

Exascale: Your Opportunity to Create a Decent HPC Language

Brad Chamberlain, Cray Inc.

PPME: Productive Programming Models for Exascale

August 14th, 2012





Prototypical Exascale Processor Technologies



Sources: http://www.zdnet.com/amds-trinity-processors-take-on-intels-ivy-bridge-3040155225/, http://insidehpc.com/2010/11/26/nvidia-reveals-details-of-echelon-gpu-designs-for-exascale/, http://itlera.com/sites/default/files/productbriefs/Tile-Gx%203036%20SB012-01.pdf





10 GbE XAUI

General Characteristics of These Architectures



- Increased hierarchy and/or sensitivity to locality
- Potentially heterogeneous processor/memory types

⇒ Exascale programmers will have a lot more to think about at the node level than in the past



Exascale Challenges



- intra-node locality and heterogeneity concerns
- limited memory bandwidth, memory::FLOP ratios
- increased resiliency concerns
- increased sensitivity to energy usage
- diversity of abstract machine models (at least initially)
- traditional programming models aren't a good fit

A time to be afraid? (very afraid?)

Or, an opportunity?

(If we have to switch to something new, we may as well take the opportunity to adopt something decent while we're at it)



A possible contrarian viewpoint: "I don't care about Exascale because...

... it's too expensive."

...it won't serve a broad enough community."

...it doesn't suit my application area."

...I don't believe it's possible/practical/worthwhile."

...anyone who can afford an Exascale machine can afford to hire an army of programmers to suffer through whatever terrible notations are required to utilize it effectively."

But look, these challenges are not Exascale-specific...





Exascale Programming Model Wishlist

general parallelism:

- data parallelism: to take advantage of SIMD HW units; for simplicity when applicable
- task parallelism: to fire asynchronous computations off to accelerators; for data-driven algorithms
- varying granularities/nestings: for generality across computations & levels of the machine
- **locality control:** to tune for locality/affinity across the machine (inter- and intra-node)

resilience-/energy-aware features: to deal with emerging issues at system scale

- user extensibility: to deal with unknowns in next-generation architectures and algorithms
- productivity features: because you know you want them







Has "productivity" run its course, thanks to HPCS?

- Maybe programmatically or politically...
- But look, who wouldn't want code that's easier to...
 - ...write ...read ...tune ...maintain

...port

.

The goal remains worthy even if the term has grown stale



How to attract new parallel programmers?



INDITEDNIETLY FLAGE FOR PARALLEL PROGRAMMING NEAREST SUPERCOMPUTING CENTER





"Decent" Parallel Languages

- What was the last parallel notation you used that felt
 - productive?
 - high-level?
 - powerful?
 - flexible?
 - effective?
 - modern?
 - fun?



• (...all the things we judge good software by...)?

(Because that's what we're competing with)





Exascale Scorecard for HPC Programming Models

	Fortran	C/C++	MPI	OpenMP	UPC
data parallelism					
task parallelism					
parallel nesting/granularities					
locality control					
resilience/energy-awareness					
user-extensibility					
productivity					





Exascale Scorecard for HPC Programming Models

	Fortran	C/C++	MPI	OpenMP	UPC
data parallelism	~			~	~
task parallelism				~	
parallel nesting/granularities				~	
locality control			~		~
resilience/energy-awareness			~		
user-extensibility					
productivity				~	





Exascale Scorecard for HPC Programming Models

	Fortran	C/C++	MPI	OpenMP	UPC
data parallelism	~	Х	Х	~	~
task parallelism	Х	Х	Х	~	Х
parallel nesting/granularities	Х	Х	Х	~	Х
locality control	Х	Х	~	Х	~
resilience/energy-awareness	Х	Х	~	Х	X
user-extensibility	Х	Х	Х	Х	X
productivity	Х	Х	Х	~	X





Q: Are we ready?A: In a nutshell, no

Q: Why?

