Replacing Pivoting in Distributed Gaussian Elimination with Randomized Techniques

Neil Lindquist, Piotr Luszczek, Jack Dongarra

{nlindqu1,luszczek,dongarra}@icl.utk.edu

ICL INNOVATIVE COMPUTING LABORATORY

THE UNIVERSITY OF TENNESSEE KNOXVILLE

Gaussian Elimination

- Solve $Ax = b$
- Factor A into L, U triangular matrices

Gaussian Elimination

- Solve $Ax = b$
- Factor A into L, U triangular matrices
- Problems:
  - Zeros on the diagonal
  - Growth factors ($\rightarrow$ cancellation in finite precision)

Gaussian Elimination

- Solve $Ax = b$
- Factor A into L, U triangular matrices
- Problems:
  - Zeros on the diagonal
  - Growth factors ($\rightarrow$ cancellation in finite precision)
- Standard solution: partial pivoting
  - Swap largest value in column onto diagonal
  - Introduces expensive data movement

Butterfly Matrix

$$B^{\langle n \rangle} = \frac{1}{\sqrt{2}} \begin{bmatrix} R_1 & R_2 \\ R_1 & -R_2 \end{bmatrix}$$

R_1, R_2 - diagonal, nonsingular matrices

ICL INNOVATIVE COMPUTING LABORATORY

THE UNIVERSITY OF TENNESSEE KNOXVILLE

Recursive Butterfly Transform

$$\begin{bmatrix} B_1^{\left\langle \frac{n}{2^{d-1}} \right\rangle} & & \\ & \ddots & \\ & & B_{2^{d-1}}^{\left\langle \frac{n}{2^{d-1}} \right\rangle} \end{bmatrix} \cdots \begin{bmatrix} B_1^{\langle n/2 \rangle} & \\ & B_2^{\langle n/2 \rangle} \end{bmatrix} B_1^{\langle n \rangle}$$

INNOVATIVE COMPUTING LABORATORY

THE UNIVERSITY OF TENNESSEE KNOXVILLE

Recursive Butterfly Transform

$$\begin{bmatrix} B_1^{\left\langle \frac{n}{2^{d-1}} \right\rangle} & & \\ & \ddots & \\ & & B_{2^{d-1}}^{\left\langle \frac{n}{2^{d-1}} \right\rangle} \end{bmatrix} \cdots \begin{bmatrix} B_1^{\langle n/2 \rangle} & \\ & B_2^{\langle n/2 \rangle} \end{bmatrix} B_1^{\langle n \rangle}$$

