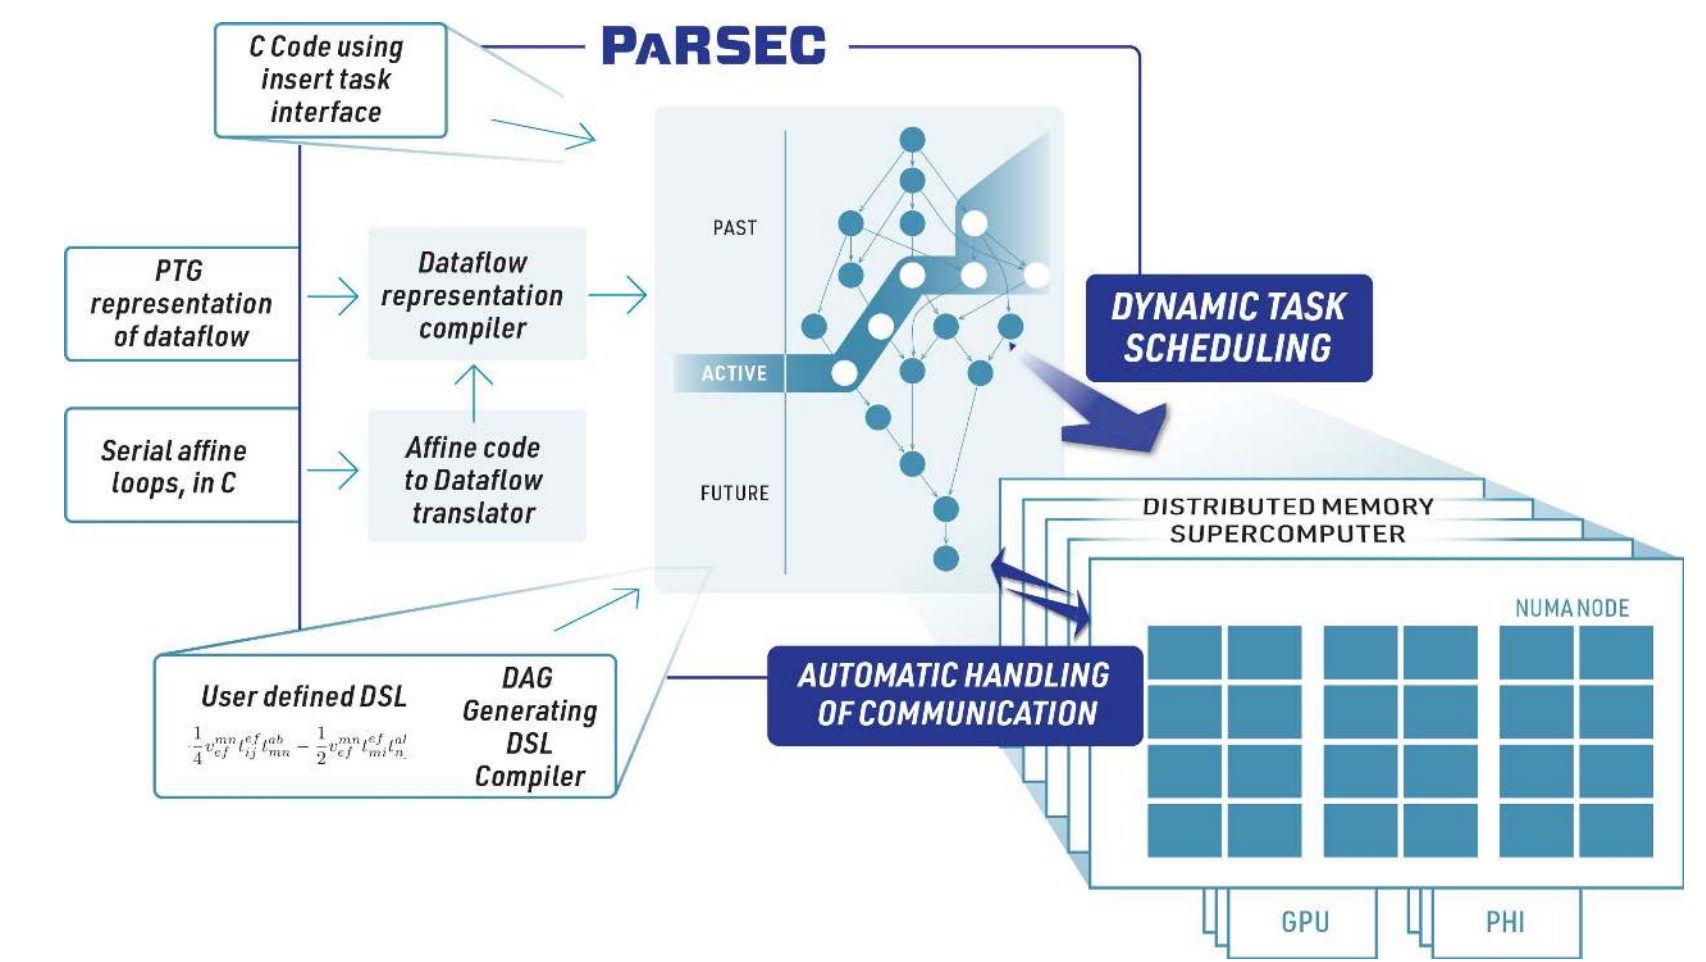


PaRSEC ENABLED LIBRARIES AND APPLICATIONS

The Distributed Tasking for Exascale (DTE) project extends the capabilities of ICL's Parallel Runtime and Execution Controller (PaRSEC) project—a generic framework for architecture-aware scheduling and management of microtasks on distributed, many-core, heterogeneous architectures. The PaRSEC environment also provides a runtime component for dynamically executing tasks on heterogeneous distributed