# **Benchmark Improvements**

Chapel Team, Cray Inc. Chapel version 1.14 October 6, 2016



COMPUTE | STORE | ANALYZE

#### **Safe Harbor Statement**

This presentation may contain forward-looking statements that are based on our current expectations. Forward looking statements may include statements about our financial guidance and expected operating results, our opportunities and future potential, our product development and new product introduction plans, our ability to expand and penetrate our addressable markets and other statements that are not historical facts. These statements are only predictions and actual results may materially vary from those projected. Please refer to Cray's documents filed with the SEC from time to time concerning factors that could affect the Company and these forward-looking statements.







- LCALS Improvements and Status
- Computer Language Benchmark Game (CLBG)
- Other Benchmark Improvements



#### **LCALS Improvements and Status**



COMPUTE | STORE | ANALYZE



# LCALS: Background

## • LCALS: Livermore Compiler Analysis Loop Suite

- Loop kernels designed to measure compiler performance
- Developed by LLNL
- https://codesign.llnl.gov/LCALS.php

LCALS Code Richard D. Hornung LCALS version 1.0 LLNL-CODE-638939 2013

#### • Three loop subsets (30 kernels total)

- Subset A: Loops representative of application codes
- Subset B: Simple, basic loops
- Subset C: Loops extracted from "Livermore Loops coded in C"
- Each kernel is run for three sizes (Short, Medium, Long)

#### • Each kernel is implemented in a number of "variants"

• RAW (traditional C usage), OpenMP, C++ template-based, etc.



# LCALS: This Effort

- All kernels/variants/sizes performance tested nightly
- Benchmark run on a 24-core Cray XC30 node
  - Compare serial performance vs. C++
  - Compare parallel performance vs. C++ with OpenMP
  - Compiled with g++ 6.2.0

#### Working toward matching or beating g++ with OpenMP

- Serial variant is on par with g++ for almost all kernels
- Parallel variant is off by 2x-3x in most cases for Long problem size
  - Best case: PIC\_2D nearly 40% faster
  - Worst case: COUPLE over 7x slower
- Short problem size parallel variant is much faster than C++ version
  - --dataParMinGranularity=1000 option limits parallelism to reasonable level Compilation/Execution commands:

chpl --fast --no-ieee-float lcals-chpl --dataParMinGranularity=1000 g++ -Ofast -fopenmp lcals.exe



# **LCALS: Improvements since 1.13**

- Removed -sassertNoSlicing config param setting
  - Subsumed by the (default-on) compiler optimization --optimize-array-indexing
- Set dataParMinGranularity to 1000
  - Small and medium loop sizes created too many tasks without this setting
- Replaced several loops over arrays with whole array assignments
- Added an output mode that is easier to parse by scripts
- Enabled performance tracking for all variants, kernels, and loop sizes
  - Now test and track the performance of all 156 variants/kernels/sizes daily
  - Graphs are online at the Chapel Performance Overview website
- Significantly increased PIC\_2D performance with improvements to atomics







ANALYZE





COMPUTE | STORE

ANALYZE





| ANALYZE

Copyright 2016 Cray Inc.

STORE

**COMPUTE** 





Copyright 2016 Cray Inc.

11





COMPUTE | STORE | ANALYZE

## LCALS Status: Parallel Performance v1.13.0

#### Long problem size





### LCALS Status: Parallel Performance v1.14.0

#### Long problem size Normalized time – Parallel Chapel vs g++/OMP parallel reference is 1.0 9 8 Normalized Time 7 6 5 4 3 2 MIT3 DEUB OUAD AP MI PIC 2D 0 PRESSURE CALC CALC CALC 2D COUPLE FIR chpl --fast Chapel parallel ■ q++ OMP --no-ieee-float g++ -Ofast -fopenmp



COMPUTE | STORE | ANALYZE

### LCALS Status: Parallel Performance v1.14.0

#### Long problem size Normalized time – Parallel Chapel vs g++/OMP parallel reference is 1.0 9 8 **Normalized Time** 7 6 5 4 3 2 PRESSURE CALC CALC CALC CALC COUPLE FIR NULADSUB AUAD RAP INT PIC 2D chpl --fast ■ q++ OMP Chapel parallel --no-ieee-float Setting -- dataParMinGranularity=1000 caused COUPLE to not use enough tasks g++ -Ofast -fopenmp



