Eigenvalue Solvers Considerations on Modern HPC Hardware Platforms

Piotr Luszczek

In collaboration with:
Mark Gates, Azzam Haidar, Jakub Kurzak, Hatem Ltaief, Stanimire Tomov, and Jack Dongarra

Software Projects – Multicores and Accelerators

netlib.org

- LAPACK
- ScaLAPACK
- BLAS
- CBLAS
- LAPACKE

classic software for multicore

icl.utk.edu/research

- PLASMA — dense linear algebra (multicore)
- MAGMA — dense linear algebra (accelerators)
- SLATE — dense linear algebra (distributed memory + accelerators)

new software for multicore and accelerators

- OpenMP — scheduling (multicore)
- PaRSEC — scheduling (distributed memory)

dynamic runtime schedulers

Basic Linear Algebra Subroutines

- Level 1 BLAS - vector ops
  - O(n) data and flops
  - Memory bound:
    - As high as 2 Gflop/s on Skylake (not per core)

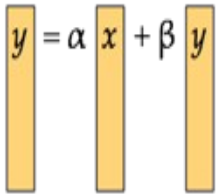

- Level 2 BLAS - mat-vec ops
  - $O(n^2)$ data and flops
  - Memory bound:
    - As high as 4 Gflop/s on Skylake (not per core)

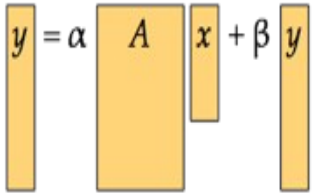

- Level 3 BLAS - mat-mat ops
  - $O(n^2)$ data, $O(n^3)$ flops
  - Surface-to-volume effect
  - Compute bound:
    - As high as 80 Gflop/s per core on Skylake

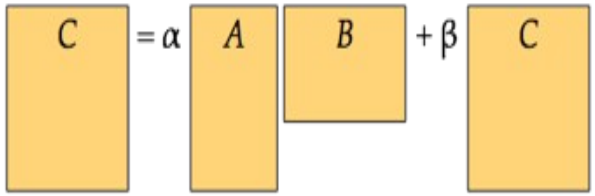

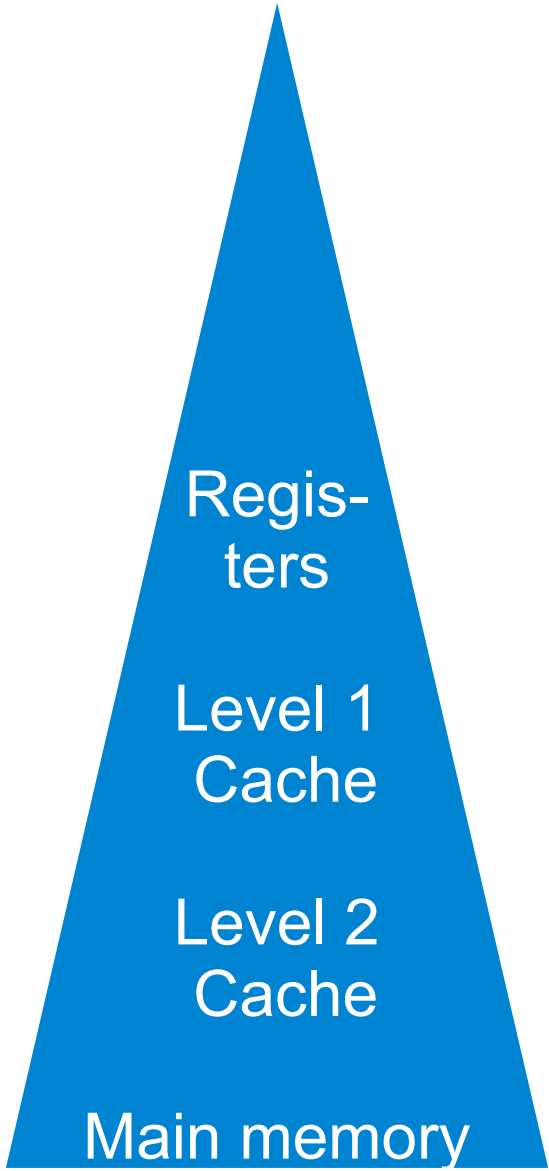

Floating Point Operations Per-Cycle Per-Core

- Most of the recent computers have FMA (Fused multiple add):
  - $x \leftarrow x + y*z$ in one cycle without intermediate rounding
- Intel Xeon (earlier models) and AMD Opteron have SSE2
  - 2 flops/cycle DP, 4 flops/cycle SP
- Intel Xeon Nehalem (’09) & Westmere (’10) have SSE4
  - 4 flops/cycle DP, 8 flops/cycle SP
- Intel Xeon Sandy Bridge(’11) and Ivy Bridge (’12) have AVX
  - 8 flops/cycle DP & 16 flops/cycle SP
- Intel Xeon Haswell (’13) & Broadwell (’14) AVX2
  - 16 flops/cycle DP & 32 flops/cycle SP
- Xeon Phi (per core) is at 16 flops/cycle DP & 32 flops/cycle SP
- Intel Xeon Skylake ('18, server) and Knight’s Landing AVX 512
  - 32 flops/cycle DP & 64 flops/cycle SP

CPU Access Latency in Clock Cycles

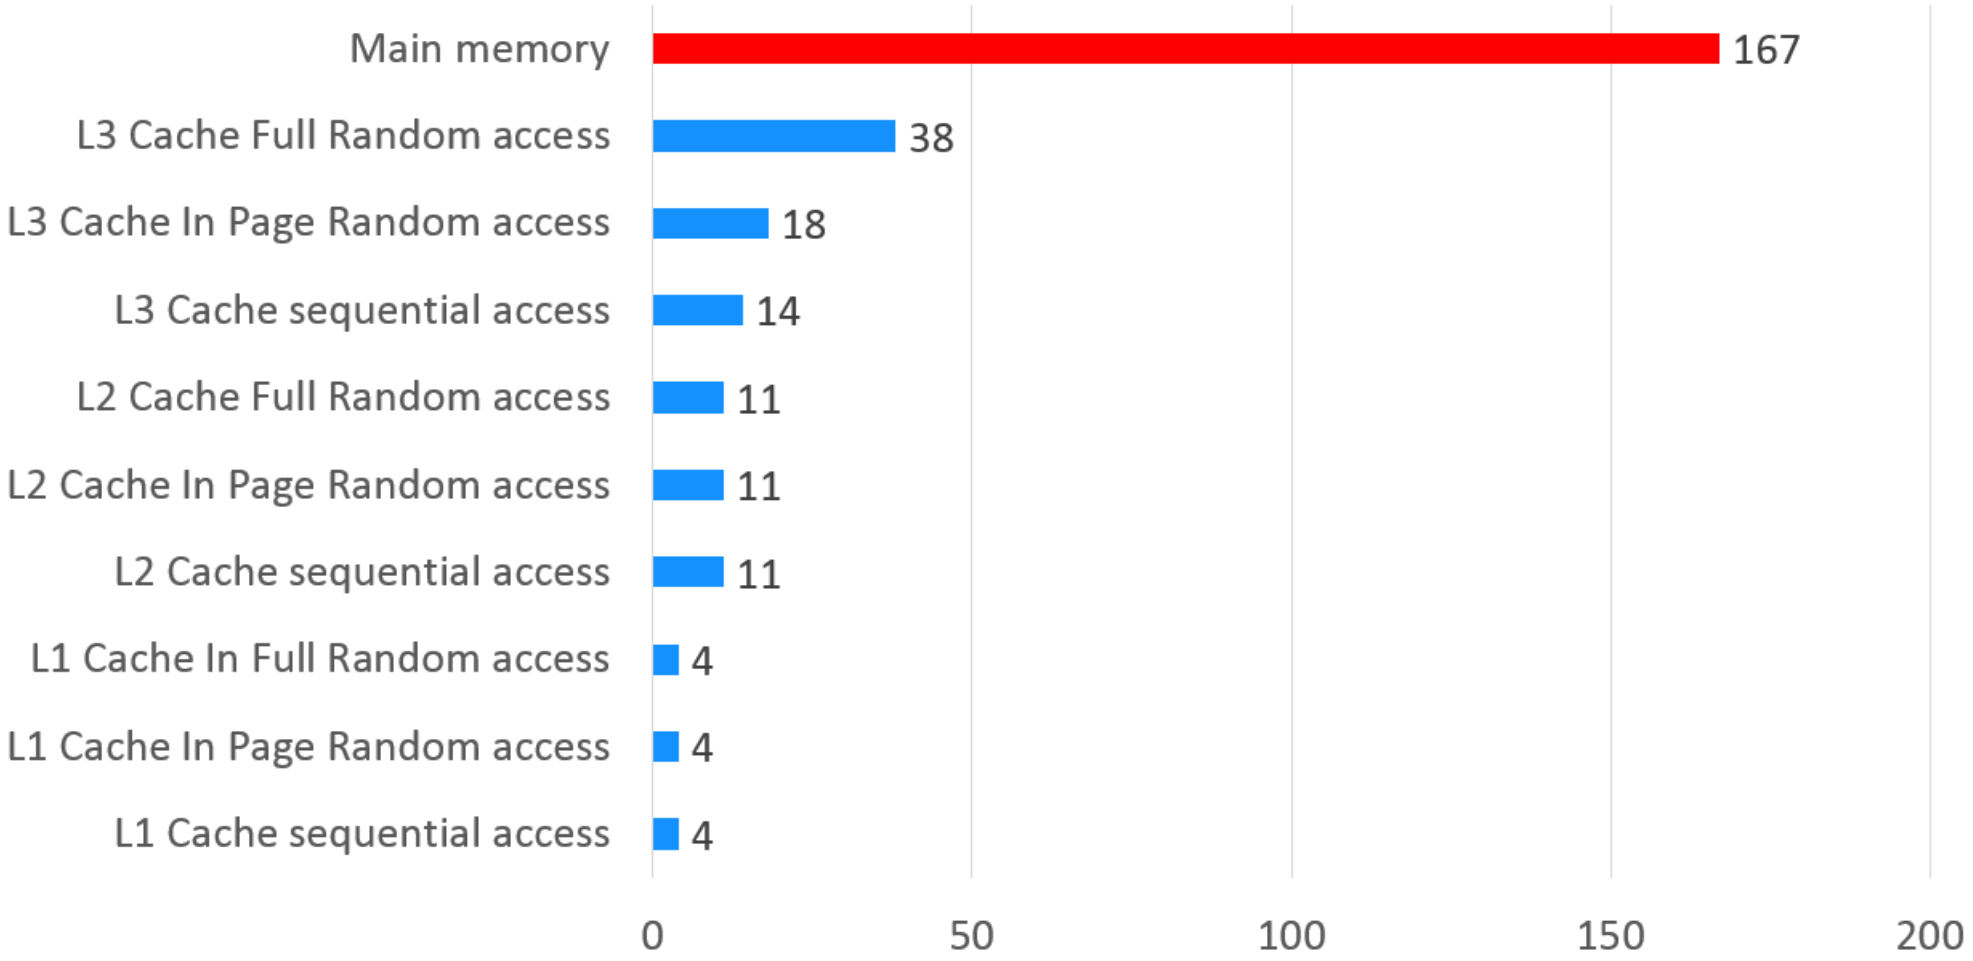

MAGMA Portability Overview

NVIDIA

GPU CENTER OF EXCELLENCE

NVIDIA's GPU Center of Excellence Program recognizes universities expanding the frontier of massively parallel computing using CUDA.

intel

Intel Parallel Computing Center

The objective of the Innovative Computing Laboratory's IPCC is the development and optimization of numerical linear algebra libraries and technologies for applications, while tackling current challenges in heterogeneous Intel® Xeon Phi™ coprocessor-based High Performance Computing.