systems along with a productivity toolbox and development framework that supports multiple domain-specific languages and extensions and tools for debugging, trace collection, and analysis.

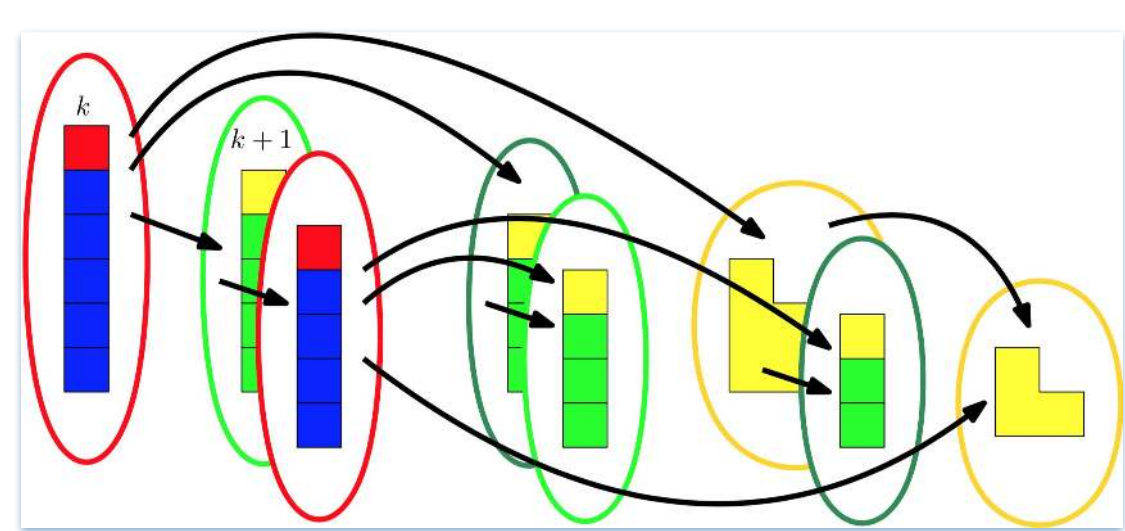


ECP SLATE

High level DAG, Cholesky factorization


Dependencies are expressed between block columns of the matrix

High level tasks insert tile-level tasks, synchronize, or insert (a)synchronous communication tasks



Cholesky Factorization (POTRF)