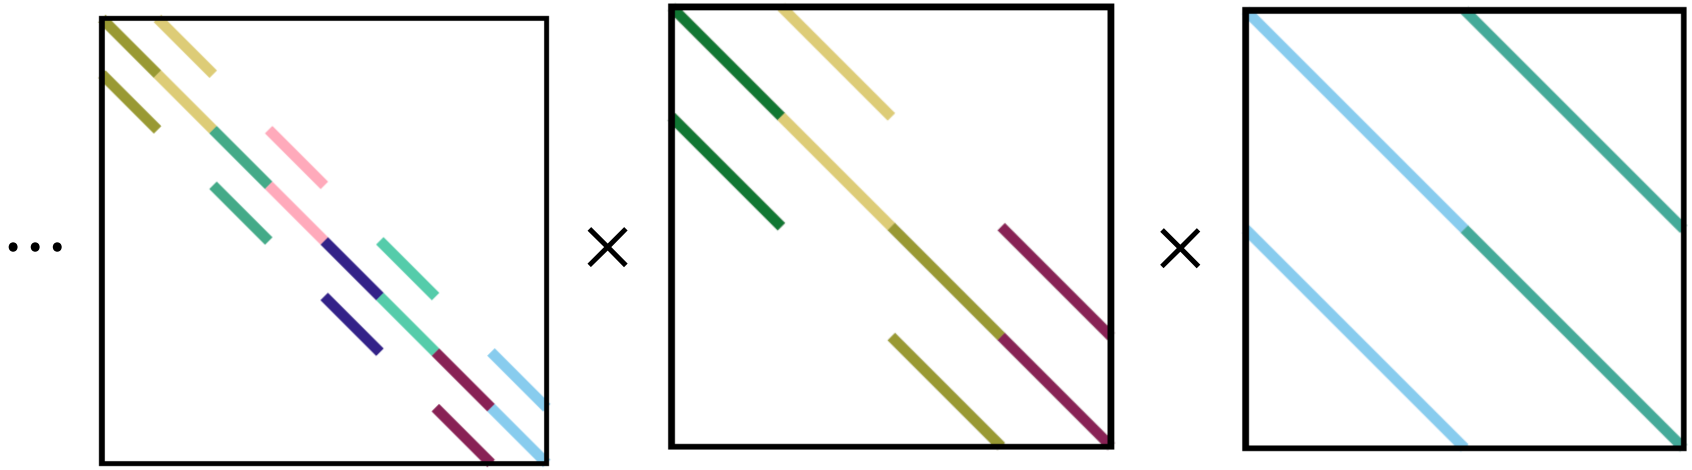

Relation to the Fast Fourier Transform

- FFT is a RBT followed by a permutation

$$B^{\langle n \rangle} = \frac{1}{\sqrt{2}} \begin{bmatrix} I & \Omega \\ I & -\Omega \end{bmatrix}$$

$$\Omega = \mathrm{diag}(1, \omega_{2n}, \omega_{2n}^{2}, \ldots, \omega_{2n}^{n-1})$$

ICL INNOVATIVE COMPUTING LABORATORY

THE UNIVERSITY OF TENNESSEE KNOXVILLE

RBT-based Solver

$\mathcal{U}, \mathcal{V}$ – recursive butterfly transforms

Write $Ax = b$ as $(\mathcal{U}^T A \mathcal{V})(\mathcal{V}^{-1} x) = (\mathcal{U}^T b)$

1. $A' = \mathcal{U}^T A \mathcal{V}$
2. $b' = \mathcal{U}^T b$
3. $LU = A'$
4. $x' = U^{-1} L^{-1} b'$
5. $x = \mathcal{V} x'$

Overheads

- $4dn^2$ FLOP to apply a 2-sided RBT
- $2dn$ FLOP to apply an RBT to a vector

- dn elements to store per transform

Implementation Details

- Using SLATE (Software for Linear Algebra Targeting Exascale)
  - Distributed, GPU-accelerated, dense linear algebra
  - Successor to ScaLAPACK
- Recursive transform depth of 2
- 1 step iterative refinement

Implementation Details

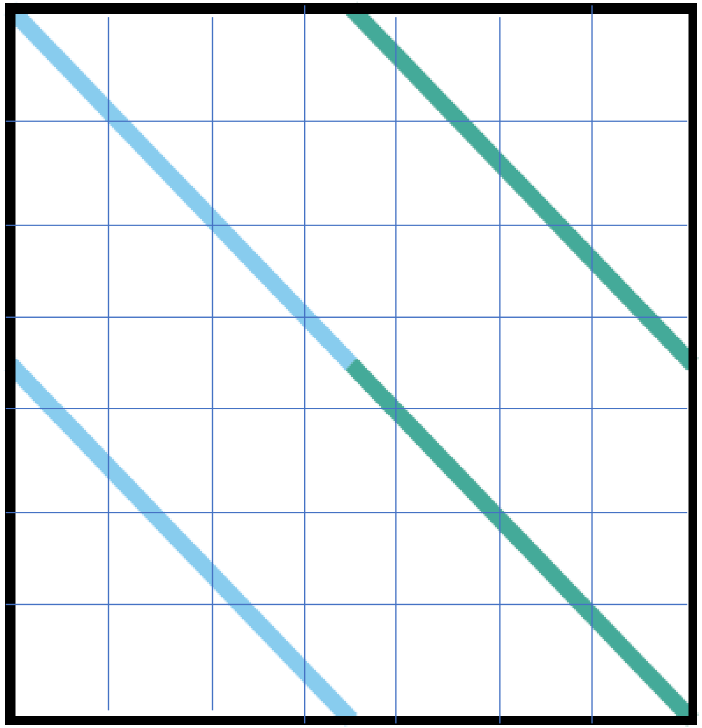

- Matrix might not be a multiple of 2^d
- Distribution may not align to butterflies

