

# HPC Challenge Benchmarks in Chapel\*

Bradford L. Chamberlain    Sung-Eun Choi    Steven J. Deitz    David Iten

Cray Inc.  
chapel.info@cray.com

## Abstract

This report presents our best to-date Chapel implementations of the global HPC Challenge benchmarks STREAM Triad, Random Access, FFT, and HPL. They improve upon our 2008 Chapel submission using the same hardware. The highlights of this year's submission include:

- Global STREAM Triad performance of 10.8 TB/s on 8192 cores of a Cray XT4 (up from 1.69 TB/s)
- EP STREAM Triad performance of 12.2 TB/s on 8192 cores of a Cray XT4.
- Random Access performance of 0.122 GUP/s on 8192 cores of a Cray XT4 (up from 0.0011 GUP/s)
- A first distributed implementation of FFT which makes use of two distinct distributions: Block and Cyclic.
- A demonstration of Chapel's portability on a Cray XT4, a Cray CX1, and an IBM pSeries 575.
- The Chapel compiler and these benchmarks are publicly available at <http://sourceforge.net/projects/chapel>.

All codes in this report compile and execute correctly with version 1.0.2 of the Chapel compiler. All reported Chapel performance results were obtained with this version of the Chapel compiler. The full code listings are provided in appendices to this report.

## 1 Overview and Contents

Chapel is a new parallel programming language under development at Cray Inc. as part of its participation in DARPA's High Productivity Computing Systems (HPCS) program.<sup>1,2</sup> The goal of the Chapel project is to improve parallel programmability, portability, and code robustness as compared to current programming models while producing programs with performance comparable to or better than MPI. Chapel is very much a work in progress, and as such, this report should be viewed as a snapshot of Chapel's current status.

In this report, we present our best to-date Chapel implementations of four HPC Challenge (HPCC) benchmarks<sup>3,4</sup>—STREAM Triad, Random Access, FFT, and HPL. We summarize the Chapel implementations by categorizing and counting the source lines of code. For each benchmark, we provide a brief overview of our Chapel implementation. For STREAM Triad and Random Access, we present performance results on up to 2048 nodes of a Cray XT4. We provide discussions of our ongoing implementation work for the FFT and HPL benchmarks.

### Contents

<i>Code Size Summary</i> .....	2	<i>Random Access Performance</i> .....	8	<i>EP STREAM Triad Code</i> .....	12
<i>Experimental Setup</i> .....	3	<i>FFT Discussion</i> .....	9	<i>Global Random Access Code</i> .....	14
<i>Global STREAM Triad Description</i> .....	4	<i>HPL Discussion</i> .....	9	<i>Global FFT Code</i> .....	15
<i>EP STREAM Triad Description</i> .....	5	<i>Portability</i> .....	10	<i>Global HPL Code</i> .....	17
<i>STREAM Triad Performance</i> .....	6	<i>Summary</i> .....	11	<i>Shared Problem Size Module Code</i> .....	19
<i>Random Access Description</i> .....	7	<i>Global STREAM Triad Code</i> .....	12		

\*This material is based upon work supported by the Defense Advanced Research Projects Agency under its Agreement No. HR0011-07-9-0001. This research used resources of the National Center for Computational Sciences at Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

<sup>1</sup><http://www.darpa.mil/IPTO/programs/hpcs/hpcs.asp>

<sup>2</sup><http://www.highproductivity.org/>

<sup>3</sup><http://icl.cs.utk.edu/hpcc/>

<sup>4</sup><http://www.hpccchallenge.org/>

## 2 Code Size Summary

The following table categorizes and counts the number of lines of code in our HPCC implementations:

<b>Code Category</b>	<i>Global STREAM Triad</i>	<i>EP STREAM Triad</i>	<i>Global Random Access</i>	<i>Global FFT</i>	<i>Global HPL</i>	<i>Problem Size</i>
Computation	2	6	7 + 25 = 34	75	50	0
Declarations	12	8	20 + 12 = 32	32	63	34
<b>Total kernel</b>	<b>14</b>	<b>14</b>	<b>27 + 37 = 66</b>	<b>107</b>	<b>113</b>	<b>34</b>
Initialization	10	10	1 + 10 = 11	26	8	0
Verification	8	10	9 + 0 = 9	11	16	0
Results/output	32	43	21 + 0 = 21	21	39	21
<b>Total benchmark</b>	<b>64</b>	<b>77</b>	<b>58 + 47 = 107</b>	<b>165</b>	<b>176</b>	<b>55</b>
Debug/test	7	0	3 + 2 = 3	5	1	0
Comments	72	131	94 + 31 = 125	140	170	39
Blank	28	29	23 + 8 = 31	47	61	8
<b>Total program</b>	<b>171</b>	<b>237</b>	<b>178 + 88 = 266</b>	<b>357</b>	<b>408</b>	<b>102</b>

The line counts for each benchmark are represented using a column of the table. The final data column represents the shared *HPCCProblemSize* module that is used by the benchmarks to automatically compute the appropriate problem size for a machine and to print it. For the Random Access benchmark, each entry is expressed as a sum—the first value represents the benchmark module itself, the second represents a helper module used to define the stream of pseudo-random update values, and the final value is the sum of the two.

The rows of the table are used to group the lines of code into various categories and running totals. The first two rows indicate the number of lines required to express the kernel of the computation and its supporting declarations, respectively. For example, in the STREAM Triad benchmark, writing the computation takes two lines of code, while its supporting variable and subroutine declarations require eleven lines of code. The next row presents the sum of these values to indicate the total number of lines required to express the kernel computation—thirteen in the case of STREAM.

The next three rows of the table count lines of code related to setup, verification, and tear-down for the benchmark. *Initialization* indicates the number of lines devoted to initializing the problem's data set, *Verification* counts the lines used to check that the computed results are correct, and *Results and Output* gives the number of lines for computing and outputting results for timing and performance. These three rows are then combined with the previous subtotal giving the number of source lines used to implement the benchmark and output its results. This subtotal should be interpreted as the SLOC (*Source Lines of Code*) count for the benchmark as specified.

The *Debug and Test* row indicates the number of lines added to make the codes more useful in our nightly regression testing system, while the *Comments* row indicates the number of comment lines and the *Blank* row indicates the number of blank lines. These values are added to the previous subtotal to give the total number of lines in each program, and they serve as a check sum against the line number labels that appear in the appendices.

### 3 Experimental Setup

Our main experimental platform was the Cray XT4 partition of Jaguar at Oak Ridge National Laboratory (ORNL). The following table provides a brief overview of these platforms:

<i>Machine Characteristics</i>	<i>XT4</i>
Number of compute nodes	7,832
Compute node processor	Quad-core AMD Opteron
Processor speed	2.1 GHz
Memory per node	8 GB

The problem sizes used for STREAM Triad and Random Access were computed by quartering the result returned by Chapel’s built-in method `physicalMemory`, which returns the amount of physical memory on a locale. (For Random Access, this was then rounded up to the nearest power of two.) For both benchmarks, this is the smallest possible problem size that exceeds 25% of the total memory as listed in `/proc/meminfo`. These same problem sizes were used for our runs of the MPI and OpenMP reference versions of these benchmarks by brutally inserting these values into the elaborate framework surrounding these codes. The following table specifies the per-node problem sizes used to obtain the results in this report:

<i>Problem Sizes Per Node</i>	<i>XT4</i>
STREAM Triad Problem Size	85983914 (1.92 GB)
Random Access Problem Size	$2^{28}$ (2.0 GB)
Random Access Number of Updates	$2^{18}$

Note that we execute a reduced number of updates ( $2^{n-10}$  instead of  $2^{n+2}$  where  $n$  is large enough to create a table that uses 25% of the total system memory) for Random Access due to long execution times. The GUP/s rate does not appear to be affected by this change.

The Chapel compiler is a source-to-C translator that invokes a C compiler to create an executable. We used the same C compiler, with the same flags, to finish the Chapel compilation and to compile the MPI and OpenMP reference implementations of the benchmarks. The Chapel runtime uses POSIX threads (*pthreads*) to implement tasks and Berkeley’s GASNet communication library<sup>5</sup> for inter-process coordination and data transfer. The following table specifies software versions and settings used to obtain the results in this report:

<i>Software</i>	<i>Flags/Settings</i>	<i>Version</i>
chpl	<code>--fast</code>	1.0.2
gcc	<code>-target=linux -O3 -std=c99 -fopenmp</code> <code>--param max-inline-insns-single=35000</code> <code>--param inline-unit-growth=10000</code> <code>--param large-function-growth=200000</code>	4.3.2
GASNet	<code>conduit=mpi, segment=fast</code>	1.12.0

On the Cray XT4, we used Cray’s *PrgEnv-gnu* programming environment module which provides a Cray C compiler wrapper around `gcc`.

The Chapel flag “`--fast`” turns off a number of runtime checks that are enabled by default for safety, including checks for out-of-bounds array accesses, null pointer dereferences, and violations of locality assertions. The flags used for the C compilation were chosen by GASNet’s auto-configuration process and were used both for the generated Chapel code and the HPC reference implementations. The GASNet conduit and segment choices are primarily used for GASNet portability. We did not use the portals conduit due to a combination of GASNet and Chapel bugs. Since the portals conduit is chosen by the Chapel system by default, we explicitly set the environment variable `CHPL_COMM_SUBSTRATE` to `mpi` and `CHPL_GASNET_SEGMENT` to `fast`.

<sup>5</sup>For details on GASNet, refer to the GASNet specification.

## 4 Global STREAM Triad Description

The core of the global version of STREAM Triad in Chapel is unchanged from last year:

```
forall (a, b, c) in (A, B, C) do
  a = b + alpha * c;
```

This pair of lines specifies parallel, element-wise iteration over the vectors  $A$ ,  $B$ , and  $C$ , referring to corresponding elements as  $a$ ,  $b$ , and  $c$  in the loop body.

The distributed implementation of these vectors and the parallel implementation of the loop are both specified by the *distribution* of  $A$ ,  $B$ , and  $C$  via a series of three declarations.

The first declaration<sup>6</sup>

```
const BlockDist = distributionValue(new Block(rank=1, bbox=[1..m], tasksPerLocale=tasksPerLocale));
```

creates a distribution *BlockDist* that maps indices across the entire set of *locales*<sup>7</sup> according to the implementation of the *Block* class. This mapping is computed by partitioning the specified bounding box,  $1 \dots m$ , across the locales using blocks of approximately equal sizes. The *tasksPerLocale* argument, which is passed the value of a *configuration constant* of the same name, specifies how many tasks should be used on each locale to implement parallel loops over the distribution's domains and arrays.

The second declaration

```
const ProblemSpace: domain(1, int(64)) distributed BlockDist = [1..m];
```

creates a *domain*—a first-class language concept representing an index set—to describe the set of indices that define the problem space. This domain, *ProblemSpace*, is declared to be a 1-dimensional domain of 64-bit integer indices, distributed using the *BlockDist* distribution. It is initialized to store the index set  $1 \dots m$  which will be divided between the locales according to the mapping defined by *BlockDist*.

