

# Strong Scalability Studies for the 2-D Poisson Equation on the Taki 2021 Cluster with Historical Comparison

Ehsan Shakeri and Matthias K. Gobbert (gobbert@umbc.edu)

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

Technical Report HPCF-2024-1, [hpcf.umbc.edu](http://hpcf.umbc.edu) > Publications

## Abstract

The new 2021 nodes in the cluster taki in the UMBC High Performance Computing Facility contain two 24-core Intel Cascade Lake CPUs and 192 GB of memory per node, connected by an high-performance InfiniBand interconnect. Parallel performance studies for the memory-bound test problem of the Poisson equation in two spatial dimensions yield several conclusions: Strong scalability studies demonstrate excellent performance 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. For the largest numbers of processes per node, the runtime for the code is not significantly reduced, a typical behavior characteristic of memory-bound code. Comparisons to results on past clusters in HPCF bring out that core-per-core performance of serial code improvements has improved again, demonstrating the quality of the newest CPUs. Also node-per-node performance of parallel code continues to improve due to the larger number of cores available on a node. albeit we have fewer nodes in taki 2021 than we had in previous partitions.

## 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](http://hpcf.umbc.edu) for more information on HPCF and the projects using its resources.

In 2021, ten faculty from 9 departments and research centers in all 3 colleges at UMBC got together and showed their community spirit by contributing funding for the extension and state-of-the-art update of HPCF. These are the faculty listed by college and department: CNMS: Joe Bennett (Chemistry), Pengwang Zhai (Physics), Larrabee Strow (Physics), Matthias Gobbert (Mathematics and Statistics); COEIT: Tyler Josephson (CBEE), Curtis Menyuk (CSEE), Chuck Eggleton (Mechanical Engineering); CAHSS: Matt Baker (GES), Lee Boot (IRC and IS). The purchase consisted of 18 compute nodes, each with two 24-core Intel Xeon Gold 6240R Cascade Lake CPUs (2.4 GHz clock speed, 35.75 MB L3 cache, 6 memory channels, 165 W power), for a total of 48 cores per node, 192 GB of memory (12 × 16 GB DDR4 at 2933 MT/s) for a total of 4 GB per core, and a 480 GB SSD drive. Figure 1.1 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 16 GB. The nodes are connected by an EDR (Enhanced Data Rate) InfiniBand network (100 Gb/s bandwidth, 90 ns latency) to a central storage of more than 750 TB. See the system description at [hpcf.umbc.edu](http://hpcf.umbc.edu) for photos, schematics, and more detailed information, also on the other portions of the taki cluster.

This report uses the same test problem that has been used repeatedly to test cluster performance, including in 2020 [14] to compare different MPI communication modes, in 2019 [4] to compare different compilers, in 2018 [3] on taki for the first time and [2] on Stampede2, in 2014 [10] and 2015 [11] on maya, in 2010 [13, 15] on tara, in 2008 [7] on hpc, and in 2003 [1] on kali. The problem is the numerical solution of the Poisson equation with homogeneous Dirichlet boundary conditions on a unit square domain in two spatial dimensions. Discretizing the spatial derivatives by the finite difference method yields a system of linear equations with a large, sparse, highly structured, symmetric positive definite system matrix. This linear system is a classical test problem for iterative solvers and contained in several textbooks including [6, 8, 9, 16]. 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

