

# HPC Challenge: Five Benchmarks of Interest in Chapel\*

Bradford L. Chamberlain    Sung-Eun Choi    Vassily Litvinov  
Tom Hildebrandt    Greg Titus  
The Chapel Team  
Cray, Inc.  
chapel.info@cray.com  
Johnathan Claridge    John G. Lewis    Kristi Maschhoff  
claridge@gmail.com    jglewis@comcast.net    kristyn@cray.com

## Abstract

This report presents our best-to-date Chapel implementations of the HPC Challenge benchmarks STREAM Triad, Random Access, and HPL, the HPC Graph Analysis benchmark Scalable Synthetic Compact Applications #2 (SSCA#2), and a Adaptive Mesh Refinement (AMR) code.

The highlights of this year's submission include:

- HPCC STREAM Triad (EP and Global)
- HPCC Global Random Access
- HPCC HPL
- A graph-type independent implementation of the HPC Graph Analysis SSCA#2 benchmark version 1.2
- A rank-independent Adaptive Mesh Refinement infrastructure code

The Chapel compiler and these benchmarks are publicly available at <http://sourceforge.net/projects/chapel>.

## 1 Overview and Contents

Chapel is a new parallel programming language under development at Cray, Inc. as part of DARPA's High Productivity Computing Systems (HPCS) program. 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 three HPC Challenge (HPCC) benchmarks—STREAM Triad, Random Access, and HPL. We also present our implementation of the (SSCA#2) graph analysis benchmark and an Adaptive Mesh Refinement infrastructure code. For each benchmark, we provide a brief overview of our Chapel implementation and source line count.

## 2 Benchmark Descriptions

In this section, we give a brief description and source line counts of each of the Chapel benchmarks. All source line counts exclude comments, but include any printing or validation required by the benchmark. All three of the HPCC benchmarks share a problem size generation module.

### **HPCC Problem Size Module Source lines of code: 56**

Complete code listings are given in the appendix, with the exception of SSCA#2 and the AMR infrastructure code due to size. For SSCA#2 we include one of the graph generation algorithms and the graph-independent portions. We refer readers to the Chapel code repository for a copy of the other codes.

---

\*This material is based upon work supported by the Defense Advanced Research Projects Agency under its Agreement No. HR0011-07-9-0001.

## 2.1 HPCC EP and Global STREAM Triad

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. 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 core computation of both versions is as follows:

```
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 global STREAM implementation uses Chapel’s Block distribution module to distribute the arrays  $A$ ,  $B$ , and  $C$ .

**EP STREAM Source lines of code: 72**

**Global STREAM Source lines of code: 68**

## 2.2 HPCC Random Access

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.

The core of the Chapel implementation is as follows:

```
forall ( , r) in (Updates, RASStream()) do
  on TableDist.idxToLocale(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.

**Global RA Source lines of code: 58**

**RA Random Stream Source lines of code: 50**

## 2.3 HPCC HPL

The Chapel version of HPL uses the Chapel Block-Cyclic module for performing the dense LU factorization.

The distributed matrix-multiply for the HPL benchmarks is as follows:

```
forall (row,col) in Rest by (blkSize, blkSize) {
  const RestLcl = Rest;
  local {
    for a in (RestLcl.dim(1))(row..#blkSize) do
      for w in 1..blkSize do
        for b in (RestLcl.dim(2))(col..#blkSize) do
          Ab[a,b] -= replA[a,w] * replB[w,b];
        }
      }
  }
```

The outer *forall* loop processes each block in the matrix and performs a local matrix-multiply using *replA* and *replB*, the necessary part of the blocks that are replicated on each locale.

**Global HPL Source lines of code: 162**

## 2.4 HPC Graph Analysis SSCA#2 version 1.2

SSCA#2 is comprised of four “kernels” that require irregular accesses to a directed, weighted graph. The first kernel builds the graph representation. We have included the code for the R-MAT (Recursive Matrix) type graph generator (Kernel 1) and its support module in the appendix. Code for three other torus graph generators are available from the Chapel source repository. Kernels 2–4 are written in a completely graph-type independent manner. That is, the same kernel code is used for all graph types. The code is written using generic programming support and iterators which are implemented by the graph support module.

Chapel’s global-view programming model results in a very clean code that is also devoid of an notion of distributed data. For example, the R-MAT graph is distributed using Chapel’s Block distribution module, but Kernels 2–4 are completely unaware of this.

**SSCA#2 R-MAT Kernel 1 Source lines of code:** 154

**SSCA#2 Kernels 2–4 Source lines of code:** 326

**SSCA#2 R-MAT Support module Source lines of code:** 56

**SSCA#2 driver:** 227

## 2.5 Adaptive Mesh Refinement

Adaptive Mesh Refinement is a methodology applied to structured grids in which the problem space can be dynamically divided up into smaller grids. Adaptive Mesh Refinement was identified as one of 13 computations that represented an important class of applications for current and future parallel computing platforms <sup>1</sup>.

The AMR infrastructure code in Chapel is completely rank-independent. Specifically, the dimensionality of the grid is not visible in the code at all. Again, Chapel’s generics support and iterators enable the code to be free of all indexing expressions. The rank of the problem is a compile-time constant that can be changed for the problem of interest.

**AMR Source lines of code:** 623

**AMR Support Source lines of code:** 663

**AMR Boundary Source lines of code:** 364

**AMR Grid Source lines of code:** 322

**AMR Level Source lines of code:** 301

## 3 Preliminary Performance Results

At the time of this writing, we are still working to get performance numbers on an Cray XE6 <sup>TM</sup>with 24-core nodes (dual 12-core AMD processors running at 2.0GHz). We are still tuning our Chapel codes, but we are also experimenting with a new (pre-release) version of the GASNet networking library that has native support for the Cray XE6 network. The release version of GASNet should be available by the date of SC11. Here we present preliminary performance results.

For EP STREAM, we ran on up to 32 nodes for both the Chapel version and the HPCC reference version 1.4.1. As the number of nodes increases, we see the Chapel version within 5% of the reference version.

For Random Access, we have been struggling with running the reference version on the Cray XE6 and have yet to collect comparison numbers. We did initial experiments on Jaguar, a Cray XT5 <sup>TM</sup>system which showed the Chapel version to perform about 2 times slower than the reference version (our 2009 submission was four times slower).

For HPL, we do not expect to be very competitive with the reference version because we have spent most of our time writing and tuning a general Chapel *domain map* (distribution) that is used in the distributed matrix multiplication section (shown

---

<sup>1</sup><http://www.eecs.berkeley.edu/Pubs/TechRpts/2006/EECS-2006-183.pdf>

nodes	Chapel EP STREAM	HPCC EP STREAM
1	38.2	39.1
2	75.3	78.3
4	151.8	156.5
8	302.7	313.2
16	606.6	626.2
32	1205.5	1251.8

Table 1: EP STREAM Triad performance for Chapel and HPCC v1.4.1 in GB/s.

above). Due to our efforts, this version is now dominated by the scalar computation. For the final submission, we may compare these portions of the HPL benchmark rather than the entire benchmark.

For SSCA#2, though this is our initial implementation and we have no reference version to compare to, we have been able to run SSCA#2 with a 8192 node R-MAT graph with maximum out degree of 126 on 20 nodes. Kernel 4, which computes *betweenness centrality*, is typically of the most interest and thus far we have spent most of our time on it. We achieved 371735.0 TEPS (Traversed Edges Per Second) for Kernel 4.

At higher node counts (which go up by a factor of two with each “scale”), our graph generation portion (Kernel 1) consumed an unreasonable amount of time, and we opted to stop the runs before getting through to Kernel 4. We now plan to focus more attention on Kernel 1.

## Acknowledgments

The authors would like to gratefully acknowledge our former team members and collaborators who assisted with our 2006, 2008, and 2009 entries in the HPCC competition: Steve Deitz, Samuel Figueroa, Mary Beth Hribar, David Iten, 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.