The third declaration

```
var A, B, C: [ProblemSpace] elemType;
```

creates the three vectors,  $A$ ,  $B$ , and  $C$ , specifying that each index in *ProblemSpace* should be mapped to a variable of type *elemType* (defined previously to be a 64-bit real floating-point value). The elements of the vectors are mapped to the same locales as the corresponding indices in the *ProblemSpace* domain.

Chapel distributions like *Block* not only map domain indices and array elements to locales, they also serve as recipes for mapping high-level operations—such as the forall loop used for the computation—down to the individual data structures and tasks that will implement the computation across the locales. In the case of a forall loop like this one, the compiler rewrites the loop using multiple iterators defined by the distribution which specify how parallel, element-wise iteration should be implemented for its domains and arrays. The distribution itself is written in Chapel using standard features such as *coforall loops* to create tasks and *on-clauses* to specify the locales on which the tasks should run. The compiler contains no semantic knowledge specific to the *Block* distribution, only that distributions implement a structural interface that it can target when lowering and optimizing high-level operations. This lack of distribution-specific knowledge is the backbone of our plan to support user-defined distributions and a large collection of provided distributions.

---

<sup>6</sup>The current distribution syntax is temporary, pending a better way to distinguish between the class that implements a distribution and the distribution value over which a domain is declared.

<sup>7</sup>A Chapel *locale* is an architectural unit of locality. Tasks running within a locale are considered to have uniform access to local data; they can also access data in other locales, but with greater overhead. On a commodity cluster, a multicore processor or SMP node would typically be considered a locale. On the Cray XT4, it is a single quadcore node.

## 5 EP STREAM Triad Description

New to this year's Chapel entry, we present both EP (embarrassingly parallel) and global versions of the STREAM Triad benchmark. The global version is far more elegant in Chapel due to its support for a global-view programming model. Nevertheless, Chapel's multi-level design allows us to fragment execution across the locales and implement an EP version of the STREAM Triad benchmark as well.

The ability to abandon Chapel's global-view array abstractions and elegantly step into an explicit SPMD-style programming model is in stark contrast to most previous languages with support for global arrays. We believe that Chapel's support for multiple levels of design is of the utmost importance for high-level languages that seek to support both programmability and performance, if for no other reason than to enable programmers to work around cases where the compiler or high-level abstractions fail them.

The core of the EP version of STREAM Triad uses a *coforall* loop over all of the locales to create a new task per locale. When coupled with the *on* statement to specify where these tasks should execute, the following loop creates remote task on each locale:

```
coforall loc in Locales do on loc {
```

Within this loop, we can create non-distributed arrays to contain the locale's portion of the STREAM vectors as follows:


```
var A, B, C: [1..m] elemType;
```

For this version of the benchmark, like the reference version, *m* refers to a per-locale problem size. The actual computation looks identical to the computation in the global version above:

```
forall (a, b, c) in (A, B, C) do  
  a = b + alpha * c;
```

In our implementation of this benchmark, we use a *local* block to guarantee that there will be no communication within this computation. This results in a faster implementation with our current version of the Chapel compiler.

## 6 STREAM Triad Performance



Locales	Global Chapel				EP Chapel 4 TPL	EP MPI				EP MPI + OpenMP			
	1 TPL	2 TPL	3 TPL	4 TPL		1 PPN	2 PPN	3 PPN	4 PPN	1 TPN	2 TPN	3 TPN	4 TPN
local	3.80	5.80	6.14	5.97	6.15								
1	3.64	5.59	6.03	6.03	6.15	4.30	6.16	6.34	6.65	4.28	6.14	6.32	6.38
2	7.26	11.09	11.97	11.69	12.26	8.62	12.37	12.69	12.79	8.56	12.30	12.71	12.74
4	11.19	17.71	21.23	21.90	24.51	17.24	24.77	25.45	25.62	17.14	24.59	25.43	25.48
8	28.85	43.51	47.44	46.62	49.04	34.46	49.50	50.83	51.29	34.26	49.17	50.63	49.65
16	57.70	87.63	94.74	96.11	96.95	68.13	98.48	101.67	103.64	67.43	97.78	101.51	100.83
32	116.22	175.60	189.52	184.23	195.89	137.71	198.02	202.99	203.06	136.96	196.70	202.37	198.60
64	231.61	354.28	378.55	369.46	392.33	275.89	397.03	405.22	405.39	269.54	390.95	406.62	404.65
128	458.08	689.86	751.34	740.20	784.63	551.10	792.50	803.33	794.76	547.72	786.69	809.55	794.79
256	918.85	1394.52	1492.43	1388.25	1569.33	1103.24	1583.39	1611.91	1605.32	1097.32	1573.16	1617.36	1587.71
512	1823.46	2747.25	2949.44	2815.62	3105.15	2185.09	3152.08	3239.64	3232.55	2164.25	3133.62	3250.04	3221.81
1024	3604.22	5371.51	5602.11	5433.64	6265.50	4401.32	6323.29	6441.16	6436.95	4330.40	6266.72	6498.65	6443.33
2048	7046.64	10278.50	11058.00	10346.20	12543.12	8814.57	12677.93	12858.80	12628.04	8767.55	12587.24	12961.92	12725.32

This efficiency graph shows the percentage efficiency of the global and EP Chapel versions and the MPI and MPI+OpenMP reference versions of the STREAM Triad benchmark. For the global Chapel version, results are reported for between 1 and 4 tasks per locale (TPL). For the MPI version, results are reported for between 1 and 4 processes per node (PPN). For the MPI+OpenMP version, results are reported for between 1 and 4 threads per node (TPN). The current implementation of Chapel does not allow for varying the number of tasks per locale for non-distributed arrays. The efficiency is taken with respect to linear scaling of the best 1-locale performance (6.64 GB/s for the MPI version with 4 MPI processes per node). The *local* results refer to compiling the Chapel codes with the “--local” flag. This flag optimizes the program for running on a single locale.

We make the following observations:

- No global version can perform as well as an EP version because there is inherent synchronization overhead in the global version that does not exist in the EP version. The global Chapel version shows little overhead and scales very well to 512 locales and reasonably well to 2048 locales.
- The EP Chapel implementation is competitive with the MPI and MPI+OpenMP implementations less a small amount of scalar overhead.
- The scalar overhead is larger for the global Chapel version than the EP Chapel version and larger with a smaller degree of intranode parallelism (especially no intranode parallelism) because the implementation is memory-bound at higher levels.
- The *local* flag removes up to 5% of the scalar overhead in the global Chapel version by eliminating the *coforall* loop over the locales and all “potentially remote” references.
- There is an oddity in the results on 4 locales that is consistent across multiple trials. However, in timings for our original submission, we did not see this behaviour.

The global Chapel implementation scales much better than our last year’s entry because we eliminated all communication after the initial remote task invocation. In last year’s entry, the remote tasks in the global Chapel version made remote accesses back to locale 0 after being spawned. At high node counts, this bottleneck became a serious issue.

## 7 Random Access Description

The Random Access benchmark computes pseudo-random updates to a large distributed table  $T$  of 64-bit unsigned integer values. As in STREAM, our distributed memory implementation uses two *Block* distributions—one to distribute the set of  $N_U$  table updates represented using a domain named *Updates*, and a second to distribute the table  $T$  and its corresponding domain.

While the official benchmark permits updates to be batched to amortize the communication overheads, in this entry, we have opted to take a pure update-at-a-time approach for the sake of elegance and to see how far we can push the performance of this implementation.

The core of the Chapel implementation can be summarized by the following three lines of code:

```
forall (_, r) in (Updates, RASstream()) do
  on T(r & indexMask) do
    T(r & indexMask) ^=r;
```

As in STREAM, we use a parallel zippered iteration to express the main computation but rather than traversing arrays, this forall loop iterates over *Updates* and *RASstream()*—an iterator defined elsewhere in the benchmark to generate the pseudo-random stream of values. Each random value is referred to as  $r$  for the purposes of the loop body while the values representing the update indices are neither named nor used (as indicated by the underscore). Since the table location corresponding to  $r$  is increasingly likely to be owned by a remote locale as the number of locales grows, we use an *on*-clause to specify that the update should be computed on the locale that owns the target table element.

To improve performance, we write the *on*-statement as follows:

```
on TableDist.ind2loc(r & indexMask) do
```

This expression says “Access *TableDist*,  $T$ ’s distribution, and call its index-to-locale mapping function to determine which locale owns the index  $r$  & *indexMask*.” With our current compiler, the *on*-statement with the array access will result in extra communication in order to determine where the array element exists in memory (not just on which locale). As the Chapel compiler and the Block distribution improves, we expect to be able to write the code with the array access and get similar performance.

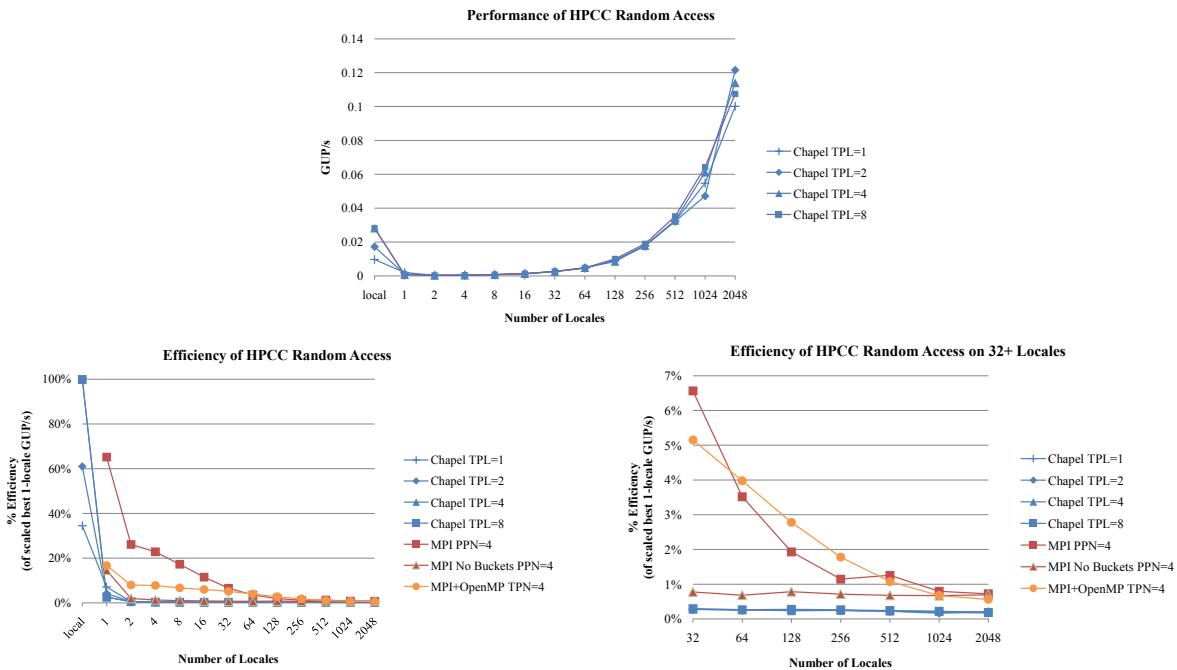
To further improve performance, we introduce a *local*-statement and write the core of the benchmark as follows:

```
forall (_, r) in (Updates, RASstream()) do
  on TableDist.ind2loc(r & indexMask) do {
    const myR = r;
    local {
      T(myR & indexMask) ^= myR;
    }
  }
```

The *local*-block tells the compiler that there will not be any communication within this block. The array access function is cloned as a result. This enables an optimization that allows us to execute the body of the *on*-statement in the GASNet handler. Without this optimization, the compiler has to generate code that will create or reuse a thread on the remote locale. Interestingly, without this optimization, using the *local* keyword does not impact performance because there is a conditional in the array access function that quickly determines the access is local.

