

Workshop on Batched, Reproducible, and Reduced Precision BLAS

Sponsored in part by



Workshop

- Extending the Basic Linear Algebra Software Library (BLAS)
 - Batched, Reproducible, and Reduced Precision
- Investigate extending the current BLAS to provide **greater parallelism for small sizes** and **reduced precision** support.
- This is the second workshop to date
 - May 18th – 19th, 2016 in Knoxville



We have been here before....

Basic Linear Algebra Subprograms Technical (BLAST) Forum
Standard

Basic Linear Algebra Subprograms Technical (BLAST) Forum

August 21, 2001

- 300 pages
- Perhaps too much
- Google
 - “BLAS Technical Forum”

BLAS Technical Forum Standard

Complete document -- pdf/ps
Individual Chapters
Chapter 1: Introduction -- pdf/ps
Chapter 2: Dense and Banded BLAS -- pdf/ps
Chapter 3: Sparse BLAS -- pdf/ps
Chapter 4: Extended and Mixed-Precision BLAS -- pdf/ps
Appendix -- pdf/ps
Legacy BLAS: C Interface to the Legacy BLAS -- pdf/ps
Journal of Development: Environmental Routines -- pdf/ps
Journal of Development: Distributed-Memory BLAS -- pdf/ps
Journal of Development: Fortran 95 Thin BLAS -- pdf/ps
Journal of Development: Interval BLAS -- pdf/ps
Bibliography -- pdf/ps
Index -- pdf/ps

Today

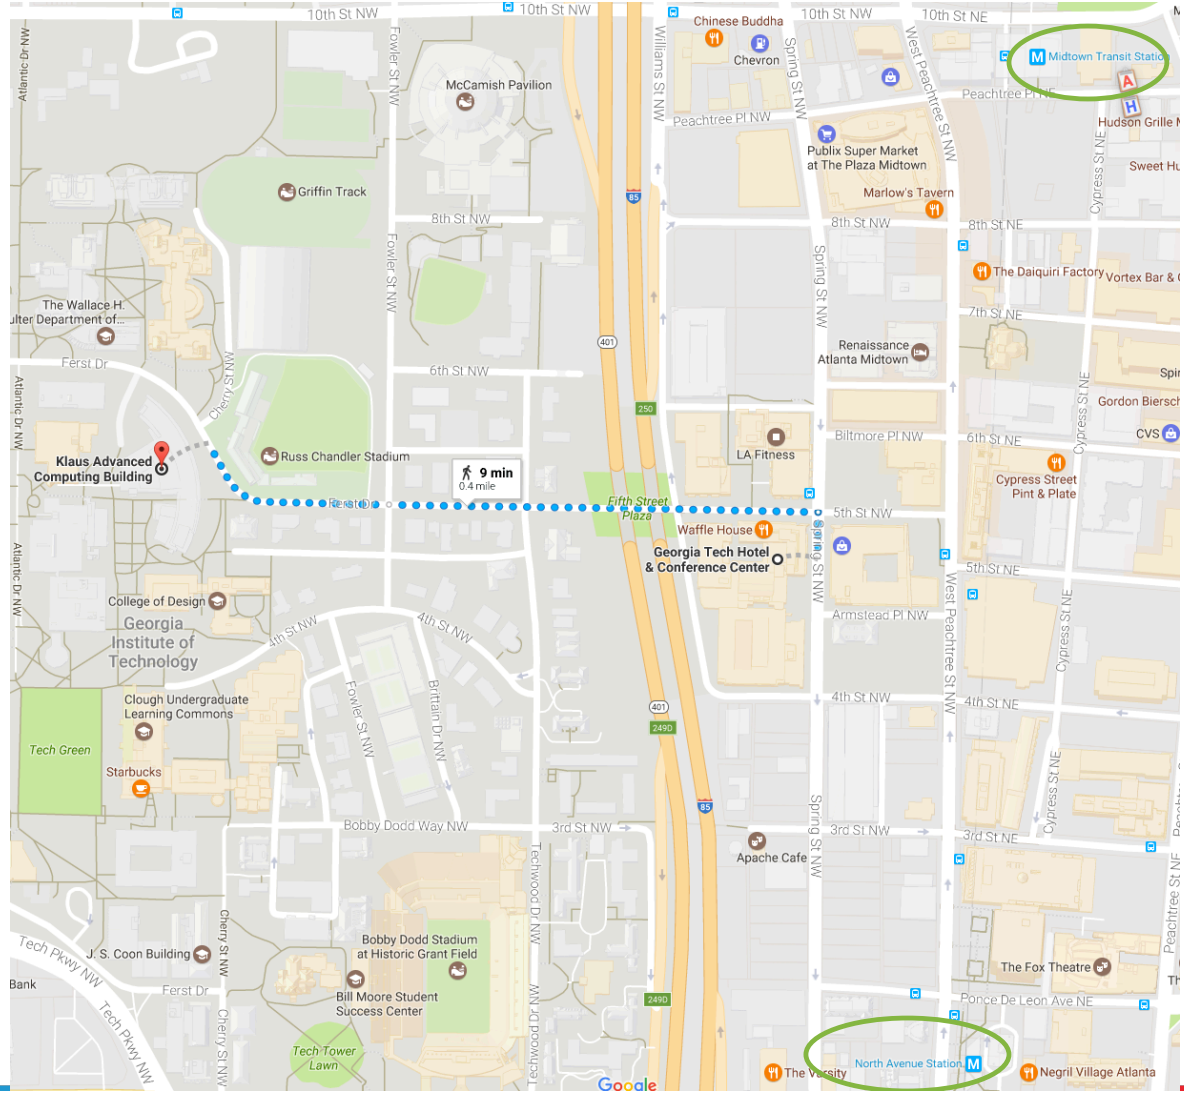
Thursday	February 23rd	
6:00 PM -8:00 PM	Reception; Georgia Tech Hotel; Conference Room E; 800 Spring St NW, Atlanta, GA; Sponsored by Intel	
Friday	February 23rd	Klaus Building; Room 1116E
8:00 AM	<i>Breakfast available</i>	
9:00 AM	Welcome& Introduction of Participants	Jack Dongarra, UTK
9:30	Report on the first Workshop, May 2016 in Knoxville, TN	Sven Hammarling, U of Manchester
10:00	Standardizing the Batched BLAS API and Memory Layout	Sam Relton, U of Manchester
10:30	Autotuning	Jakub Kurzak and Piotr Luszczek, UTK
11:00	<i>Break</i>	Room
11:30	A proposed modification to the batch BLAS interface	Ahmad Ahmad, UTK
12:00	Reproducible BLAS & Discussion	Jim Demmel, UCB
12:30	Update on the XBLAS	Greg Henry, Intel
1:00	<i>Lunch provided</i>	Room
1:30	MAGMA Batched Computations: Approaches and Applications	Stan Tomov, UTK
2:00	High Performance Design of Batched Tensor Computations: Performance Analysis, Modeling, Tuning and Optimization	Azzam Haidar, UTK
2:30	Tensor Batched BLAS	Paul Springer, Aachen
3:00	Sparse Matrix-Matrix to Batched BLAS	Siva Rajamanickam, SNL
3:30	<i>Break</i>	Room
4:00	Status of the NLAFFET Project	Bo Kagstrom, Umea U
4:30	Integer GEMM	Maurat Gurney, Intel
7:00	<i>Dinner Alma Cocina; 191 Peachtree Street NE, Atlanta</i>	Sponsored by Intel