A: We've built too many assumptions about our target architectures into our programming models

- granularity and style of parallelism
- mode of communication
- single level of locality, if any at all
- inflexible, non-productive base languages

productive programming models should be focused more on the user's application than on the hardware





STREAM Triad: a trivial parallel computation

Given: *m*-element vectors *A*, *B*, *C*

Compute: $\forall i \in 1..m, A_i = B_i + \alpha \cdot C_i$

In pictures:







STREAM Triad: a trivial parallel computation

Given: *m*-element vectors *A*, *B*, *C*

Compute: $\forall i \in 1..m, A_i = B_i + \alpha \cdot C_i$

In pictures, in parallel (distributed memory multicore):





STREAM Triad: MPI



```
MPI
#include <hpcc.h>
                                                        if (!a || !b || !c) {
                                                          if (c) HPCC free(c);
                                                          if (b) HPCC free(b);
                                                          if (a) HPCC free(a);
static int VectorSize;
                                                          if (doIO) {
static double *a, *b, *c;
                                                            fprintf( outFile, "Failed to allocate memory
                                                          (%d).\n", VectorSize );
int HPCC StarStream(HPCC Params *params) {
                                                            fclose( outFile );
  int myRank, commSize;
 int rv, errCount;
                                                          return 1;
 MPI Comm comm = MPI COMM WORLD;
 MPI Comm size ( comm, &commSize );
 MPI Comm rank( comm, &myRank );
  rv = HPCC Stream( params, 0 == myRank);
                                                        for (j=0; j<VectorSize; j++) {</pre>
 MPI Reduce ( &rv, &errCount, 1, MPI INT, MPI SUM,
                                                         b[i] = 2.0;
   0, comm );
                                                          c[j] = 0.0;
  return errCount;
                                                        scalar = 3.0;
int HPCC Stream(HPCC Params *params, int doIO) {
  register int j;
  double scalar;
                                                        for (j=0; j<VectorSize; j++)</pre>
 VectorSize = HPCC LocalVectorSize( params, 3,
                                                          a[j] = b[j]+scalar*c[j];
   sizeof(double), 0 );
                                                        HPCC free(c);
  a = HPCC XMALLOC( double, VectorSize );
                                                        HPCC free(b);
  b = HPCC XMALLOC( double, VectorSize );
                                                        HPCC free(a);
  c = HPCC XMALLOC( double, VectorSize );
                                                        return 0;
```



STREAM Triad: MPI+OpenMP

#include <hpcc.h>



```
#ifdef OPENMP
#include <omp.h>
#endif
static int VectorSize;
static double *a, *b, *c;
int HPCC StarStream(HPCC Params *params) {
  int myRank, commSize;
 int rv, errCount;
 MPI Comm comm = MPI COMM WORLD;
 MPI Comm size( comm, &commSize );
 MPI Comm rank( comm, &myRank );
  rv = HPCC Stream( params, 0 == myRank);
 MPI Reduce ( &rv, &errCount, 1, MPI INT, MPI SUM,
   0, comm );
  return errCount;
int HPCC Stream(HPCC Params *params, int doIO) {
  register int j;
  double scalar;
 VectorSize = HPCC LocalVectorSize( params, 3,
   sizeof(double), 0 );
  a = HPCC XMALLOC( double, VectorSize );
  b = HPCC XMALLOC( double, VectorSize );
```

```
c = HPCC_XMALLOC( double, VectorSize );
```

```
MPI + OpenMP
```

```
if (!a || !b || !c) {
    if (c) HPCC_free(c);
    if (b) HPCC_free(b);
    if (a) HPCC_free(a);
    if (doIO) {
      fprintf( outFile, "Failed to allocate memory
    (%d).\n", VectorSize );
      fclose( outFile );
    }
    return 1;
```

```
#ifdef _OPENMP
#pragma omp parallel for
#endif
for (j=0; j<VectorSize; j++) {
    b[j] = 2.0;
    c[j] = 0.0;
}</pre>
```

scalar = 3.0;

```
HPCC_free(c);
HPCC_free(b);
HPCC_free(a);
```

```
return 0;
```