Another interesting observation pertains to the declaration of the constant *myR*. This constant is declared on the remote locale so it can be accessed in the *local*-block with impunity. The index  $r$ , on the other hand, exists on the locale that executes the forall loop. Another optimization ensures that we pass the value of  $r$  to the remote locale that executes the *on*-statement, rather than a reference to  $r$  which would then require a remote read. This optimization applies because  $r$  is not changed on the remote locale and because there is no synchronization with another thread before the value is read. That said, the semantics of the language, and thus the *local*-block, assume that  $r$  is not necessarily local and require the copy.

## 8 Random Access Performance



Locales	Chapel				MPI 4 PPN	MPI (No Buckets) 4 PPN	MPI + OpenMP 4 TPN
	1 TPL	2 TPL	4 TPL	8 TPL			
local	0.00972525	0.0172033	0.0281936	0.0281603			
1	0.00200099	0.00116716	0.000714489	0.000741223	0.018384637	0.00410575	0.004713876
2	0.000317063	0.000315868	0.000293979	0.00032569	0.014732936	0.001107038	0.004503648
4	0.000413006	0.000410412	0.000423378	0.000427791	0.025734469	0.001485908	0.008797149
8	0.000706298	0.000710561	0.000710417	0.000742912	0.038924038	0.002301621	0.015105049
16	0.00125442	0.00132351	0.00131812	0.00137499	0.051985597	0.003872435	0.02702586
32	0.00247275	0.00254758	0.00255019	0.00265785	0.059201854	0.007007821	0.046489749
64	0.00456788	0.00463443	0.00483206	0.00476748	0.063466622	0.012337323	0.07174715
128	0.00930662	0.00904224	0.00852665	0.00999743	0.069669644	0.028327548	0.100320905
256	0.0175101	0.0180818	0.017865	0.0189707	0.082709838	0.051622158	0.128392311
512	0.0323605	0.0321982	0.032745	0.0351792	0.181213405	0.098146035	0.155228153
1024	0.0547629	0.0471243	0.0611901	0.0642561	0.229965659	0.193492656	0.194191687
2048	0.100144	0.121527	0.114039	0.10749	0.41723	0.402235	0.325563

The performance graph on top shows the raw performance (GUP/s) of the Chapel version of the Random Access benchmark varying the number of tasks per locale between 1 and 8. The efficiency graphs below show the percentage efficiency of the Chapel version and MPI and MPI+OpenMP reference versions of the Random Access benchmark. For the MPI versions, 4 MPI processes were used per node. For the MPI+Open version, 4 tasks were used per node. The “No Bucket” variation on the MPI reference version shows performance without bucketing. In some ways, this is a purer implementation of random accesses and it also more closely matches the Chapel implementation. The efficiency is taken with respect to linear scaling of the best 1-locale performance (0.028 GUP/s for the Chapel version compiled with the “--local” flag and executed with 4 tasks per node).

We make the following observations:

- On high node counts, the Chapel implementation exhibits about a quarter of the performance of the reference version.
- The reference version with bucketing does not scale because the buckets decrease in size.
- The “--local” flag produces scalar performance that exceeds that of an MPI implementation.

The Chapel implementation performs far better than last year’s entry because of eliminated communication (as described in Section 6) and because of the optimization of on-statements described in Section 7.



## 9 FFT Discussion

Although last year’s HPC entry made use of a Block distribution for STREAM and RA, support for slicing was incomplete. Since our Chapel implementation of FFT makes extensive use of array slicing, last year’s HPC FFT entry only executed on a single locale. With array slicing now supported through some Chapel distributions, this year’s HPC FFT entry can run on multiple locales. The Chapel FFT implementation makes use of a `localSlice()` method that is useful when a slice’s elements are known to be local to the current locale. This permits a user to store an alias to the local portion of a distributed array while only paying local array overheads for its operations. This can be thought of, in Fortran terms, as creating a local dope vector for the alias portion of the array.


The Chapel FFT implementation makes use of the Cyclic distribution as well as the Block distribution, both standard to Chapel version 1.0.2. The Cyclic distribution maps indices to locales in a round-robin fashion starting from a user-specified index, or 0 by default. The Chapel implementation uses a Block-distributed vector for the first half of the computation and then a Cyclic-distributed vector for the second half. When the number of locales is a power of four, all of the butterflies within the radix-4 implementation only access local data. The only communication in the main computation is in the assignment between the two vectors; this requires all-to-all communication.

The following code excerpts the main structure of the Chapel version of FFT:

```
for (str, span) in genDFTStrideSpan(numElements, cyclicPhase) {
  forall (bankStart, twidIndex) in (ADom by 2*span, 0..) {
    var wk2 = W(twidIndex),
        wk1 = W(2*twidIndex),
        wk3 = (wk1.re - 2 * wk2.im * wk1.im,
              2 * wk2.im * wk1.re - wk1.im):elemType;
    forall lo in bankStart..#str do
      on ADom.dist.ind2loc(lo) do
        local butterfly(wk1, wk2, wk3, A.localSlice(lo..by str #radix));
  }
}
```

The outer *for* loop iterates over the serial phases of the algorithm, each of which has a unique stride and distance between its butterflies. The outer *forall* loop is used to create batches of butterflies that share the same twiddle factors while the inner *forall* loop describes that set of butterflies. Nested parallelism is beneficial because the trip counts of the inner and outer loops vary dramatically between the earlier and later phases.

This implementation pushes on Chapel features that have not been optimized for performance. Most egregious, assignment between vectors of different distributions uses a naive and slow implementation. (This will eventually be optimized within the Chapel distributions.) The following graph shows the performance that we are currently getting on 1–16 locales of a Cray XT4:




## 10 HPL Discussion

Although an initial Block-Cyclic distribution (which is a good match for HPL) is part of Chapel v1.0.2, its implementation is incomplete. In particular, the Chapel implementation of HPL relies on rank change and reindexing, but these operations are not yet implemented for any Chapel distribution. Thus there is not yet a distributed-memory implementation of HPL in Chapel.

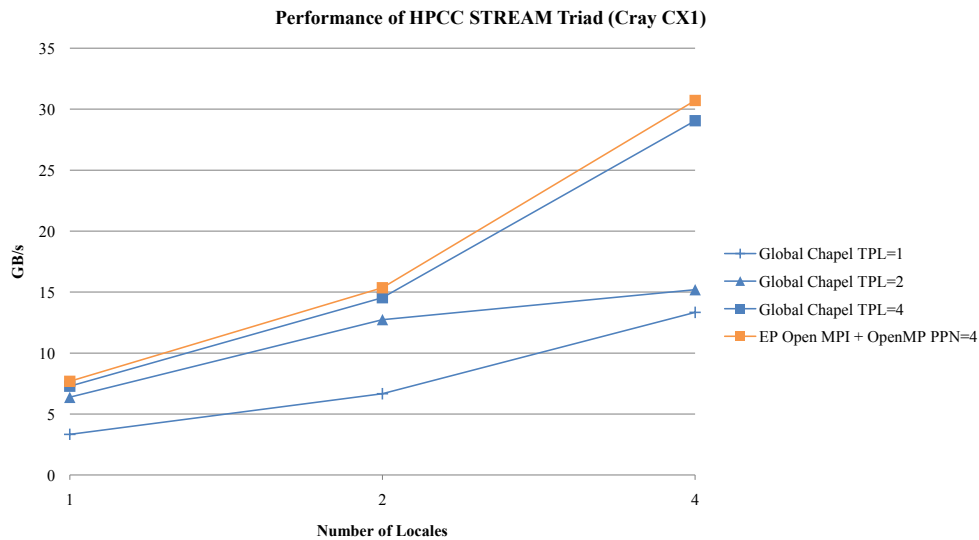
The single-locale implementation presented in this report continues to execute correctly. It is largely unchanged from last year’s entry. However, this year it does take advantage of parallelism within a locale because, as of Chapel v1.0, data parallel constructs like *forall* loops now use multiple tasks within a locale.

# 11 Portability

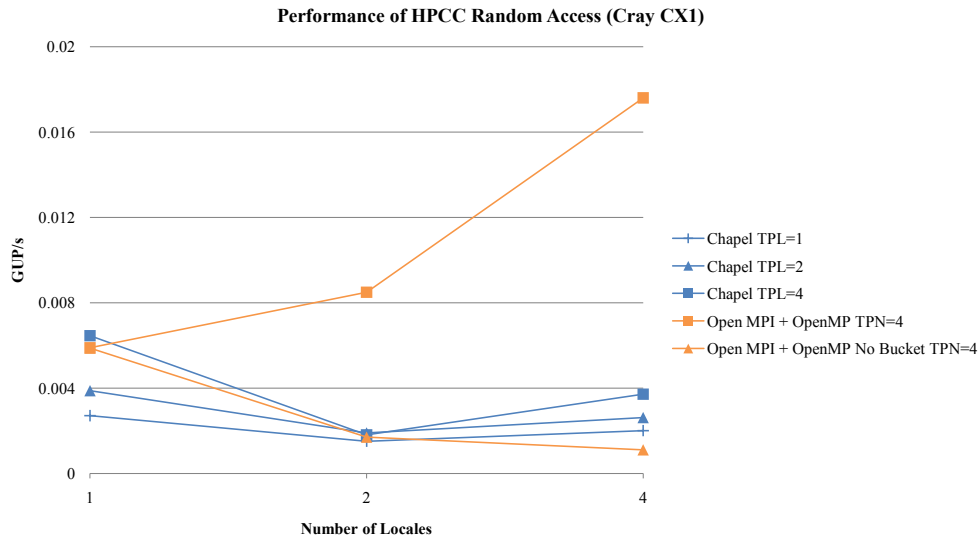
Chapel's source-to-source compilation and use of GASNet and pthreads makes our current implementation of Chapel highly portable. To demonstrate this, we have included some graphs showing performance numbers for the STREAM Triad and Random Access benchmarks on an IBM pSeries 575 and a Cray CX1. These graphs demonstrate portability of the Chapel implementation, but the results, gathered without serious exploration, are not representative of these systems. On the IBM platform, we met our quota of CPU usage before having a change to run the Random Access benchmark.