# A EP STREAM Triad

```

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;
23 //
24 //
25 // Use shared user module for computing HPC problem sizes
26 //
27 use HPCProblemSize;
28 //
29 //
30 // The number of vectors and element type of those vectors
31 //
32 const numVectors = 3;
33 type elemType = real(64);
34 //
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;
45 //
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;
52 //
53 //
54 // Configuration constants to indicate whether or not to use a
55 // pseudo-random seed (based on the clock) or a fixed seed; and to
56 // specify the fixed seed explicitly
57 //
58 config const useRandomSeed = true,
59 seed = if useRandomSeed then SeedGenerator.currentTime else 3141592653;
60 //
61 //
62 // *** To ensure determinism of output, there is no more printing of the
63 // *** arrays in initVectors and verifyResults.
64 //
65 //
66 // Configuration constants to control what's printed -- benchmark
67 // parameters and/or statistics
68 //
69 config const printParams = true,
70 printStats = true;
71 //
72 //
73 // The program entry point
74 //
75 proc main() {
76   printConfiguration(); // print the problem size, number of trials, etc.
77 //
78 //
79 // *** Aggregates for collecting per-locale results for the minimum
80 // *** execution time per trial, and whether verification passed
81 //
82 var minTimes: [LocaleSpace] real;
83 var validAnswers: [LocaleSpace] bool;
84 //
85 //
86 // *** Fragment control so that we have a single task running on
87 // *** every locale.
88 //
89 coforall loc in Locales do on loc {
90 //
91 //
92 // *** We declare these variables outside of the local block since
93 // *** we'll need to access them when we write back to the global
94 // *** aggregates declared above.
95 //
96 var validAnswer: bool;
97 var execTime: [1..numTrials] real;
98 //
99 //
100 // *** Indicates that all of the code in this block is local to
101 // *** this locale. There is no communication. A violation will
102 // *** result in an error, though error checking is disabled with
103 // *** --fast or --no-checks.
104 //
105 local {
106 //
107 //
108 // *** A, B, and C are the three local vectors
109 //
110 var A, B, C: [1..m] elemType;
111 //
112 initVectors(B, C); // Initialize the input vectors, B and C
113 //
114 for trial in 1..numTrials { // loop over the trials
115   const startTime = getCurrentTime(); // capture the start time
116 //
117 //
118 // *** The main loop looks identical to stream.chpl. However,
119 // *** in this version we are iterating over arrays that are
120 // *** not distributed.
121 //
122 forall (a, b, c) in (A, B, C) do
123   a = b + alpha * c;
124 //
125   execTime(trial) = getCurrentTime() - startTime; // store the elapsed time
126 }
127 //
128 validAnswer = verifyResults(A, B, C); // verify...
129 }
130 //
131 //
132 // *** Write times and verification result into aggregates
133 // *** declared above. These are declared over LocaleSpace so we
134 // *** can write to them in parallel.
135 //
136 minTimes[here.id] = min reduce execTime;
137 validAnswers[here.id] = validAnswer;
138 }
139 //
140 //
141 // *** Pass minimum, average, and maximum times to printResults
142 //
143 printResults(%& reduce validAnswers, minTimes);
144 }
145 //
146 //
147 // Print the problem size and number of trials
148 //
149 proc printConfiguration() {
150   if (printParams) {
151     //
152     // *** Here we multiply m by the number of locales so that we can
153     // *** print out the global problem size.
154     //
155     printProblemSize(elemType, numVectors, m * numLocales);
156     writeln("Number of trials = ", numTrials, "\n");
157   }
158 }
159 //
160 //
161 // *** Both initVectors and verifyResults are almost identical to
162 // *** stream.chpl even though they are called with arrays that are
163 // *** not distributed. For initialization, the same random stream is
164 // *** used on each locale. In the global version, a single logical
165 // *** stream of random numbers is used across all of the locales.
166 // ***
167 // *** In this version, we've omitted a way to print the arrays. This
168 // *** ensures determinism of output. Printing the arrays also
169 // *** violates the locality constraint imposed by the local block
170 // *** from which these functions are called.
171 //
172 // Initialize vectors B and C using a random stream of values
173 //
174 proc initVectors(B, C) {
175   var randlist = new RandomStream(seed);
176 //
177   randlist.fillRandom(B);
178   randlist.fillRandom(C);
179 //
180   delete randlist;
181 }
182 //
183 //
184 // Verify that the computation is correct
185 //
186 proc verifyResults(A, B, C) {
187 //
188 //
189 // recompute the computation, destructively storing into B to save space
190 //
191 forall (b, c) in (B, C) do
192   b += alpha * c;
193 //
194 //
195 // Compute the infinity-norm by computing the maximum reduction of the
196 // absolute value of A's elements minus the new result computed in B.
197 // "[i in I]" represents an expression-level loop: "forall i in I"
198 //
199 const infNorm = max reduce [(a,b) in (A,B)] abs(a - b);
200 //
201 return (infNorm <= epsilon); // return whether the error is acceptable
202 }
203 //
204 //
205 // Print out success/failure, the timings, and the GB/s value.
206 //
207 // *** Here we report maximum, average, and minimum times instead of
208 // *** total, average, and minimum.
209 //
210 proc printResults(successful, minTimes) {
211   writeln("Validation: ", if successful then "SUCCESS" else "FAILURE");
212   if (printStats) {
213     const maxTime = max reduce minTimes,
214     avgTime = + reduce minTimes / numLocales,

```

```

215         minTime = min reduce minTimes;
216         writeln("Execution time:");
217         writeln("  max (seconds) = ", maxTime);
218         writeln("  avg (seconds) = ", avgTime);
219         writeln("  min (seconds) = ", minTime);

221     const maxGBPerSec = numVectors * numBytes(elemType) * (m / minTime) * 1e-9,
222           avgGBPerSec = numVectors * numBytes(elemType) * (m / avgTime) * 1e-9,
223           minGBPerSec = numVectors * numBytes(elemType) * (m / maxTime) * 1e-9;
224     writeln("Performance (GB/s):");
225     writeln("  max = ", maxGBPerSec);
226     writeln("  avg = ", avgGBPerSec);
227     writeln("  min = ", minGBPerSec);
228   }
229 }

```

## B Global STREAM Triad

```
1 //
2 // Use standard modules for Block distributions, Timing routines, Type
3 // utility functions, and Random numbers
4 //
5 use BlockDist, Time, Types, Random;

7 //
8 // Use shared user module for computing HPC problem sizes
9 //
10 use HPCProblemSize;

12 //
13 // The number of vectors and element type of those vectors
14 //
15 const numVectors = 3;
16 type elemType = real(64);

18 //
19 // Configuration constants to set the problem size (m) and the scalar
20 // multiplier, alpha
21 //
22 const const m = computeProblemSize(numVectors, elemType),
23         alpha = 3.0;

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 const const numTrials = 10,
30         epsilon = 0.0;

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 const const useRandomSeed = true,
38         seed = if useRandomSeed then SeedGenerator.currentTime else 31415926539;

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

48 //
49 // The program entry point
50 //
51 proc main() {
52   printConfiguration(); // print the problem size, number of trials, etc.

54   //
55   // ProblemSpace describes the index set for the three vectors. It
56   // is a 1D domain storing 64-bit ints and is distributed according
57   // to a Block distribution. In this case, the Block distribution is 1D
58   // distribution computed by blocking the bounding box 1..m across the set
59   // of locales. The ProblemSpace domain contains the indices 1..m.
60   //
61   const ProblemSpace:
62     domain(1, int(64)) dmapped Block(boundingBox=[1..m]) = [1..m];

64   //
65   // A, B, and C are the three distributed vectors, declared to store
66   // a variable of type elemType for each index in ProblemSpace.
67   //
68   var A, B, C: [ProblemSpace] elemType;

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

72   var execTime: [1..numTrials] real; // an array of timings

74   for trial in 1..numTrials { // loop over the trials
75     const startTime = getCurrentTime(); // capture the start time

77     //
78     // The main loop: Iterate over the vectors A, B, and C in a
79     // parallel, zippered manner storing the elements as a, b, and c.
80     // Compute the multiply-add on b and c, storing the result to a.
81     //

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

85     execTime(trial) = getCurrentTime() - startTime; // store the elapsed time
86   }

88   const validAnswer = verifyResults(A, B, C); // verify...
89   printResults(validAnswer, execTime); // ...and print the results
90 }

92 //
93 // Print the problem size and number of trials
94 //
95 proc printConfiguration() {
96   if (printParams) {
97     if (printStats) then printLocalesTasks();
98     printProblemSize(elemType, numVectors, m);
99     writeln("Number of trials = ", numTrials, "\n");
100  }
101 }

103 //
104 // Initialize vectors B and C using a random stream of values and
105 // optionally print them to the console
106 //
107 proc initVectors(B, C) {
108   var randlist = new RandomStream(seed);

110   randlist.fillRandom(B);
111   randlist.fillRandom(C);

113   if (printArrays) {
114     writeln("B is: ", B, "\n");
115     writeln("C is: ", C, "\n");
116   }

118   delete randlist;
119 }

121 //
122 // Verify that the computation is correct
123 //
124 proc verifyResults(A, B, C) {
125   if (printArrays) then writeln("A is: ", A, "\n"); // optionally print A

127   //
128   // recompute the computation, destructively storing into B to save space
129   //
130   forall (b, c) in (B, C) do
131     b += alpha * c;

133   if (printArrays) then writeln("A-hat is: ", B, "\n"); // and A-hat too

135   //
136   // Compute the infinity-norm by computing the maximum reduction of the
137   // absolute value of A's elements minus the new result computed in B.
138   // "[i in I]" represents an expression-level loop: "forall i in I"
139   //
140   const infNorm = max reduce [(a,b) in (A,B)] abs(a - b);

142   return (infNorm <= epsilon); // return whether the error is acceptable
143 }

145 //
146 // Print out success/failure, the timings, and the GB/s value
147 //
148 proc printResults(successful, execTimes) {
149   writeln("Validation: ", if successful then "SUCCESS" else "FAILURE");
150   if (printStats) {
151     const totalTime = + reduce execTimes,
152           avgTime = totalTime / numTrials,
153           minTime = min reduce execTimes;
154     writeln("Execution time:");
155     writeln("  tot = ", totalTime);
156     writeln("  avg = ", avgTime);
157     writeln("  min = ", minTime);

159     const GBPerSec = numVectors * numBytes(elemType) * (m / minTime) * 1e-9;
160     writeln("Performance (GB/s) = ", GBPerSec);
161   }
162 }
```

# C Global Random Access

## 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 HPCC problem sizes and for
8 // defining RA's random stream of values
9 //
10 use HPCCProblemSize, 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 // Configuration constants to control what's printed -- benchmark
43 // parameters, input and output arrays, and/or statistics
44 //
45 config const printParams = true,
46     printArrays = false,
47     printStats = true;

49 //
50 // TableDist is a 1D block distribution for domains storing indices
51 // of type "indexType", and it is computed by blocking the bounding
52 // box 0..m-1 across the set of locales. UpdateDist is a similar
53 // distribution that is computed by blocking the indices 0..N_U-1
54 // across the locales.
55 //
56 const TableDist = new dmap(new Block(boundingBox=[0..m-1])),
57     UpdateDist = new dmap(new Block(boundingBox=[0..N_U-1]));

59 //
60 // TableSpace describes the index set for the table. It is a 1D
61 // domain storing indices of type indexType, it is distributed
62 // according to TableDist, and it contains the indices 0..m-1.
63 // Updates is an index set describing the set of updates to be made.
64 // It is distributed according to UpdateDist and contains the
65 // indices 0..N_U-1.
66 //
67 const TableSpace: domain(1, indexType) dmapped TableDist = [0..m-1],
68     Updates: domain(1, indexType) dmapped UpdateDist = [0..N_U-1];

71 //
72 // The program entry point
73 //
74 proc main() {
75 //
76 // T is the distributed table itself, storing a variable of type
77 // elemType for each index in TableSpace.
78 //
79 var T: [TableSpace] elemType;

81 printConfiguration(); // print the problem size, number of trials, etc.