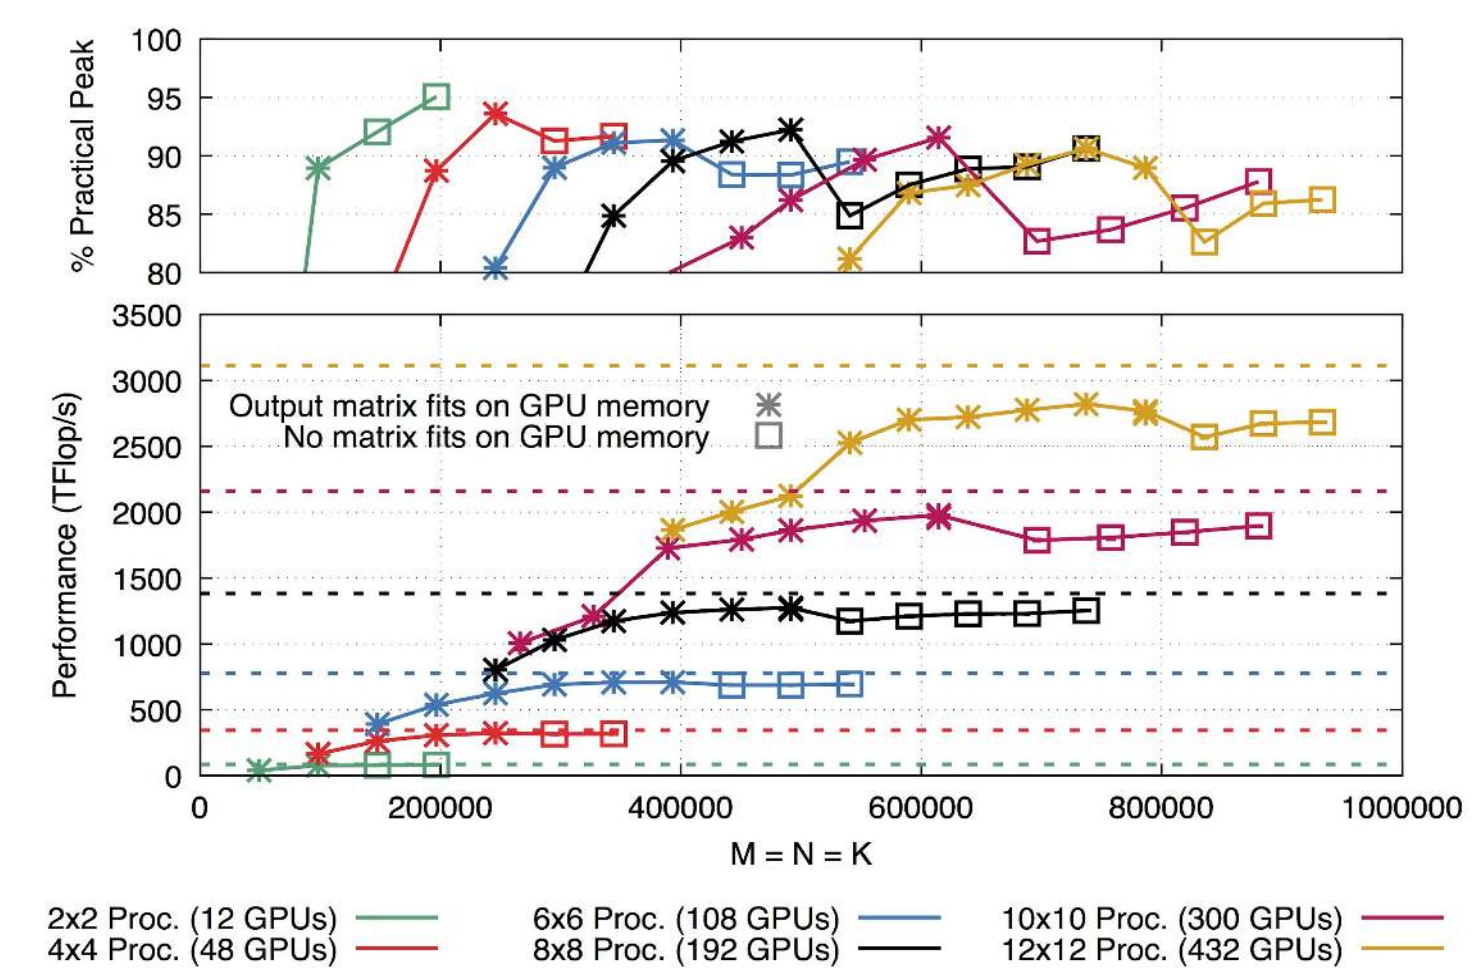
Double Precision • 64 nodes (8x8), 32 cores each
tiles of 1024x1024 doubles