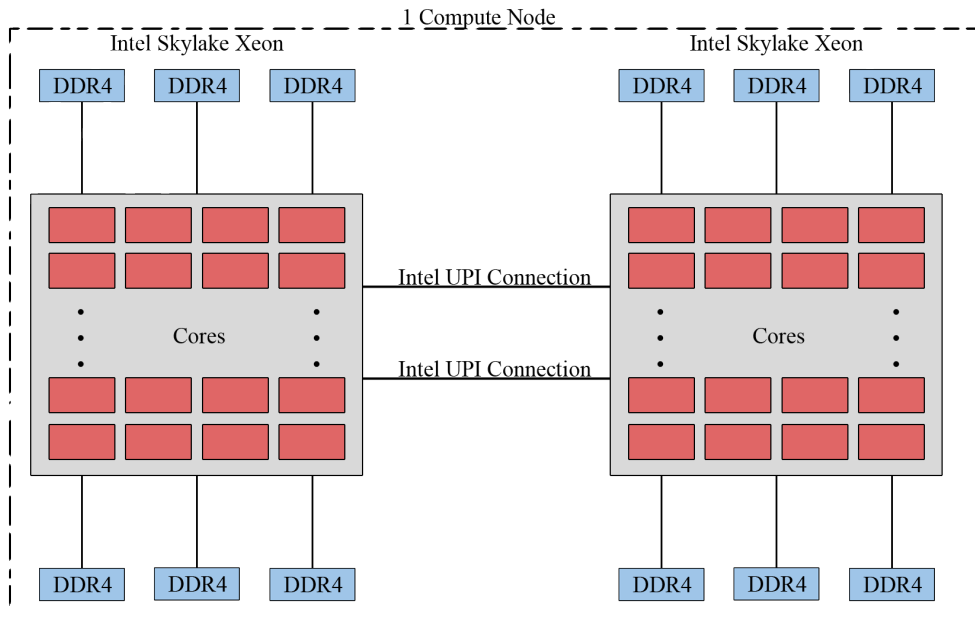


Figure 1.1: Schematic of a compute node with two Intel Cascade Lake Xeon CPUs.

language, with MPI (Message Passing Interface) for communications between distributed-memory cluster nodes, using the default Intel compiler and Intel MPI.

The studies in this report allow for a number of conclusions:

(1) The strong scalability studies in Section 4 demonstrate excellent performance when using multiple nodes: As we double the number of nodes, the runtime is almost halved. This observation confirms the quality of the InfiniBand interconnect. By progressively doubling the number of processes per node, the runtime is also halved for the smaller numbers of 1, 2, 4, and 8 processes per node. However, for the largest numbers of 16 and 32 processes per node, the runtime for the code is not significantly reduced. This behavior is a typical characteristic of memory-bound code such as this. The limiting factor in the performance of memory-bound code is memory access. Therefore, 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, as is indicated in the schematic of a compute node in Figure 1.1.

(2) The historical comparison in Section 5 of the same test problem on various past clusters at UMBC bring out that core-per-core performance of serial code improvements has improved again, demonstrating the quality of the newest CPUs. Also node-per-node performance of parallel code continues to improve due to the larger number of cores available on a node, albeit we have fewer nodes in taki 2021 than we had in previous partitions.

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 2021 using the default Intel compiler and Intel MPI. Section 5 provides a historical comparison of performance of the results in this report and previous clusters in HPCF.

## 2 The Elliptic Test Problem

We consider the classical elliptic test problem of the Poisson equation with homogeneous Dirichlet boundary conditions (see, e.g., [16, Chapter 8])

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

on the unit square domain  $\Omega = (0, 1) \times (0, 1) \subset \mathbb{R}^2$ . Here,  $\partial\Omega$  denotes the boundary of the domain  $\Omega$  and the Laplace operator in is defined as  $\Delta u = \frac{\partial^2 u}{\partial x_1^2} + \frac{\partial^2 u}{\partial x_2^2}$ . 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  $(x_{k_1}, x_{k_2}) \in \bar{\Omega} \subset \mathbb{R}^2$  with  $x_{k_i} = h k_i$ ,

$k_i = 0, 1, \dots, N, N + 1$ , in each dimension  $i = 1, 2$ . Denote the approximations to the solution at the mesh points by  $u_{k_1, k_2} \approx u(x_{k_1}, x_{k_2})$ . Then approximate the second-order derivatives in the Laplace operator at the  $N^2$  interior mesh points by

$$\frac{\partial^2 u(x_{k_1}, x_{k_2})}{\partial x_1^2} + \frac{\partial^2 u(x_{k_1}, x_{k_2})}{\partial x_2^2} \approx \frac{u_{k_1-1, k_2} - 2u_{k_1, k_2} + u_{k_1+1, k_2}}{h^2} + \frac{u_{k_1, k_2-1} - 2u_{k_1, k_2} + u_{k_1, k_2+1}}{h^2} \quad (2.2)$$

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

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

$$A = \begin{bmatrix} S & T & & & \\ T & S & T & & \\ & \ddots & \ddots & \ddots & \\ & & T & S & T \\ & & & T & S \end{bmatrix} \in \mathbb{R}^{N^2 \times N^2} \quad (2.3)$$

with the tri-diagonal matrix  $S = \text{tridiag}(-1, 4, -1) \in \mathbb{R}^{N \times N}$  for the diagonal blocks of  $A$  and with  $T = -I \in \mathbb{R}^{N \times N}$  denoting a negative identity matrix for the off-diagonal blocks of  $A$ .