STORE Copyright 2016 Cray Inc. ANALYZE

**COMPUTE** 

# **LCALS Status: Parallel Performance**

#### Short problem size Normalized time – Parallel Chapel vs g++/OMP parallel reference is 1.0 1.2 1 Normalized Time 0.8 0.6 0.4 0.2 INTS DEUB OURD TRAP INT PIC 20 0 PRESSURE CALC CALC CALC 2D COUPLE FIR chpl --fast ■ q++ OMP Chapel parallel --no-ieee-float g++ -Ofast -fopenmp



COMPUTE | STORE | ANALYZE

# **LCALS: Next Steps**

#### Continue optimization effort for parallel kernels

- Understand the remaining parallel performance gaps
- Bring all parallel kernels in line with reference versions
- Avoid need for minGranularity setting
- Optimize the serial kernels that are still lagging
  IF\_QUAD, HYDRO\_2D still need some effort

#### • Explore more elegant Chapel loop expressions

• Make further use of whole-array operations, array slicing, etc.



#### **Computer Language Benchmark Game (CLBG)**



COMPUTE | STORE | ANALYZE



# **CLBG: Background**

#### **CLBG:** Website supporting cross-language comparisons

- based on 13 simple serial/shared-memory computations:
  - binary-trees: memory management stressor
  - chameneos-redux: tasking coordination via a shared resource
  - fannkuch-redux: compute permutations on a small array
  - fasta, k-nucleotide, regex-dna, reverse-complement: string manipulation
  - mandelbrot: compute the Mandelbrot set
  - meteor: solve a puzzle (program startup is the bottleneck for Chapel)
  - **n-body:** simulate the solar system's largest bodies (wants vectorization)
  - pidigits: compute pi (wants GMP or equivalent)
  - **spectral-norm:** compute matrix-vector operations
  - thread-ring: pass token between tasks as quickly as possible
- must follow prescribed algorithm (except meteor, which is free-form)
- website's summaries don't include three benchmarks:
  - chameneos-redux, thread-ring, meteor
  - the first two are parallel, which makes them of interest to us
  - program startup time is our bottleneck for the third, so we include it as well



#### **CLBG: Website**

#### Can sort results by execution time, code size, memory or CPU use:

The Computer Language Benchmarks Game

n-body

description

#### program source code, command-line and measurements

×	source	secs	KB	gz	cpu	cpu load
1.0	<b>C++</b> g++ #3	9.30	1,712	1763	9.29	100% 1% 1% 0%
1.0	<u>C++ g++ #8</u>	9.37	1,152	1544	9.36	1% 0% 100% 1%
1.0	<b>C</b> gcc #4	9.56	1,000	1490	9.56	1% 100% 1% 1%
1.0	<u>C++ g++ #7</u>	9.65	940	1545	9.64	100% 1% 1% 1%
1.1	Fortran Intel #5	9.79	516	1659	9.78	1% 0% 1% 100%
1.2	Ada 2005 GNAT #2	10.99	1,952	2604	10.99	0% 1% 100% 1%
1.3	<u>C++ g++ #5</u>	11.76	1,728	1749	11.75	1% 100% 1% 0%
1.9	Ada 2005 GNAT #5	18.00	2,028	2436	18.00	0% 0% 100% 1%
2.1	<u>C++ g++ #6</u>	19.20	1,096	1668	19.19	100% 1% 0% 0%
2.1	<u>C++ g++</u>	19.37	1,056	1659	19.36	1% 100% 0% 0%
2.1	Fortran Intel #2	19.84	508	1496	19.83	1% 0% 1% 100%
2.1	Fortran Intel	19.96	512	1389	19.95	0% 1% 0% 100%
2.2	<u>C++ g++ #4</u>	20.20	684	1428	20.19	0% 0% 1% 100%
2.3	C gcc #3	20.97	952	1208	20.96	1% 1% 1% 100%