AMD

Long-term collaboration and support on the development of clMAGMA, the OpenCL™ port of MAGMA.

FEATURES AND SUPPORT

- **MAGMA 2.2** FOR **CUDA**
- **clMAGMA 1.4** FOR **OpenCL**
- **MAGMA MIC 1.4** FOR **Intel Xeon Phi**

CUDA	OpenCL	Intel Xeon Phi	
●	●	●	Linear system solvers
●	●	●	Eigenvalue problem solvers
●	●		Auxiliary BLAS
●			Batched LA
●		●	Sparse LA
●	●	●	CPU Interface
●	●	●	GPU Interface
●	●	●	Multiple precision support
●			Non-GPU-resident factorizations
●	●	●	Multicore and multi-GPU support
●	●	●	LAPACK testing
●	●	●	Linux
●	●		Windows
●	●		Mac OS

MAGMA Performance – LU Factorization

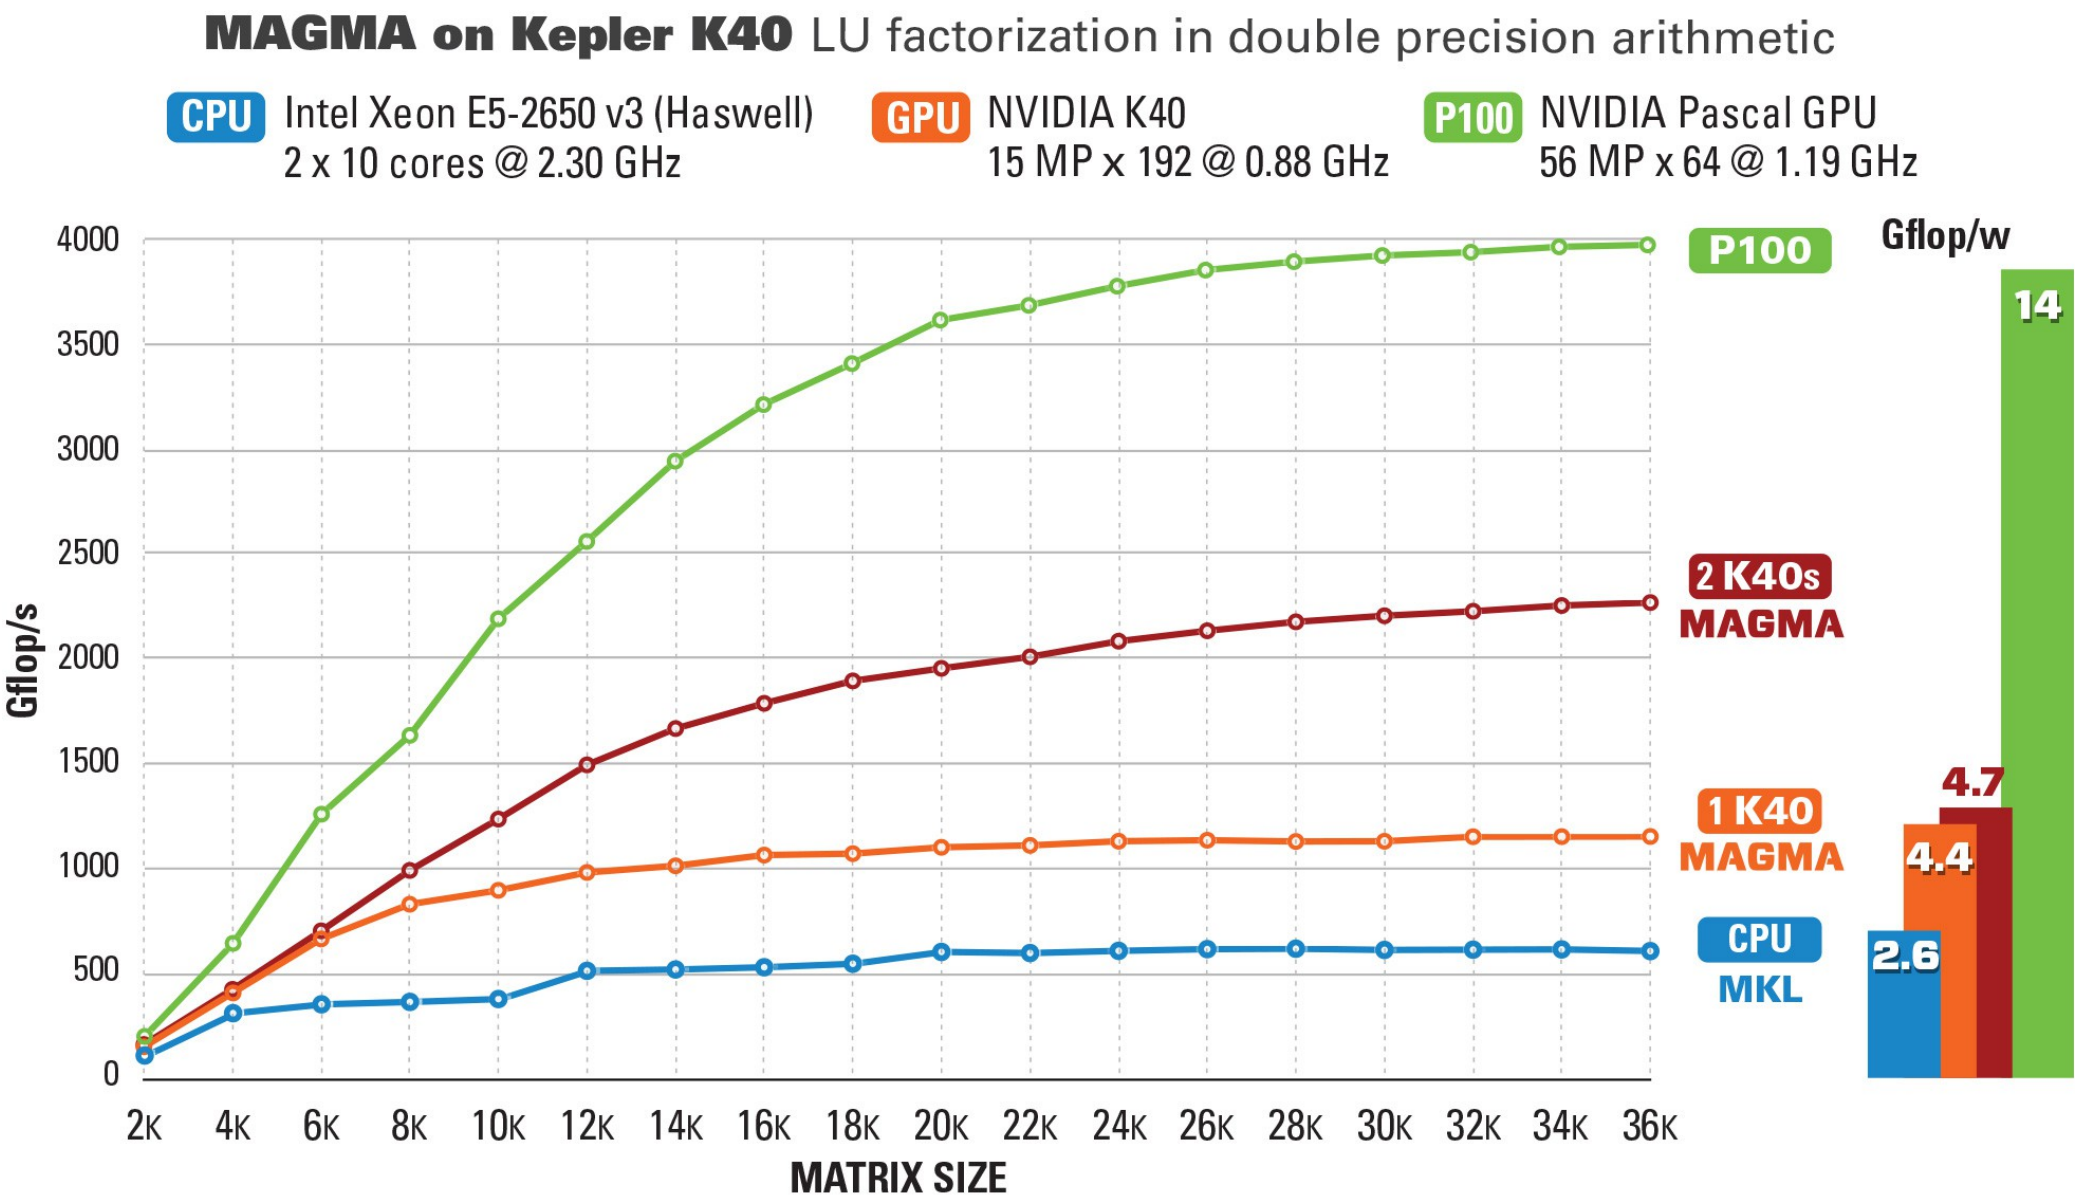

MAGMA Performance on NVIDIA Volta