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` [12]. 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` [12].

### 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 square  $\Omega = (0, 1) \times (0, 1)$  with right-hand side function  $f(x_1, x_2) = (-2\pi^2) (\cos(2\pi x_1) \sin^2(\pi x_2) + \sin^2(\pi x_1) \cos(2\pi x_2))$ , for which the true analytic solution in closed form  $u(x_1, x_2) = \sin^2(\pi x_1) \sin^2(\pi x_2)$  is known. On a mesh with  $32 \times 32$  interior points and mesh spacing  $h = 1/33 \approx 0.030303$ , the numerical solution  $u_h(x_1, x_2)$  can be plotted vs.  $(x_1, x_2)$  as a mesh plot as in Figure 3.1 (a). The shape of the solution clearly agrees with the true solution  $u(x_1, x_2)$  of the problem. At each mesh point, an error is incurred compared to the true solution  $u(x_1, x_2)$ . A mesh plot of the error  $u - u_h$  vs.  $(x_1, x_2)$  is shown in Figure 3.1 (b). We see that the maximum error occurs at the center of the domain of size about  $3 \times 10^{-3}$  (note the scale on the vertical axis), which compares well to the order of magnitude  $h^2 \approx 10^{-3}$  of the theoretically predicted error.

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  $N \times N$  mesh, the number of degrees of freedom  $N^2$  (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  $4N^2$  double-precision numbers needed to store the four vectors of significant length  $N^2$  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 [5, 9]. The fact that this convergence order is attained also




Figure 3.1: Mesh plots of (a) the numerical solution  $u_h$  vs.  $(x_1, x_2)$  and (b) the error  $u - u_h$  vs.  $(x_1, x_2)$ .

Table 3.1: Convergence study of the finite difference method with serial code.

$N$	DOF	$\ u - u_h\ $	Ratio	#iter	wall clock time		memory usage (GB)	
					HH:MM:SS	seconds	predicted	observed
1024	1,048,576	3.1266e-06	—	1,581	00:00:08	8.41	< 1	< 1
2048	4,194,304	7.8019e-07	4.01	3,192	00:01:22	82.34	< 1	< 1
4096	16,777,216	1.9366e-07	4.03	6,452	00:11:44	703.91	< 1	< 1
8192	67,108,864	4.7375e-08	4.09	13,033	01:36:08	5,768.04	2	2.17
16384	268,435,456	1.1545e-08	4.10	26,316	13:00:45	46,844.64	8	8.17

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, see [3] and the references therein. 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 2021 Using MPI-Only Code

This section presents the strong scalability studies using MPI-only code on taki 2021 using the default Intel compiler and Intel MPI. The Intel compiler `icc` and the Intel MPI implementation, currently version 2021.4.0 20210910, are accessed on taki through the wrapper `mpicc`. Since the compiler and MPI implementation are the defaults, they are available after the `module load default-environment` command in the `.bashrc` file in the user’s home directory that is automatically executed upon login to taki. We use the compiler options `-O3 -std=c99 -Wall` for the Intel Cascade Lake CPUs. We actually compile hybrid MPI+OpenMP code by including OpenMP multi-threading with the compiler option `-qopenmp`, but set the environment variable `OMP_NUM_THREAD` to 1 at run time. This limits the MPI+OpenMP code to 1 thread per MPI process, which is equivalent to running MPI-only code compiled without OpenMP.

HPCF uses the slurm workload manager ([slurm.schedmd.com](http://slurm.schedmd.com)) for job scheduling. The slurm submission script uses the `mpirun` command to start the job, with the option `-print-rank-map` that is supposed to print the MPI’s rank mapping. 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 by using the `--exclusive` flag. Correspondingly, we request all memory of the node for the job by `--mem=MaxMemPerNode`.

