous observation that performance improvement is very good from 1 to 2 processes, from 2 to 4 processes, from 4 to 8 processes but not quite as good from 8 to 16 processes and 16 to 32 processes. The speedup numbers in Table 4.2 (b) help reach the same conclusions when speedup is near-optimal with $S_p \approx p$ for $p \leq 8$. For $p = 16$ and $p = 32$, sub-optimal speedup is clear. The speedup numbers also indicate sub-optimal speedup for $p > 32$, but recall that the runtimes clearly showed halving from each column to the next

one; the speedup numbers can only give this indication qualitatively. The efficiency data in Table 4.2 (c) can bring out these effects more quantitatively, namely efficiency is near-optimal $E_p \approx 1$ for $p \leq 8$, then clearly identifies the efficiency drop taking place from $p = 8$ to $p = 16$ and from $p = 16$ to $p = 32$. But for $p > 32$, the efficiency numbers decrease only slowly, which confirms quantitatively the aforementioned halving of runtimes from each column to the next one for these columns.

The plots in Figures 4.2 (a) and (b) visualize the numbers in Tables 4.2 (b) and (c), respectively. These plots do not provide new data but simply provide a graphical representation of the results in Table 4.2. It is customary in results for fixed problem sizes that the speedup is better for larger problems, since the increased communication time for more parallel processes does not dominate over the calculation time as quickly as it does for small problems. This is born out generally by both plots in Figure 4.2. We can indeed see in both Figures 4.2 (a) and (b) that both speedup and efficiency get better, the larger N gets. The efficiency plot shows how bad the behavior is for $p \geq 8$, which supports the conclusion from earlier.

Table 4.2: Strong scalability study on taki 2018 using the default Intel compiler and Intel MPI, version 18.0.3.

(a) Wall clock time T_p in HH:MM:SS											
N	$p = 1$	$p = 2$	$p = 4$	$p = 8$	$p = 16$	$p = 32$	$p = 64$	$p = 128$	$p = 256$	$p = 512$	$p = 1024$
256	00:01:09	00:00:35	00:00:18	00:00:10	00:00:08	00:00:07	00:00:03	00:00:02	00:00:01	N/A	N/A
512	00:19:21	00:09:55	00:05:00	00:02:38	00:01:52	00:01:34	00:00:51	00:00:27	00:00:16	00:00:09	N/A
1024	05:34:36	02:50:40	01:27:40	00:43:58	00:29:57	00:25:04	00:12:53	00:06:44	00:03:56	00:02:05	00:01:13
2048	116:50:39	47:04:26	23:32:45	11:53:27	08:04:29	06:44:52	03:24:50	01:44:28	00:54:07	00:29:04	00:16:29
(b) Observed speedup $S_p = T_1/T_p$											
N	$p = 1$	$p = 2$	$p = 4$	$p = 8$	$p = 16$	$p = 32$	$p = 64$	$p = 128$	$p = 256$	$p = 512$	$p = 1024$
256	1.00	2.00	3.92	6.82	8.58	9.88	20.41	36.71	68.70	N/A	N/A
512	1.00	1.95	3.87	7.37	10.41	12.39	22.73	43.49	73.03	128.03	N/A
1024	1.00	1.96	3.82	7.61	11.17	13.34	25.98	49.64	85.23	161.12	275.16
2048	1.00	2.48	4.96	9.83	14.47	17.32	34.23	67.11	129.56	241.21	425.41
(c) Observed efficiency $E_p = S_p/p$											
N	$p = 1$	$p = 2$	$p = 4$	$p = 8$	$p = 16$	$p = 32$	$p = 64$	$p = 128$	$p = 256$	$p = 512$	$p = 1024$
256	1.00	1.00	0.98	0.85	0.54	0.31	0.32	0.29	0.27	N/A	N/A
512	1.00	0.98	0.97	0.92	0.65	0.39	0.36	0.34	0.29	0.25	N/A
1024	1.00	0.98	0.95	0.95	0.70	0.42	0.41	0.39	0.33	0.31	0.27
2048	1.00	1.24	1.24	1.23	0.90	0.54	0.53	0.52	0.51	0.47	0.42




Figure 4.2: Strong scalability study on taki 2018 using the default Intel compiler and Intel MPI, version 18.0.3.

Acknowledgments

The hardware used in the computational studies is part of the UMBC High Performance Computing Facility (HPCF). The facility is supported by the U.S. National Science Foundation through the MRI program (grant nos. CNS-0821258, CNS-1228778, and OAC-1726023) and the SCREMS program (grant no. DMS-0821311), with additional substantial support from the University of Maryland, Baltimore County (UMBC). See hpcf.umbc.edu for more information on HPCF and the projects using its resources. Co-author Carlos Barajas was supported by UMBC as HPCF RA.

References

- [1] Kritesh Arora, Carlos Barajas, and Matthias K. Gobbert. Parallel performance studies for an elliptic test problem on the Stampede2 cluster and comparison of networks. Technical Report HPCF-2018-10, UMBC High Performance Computing Facility, University of Maryland, Baltimore County, 2018.
- [2] Carlos Barajas and Matthias K. Gobbert. Parallel performance studies for an elliptic test problem on the 2018 portion of the Taki cluster. Technical Report HPCF-2018-18, UMBC High Performance Computing Facility, University of Maryland, Baltimore County, 2018.
- [3] Dietrich Braess. *Finite Elements*. Cambridge University Press, third edition, 2007.
- [4] James W. Demmel. *Applied Numerical Linear Algebra*. SIAM, 1997.
- [5] Anne Greenbaum. *Iterative Methods for Solving Linear Systems*, vol. 17 of *Frontiers in Applied Mathematics*. SIAM, 1997.
- [6] Arieh Iserles. *A First Course in the Numerical Analysis of Differential Equations*. Cambridge Texts in Applied Mathematics. Cambridge University Press, second edition, 2009.
- [7] Peter S. Pacheco. *Parallel Programming with MPI*. Morgan Kaufmann, 1997.
- [8] David Stonko, Samuel Khuvis, and Matthias K. Gobbert. Numerical methods to solve 2-D and 3-D elliptic partial differential equations using Matlab on the cluster maya. Technical Report HPCF-2014-9, UMBC High Performance Computing Facility, University of Maryland, Baltimore County, 2014.
- [9] Guan Wang and Matthias K. Gobbert. Performance comparison between blocking and non-blocking communications for a three-dimensional Poisson problem. Technical Report HPCF-2009-5, UMBC High Performance Computing Facility, University of Maryland, Baltimore County, 2009.
- [10] David S. Watkins. *Fundamentals of Matrix Computations*. Wiley, third edition, 2010.