DPLASMA

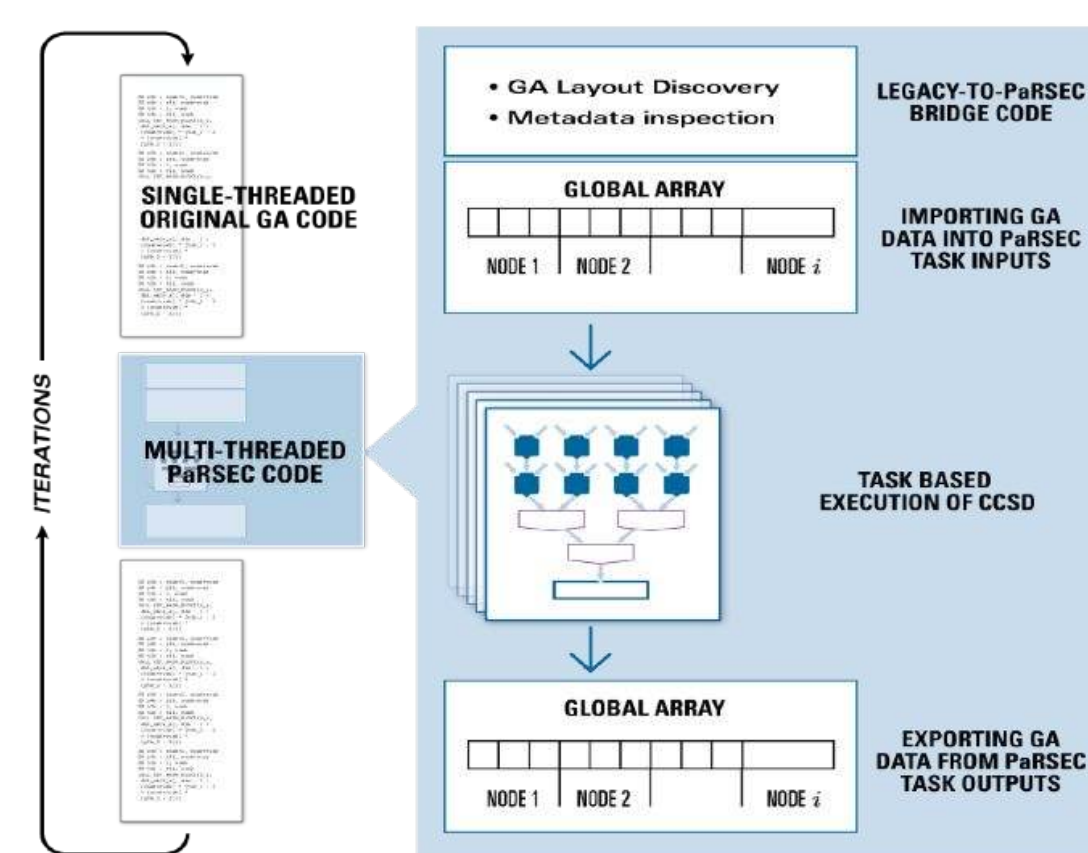

Hybrid Matrix-Matrix Multiply (GEMM)

Double precision (dgemm) / 2-72 Nodes of Summit (40 cores + 6 V100s)
Tiled Algorithm, with tiles of 1024x1024 doubles