We conduct complete performance studies of the test problem for five progressively finer meshes of  $N = 1024, 2048, 4096, 8192, 16384$ . These studies result in progressively larger systems of linear equations with system dimensions ranging from about 1 million for  $N = 1024$  to over 250 million for  $N = 16384$ .

Table 4.1 collects the results of the performance studies on the 2021 portion of the CPU cluster in taki. For each mesh resolution of the five meshes with  $N = 1024, 2048, 4096, 8192, 16384$ , the parallel implementation of the test problem is run on all possible combinations of nodes from 1 to 16 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.,

Table 4.1: Wall clock time in HH:MM:SS using MPI-only code on taki 2021 using the Intel compiler and Intel MPI.

(a) Mesh resolution $N \times N = 1024 \times 1024$ , system dimension 1048576					
	1 node	2 nodes	4 nodes	8 nodes	16 nodes
1 process per node	00:00:08	00:00:03	00:00:02	00:00:01	00:00:00
2 processes per node	00:00:03	00:00:02	00:00:01	00:00:00	00:00:00
4 processes per node	00:00:02	00:00:01	00:00:01	00:00:00	00:00:00
8 processes per node	00:00:01	00:00:00	00:00:00	00:00:00	00:00:00
16 processes per node	00:00:00	00:00:00	00:00:00	00:00:00	00:00:00
32 processes per node	00:00:00	00:00:00	00:00:00	00:00:00	00:00:00
(b) Mesh resolution $N \times N = 2048 \times 2048$ , system dimension 4194304					
	1 node	2 nodes	4 nodes	8 nodes	16 nodes
1 process per node	00:01:22	00:00:39	00:00:18	00:00:07	00:00:03
2 processes per node	00:00:39	00:00:18	00:00:07	00:00:03	00:00:02
4 processes per node	00:00:18	00:00:07	00:00:03	00:00:02	00:00:01
8 processes per node	00:00:10	00:00:04	00:00:02	00:00:01	00:00:00
16 processes per node	00:00:06	00:00:02	00:00:01	00:00:00	00:00:00
32 processes per node	00:00:04	00:00:01	00:00:00	00:00:00	00:00:00
(c) Mesh resolution $N \times N = 4096 \times 4096$ , system dimension 16777216					
	1 node	2 nodes	4 nodes	8 nodes	16 nodes
1 process per node	00:11:44	00:05:52	00:02:50	00:01:20	00:00:36
2 processes per node	00:05:49	00:02:51	00:01:21	00:00:37	00:00:15
4 processes per node	00:02:50	00:01:21	00:00:37	00:00:15	00:00:07
8 processes per node	00:01:34	00:00:44	00:00:19	00:00:08	00:00:04
16 processes per node	00:01:00	00:00:29	00:00:11	00:00:04	00:00:02
32 processes per node	00:00:52	00:00:26	00:00:08	00:00:03	00:00:02
(d) Mesh resolution $N \times N = 8192 \times 8192$ , system dimension 67108864					
	1 node	2 nodes	4 nodes	8 nodes	16 nodes
1 process per node	01:36:08	00:48:24	00:24:04	00:12:02	00:05:51
2 processes per node	00:47:21	00:24:20	00:12:11	00:05:56	00:02:48
4 processes per node	00:24:30	00:12:12	00:06:00	00:02:47	00:01:16
8 processes per node	00:13:09	00:06:32	00:03:13	00:01:29	00:00:41
16 processes per node	00:09:03	00:04:26	00:02:13	00:01:01	00:00:24
32 processes per node	00:07:41	00:03:55	00:01:55	00:00:54	00:00:18
(e) Mesh resolution $N \times N = 16384 \times 16384$ , system dimension 268435456					
	1 node	2 nodes	4 nodes	8 nodes	16 nodes
1 process per node	13:00:45	06:29:59	03:16:26	01:37:56	00:49:10
2 processes per node	06:31:07	03:20:06	01:40:00	00:49:40	00:24:59
4 processes per node	03:15:49	01:40:39	00:49:47	00:24:53	00:12:08
8 processes per node	01:45:44	00:53:18	00:26:39	00:13:18	00:06:35
16 processes per node	01:11:44	00:36:33	00:18:14	00:09:19	00:04:34
32 processes per node	01:04:36	00:32:10	00:16:06	00:08:17	00:04:07

the serial run, of the code for that particular mesh. The lower-right entry of each subtable lists the runtime using 32 cores on 16 nodes for a total of 512 parallel processes working together to solve the problem. Notice that each node has two 24-core CPUs for a total of 48 cores, so even with 32 processes per node, many cores are not used by our job and remain available for the operating system and other system tasks.