Implementation Details

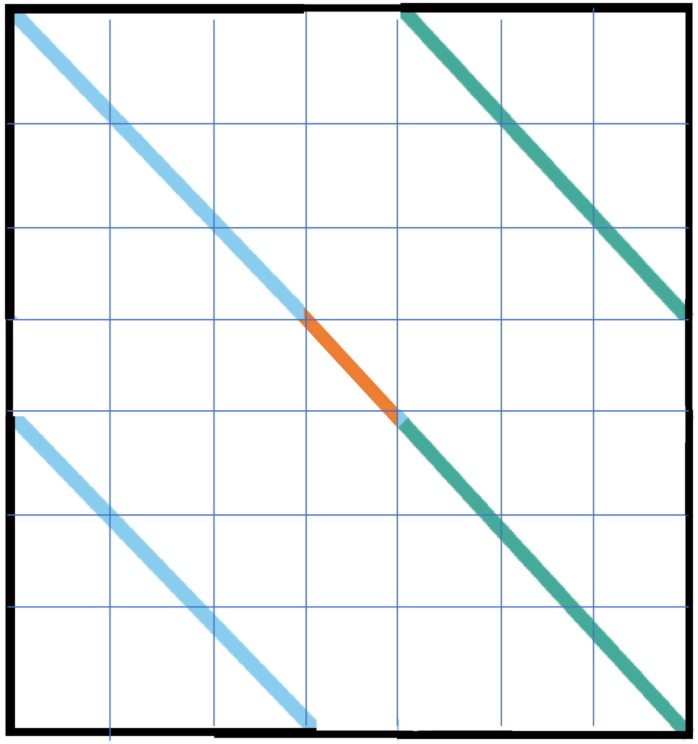

- Matrix might not be a multiple of 2^d
- Distribution may not align to butterflies

- Minor modification to avoid

Accuracy results

$$n = 100\,000, \ \text{error} = \frac{\|r\|_1}{\|A\|_1 \cdot \|x\|_1}$$

	Partial pivoting	RBT Solver Refined	RBT Solver	No pivoting Refined	No pivoting
Random [0,1]	2.34×10^{-15}	2.66×10^{-17}	3.97×10^{-12}	2.23×10^{-05}	1.10×10^{-06}
Random [-1, 1]	3.23×10^{-15}	3.18×10^{-17}	1.43×10^{-11}	9.08×10^{-17}	7.94×10^{-11}
Random Normal	4.74×10^{-15}	2.74×10^{-17}	2.05×10^{-11}	3.01×10^{-05}	7.87×10^{-05}
Random {0,1}	3.39×10^{-15}	2.37×10^{-17}	1.84×10^{-11}	NA	NA
circul	1.28×10^{-17}	9.70×10^{-19}	6.47×10^{-18}	9.85×10^{-19}	1.69×10^{-14}
fiedler	1.01×10^{-18}	4.59×10^{-19}	1.99×10^{-17}	NA	NA
gfpp	NA	2.79×10^{-19}	5.06×10^{-18}	NA	NA
orthog	5.70×10^{-16}	1.06×10^{-2}	1.14×10^{-2}	6.16×10^{-2}	7.59×10^{-2}
riemann	3.17×10^{-17}	4.17×10^{-11}	2.08×10^{-8}	4.47×10^{-19}	9.72×10^{-16}
ris	1.53×10^{-15}	1.23×10^{-1}	1.23×10^{-1}	1.23×10^{-1}	1.23×10^{-1}