83 //
84 // In parallel, initialize the table such that each position
85 // contains its index. "[i in TableSpace]" is shorthand for "forall
86 // i in TableSpace"

87 //
88 [i in TableSpace] T(i) = i;

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

92 //
93 // The main computation: Iterate over the set of updates and the
94 // stream of random values in a parallel, zippered manner, dropping
95 // the update index on the ground and storing the random value
96 // in r. Use an on-clause to force the table update to be executed on
97 // the locale which owns the table element in question to minimize
98 // communications. Compute the update using r both to compute the
99 // index and as the update value.
100 //
101 forall ( , r) in (Updates, RAStrm()) do
102 on TableDist.idxToLocale(r & indexMask) do {
103     const myR = r;
104     local {
105         T(myR & indexMask) ^= myR;
106     }
107 }

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

111 const validAnswer = verifyResults(T); // verify the updates
112 printResults(validAnswer, execTime); // print the results
113 }

115 //
116 // Print the problem size and number of updates
117 //
118 proc printConfiguration() {
119     if (printParams) {
120         if (printStats) then printLocalesTasks();
121         printProblemSize(elemType, numTables, m);
122         writeln("Number of updates = ", N_U, "\n");
123     }
124 }

126 //
127 // Verify that the computation is correct
128 //
129 proc verifyResults(T) {
130 //
131 // Print the table, if requested
132 //
133 if (printArrays) then writeln("After updates, T is: ", T, "\n");

135 //
136 // Reverse the updates by recomputing them, this time using an
137 // atomic statement to ensure no conflicting updates
138 //
139 forall ( , r) in (Updates, RAStrm()) do
140 on TableDist.idxToLocale(r & indexMask) do
141     atomic T(r & indexMask) ^= r;

143 //
144 // Print the table again after the updates have been reversed
145 //
146 if (printArrays) then writeln("After verification, T is: ", T, "\n");

148 //
149 // Compute the number of table positions that weren't reverted
150 // correctly. This is an indication of the number of conflicting
151 // updates.
152 //
153 const numErrors = + reduce [i in TableSpace] (T(i) != i);
154 if (printStats) then writeln("Number of errors is: ", numErrors, "\n");

156 //
157 // Return whether or not the number of errors was within the benchmark's
158 // tolerance.
159 //
160 return numErrors <= (errorTolerance * N_U);
161 }

163 //
164 // Print out success/failure, the execution time, and the GUPS value
165 //
166 proc printResults(successful, execTime) {
167     writeln("Validation: ", if successful then "SUCCESS" else "FAILURE");
168     if (printStats) {
169         writeln("Execution time = ", execTime);
170         writeln("Performance (GUPS) = ", (N_U / execTime) * 1e-9);
171     }
172 }
```

## 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
8 }

9 //

10 // m2 is a table (tuple) of helper values used to fast-forward
11 // through the random stream.
12 //
13 const m2: randWidth*randType = computeM2Vals();

15 //
16 // A serial iterator for the random stream that resets the stream
17 // to its 0th element and yields values endlessly.
18 //
```



```

19 iter RASStream() {
20   var val = getNextRandom(0);
21   while (1) {
22     getNextRandom(val);
23     yield val;
24   }
25 }

27 //
28 // A "follower" iterator for the random stream that takes a range of
29 // 0-based indices (follower) and yields the pseudo-random values
30 // corresponding to those indices. Follower iterators like these
31 // are required for parallel zippered iteration.
32 //
33 iter RASStream(param tag: iterKind, followThis) where tag == iterKind.follower {
34   if followThis.size != 1 then
35     halt("RASStream cannot use multi-dimensional iterator");
36   var val = getNextRandom(followThis(1).low);
37   for followThis {
38     getNextRandom(val);
39     yield val;
40   }
41 }

43 //
44 // A helper function for "fast-forwarding" the random stream to
45 // position n in O(log2(n)) time
46 //
47 proc getNextRandom(in n: uint(64)) {
48   param period = 0x7fffffff7;
49
50   n %= period;
51   if (n == 0) then return 0x1;
52   var ran: randType = 0x2;
53   for i in 0..log2(n)-1 by -1 {

54     var val: randType = 0;
55     for j in 0..randWidth do
56       if ((ran >> j) & 1) then val ^= m2(j+1);
57       ran = val;
58       if ((n >> i) & 1) then getNextRandom(ran);
59     }
60     return ran;
61 }

63 //
64 // A helper function for advancing a value from the random stream,
65 // x, to the next value
66 //
67 proc getNextRandom(inout x) {
68   param POLY = 0x7;
69   param hiRandBit = 0x1:randType << (randWidth-1);
70
71   x = (x << 1) ^ (if (x & hiRandBit) then POLY else 0);
72 }

74 //
75 // A helper function for computing the values of the helper tuple, m2
76 //
77 proc computeM2Vals() {
78   var m2tmp: randWidth*randType;
79   var nextVal = 0x1: randType;
80   for i in 1..randWidth {
81     m2tmp(i) = nextVal;
82     getNextRandom(nextVal);
83     getNextRandom(nextVal);
84   }
85   return m2tmp;
86 }
87 }

