

Global HPCC Benchmarks in Chapel: STREAM Triad, Random Access, and FFT

(Revision 1.3 — CUG 2007)

Bradford L. Chamberlain, Steven J. Deitz, Mary Beth Hribar, Wayne A. Wong
Chapel Team, Cray Inc.
chapel_info@cray.com

Abstract—Chapel is a new parallel programming language being developed by Cray Inc. as part of its participation in DARPA’s High Productivity Computing Systems program. In this article, we describe Chapel implementations of the global HPC Challenge (HPCC) benchmarks for the STREAM Triad, Random Access, and FFT computations. The Chapel implementations use 5–13× fewer lines of code than the reference implementations supplied by the HPCC team. All codes in this article compile and run with the current version of the Chapel compiler. We provide an introduction to Chapel, highlight key features used in our HPCC implementations, and discuss our plans and challenges for obtaining performance for these codes as our compiler development progresses. The full codes are listed in appendices to this article.

I. INTRODUCTION

Chapel is a new parallel programming language being developed by Cray Inc. as part of its participation in DARPA’s High Productivity Computing Systems program (HPCS). The Chapel team is working to design, implement, and demonstrate a language that improves parallel programmability, portability, and code robustness as compared to current practice while producing programs whose performance is comparable to or better than MPI.

In this article, we present Chapel implementations of three of the global HPC Challenge (HPCC) benchmarks—STREAM Triad, Random Access (RA), and Fast Fourier Transform (FFT). These benchmarks were developed as a means of measuring sustained parallel performance by stressing traditional system bottlenecks. We describe our Chapel implementations of the benchmarks, highlighting key language concepts used to implement them. We submitted an earlier version of these codes to the second annual 2006 HPC Challenge competition for “most productive” implementations, judged 50% on elegance and 50% on performance. We were selected as one of six finalists to present at the HPCC BOF at SC06.

The codes presented in this article compile and execute with our current Chapel compiler (version 0.4.550). To date, our compiler development effort has focused on providing a single-node prototype of Chapel’s features in order to support experimentation with the language and generate feedback on its specification. While our language specification and compiler architecture have both been designed with distributed-memory

execution and performance optimizations in mind, minimal development effort has been invested toward these activities to date. As a result, this article does not contain performance results since they would not reflect our team’s emphasis. Beginning in 2007, we have focused our implementation effort on serial performance optimizations and distributed-memory execution, and we expect to have publishable performance results for our codes within the coming year.

Though this article does not report on performance, we did write our benchmark implementations to be portable, performance-minded parallel codes rather than simply optimizing for elegance, clarity, and size. Our goal was to write versions of the benchmarks that would serve as clear reference implementations for future HPCC competitors while also producing codes that should result in good performance as our compiler matures. In lieu of performance results, this article contains a discussion of the top performance-related issues for each benchmark and our strategies for achieving good performance in the future.

The rest of this paper is organized as follows: In Section II we give an overview of our results. In Section III we provide an introduction to Chapel’s motivating themes. In Section IV we describe the coding conventions that we adopted in writing these benchmarks. Sections V, VI, and VII each describe one of the three benchmarks, providing a high-level overview of our approach in Chapel as well as performance-related issues that highlight Chapel’s benefits and challenges. Finally in Section VIII, we summarize and provide a brief status report for the Chapel project. Our complete source listings are provided in the appendices.

II. OVERVIEW OF RESULTS

Table I categorizes and counts the number of lines of code utilized by our HPCC implementations. The line counts for each benchmark are represented in a column of the table. The fourth data column represents the common *HPCCProblemSize* module that is shared by the benchmarks to compute and print the problem size. For the Random Access benchmark, each entry is expressed as a sum—the first value represents the benchmark module itself, the second represents an additional module used to generate the random number stream, and the final value is the sum for both modules.

The rows of the table are used to group the lines of code into various categories and running totals. The first two rows

TABLE I
LINE COUNTS AND CLASSIFICATIONS FOR HPCC CODES IN CHAPEL

| line count type | Benchmark Code | | | |
|------------------------|---------------------|----------------------|------------|---------------|
| | <i>STREAM Triad</i> | <i>Random Access</i> | <i>FFT</i> | <i>Common</i> |
| Kernel computation | 1 | 4 + 19 = 23 | 48 | 0 |
| Kernel declarations | 9 | 14 + 21 = 35 | 25 | 20 |
| <i>Total kernel</i> | <i>10</i> | <i>18 + 40 = 58</i> | <i>73</i> | <i>20</i> |
| Initialization | 9 | 1 + 0 = 1 | 21 | 0 |
| Verification | 6 | 15 + 0 = 15 | 13 | 0 |
| Results and Output | 29 | 18 + 0 = 18 | 16 | 12 |
| Total Benchmark | 54 | 52 + 40 = 92 | 123 | 32 |
| Debug and Test | 9 | 6 + 0 = 6 | 8 | 3 |
| Blank | 27 | 26 + 12 = 38 | 50 | 9 |
| <i>Total Program</i> | <i>90</i> | <i>84 + 52 = 136</i> | <i>181</i> | <i>44</i> |

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 a single line of code, while its supporting variable and subroutine declarations require 9 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.

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 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 in the next row to indicate 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 *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 are provided to serve as a checksum against the line number labels that appear in the appendices.

Table II compares the total SLOC for our Chapel codes with the standard HPCC reference implementations. The Chapel result for each code is obtained by summing its *Total Benchmark* result from Table I with that of the common module. The reference results are the sum of the *Framework* and *Parallel* numbers reported in the table from the HPCC website's FAQ.¹

This table shows that our Chapel codes are approximately 5–13× smaller than the reference implementations. While shorter codes are not necessarily better or easier to understand, we believe that our Chapel implementations are not only succinct, but also clear representations of the benchmarks that will perform well as our compiler matures. The rest of this paper examines the codes qualitatively to complement the quantitative results in this section.

¹<http://www.hpcchallenge.org/faq/index.html>

TABLE II
SLOC COMPARISON BETWEEN REFERENCE AND CHAPEL HPCC CODES

| | Benchmark Code | | |
|-----------------------------------|---------------------|----------------------|-------------|
| | <i>STREAM Triad</i> | <i>Random Access</i> | <i>FFT</i> |
| SLOC for Reference Implementation | 433 | 1668 | 1406 |
| SLOC for Chapel Implementation | 86 | 124 | 155 |
| <i>SLOC Ratio</i> | <i>5.03</i> | <i>13.45</i> | <i>9.07</i> |

III. CHAPEL'S MOTIVATING THEMES

One of Chapel's primary themes is to support general parallel programming using high-level abstractions. Chapel does this through its support for a *global-view programming model* that raises the level of abstraction for both data structures and control flow as compared to parallel programming models currently used in production.