STREAM Triad: MPI+OpenMP vs. CUDA



MPI + OpenMP CUDA #include <hpcc.h> #ifdef OPENMP #define N 2000000 #include <omp.h> #endif int main() { static int VectorSize; static double *a, *b, *c; float *d a, *d b, *d c; float scalar; int HPCC StarStream (HPCC Params *params) int myRank, commSize; int rv, errCount; MPI Comm comm = MPI COMM WORLD; cudaMalloc((void**)&d a, sizeof(float)*N); cudaMalloc((void**)&d b, sizeof(float)*N); MPI Comm_size(comm, &commSize); MPI_Comm_rank(comm, &myRank); cudaMalloc((void**)&d c, sizeof(float)*N); rv = HPCC Stream(params, 0 == myRank); MPI Reduce(&rv, &errCount, 1, MPI INT, MPI SUM, 0, comm); HPC suffers from too many distinct notations for expressing parallelism and locality register int j; double scalar; Hybrid solutions are our only recourse to program Exascale at all; and while they're reasonable as conservative approaches, they're far from ideal if (c) HPCC_free(c); if (b) HPCC free(b); cudaThreadSynchronize(); if (a) HPCC_free(a); if (doIO) { fprintf(outFile, "Failed to allocate memory (%d).\n", VectorSize); cudaFree(d a); fclose(outFile); cudaFree(d b); return 1; cudaFree(d c); #ifdef OPENMP #pragma omp parallel for #endif global void set array(float *a, float value, int len) { for (j=0; j<VectorSize; j++) {</pre> b[j] = 2.0;int idx = threadIdx.x + blockIdx.x * blockDim.x; c[j] = 0.0;if (idx < len) a[idx] = value; scalar = 3.0;#ifdef OPENMP void STREAM Triad (float *a, float *b, float *c, #pragma omp parallel for global #endif float scalar, int len) { int idx = threadIdx.x + blockIdx.x * blockDim.x; HPCC free(c); if (idx < len) c[idx] = a[idx]+scalar*b[idx];</pre> HPCC_free(b); HPCC free(a); return 0;



By Analogy: Let's Cross the United States!







By Analogy: Let's Cross the United States!



STREAM Triad: Chapel





HPCC_fre HPCC_fre HPCC_fre

return (

<u>Philosophy:</u> Good language design can tease details of locality and parallelism away from an algorithm, permitting the compiler, runtime, applied scientist, and HPC expert each to focus on their strengths.

HAPEL



A Common Question

Q: Didn't we try this before with HPF?

Q': Orville, didn't Percy Pilcher die in *his* prototype powered aircraft?



A': No Wilbur, he died in a glider; and even if it had been in his prototype, that doesn't mean we're doomed to fail.





Q: How Can Chapel Succeed When HPF Failed?

A: Chapel has had the chance to learn from HPF's mistakes (and other languages' successes and failures)

- Why did HPF fail?
 - lack of sufficient performance soon enough
 - vagueness in execution/implementation model
 - only supported a single level of data parallelism, no task/nested
 - inability to drop to lower levels of control
 - fixed set of limited distributions on dense arrays
 - lacked richer data parallel abstractions
 - lacked an open source implementation
 - too Fortran-based for modern programmers
 - ...?
- The failure of one language---even a well-funded, US-backed one---does not dictate the failure of all future languages

(For more on this topic see https://www.ieeetcsc.org/activities/blog/myths_about_scalable_parallel_programming_languages_part2)





Chapel in a Nutshell

Chapel: a parallel language that has emerged from DARPA HPCS

• general parallelism:

- data-, task-, and nested parallelism
- highly dynamic multithreading or static SPMD-style
- multiresolution philosophy: high-level features built on low-level
 - to provide "manual overrides"
 - to support a separation of concerns (application vs. parallel experts)
- Iocality control:
 - explicit or data-driven placement of data and tasks
 - locality expressed distinctly from parallelism
- features for productivity: type inference, iterators, rich array types
- **portable:** designed and implemented to support diverse systems
- open source: developed and distributed under the BSD license
- plausibly adoptable: forward-thinking HPC users want a mature version





Chapel: Well-Positioned for Exascale

data parallelism	1
task parallelism	 Image: A start of the start of
parallel nesting/granularities	1
locality control	~
resilience/energy-awareness	X
user-extensibility	 Image: A start of the start of
productivity	1





Chapel limitations for Exascale, today:

- locales only support a single level of hierarchy
 - useful for horizontal (inter-node) locality
 - less so for describing additional hierarchy within a node
- lack of fault tolerance/error handling features

In Chapel's original design, these were both considered "version 2.0" features due to...