```

# D Global HPL

```
1 //
2 // Use standard modules for vector and matrix Norms, Random numbers
3 // and Timing routines
4 //
5 use Norm, Random, Time;
6
7 //
8 // Use the user module for computing HPCC problem sizes
9 //
10 use HPCCProblemSize;
11
12 //
13 // Use the distributions we need for this computation
14 //
15 use d, r, f;
16
17 //
18 // The number of matrices and the element type of those matrices
19 // indexType can be int or int(64), elemType can be real or complex
20 //
21 const numMatrices = 1;
22 type indexType = int,
23      elemType = real;
24
25 //
26 // Configuration constants indicating the problem size (n) and the
27 // block size (blkSize)
28 //
29 config const n = computeProblemSize(numMatrices, elemType, rank=2,
30                                     memFraction=2, retType=indexType),
31      blkSize = 8;
32
33 //
34 // Configuration constant used for verification thresholds
35 //
36 config const epsilon = 2.0e-15;
37
38 //
39 // Configuration constants to indicate whether or not to use a
40 // pseudo-random seed (based on the clock) or a fixed seed; and to
41 // specify the fixed seed explicitly
42 //
43 config const useRandomSeed = true,
44      seed = if useRandomSeed then SeedGenerator.currentTime else 31415;
45
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;
53
54 // These are solely to make the testing system happy given the COMPOPTS file.
55 // To be removed once COMPOPTS becomes a non-issue.
56 config var reproducible = false, verbose = false;
57
58 //
59 // The program entry point, currently module initialization
60 //
61 printConfiguration();
62
63 //
64 // Compute targetLocales - required for Dimensional.
65 // We hard-code 2 dimensions.
66 //
67 var targetIds: domain(2);
68 var targetLocales: [targetIds] locale;
69 setupTargetLocalesArray(targetIds, targetLocales, Locales);
70
71 // Here are the dimensions of our grid of locales.
72 const t11 = targetIds.dim(1).length,
73      t12 = targetIds.dim(2).length;
74 if printParams && printStats then
75     writeln("target locales ", t11, " x ", t12);
76
77 // Create the dimensional descriptors
78 const
79     bdim1 = new idist(lowIdx=1, blockSize=blkSize, numLocales=t11),
80     rdim1 = new vdist(t11),
81
82     bdim2 = new idist(lowIdx=1, blockSize=blkSize, numLocales=t12),
83     rdim2 = new vdist(t12);
84
85 //
86 // MatVectSpace is a 2D domain of type indexType that represents the
87 // n x n matrix adjacent to the column vector b. MatrixSpace is a
88 // subdomain that is created by slicing into MatVectSpace,
89 // inheriting all of its rows and its low column bound. As our
90 // standard distribution library is filled out, MatVectSpace will be
91 // distributed using a BlockCyclic(blkSize) distribution.
92 //
93 // We use 'AbD' instead of 'MatVectSpace' throughout.
94 //
95 const AbD: domain(2, indexType)
96     dmapped DimensionalDist(targetLocales, bdim1, bdim2, "dim")
97     = [1..n, 1..n+1],
98     MatrixSpace = AbD[...];
99
100 var Ab : [AbD] elemType, // the matrix A and vector b
101     piv : [1..n] indexType; // a vector of pivot values
102
103 //
104 // Create the 1-d replicated arrays for schurComplement().
105 //
106 const
107     replAD = [1..n, 1..blkSize]
```

```
108     dmapped DimensionalDist(targetLocales, bdim1, rdim2, "distBR"),
109     replBD = [1..blkSize, 1..n+1]
110     dmapped DimensionalDist(targetLocales, rdim1, bdim2, "distRB");
111
112 var replA: [replAD] elemType,
113     replB: [replBD] elemType;
114
115 initAB();
116
117 const startTime = getCurrentTime(); // capture the start time
118
119 LUFactorize(n, piv); // compute the LU factorization
120
121 var x = backwardSub(n); // perform the back substitution
122
123 const execTime = getCurrentTime() - startTime; // store the elapsed time
124
125 //
126 // Validate the answer and print the results
127 const validAnswer = verifyResults(x);
128 printResults(validAnswer, execTime);
129
130 //
131 // blocked LU factorization with pivoting for matrix augmented with
132 // vector of RHS values.
133 //
134 proc LUFactorize(n: indexType,
135                 piv: [1..n] indexType) {
136
137     // Initialize the pivot vector to represent the initially unpivoted matrix.
138     piv = 1..n;
139
140     /* The following diagram illustrates how we partition the matrix.
141     Each iteration of the loop increments a variable blk by blkSize;
142     point (blk, blk) is the upper-left location of the currently
143     unfactored matrix (the dotted region represents the areas
144     factored in prior iterations). The unfactored matrix is
145     partitioned into four subdomains: t1, tr, b1, and br, and an
146     additional domain (not shown), l, that is the union of t1 and b1.
147
148     (point blk, blk)
149     +-----+
150     |.....|
151     |.....|
152     |.....|
153     |.....|
154     |.....| t1 |      tr      |
155     |.....|
156     |.....|
157     |.....|
158     |.....|
159     |.....| b1 |      br      |
160     |.....|
161     |.....|
162     +-----+
163
164     */
165     for blk in 1..n by blkSize {
166         if printStats then writeln("processing block ", blk);
167
168         const t1 = AbD[blk..#blkSize, blk..#blkSize],
169             tr = AbD[blk..#blkSize, blk+blkSize..],
170             b1 = AbD[blk+blkSize.., blk..#blkSize],
171             br = AbD[blk+blkSize.., blk+blkSize..],
172             l = AbD[blk.., blk..#blkSize];
173
174         //
175         // Now that we've sliced and diced Ab properly, do the blocked-LU
176         // computation:
177         //
178         panelSolve(l, piv);
179         updateBlockRow(t1, tr);
180
181         //
182         // update trailing submatrix (if any)
183         //
184         schurComplement(b1, tr, br);
185     }
186
187 //
188 // Distributed matrix-multiply for HPL. The idea behind this algorithm is that
189 // some point the matrix will be partitioned as shown in the following diagram:
190 //
191 //
192 // [1]-----+
193 // |          |bbbb|bbbb|bbbb| Solve for the dotted region by
194 // |          |bbbb|bbbb|bbbb| multiplying the 'a' and 'b' region.
195 // |          |bbbb|bbbb|bbbb| The 'a' region is a block column, the
196 // +-----+-----+-----+ 'b' region is a block row.
197 // |aaaaa|.....|.....|.....|
198 // |aaaaa|.....|.....|.....| The 'a' region was 'b1' in the calling
199 // |aaaaa|.....|.....|.....| function but called AD here. Similarly,
200 // +-----+-----+-----+ 'b' was 'tr' in the calling code, but BD
201 // |aaaaa|.....|.....|.....| here.
202 // |aaaaa|.....|.....|.....|
203 // |aaaaa|.....|.....|.....|
204 // +-----+-----+-----+
205 //
206 // Every locale with a block of data in the dotted region updates
207 // itself by multiplying the neighboring a-region block to its left
208 // with the neighboring b-region block above it and subtracting its
209 // current data from the result of this multiplication. To ensure that
210 // all locales have local copies of the data needed to perform this
211 // multiplication we copy the data A and B data into the replA and
212 // replB arrays, which will use a dimensional (block-cyclic,
213 // replicated-block) distribution (or vice-versa) to ensure that every
214 // locale only stores one copy of each block it requires for all of
215 // its rows/columns.
```

```

215 //
216 proc schurComplement(AD: domain, BD: domain, Rest: domain) {
217
218 // Prevent replication of unequal-sized slices
219 if Rest.numIndices == 0 then return;
220
221 //
222 // Copy data into replicated arrays so every processor has a local copy
223 // of the data it will need to perform a local matrix-multiply.
224 //
225 coforall dest in targetLocales[targetIds.dim(1).high, targetIds.dim(2)] do
226 on dest do
227 // replA on tgLocales[d1,i] gets a copy of Ab from tgLocales[d1,..]
228 replA = Ab[1..n, AD.dim(2)];
229
230 coforall dest in targetLocales[targetIds.dim(1), targetIds.dim(2).high] do
231 on dest do
232 // replB on tgLocales[i,d2] gets a copy of Ab from tgLocales[..,d2]
233 replB = Ab[BD.dim(1), 1..n+1];
234
235 // do local matrix-multiply on a block-by-block basis
236 forall (row,col) in Rest by (blkSize, blkSize) {
237 // Workaround: localize Rest explicitly
238 const RestLcl = Rest;
239
240 local {
241 for a in (RestLcl.dim(1)) (row..#blkSize) do
242 for w in 1..blkSize do
243 for b in (RestLcl.dim(2)) (col..#blkSize) do
244 Ab[a,b] -= replA[a,w] * replB[w,b];
245 }
246 }
247 }
248
249 //
250 // do unblocked-LU decomposition within the specified panel, update the
251 // pivot vector accordingly
252 //
253 proc panelSolve(
254 panel: domain,
255 piv: [] indexType) {
256
257 for k in panel.dim(2) { // iterate through the columns
258 const col = panel[k.., k..k];
259
260 // If there are no rows below the current column return
261 if col.numIndices == 0 then return;
262
263 // Find the pivot, the element with the largest absolute value.
264 const ( , (pivotRow, )) = maxloc reduce(abs(Ab(col)), col);
265
266 // Capture the pivot value explicitly (note that result of maxloc
267 // is absolute value, so it can't be used directly).
268 //
269 const pivotVal = Ab[pivotRow, k];
270
271 // Swap the current row with the pivot row and update the pivot vector
272 // to reflect that
273 Ab[k..k, ..] <=> Ab[pivotRow..pivotRow, ..];
274 piv[k] <=> piv[pivotRow];
275
276 if (pivotVal == 0) then
277 halt("Matrix cannot be factorized");
278
279 // divide all values below and in the same col as the pivot by
280 // the pivot value
281 Ab[k+1.., k..k] /= pivotVal;
282
283 // update all other values below the pivot
284 forall (i,j) in panel[k+1.., k+1..] do
285 Ab[i,j] -= Ab[i,k] * Ab[k,j];
286 }
287 }
288
289 //
290 // Update the block row (tr for top-right) portion of the matrix in a
291 // blocked LU decomposition. Each step of the LU decomposition will
292 // solve a block (t1 for top-left) portion of a matrix. This function
293 // solves the rows to the right of the block.
294 //
295 proc updateBlockRow(
296 t1: domain,
297 tr: domain) {
298
299 for row in tr.dim(1) {
300 const activeRow = tr[row..row, ..],
301 prevRows = tr.dim(1).low..row-1;
302
303 forall (i,j) in activeRow do
304 for k in prevRows do
305 Ab[i, j] -= Ab[i, k] * Ab[k, j];
306 }
307 }
308
309 //
310 // compute the backwards substitution
311 //
312 proc backwardSub(n: indexType) {
313 const bd = Ab.domain.dim(1); // or simply 1..n
314 var x: [bd] elemType;
315
316 for i in bd by -1 do
317 x[i] = (Ab[i,n+1] - (+ reduce [j in i+1..bd.high] (Ab[i,j] * x[j])))
318 / Ab[i,i];
319
320 return x;
321 }
322
323 //
324 // print out the problem size and block size if requested
325 //
326 proc printConfiguration() {
327 if (printParams) {
328 if (printStats) then printLocalesTasks();
329 printProblemSize(elemType, numMatrices, n, rank=2);
330 writeln("block size = ", blkSize, "\n");
331 }
332 }
333
334 //
335 // construct an n by n+1 matrix filled with random values and scale
336 // it to be in the range -1.0..1.0
337 //
338 proc initAB() {
339 fillRandom(Ab, seed);
340 Ab = Ab * 2.0 - 1.0;
341 }
342
343 //
344 // calculate norms and residuals to verify the results
345 //
346 proc verifyResults(x) {
347 initAB();
348
349 const axmbNorm = norm(gaxpyMinus(Ab[... 1..n], x, Ab[... n+1..n+1]), normType.normInf);
350
351 const alnorm = norm(Ab[... 1..n], normType.norm1),
352 aInfNorm = norm(Ab[... 1..n], normType.normInf),
353 x1Norm = norm(Ab[... n+1..n+1], normType.norm1),
354 xInfNorm = norm(Ab[... n+1..n+1], normType.normInf);
355
356 const resid1 = axmbNorm / (epsilon * alnorm * n),
357 resid2 = axmbNorm / (epsilon * alnorm * x1Norm),
358 resid3 = axmbNorm / (epsilon * aInfNorm * xInfNorm);
359
360 if (printStats) {
361 writeln("resid1: ", resid1);
362 writeln("resid2: ", resid2);
363 writeln("resid3: ", resid3);
364 }
365
366 return max(resid1, resid2, resid3) < 16.0;
367 }
368
369 //
370 // print success/failure, the execution time and the Gflop/s value
371 //
372 proc printResults(successful, execTime) {
373 writeln("Validation: ", if successful then "SUCCESS" else "FAILURE");
374 if (printStats) {
375 writeln("Execution time = ", execTime);
376 const GflopPerSec = ((2.0/3.0) * n**3 + (3.0/2.0) * n**2) / execTime * 10e-9;
377 writeln("Performance (Gflop/s) = ", GflopPerSec);
378 }
379 }
380
381 //
382 // simple matrix-vector multiplication, solve equation A*x-y
383 //
384 proc gaxpyMinus(A: [],
385 x: [?xD],
386 y: [?yD]) {
387 var res: [1..n] elemType;
388
389 forall i in 1..n do
390 res[i] = (+ reduce [j in xD] (A[i,j] * x[j])) - y[i,n+1];
391
392 return res;
393 }

```

## E Shared Problem Size Module Code for HPCC Benchmarks

```
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 proc computeProblemSize(numArrays: int, // #arrays in the benchmark
15 type 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 proc printLocalesTasks() {
77 writeln("Number of Locales = ", numLocales);
78 writeln("Tasks per locale = ", dataParTasksPerLocale);
79 }
80
81 //
82 // Print out the problem size, #bytes per array, and total memory
83 // required by the arrays
84 //
85 proc printProblemSize(type elemType, numArrays, problemSize: ?psType,
86 param rank=1, problemSize2=problemSize) {
87 const bytesPerArray = problemSize**rank * numBytes(elemType),
88 totalMemInGB = (numArrays * bytesPerArray:real) / (1024**3),
89 lgProbSize = log2(problemSize):psType,
90 lgProbSize2 = log2(problemSize2):psType;
91
92 write("Problem size = ", problemSize);
93 for i in 2..rank do write(" x ", problemSize2);
94 if (2**lgProbSize == problemSize && 2**lgProbSize2 == problemSize2) {
95 write(" (2**", lgProbSize);
96 for i in 2..rank do write(" x 2**", lgProbSize2);
97 write(")");
98 }
99 writeln();
100 writeln("Bytes per array = ", bytesPerArray);
101 writeln("Total memory required (GB) = ", totalMemInGB);
102 }
103 }
```