#### The Computer Language Benchmarks Game

n-body

description

#### program source code, command-line and measurements

×	source		secs	KB	gz	cpu	cpu load
1.0	Chapel		21.55	20,688	962	21.55	100% 0% 1% 0%
1.1	Hack #3		24 min	114,024	1080	24 min	54% 47% 1% 1%
1.1	<b>PHP</b> #3		7 min	7,928	1082	7 min	0% 100% 0% 1%
1.2	Ruby #2		11 min	9,144	1137	11 min	0% 1% 100% 0%
1.2	Ruby JRuby	#2	292.96	706,256	1137	5 min	26% 29% 25% 24%
1.2	Fortran Inte	el #4	21.91	512	1172	21.91	1% 0% 0% 100%
1.2	C gcc		21.15	1,028	1173	21.14	1% 100% 0% 2%
1.2	C gcc #6		21.14	956	1180	21.13	100% 1% 1% 0%
1.2	Python 3	Python 3 15 min			1181	15 min	44% 13% 0% 44%
1.2	Swift #2		32.09	4,376	1192	32.08	1% 0% 100% 0%
1.2	Lua #2		8 min	2,140	1193	8 min	0% 1% 100% 0%
1.2	Swift #6	dz =	1% 0% 100% 0%				
1.2	Lua	strip comments and extra whitespace, then gzip					1% 0% 0% 100%
1.3	Swift #3						100% 0% 1% 0%



#### COMPUTE |

STORE

ANALYZE

#### **CLBG: Website**

#### Can also compare languages pair-wise:

COM PUTE

The Computer Language Benchmarks Game Chapel programs versus Swift all other Chapel programs & measurements by benchmark task performance									
regex-dn	a								
source	secs	KB	gz	cpu	cpu load				
Chapel	9.35	1,787,668	468	18.46	100% 15% 15% 69%				
Swift	103.03	270,796	712	102.95	1% 1% 0% 100%				
pidigits source Chapel Swift	secs <b>1.60</b> 5.02	KB 22,088 7,524	gz 501 1017	сри 1.60 5.01	cpu load 99% 4% 2% 1% 100% 1% 1% 0%				
<u>n-body</u> source	secs	KB	gz	cpu	cpu load				
Swift	23.66	20,000	1253	23.65	100% 0% 0% 1%				
Swift	23.66	4,388	1253	23.65	100% 0% 0% 1%				



Copyright 2016 Cray Inc.

STORE

ANALYZE

# **CLBG: Website**

#### • site has a sound philosophy about too-easy answers

We want easy answers, but easy answers are often incomplete or wrong. You and I know, there's more we should understand:

details

stories

fast? conclusions

#### • yet, most readers probably still jump to conclusions

- execution time dominates default/only views of results
- it's human nature

#### • we're interested in elegance as well as performance

- elegance is obviously in the eye of the beholder
  - we compare source codes manually
  - but then use CLBG's code size metric as a quantitative stand-in
- want to be able to compare both axes simultaneously
- to that end, we used scatter plots to compare implementations



# **CLBG: Background**

#### **Background:**

- we ported these to evaluate Chapel serial/tasking performance
  - many "top" entries are more heroic than typical programmers would write
  - we strived for implementations that balance elegance with speed
- FAQ isn't particularly encouraging of adding new languages:

```
Why don't you include language X?
... my favorite language implementation?
... Microsoft® Windows® ?
```

Because I know it will take more time than I choose. Been there; done that.

Measurements of "proggit popular" language implementations like <u>Crystal</u> and <u>Nim</u> and <u>Julia</u> will attract attention and be the basis of yet another successful website (unlike more Fortran or Ada or Pascal or Lisp). So make those repeated measurements at multiple workloads, and publish them and promote them.

If you're interested in something not shown on the benchmarks game website then please take the program source code and the measurement scripts and **publish your own measurements**.



# **CLBG: This Effort**

Feb 2016: Inquired about submitting a Chapel entry

Apr 2016: Got a positive response

May 2016: Submitted first program

Sep 2016: Submitted final program

Listed on the front page:

Oct 2017: Upgraded to 1.14

The Computer Language Benchmarks Game

#### 64-bit quad core data set

Will your toy benchmark program be faster if you write it in a different programming language? It depends how you write it!

#### Which programs are fast?

Which are succinct? Which are efficient?

Ada	Ada <u>C</u> Chape		el	Clojure		<u>C#</u>	<u>C++</u>	
Dart	Er]	lang	<u>F#</u>	For	tran	Go	Hack	
Haskell		Java	Jav	JavaScript		Lisp	Lua	
OCaml		Pascal F		Perl	PHP	Py	Python	
Racket		Ruby	JRuby		Rus	t S	cala	
Smalltalk			Swi	.ft	TypeScript		-	