This performance graph shows raw GB/s performance of the Chapel implementation of STREAM Triad on an IBM pSeries 575 varying the tasks per locale between 1 and 64. The IBM pSeries 575 on which these results were obtained has 104 nodes. Each node has 16 dual core IBM Power6 processors running at 4.7 GHz and either 128 GB or 256 GB of memory. The benchmark was run with a problem size of 132146193 per locale (2.95 GB/locale) and the locales were mapped to nodes.



This performance graph shows raw GB/s performance of the Chapel implementation of STREAM Triad on a Cray CX1 varying the tasks per locale between 1 and 4. The Cray CX1 on which these results were obtained has 8 nodes and an Infiniband network. Each node contains 4 dual core Intel Xeon processors running at 3 GHz and 16 GB of memory. Chapel locales were mapped to nodes on this machine.



This performance graph shows raw GUP/s performance of the Chapel implementation of Random Access on a Cray CX1 (see characteristics above) varying the tasks per locale between 1 and 4.

## 12 Summary

We hope this report provides a glimpse into the Chapel implementation, our progress and the language design. The performance of the Chapel implementation is improving steadily. Both STREAM Triad and Random Access are much closer to optimal than a year ago. The FFT benchmark is now running on multiple locales on small problem sizes, and we've made significant progress on the multi-locale implementation of HPL.

## Acknowledgments

The authors would like to gratefully acknowledge our former team members and collaborators who assisted with our 2006 and 2008 entries in the HPCC competition: Samuel Figueroa, Mary Beth Hribar, John Lewis, Andrew Stone, Adrian Tate, and Wayne Wong. In addition, we thank all of Chapel's past contributors and early users for helping us reach this stage.

We would like to thank Oak Ridge National Laboratory and the National Center for Computational Sciences for the computing time on the Cray XT4 partition of Jaguar. We would also like to thank SARA for their generous donation of time and user support on the Huygens IBM pSeries 575 clustered SMP system, and by extension the Netherlands National Computing Facilities foundation (NCF) who funds Huygens.

## A Global STREAM Triad Code

```
1 //
2 // Use standard modules for Block distributions, Timing routines, Type
3 // utility functions, and Random numbers
4 //
5 use BlockDist, Time, Types, Random;
6
7 //
8 // Use shared user module for computing HPC problem sizes
9 //
10 use HPCProblemSize;
11
12 //
13 // The number of vectors and element type of those vectors
14 //
15 const numVectors = 3;
16 type elemType = real(64);
17
18 //
19 // Configuration constants to set the problem size (m) and the scalar
20 // multiplier, alpha
21 //
22 config const m = computeProblemSize(numVectors, elemType),
23             alpha = 3.0;
24
25 //
26 // Configuration constants to set the number of trials to run and the
27 // amount of error to permit in the verification
28 //
29 config const numTrials = 10,
30             epsilon = 0.0;
31
32 //
33 // The number of tasks to use per Chapel locale
34 //
35 config const tasksPerLocale = here.numCores;
36
37 //
38 // Configuration constants to indicate whether or not to use a
39 // pseudo-random seed (based on the clock) or a fixed seed; and to
40 // specify the fixed seed explicitly
41 //
42 config const useRandomSeed = true,
43             seed = if useRandomSeed then SeedGenerator.clockMS else 314159265;
44
45 //
46 // Configuration constants to control what's printed -- benchmark
47 // parameters, input and output arrays, and/or statistics
48 //
49 config const printParams = true,
50             printArrays = false,
51             printStats = true;
52
53 //
54 // The program entry point
55 //
56 def main() {
57     printConfiguration(); // print the problem size, number of trials, etc.
58
59     //
60     // BlockDist is a 1D block distribution that is computed by blocking
61     // the bounding box 1..m across the set of locales
62     //
63     const BlockDist = distributionValue(
64         new Block(rank=1, bbox=[1..m], tasksPerLocale=tasksPerLocale));
65
66     //
67     // ProblemSpace describes the index set for the three vectors. It
68     // is a 1D domain storing 64-bit ints and is distributed according
69     // to BlockDist. It contains the indices 1..m.
70     //
71     const ProblemSpace: domain(1, int(64)) distributed BlockDist = [1..m];
72
73     //
74     // A, B, and C are the three distributed vectors, declared to store
75     // a variable of type elemType for each index in ProblemSpace.
76     //
77     var A, B, C: [ProblemSpace] elemType;
78
79     initVectors(B, C); // Initialize the input vectors, B and C
80
81     var execTime: [1..numTrials] real; // an array of timings
82
83     for trial in 1..numTrials {
84         const startTime = getCurrentTime(); // loop over the trials
85         // capture the start time
86         //
87         // The main loop: Iterate over the vectors A, B, and C in a
88         // parallel, zippered manner storing the elements as a, b, and c.
89         // Compute the multiply-add on b and c, storing the result to a.
90         //
91         forall (a, b, c) in (A, B, C) do
92             a = b + alpha * c;
93
94         execTime(trial) = getCurrentTime() - startTime; // store the elapsed time
95     }
96
97     const validAnswer = verifyResults(A, B, C); // verify...
98     printResults(validAnswer, execTime); // ...and print the results
99 }
100
101 //
102 // Print the problem size and number of trials
103 //
104 def printConfiguration() {
105     if (printParams) {
106         if (printStats) then printLocalesTasks(tasksPerLocale);
107         printProblemSize(elemType, numVectors, m);
108         writeln("Number of trials = ", numTrials, "\n");
109     }
110 }
111
112 //
113 // Initialize vectors B and C using a random stream of values and
114 // optionally print them to the console
115 //
116 def initVectors(B, C) {
117     var randlist = new RandomStream(seed);
118
119     randlist.fillRandom(B);
120     randlist.fillRandom(C);
121
122     if (printArrays) {
123         writeln("B is: ", B, "\n");
124         writeln("C is: ", C, "\n");
125     }
126
127     delete randlist;
128 }
129
130 //
131 // Verify that the computation is correct
132 //
133 def verifyResults(A, B, C) {
134     if (printArrays) then writeln("A is: ", A, "\n"); // optionally print A
135
136     //
137     // recompute the computation, destructively storing into B to save space
138     //
139     forall (b, c) in (B, C) do
140         b += alpha * c;
141
142     if (printArrays) then writeln("A-hat is: ", B, "\n"); // and A-hat too
143
144     //
145     // Compute the infinity-norm by computing the maximum reduction of the
146     // absolute value of A's elements minus the new result computed in B.
147     // "[i in I]" represents an expression-level loop: "forall i in I"
148     //
149     const infNorm = max reduce [(a,b) in (A,B)] abs(a - b);
150
151     return (infNorm <= epsilon); // return whether the error is acceptable
152 }
153
154 //
155 // Print out success/failure, the timings, and the GB/s value
156 //
157 def printResults(successful, execTimes) {
158     writeln("Validation: ", if successful then "SUCCESS" else "FAILURE");
159     if (printStats) {
160         const totalTime = + reduce execTimes,
161             avgTime = totalTime / numTrials,
162             minTime = min reduce execTimes;
163         writeln("Execution time:");
164         writeln(" tot = ", totalTime);
165         writeln(" avg = ", avgTime);
166         writeln(" min = ", minTime);
167
168         const GBPerSec = numVectors * numBytes(elemType) * (m / minTime) * 1e-9;
169         writeln("Performance (GB/s) = ", GBPerSec);
170     }
171 }
```

## B EP STREAM Triad Code

```
1 //
2 // Embarassingly Parallel Implementation of STREAM Triad
3 //
4 // This version of the stream benchmark is not as elegant as
5 // stream.chpl. It is a per-locale code with no communication in the
6 // actual computation. It highlights the ability of a Chapel
7 // programmer to escape the global-view programming model and write
8 // codes with a fragmented, per-locale model.
9 //
10 // Comments marked with '***' point out differences with the global
11 // version of this benchmark in stream.chpl.
12 //
13 //
14 //
15 // *** Since this benchmark is written on a per-locale basis, there is no
16 // *** need to use the Block distribution. The following use omits
17 // *** BlockDist, included in stream.chpl.
18 //
19 // Use standard modules for Timing routines, Type utility functions,
20 // and Random numbers
21 //
22 use Time, Types, Random;
```

```

24 //
25 // Use shared user module for computing HPC problem sizes
26 //
27 use HPCProblemSize;

29 //
30 // The number of vectors and element type of those vectors
31 //
32 const numVectors = 3;
33 type elemType = real(64);

35 //
36 // *** To ensure a local problem size, we spoof the number of vectors
37 // *** passed to the computeProblemSize function to be the number of
38 // *** vectors times the number of locales.
39 //
40 // Configuration constants to set the problem size (m) and the scalar
41 // multiplier, alpha
42 //
43 config const m = computeProblemSize(numVectors*numLocales, elemType),
44             alpha = 3.0;

46 //
47 // Configuration constants to set the number of trials to run and the
48 // amount of error to permit in the verification
49 //
50 config const numTrials = 10,
51             epsilon = 0.0;

53 //
54 // *** There isn't (yet) a way to set the number of tasks to use for
55 // *** implementing a forall loop over a default array. When there is
56 // *** such a way, we will want to set it via this configuration
57 // *** constant to get functionality like we have with stream.chpl.
58 //
59 // config const tasksPerLocale = here.numCores;

61 //
62 // Configuration constants to indicate whether or not to use a
63 // pseudo-random seed (based on the clock) or a fixed seed; and to
64 // specify the fixed seed explicitly
65 //
66 config const useRandomSeed = true,
67             seed = if useRandomSeed then SeedGenerator.clockMS else 314159265;

69 //
70 // *** To ensure determinism of output, there is no more printing of the
71 // *** arrays in initVectors and verifyResults.
72 //
73 //
74 // Configuration constants to control what's printed -- benchmark
75 // parameters and/or statistics
76 //
77 config const printParams = true,
78             printStats = true;

80 //
81 // The program entry point
82 //
83 def main() {
84   printConfiguration(); // print the problem size, number of trials, etc.

86 //
87 // *** Aggregates for collecting per-locale results for the minimum
88 // *** execution time per trial, and whether verification passed
89 //
90 var minTimes: [LocaleSpace] real;
91 var validAnswers: [LocaleSpace] bool;

93 //
94 // *** Fragment control so that we have a single task running on
95 // *** every locale.
96 //
97 coforall loc in Locales do on loc {

99 //
100 // *** We declare these variables outside of the local block since
101 // *** we'll need to access them when we write back to the global
102 // *** aggregates declared above.
103 //
104 var validAnswer: bool;
105 var execTime: [1..numTrials] real;

107 //
108 // *** Indicates that all of the code in this block is local to
109 // *** this locale. There is no communication. A violation will
110 // *** result in an error, though error checking is disabled with
111 // *** --fast or --no-checks.
112 //
113 local {

115 //
116 // *** A, B, and C are the three local vectors
117 //
118 var A, B, C: [1..m] elemType;

120   initVectors(B, C); // Initialize the input vectors, B and C

122   for trial in 1..numTrials { // loop over the trials
123     const startTime = getCurrentTime();

125 //
126 // *** The main loop looks identical to stream.chpl. However,
127 // *** in this version we are iterating over arrays that are
128 // *** not distributed.
129 //
130     forall (a, b, c) in (A, B, C) do