## F SSCA#2 R-MAT graph creation (Kernel 1)

```

1 module SSCA2_RMAT_graph_generator
2
3 {
4 // +-----+
5 // |          RMAT approximate power law graph generator          |
6 // |          and                                               |
7 // |          SSCA #2 Kernel 1                                 |
8 // |          |
9 // |          Generate approximate power law graph              |
10 // |          |
11 // |          This procedure accepts a generic graph representation which must
12 // |          provide only the capabilities to:
13 // |          1. add a vertex to a neighbor list for some other vertex
14 // |          2. assign an integer weight to an entry in an associated weight list
15 // |          |
16 // |          The RMAT procedure implicitly assumes that vertices are integers
17 // |          in the range [0, 2^SCALE]. So vertices are integers in this
18 // |          procedure and edges are pairs of integers.
19 // |          |
20 // |          This implementation first generates the RMAT graph as a list of
21 // |          triples, using an array of edges and an associated array of weights.
22 // |          These are transformed into the Chapel graph representation using
23 // |          native Chapel syntax for associative and sparse domains and arrays.
24 // |          That transformation is Kernel 1 of the SSCA #2 benchmark.
25 // |          |
26 // |          The code uses two auxiliary procedures for clarity and to help the
27 // |          compiler use more abstract, higher level, code.
28 // +-----+
29
30 use SSCA2_compilation_config_params;
31
32 record directed_vertex_pair {
33   var start = 1: int;
34   var end   = 1: int;
35 }
36
37 proc +(l: directed_vertex_pair, r: directed_vertex_pair)
38   return new directed_vertex_pair (l.start + r.start, l.end + r.end);
39
40 // =====
41 // Quadrant selection algorithm
42 // =====
43
44 proc assign_quadrant ( u : real, a : real, b : real, c : real, d : real,
45                      bit : int ) : directed_vertex_pair
46 {
47   // -----
48   // The heart of the RMAT random graph generator is a procedure that
49   // assigns a single bit of the edge starting and ending vertex by
50   // assigning the edge with specified probability to one of the four
51   // quadrants of a 2 x 2 grid.
52   //
53   // Chapel is able to promote this conditional-based scalar procedure to
54   // array operations where it is not able to promote conditional code
55   // directly.
56   //
57   // Determine randomly in which quadrant of the grid a point lies,
58   // based on the following picture:
59   //
60   //
61   // +-----+
62   // | u < a          | u < a+b |
63   // |-----+-----|
64   // | u < a + b + c | otherwise |
65   // +-----+
66   //
67   // The probability of the edge falling in the upper left quadrant is a,
68   // in the upper right quadrant, b, the lower left quadrant c
69   // and the lower right quadrant d, where the probabilities are
70   // normalized to sum to one.
71   // -----
72
73   var start_inc = 0;
74   var end_inc   = 0;
75   var edge : directed_vertex_pair;
76
77   if u <= a then
78     {}
79   else if u <= a + b then
80     { end_inc = 1; }
81   else if u <= a + b + c then
82     { start_inc = 1; }
83   else
84     { start_inc = 1; end_inc = 1; };
85
86   edge.start = bit * start_inc;
87   edge.end   = bit * end_inc;
88   return ( edge );
89 }
90
91 // =====
92 // Main RMAT Graph Generation Procedure
93 // =====
94
95 proc Gen_RMAT_graph ( a : real,
96                      b : real,
97                      c : real,
98                      d : real,
99                      SCALE : int,
100                     N_VERTICES : int,
101                     n_raw_edges : int,
102                     MAX_EDGE_WEIGHT : int,
103                     G )
104 {
105   use Random;
106
107   const vertex_range = 1..N_VERTICES,
108         edge_range   = 1..n_raw_edges,
109         rand_range   = 1..n_raw_edges + 1;
110
111   // Random Numbers return in the range [0.0, 1.0)
112
113   var Rand_Gen = if REPRODUCIBLE_PROBLEMS then
114     new RandomStream (seed = 0556707007)
115   else
116     new RandomStream ();
117
118   var Noisy_a : [edge_range] real,
119       Noisy_b : [edge_range] real,
120       Noisy_c : [edge_range] real,
121       Noisy_d : [edge_range] real,
122       norm    : [edge_range] real,
123       Unif_Random : [edge_range] real,
124       Edges    = [edge_range] new directed_vertex_pair ();
125
126   // -----
127   // The RMAT algorithm is based on recursively sub-dividing a grid.
128   // In this case, the grid is square of order 2^SCALE by 2^SCALE.
129   // So exactly "SCALE" steps of recursion will be required and the
130   // recursion can be implemented directly by iteration.
131   // -----
132
133   // Note on random number generators -- the RMAT generator creates
134   // 5*SCALE vectors of length 2*SCALE. The dependence on powers
135   // of two provides an opportunity to expose statistical correlations
136   // in the pseudo-random numbers. This certainly occurs with the
137   // current Chapel Random module. The "skips" in the random number
138   // sequence dramatically change the results. Without them, the
139   // Chapel RMAT matrices are inconsistent with what is seen in
140   // other implementations.
141   // -----
142
143   writeln ("Random graph generated by stride of 1 in one random stream",
144           " with skips" );
145
146   var bit = 1 << SCALE;
147   var skip : real;
148
149   for depth in 1..SCALE do {
150     bit >>= 1;
151
152     // randomize the coefficients, tweaking them by numbers in [-.05, .05)
153
154     skip = Rand_Gen.getNext ();
155     Rand_Gen.fillRandom (Unif_Random);
156     Noisy_a = a * (0.95 + 0.1 * Unif_Random);
157
158     skip = Rand_Gen.getNext ();
159     Rand_Gen.fillRandom (Unif_Random);
160     Noisy_b = b * (0.95 + 0.1 * Unif_Random);
161
162     skip = Rand_Gen.getNext ();
163     Rand_Gen.fillRandom (Unif_Random);
164     Noisy_c = c * (0.95 + 0.1 * Unif_Random);
165
166     skip = Rand_Gen.getNext ();
167     Rand_Gen.fillRandom (Unif_Random);
168     Noisy_d = d * (0.95 + 0.1 * Unif_Random);
169
170     norm = 1.0 / (Noisy_a + Noisy_b + Noisy_c + Noisy_d);
171     Noisy_a *= norm;
172     Noisy_b *= norm;
173     Noisy_c *= norm;
174     Noisy_d *= norm;
175
176     skip = Rand_Gen.getNext ();
177     Rand_Gen.fillRandom (Unif_Random);
178
179     Edges += assign_quadrant ( Unif_Random, Noisy_a, Noisy_b,
180                             Noisy_c, Noisy_d, bit );
181
182   };
183
184   // -----
185   // Assign weights to edges randomly, then randomly relabel the vertices
186   // to avoid locality from the obvious imbalance that will arise when
187   // one of the coefficients is clearly larger than the others
188   // -----
189
190   var permutation : [vertex_range] int = vertex_range;
191   var Edge_Weight : [edge_range] int;
192
193   Rand_Gen.fillRandom ( Unif_Random );
194   Edge_Weight = floor (1 + Unif_Random * MAX_EDGE_WEIGHT) : int;
195
196   Rand_Gen.fillRandom ( Unif_Random (vertex_range) );
197
198   for v in vertex_range do
199     { var new_id : int;
200       new_id = floor (1 + Unif_Random (v) * N_VERTICES) : int;
201       permutation (v) <=> permutation (new_id);
202     };
203
204   var node_count : [vertex_range] int = 0;
205
206   // for e in edge_range do {
207   //   Edges(e).start = permutation (Edges(e).start);
208   //   Edges(e).end   = permutation (Edges(e).end );
209   // }
210   Edges.start = permutation (Edges.start);
211   Edges.end   = permutation (Edges.end );
212
213 }

```

```

215   if DEBUG_GRAPH_GENERATOR || DEBUG_WEIGHT_GENERATOR then {
216     writeln ();
217     for e in edge_range do
218       writeln ("edge (" , e, ") : (" , Edges(e), " , " ,
219         Edge_Weight (e), " )" );
220   }
221
222   // -----
223   // Graph consistency checking
224   // -----
225
226   assert ( Edges.start > 0,
227     "edge start vertices out of low end of range");
228
229   assert ( Edges.end > 0,
230     "edge end vertices out of low end of range");
231
232   assert ( Edges.start <= 2**SCALE,
233     "edge start vertices out of high end of range");
234
235   assert ( Edges.end <= 2**SCALE,
236     "edge end vertices out of high end of range");
237
238   assert ( Edge_Weight > 0,
239     "edge weightd out of low end of range");
240
241   assert ( Edge_Weight <= MAX_EDGE_WEIGHT,
242     "edge weightd out of high end of range");
243
244   writeln (); writeln ("Vertex Set in G:", G.vertices);
245
246   // -----
247   // Kernel 1: assemble graph from list of triples. Include
248   // only non-self incident edges. in case of duplicates, last
249   // in wins (+= instead of = works to take sum of weights)
250   // -----
251
252   var collisions = 0, self_edges = 0;
253   for e in edge_range do {
254     var u = Edges (e).start;
255     var v = Edges (e).end ;
256
257     if ( v != u ) then {
258       if G.Neighbors (u).member(v) then {
259         collisions += 1;
260       }
261     }
262     else {
263       G.Neighbors (u).add (v);
264       G.Row(u).Weight (v) = Edge_Weight (e);
265     }
266   }
267   else
268     self_edges += 1;
269 }
270
271 writeln ( "# of raw edges generated " , n_raw_edges );
272 writeln ( "# of duplicate edges " , collisions );
273 writeln ( "# of self edges " , self_edges );
274 writeln ( "# of edges in final graph " ,
275   + reduce [v in G.vertices] G.n_Neighbors (v) );
276
277 if DEBUG_GRAPH_GENERATOR || DEBUG_WEIGHT_GENERATOR then {
278   writeln ();
279   writeln ("tuples denote (edge, weight)");
280   writeln ();
281   for u in G.vertices do {
282     write ( "row " , u, " : [" , G.n_Neighbors (u), "]" );
283     for (v,w) in (G.Neighbors (u), G.edge_weight (u) ) do
284       write ( (v, w) );
285     writeln ();
286   }
287 }
288
289 var max_edges = max reduce [v in vertex_range] G.n_Neighbors (v);
290
291 var edge_count : [0..max_edges] int = 0;
292
293 for v in G.vertices do
294   edge_count (G.n_Neighbors (v)) += 1;
295
296 writeln ("histograph of node distributions by number of outgoing edges");
297 writeln ( "# edges # nodes");
298 for count in 0..max_edges do
299   writeln ( count, " " , edge_count (count) );
300
301 writeln ();
302 }
303 }