STORE

ANALYZE

# **CLBG: Improvements due to 1.14**

### 1.14 improved many benchmarks with no code changes:

- thread-ring: benefitted from qthread native sync variables
  - climbed ~16 slots  $\Rightarrow$  now 5<sup>th</sup> fastest after Haskell, Go, F#, Scala
  - 1<sup>st</sup> most compact code followed by Ruby, Racket, Erlang, Ocaml, Python
- fannkuch-redux: benefitted from optimized array accesses
  - climbed from ~#22 to #6 in performance
  - ~1.5–2x more compact than most other top entries
- **chameneos-redux:** benefitted from tasking improvements
  - climbed from ~#11 to #8 in terms of performance
- **binary-trees:** benefitted from jemalloc improvements
  - climbed ~2 performance slots as a result
  - still ~5x off from top entries which use explicit memory pools
- **n-body:** saw marginal improvements, but climbed ~17 slots
- regex-dna, revcomp: saw marginal improvements, climbed ~3 slots
- meteor: saw marginal improvements, climbed ~1 slot
- fasta: saw marginal improvements, no change in rank



# **CLBG: Improvements due to 1.14**

#### **1.14 enabled code improvements for other benchmarks:**

- pidigits: created versions that use the 'bigint' type
  - pidigits: uses operators everywhere
  - pidigits-fast: uses methods to avoid assigning returned records
- **knucleotide:** needed bug fix due to buggy auto-'use' of 'Sort' in 1.13
  - updated to new 'Sort' interfaces while here
  - also saw performance improvements from optimized array accesses
- **binary-trees:** created an initializer-based implementation
- mandelbrot: used complex values and the dynamic domain iterator
- fasta: removed a downcast on ascii(); simplified I/O due to a bug fix
- **meteor:** made use of enum.size
- **revcomp:** removed downcasts on ascii()

#### To date, have only submitted two of the above cases

- trying to avoid maintenance fatigue
- latest versions available in <u>examples/benchmarks/shootout</u> on GitHub



# **CLBG: Status**

#### Chapel entry highlights (as of Oct 17<sup>th</sup>):

- performance rankings:
  - top entries: pidigits
  - top-5 entries: meteor-contest, thread-ring
  - top-10 entries: fannkuch-redux, chameneos-redux
  - top-20 entries: n-body, spectral-norm, binary-trees
- code compactness rankings:
  - top entries: n-body, thread-ring
  - top-5 entries: spectral-norm, pidigits
  - top-20 entries: mandelbrot, regex-dna, chameneos-redux, meteor





#### **CLBG Scatter Plots**



COMPUTE | STORE | ANALYZE



# **CLBG Scatter Plots**

#### • Made scatter plots to compare performance and code size

- created these to help us understand where we're falling short
  - e.g., helps us identify performance outliers
- compared with the languages of highest interest to our team:
  - traditional: C, C++, Fortran, Java
  - productive: Python
  - modern: Scala, Go, Rust, Swift
- each program is scaled by the fastest/smallest entries for that axis
  - In *any* language -- not restricted to the subset we're focusing on here
  - e.g., binary trees is scaled by C's time (fastest) and Ruby's size (smallest)
- data is from the CLBG repository on Oct 18th

#### • Notes:

- these only characterize the submitted programs
  - i.e., better versions could potentially be written in each language
  - however, this is the data we have to work with
- not all languages have entries for all programs



#### **CLBG Scatter Plots: Chapel Programs**



# Pairwise Language Comparison Graphs

#### The first series of graphs compares languages pairwise

• For each language, we plot...

...the fastest version of each benchmark as a circle

- ...the smallest version of each benchmark as a square
- ...the mean of each set of benchmarks as a larger square/circle
- We also plot an oval at  $1\sigma$  (a standard deviation away from the mean)
  - this provides an overall "profile" for the language's fastest/smallest entries
- The axis scales are fixed across the graphs
  - in a few cases we zoom out in the subsequent slide to display outliers



#### **CLBG Scatter Plots: Chapel vs. C**



#### **CLBG Scatter Plots: Chapel vs. C++**



#### **CLBG Scatter Plots: Chapel vs. Fortran**



(34

#### **CLBG Scatter Plots: Chapel vs. Java**



#### **CLBG Scatter Plots: Chapel vs. Java (zoom)**



#### **CLBG Scatter Plots: Chapel vs. Python**



#### CLBG Fastest Entries: Chapel vs. Python (zoom)



#### **CLBG Scatter Plots: Chapel vs. Scala**



#### **CLBG Scatter Plots: Chapel vs. Go**



#### **CLBG Scatter Plots: Chapel vs. Rust**



#### **CLBG Scatter Plots: Chapel vs. Rust (zoom)**



Ĉ

#### **CLBG Scatter Plots: Chapel vs. Swift**



#### **CLBG Scatter Plots: Chapel vs. Swift (zoom)**



# **Language Summary Plots**

- The following two graphs plot the means for all languages first: across the set of fastest entries then: across the set of most compact entries
- Note that the y-axis is logarithmic
  - otherwise, Python's inclusion flattens all data along the x-axis



#### CLBG Fastest Codes: Averages (log scale perf)





#### CLBG Smallest Codes: Averages (log scale perf)



Copyright 2016 Cray Inc.