Global-view data structures are arrays and other data aggregates whose size and indices are expressed globally in spite of the fact that their implementations may be distributed across the memories of multiple nodes or *locales*.² This contrasts with most parallel languages used in practice, which require users to partition distributed data aggregates into per-processor chunks, either manually or using language abstractions. HPF and ZPL are two other recent parallel languages that support global-view data structures [15], [6], though in a more restricted form than Chapel.

A global view of control means that a user's program commences execution with a single logical thread of control and that additional parallelism is introduced through the use of specific language concepts. All parallelism in Chapel is implemented via multithreading, though these threads are created via high-level language concepts and managed by the compiler and runtime, rather than through explicit fork/join-style programming. An impact of this approach is that Chapel can express parallelism that is more general than the Single Program, Multiple Data (SPMD) model that today's most common parallel programming approaches use as the basis for their programming and execution models. Examples include Co-Array Fortran, Unified Parallel C (UPC), Titanium, HPF, ZPL, SHMEM, and typical uses of MPI [18], [11], [20], [15], [6], [2], [19], [12]. Our multithreaded execution model is perhaps most similar to that which is supported by the Cilk language or the Cray MTA's runtime libraries [13], [1]. Moreover, Chapel's general support for parallelism does not preclude the user from coding in an SPMD style if they wish.

Supporting general parallel programming also means targeting a broad range of parallel architectures. Chapel is designed to target a wide spectrum of HPC hardware including clusters of commodity processors and SMPs; vector, multithreading, and multicore processors; distributed memory, shared address space, and shared memory architectures; networks of any topology; and custom vendor architectures. Our portability

²A *locale* in Chapel is a unit of the target architecture that supports computation and data storage. Locales are defined for each architecture such that a locale's threads will all have similar access times to any specific memory address. For commodity clusters, each of their (single-core) processors, multicore processors, or SMP nodes would be considered a locale.

goal is to have any legal Chapel program run correctly on all of these architectures, and for Chapel programs that express parallelism in an architecturally-neutral way to perform reasonably on all of them. Naturally, Chapel programmers can tune their codes to more closely match a particular machine's characteristics, though doing so may cause the program to be a poor match for other architectures. In this article we present codes written in an architecturally-neutral manner and note places where they could be tuned to better match specific architectural characteristics, including memory models and network capabilities.

A second theme in Chapel is to allow the user to optionally and incrementally specify where data and computation should be placed on the physical machine. We consider this control over program locality to be essential for achieving scalable performance on large machine sizes given current architectural trends. Such control contrasts with shared-memory programming models like OpenMP [7] which present the user with a flat memory model. It also contrasts with SPMD-based programming models in which such details are explicitly specified by the programmer on a process-by-process basis via the multiple cooperating program instances.

A third theme in Chapel is support for object-oriented programming (OOP), which has been instrumental in raising productivity in the mainstream programming community. Chapel supports traditional reference-based classes as well as value classes. Programmers are not required to use an object-oriented style in their code, so that traditional Fortran and C programmers need not adopt a new programming paradigm in order to use Chapel effectively.³

Chapel's fourth theme is support for generic programming and polymorphism, allowing code to be written in a style that is generic across types and thereby applicable to variables of various types, sizes, and precisions. The goal of these features is to support exploratory programming as in popular interpreted and scripting languages, and to support code reuse by allowing algorithms to be expressed without explicitly replicating them for each possible type. This flexibility at the source level is implemented by having the compiler create versions of the code for each required type signature rather than by relying on dynamic typing which would incur unacceptable runtime overheads.

Chapel's first two themes are designed to provide support for general, performance-oriented parallel programming through high-level abstractions. The second two themes are supported to help narrow the gulf that exists between parallel programming languages and popular mainstream programming or scripting languages. The benchmarks in this paper illustrate Chapel's support for global-view programming, for locality control, and for generic programming. Due to the relatively straightforward style of parallel computation required for these benchmarks, this article does not demonstrate many of Chapel's features for more general parallelism and locality control. We also chose not to utilize OOP concepts in our

benchmark codes since we do not believe that the introduction of objects would greatly improve the benchmarks' clarity, generality, organization, or performance.

For a more complete introduction to Chapel, the reader is referred to our project's website⁴, overview publications, and draft language specification [8], [5], [3].

IV. CODING CONVENTIONS

In writing these codes, we used the HPC Challenge Class 2 Official Specification as our primary guide for defining the computations [10]. We studied and benefited from the HPCC reference implementations as well as the 2005 finalist codes, but typically chose to express the benchmarks in our own style rather than trying to mimic pre-existing implementations.

In particular, we chose names for our variables and sub-routines that we found descriptive and appealing rather than trying to adhere to the naming conventions of previous implementations. The primary exception to this rule is for variables named in the written specification, such as m , n , and N_U . For these variables, we adopted the specification's names in order to clarify the ties between our code and the official description.

Several concerns directed our coding design (in roughly this order):

- faithfulness to the written specification
- ability to generate code with good performance
- clarity and elegance of the solution, emphasizing readability over minimization of code size
- appropriate and demonstrative uses of Chapel features
- implementations that are generic with respect to types and problem parameters
- support for execution-time control over key program parameters like problem sizes
- ability to be tested in our nightly regression suite

Some of these motivations, particularly the last three, cause our programs to be slightly more general than required by the written specification. However, we believe that they also result in more interesting and realistic application codes.

Structurally, we tried to keep the timed kernel of the computation in the program's `main()` procedure, moving other functionality such as initialization, verification, and I/O into separate routines. In the Random Access and FFT benchmarks, we also abstracted kernels of the computation into helper iterators and routines to improve abstraction and reuse.

Stylistically, we tend to use mixed-case names to express multi-word identifiers, rather than underscores. We typically use an initial lower-case letter when naming procedures and non-distributed variables, while domains, distributed arrays, and class instances start with an upper-case letter.

In our code listings, **boldface** text is used to indicate Chapel's reserved words and standard types, while “...” represents code that is elided in an excerpt for brevity.

As mentioned previously, we approached these codes as we would for a large-scale parallel machine—thus they contain distributed data structures, parallel loops, and array-based parallelism. Since our current compiler only supports execution

³Note that many of Chapel's standard library capabilities are implemented using objects, so these may require Fortran and C programmers to utilize a method-invocation style of syntax when using them. However, such use does not necessitate broader adoption of OOP methodologies.

⁴<http://chapel.cs.washington.edu>

```

const ProblemSpace: domain(1) distributed(Block) = [1..m];
var A, B, C: [ProblemSpace] elemType;

A = B + alpha * C;

```

Fig. 1. STREAM Triad Kernel in Chapel

on a single locale, these constructs will necessarily fall into the degenerate cases of allocating data structures on a single locale and, for the single-core processors we use, executing parallel loops and array statements using a single thread.