```

## G SSCA#2 Kernels 2-4

```

1 module SSCA2_kernels
2
3 // +-----
4 // | Polymorphic Implementation of SSCA #2, Kernels 2-4
5 // |
6 // | Each kernel takes a graph argument which provides for each vertex
7 // | 1. an iterator for its set of neighbors
8 // | 2. a parallel integer array of edge weights, which can be zipper
9 // | iterated with the set of neighbors
10 // | 3. the number of neighbors
11 // |
12 // | These are the only requirements on the representation of the graph.
13 // |
14 // | Filtering in Kernel 4 is turned on or off by a compilation time param.
15 // +-----
16
17 {
18 use SSCA2_compilation_config_params, Time;
19
20 var stopwatch : Timer;
21
22 // =====
23 // KERNEL 2:
24 // =====
25 // Find the edges with the largest edges. Return a list of
26 // edges, all of which have the largest weight.
27 // =====
28
29 proc largest_edges ( G, heavy_edge_list : domain )
30
31 // edge_weights can be either an array over an associative
32 // domain or over a sparse domain. the output heavy_edge_list
33 // can either kind of domain or something else purpose-built
34 // for this task.
35 {
36 if PRINT_TIMING_STATISTICS then stopwatch.start ();
37 var heaviest_edge_weight$ : sync int = 0;
38
39 // -----
40 // find heaviest edge weight in a single pass over all edges
41 // -----
42
43 // heaviest_edge_weight = max reduce [ s in G.vertices ]
44 // [ w in G.edge_weight ( s ) ] w;
45
46 forall s in G.vertices do
47 forall w in G.edge_weight ( s ) do
48 heaviest_edge_weight$ = max ( w, heaviest_edge_weight$ );
49
50 // -----
51 // in a second pass over all edges, extract list
52 // of all edges matching the heaviest weight
53 // -----
54
55 forall s in G.vertices do
56 forall ( t, w ) in ( G.Neighbors ( s ), G.edge_weight ( s ) ) do
57
58 // should be forall, requires a custom parallel iterator in the
59 // random graph case and zippering for associative domains may
60 // also present a problem
61
62 if w == heaviest_edge_weight$.readXX () then {
63 heavy_edge_list.add ( ( s,t ) );
64 };
65
66 if PRINT_TIMING_STATISTICS then {
67 stopwatch.stop ();
68 writeln ( "Elapsed time for Kernel 2: ", stopwatch.elapsed (),
69 " seconds");
70 stopwatch.clear ();
71 }
72
73 // -----
74 // should be able to write a user-defined "maxlocs"
75 // reduction more efficiently than this scheme
76 // -----
77
78 if DEBUG_KERNEL2 then {
79 writeln ();
80 writeln ( "Heaviest weight : ",
81 heaviest_edge_weight$.readFF ());
82 writeln ( "Number of heavy edges:", heavy_edge_list.numIndices );
83 writeln ();
84 writeln ( "Edges with largest weight and other neighbors:" );
85 for ( s,t ) in heavy_edge_list do {
86 writeln ( "edge ", ( s,t );
87 for ( v,w ) in G.Neighbors ( s ), G.edge_weight ( s ) do
88 writeln ( " ", v, " ", w );
89 }
90 };
91
92 // =====
93 // KERNEL 3:
94 // =====
95 // For each root (heavy) edge, find the subgraph (vertices and edges)
96 // defined by directed paths of length no greater than max_path_length
97 // in which the first edge traversed is the root edge
98 // =====
99
100 proc rooted_heavy_subgraphs ( G,
101 Heavy_Edge_List : domain,
102 Heavy_Edge_Subgraph : [],
103 in max_path_length : int )
104
105 // -----
106 // there is a classic space versus time tradeoff. if the subgraphs expanded
107
108 // by breadth first search are small, it would make sense to use a hash
109 // table or an associative domain to represent the search. If the subgraphs
110 // are large, using a full length vector to represent the search is more
111 // appropriate. We expect small diameters for power law graphs, so we
112 // expect large subgraphs.
113 // -----
114
115 {
116 if PRINT_TIMING_STATISTICS then stopwatch.start ();
117
118 const vertex_domain = G.vertices;
119
120 forall ( x, y ) in Heavy_Edge_List do {
121 var Active_Level, Next_Level : domain ( index ( vertex_domain ) );
122 var min_distance$ : [ vertex_domain ] sync int = -1;
123
124 if DEBUG_KERNEL3 then
125 writeln ( " Building heavy edge subgraph from pair:", ( x,y ) );
126 Active_Level.add ( y );
127 Next_Level.clear ();
128 Heavy_Edge_Subgraph ( ( x, y ) ).nodes.clear ();
129 Heavy_Edge_Subgraph ( ( x, y ) ).edges.clear ();
130 min_distance$ ( y ).writeFF ( 0 );
131
132 Heavy_Edge_Subgraph ( ( x, y ) ).edges.add ( ( x, y ) );
133 Heavy_Edge_Subgraph ( ( x, y ) ).nodes.add ( x );
134 Heavy_Edge_Subgraph ( ( x, y ) ).nodes.add ( y );
135
136 for path_length in 1 .. max_path_length do {
137
138 forall v in Active_Level do {
139
140 forall w in G.Neighbors ( v ) do { // eventually, will be forall
141
142 if min_distance$ ( w ).readXX () < 0 then {
143
144 if min_distance$ ( w ).readFE () < 0 then {
145 Next_Level.add ( w );
146 Heavy_Edge_Subgraph ( ( x, y ) ).nodes.add ( w );
147 min_distance$ ( w ).writeEF ( path_length );
148 }
149 else
150 min_distance$ ( w ).writeEF ( path_length );
151 }
152
153 // min_distance$ must have been set by some thread by now
154
155 if min_distance$ ( w ).readFF () == path_length then {
156 Heavy_Edge_Subgraph ( ( x, y ) ).edges.add ( ( v, w ) );
157 }
158 }
159 }
160
161 if path_length < max_path_length then {
162 Active_Level = Next_Level;
163 Next_Level.clear ();
164 }
165 }
166 }
167
168 if PRINT_TIMING_STATISTICS then {
169 stopwatch.stop ();
170 writeln ( "Elapsed time for Kernel 3: ", stopwatch.elapsed (),
171 " seconds");
172 stopwatch.clear ();
173 }
174 } // end of rooted_heavy_subgraphs
175
176 // =====
177 // generic class structure must be defined outside of
178 // generic procedure. used by Betweenness Centrality kernel 4.
179 // =====
180 // The set of vertices at a particular distance from s form a
181 // level set. The class allows the full set of vertices to be
182 // partitioned into a linked list of level sets. Each instance
183 // of the outer loop in kernel 4 creates such a partitioning.
184 // =====
185
186 class Level_Set {
187 type Sparse_Vertex_List;
188 var Members : Sparse_Vertex_List;
189 var previous : Level_Set ( Sparse_Vertex_List );
190 }
191
192 // sungeun: 8/2011
193 // Added replicated level sets
194 //
195 // Each locale will have its own level sets. A locale's level set
196 // will only contain nodes that are physically allocated on that
197 // particular locale. We implement this using the replicated
198 // distribution.
199
200 use ReplicatedDist;
201
202 // =====
203 // KERNEL 4
204 // =====
205 // Calculate Betweenness Centrality for simple unweighted directed or
206 // undirected graphs, using Madduri, et.al.'s modification of
207 // Brandes's 2001 algorithm
208 // =====
209
210 proc approximate_betweenness_centrality ( G, starting_vertices,
211 Between_Cent : [] real,
212 out Sum_Min_Dist : real )