...our focus on petascale systems within HPCS

...the knowledge that our plate was already quite full





Current Work: Hierarchical Locales

Concept:

 Support locales within locales to describe architectural sub-structures within a node



- As with traditional locales, on-clauses and domain maps can be used to map tasks and variables to a sub-locale's memory and processors
- Locale structure is defined as Chapel code
 - permits implementation policy to be specified in-language
 - introduces a new Chapel role: architectural modeler



Sublocale Examples: Tiled Processors

```
class locale: AbstractLocale {
  const xt = 6, yt = xTiles;
  const sublocGrid: [0..#xt, 0..#yt] tiledLoc = ...;
  const allSublocs: [0..#xt*yt] tiledLoc = ...;
  // tasking interface
  // memory interface
```

class tiledLoc: AbstractLocale {
 // tasking interface
 // memory interface



Tilera Tile-Gx



Sublocale Examples: Hybrid Processors

```
class locale: AbstractLocale {
  const numCPUs = 2, numGPUs = 2;
  const cpus: [0..#numCPUs] cpuLoc = ...;
  const gpus: [0..#numGPUs] gpuLoc = ...;
  // tasking interface
  // memory interface
```

class cpuLoc: AbstractLocale { ... }

class gpuLoc: AbstractLocale {
 // sublocales for different
 // memory types, thread blocks...?
 // tasking, memory interfaces







"How can I help Chapel?"

- Evolve from "How will they do it?" to "How will we do it?"
- Let our management (and yours) know what you think about it
- Help find funding to pursue Exascale-/DOE-specific challenges
- Help us tap into the mainstream/open-source community
- Become a collaborator who is truly interested in pitching in
 - Computer scientists with mutually beneficial technologies
 - Industry partners
 - Application experts interested in pair programming studies





Sample Pair-Programming Study: LLNL/LULESH

Apr 2011: LLNL expresses interest in Chapel at Salishan

made us aware of the LULESH benchmark from DARPA UHPC

Summer 2011: Cray intern ports LULESH to Chapel

• *caveat:* used structured mesh to represent data arrays

Nov 2011: Chapel team tunes LULESH for single-node performance

Dec 2011: Chapel team visits LLNL (talk, tutorial, 1-on-1 sessions)

Mar 2012: Jeff Keasler (LLNL) visits Cray to pair-program

- in one afternoon, converted from structured to unstructured mesh
- impact on code minimal (mostly in declarations) due to:
 - domains/arrays/iterators
 - rank-independent features

Apr 2012: LLNL reports on collaboration at Salishan

Apr 2012: Chapel 1.5.0 release includes current version of LULESH

Next steps: distributed sparse domains, improved scalability





Q: If, as a community, we/you were to do the audacious and produce a productive/decent parallel language, what would that process look like?





Productivity continues to be worth striving for

- to improve the lives of existing programmers, especially as the machines get more complex
- to improve our chances of attracting new users

Current programming models aren't ideal for Exascale

- and they will be difficult to fix
- however, that's not to say we should just abandon them

Past failures do not dictate future ones

Exascale represents an opportunity to move to more ideal programming models

• more productive, higher-level, machine-independent, ...



The Chapel Team (Summer 2012)









Chapel project page: <u>http://chapel.cray.com</u>

• overview, papers, presentations, language spec, ...

Chapel SourceForge page: https://sourceforge.net/projects/chapel/

release downloads, public mailing lists, code repository, ...

Blog Series (for more opinions like these):

Myths About Scalable Programming Languages:

https://www.ieeetcsc.org/activities/blog/

Upcoming Events: PGAS/PGAS-X (talks): October 10th-12th SC12 (tutorial, BoFs?): November 12th-16th







http://chapel.cray.com chapel_info@cray.com http://sourceforge.net/projects/chapel/