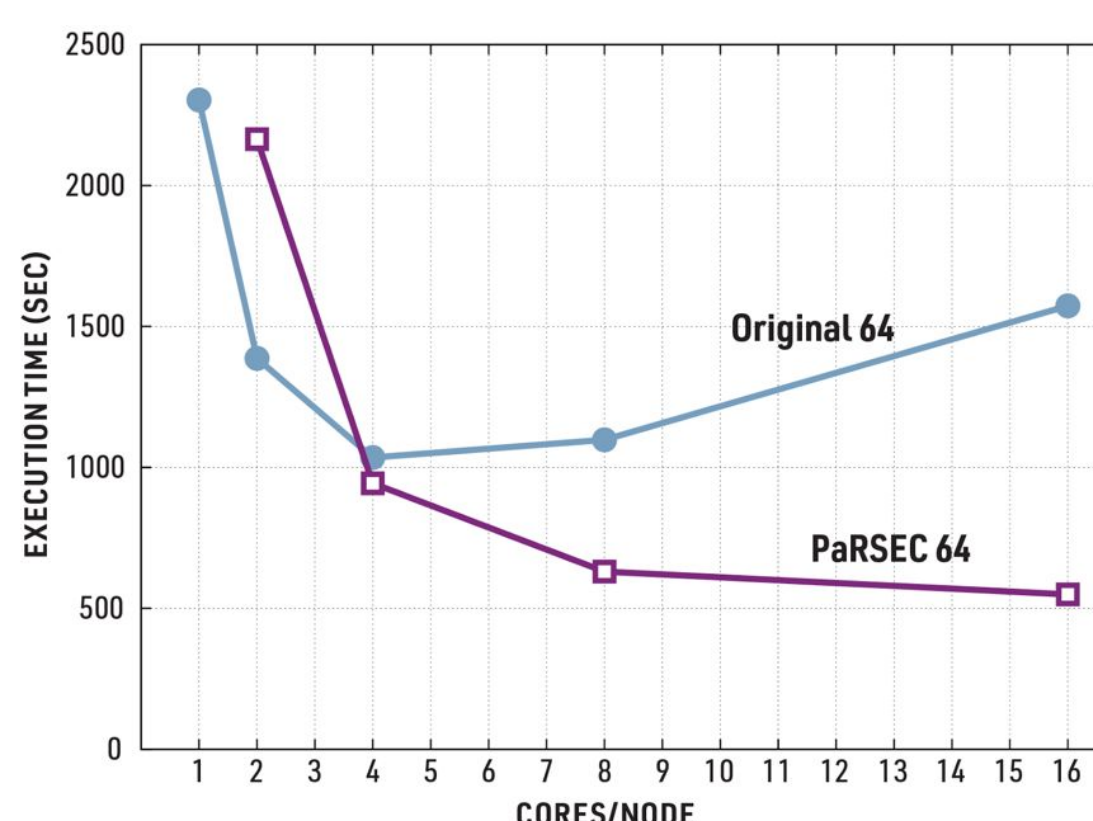


NWChem INTEGRATION

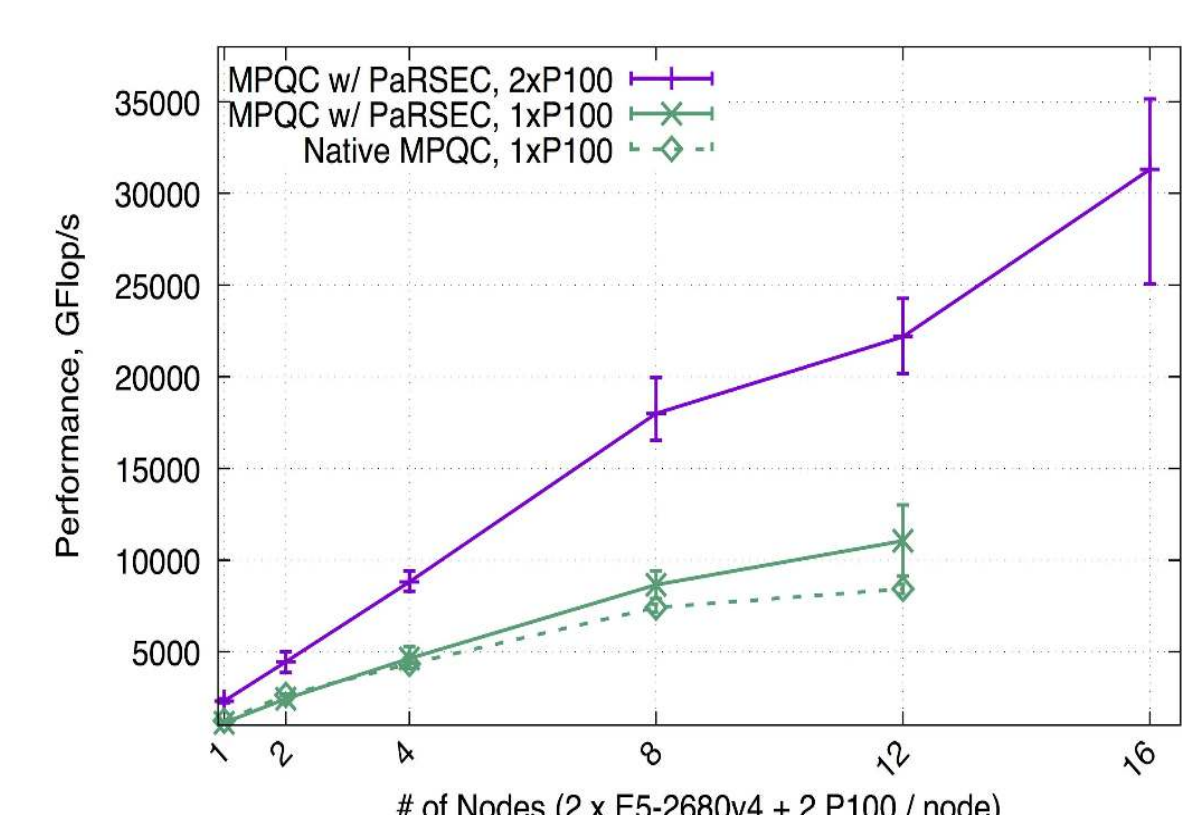
PaRSEC Kernel inserted into existing NWChem codebase improves manycore scalability