ICL INNOVATIVE COMPUTING LABORATORY

THE UNIVERSITY OF TENNESSEE KNOXVILLE

Performance results

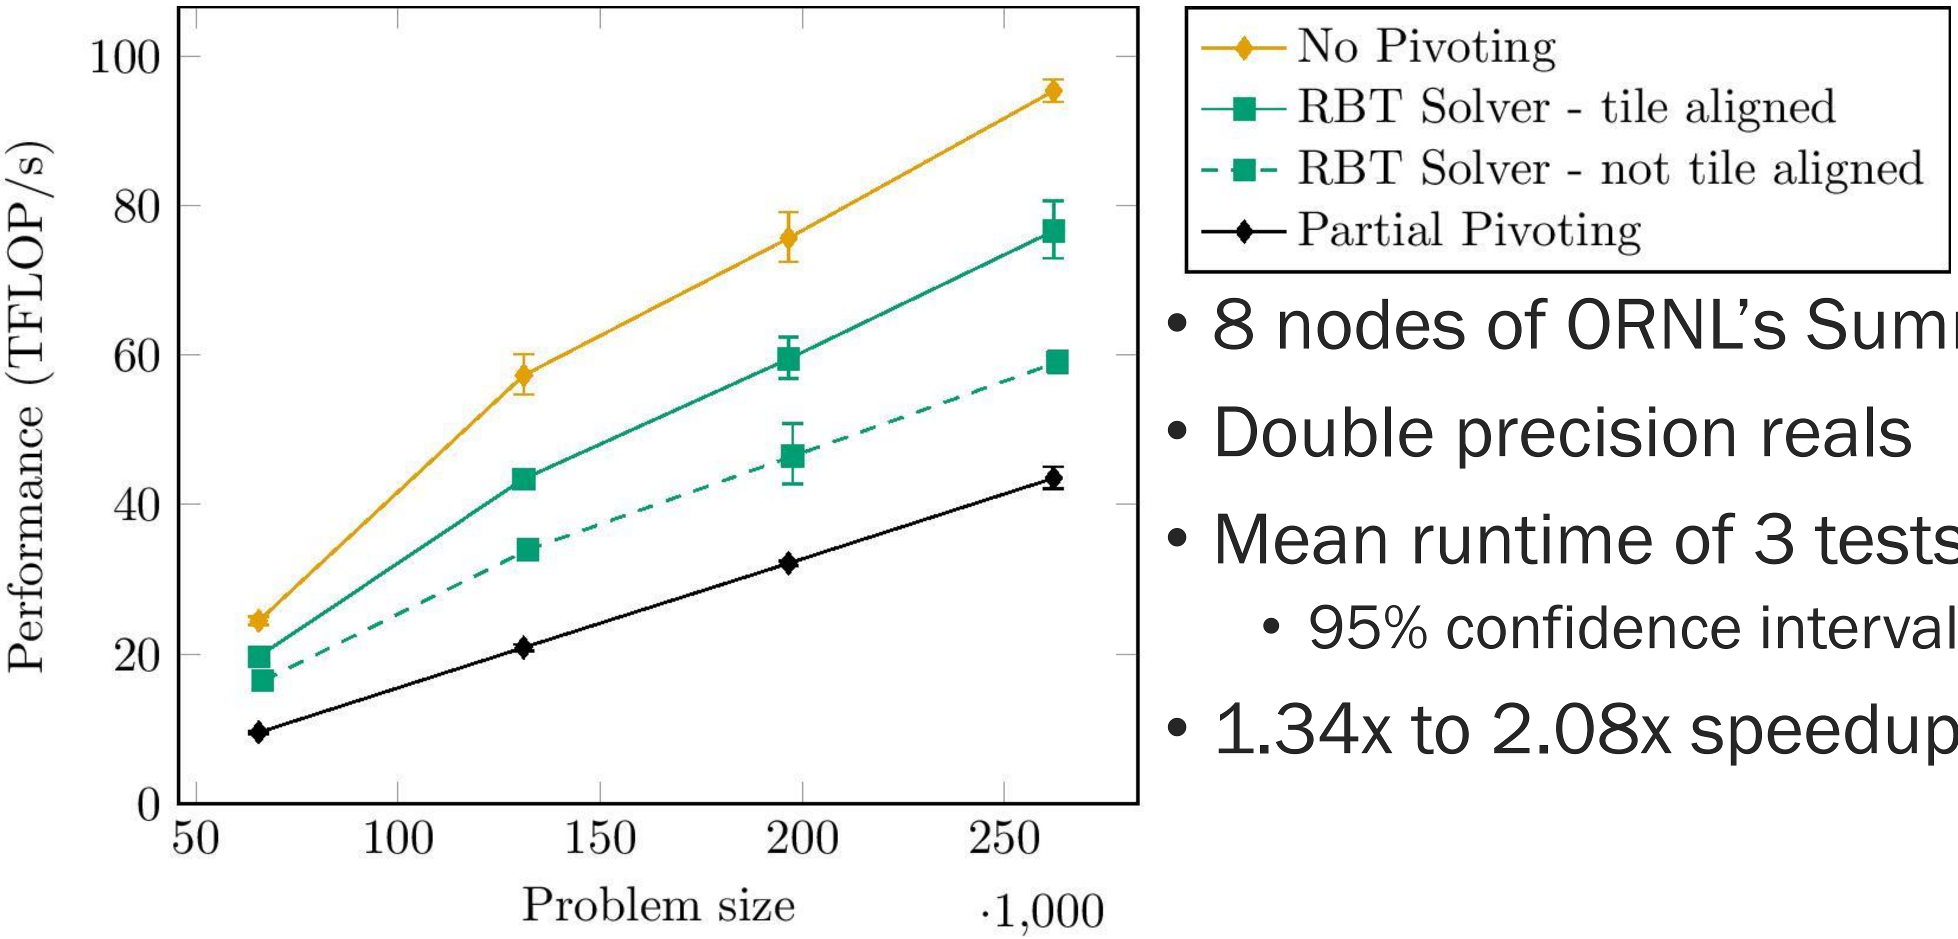

- 8 nodes of ORNL's Summit
- Double precision reals
- Mean runtime of 3 tests
  - 95% confidence interval
- 1.34x to 2.08x speedup

ICL INNOVATIVE COMPUTING LABORATORY

THE UNIVERSITY OF TENNESSEE KNOXVILLE

Conclusions

- Recursive Butterfly Transforms can replace pivoting in Gaussian Elimination
  - Often as accurate
  - 1.34x to 2.08x speedup

Acknowledgements

- This material is based upon work partially supported by the National Science Foundation under OAC Grant No. 2004541.
- This research used the computational resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725, provided by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration.

ICL INNOVATIVE COMPUTING LABORATORY

THE UNIVERSITY OF TENNESSEE KNOXVILLE

ICL
INNOVATIVE
COMPUTING LABORATORY

THE UNIVERSITY OF
TENNESSEE
KNOXVILLE

Experiment Configuration

- 8 nodes of Summit
- Each node
  - 2 processes
  - 2 22-core IBM POWER 9 CPUs
  - 6 NVIDIA Volta V100 GPUs
- Double precision reals
- Spectrum MPI 10.3.1.2, ESSL 6.1.0-2
- GCC 8.1.1, CUDA 10.1.243