```

```

215 // -----
216 // The betweenness centrality metric for a given node v is defined
217 // as the double sum over s not equal to v and t not equal to
218 // either s or v of the ratio of the number of shortest paths from s to t
219 // passing through v to the number of shortest paths from s to t.
220 // -----
221 // Brandes's algorithm decomposes the computation of this metric into,
222 // first, separate sums for each vertex s, which can be computed
223 // independently in parallel, and
224 // two, a recursive, tree-based, calculation of the path counts for
225 // any particular s.
226 // The complexity of this algorithm is O ( |V||E| ) time for an unweighted
227 // graph. The algorithm requires O ( |V| ) temporary space for each
228 // process that executes instances of the outermost loop.
229 // -----
230 {
231   const vertex_domain = G.vertices;
232
233   // Had to change declaration below
234   // type Sparse_Vertex_List = sparse subdomain ( G.vertices );
235   // to accommodate block distribution of G.vertices
236
237   type Sparse_Vertex_List = domain(index(vertex_domain));
238
239   var Between_Cent$ : [vertex_domain] sync real = 0.0;
240   var Sum_Min_Dist$ : sync real = 0.0;
241
242   // -----
243   // Each iteration of the outer loop of Brandes's algorithm
244   // computes the contribution (the "dependency" metric) for
245   // one particular vertex (s) independently.
246   // -----
247
248   if PRINT_TIMING_STATISTICS then stopwatch.start ();
249
250   forall s in starting_vertices do {
251
252     if DEBUG_KERNEL4 then writeln ( "expanding from starting node ", s );
253
254     // sungeun: 8/2011
255     // Privatization of the following distributed arrays may
256     // be of concern.
257
258     // -----
259     // all locally declared variables become private data
260     // for each instance of the parallel for loop
261     // -----
262
263     var min_distance$ : [vertex_domain] sync int = -1;
264     var path_count$ : [vertex_domain] sync real (64) = 0.0;
265     var depend : [vertex_domain] real = 0.0;
266     var Lcl_Sum_Min_Dist: sync real = 0.0;
267
268     // The structure of the algorithm depends on a breadth-first
269     // traversal. Each vertex will be marked by the length of
270     // the shortest path (min_distance$) from s to it. The array
271     // path_count$ will hold a count of the number of shortest
272     // paths from s to this node. The number of paths in moderate
273     // sized tori exceeds 2**64.
274
275     // Used to check termination of the forward pass
276     // -----
277     // sungeun: 8/2011
278     // Can possibly use a replicated bool with rcCollect(), but
279     // not sure of the performance implications for large numbers
280     // of locales.
281     var Active_Remaining: [LocaleSpace] bool = true;
282     var remaining = true;
283
284     // Replicated level sets
285     var Active_Level : [rcDomain] Level_Set (Sparse_Vertex_List);
286     var Next_Level : [rcDomain] Level_Set (Sparse_Vertex_List);
287     coforall loc in Locales do on loc {
288       rLocal(Active_Level) = new Level_Set (Sparse_Vertex_List);
289       rLocal(Active_Level).previous = nil;
290       rLocal(Next_Level) = new Level_Set (Sparse_Vertex_List);
291       rLocal(Next_Level).previous = rLocal(Active_Level);
292     }
293
294     var current_distance : int = 0;
295
296     // establish the initial level sets for the
297     // breadth-first traversal from s
298
299     on s {
300       rLocal(Active_Level).Members.add ( s );
301       rLocal(Next_Level).Members.clear ();
302       min_distance$ ( s ) . writeFF ( 0 );
303       path_count$ ( s ) . writeFF ( 1 );
304     }
305
306     while remaining do {
307
308       // -----
309       // expand the neighbor sets for all vertices at the
310       // current distance from the starting vertex s
311       // -----
312
313       current_distance += 1;
314
315       // sungeun: 8/2011
316       // basic single use barrier
317       var count: sync int = numLocales;
318       var barrier: single bool;
319
320       // sungeun: 8/2011
321       // Copy this value to a constant to enable remote value
322       // forwarding optimization.
323       const current_distance_c = current_distance;
324       coforall loc in Locales do on loc {
325         forall u in rLocal(Active_Level).Members do { // sparse
326
327           // -----
328           // should be forall, requires a custom parallel iterator in the
329           // random graph case and zippering for associative domains
330           // -----
331           // -----
332           // add any unmarked neighbors to the next level
333           // -----
334           // -----
335           if ( FILTERING && w % 8 != 0 ) || !FILTERING then
336             if min_distance$ ( v ) . readXX ( ) < 0 then
337               {
338                 if min_distance$ ( v ) . readFE ( ) < 0 then
339                   {
340                     min_distance$ ( v ) . writeEF ( current_distance_c );
341                     rLocal(Next_Level).Members.add ( v );
342                     if VALIDATE_BC then
343                       Lcl_Sum_Min_Dist += current_distance_c;
344                   }
345                 else
346                   // could min_distance$(v) be < current_distance?
347                   min_distance$ ( v ) . writeEF ( current_distance_c );
348               }
349
350           // -----
351           // only neighbors of u that are in the next level
352           // are on shortest paths from s through v. Some
353           // thread will have set min_distance$ (v) by the
354           // time this code is reached, whether v lies in
355           // the previous, the current or the next level.
356           // -----
357
358           if min_distance$ ( v ) . readFF ( ) == current_distance_c
359             then
360             path_count$ ( v ) += path_count$ ( u ) . readFF ( );
361           }
362         };
363
364         // sungeun: 8/2011
365         // This (split-phase) barrier is needed to insure all updates
366         // to Next_Level are completed before creating the next
367         // Next_Level.
368         var myc = count;
369         if myc==1 {
370           count = myc-1; // release the lock
371           // do some work while we wait
372           rLocal(Active_Level) = rLocal(Next_Level);
373
374           barrier; // wait for everyone
375         } else { // last one here
376           barrier=true; // release everyone first
377           rLocal(Active_Level) = rLocal(Next_Level);
378         }
379
380         rLocal(Next_Level) = new Level_Set (Sparse_Vertex_List);
381         rLocal(Next_Level).previous = rLocal(Active_Level);
382         Active_Remaining[here.id] =
383           rLocal(Active_Level).Members.numIndices:bool;
384       }
385
386       remaining = || reduce Active_Remaining;
387     }; // end forward pass
388
389     if VALIDATE_BC then
390       Sum_Min_Dist$ += Lcl_Sum_Min_Dist;
391
392     // -----
393     // compute the dependencies recursively, traversing the vertices
394     // of the graph in non-increasing order of distance (reverse
395     // ordering from the initial traversal)
396     // -----
397     var graph_diameter = current_distance - 1;
398
399     if DEBUG_KERNEL4 then
400       writeln ( " graph diameter from starting node ", s,
401         " is ", graph_diameter );
402
403     // sungeun: 8/2011
404     // basic single use barrier
405     var count: sync int = numLocales;
406     // to simplify synchronization between multiple barriers
407     var barrier: [2..graph_diameter] single bool;
408
409     coforall loc in Locales do on loc {
410       delete rLocal(Next_Level); // it's empty
411       rLocal(Next_Level) = rLocal(Active_Level).previous; // back up to last level
412       delete rLocal(Active_Level);
413       rLocal(Active_Level) = rLocal(Next_Level);
414
415       for current_distance in 2 .. graph_diameter by -1 do {
416
417         rLocal(Next_Level) = rLocal(Active_Level).previous;
418         delete rLocal(Active_Level);
419         rLocal(Active_Level) = rLocal(Next_Level);
420
421         // inner reduction should parallelize eventually; compiler
422         // serializes it today (and warns us that it did)
423
424         forall u in rLocal(Active_Level).Members do
425           {
426             depend (u) = + reduce
427               [ ( v, w ) in ( G.Neighbors ( u ), G.edge_weight ( u ) ) ]
428               if ( min_distance$ ( v ) . readFF ( ) == current_distance ) &&
429                 ( ( FILTERING && w % 8 != 0 ) || !FILTERING ) )
430             then
431               ( path_count$ ( u ) . readFF ( ) /
432                 path_count$ ( v ) . readFF ( ) ) *
433               ( 1.0 + depend ( v ) );
434           }
435
436           // -----
437           // -----
438           // -----
439           // -----
440           // -----
441           // -----
442           // -----
443           // -----
444           // -----
445           // -----
446           // -----
447           // -----
448           // -----
449           // -----
450           // -----
451           // -----
452           // -----
453           // -----
454           // -----
455           // -----
456           // -----
457           // -----
458           // -----
459           // -----
460           // -----
461           // -----
462           // -----
463           // -----
464           // -----
465           // -----
466           // -----
467           // -----
468           // -----
469           // -----
470           // -----
471           // -----
472           // -----
473           // -----
474           // -----
475           // -----
476           // -----
477           // -----
478           // -----
479           // -----
480           // -----
481           // -----
482           // -----
483           // -----
484           // -----
485           // -----
486           // -----
487           // -----
488           // -----
489           // -----
490           // -----
491           // -----
492           // -----
493           // -----
494           // -----
495           // -----
496           // -----
497           // -----
498           // -----
499           // -----
500           // -----
501           // -----
502           // -----
503           // -----
504           // -----
505           // -----
506           // -----
507           // -----
508           // -----
509           // -----
510           // -----
511           // -----
512           // -----
513           // -----
514           // -----
515           // -----
516           // -----
517           // -----
518           // -----
519           // -----
520           // -----
521           // -----
522           // -----
523           // -----
524           // -----
525           // -----
526           // -----
527           // -----
528           // -----
529           // -----
530           // -----
531           // -----
532           // -----
533           // -----
534           // -----
535           // -----
536           // -----
537           // -----
538           // -----
539           // -----
540           // -----
541           // -----
542           // -----
543           // -----
544           // -----
545           // -----
546           // -----
547           // -----
548           // -----
549           // -----
550           // -----
551           // -----
552           // -----
553           // -----
554           // -----
555           // -----
556           // -----
557           // -----
558           // -----
559           // -----
560           // -----
561           // -----
562           // -----
563           // -----
564           // -----
565           // -----
566           // -----
567           // -----
568           // -----
569           // -----
570           // -----
571           // -----
572           // -----
573           // -----
574           // -----
575           // -----
576           // -----
577           // -----
578           // -----
579           // -----
580           // -----
581           // -----
582           // -----
583           // -----
584           // -----
585           // -----
586           // -----
587           // -----
588           // -----
589           // -----
590           // -----
591           // -----
592           // -----
593           // -----
594           // -----
595           // -----
596           // -----
597           // -----
598           // -----
599           // -----
600           // -----
601           // -----
602           // -----
603           // -----
604           // -----
605           // -----
606           // -----
607           // -----
608           // -----
609           // -----
610           // -----
611           // -----
612           // -----
613           // -----
614           // -----
615           // -----
616           // -----
617           // -----
618           // -----
619           // -----
620           // -----
621           // -----
622           // -----
623           // -----
624           // -----
625           // -----
626           // -----
627           // -----
628           // -----
629           // -----
630           // -----
631           // -----
632           // -----
633           // -----
634           // -----
635           // -----
636           // -----
637           // -----
638           // -----
639           // -----
640           // -----
641           // -----
642           // -----
643           // -----
644           // -----
645           // -----
646           // -----
647           // -----
648           // -----
649           // -----
650           // -----
651           // -----
652           // -----
653           // -----
654           // -----
655           // -----
656           // -----
657           // -----
658           // -----
659           // -----
660           // -----
661           // -----
662           // -----
663           // -----
664           // -----
665           // -----
666           // -----
667           // -----
668           // -----
669           // -----
670           // -----
671           // -----
672           // -----
673           // -----
674           // -----
675           // -----
676           // -----
677           // -----
678           // -----
679           // -----
680           // -----
681           // -----
682           // -----
683           // -----
684           // -----
685           // -----
686           // -----
687           // -----
688           // -----
689           // -----
690           // -----
691           // -----
692           // -----
693           // -----
694           // -----
695           // -----
696           // -----
697           // -----
698           // -----
699           // -----
700           // -----
701           // -----
702           // -----
703           // -----
704           // -----
705           // -----
706           // -----
707           // -----
708           // -----
709           // -----
710           // -----
711           // -----
712           // -----
713           // -----
714           // -----
715           // -----
716           // -----
717           // -----
718           // -----
719           // -----
720           // -----
721           // -----
722           // -----
723           // -----
724           // -----
725           // -----
726           // -----
727           // -----
728           // -----
729           // -----
730           // -----
731           // -----
732           // -----
733           // -----
734           // -----
735           // -----
736           // -----
737           // -----
738           // -----
739           // -----
740           // -----
741           // -----
742           // -----
743           // -----
744           // -----
745           // -----
746           // -----
747           // -----
748           // -----
749           // -----
750           // -----
751           // -----
752           // -----
753           // -----
754           // -----
755           // -----
756           // -----
757           // -----
758           // -----
759           // -----
760           // -----
761           // -----
762           // -----
763           // -----
764           // -----
765           // -----
766           // -----
767           // -----
768           // -----
769           // -----
770           // -----
771           // -----
772           // -----
773           // -----
774           // -----
775           // -----
776           // -----
777           // -----
778           // -----
779           // -----
780           // -----
781           // -----
782           // -----
783           // -----
784           // -----
785           // -----
786           // -----
787           // -----
788           // -----
789           // -----
790           // -----
791           // -----
792           // -----
793           // -----
794           // -----
795           // -----
796           // -----
797           // -----
798           // -----
799           // -----
800           // -----
801           // -----
802           // -----
803           // -----
804           // -----
805           // -----
806           // -----
807           // -----
808           // -----
809           // -----
810           // -----
811           // -----
812           // -----
813           // -----
814           // -----
815           // -----
816           // -----
817           // -----
818           // -----
819           // -----
820           // -----
821           // -----
822           // -----
823           // -----
824           // -----
825           // -----
826           // -----
827           // -----
828           // -----
829           // -----
830           // -----
831           // -----
832           // -----
833           // -----
834           // -----
835           // -----
836           // -----
837           // -----
838           // -----
839           // -----
840           // -----
841           // -----
842           // -----
843           // -----
844           // -----
845           // -----
846           // -----
847           // -----
848           // -----
849           // -----
850           // -----
851           // -----
852           // -----
853           // -----
854           // -----
855           // -----
856           // -----
857           // -----
858           // -----
859           // -----
860           // -----
861           // -----
862           // -----
863           // -----
864           // -----
865           // -----
866           // -----
867           // -----
868           // -----
869           // -----
870           // -----
871           // -----
872           // -----
873           // -----
874           // -----
875           // -----
876           // -----
877           // -----
878           // -----
879           // -----
880           // -----
881           // -----
882           // -----
883           // -----
884           // -----
885           // -----
886           // -----
887           // -----
888           // -----
889           // -----
890           // -----
891           // -----
892           // -----
893           // -----
894           // -----
895           // -----
896           // -----
897           // -----
898           // -----
899           // -----
900           // -----
901           // -----
902           // -----
903           // -----
904           // -----
905           // -----
906           // -----
907           // -----
908           // -----
909           // -----
910           // -----
911           // -----
912           // -----
913           // -----
914           // -----
915           // -----
916           // -----
917           // -----
918           // -----
919           // -----
920           // -----
921           // -----
922           // -----
923           // -----
924           // -----
925           // -----
926           // -----
927           // -----
928           // -----
929           // -----
930           // -----
931           // -----
932           // -----
933           // -----
934           // -----
935           // -----
936           // -----
937           // -----
938           // -----
939           // -----
940           // -----
941           // -----
942           // -----
943           // -----
944           // -----
945           // -----
946           // -----
947           // -----
948           // -----
949           // -----
950           // -----
951           // -----
952           // -----
953           // -----
954           // -----
955           // -----
956           // -----
957           // -----
958           // -----
959           // -----
960           // -----
961           // -----
962           // -----
963           // -----
964           // -----
965           // -----
966           // -----
967           // -----
968           // -----
969           // -----
970           // -----
971           // -----
972           // -----
973           // -----
974           // -----
975           // -----
976           // -----
977           // -----
978           // -----
979           // -----
980           // -----
981           // -----
982           // -----
983           // -----
984           // -----
985           // -----
986           // -----
987           // -----
988           // -----
989           // -----
990           // -----
991           // -----
992           // -----
993           // -----
994           // -----
995           // -----
996           // -----
997           // -----
998           // -----
999           // -----
1000          // -----