CCSD module on 64 nodes



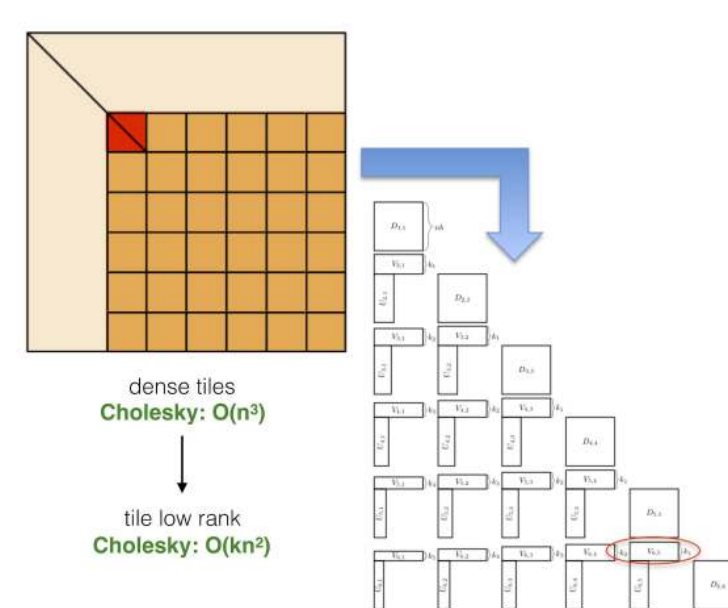
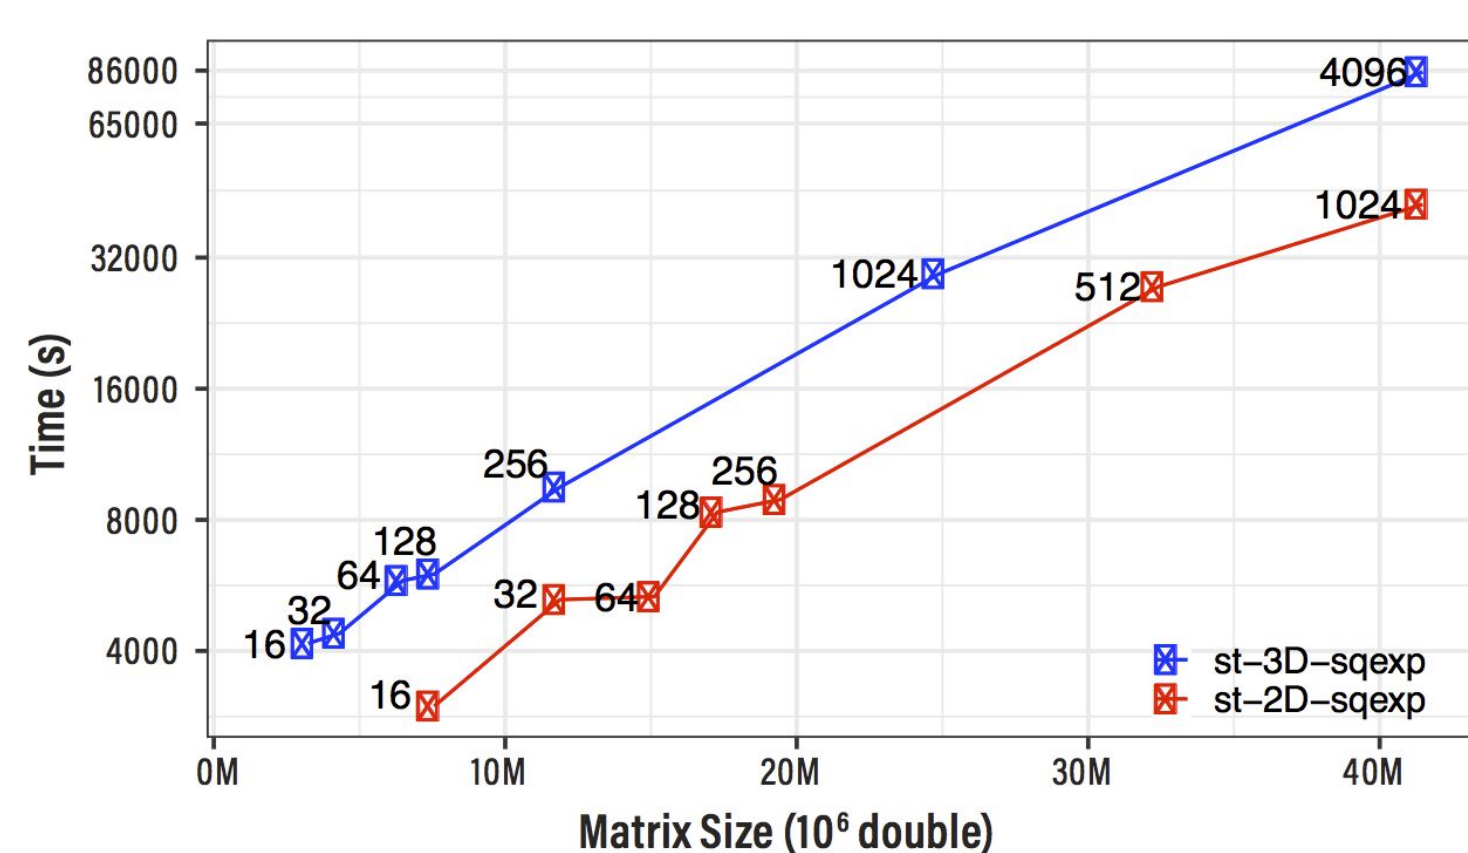
Massively Parallel Quantum Chemistry (MPQC) on Hybrid System



HICMA Hierarchical Computations on Manycore Architectures

Tile, Low-Rank, Cholesky Factorization for Large Matrices

Shaheen II: 4096 nodes (32 cores each @ 2.30 GHz (Intel Haswell))



Tiles of the matrix are communicated under their low-rank representation (at most $2n$ vs. n^2)

Kernels (operations on tiles) either decompress the tile locally then re-compress, or operate directly on the low-rank representation

LEFT: problem size too large to obtain result with ScaLAPACK.

Numbers on points represent the number of Shaheen II nodes used to compute the factorization

RIGHT: time of HICMA-TLR includes time to compute the low rank representation, execute the operation, decompress the final representation and refine result to the ScaLAPACK precision using 16 nodes, and is compared to runs using ScaLAPACK from 16 to 256 nodes.

Comparison with ScaLAPACK

Shaheen II: 4096 nodes (32 cores each @ 2.30 GHz (Intel Haswell))