ICL
INNOVATIVE COMPUTING LABORATORY
THE UNIVERSITY OF TENNESSEE KNOXVILLE

Performance results

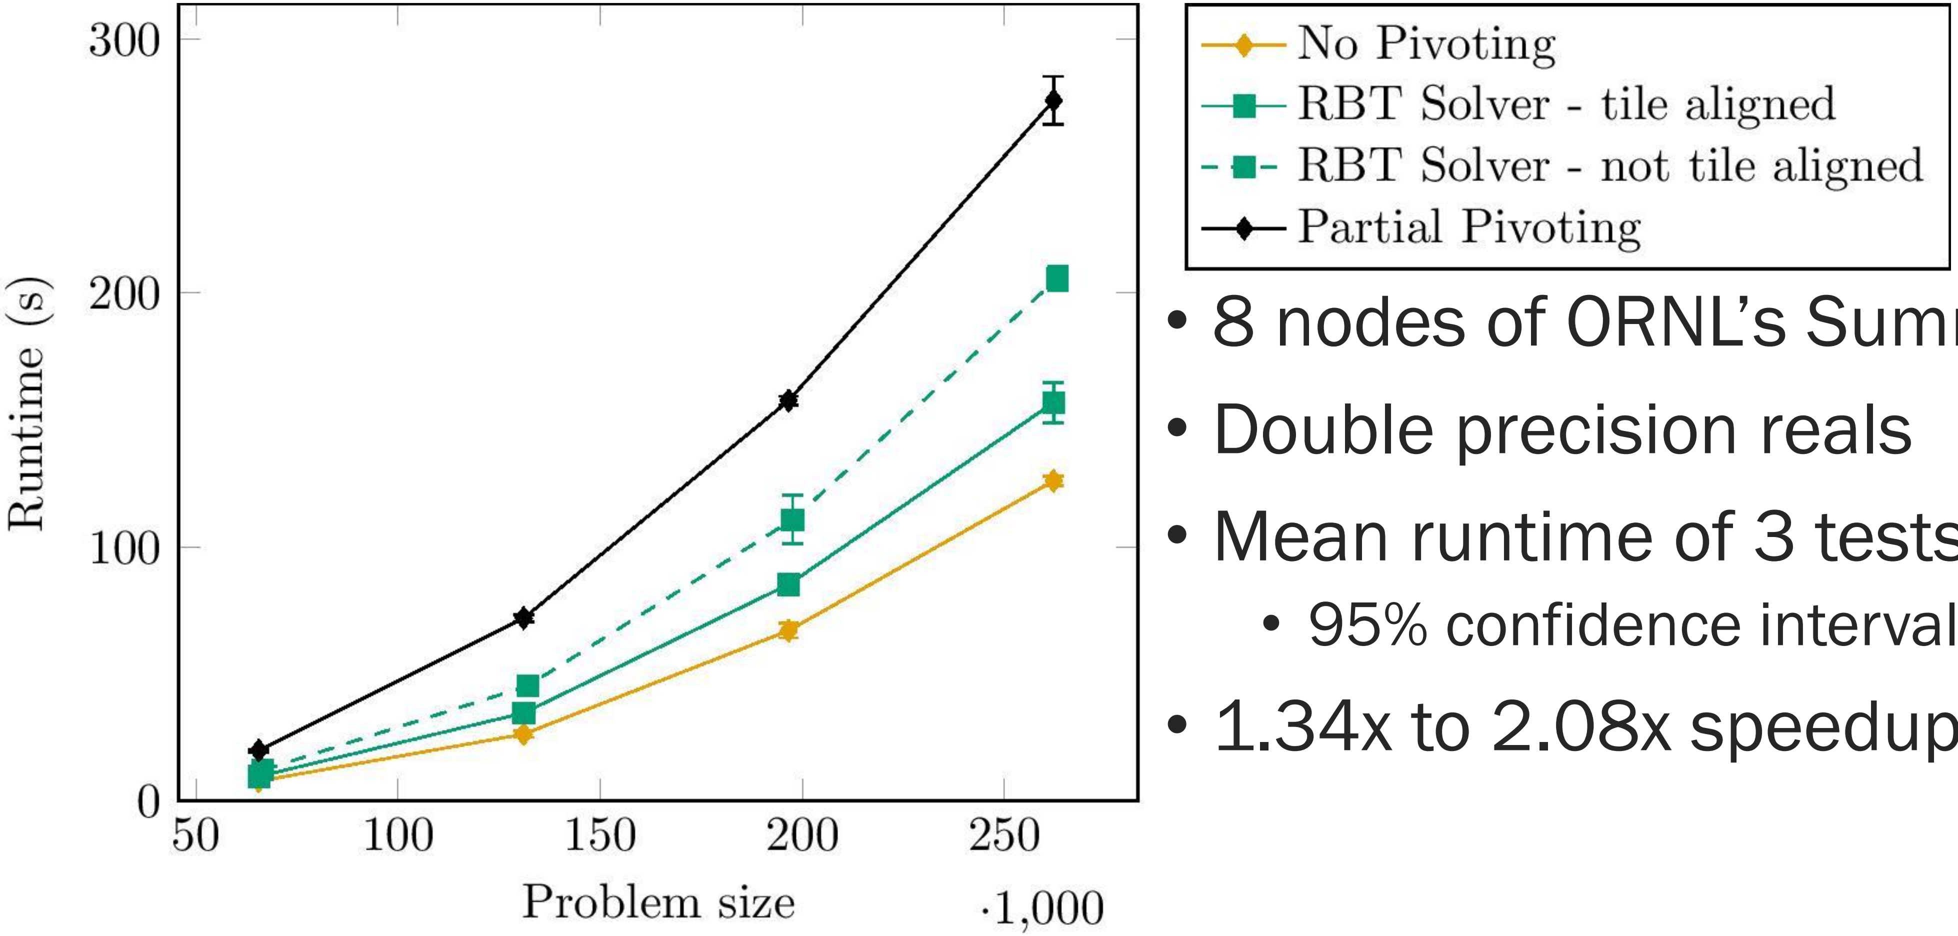

- 8 nodes of ORNL's Summit
- Double precision reals
- Mean runtime of 3 tests
  - 95% confidence interval
- 1.34x to 2.08x speedup

ICL INNOVATIVE COMPUTING LABORATORY

THE UNIVERSITY OF TENNESSEE KNOXVILLE