V. STREAM TRIAD

The STREAM Triad benchmark asks the programmer to take two vectors of random 64-bit floating-point values, b and c , and to use them to compute $a = b + \alpha \cdot c$ for a given scalar value α . As with all of the HPC benchmarks, the problem size for the vectors must be chosen such that they consume $1/4 - 1/2$ of the system memory. STREAM Triad is designed to stress local memory bandwidth since the vectors may be allocated in an aligned manner such that no communication is required to perform the computation.

A. Overview of STREAM Triad in Chapel

Our approach to the STREAM Triad benchmark is summarized by the lines of code in Figure 1. This code excerpt presupposes the definition of two named values, m defining the problem size and $alpha$ defining the scalar multiplication value for the Triad computation (α). It also refers to a named type, $elemType$, that represents the element type to be stored in the vectors.

The first line declares a constant named *ProblemSpace* that is defined to be a *domain*—a first-class representation of an index set, potentially distributed across the memories of multiple locales. In this instance, the index set is declared to be 1-dimensional and to describe the indices $\{1, 2, \dots, m\}$. *ProblemSpace*'s definition also specifies that it should be distributed across the user's locale set using the "Block" distribution. Such distributions specify how a domain's indices should be mapped to a set of locales and how they should be represented within each locale's memory. In this case, the *Block* distribution is a standard Chapel distribution that assigns contiguous blocks of indices to each locale.

The next line uses the *ProblemSpace* domain to declare three arrays— A , B , and C —used to represent the vectors, a , b , and c from the written specification. The domain's index set defines the size and shape of these arrays, and its distribution specifies the arrays' distributed implementation across the locales and within their local memories. The identifier *elemType* specifies the type of each array element (defined to be a 64-bit floating-point value in our implementation).

The final line expresses the computation itself, using whole-array syntax to specify the elementwise multiplications, additions, and assignments needed to perform the Triad computation. Whole-array operations like this one are implicitly parallel and each locale will perform the operations for the array elements that it owns, as defined by *ProblemSpace*'s

distribution (since that was the domain used to define all three arrays).

B. STREAM Triad Performance Notes

Once we have a distributed-memory implementation of Chapel and some additional scalar optimizations, we expect our STREAM Triad implementation to perform as well as the target architecture allows—constrained only by the rate of a locale's local memory accesses. The timed Triad computation is embarrassingly parallel, and its expression in Chapel makes this obvious through the use of whole-array operations on arrays declared using a single domain. This demonstrates one of the many advantages of domains, which is that they often provide the compiler with valuable information about the relative alignment of arrays in distributed memory [4].

Our main implementation challenge in obtaining good performance relates to the definition of the *Block* distribution itself. While our team has extensive experience implementing block distributions whose performance competes with or outperforms hand-coded benchmarks [4], [9], a major goal for Chapel is to write all of its standard distributions using the same mechanism that advanced Chapel users would use to author their own distributions. This contrasts with prior language work in which compilers typically have special knowledge of standard distributions. Our challenge is to design a distribution interface that allows the compiler to implement computations using standard distributions without compromising performance. If we fail to meet this challenge, we can always fall back on the approach of embedding knowledge of standard distributions into the compiler; however, this would be unfortunate because it suggests that users may face a performance penalty if they need to use a distribution that Chapel does not provide.

Our final performance-related note for STREAM Triad relates to the fact that Chapel uses a multithreaded execution model to implement its program semantics. Users who are accustomed to an SPMD execution model may worry that for embarrassingly parallel applications like STREAM Triad, the overhead of supporting a general multithreaded execution environment may be overkill since simpler SPMD execution is sufficient. While we believe that multithreaded execution is required in the next generation of parallel languages to support general parallel programming, we also recognize the efficiency advantages of the SPMD model due to its simplicity. For this reason, we expect to optimize our implementation for phases of computation in which statically scheduled threads suffice, either in the compiler-generated code or in its runtime support system.

Using STREAM Triad as a specific example, when the compiler analyzes the source code, it can trivially see that

```

const TableSpace: domain(1, indexType) distributed(Block) = [0..m];
var T: [TableSpace] elemType;

const UpdateSpace: domain(1, indexType) distributed(Block) = [0..N_U];

[i in TableSpace] T(i) = i;

forall block in UpdateSpace.subBlocks do
  for r in RAStrstream(block) do
    T(r & indexMask) ^= r;
...

iterator RAStrstream(block) {
  var val = getNthRandom(block.low);
  for i in block {
    getNextRandom(val);
    yield val;
  }
}

```

Fig. 2. Random Access Kernel and Stream Iterator in Chapel

there is neither task parallelism nor explicit parallelism specified in the code. All parallelism comes from performing implicitly parallel operations on distributed arrays—whole-array operations, pseudo-random fills, and reductions. Given such information, the compiler may choose to generate code that schedules a single thread per core across the distributed machine, using them in an SPMD fashion. Recall that Chapel’s global-view model of control specifies that the entry point is executed by a single *logical* thread, allowing an SPMD execution strategy to be used as long as the illusion of a single startup thread is maintained and the program’s semantics allow it. Similarly, the compiler may use an SPMD execution strategy for a single phase of the user’s program and more general techniques for other phases when it can prove such an approach is legal.

Alternatively, Chapel’s runtime support for multithreading may be tuned so that when a single thread is running per core, the overheads required by more general multithreading capabilities are minimized. In the limit, this may allow multithreaded Chapel codes executing in an SPMD style to be competitive with explicit SPMD execution models. Even in this scenario it should be expected that compiler-generated information about the threading and communication requirements of a computation phase would help Chapel’s runtime libraries shut down unnecessary sources of runtime overhead, such as services for creating remote threads or servicing one-sided communication requests on a commodity cluster.

VI. RANDOM ACCESS (RA)

The Random Access benchmark computes pseudo-random updates to a large distributed table of 64-bit integer values, T . Each table location T_i is initialized to store its index i in order to support verification of the computation afterwards. The pseudo-random values that are generated are used both to specify the table location to be updated and the value to use for the update. The updates themselves are performed by xor-ing the random value into the table location. Because multiple updates may attempt to modify the same table location simultaneously, the benchmark permits a certain fraction

of such conflicting updates to be lost. Random Access is designed to test a system’s ability to randomly access memory at a fine granularity. The benchmark permits updates to be performed in batches of a certain size as a concession to architectures and programming models that do not support fine-grain communication.