```



```

439         // do not need conditional u != s
441         Between_Cent$ (u) += depend (u);
442     }
443     // sungeun: 8/2011
444     // This barrier is needed to insure all updates to depend are
445     // complete before the next pass.
446     var myc = count;
447     if myc==1 {
448         count = numLocales;
449         barrier[current_distance] = true;
450     } else {
451         count = myc-1;
452         barrier[current_distance];
453     }
454 };
455 delete rcLocal(Active_Level);
456 }
458 }; // closure of outer embarassingly parallel forall
461 if PRINT_TIMING_STATISTICS then {
462     stopwatch.stop ();
463     var K4_time = stopwatch.elapsed ();
464     stopwatch.clear ();
465     writeln ( "Elapsed time for Kernel 4: ", K4_time, " seconds");
467     var n0 = + reduce [v in vertex_domain] (G.n_Neighbors (v) == 0);
468     var n_edges = + reduce [v in vertex_domain] G.n_Neighbors (v);
469     var N_VERTICES = vertex_domain.numIndices;
470     var TEPS = 7.0 * N_VERTICES * (N_VERTICES - n0) / K4_time;
471     var Adjusted_TEPS = n_edges * (N_VERTICES - n0) / K4_time;
473     writeln ( " TEPS: ", TEPS );
474     writeln ( " edge count adjusted TEPS: ", Adjusted_TEPS );
475 }
477 if VALIDATE_BC then
478     Sum_Min_Dist = Sum_Min_Dist$;
480     Between_Cent = Between_Cent$;
482 } // end of Brandes' betweenness centrality calculation
484 }

```

## H SSCA#2 R-MAT graph support code

```

1 module analyze_RMAT_graph_associative_array {
2
3 // +=====
4 // | Define associative array-based representations for general sparse |
5 // | graphs. Provide execution template to generate a random RMAT graph |
6 // | of a specified size and execute and verify SSCA2 kernels 2 through 4. |
7 // +=====
8
9 proc generate_and_analyze_associative_array_RMAT_graph_representation {
10
11 // -----
12 // compute a random power law graph with 2^SCALE vertices, using
13 // the RMAT generator. Initially generate a list of triples.
14 // Then convert it to a Chapel representation of a sparse graph,
15 // timing this step (Kernel 1). Finally, execute Kernels 2, 3 and 4
16 // of SSCA #2, using identically the same code as in the various
17 // torus cases.
18 // -----
19
20 use SSCA2_compilation_config_params, SSCA2_execution_config_consts;
21
22 use SSCA2_driver, SSCA2_RMAT_graph_generator;
23
24 use BlockDist;
25
26 var n_raw_edges = 8 * N_VERTICES;
27
28 assert ( SCALE > 1, "SCALE must be greater than 1");
29
30 select SCALE {
31   when 2 do n_raw_edges = N_VERTICES / 2;
32   when 3 do n_raw_edges = N_VERTICES;
33   when 4 do n_raw_edges = 2 * N_VERTICES;
34   when 5 do n_raw_edges = 4 * N_VERTICES;
35 }
36
37 writeln ('-----');
38 writeln ('Order of RMAT generated graph:', N_VERTICES);
39 writeln ('          number of raw edges:', n_raw_edges);
40 writeln ('-----');
41 writeln ();
42
43 // -----
44 // The data structures below are chosen to implement an irregular (sparse)
45 // graph using associative domains and arrays.
46 // Each node in the graph has a list of neighbors and a corresponding list
47 // of (integer) weights for the implicit edges.
48 // The list of neighbors is really just a set; the only properties we need
49 // are that we be able to build it (add vertices to it) and that we be
50 // able to iterate over it. Those properties are satisfied by Chapel's
51 // associative domains, so each neighbor set is represented by an
52 // associative domain. The weights are an integer array over the
53 // neighbor domain.
54 //
55 // We would have liked to have defined the global set of neighbors as
56 // an array of associative domains, but that is not supported in Chapel.
57 // Consequently we build an array of records, where each record provides
58 // the neighbor set and the weights for a particular node. The name
59 // "row_struct" anticipates the planned use of sparse matrices for this same
60 // kind of graph structure.
61 // -----
62
63 const vertex_domain =
64   if DISTRIBUTION_TYPE == "BLOCK" then
65     [1..N_VERTICES] dmapped Block ( [1..N_VERTICES] )
66   else
67     [1..N_VERTICES] ;
68
69 record row_struct {
70   type vertex;
71   var Row_Neighbors : domain (vertex);
72   var Weight       : [Row_Neighbors] int;
73 }
74
75 class Associative_Graph {
76   const vertices;
77   var Row       : [vertices] row_struct (index (vertices));
78
79   proc Neighbors ( v : index (vertices) ) {return Row (v).Row_Neighbors;}
80
81   iter edge_weight ( v : index (vertices) ) var {
82     for w in Row (v).Weight do
83       yield w;} // var iterator to avoid a copy
84
85   // Simply forward the domain's parallel iterator
86   // FYI: no fast follower opt
87   iter edge_weight ( v : index (vertices), param tag: iterKind)
88   where tag == iterKind.leader {
89     for block in Row(v).Weight._value.these(tag) do
90       yield block;
91   }
92
93   iter edge_weight ( v : index (vertices), param tag: iterKind, followThis)
94   where tag == iterKind.follower {
95     for elem in Row(v).Weight._value.these(tag, followThis) do
96       yield elem;
97   }
98
99   proc n_Neighbors ( v : index (vertices) )
100   {return Row (v).Row_Neighbors.numIndices;}
101 }
102
103 var G = new Associative_Graph (vertex_domain);
104
105 // -----
106 // generate RMAT graph of the specified size, based on input config
107 // values for quadrant assignment.
108 // -----
109
110 Gen_RMAT_graph ( RMAT_a, RMAT_b, RMAT_c, RMAT_d,
111                 SCALE, N_VERTICES, n_raw_edges, MAX_EDGE_WEIGHT, G );
112
113 execute_SSCA2 ( G );
114 writeln (); writeln ();
115 delete G;
116 }

```