```

```

131       a = b + alpha * c;

133       execTime(trial) = getCurrentTime() - startTime;
134     }

136     validAnswer = verifyResults(A, B, C); // verify...
137   }

139 //
140 // *** Write times and verification result into aggregates
141 // *** declared above. These are declared over LocaleSpace so we
142 // *** can write to them in parallel.
143 //
144 minTimes[here.id] = min reduce execTime;
145 validAnswers[here.id] = validAnswer;
146 }

148 //
149 // *** Pass minimum, average, and maximum times to printResults
150 //
151 printResults(&& reduce validAnswers, minTimes);
152 }

154 //
155 // Print the problem size and number of trials
156 //
157 def printConfiguration() {
158   if (printParams) {
159     //
160     // *** Here we multiply m by the number of locales so that we can
161     // *** print out the global problem size.
162     //
163     printProblemSize(elemType, numVectors, m * numLocales);
164     writeln("Number of trials = ", numTrials, "\n");
165   }
166 }

168 //
169 // *** Both initVectors and verifyResults are almost identical to
170 // *** stream.chpl even though they are called with arrays that are
171 // *** not distributed. For initialization, the same random stream is
172 // *** used on each locale. In the global version, a single logical
173 // *** stream of random numbers is used across all of the locales.
174 //
175 // *** In this version, we've omitted a way to print the arrays. This
176 // *** ensures determinism of output. Printing the arrays also
177 // *** violates the locality constraint imposed by the local block
178 // *** from which these functions are called.
179 //
180 // Initialize vectors B and C using a random stream of values
181 //
182 def initVectors(B, C) {
183   var randlist = new RandomStream(seed);

185   randlist.fillRandom(B);
186   randlist.fillRandom(C);

188   delete randlist;
189 }

191 //
192 // Verify that the computation is correct
193 //
194 def verifyResults(A, B, C) {

196 //
197 // recompute the computation, destructively storing into B to save space
198 //
199 forall (b, c) in (B, C) do
200   b += alpha * c;

202 //
203 // Compute the infinity-norm by computing the maximum reduction of the
204 // absolute value of A's elements minus the new result computed in B.
205 // "[i in I]" represents an expression-level loop: "forall i in I"
206 //
207 const infNorm = max reduce [(a,b) in (A,B)] abs(a - b);

209   return (infNorm <= epsilon); // return whether the error is acceptable
210 }

212 //
213 // Print out success/failure, the timings, and the GB/s value.
214 //
215 // *** Here we report maximum, average, and minimum times instead of
216 // *** total, average, and minimum.
217 //
218 def printResults(successful, minTimes) {
219   writeln("Validation: ", if successful then "SUCCESS" else "FAILURE");
220   if (printStats) {
221     const maxTime = max reduce minTimes,
222           avgTime = + reduce minTimes / numLocales,
223           minTime = min reduce minTimes;
224     writeln("Execution time:");
225     writeln(" max = ", maxTime);
226     writeln(" avg = ", avgTime);
227     writeln(" min = ", minTime);

229     const maxGBPerSec = numVectors * numBytes(elemType) * (m / minTime) * 1e-9,
230           avgGBPerSec = numVectors * numBytes(elemType) * (m / avgTime) * 1e-9,
231           minGBPerSec = numVectors * numBytes(elemType) * (m / maxTime) * 1e-9;
232     writeln("Performance (GB/s):");
233     writeln(" max = ", maxGBPerSec);
234     writeln(" avg = ", avgGBPerSec);
235     writeln(" min = ", minGBPerSec);
236   }
237 }

```

# C Global Random Access Code

## C.1 Benchmark Code

```
1 //
2 // Use standard modules for Block distributions and Timing routines
3 //
4 use BlockDist, Time;

6 //
7 // Use the user modules for computing HPC problem sizes and for
8 // defining RA's random stream of values
9 //
10 use HPCProblemSize, RARandomStream;

12 //
13 // The number of tables as well as the element and index types of
14 // that table
15 //
16 const numTables = 1;
17 type elemType = randType,
18     indexType = randType;

20 //
21 // Configuration constants defining log2(problem size) -- n -- and
22 // the number of updates -- N_U
23 //
24 config const n = computeProblemSize(numTables, elemType,
25     returnLog2=true, retType=indexType),
26     N_U = 2*(n+2);

28 //
29 // Constants defining the problem size (m) and a bit mask for table
30 // indexing
31 //
32 const m = 2**n,
33     indexMask = m-1;

35 //
36 // Configuration constant defining the number of errors to allow (as a
37 // fraction of the number of updates, N_U)
38 //
39 config const errorTolerance = 1e-2;

41 //
42 // The number of tasks to use per Chapel locale
43 //
44 config const tasksPerLocale = here.numCores;

46 //
47 // Configuration constants to control what's printed -- benchmark
48 // parameters, input and output arrays, and/or statistics
49 //
50 config const printParams = true,
51     printArrays = false,
52     printStats = true;

54 //
55 // TableDist is a 1D block distribution for domains storing indices
56 // of type "indexType", and it is computed by blocking the bounding
57 // box 0..m-1 across the set of locales. UpdateDist is a similar
58 // distribution that is computed by blocking the indices 0..N_U-1
59 // across the locales.
60 //
61 const TableDist = distributionValue(new Block(1, indexType, bbox=[0..m-1],
62     tasksPerLocale=tasksPerLocale)),
63     UpdateDist = distributionValue(new Block(1, indexType, bbox=[0..N_U-1],
64     tasksPerLocale=tasksPerLocale));

66 //
67 // TableSpace describes the index set for the table. It is a 1D
68 // domain storing indices of type indexType, it is distributed
69 // according to TableDist, and it contains the indices 0..m-1.
70 // Updates is an index set describing the set of updates to be made.
71 // It is distributed according to UpdateDist and contains the
72 // indices 0..N_U-1.
73 //
74 const TableSpace: domain(1, indexType) distributed TableDist = [0..m-1],
75     Updates: domain(1, indexType) distributed UpdateDist = [0..N_U-1];

77 //
78 // T is the distributed table itself, storing a variable of type
79 // elemType for each index in TableSpace.
80 //
81 var T: [TableSpace] elemType;

83 //
84 // The program entry point
85 //
86 def main() {
87     printConfiguration(); // print the problem size, number of trials, etc.
88 }

90 // In parallel, initialize the table such that each position
91 // contains its index. "[i in TableSpace]" is shorthand for "forall
92 // i in TableSpace"
93 //
94 [i in TableSpace] T(i) = i;

96 const startTime = getCurrentTime(); // capture the start time

98 //
99 // The main computation: Iterate over the set of updates and the
100 // stream of random values in a parallel, zippered manner, dropping
101 // the update index on the ground ("_") and storing the random value
102 // in r. Use an on- clause to force the table update to be executed on
103 // the locale which owns the table element in question to minimize
104 // communications. Compute the update using r both to compute the
105 // index and as the update value.
106 //
107 forall (_, r) in (Updates, RStream()) do
108     on TableDist.ind2loc(r & indexMask) do {
109         const myR = r;
110         local {
111             T(myR & indexMask) ^= myR;
112         }
113     }

115 const execTime = getCurrentTime() - startTime; // capture the elapsed time

117 const validAnswer = verifyResults(); // verify the updates
118 printResults(validAnswer, execTime); // print the results
119 }

121 //
122 // Print the problem size and number of updates
123 //
124 def printConfiguration() {
125     if (printParams) {
126         if (printStats) then printLocalesTasks(tasksPerLocale);
127         printProblemSize(elemType, numTables, m);
128         writeln("Number of updates = ", N_U, "\n");
129     }
130 }

132 //
133 // Verify that the computation is correct
134 //
135 def verifyResults() {
136     //
137     // Print the table, if requested
138     //
139     if (printArrays) then writeln("After updates, T is: ", T, "\n");

141     //
142     // Reverse the updates by recomputing them, this time using an
143     // atomic statement to ensure no conflicting updates
144     //
145     forall (_, r) in (Updates, RStream()) do
146         on TableDist.ind2loc(r & indexMask) do
147             atomic T(r & indexMask) ^= r;

149     //
150     // Print the table again after the updates have been reversed
151     //
152     if (printArrays) then writeln("After verification, T is: ", T, "\n");

154     //
155     // Compute the number of table positions that weren't reversed
156     // correctly. This is an indication of the number of conflicting
157     // updates.
158     //
159     const numErrors = + reduce [i in TableSpace] (T(i) != i);
160     if (printStats) then writeln("Number of errors is: ", numErrors, "\n");

162     //
163     // Return whether or not the number of errors was within the benchmark's
164     // tolerance.
165     //
166     return numErrors <= (errorTolerance * N_U);
167 }

169 //
170 // Print out success/failure, the execution time, and the GUPS value
171 //
172 def printResults(successful, execTime) {
173     writeln("Validation: ", if successful then "SUCCESS" else "FAILURE");
174     if (printStats) {
175         writeln("Execution time = ", execTime);
176         writeln("Performance (GUPS) = ", (N_U / execTime) * 1e-9);
177     }
178 }
```