A. Overview of Random Access in Chapel

Figure 2 shows an excerpt from the Chapel implementation of the Random Access benchmark that summarizes the computation. It assumes the definition of three named values: m which defines the problem size, N_U which represents the number of updates (N_U), and $indexMask$ which stores the bitmask used to create legal indices into the table. It also refers to two named types, $indexType$ and $elemType$, used to represent the types that should be used to store the indices and elements of table T , respectively.

The first two lines define a domain and array used to represent the distributed table, T , that is randomly accessed by the benchmark. The next line defines a second domain, $UpdateSpace$ that represents the distribution of the table update work across the locale set. Note that $UpdateSpace$ is not used to allocate any arrays, only to distribute the computation space.

Both of the domain declarations in this excerpt specify a second domain argument in addition to the rank, $indexType$. This indicates the integer type to be used for representing the domain’s indices which defaults to `int` if left unspecified. In this benchmark, we use a 64-bit unsigned integer as the index type to support large tables and to simplify the indexing operations using 64-bit pseudo-random values. Note furthermore that the domain initializers use a slightly different domain specifier than in the STREAM Triad benchmark: $[lo..hi]$. This domain syntax is meant to suggest open intervals in math, and defines the index set $\{lo, lo + 1, \dots, hi - 1\}$.

The computation starts on the fourth line of code which initializes the table using a *forall expression*, representing the parallel loop: “for all indices i in $TableSpace$, assign T_i the value i .”

The main loop in the code computes the random accesses using a nested loop which is defined by invoking two *iterators*. In Chapel, an iterator is like a function, but rather than creating a single result via a `return`, it generates a stream of results via a `yield` keyword. Chapel’s `yield` statement is like a traditional `return` except that the iterator’s execution conceptually continues onwards from that point. Iterators are typically used to specify loops, and can be viewed as a concept for abstracting complex control flow away from loop headers just as functions are used to abstract code away from an invocation point. As an example, the complicated loop nest used to implement a multidimensional tiled iteration can be specified once using an iterator and then invoked multiple times to tile many loops in a clean, modular manner.

The outer loop used to compute the random updates invokes a standard iterator defined by the *Block* distribution to generate sub-blocks of work that can be performed in parallel. The number of sub-blocks created will be determined by the degree of parallelism supported by the target architecture. For example, on a machine with 16 locales of 4 cores each, 64 sub-blocks would be generated by the iterator, one for each core.

The inner loop uses a user-defined iterator `RAStream()` that is shown at the bottom of Figure 2. This iterator is written to generate a stream of values corresponding to the indices specified by its argument *block*. It is written using two helper functions not shown here, `getNthRandom()` and `getNextRandom()`. Upon computing each value, the iterator *yields* it and then goes on to compute the next one. After exiting the `for` loop, the end of the iterator is reached, and it terminates.

The body of the loop uses the values yielded by the iterator to update the table using Chapel’s xor operator (`^`) used in its read-modify-write form (`^=`). Since multiple tasks created by the outer loop may attempt to update a given table position simultaneously, this code contains a race condition as permitted by the benchmark specification. In order to avoid the race condition (as we must during the benchmark’s verification stage) the loop body can be prefixed by Chapel’s `atomic` keyword to specify that the update must be computed as though it was instantaneous from the point-of-view of other threads. In such a case, the loop body would be written:

```
atomic T(r & indexMask) ^= r;
```

By implementing the random stream using an iterator, our code allows different pseudo-random number generation techniques to be swapped in and out without modifying the computation itself. Moreover, the use of an iterator allows additional parallelism to be introduced by rewriting the `RAStream()` iterator to return multiple values rather than singletons. Such an implementation would cause the table update operation to be promoted (as in the STREAM Triad benchmark) resulting in implicit parallelism and the potential for batching communication.

B. Random Access Performance Notes

For most platforms, the limiting factor for Random Access performance is the hardware rather than the software. We anticipate this to be the case for our parallel Chapel

implementation as well. Without significant optimization, the default implementation that we provide will utilize the full parallelism of a machine’s processors, but will make updates to the global table one element at a time, requiring a lot of fine-grain communication. For architectures that are tuned for message-passing workloads this is a worst-case scenario, while for others, having many small messages in flight is not an issue. As we noted in the previous section, the iterators used to describe the computation can be modified to create chunks of work of varying sizes to amortize these communication overheads or to use vector scatter technologies (if available on the architecture). In addition, a programmer who is committed to tuning the computation for a specific architecture could restructure the loops, add new levels of parallelism, or batch updates manually in order to make the computation match the architecture’s characteristics. For example, on an architecture with single-threaded processors and poor fine-grain communication, the programmer might choose to rewrite the outer loop to oversubscribe each locale since the communication overheads are likely to dominate the costs of software multithreading.

Our second performance-related note for Random Access relates to iterators. Iterators are one of several language concepts in Chapel that define semantics without specifying a specific implementation approach. This provides the user with a nice abstraction while allowing the compiler to consider a palette of implementation options and select one that is well-suited for the specific case. In particular, an iterator’s definition and uses can be analyzed to select an implementation technique that is appropriate for each invocation. While iterator bodies can contain complicated control structures and multiple `yield` statements in the general case, and while they can be invoked in complicated ways such as *zipped iteration*, many common cases simply use iterators as a clean and modular way of expressing a loop. The `RAStream()` iterator in our Random Access code is one such example, defining a straightforward loop structure and invoking it in a straightforward way. In such cases, the compiler can simply inline the iterator’s definition in place of the loop statement and then inline the loop’s body into the iterator’s `yield` statement. This results in no runtime overhead compared to writing the loop out explicitly, while supporting loop-level modularity and the ability for the iterator to be invoked in more complicated ways in other loops. For an overview of some of our early and more general implementation approaches for iterators, please refer to our previous work [14].

Another Chapel concept that defines semantics without specifying an implementation mechanism is the *atomic section*. Atomic sections that contain arbitrary code can be implemented using a number of sophisticated techniques, including the body of research known as *Software* and *Hardware Transactional Memory* in which groups of memory operations appear to execute atomically from the point of view of other threads. Implementing Software Transactional Memory effectively on a distributed memory machine remains an open problem, and one that we are actively investigating. For simple atomic statements like the one contained in this code, however, simpler implementation techniques can be utilized. For exam-

```

for i in [2..log2(numElements)] by 2 {
  const m = radix*span,
        m2 = 2*m;

  forall (k,k1) in (ADom by m2, 0..) {
    var wk2 = W(k1),
        wk1 = W(2*k1),
        wk3 = (wk1.re - 2 * wk2.im * wk1.im,
              2 * wk2.im * wk1.re - wk1.im):elemType;

    forall j in [k..k+span) do
      butterfly(wk1, wk2, wk3, A[j..j+3*span by span]);

    wk1 = W(2*k1+1);
    wk3 = (wk1.re - 2 * wk2.re * wk1.im,
          2 * wk2.re * wk1.re - wk1.im):elemType;
    wk2 *= 1.0i;

    forall j in [k+m..k+m+span) do
      butterfly(wk1, wk2, wk3, A[j..j+3*span by span]);
  }
  span *= radix;
}
...

def butterfly(wk1, wk2, wk3, inout A:[1..radix]) { ... }

```

Fig. 3. 1D Radix-4 FFT in Chapel

ple, by guarding the verification loop containing the atomic updates with the proper synchronization, the compiler can use locks to ensure that no two threads try to access the same table element simultaneously. In the common case where there are no conflicts, the overhead of these locks should be negligible compared to the overhead of the fine-grain remote accesses. Moreover, on architectures with support for *atomic memory operations (AMOs)*, *extended memory semantics (EMS)*, or *compare-and-swap (CAS)* instructions, the Chapel compiler can utilize such capabilities to further reduce overheads.