We choose the mesh resolution of  $16384 \times 16384$  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 runtime from 16 to 32 processes is small, the runtimes 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 indicated in Figure 1.1.

Reading along each row of the  $16384 \times 16384$  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 16 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 for instance, the runtime for 2 nodes with 1 processes per node is almost same as for 1 node with 2 processes per node. Thus, it is advisable to use the smallest number of nodes with the largest number of processes per node.

When comparing now all subtables in Table 4.1, we observe that when we double the size of the mesh from one subtable to the next, the runtimes increase by a factor of about 8 to 10 for corresponding entries. The relative performance in each of the subtables in Table 4.1 exhibits largely analogous behavior to the  $16384 \times 16384$  mesh, in particular the  $8192 \times 8192$  mesh subtable. For smaller meshes, some times for larger numbers of nodes are eventually so fast that improvement is small with more processes per node, but behavior is analogous for the more significant times.

## 5 Historical Comparison of Performance Studies on Different Clusters

This section contains a comparison of performance studies for the test problem obtained on different distributed-memory clusters that existed at UMBC over time. The oldest cluster kali from 2003 was the first cluster with a high-performance interconnect at UMBC, that is, it was the first to have a second network dedicated to MPI communications, in addition to the primary management network. These clusters, acquired over the years, give an impression what was the state-of-the-art CPU model and what constituted a high-performance interconnect at that time, on the concrete example of the test problem used in this report. Additionally, we include the cluster Stampede2 in the Texas Advanced Computing Center (TACC), a nationally funded supercomputer, which had at the time the same generation CPU as taki 2018.

Table 5.1 reports results for the test problem on the historical mesh resolution of  $N = 4096$ , which was the largest resolution that could be solved on kali in 2003. The following list briefly describes each cluster in the table, see the cited references for detailed descriptions.

**Cluster kali (2003) [1]:** This cluster was a 33-node distributed-memory cluster with 32 compute nodes including a storage node (with extended memory of 4 GB) containing the 0.5 TB central storage, each with two (single-core) Intel Xeon processors (2.0 GHz clock speed) and 1 GB of memory per node, connected by a Myrinet interconnect, plus 1 combined user/management node. Note that for the case of all cores on 1 node, that is, for the case of both (single-core) CPUs used simultaneously, the performance was worse than for 1 CPU and hence the results were not recorded at the time. The cluster kali was funded by a SCREMS grant from the National Science Foundation and owned by the Department of Mathematics and Statistics at UMBC.

**Cluster hpc (2008) [7]:** The cluster hpc was a 35-node distributed-memory cluster with 33 compute nodes plus 1 development and 1 combined user/management node, each with two dual-core AMD Opteron processors and at least 13 GB of memory, connected by a DDR InfiniBand network and with an InfiniBand-accessible 14 TB parallel file system. The cluster hpc was the first cluster of HPCF, when HPCF was formed by seed funding from UMBC and contributions of faculty in 2008.

**Cluster tara (2009) [13]:** The cluster tara was an 86-node distributed-memory cluster with two quad-core Intel Nehalem X5550 processors (2.6 GHz clock speed, 8 MB L3 cache, 3 memory channels) and 24 GB memory per node, a QDR InfiniBand interconnect, and 160 TB central storage. This cluster was the first one principally funded by an MRI grant from the National Science Foundation to the user community of HPCF.