## C.2 Supporting Module Code

```
1 //
2 // A helper module for the RA benchmark that defines the random stream
3 // of values
4 //
5 module RARandomStream {
6     param randWidth = 64; // the bit-width of the random numbers

7     type randType = uint(randWidth); // the type of the random numbers

9     //
10    // bitDom is a non-distributed domain whose indices correspond to
11    // the bit positions in the random values. m2 is a table of helper
12    // values used to fast-forward through the random stream.
```

```

13 //
14 const m2: randWidth*randType = computeM2Vals();
15
16 //
17 // A serial iterator for the random stream that resets the stream
18 // to its 0th element and yields values endlessly.
19 //
20 def RAStream() {
21   var val = getNextRandom(0);
22   while (1) {
23     getNextRandom(val);
24     yield val;
25   }
26 }
27
28 //
29 // A "follower" iterator for the random stream that takes a range of
30 // 0-based indices (follower) and yields the pseudo-random values
31 // corresponding to those indices. Follower iterators like these
32 // are required for parallel zippered iteration.
33 //
34 def RAStream(param tag: iterator, follower) where tag == iterator.follower {
35   if follower.size != 1 then
36     halt("RAStream cannot use multi-dimensional iterator");
37   var val = getNextRandom(follower(1).low);
38   for follower {
39     getNextRandom(val);
40     yield val;
41   }
42 }
43
44 //
45 // A helper function for "fast-forwarding" the random stream to
46 // position n in O(log2(n)) time
47 //
48 def getNextRandom(in n: uint(64)) {
49   param period = 0x7fffffff7fffffff;
50
51   n %= period;
52   if (n == 0) then return 0x1;
53   var ran: randType = 0x2;
54   for i in 0..log2(n)-1 by -1 {
55     var val: randType = 0;
56     for j in 0..randWidth do
57       if ((ran >> j) & 1) then val ^= m2(j+1);
58       ran = val;
59     if ((n >> i) & 1) then getNextRandom(ran);
60   }
61   return ran;
62 }
63
64 //
65 // A helper function for advancing a value from the random stream,
66 // x, to the next value
67 //
68 def getNextRandom(inout x) {
69   param POLY = 0x7;
70   param hiRandBit = 0x1:randType << (randWidth-1);
71   x = (x << 1) ^ (if (x & hiRandBit) then POLY else 0);
72 }
73
74 //
75 // A helper function for computing the values of the helper tuple, m2
76 //
77 //
78 def computeM2Vals() {
79   var m2tmp: randWidth*randType;
80   var nextVal = 0x1: randType;
81   for param i in 1..randWidth {
82     m2tmp[i] = nextVal;
83     getNextRandom(nextVal);
84     getNextRandom(nextVal);
85   }
86   return m2tmp;
87 }
88 }

```

## D Global FFT Code

```

1 /* This implementation of the FFT benchmark uses radix-4 butterflies
2 and is divided into two main phases: one which uses a Block
3 distribution and the second which uses a Cyclic distribution. When
4 run on 4**k locales, this guarantees that each butterfly will only
5 access local data. In an optimized implementation, this should
6 cause most of the communication to occur when copying the vector
7 between Block and Cyclic storage formats.
8 */
9
10 //
11 // Use standard modules for Bit operations, Random numbers, Timing, and
12 // Block and Cyclic distributions
13 //
14 use BitOps, Random, Time, BlockDist, CyclicDist;
15
16 //
17 // Use shared user module for computing HPC problem sizes
18 //
19 use HPCProblemSize;
20
21 const radix = 4; // the radix of this FFT implementation
22
23 const numVectors = 2; // the number of vectors to be stored
24 type elemType = complex(128); // the element type of the vectors
25 type idxType = int(64); // the index type of the vectors
26
27 //
28 // A configuration constant defining log2(problem size) -- n -- and a
29 // constant defining the problem size itself -- m
30 //
31 config const n = computeProblemSize(numVectors, elemType, returnLog2 = true);
32 const m = 2**n;
33
34 //
35 // The number of tasks to use per Chapel locale in parallel loops
36 //
37 config const tasksPerLocale = here.numCores;
38
39 //
40 // Configuration constants defining the epsilon and threshold values
41 // used to verify the result
42 //
43 config const epsilon = 2.0 ** -51.0,
44 threshold = 16.0;
45
46 //
47 // Configuration constants to indicate whether or not to use a
48 // pseudo-random seed (based on the clock) or a fixed seed; and to
49 // specify the fixed seed explicitly
50 //
51 config const useRandomSeed = true,
52 seed = if useRandomSeed then SeedGenerator.clockMS else 314159265;
53
54 //
55 // Configuration constants to control what's printed -- benchmark
56 // parameters, input and output arrays, and/or statistics
57 //
58 config const printParams = true,
59 printArrays = false,
60 printStats = true;
61
62 //
63 // The program entry point
64 //
65 def main() {
66   printConfiguration(); // print the problem size
67
68   //
69   // This implementation assumes 4**k locales due to its assertion that
70   // all butterflies are local to a given locale
71   //
72   assert(4**log4(numLocales) == numLocales,
73         "numLocales must be a power of 4 for this fft implementation");
74
75   //
76   // TwiddleDom describes the index set used to define the vector of
77   // twiddle values and is a 1D domain indexed by 64-bit ints from 0
78   // to m/4-1 stored using the block distribution TwiddleDist.
79   // Twiddles is the vector of twiddle values.
80   //
81   const TwiddleDist = distributionValue(new Block(1, idxType, bbox=[0..m/4-1]));
82   const TwiddleDom: domain(1, idxType) distributed TwiddleDist = [0..m/4-1];
83   var Twiddles: [TwiddleDom] elemType;
84
85   //
86   // ProblemSpace describes the abstract problem space used for the
87   // FFT benchmark: the indices 0..m-1
88   //
89   const ProblemSpace = [0..m-1];
90
91   //
92   // BlkDist describes the problem space as distributed in a Block
93   // manner between the Locales where ProblemSpace defines the
94   // bounding box used to compute the blocking. BlkDom defines the
95   // Block-distributed problem space and is used to define the vectors
96   // z (used to store the input vector) and Zblk (used for the first
97   // half of the FFT phases).
98   //
99   const BlkDist = distributionValue(new Block(1, idxType, bbox=ProblemSpace,
100     tasksPerLocale=tasksPerLocale));
101   const BlkDom: domain(1, idxType) distributed BlkDist = ProblemSpace;
102   var Zblk, z: [BlkDom] elemType;
103
104   //
105   // CycDist describes the problem space as distributed in a Cyclic
106   // manner between the Locales where locale #0 stores element 0.
107   // CycDom defines the Cyclic-distributed problem space and is used
108   // to define the Zcyc vector, used for the second half of the FFT
109   // phases.
110   //
111   const CycDist = distributionValue(new Cyclic(1, idxType,
112     tasksPerLocale=tasksPerLocale));
113   const CycDom: domain(1, idxType) distributed CycDist = ProblemSpace;
114   var Zcyc: [CycDom] elemType;
115
116   initVectors(Twiddles, z); // initialize twiddles and input vector z
117
118   const startTime = getCurrentTime(); // capture the start time
119
120   [(a,b) in (Zblk, z)] a = conjg(b); // store the conjugate of z in Zblk
121   bitReverseShuffle(Zblk); // permute Zblk
122
123   dfft(Zblk, Twiddles, cyclicPhase=false); // compute the DFFT, block phases
124
125   forall (b, c) in (Zblk, Zcyc) do // copy vector to Cyclic storage
126     c = b;
127
128   dfft(Zcyc, Twiddles, cyclicPhase=true); // compute the DFFT, cyclic phases

```

```

130 forall (b, c) in (Zblk, Zcyc) do // copy vector back to Block storage
131     b = c;
132
133 const execTime = getCurrentTime() - startTime; // store the elapsed time
134
135 const validAnswer = verifyResults(z, Zblk, Zcyc, Twiddles); // validate answer
136 printResults(validAnswer, execTime); // print the results
137 }
138
139 //
140 // compute the discrete fast Fourier transform of a vector A declared
141 // over domain ADom using twiddle vector W
142 //
143 def dfft(A: [?ADom], W, cyclicPhase) {
144     const numElements = A.numElements;
145     //
146     // loop over the phases of the DFT sequentially using custom
147     // iterator genDFTStrideSpan that yields the stride and span for
148     // each bank of butterfly calculations
149     //
150     for (str, span) in genDFTStrideSpan(numElements, cyclicPhase) {
151         //
152         // loop in parallel over each of the banks of butterflies with
153         // shared twiddle factors, zippering with the unbounded range
154         // 0.. to get the base twiddle indices
155         //
156         forall (bankStart, twidIndex) in (ADom by 2*span, 0..) {
157             //
158             // compute the first set of multipliers for the low bank
159             //
160             var wk2 = W(twidIndex),
161                 wk1 = W(2*twidIndex),
162                 wk3 = (wk1.re - 2 * wk2.im * wk1.im,
163                     2 * wk2.im * wk1.re - wk1.im):elemType;
164             //
165             // loop in parallel over the low bank, computing butterflies
166             // Note: lo..#num == lo, lo+1, lo+2, ..., lo+num-1
167             // lo.. by str #num == lo, lo+str, lo+2*str, ... lo+(num-1)*str
168             //
169             forall lo in bankStart..#str do
170                 on ADom.dist.ind2loc(lo) do
171                     local butterfly(wk1, wk2, wk3, A.localSlice(lo..by str #radix));
172
173             //
174             // update the multipliers for the high bank
175             //
176             wk1 = W(2*twidIndex+1);
177             wk3 = (wk1.re - 2 * wk2.re * wk1.im,
178                 2 * wk2.re * wk1.re - wk1.im):elemType;
179             wk2 *= 1.0i;
180
181             //
182             // loop in parallel over the high bank, computing butterflies
183             //
184             forall lo in bankStart+span..#str do
185                 on ADom.dist.ind2loc(lo) do
186                     local butterfly(wk1, wk2, wk3, A.localSlice(lo.. by str #radix));
187             }
188         }
189
190 if cyclicPhase {
191     //
192     // Do the last set of butterflies...
193     //
194     const str = radix**log4(numElements-1);
195     //
196     // ...using the radix-4 butterflies with 1.0 multipliers if the
197     // problem size is a power of 4
198     //
199     if (str*radix == numElements) {
200         forall lo in 0..#str do
201             on ADom.dist.ind2loc(lo) do
202                 local butterfly(1.0, 1.0, 1.0, A.localSlice(lo.. by str #radix));
203         }
204         //
205         // ...otherwise using a simple radix-2 butterfly scheme
206         //
207         else
208             forall lo in 0..#str do
209                 on ADom.dist.ind2loc(lo) do
210                     local {
211                         const a = A(lo),
212                             b = A(lo+str);
213                         A(lo) = a + b;
214                         A(lo+str) = a - b;
215                     }
216             }
217     }
218
219 //
220 // this is the radix-4 butterfly routine that takes multipliers wk1,
221 // wk2, and wk3 and a 4-element array (slice) A.
222 //
223 def butterfly(wk1, wk2, wk3, X:[0..3]) {
224     var x0 = X(0) + X(1),
225         x1 = X(0) - X(1),
226         x2 = X(2) + X(3),
227         x3rot = (X(2) - X(3))*1.0i;
228
229     X(0) = x0 + x2; // compute the butterfly in-place on X
230     x0 -= x2;
231     X(2) = wk2 * x0;
232     x0 = x1 + x3rot;
233     X(1) = wk1 * x0;
234     x0 = x1 - x3rot;
235     X(3) = wk3 * x0;
236 }
237
238 //
239 // this iterator generates the stride and span values for the phases
240 // of the DFFT simply by yielding tuples: (radix**i, radix**(i+1))
241 //
242 def genDFTStrideSpan(numElements, cyclicPhase) {
243     const (start, end) = if !cyclicPhase then (1, numLocales:idxType)
244         else (numLocales, numElements-1);
245     var stride = start;
246     for i in log4(start)+1..log4(end):int {
247         const span = stride * radix;
248         yield (stride, span);
249         stride = span;
250     }
251 }
252
253 //
254 // Print the problem size
255 //
256 def printConfiguration() {
257     if (printParams) {
258         if (printStats) then printLocalesTasks(tasksPerLocale=tasksPerLocale);
259         printProblemSize(elemType, numVectors, m);
260     }
261 }
262
263 //
264 // Initialize the twiddle vector and random input vector and
265 // optionally print them to the console
266 //
267 def initVectors(Twiddles, z) {
268     computeTwiddles(Twiddles);
269     bitReverseShuffle(Twiddles);
270
271     fillRandom(z, seed);
272
273     if (printArrays) {
274         writeln("After initialization, Twiddles is: ", Twiddles, "\n");
275         writeln("z is: ", z, "\n");
276     }
277 }
278
279 //
280 // Compute the twiddle vector values
281 //
282 def computeTwiddles(Twiddles) {
283     const numTwids = Twiddles.numElements,
284         delta = 2.0 * atan(1.0) / numTwids;
285
286     Twiddles(0) = 1.0;
287     Twiddles(numTwids/2) = let x = cos(delta * numTwids/2)
288         in (x, x):elemType;
289     forall i in 1..numTwids/2-1 {
290         const x = cos(delta*i),
291             y = sin(delta*i);
292         Twiddles(i) = (x, y):elemType;
293         Twiddles(numTwids - i) = (y, x):elemType;
294     }
295 }
296
297 //
298 // Perform a permutation of the argument vector by reversing the bits
299 // of the indices
300 //
301 def bitReverseShuffle(Vect: [?Dom]) {
302     const numBits = log2(Vect.numElements),
303         Perm: [i in Dom] Vect.eltType = Vect(bitReverse(i, revBits=numBits));
304     Vect = Perm;
305 }
306
307 //
308 // Reverse the low revBits bits of val
309 //
310 def bitReverse(val: ?valType, revBits = 64) {
311     param mask = 0x0102040810204080;
312     const valReverse64 = bitMatMultOr(mask, bitMatMultOr(val:uint(64), mask)),
313         valReverse = bitRotLeft(valReverse64, revBits);
314     return valReverse: valType;
315 }
316
317 //
318 // Compute the log base 4 of x
319 //
320 def log4(x) return logBasePow2(x, 2);
321
322 //
323 // verify that the results are correct by reapplying the dfft and then
324 // calculating the maximum error, comparing against epsilon
325 //
326 def verifyResults(z, Zblk, Zcyc, Twiddles) {
327     if (printArrays) then writeln("After FFT, Z is: ", Zblk, "\n");
328
329     [z in Zblk] z = conjg(z) / m;
330     bitReverseShuffle(Zblk);
331     dfft(Zblk, Twiddles, cyclicPhase=false);
332     forall (b, c) in (Zblk, Zcyc) do
333         c = b;
334         dfft(Zcyc, Twiddles, true);
335         forall (b, c) in (Zblk, Zcyc) do
336             b = c;
337         }
338
339     if (printArrays) then writeln("After inverse FFT, Z is: ", Zblk, "\n");
340
341     var maxerr = max reduce sqrt((z.re - Zblk.re)**2 + (z.im - Zblk.im)**2);
342     maxerr /= (epsilon * n);
343     if (printStats) then writeln("error = ", maxerr);
344
345     return (maxerr < threshold);
346 }
347
348 //
349 // print out success/failure, the timing, and the GFlop/s value
350 //
351 def printResults(successful, execTime) {
352     writeln("Validation: ", if successful then "SUCCESS" else "FAILURE");