VII. FAST FOURIER TRANSFORM (FFT)

The FFT benchmark requires the user to compute a 1D discrete fourier transform on a vector of pseudo-random values. The user is given a fair amount of latitude in choosing an FFT algorithm and approach as long as the outcome is correct. Depending on how it is written, the FFT benchmark could stress a number of system parameters including floating point operations, local memory bandwidth, cache size, and network.

A. Overview of FFT in Chapel

We chose to implement a radix-4 FFT in order to take advantage of its improved Flops-to-MemOps ratio. Our main loop for FFT is shown in Figure 3. The outer loop iterates over the FFT phases using Chapel’s support for *strided ranges* of integer values. The next loop uses *zippered iteration* to traverse two iteration spaces simultaneously, the second of which is defined using the *indefinite range* “0..” to indicate that it should start counting at zero and iterate as many times as the iterator with which it is zippered. The iterator indices, k and $k1$ are defined using a tuple-style variable declaration.

The inner loops describe the parallelizable calls to the `butterfly()` routine, which is written to take three twiddle

values of unspecified type that it uses to update values within a 4-element array, A . The argument types are unspecified and are typically represented using Chapel’s 128-bit complex type. However, by omitting these argument types, programmers can optimize calls that they know to use twiddle factors with zero components by passing in values of Chapel’s real or pure imaginary types. Such calls would cause the compiler to instantiate the butterfly routine for these types, thereby eliminating the extraneous floating point operations against zero values that would be incurred by a routine written specifically for complex values.

The calls to `butterfly()` are also of interest because they pass a strided slice of A to the formal vector argument that was defined using the anonymous domain “[1..radix]”. This results in the creation of an *array view*, allowing the original array elements to be accessed within the routine using the local indices $1, 2, \dots, \text{radix}$.

B. FFT Performance Notes

It is generally known to be very challenging to express a global 1D FFT that is clear while also being well-tuned for general parallel architectures—there are simply too many degrees of freedom and too many architectural characteristics that can impact performance. In our approach, we have implemented the computation cleanly and expressed both outer- and inner-loop concurrency in order to maximize the parallelism available to the compiler. An aggressive compiler might recognize that the trip counts of these loops vary with the FFT’s phases and create specialized versions of `dfft()`’s loops to optimize for outer- and inner-loop parallelism, or a blend of the two. However, even in such a case, the target architecture’s memory model and network can still play a

key role in how the performance-minded user might express additional optimizations to maximize efficiency.

At one end of the spectrum, even on a parallel machine with a flat shared memory like the Cray MTA, the clean expression of the computation as given here is unlikely to perform optimally. In our 2006 HPCC entry, we provided an alternate Chapel implementation of FFT that is based on the 2005 HPCC submission for the MTA2 [16]. In that implementation, several specialized versions of the inner loops were created in order to maximize outer-loop parallelism, take advantage of twiddle values with zero components and tune for the MTA's characteristics. While we are optimistic that our simpler implementation of the benchmark would come close to achieving the performance of this specialized version, it will always be the case that programmers willing to expend additional effort to tune for a given platform are likely to produce better performance. Optimizations such as those expressed in this implementation would be likely to improve performance for flat shared-memory contexts such as an SMP node, multicore processor, or the MTA.

On the other end of the spectrum is the distributed-memory case. Here, our implementation faces the problem that in later phases of the computation, the array slices become increasingly spread out, accessing values stored in the memories of one or more remote locales. As in the Random Access benchmark, with no optimization this is likely to result in fine-grain messages that exercise the machine's ability to handle such communications efficiently. Recent work by John Mellor-Crummey's team on the Rice D-HPF compiler strives to statically analyze loops similar to the ones in our solution so that the programmer can express the computation naturally while having the compiler generate an efficient FFT implementation [17]. In their work, they analyze loop bounds and array access patterns in order to create specialized versions of the loops for butterflies that are completely local or distributed. In the latter case, communication is overlapped with computation in order to bring remote values into a locale's memory in a way that hides the communication latency. We hope to duplicate their success in the Chapel compiler and believe that Chapel's support for domains and index types should support such analysis cleanly.

Alternatively, the Chapel programmer can explicitly manage the array's distribution in order to minimize communication. For example, by *redistributing* the domain representing the problem space from a *Block* to a *Cyclic* distribution at the appropriate stage in the computation, accesses to values that would have been remote in the later phases can be localized (given an appropriate number of processors). This approach effectively implements the common technique of computing the 1D FFT using a 2D representation that is blocked in one dimension, and then transposing it partway through the computation to localize accesses.

Of course, the 2D approach to 1D FFTs can also be coded directly in Chapel by declaring 2D domains and arrays that are distributed in one dimension and local in the other. After the local butterflies have been exhausted, the array can be transposed or the domain can be redistributed in order to localize the butterflies in the second dimension. We have coded

several of these variations in Chapel, but omit them from this article due to space limitations—such an exploration merits a paper of its own.

We conclude this discussion of FFT performance notes by echoing our opening statement: The FFT is a rich computation with many possible approaches that are sensitive to architectural details. As we target large-scale machines, our hope is that as we can achieve acceptable portable performance from our baseline implementation without requiring the algorithm to be cluttered beyond recognition. We also expect to demonstrate the ability for a user to further optimize performance for a particular architecture by performing the various rewrites described in this section.

VIII. SUMMARY AND FUTURE WORK

In summary, we have created implementations of the HPCC STREAM Triad, Random Access, and FFT benchmarks that we believe represent the computations cleanly, succinctly, and in a manner that will support good performance as our compiler development effort proceeds. As stated in the introduction, all codes listed in the appendices compile and run with our current compiler. The Chapel compiler is written in C++ and utilizes a Chapel-to-C compilation strategy for portability, allowing us to target any UNIX-like environment that supports a standard C compiler. Our compiler currently builds, compiles, and generates executables for Linux, 64-bit Linux, Cygwin, SunOS, and Mac OS X platforms. We run a nightly regression test suite of over 1,685 tests to ensure forward progress.