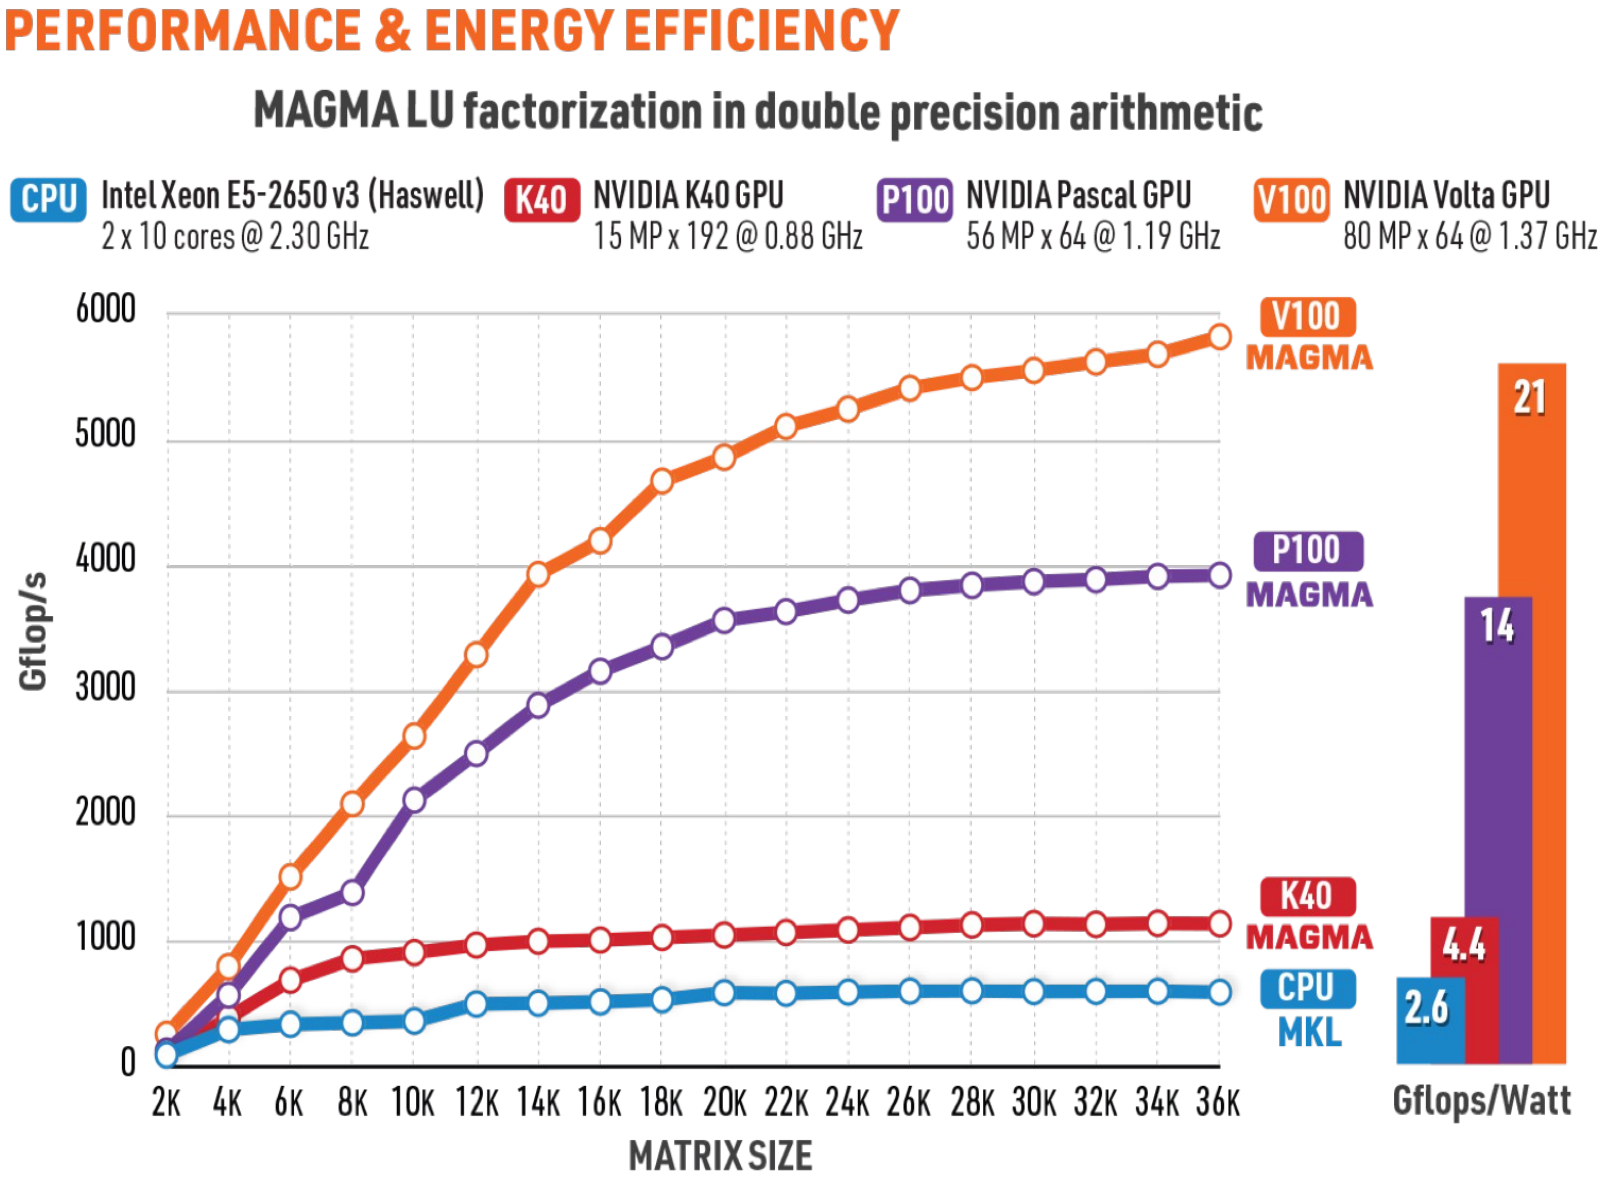

MAGMA's DGETRF Performance on Kepler K20c

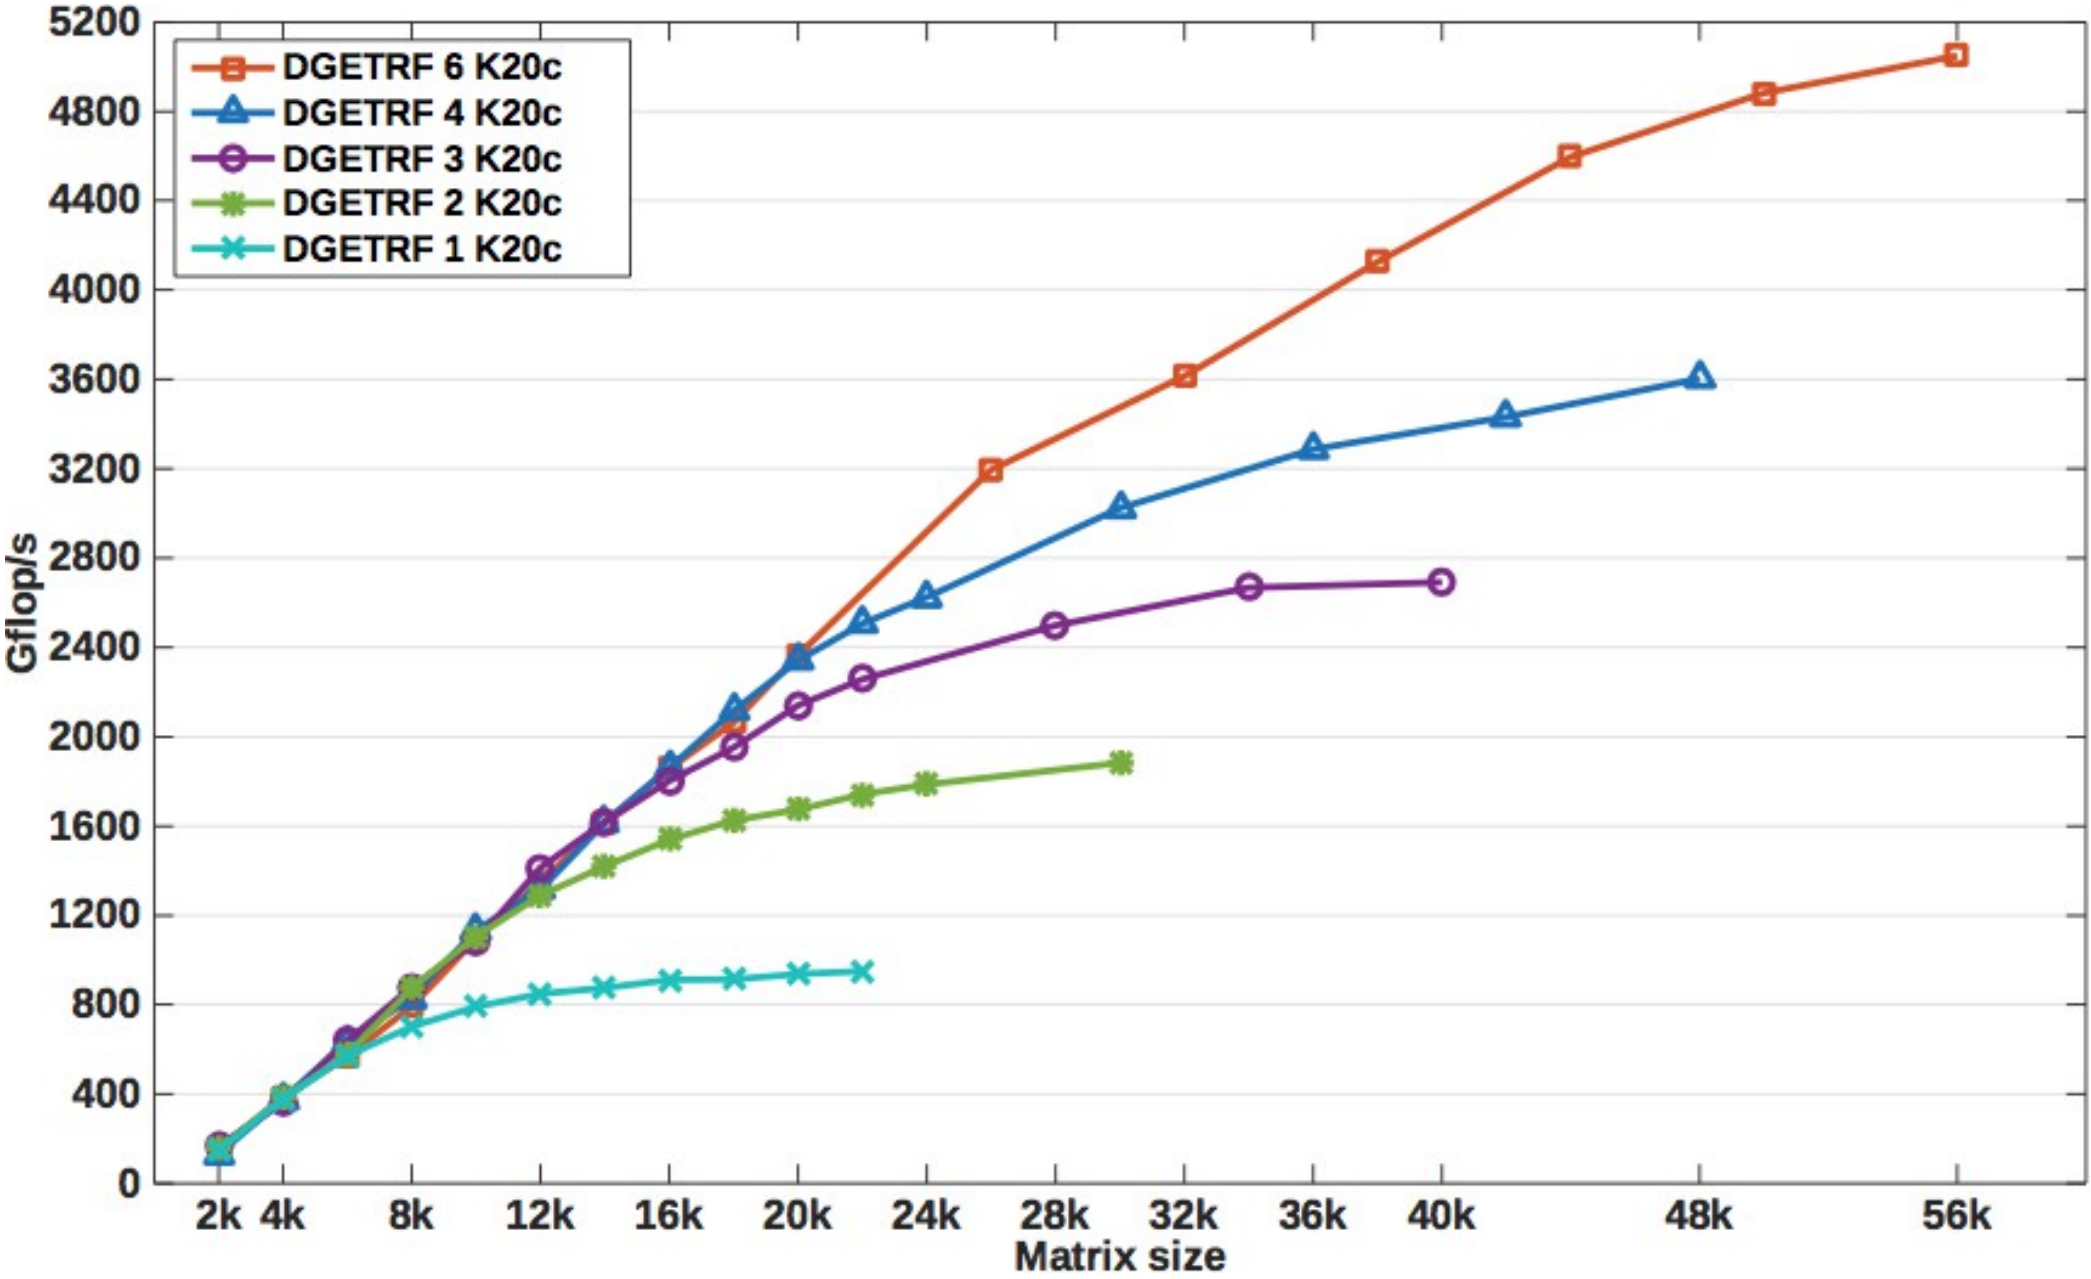

MAGMA's DGETRF Performance on Phi KNC

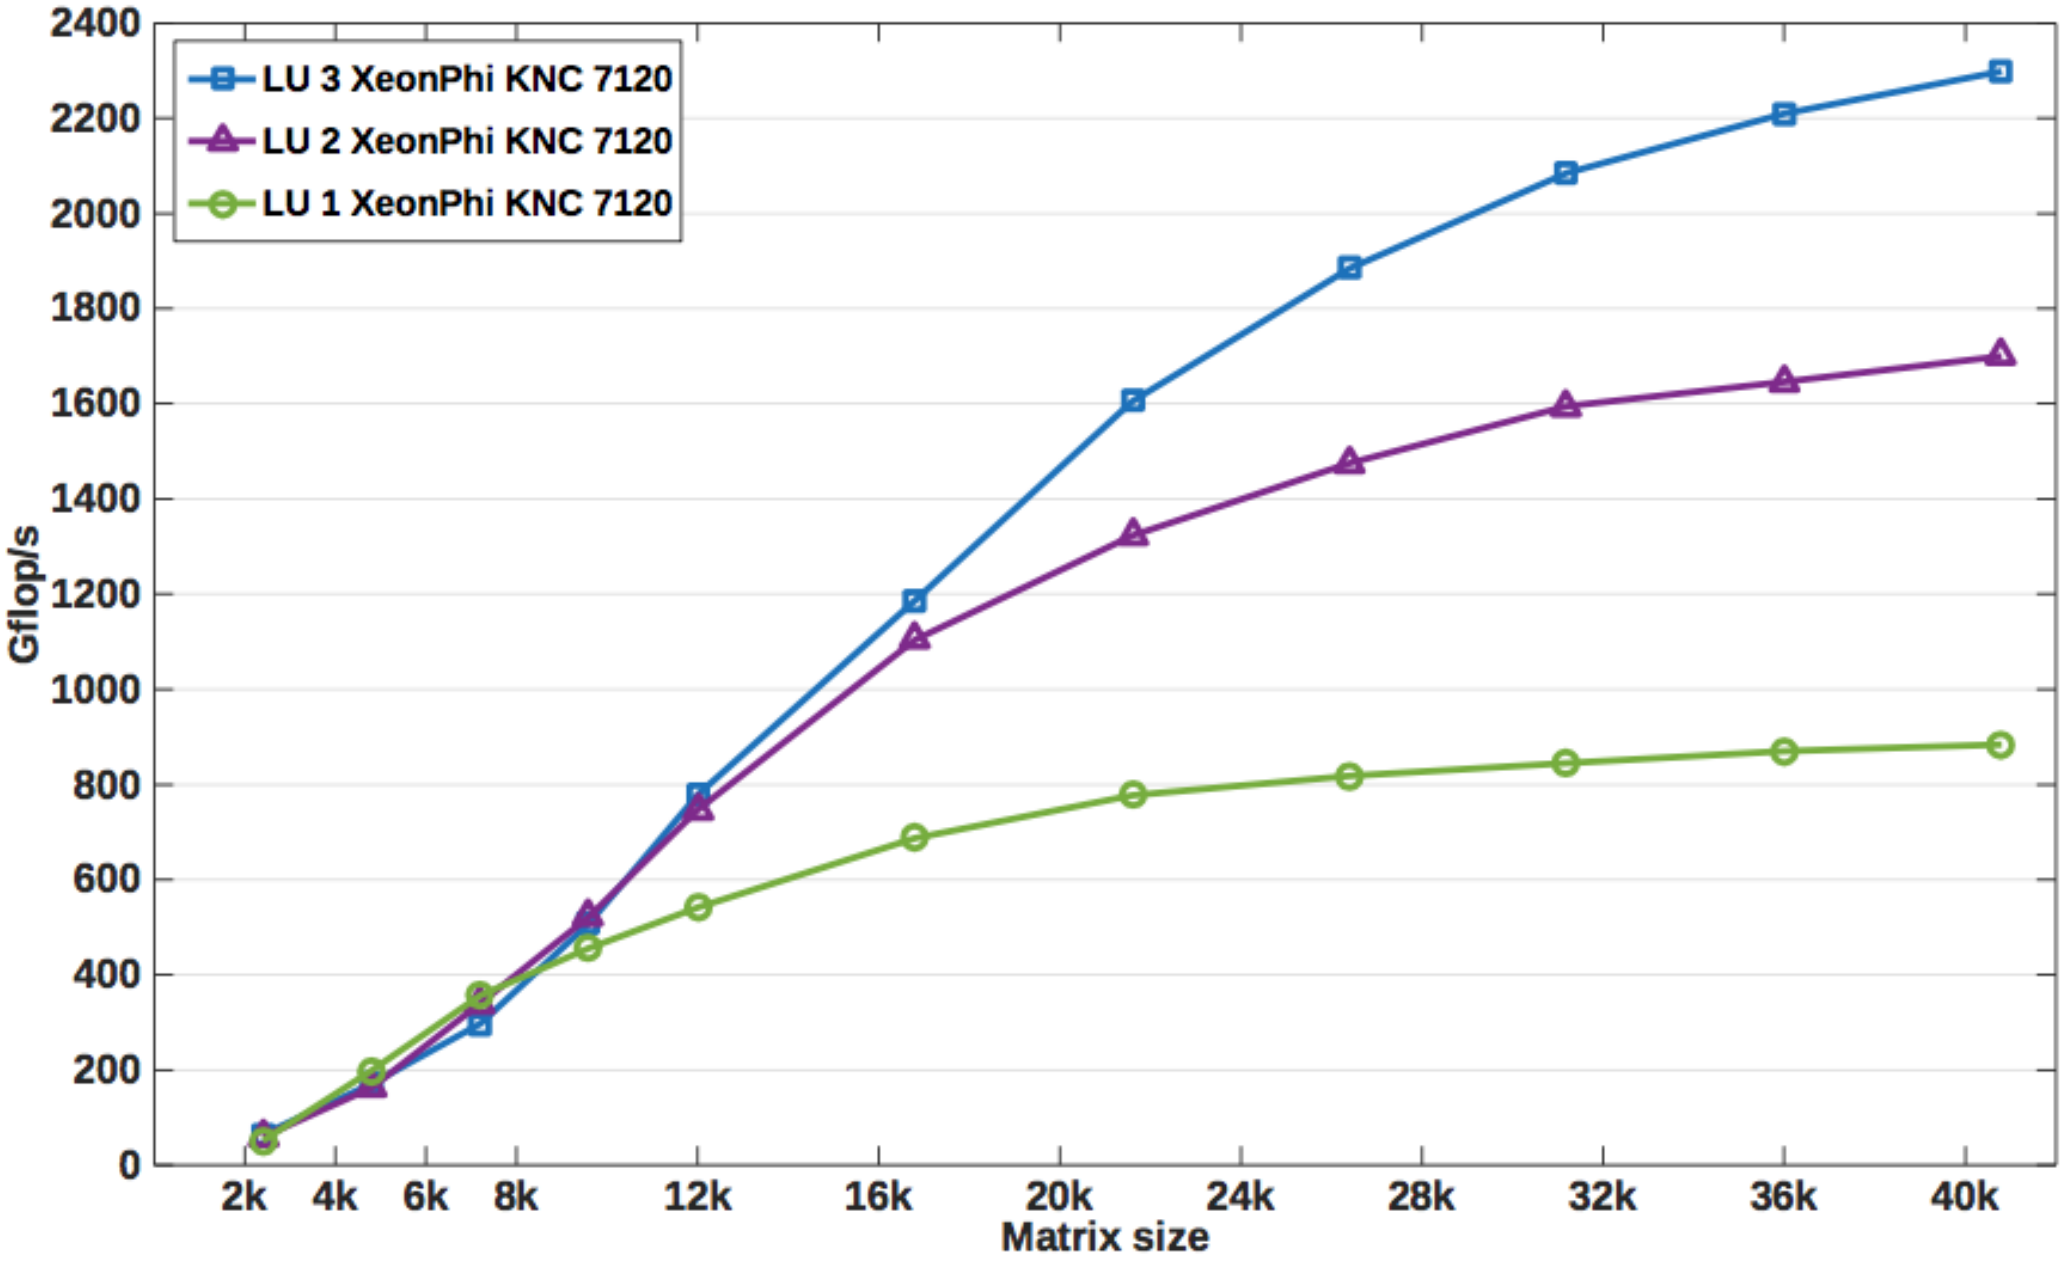

MAGMA Batched Performance

ROUTINES

- LU, QR, and Cholesky ✓
- Solvers and matrix inversion ✓
- All BLAS 3 (fixed + var) ✓
- SYMV, GEMV (fixed + var) ✓