# **Per-Benchmark Comparison Graphs**

#### The next graphs compare benchmarks across languages

- For each language, we plot its fastest and most compact version
- We connect these versions with a line to help the eye associate them
  - note that this line doesn't imply a bound or a constraint, just a visual link
- Note that the y-axis is logarithmic



## **Binary-Trees: Language Comparisons**



#### **Chameneos-redux: Language Comparisons**



#### Fannkuch-redux: Language Comparisons



## **Fasta: Language Comparisons**



Copyright 2016 Cray Inc.

## **K-Nucleotide: Language Comparisons**



### Mandelbrot: Language Comparisons



### **Meteor: Language Comparisons**



### **N-Body: Language Comparisons**



## **Pidigits: Language Comparisons**



#### **Regexdna: Language Comparisons**



#### **Reverse-Complement: Language Comparisons**



## **Spectral-Norm: Language Comparisons**



## **Thread-Ring: Language Comparisons**



#### **CLBG Comparison Graphs**



COMPUTE | STORE | ANALYZE



#### • The following graphs compare Chapel to C/C++ versions

- an update to what we did for the 1.11 release notes
- run using our team's systems:
  - 2 x 12-core Intel Xeon
  - gcc/g++
- reflects more recent HW/compiler than the official CLBG system
- yet imperfect:
  - some C/C++ entries are tuned specifically for the CLBG system
  - others rely on libraries that are not installed on our system



#### chameneosredux, all versions









mandelbrot, all versions



# **CLBG: Future Work**

- Continue improving Chapel and our entries:
  - fasta: parallelize and optimize our current version
  - **string benchmarks:** improve string operations and performance
  - **pidigits:** optimize "assign returned record" idioms
  - meteor: reduce startup time in qthreads/hwloc
  - **n-body:** enable vectorization
  - **k-nucleotide:** improve associative domain performance and features
  - mandelbrot: consider adding 'unroll' keyword to for loops
  - continue to study outliers and work on improving them

#### Publish studies that dive beyond the superficial

- e.g., show heroism of top versions, compare with more typical ones
- Consider submitting more heroic Chapel versions

# Don't lose sight of multi-locale performance work

encourage HPC community to establish a CLBG equivalent



#### **Other Benchmark Improvements**



COMPUTE | STORE | ANALYZE



#### **Other Benchmark Improvements**

#### Switched ISx to use the low-level PCG interface

• results in identical data sets as the reference version



COMPUTE | STORE | ANALYZE



## **Legal Disclaimer**

Information in this document is provided in connection with Cray Inc. products. No license, express or implied, to any intellectual property rights is granted by this document.

Cray Inc. may make changes to specifications and product descriptions at any time, without notice.

All products, dates and figures specified are preliminary based on current expectations, and are subject to change without notice.

Cray hardware and software products may contain design defects or errors known as errata, which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Cray uses codenames internally to identify products that are in development and not yet publically announced for release. Customers and other third parties are not authorized by Cray Inc. to use codenames in advertising, promotion or marketing and any use of Cray Inc. internal codenames is at the sole risk of the user.

Performance tests and ratings are measured using specific systems and/or components and reflect the approximate performance of Cray Inc. products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance.

The following are trademarks of Cray Inc. and are registered in the United States and other countries: CRAY and design, SONEXION, and URIKA. The following are trademarks of Cray Inc.: ACE, APPRENTICE2, CHAPEL, CLUSTER CONNECT, CRAYPAT, CRAYPORT, ECOPHLEX, LIBSCI, NODEKARE, THREADSTORM. The following system family marks, and associated model number marks, are trademarks of Cray Inc.: CS, CX, XC, XE, XK, XMT, and XT. The registered trademark LINUX is used pursuant to a sublicense from LMI, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis. Other trademarks used in this document are the property of their respective owners.