A limited initial release of the Chapel compiler was made available to members of the government High Productivity Language Systems (HPLS) team in December 2006. This release was designed to demonstrate many of Chapel's features in a shared-memory context and will enable experimentation with the Chapel language by Cray developers and external reviewers. Our implementation priorities for 2007 are to support multi-locale execution for distributed memory machines and to improve the performance of single-thread execution. We plan to minimize hardware and software assumptions during this phase in order to support portability of the implementation (*e.g.*, commodity Linux clusters are a primary target). We plan to release updated versions of the compiler in 2007.

In concluding this article, it is worth noting that while the HPCC benchmarks have demonstrated many of Chapel's productivity features for global-view programming and software engineering, they remain rather restricted in terms of the parallel concepts that they exercise. In particular, none of these benchmarks required any significant task parallelism, thread synchronization, or nested parallelism. Because the computations were typically driven by a single distributed problem vector, there was no need for Chapel's features for explicit locality control. And even within the data parallel space, all of these benchmarks used only dense 1D vectors of data, leaving Chapel's support for multidimensional, strided, and sparse arrays; associative and opaque arrays; and array composition on the table. In future work, we intend to undertake similar studies for computations that exercise a broader range of Chapel's features.

ACKNOWLEDGMENTS

The authors would like to thank the 2005 HPC Challenge finalists who shared their implementations with us and fielded questions about them—Petr Konecny, Nathan Wichmann, Calin Cascaval, and particularly John Feo who helped us learn a lot about FFT in a short amount of time. We would also like to recognize the following people for their efforts and impact on the Chapel language and its implementation—David Callahan, Hans Zima, John Plevyak, David Iten, Shannon Hoffswell, Roxana Diaconescu, Mark James, Mackale Joyner, and Robert Bocchino.

REFERENCES

- [1] Gail Alverson, Simon Kahan, Richard Korry, Cathy McCann, and Burton Smith. Scheduling on the Tera MTA. In D. G. Feitelson and L. Rudolph, editors, *Job Scheduling Strategies for Parallel Processing*, volume 949 of *LNCS*, pages 19–44. Springer Verlag, April 1995.
- [2] Ray Barriuso and Allan Knies. SHMEM user’s guide for C. Technical report, Cray Research Inc., June 1994.
- [3] David Callahan, Bradford L. Chamberlain, and Hans Zima. The Cascade High Productivity Language. In *Ninth International Workshop on High-Level Parallel Programming Models and Supportive Environments (HIPS’04)*, pages 52–60. April 2004.
- [4] Bradford L. Chamberlain. *The Design and Implementation of a Region-Based Parallel Language*. PhD thesis, University of Washington, November 2001.
- [5] Bradford L. Chamberlain, David Callahan, and Hans P. Zima. Parallel programmability and the Chapel language. *to appear in an upcoming special issue of the International Journal of High Performance Computing Applications on High Productivity Programming Languages and Models*, 2007.
- [6] Bradford L. Chamberlain, Sung-Eun Choi, Steven J. Deitz, and Lawrence Snyder. The high-level parallel language ZPL improves productivity and performance. In *IEEE International Workshop on Productivity and Performance in High-End Computing*, 2004.
- [7] Rohit Chandra, Ramesh Menon, Leo Dagum, David Kohr, Dror Maydan, and Jeff McDonald. *Parallel Programming in OpenMP*. Morgan Kaufmann, October 2000.
- [8] Cray Inc., Seattle, WA. *Chapel Specification*. (<http://chapel.cs.washington.edu>).
- [9] Steven J. Deitz. *High-Level Programming Language Abstractions for Advanced and Dynamic Parallel Computations*. PhD thesis, University of Washington, 2005.
- [10] Jack Dongarra and Piotr Luszczek. HPC challenge awards: Class 2 specification. Available at: <http://www.hpcchallenge.org>, June 2005.
- [11] Tarek El-Ghazawi, William Carlson, Thomas Sterling, and Katherine Yelick. *UPC: Distributed Shared-Memory Programming*. Wiley-Interscience, June 2005.
- [12] William Gropp, Steven Huss-Lederman, Andrew Lumsdaine, Ewing Lusk, Bill Nitzberg, William Saphir, and Marc Snir. *MPI: The Complete Reference, volume 2*. Scientific and Engineering Computation. MIT Press, September 1998.
- [13] Supercomputing Technologies Group. *Cilk 5.3.2 Reference Manual*. MIT Laboratory for Computer Science, November 2001.
- [14] Mackale Joyner, Steven J. Deitz, and Bradford L. Chamberlain. Iterators in Chapel. In *Eleventh International Workshop on High-Level Parallel Programming Models and Supportive Environments (HIPS’06)*. April 2006.
- [15] Charles H. Koelbel, David B. Loveman, Robert S. Schreiber, Guy L. Steele, Jr., and Mary E. Zosel. *The High Performance Fortran Handbook*. Scientific and Engineering Computation. MIT Press, September 1996.
- [16] Petr Konecny, Simon Kahan, and John Feo. SC05 HPCChallenge class 2 award—Cray MTA2. November 2005.
- [17] John Mellor-Crummey. Personal communication.
- [18] Robert W. Numerich and John Reid. Co-array Fortran for parallel programming. *SIGPLAN Fortran Forum*, 17(2):1–31, 1998.
- [19] Marc Snir, Steve Otto, Steve Huss-Lederman, David Walker, and Jack Dongarra. *MPI: The Complete Reference, volume 1*. Scientific and Engineering Computation. MIT Press, 2nd edition, September 1998.
- [20] Kathy Yelick, Luigi Semenzato, Geoff Pike, Carleton Miyamoto, Ben Liblit, Arvind Krishnamurthy, Paul Hilfinger, Susan Graham, David Gay, Phil Colella, and Alex Aiken. Titanium: A high-performance Java dialect. *Concurrency: Practice and Experience*, 10(11–13):825–836, September 1998.