Tomorrow

Saturday	February 25th	February 25 th ; Room 1116E
8:00 AM	<i>Breakfast available</i>	
9:00 AM	Half Precision Benchmarks	Piotr Luszczek, UTK
9:30	Batched Gauss-Jordan and batched Gauss-Huard for Block-Jacobi Preconditioning	Hartwig Anzt, UTK
10:00	Exploiting Batched Operation in Applications	David Keyes and Hatem Ltaief, KAUST
10:30	Auto-tuning Work for the QR Factorization Kernel	Wissam Lakhdar, Texas A&M
11:00	<i>Break</i>	Room
	<i>Vendor presentations</i>	
11:30	Intel, Compact Batched BLAS	Tim Costa, Intel
12:00	ARM	Chris Goodyer, ARM
12:30	<i>Lunch provided</i>	Room
	<i>Vendor presentations continued</i>	
1:30	NAG	Mike Dewar, NAG
2:00	MathWorks	Pat Quinllen, MathWorks
2:30	Nvidia	Sharan Chetlur, Nvidia
3:00	Cray	Aaron Collier, Cray
3:30	Adobe	Shoaib Kamil, Adobe
4:00	Wrap up	Jack Dongarra

General Information

- Webpage: <http://bit.ly/Batch-BLAS-2017>
- Thanks again to Intel for financial sponsorship and GATech & UTK/ ICL
 - Pradeep Dubey, Jason Riedy, Anna Stroup, and my team back at UTK
- Keep it informal, ask questions
- Wifi
 - Eduroam
 - Or see Anna for login information
- Group picture before lunch just outside this room
- Introductions



Today

Thursday	February 23rd	
6:00 PM -8:00 PM	Reception; Georgia Tech Hotel; Conference Room E; 800 Spring St NW, Atlanta, GA; Sponsored by Intel	
Friday	February 23rd	Klaus Building; Room 1116E
8:00 AM	<i>Breakfast available</i>	
9:00 AM	Welcome& Introduction of Participants	Jack Dongarra, UTK
9:30	Report on the first Workshop, May 2016 in Knoxville, TN	Sven Hammarling, U of Manchester
10:00	Standardizing the Batched BLAS API and Memory Layout	Sam Relton, U of Manchester
10:30	Autotuning	Jakub Kurzak and Piotr Luszczek, UTK
11:00	<i>Break</i>	Room
11:30	A proposed modification to the batch BLAS interface	Ahmad Ahmad, UTK
12:00	Reproducible BLAS & Discussion	Jim Demmel, UCB
12:30	Update on the XBLAS	Greg Henry, Intel
1:00	<i>Lunch provided</i>	Room
1:30	MAGMA Batched Computations: Approaches and Applications	Stan Tomov, UTK
2:00	High Performance Design of Batched Tensor Computations: Performance Analysis, Modeling, Tuning and Optimization	Azzam Haidar, UTK
2:30	Tensor Batched BLAS	Paul Springer, Aachen
3:00	Sparse Matrix-Matrix to Batched BLAS	Siva Rajamanickam, SNL
3:30	<i>Break</i>	Room
4:00	Status of the NLAFFET Project	Bo Kagstrom, Umea U
4:30	Integer GEMM	Maurat Gurney, Intel
7:00	<i>Dinner Alma Cocina; 191 Peachtree Street NE, Atlanta</i>	Sponsored by Intel

**Dinner
Tonight
7:00 pm**

***Alma Cocina;
191 Peachtree
Street NE,
Atlanta***