**Clusters maya 2009, 2010, 2013 (2015) [11]:** The cluster tara became part of the cluster maya as maya 2009, and its QDR InfiniBand network was extended to the 2013 portion of maya. The results for maya 2009 used the identical hardware as the results for tara, but new middleware and new versions of the compiler and MPI implementation from 2013. The cluster maya 2010 was a 168-node distributed-memory cluster with two quad-core Intel Nehalem X5560 processors (2.8 GHz clock speed, 8 MB L3 cache, 3 memory channels) and 24 GB memory per node, connected by a DDR InfiniBand interconnect. The cluster maya 2013 was a 72-node cluster with two eight-core Intel E5-2650v2 Ivy Bridge CPUs (2.6 GHz clock speed, 20 MB L3 cache, 4 memory channels) and 64 GB memory per node, connected by same QDR InfiniBand as maya 2009. The cluster maya had a central storage of more than 750 TB.

**Cluster taki 2018 (2019) [4]:** The 2018 portion of the CPU cluster in taki 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

Table 5.1: Runtimes (speedup) for  $N = 4096$  on the clusters kali, hpc, tara, maya, taki 2018, taki 2021, and Stampede2. “All nodes” refers to 32 nodes, except for taki 2021 with 16 nodes.

Cluster (year)	serial	1 node	all nodes	all nodes
	(1 core) time	all cores time (speedup)	1 core per node time (speedup)	all cores time (speedup)
kali (2003) [1]	02:00:49	N/A (N/A)	00:04:05 (29.59)	00:04:49 (25.08)
hpc (2008) [7]	01:51:29	00:32:37 (3.42)	00:03:23 (32.95)	00:01:28 (76.01)
tara (2009) [13]	00:31:16	00:06:39 (4.70)	00:01:05 (28.86)	00:00:09 (208.44)
maya 2009 (2015) [11]	00:17:05	00:05:55 (2.89)	00:00:35 (29.29)	00:00:08 (128.13)
maya 2010 (2015) [11]	00:17:00	00:05:48 (2.93)	00:00:34 (30.00)	00:00:06 (170.00)
maya 2013 (2015) [11]	00:12:26	00:02:44 (4.55)	00:00:15 (49.07)	00:00:12 (62.17)
taki 2018 (2019) [4]	00:13:38	00:01:03 (12.98)	00:00:14 (58.43)	00:00:01 (818.00)
taki 2021 (Section 4)	00:11:44	00:00:52 (13.54)	00:00:36 (19.56)	00:00:02 (352.00)
Stampede2 (2018) [2]	00:13:04	00:01:00 (13.07)	00:00:16 (49.00)	00:00:01 (784.00)

power), for a total of 36 cores per node, and 384 GB memory per node ( $12 \times 32$  GB DDR4 at 2666 MT/s), for a total of 10.6 GB per core, connected by an EDR InfiniBand interconnect.

**Cluster taki 2021:** The 2021 portion of the CPU cluster in taki consists of 18 compute nodes, each with two 24-core Intel Xeon Gold 6240R Cascade Lake CPUs (2.4 GHz clock speed, 35.75 MB L3 cache, 6 memory channels, 165 W power), for a total of 48 cores per node, and 192 GB of memory ( $12 \times 16$  GB DDR4 at 2933 MT/s) for a total of 4 GB per core, connected by an EDR InfiniBand interconnect.

**Cluster Stampede2 (2018) [2]:** The last row in the table contains results for the Skylake portion of Stampede2 with each compute node having two 24-core Intel Xeon Platinum 8160 Skylake CPUs (2.1 GHz clock speed, 33 MB L3 cache, 6 memory channels) and 192 GB memory per node, connected by an Omni-Path network.

**Results on kali:** On the cluster kali from 2003, we observe a factor of approximately 30 speedup by increasing the number of nodes from 1 to 32. However by using both cores on each node we only see a factor of approximately 25 speedup. We do not observe the expected 64 factor speedup, since both CPUs on the node share a bus connection to the memory, which leads to contention in essentially synchronized algorithms like Krylov subspace methods. Hence, it is actually faster to leave the second CPU idling rather than to use both [1].

**Results on hpc:** Note that there are four cores on each node of cluster hpc from 2008, compared to just two on the cluster kali, since the CPUs are dual-core. We observe approximately fourfold speedup that we would expect by running it on four cores rather than one. By running on 32 nodes with one core per node we observe the expected speedup of approximately 32; more in detail, the speedup is slightly better than optimal, which is explained by the smaller portions of the subdivided problem on each node fitting better into the cache of the processors. We see this for the first time here, but it is a typical effect in strong performance studies, in which a problem that already fits on one node is divided into smaller and smaller pieces as the number of nodes grows. Finally, by using all cores on 32 nodes we observe a speedup of 76.01, less than the optimal speedup of 128, for a still respectable efficiency of 59.38% [7].