APPENDIX A STREAM TRIAD CHAPEL CODE

```

1 use Time;
2 use Types;
3 use Random;

5 use HPCCProblemSize;

8 param numVectors = 3;
9 type elemType = real(64);

11 config const m = computeProblemSize(elemType, numVectors),
12             alpha = 3.0;

14 config const numTrials = 10,
15             epsilon = 0.0;

17 config const useRandomSeed = true,
18             seed = if useRandomSeed then SeedGenerator.clockMS else 314159265;

20 config const printParams = true,
21             printArrays = false,
22             printStats = true;

25 def main() {
26   printConfiguration();

28   const ProblemSpace: domain(1) distributed(Block) = [1..m];
29   var A, B, C: [ProblemSpace] elemType;

31   initVectors(B, C);

33   var execTime: [1..numTrials] real;

35   for trial in 1..numTrials {
36     const startTime = getCurrentTime();
37     A = B + alpha * C;
38     execTime(trial) = getCurrentTime() - startTime;
39   }

41   const validAnswer = verifyResults(A, B, C);
42   printResults(validAnswer, execTime);
43 }

46 def printConfiguration() {
47   if (printParams) {
48     printProblemSize(elemType, numVectors, m);
49     writeln("Number of trials = ", numTrials, "\n");
50   }
51 }

54 def initVectors(B, C) {
55   var randlist = RandomStream(seed);

57   randlist.fillRandom(B);
58   randlist.fillRandom(C);

60   if (printArrays) {
61     writeln("B is: ", B, "\n");
62     writeln("C is: ", C, "\n");
63   }
64 }

67 def verifyResults(A, B, C) {
68   if (printArrays) then writeln("A is: ", A, "\n");

70   const infNorm = max reduce [i in A.domain] abs(A(i) - (B(i) + alpha * C(i)));

72   return (infNorm <= epsilon);
73 }

```

```
76 def printResults(successful, execTimes) {
77   writeln("Validation: ", if successful then "SUCCESS" else "FAILURE");
78   if (printStats) {
79     const totalTime = + reduce execTimes,
80       avgTime = totalTime / numTrials,
81       minTime = min reduce execTimes;
82     writeln("Execution time:");
83     writeln("  tot = ", totalTime);
84     writeln("  avg = ", avgTime);
85     writeln("  min = ", minTime);
86
87     const GBPerSec = numVectors * numBytes(elemType) * (m/minTime) * 1.0e-9;
88     writeln("Performance (GB/s) = ", GBPerSec);
89   }
90 }
```

APPENDIX B RANDOM ACCESS CHAPEL CODE

A. Random Access Computation

```

1 use Time;

3 use HPCCPProblemSize;
4 use RARandomStream;

7 param numTables = 1;
8 type elemType = randType,
9     indexType = randType;

11 config const n = computeProblemSize(elemType, numTables,
12                                     returnLog2=true): indexType,
13     N_U = 2**(n+2);

15 const m = 2**n,
16     indexMask = m-1;

18 config const sequentialVerify = (numLocales < log2(m)),
19     errorTolerance = 1.0e-2;

21 config const printParams = true,
22     printArrays = false,
23     printStats = true;

26 def main() {
27     printConfiguration();

29     const TableSpace: domain(1, indexType) distributed(Block) = [0..m];
30     var T: [TableSpace] elemType;

32     const UpdateSpace: domain(1, indexType) distributed(Block) = [0..N_U];

34     const startTime = getCurrentTime();

36     [i in TableSpace] T(i) = i;

38     forall block in UpdateSpace.subBlocks do
39         for r in RASstream(block) do
40             T(r & indexMask) ^= r;

42     const execTime = getCurrentTime() - startTime;

44     const validAnswer = verifyResults(T, UpdateSpace);
45     printResults(validAnswer, execTime);
46 }

49 def printConfiguration() {
50     if (printParams) {
51         printProblemSize(elemType, numTables, m);
52         writeln("Number of updates = ", N_U, "\n");
53     }
54 }

57 def verifyResults(T: [?TDom], UpdateSpace) {
58     if (printArrays) then writeln("After updates, T is: ", T, "\n");

60     if (sequentialVerify) then
61         for r in RASstream([0..N_U]) do
62             T(r & indexMask) ^= r;
63     else
64         forall i in UpdateSpace {
65             const r = getNthRandom(i+1);
66             atomic T(r & indexMask) ^= r;
67         }

69     if (printArrays) then writeln("After verification, T is: ", T, "\n");

71     const numErrors = + reduce [i in TDom] (T(i) != i);
72     if (printStats) then writeln("Number of errors is: ", numErrors, "\n");

```

```

74  return numErrors <= (errorTolerance * N_U);
75 }

78 def printResults(successful, execTime) {
79  writeln("Validation: ", if successful then "SUCCESS" else "FAILURE");
80  if (printStats) {
81    writeln("Execution time = ", execTime);
82    writeln("Performance (GUPS) = ", N_U / execTime * 1.0e-9);
83  }
84 }

```

B. Random Access: Random Value Generation Module

```

1  module RARandomStream {
2  param randWidth = 64;
3  type randType = uint(randWidth);

5  const bitDom = [0..randWidth),
6         m2: [bitDom] randType = computeM2Vals(randWidth);

9  iterator RAStrStream(block) {
10  var val = getNthRandom(block.low);
11  for i in block {
12    getNextRandom(val);
13    yield val;
14  }
15 }

18 def getNthRandom(in n) {
19  param period = 0x7fffffffffffffff/7;

21  n %= period;
22  if (n == 0) then return 0x1;

24  var ran: randType = 0x2;
25  for i in [0..log2(n)) by -1 {
26    var val: randType = 0;
27    for j in bitDom do
28      if ((ran >> j) & 1) then val ^= m2(j);
29    ran = val;
30    if ((n >> i) & 1) then getNextRandom(ran);
31  }
32  return ran;
33 }

36 def getNextRandom(inout x) {
37  param POLY:randType = 0x7;
38  param hiRandBit = 0x1:randType << (randWidth-1);

40  x = (x << 1) ^ (if (x & hiRandBit) then POLY else 0);
41 }

44 iterator computeM2Vals(numVals) {
45  var nextVal = 0x1:randType;
46  for i in 1..numVals {
47    yield nextVal;
48    getNextRandom(nextVal);
49    getNextRandom(nextVal);
50  }
51 }
52 }