```



```

353 if (printStats) {
354     writeln("Execution time = ", execTime);

```

```

355     writeln("Performance (Gflop/s) = ", 5 * (m * n / execTime) * 1e-9);
356 }
357 }

```

## E Global HPL Code

```

1 //
2 // Use standard modules for vector and matrix Norms, Random numbers
3 // and Timing routines
4 //
5 use Norm, Random, Time;

7 //
8 // Use the user module for computing HPCC problem sizes
9 //
10 use HPCCProblemSize;

12 //
13 // The number of matrices and the element type of those matrices
14 //
15 const numMatrices = 1;
16 type indexType = int,
17     elemType = real;

19 //
20 // Configuration constants indicating the problem size (n) and the
21 // block size (blkSize)
22 //
23 config const n = computeProblemSize(numMatrices, elemType, rank=2,
24                                     memFraction=2, retType=indexType),
25     blkSize = 5;

27 //
28 // Configuration constant used for verification thresholds
29 //
30 config const epsilon = 2.0e-15;

32 //
33 // Configuration constants to indicate whether or not to use a
34 // pseudo-random seed (based on the clock) or a fixed seed; and to
35 // specify the fixed seed explicitly
36 //
37 config const useRandomSeed = true,
38     seed = if useRandomSeed then SeedGenerator.clockMS else 31415;

40 //
41 // Configuration constants to control what's printed -- benchmark
42 // parameters, input and output arrays, and/or statistics
43 //
44 config const printParams = true,
45     printArrays = false,
46     printStats = true;

48 //
49 // The program entry point
50 //
51 def main() {
52     printConfiguration();

54 //
55 // MatVectSpace is a 2D domain of type indexType that represents the
56 // n x n matrix adjacent to the column vector b. MatrixSpace is a
57 // subdomain that is created by slicing into MatVectSpace,
58 // inheriting all of its rows and its low column bound. As our
59 // standard distribution library is filled out, MatVectSpace will be
60 // distributed using a BlockCyclic(blkSize) distribution.
61 //
62 const MatVectSpace: domain(2, indexType) = [1..n, 1..n+1],
63     MatrixSpace = MatVectSpace[...];

65 var Ab : [MatVectSpace] elemType, // the matrix A and vector b
66     piv: [1..n] indexType, // a vector of pivot values
67     x : [1..n] elemType; // the solution vector, x

69 var A => Ab[MatrixSpace], // an alias for the Matrix part of Ab
70     b => Ab[...; n+1]; // an alias for the last column of Ab

72 initAB (Ab);

74 const startTime = getCurrentTime(); // capture the start time

76 LUFactorize(n, Ab, piv); // compute the LU factorization

78 x = backwardSub(n, A, b); // perform the back substitution

80 const execTime = getCurrentTime() - startTime; // store the elapsed time

82 //
83 // Validate the answer and print the results
84 const validAnswer = verifyResults(Ab, MatrixSpace, x);
85 printResults(validAnswer, execTime);
86 }

88 //
89 // blocked LU factorization with pivoting for matrix augmented with
90 // vector of RHS values.
91 //
92 def LUFactorize(n: indexType, Ab: [1..n, 1..n+1] elemType,
93               piv: [1..n] indexType) {
94     const AbD = Ab.domain; // alias Ab.domain to save typing

96 // Initialize the pivot vector to represent the initially unpivoted matrix.
97 piv = 1..n;

99 /* The following diagram illustrates how we partition the matrix.

```

```

100 Each iteration of the loop increments a variable blk by blkSize;
101 point (blk, blk) is the upper-left location of the currently
102 unfactored matrix (the dotted region represents the areas
103 factored in prior iterations). The unfactored matrix is
104 partitioned into four subdomains: tl, tr, bl, and br, and an
105 additional domain (not shown), l, that is the union of tl and bl.

```

```

107     (point blk, blk)
108     +-----+-----+
109     |.....|.....|
110     |.....|.....|
111     |...+-----|
112     |...| | | |
113     |...| tl | | tr | |
114     |...| | | |
115     |...+-----|
116     |...| | | |
117     |...| | | |
118     |...| bl | | br | |
119     |...| | | |
120     |...| | | |
121     +-----+-----+
122 */
123 for blk in 1..n by blkSize {
124     const tl = AbD[blk..#blkSize, blk..#blkSize],
125           tr = AbD[blk..#blkSize, blk+blkSize..],
126           bl = AbD[blk+blkSize.., blk..#blkSize],
127           br = AbD[blk+blkSize.., blk+blkSize..],
128           l = AbD[blk.., blk..#blkSize];

130 //
131 // Now that we've sliced and diced Ab properly, do the blocked-LU
132 // computation:
133 //
134 panelSolve(Ab, l, piv);
135 if (tr.numIndices > 0) then
136     updateBlockRow(Ab, tl, tr);

138 //
139 // update trailing submatrix (if any)
140 //
141 if (br.numIndices > 0) then
142     schurComplement(Ab, blk);
143 }
144 }

146 //
147 // Distributed matrix-multiply for HPL. The idea behind this algorithm is that
148 // some point the matrix will be partitioned as shown in the following diagram:
149 //
150 // [1]----+-----+-----+
151 // | | bbbbb|bbbb|bbbb| Solve for the dotted region by
152 // | | bbbbb|bbbb|bbbb| multiplying the 'a' and 'b' region.
153 // | | bbbbb|bbbb|bbbb| The 'a' region is a block column, the
154 // +-----[2]-----+-----+ 'b' region is a block row.
155 // |aaaaa|.....|.....|.....|
156 // |aaaaa|.....|.....|.....| The vertex labeled [1] is location
157 // |aaaaa|.....|.....|.....| (ptOp, ptOp) in the code below.
158 // +-----+-----+-----+
159 // |aaaaa|.....|.....|.....| The vertex labeled [2] is location
160 // |aaaaa|.....|.....|.....| (ptSol, ptSol)
161 // |aaaaa|.....|.....|.....|
162 // +-----+-----+-----+
163 //
164 // Every locale with a block of data in the dotted region updates
165 // itself by multiplying the neighboring a-region block to its left
166 // with the neighboring b-region block above it and subtracting its
167 // current data from the result of this multiplication. To ensure that
168 // all locales have local copies of the data needed to perform this
169 // multiplication we copy the data A and B data into the replA and
170 // replB arrays, which will use a dimensional (block-cyclic,
171 // replicated-block) distribution (or vice-versa) to ensure that every
172 // locale only stores one copy of each block it requires for all of
173 // its rows/columns.
174 //
175 def schurComplement(Ab: [1..n, 1..n+1] elemType, ptOp: indexType) {
176     const AbD = Ab.domain;

178 //
179 // Calculate location of ptSol (see diagram above)
180 //
181 const ptSol = ptOp+blkSize;

183 //
184 // Copy data into replicated array so every processor has a local copy
185 // of the data it will need to perform a local matrix-multiply. These
186 // replicated distributions aren't implemented yet, but imagine that
187 // they look something like the following:
188 //
189 //var replAbD: domain(2)
190 //     distributed new Dimensional(BlkCyc(blkSize, Replicated))
191 //     = AbD[ptSol.., 1..#blkSize];
192 //
193 const replAD: domain(2) = AbD[ptSol.., ptOp..#blkSize],
194     replBD: domain(2) = AbD[ptOp..#blkSize, ptSol..];

196 const replA : [replAD] elemType = Ab[ptSol.., ptOp..#blkSize],
197     replB : [replBD] elemType = Ab[ptOp..#blkSize, ptSol..];