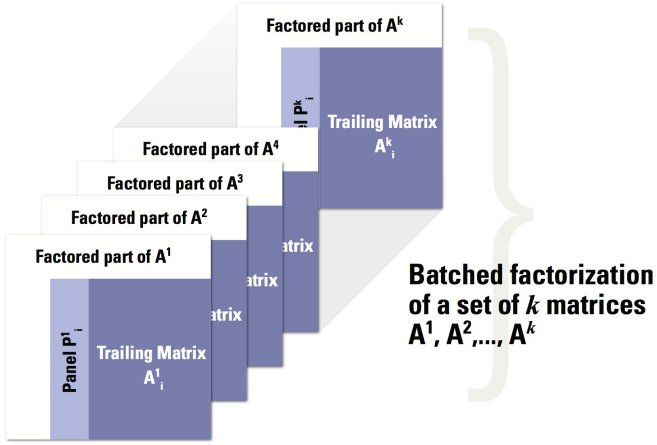

Batched factorization of a set of k matrices $A^1, A^2, \ldots, A^k$

PERFORMANCE OF BATCHED LU

in double precision arithmetic on 1 million matrices

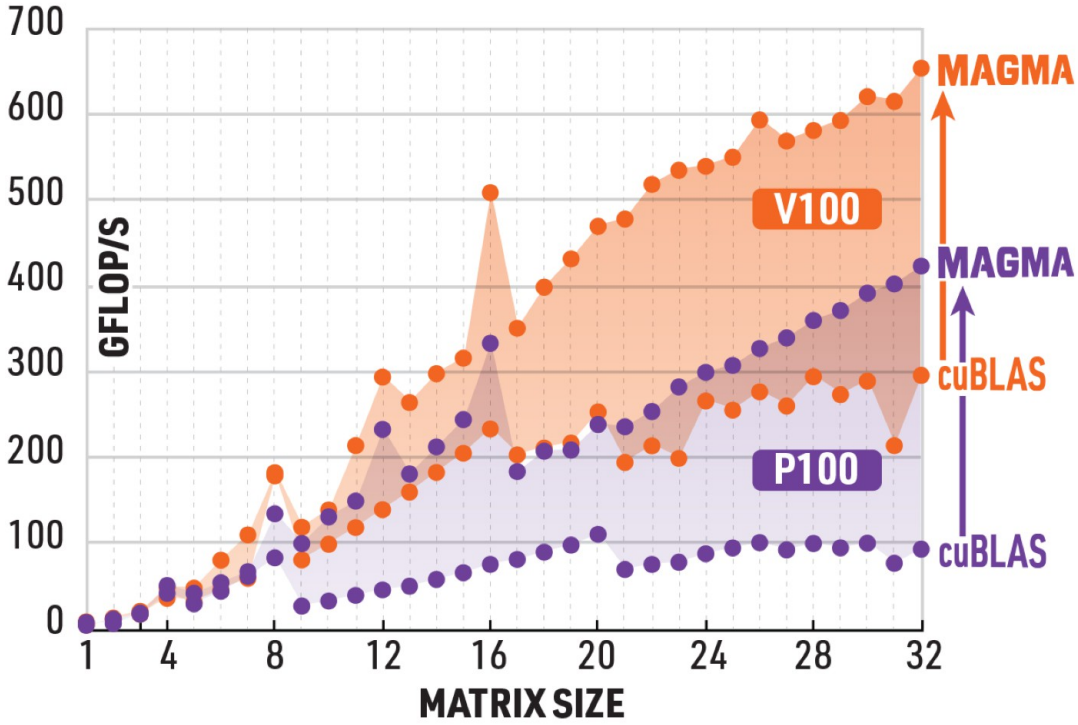

MAGMA Algorithmic Overview

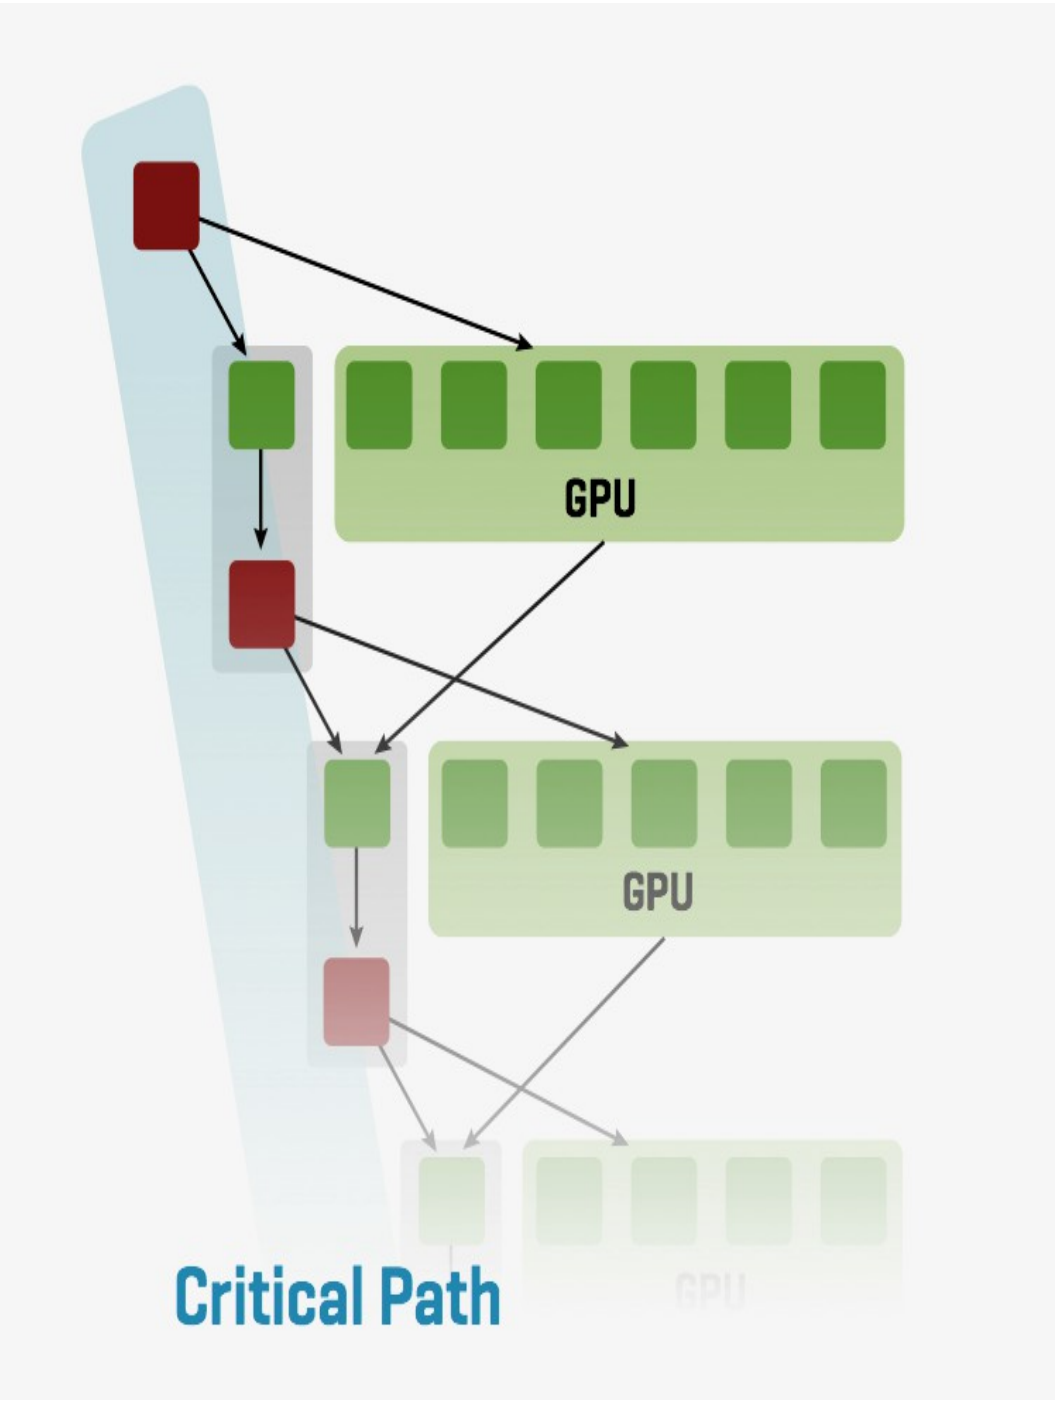

- dense linear algebra for accelerators
  - NVIDIA using CUDA
  - AMD using OpenCL
  - Intel Xeon Phi
- hybrid, CPU-GPU implementations
  - single-GPU
  - multi-GPU
  - OO-GPU-memory
- managing data transfers
- some batched routines
- some sparse solvers

Recent Eigenvalue and SVD Algorithms

Classic (old LAPACK, ScaLAPACK)