```

APPENDIX C
FFT CHAPEL CODE (1D RADIX-4)

```

1 use BitOps;
2 use Random;
3 use Time;

5 use HPCProblemSize;

8 param radix = 4;

10 param numVectors = 2;
11 type elemType = complex(128);

14 config const n = computeProblemSize(elemType, numVectors, returnLog2 = true);
15 const m = 2**n;

17 config const epsilon = 2.0 ** -51.0,
18             threshold = 16.0;

20 config const useRandomSeed = true,
21             seed = if useRandomSeed then SeedGenerator.clockMS else 314159265;

23 config const printParams = true,
24             printArrays = false,
25             printStats = true;

28 def main() {
29   printConfiguration();

31   const TwiddleDom: domain(1) distributed(Block) = [0..m/4];
32   var Twiddles: [TwiddleDom] elemType;

34   const ProblemDom: domain(1) distributed(Block) = [0..m];
35   var Z, z: [ProblemDom] elemType;

37   initVectors(Twiddles, z);

39   const startTime = getCurrentTime();

41   Z = conjg(z);
42   bitReverseShuffle(Z);
43   dfft(Z, Twiddles);

45   const execTime = getCurrentTime() - startTime;

47   const validAnswer = verifyResults(z, Z, Twiddles);
48   printResults(validAnswer, execTime);
49 }

52 def printConfiguration() {
53   if (printParams) then printProblemSize(elemType, numVectors, m);
54 }

57 def initVectors(Twiddles, z) {
58   computeTwiddles(Twiddles);
59   bitReverseShuffle(Twiddles);

61   fillRandom(z, seed);

63   if (printArrays) {
64     writeln("After initialization, Twiddles is: ", Twiddles, "\n");
65     writeln("z is: ", z, "\n");
66   }
67 }

70 def computeTwiddles(Twiddles) {
71   const numTwiddles = Twiddles.numElements,
72         delta = 2.0 * atan(1.0) / numTwiddles;

```

```

74 Twiddles(0) = 1.0;
75 Twiddles(numTwdls/2) = let x = cos(delta * numTwdls/2)
76     in (x, x):elemType;
77 forall i in [1..numTwdls/2) {
78     const x = cos(delta*i),
79           y = sin(delta*i);
80     Twiddles(i) = (x, y):elemType;
81     Twiddles(numTwdls - i) = (y, x):elemType;
82 }
83 }

86 def bitReverseShuffle(Vect: [?Dom]) {
87     const numBits = log2(Vect.numElements),
88           Perm: [i in Dom] Vect.eltType = Vect(bitReverse(i, revBits=numBits));
89     Vect = Perm;
90 }

93 def bitReverse(val: ?valType, revBits = 64) {
94     param mask = 0x0102040810204080;
95     const valReverse64 = bitMatMultOr(mask, bitMatMultOr(val:uint(64), mask)),
96           valReverse = bitRotLeft(valReverse64, revBits);
97     return valReverse: valType;
98 }

101 def dfft(A: [?ADom], W) {
102     const numElements = A.numElements;
103     var span = 1;

105     for i in [2..log2(numElements)) by 2 {
106         const m = radix*span,
107               m2 = 2*m;

109         forall (k,k1) in (ADom by m2, 0..) {
110             var wk2 = W(k1),
111                   wk1 = W(2*k1),
112                   wk3 = (wk1.re - 2 * wk2.im * wk1.im,
113                         2 * wk2.im * wk1.re - wk1.im):elemType;

115             forall j in [k..k+span) do
116                 butterfly(wk1, wk2, wk3, A[j..j+3*span by span]);

118             wk1 = W(2*k1+1);
119             wk3 = (wk1.re - 2 * wk2.re * wk1.im,
120                 2 * wk2.re * wk1.re - wk1.im):elemType;
121             wk2 *= 1.0i;

123             forall j in [k+m..k+m+span) do
124                 butterfly(wk1, wk2, wk3, A[j..j+3*span by span]);
125         }
126         span *= radix;
127     }

129     if ((span*radix) == numElements) then
130         forall j in [0..span) do
131             butterfly(1.0, 1.0, 1.0, A[j..j+3*span by span]);
132     else
133         forall j in [0..span) {
134             const a = A(j),
135                   b = A(j+span);
136             A(j) = a + b;
137             A(j+span) = a - b;
138         }
139     }

142 def butterfly(wk1, wk2, wk3, inout A:[1..radix]) {
143     var x0 = A(1) + A(2),
144           x1 = A(1) - A(2),
145           x2 = A(3) + A(4),
146           x3rot = (A(3) - A(4))*1.0i;

148     A(1) = x0 + x2;
149     x0 -= x2;
150     A(3) = wk2 * x0;

```

```
151  x0 = x1 + x3rot;
152  A(2) = wk1 * x0;
153  x0 = x1 - x3rot;
154  A(4) = wk3 * x0;
155  }

158  def verifyResults(z, Z, Twiddles) {
159    if (printArrays) then writeln("After FFT, Z is: ", Z, "\n");

161    Z = conjg(Z) / m;
162    bitReverseShuffle(Z);
163    dfft(Z, Twiddles);

165    if (printArrays) then writeln("After inverse FFT, Z is: ", Z, "\n");

167    var maxerr = max reduce sqrt((z.re - Z.re)**2 + (z.im - Z.im)**2);
168    maxerr /= (epsilon * n);
169    if (printStats) then writeln("error = ", maxerr);

171    return (maxerr < threshold);
172  }

175  def printResults(successful, execTime) {
176    writeln("Validation: ", if successful then "SUCCESS" else "FAILURE");
177    if (printStats) {
178      writeln("Execution time = ", execTime);
179      writeln("Performance (Gflop/s) = ", 5.0 * (m * n) / execTime / 1.0e-9);
180    }
181  }
```


APPENDIX D
HPCC PROBLEM SIZE COMPUTATION CODE

```

1 module HPCCProblemSize {
2   use Memory;
3   use Types;

4
5   config const memRatio = 4;

6
7   def computeProblemSize(type elemType, numArrays, returnLog2 = false) {
8     const totalMem = + reduce Locale.physicalMemory(unit = Bytes),
9           memoryTarget = totalMem / memRatio,
10          bytesPerIndex = numArrays * numBytes(elemType);

11
12    var numIndices = (memoryTarget / bytesPerIndex): int;

13
14    var lgProblemSize = log2(numIndices);
15    if (returnLog2) {
16      numIndices = 2**lgProblemSize;
17      if (numIndices * bytesPerIndex <= memoryTarget) {
18        numIndices *= 2;
19        lgProblemSize += 1;
20      }
21    }

22
23    const smallestMem = min reduce Locale.physicalMemory(unit = Bytes);
24    if ((numIndices * bytesPerIndex)/numLocales > smallestMem) then
25      halt("System is too heterogeneous: blocked data won't fit into memory");

26
27    return if returnLog2 then lgProblemSize else numIndices;
28  }

29
30
31  def printProblemSize(type elemType, numArrays, problemSize: ?psType) {
32    const bytesPerArray = problemSize * numBytes(elemType),
33          totalMemInGB = (numArrays * bytesPerArray:real) / (1024**3),
34          lgProbSize = log2(problemSize):psType;

35
36    write("Problem size = ", problemSize);
37    if (2**lgProbSize == problemSize) {
38      write(" (2**", lgProbSize, ")");
39    }
40    writeln();
41    writeln("Bytes per array = ", bytesPerArray);
42    writeln("Total memory required (GB) = ", totalMemInGB);
43  }
44 }

```