**Results on tara:** On the cluster tara from 2009, we observe a less than optimal speedup of approximately 5 by running on all 8 cores rather than on one, caused by the cores of a CPU competing for memory access. By running on 32 nodes with one core per node we observe a speedup of approximately 30. Finally, by using all 8 cores on 32 nodes we observe a speedup of 208, less than the optimal speedup of 256, but an excellent efficiency of 81.25% [13].

**Results on maya:** On maya 2009, we observe that by running on all 8 cores on a single node rather than 1 core there is a speedup of approximately 3 rather than the optimal speedup of 8. By running on 32 nodes with one core per node we observe a speedup of approximately 29 of the possible 32, which is very good. Finally, by using all 8 cores on 32 nodes we observe a speedup of 128, half of the optimal speedup of 256, or an efficiency of 50.00% [11]. On maya 2010, we observe that by running on all 8 cores on a single node rather than 1 core there is a speedup of approximately 3 rather than the optimal speedup of 8. By running on 32 nodes with one core per node we observe a speedup of approximately 30. When combining the use of all cores with the use of 32 nodes, we observe a speedup of 170, short of the optimal speedup of 256, or an efficiency of 66.41% [11]. On maya 2013, we observe that by running on all 16 cores on a single node rather than on one core there is a speedup of approximately 5 rather than the expected speedup of 16. We observe a greater than optimal speedup of 49.07 by running on 32 nodes with one process per node; this is caused by the relatively small problem fitting into cache after dividing it onto 32 nodes, together with the quality of the QDR InfiniBand interconnect. [11].



**Results on taki 2018:** Recall that each node in taki 2018 has two 18-core Intel Skylake CPUs, for a total of 36 cores available on each node. The data for “all cores” on 1 node and on 32 nodes uses actually 32 cores per node. We can observe by reading down the first column of Table 5.1 that in the past the core per core performance improved each time. But this is not true any more now, as serial code on a single core in the Skylake processor is not faster than in the Ivy Bridge CPUs from 2013; compare the clock speeds of each CPU specified above. This demonstrates that using multi-threading and/or MPI on a shared-memory node is absolutely vital to achieving optimal and scalable performance today and in the future. The possibility of this can be seen by the 13-fold speedup of the code that is possible when using “all cores” of 1 node, which is an efficiency of only 41%, but still yields a time that is more than twice as fast than using all cores of maya 2013. In fact, when comparing the “all cores” results of taki 2018 to the previous rows, we see that they are the fastest in each column and an improvement over all past results. We also see that the EDR InfiniBand on taki 2018 provides for excellent speedup, as the 32-node results have better than optimal speedup, both comparing the 1-core and the 32-core cases, namely 58.43 instead of nominal 32 for the 1-core comparison, for instance. Then, for the all core comparison, the speedup of 818.00 of the nominal 1024 for 32 nodes with 32 cores gives an efficiency of 80% compared to the serial run, but when compared to all cores on 1 node it is a speedup of 63.00, which is actually better than the nominal 32 improvement to be expected. We note that these better than optimal results typically result from the problem being divided into such small chunks on each core that they fit in cache there, but the multi-node performance is still a testament to the low latency of the network.

**Results on taki 2021:** Recall that each node in taki 2021 has two 24-core Intel Skylake CPUs, for a total of 48 cores available on each node, significantly more than in taki 2018; however, sticking with powers of 2, “all cores” still refers to 32 cores per node, the same as for taki 2018. Also recall that we only have 18 nodes in taki 2021, so “all nodes” refers to 16 nodes for taki 2021, which used to refer to 32 nodes in all other rows of the table. Reading down the first column of Table 5.1, we can see that the single-core performance of taki 2021 is the fastest now, demonstrating the quality of newer CPUs. The speedup to 32 cores per node is still about 13, just like in taki 2018, thus yielding a faster runtime also when using all cores in taki 2021. When noting the longer time reported for “all nodes” in taki 2021, remember again that these are only 16 nodes, half as many as in the other rows. Then, the observed runtime cannot be criticized, and in fact, the speedup of about 20 for 16 nodes over 1 node with 1 process per node each is better than the expected factor of 16. Thus, we can conclude that the 2021 purchase resulted in quality hardware that stands up well to historical comparisons over the years.