```

```

199 // do local matrix-multiply on a block-by-block basis
200 forall (row,col) in Abd[ptSol..., ptSol...] by (blkSize, blkSize) {
201 //
202 // At this point, the dgemms should all be local, so assert that
203 // fact
204 //
205 local {
206 const aBlkD = replAD[row..#blkSize, ptOp..#blkSize],
207        bBlkD = replBD[ptOp..#blkSize, col..#blkSize],
208        cBlkD = Abd[row..#blkSize, col..#blkSize];
209
210 dgemm(aBlkD.dim(1).length,
211        aBlkD.dim(2).length,
212        bBlkD.dim(2).length,
213        replA(aBlkD),
214        replB(bBlkD),
215        Ab(cBlkD));
216 }
217 }
218 }
219
220 //
221 // calculate C = C - A * B.
222 //
223 def dgemm(p: indexType, // number of rows in A
224           q: indexType, // number of cols in A, number of rows in B
225           r: indexType, // number of cols in B
226           A: [l..p, l..q] ?t,
227           B: [l..q, l..r] t,
228           C: [l..p, l..r] t) {
229 // Calculate (i,j) using a dot product of a row of A and a column of B.
230 for i in l..p do
231   for j in l..r do
232     for k in l..q do
233       C[i,j] -= A[i, k] * B[k, j];
234 }
235
236 //
237 // do unblocked-LU decomposition within the specified panel, update the
238 // pivot vector accordingly
239 //
240 def panelSolve(Ab: [] ?t,
241               panel: domain(2), indexType,
242               piv: [] indexType) {
243 const pnlRows = panel.dim(1),
244        pnlCols = panel.dim(2);
245
246 //
247 // Ideally some type of assertion to ensure panel is embedded in Ab's
248 // domain
249 //
250 assert(piv.domain.dim(1) == Ab.domain.dim(1));
251
252 if (pnlCols.length == 0) then return;
253
254 for k in pnlCols { // iterate through the columns
255   var col = panel[k.., k..k];
256
257 // If there are no rows below the current column return
258 if col.dim(1).length == 0 then return;
259
260 // Find the pivot, the element with the largest absolute value.
261 const (_, (pivotRow, _)) = maxloc reduce(abs(Ab(col)), col),
262        pivot = Ab[pivotRow, k];
263
264 // Swap the current row with the pivot row
265 piv[k] <=> piv[pivotRow];
266
267 Ab[k, ..] <=> Ab[pivotRow, ..];
268
269 if (pivot == 0) then
270   halt("Matrix can not be factorized");
271
272 // divide all values below and in the same col as the pivot by
273 // the pivot
274 if k+1 <= pnlRows.high then
275   Ab(col)[k+1.., k..k] /= pivot;
276
277 // update all other values below the pivot
278 if k+1 <= pnlRows.high && k+1 <= pnlCols.high then
279   forall (i,j) in panel[k+1.., k+1..] do
280     Ab[i,j] -= Ab[i,k] * Ab[k,j];
281 }
282 }
283
284 //
285 // Update the block row (tr for top-right) portion of the matrix in a
286 // blocked LU decomposition. Each step of the LU decomposition will
287 // solve a block (tl for top-left) portion of a matrix. This function
288 // solves the rows to the right of the block.
289 //
290 def updateBlockRow(Ab: [] ?t, tl: domain(2), tr: domain(2)) {
291 const tlRows = tl.dim(1),
292        tlCols = tl.dim(2),
293        trRows = tr.dim(1),
294        trCols = tr.dim(2);
295
296 assert(tlCols == trRows);
297
298 //
299 // Ultimately, we will probably want to do some replication of the
300 // tl block in order to make this operation completely localized as
301 // in the dgemm. We have not yet undertaken that optimization.
302 //
303 for i in trRows do
304   forall j in trCols do
305     for k in tlRows.low..i-1 do
306       Ab[i, j] -= Ab[i, k] * Ab[k, j];
307 }
308
309 //
310 // compute the backwards substitution
311 //
312 def backwardSub(n: int,
313                A: [l..n, l..n] elemType,
314                b: [l..n] elemType) {
315 var x: [b.domain] elemType;
316
317 for i in [b.domain by -1] {
318   x[i] = b[i];
319
320   for j in [i+1..b.domain.high] do
321     x[i] -= A[i,j] * x[j];
322
323   x[i] /= A[i,i];
324 }
325
326 return x;
327 }
328
329 //
330 // print out the problem size and block size if requested
331 //
332 def printConfiguration() {
333 if (printParams) {
334   if (printStats) then printLocalesTasks(tasksPerLocale=1);
335   printProblemSize(elemType, numMatrices, n, rank=2);
336   writeln("block size = ", blkSize, "\n");
337 }
338 }
339
340 //
341 // construct an n by n+1 matrix filled with random values and scale
342 // it to be in the range -1.0..1.0
343 //
344 def initAB(Ab: [] elemType) {
345 fillRandom(Ab, seed);
346 Ab = Ab * 2.0 - 1.0;
347 }
348
349 //
350 // calculate norms and residuals to verify the results
351 //
352 def verifyResults(Ab, MatrixSpace, x) {
353 var A => Ab[MatrixSpace],
354     b => Ab[.., n+1];
355
356 initAB(Ab);
357
358 const axmbNorm = norm(gaxpyMinus(n, n, A, x, b), normType.normInf);
359
360 const aInorm = norm(A, normType.norm1),
361        aInfNorm = norm(A, normType.normInf),
362        x1Norm = norm(x, normType.norm1),
363        xInfNorm = norm(x, normType.normInf);
364
365 const resid1 = axmbNorm / (epsilon * aInorm * n),
366        resid2 = axmbNorm / (epsilon * aInorm * x1Norm),
367        resid3 = axmbNorm / (epsilon * aInfNorm * xInfNorm);
368
369 if (printStats) {
370   writeln("resid1: ", resid1);
371   writeln("resid2: ", resid2);
372   writeln("resid3: ", resid3);
373 }
374
375 return max(resid1, resid2, resid3) < 16.0;
376 }
377
378 //
379 // print success/failure, the execution time and the Gflop/s value
380 //
381 def printResults(successful, execTime) {
382 writeln("Validation: ", if successful then "SUCCESS" else "FAILURE");
383 if (printStats) {
384   writeln("Execution time = ", execTime);
385   const GflopPerSec = ((2.0/3.0) * n**3 + (3.0/2.0) * n**2) / execTime * 10e-9;
386   writeln("Performance (Gflop/s) = ", GflopPerSec);
387 }
388 }
389
390 //
391 // simple matrix-vector multiplication, solve equation A*x=y
392 //
393 def gaxpyMinus(n: indexType,
394               m: indexType,
395               A: [l..n, l..m],
396               x: [l..m],
397               y: [l..n]) {
398 var res: [l..n] elemType;
399
400 for i in l..n do
401   for j in l..m do
402     res[i] += A[i,j]*x[j];
403
404 for i in l..n do
405   res[i] -= y[i];
406
407 return res;
408 }

```

## F Shared Problem Size Module Code

```
1 //
2 // A shared module for computing the appropriate problem size for the
3 // HPCC benchmarks
4 //
5 module HPCCProblemSize (
6 //
7 // Use the standard modules for reasoning about Memory and Types
8 //
9 use Memory, Types;
10
11 //
12 // The main routine for computing the problem size
13 //
14 def computeProblemSize(numArrays: int, // #arrays in the benchmark
15     elemType, // the element type of those arrays
16     rank=1, // rank of the arrays
17     returnLog2=false, // whether to return log2(probSize)
18     memFraction=4, // fraction of mem to use (eg, 1/4)
19     type retType = int(64)): retType { // type to return
20 //
21 // Compute the total memory available to the benchmark using a sum
22 // reduction over the amount of physical memory (in bytes) owned
23 // by the set of locales on which we're running. Then compute the
24 // number of bytes we want to use as defined by memFraction and the
25 // number that will be required by each index in the problem size.
26 //
27 const totalMem = + reduce Locales.physicalMemory(unit = MemUnits.Bytes),
28     memoryTarget = totalMem / memFraction,
29     bytesPerIndex = numArrays * numBytes(elemType);
30
31 //
32 // Use these values to compute a base number of indices
33 //
34 var numIndices = memoryTarget / bytesPerIndex;
35
36 //
37 // If the user requested a 2**n problem size, compute appropriate
38 // values for numIndices and lgProblemSize
39 //
40 var lgProblemSize = log2(numIndices);
41 if (returnLog2) {
42     if rank != 1 then
43         halt("computeProblemSize() can't compute 2D 2**n problem sizes yet");
44     numIndices = 2**lgProblemSize;
45     if (numIndices * bytesPerIndex <= memoryTarget) {
46         numIndices *= 2;
47         lgProblemSize += 1;
48     }
49 }
50 //
51
52 // Compute the smallest amount of memory that any locale owns
53 // using a min reduction and ensure that it is sufficient to hold
54 // an even portion of the problem size.
55 //
56 const smallestMem = min reduce Locales.physicalMemory(unit = MemUnits.Bytes);
57 if ((numIndices * bytesPerIndex)/numLocales > smallestMem) then
58     halt("System is too heterogeneous: blocked data won't fit into memory");
59
60 //
61 // return the problem size as requested by the callee
62 //
63 if returnLog2 then
64     return lgProblemSize: retType;
65 else
66     select rank {
67         when 1 do return numIndices: retType;
68         when 2 do return ceil(sqrt(numIndices)): retType;
69         otherwise halt("Unexpected rank in computeProblemSize");
70     }
71 }
72
73 //
74 // Print out the machine configuration used to run the job
75 //
76 def printLocalesTasks(tasksPerLocale=1) {
77     writeln("Number of Locales = ", numLocales);
78     writeln("Tasks per locale = ", tasksPerLocale);
79 }
80
81 //
82 // Print out the problem size, #bytes per array, and total memory
83 // required by the arrays
84 //
85 def printProblemSize(type elemType, numArrays, problemSize: ?psType,
86     param rank=1) {
87     const bytesPerArray = problemSize**rank * numBytes(elemType),
88         totalMemInGB = (numArrays * bytesPerArray:real) / (1024**3),
89         lgProbSize = log2(problemSize):psType;
90
91     write("Problem size = ", problemSize);
92     for i in 2..rank do write(" x ", problemSize);
93     if (2**lgProbSize == problemSize) {
94         write(" (2**", lgProbSize);
95         for i in 2..rank do write(" x 2**", lgProbSize);
96         write(")");
97     }
98     writeln();
99     writeln("Bytes per array = ", bytesPerArray);
100    writeln("Total memory required (GB) = ", totalMemInGB);
101 }
102 }
```