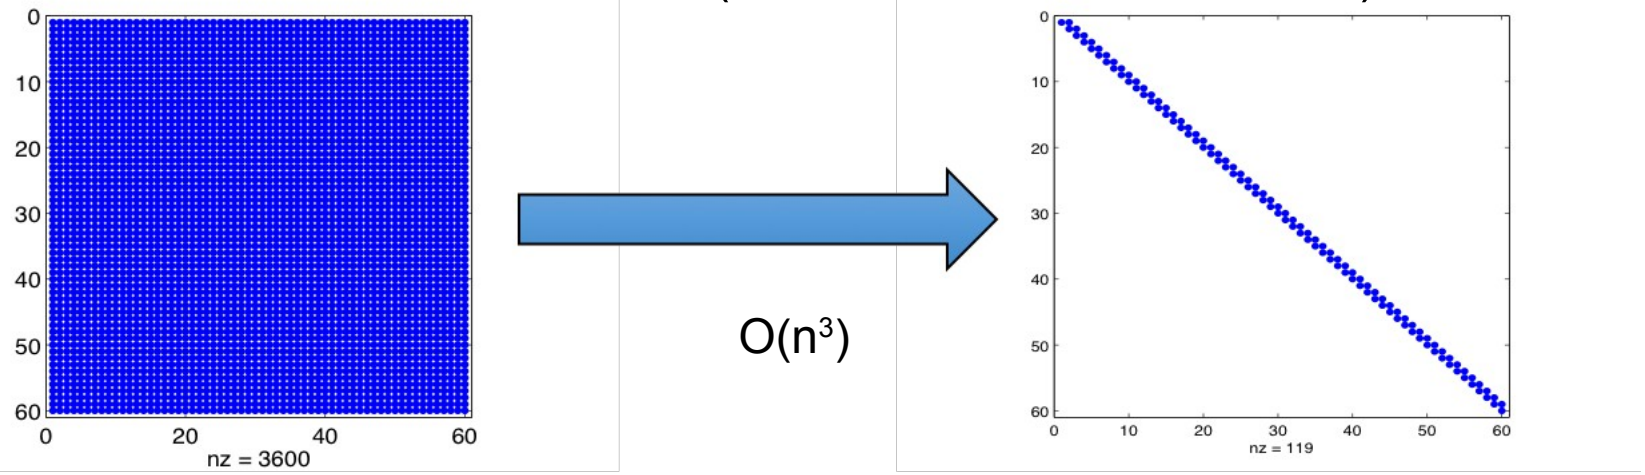

New (multicore libraries, ELPA)

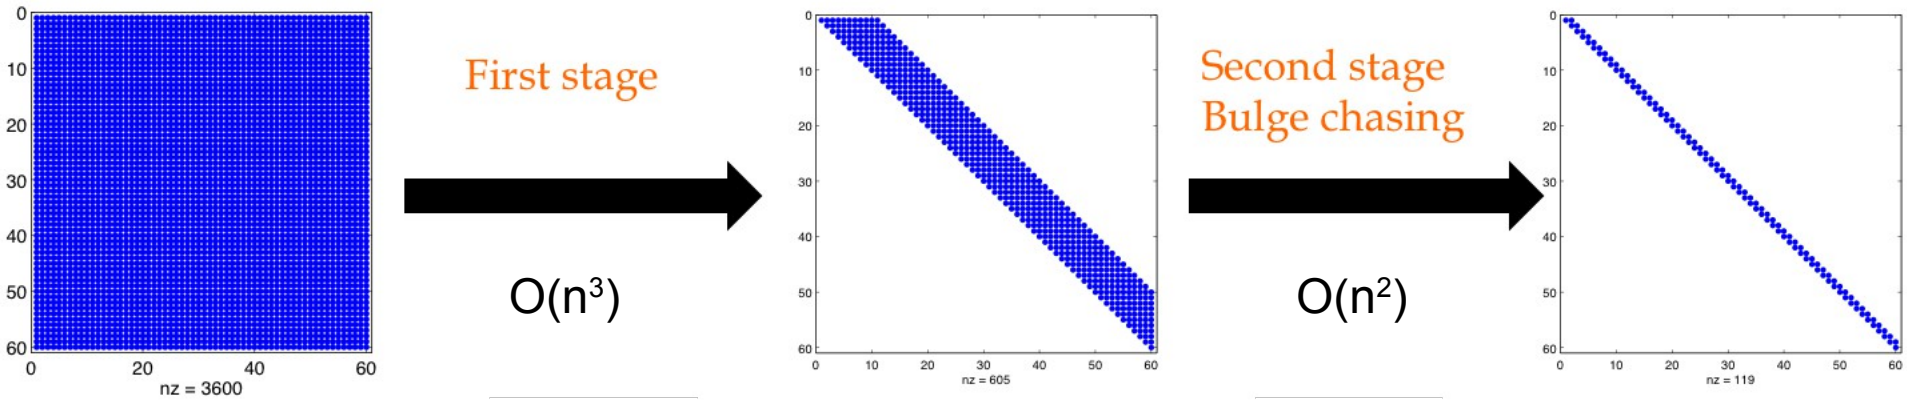

Recent Eigenvalue, SVD Algorithms: 1st Stage

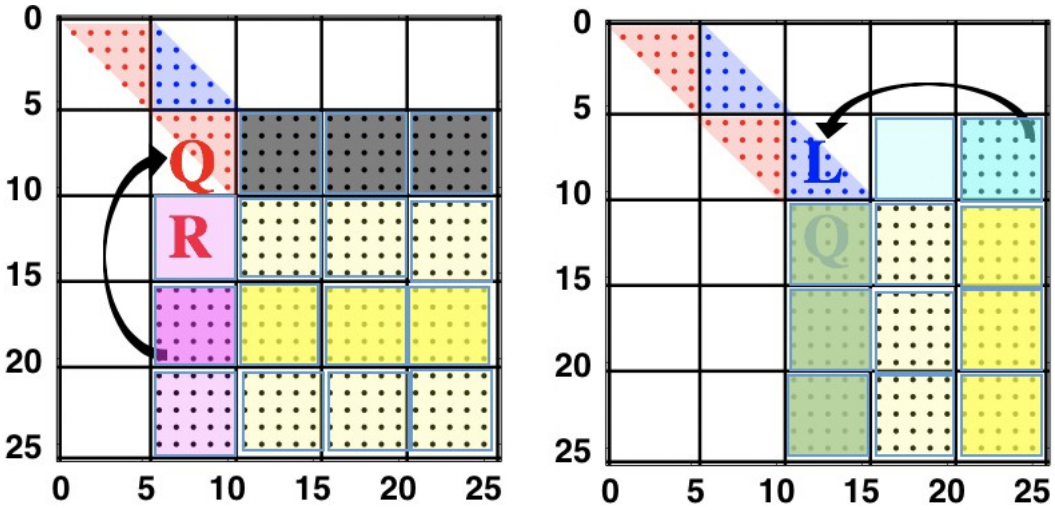

New (multicore libraries)

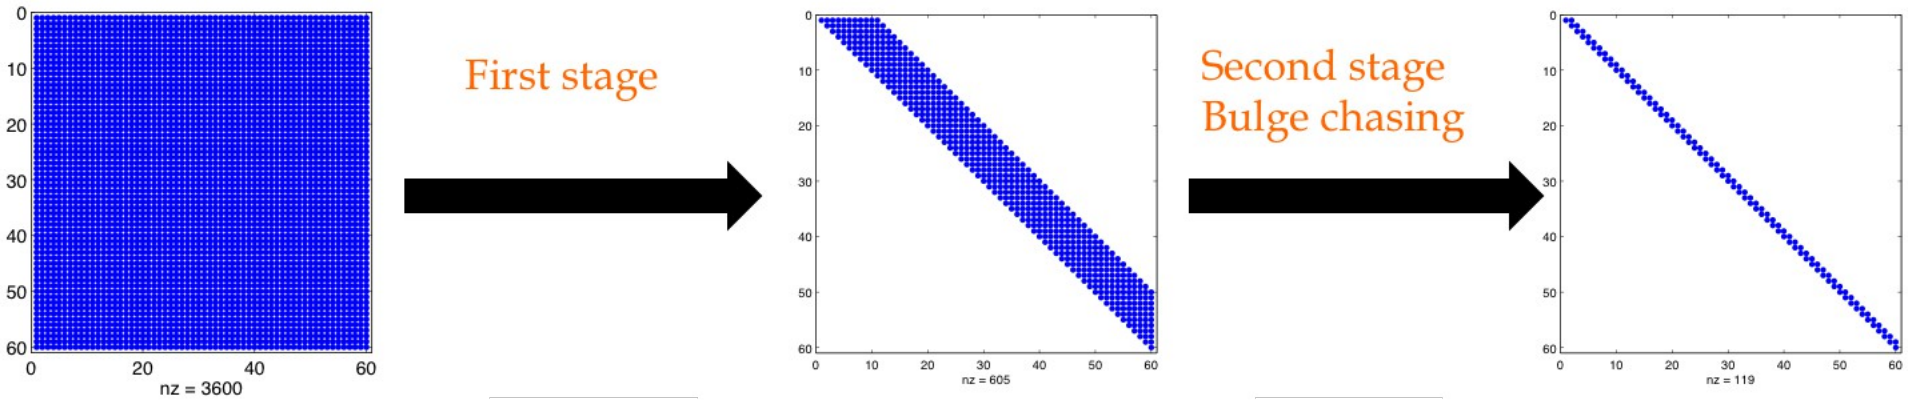

Recent Eigenvalue, SVD Algorithms: 2nd Stage

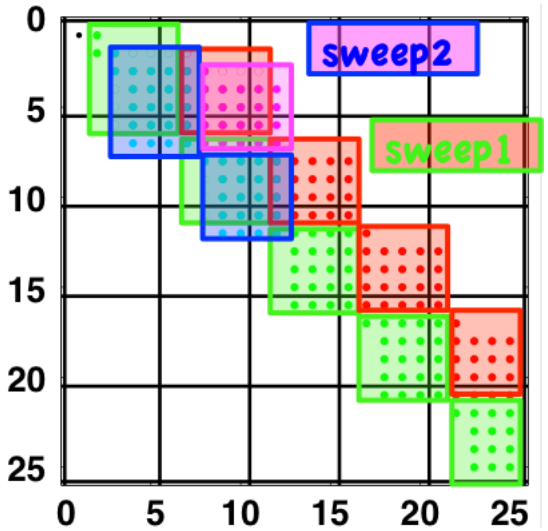

New (multicore libraries)