**Results on Stampede2:** The last row of the table shows the results for the Stampede2 cluster, which has the same generation Intel Skylake CPUs as taki 2018, but with 24 cores each instead of 18, for a total of 48 cores per node. Like for taki 2018, the data for “all cores” on 1 node and on 32 nodes uses actually 32 cores per node of the available 48 cores. The observations for taki 2018 also apply here in essentially the same way. It is notable that the serial run is actually slightly faster than on taki 2018, despite the nominally lower clock rate on Stampede2; see above. But the 1-core on 32 nodes result is not faster, possibly owing to a different network, and none of the differences between taki 2018 and Stampede2 are significant.

## 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, OAC-1726023, and CNS-1920079) and the SCREMS program (grant no. DMS-0821311), with additional substantial support from the University of Maryland, Baltimore County (UMBC). See [hpcf.umbc.edu](http://hpcf.umbc.edu) for more information on HPCF and the projects using its resources. Co-author Ehsan Shakeri was supported by UMBC as HPCF RA.

## References

- [1] Kevin P. Allen. Efficient parallel computing for solving linear systems of equations. *UMBC Review: Journal of Undergraduate Research and Creative Works*, vol. 5, pp. 8–17, 2004.
- [2] 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.

- [3] 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.
- [4] Carlos Barajas and Matthias K. Gobbert. Strong and weak scalability studies for the 2-D Poisson equation on the Taki 2018 cluster. Technical Report HPCF–2019–1, UMBC High Performance Computing Facility, University of Maryland, Baltimore County, 2019.
- [5] Dietrich Braess. *Finite Elements*. Cambridge University Press, third edition, 2007.
- [6] James W. Demmel. *Applied Numerical Linear Algebra*. SIAM, 1997.
- [7] Matthias K. Gobbert. Parallel performance studies for an elliptic test problem. Technical Report HPCF–2008–1, UMBC High Performance Computing Facility, University of Maryland, Baltimore County, 2008.
- [8] Anne Greenbaum. *Iterative Methods for Solving Linear Systems*, vol. 17 of *Frontiers in Applied Mathematics*. SIAM, 1997.
- [9] Arieh Iserles. *A First Course in the Numerical Analysis of Differential Equations*. Cambridge Texts in Applied Mathematics. Cambridge University Press, second edition, 2009.
- [10] Samuel Khuvis and Matthias K. Gobbert. Parallel performance studies for an elliptic test problem on maya 2013. Technical Report HPCF–2014–6, UMBC High Performance Computing Facility, University of Maryland, Baltimore County, 2014.
- [11] Samuel Khuvis and Matthias K. Gobbert. Parallel performance studies for an elliptic test problem on the cluster maya. Technical Report HPCF–2015–6, UMBC High Performance Computing Facility, University of Maryland, Baltimore County, 2015.
- [12] Peter S. Pacheco. *Parallel Programming with MPI*. Morgan Kaufmann, 1997.
- [13] Andrew M. Raim and Matthias K. Gobbert. Parallel performance studies for an elliptic test problem on the cluster tara. Technical Report HPCF–2010–2, UMBC High Performance Computing Facility, University of Maryland, Baltimore County, 2010.
- [14] Jinglong Sun, Ehsan Shakeri, and Matthias K. Gobbert. Strong scalability studies for the 2-D and 3-D Poisson equations on the Taki 2018 cluster. Technical Report HPCF–2020–1, UMBC High Performance Computing Facility, University of Maryland, Baltimore County, 2020.
- [15] Hafez Tari and Matthias K. Gobbert. A comparative study of the parallel performance of the blocking and non-blocking MPI communication commands on an elliptic test problem on the cluster tara. Technical Report HPCF–2010–6, UMBC High Performance Computing Facility, University of Maryland, Baltimore County, 2010.
- [16] David S. Watkins. *Fundamentals of Matrix Computations*. Wiley, third edition, 2010.