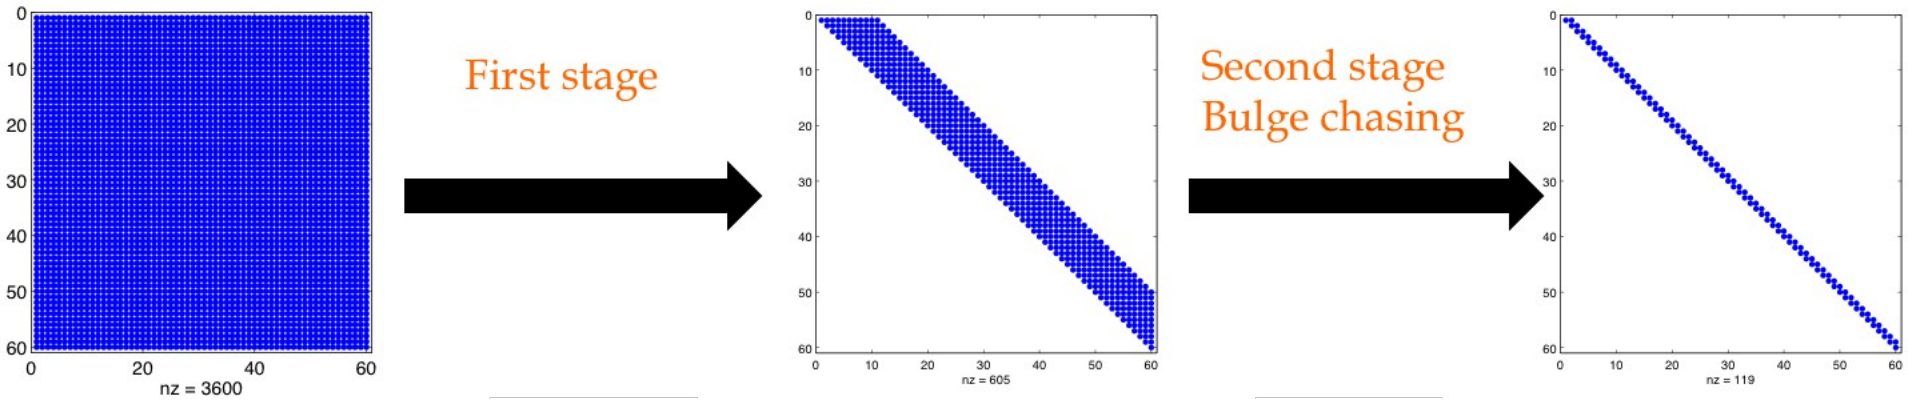

Multicore Performance Results

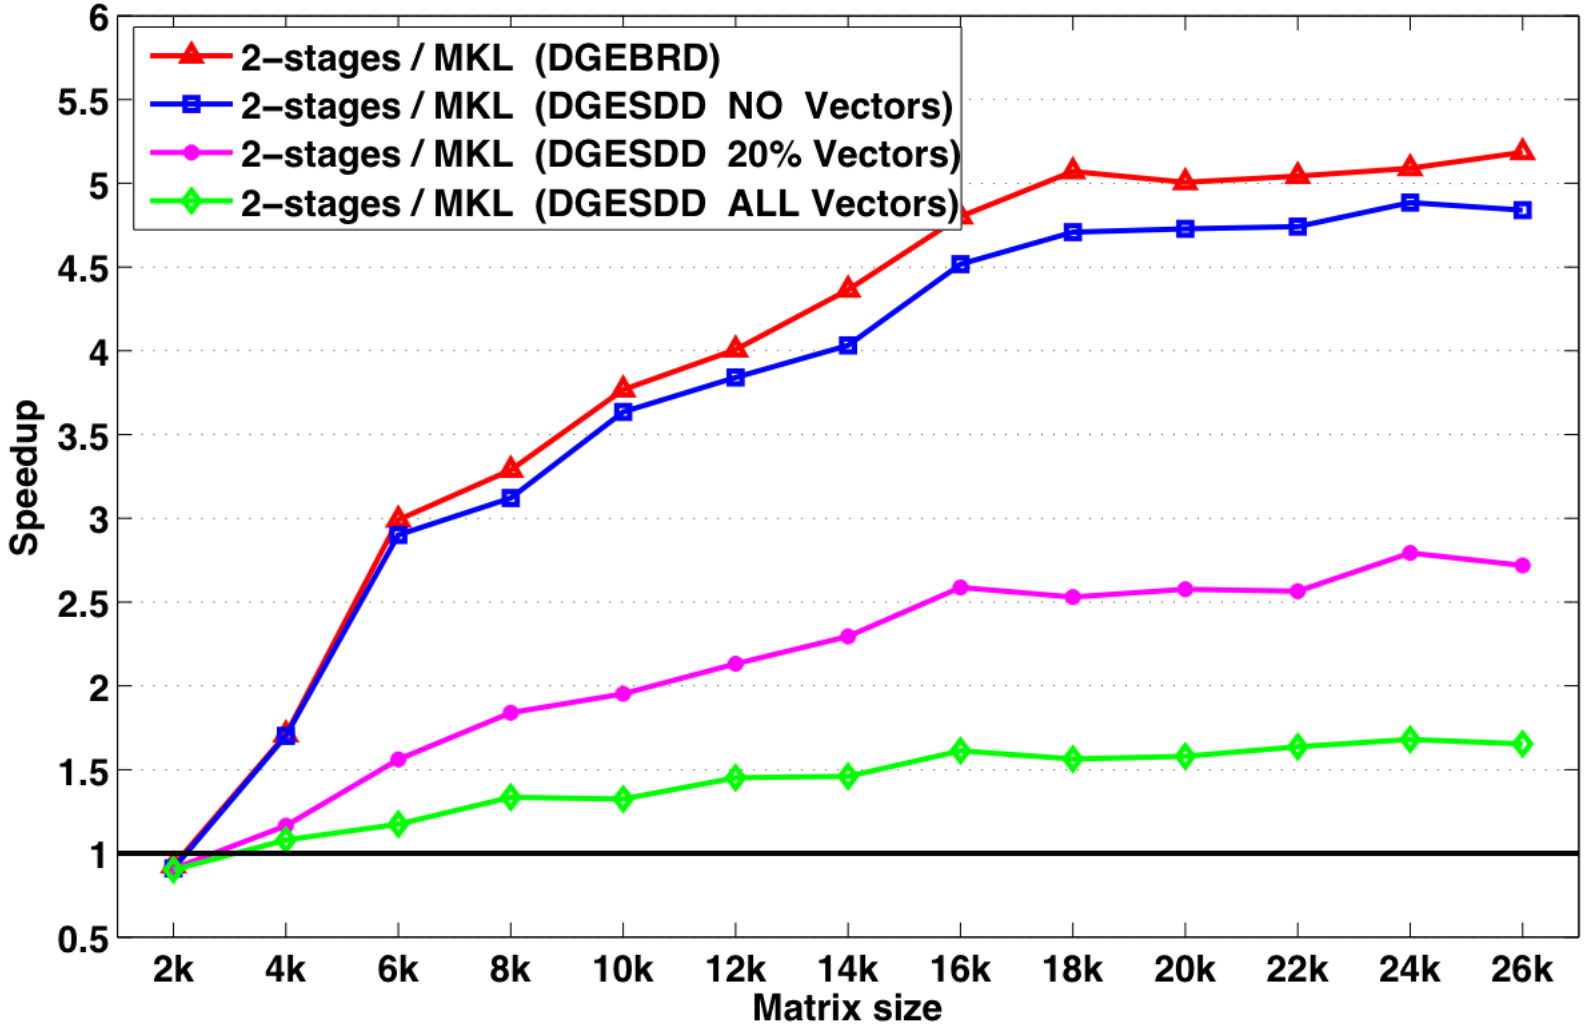

2x8 core Intel Xeon E5-2670 (Sandy Bridge) @ 2.6 GHz

Distributed Memory Performance

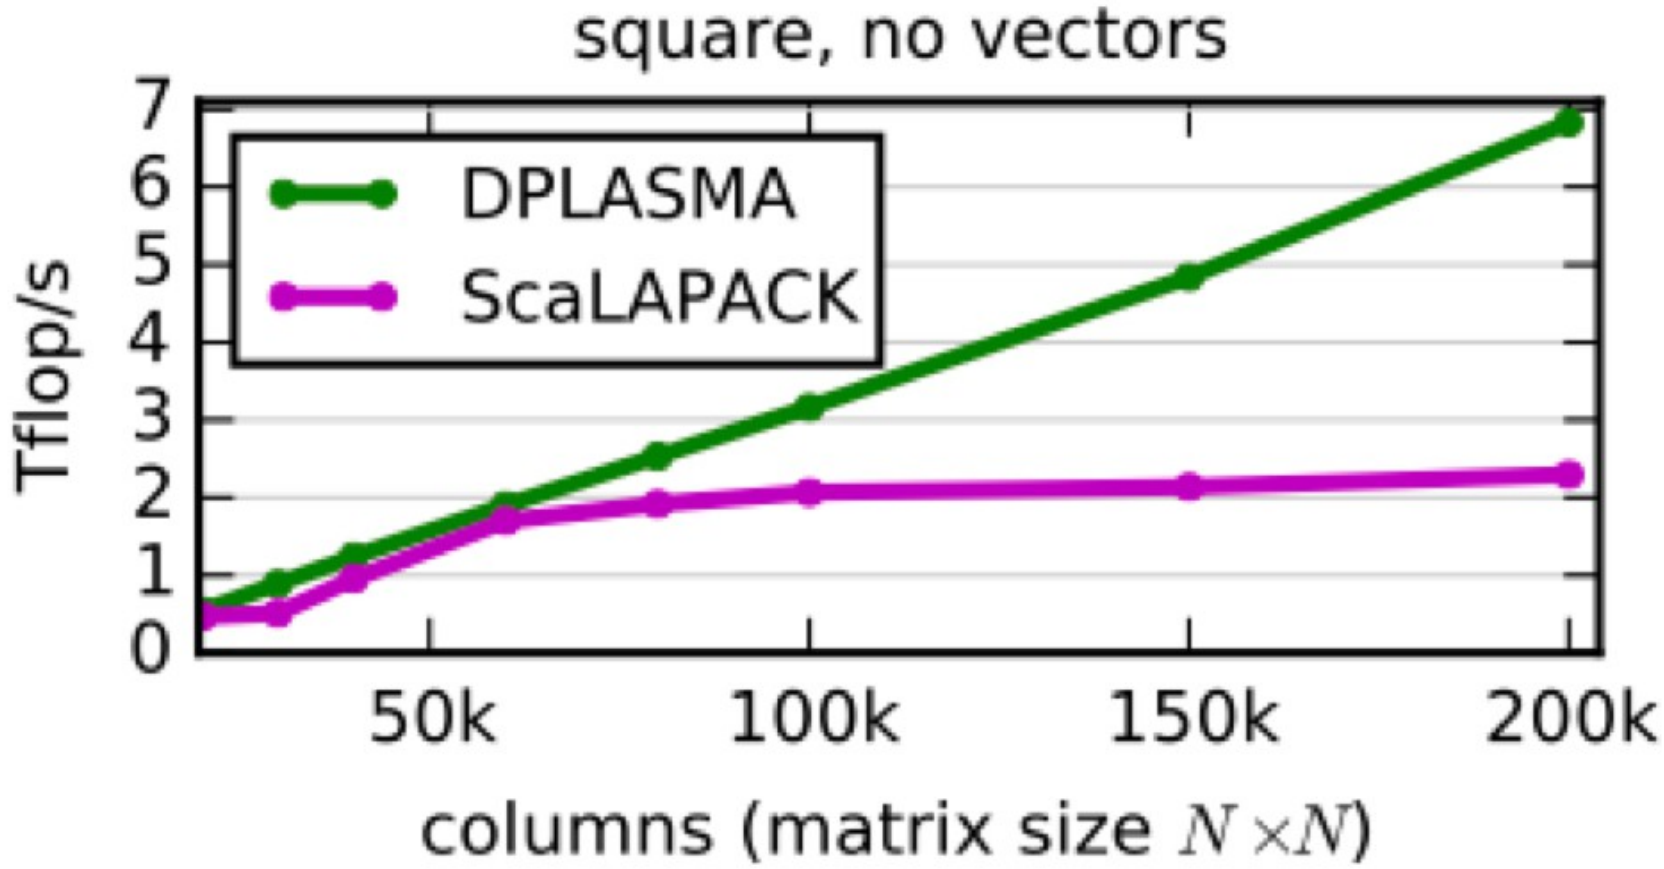

1764 Broadwell cores (49*36)

Low-precision Considerations

1.88 Exa-Ops on Summit

- Algorithm packaged in the comparative genomics application Combinatorial Metrics (CoMet)
  - Custom Correlation Coefficient method, specializes in comparing variations of the same genes (alleles) present in a given population
  - Analyzes datasets composed of millions of genomes (previously impossible)
  - Study variations between all possible combinations of two or three alleles at a time.
  - Uncovers hidden networks of genes in plants and animals that contribute to observable traits
    - Biomarkers for drought-resistance in plants or disease in humans
- Principal investigators
  - Daniel Jacobson (Computational Biologist) and
  - Wayne Joubert (computational scientist)

Floating Point Formats from IEEE 754 (2008)

Precision	Width	Exponent bits	Mantissa bits	Epsilon	Max
Quadruple	128	15	112	$O(10^{-34})$	$1.2x10^{4932}$
Extended	80	15	64	$O(10^{-19})$	
Double	64	11	52	$O(10^{-16})$	$1.8x10^{308}$
Single	32	8	23	$O(10^{-7})$	$3.4x10^{38}$
Half*	16	5	10	$O(10^{-3})$	65504

*Only storage format is specified. IEEE 2018 covers the compute rules.

Best Performers for FP16, FP32, FP64 (Pascal)

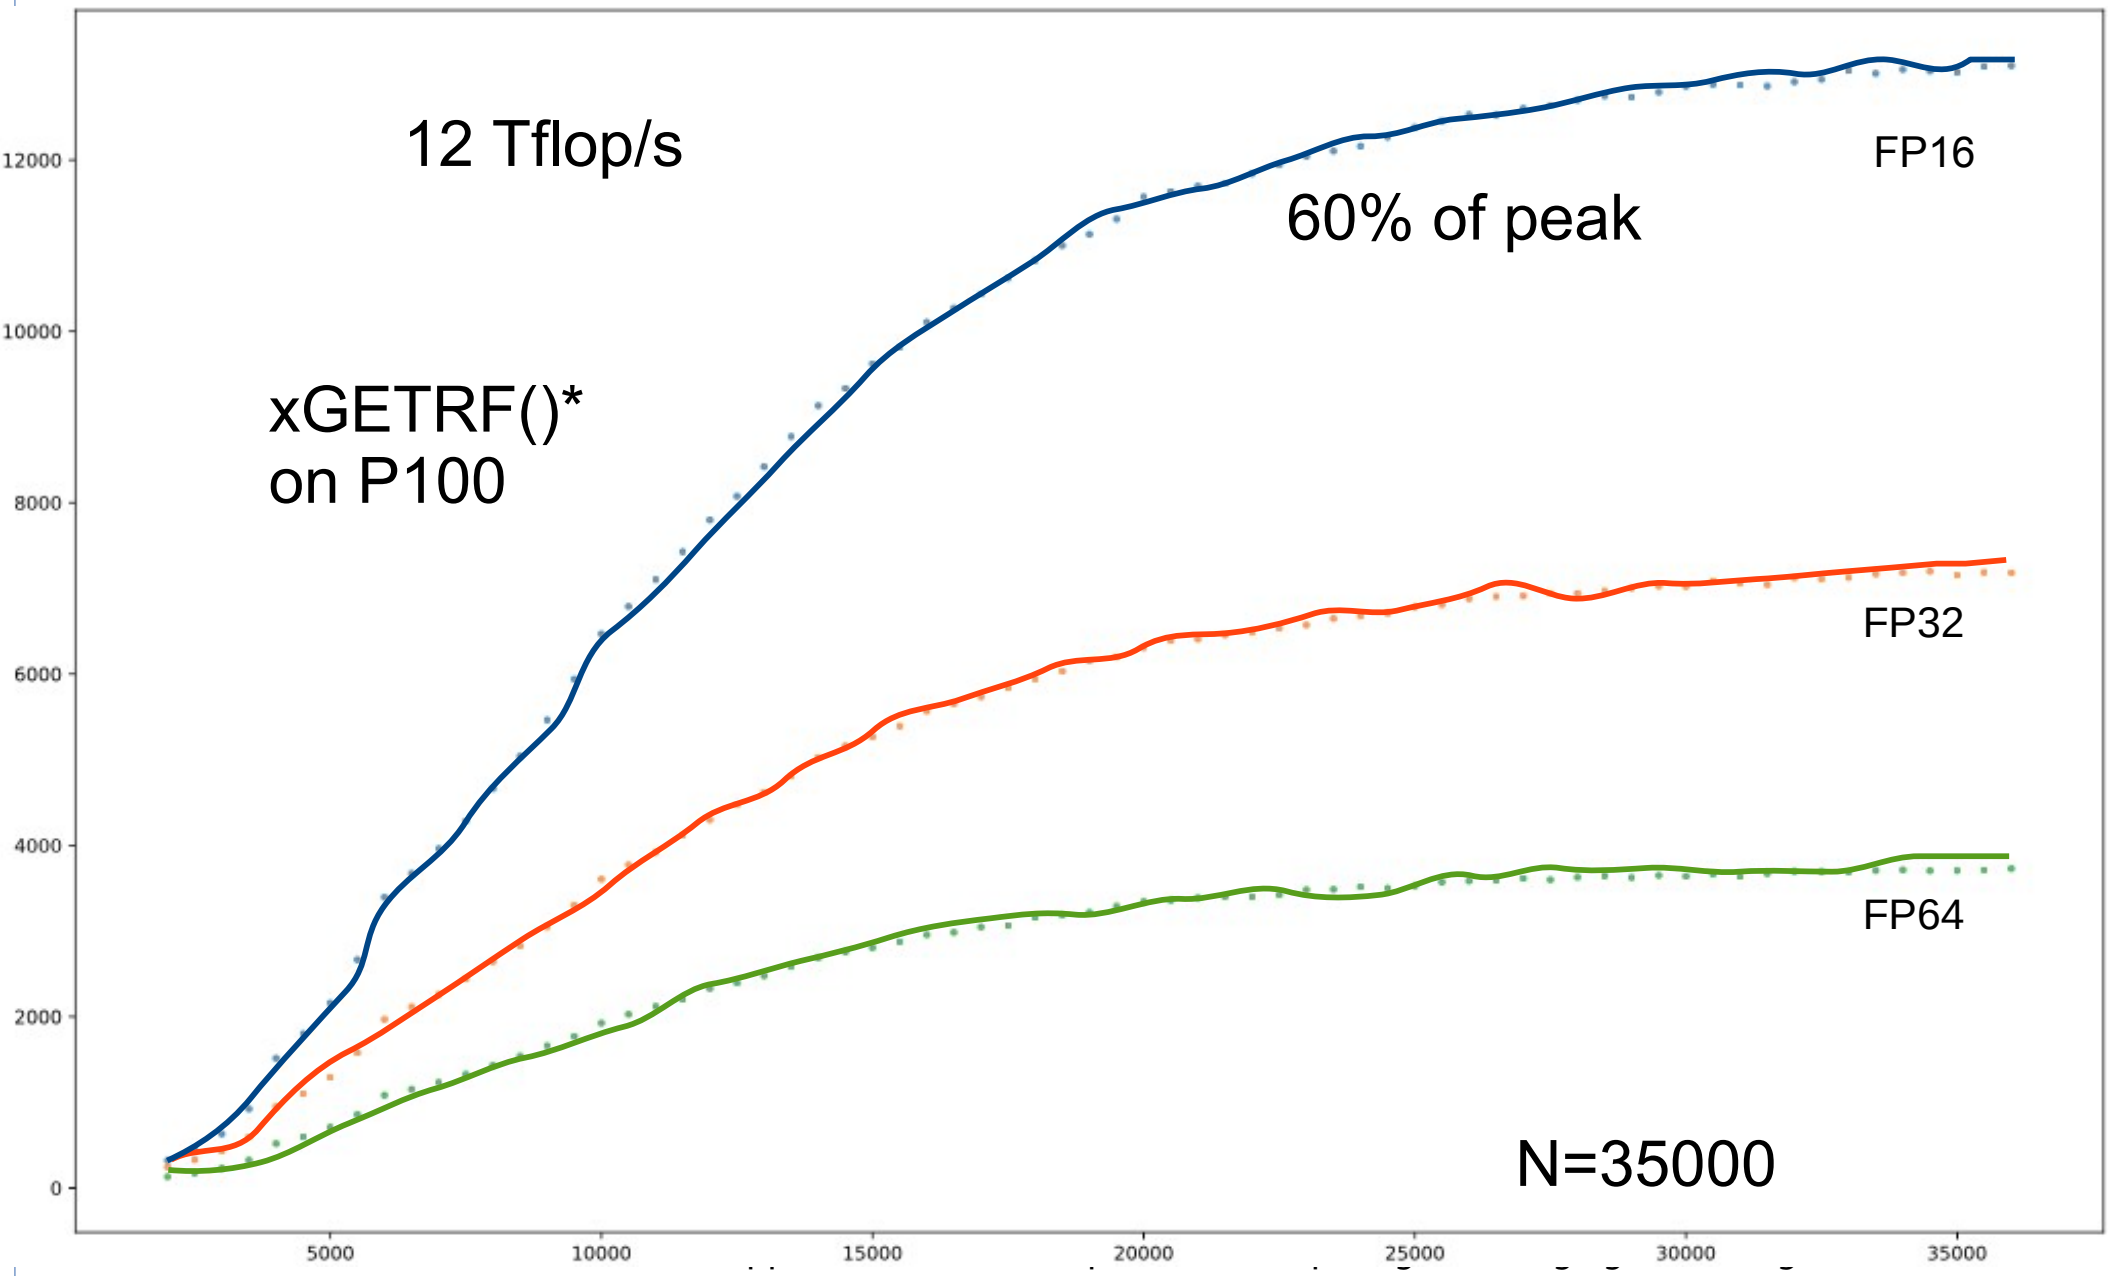

Convergence Results: All Precisions