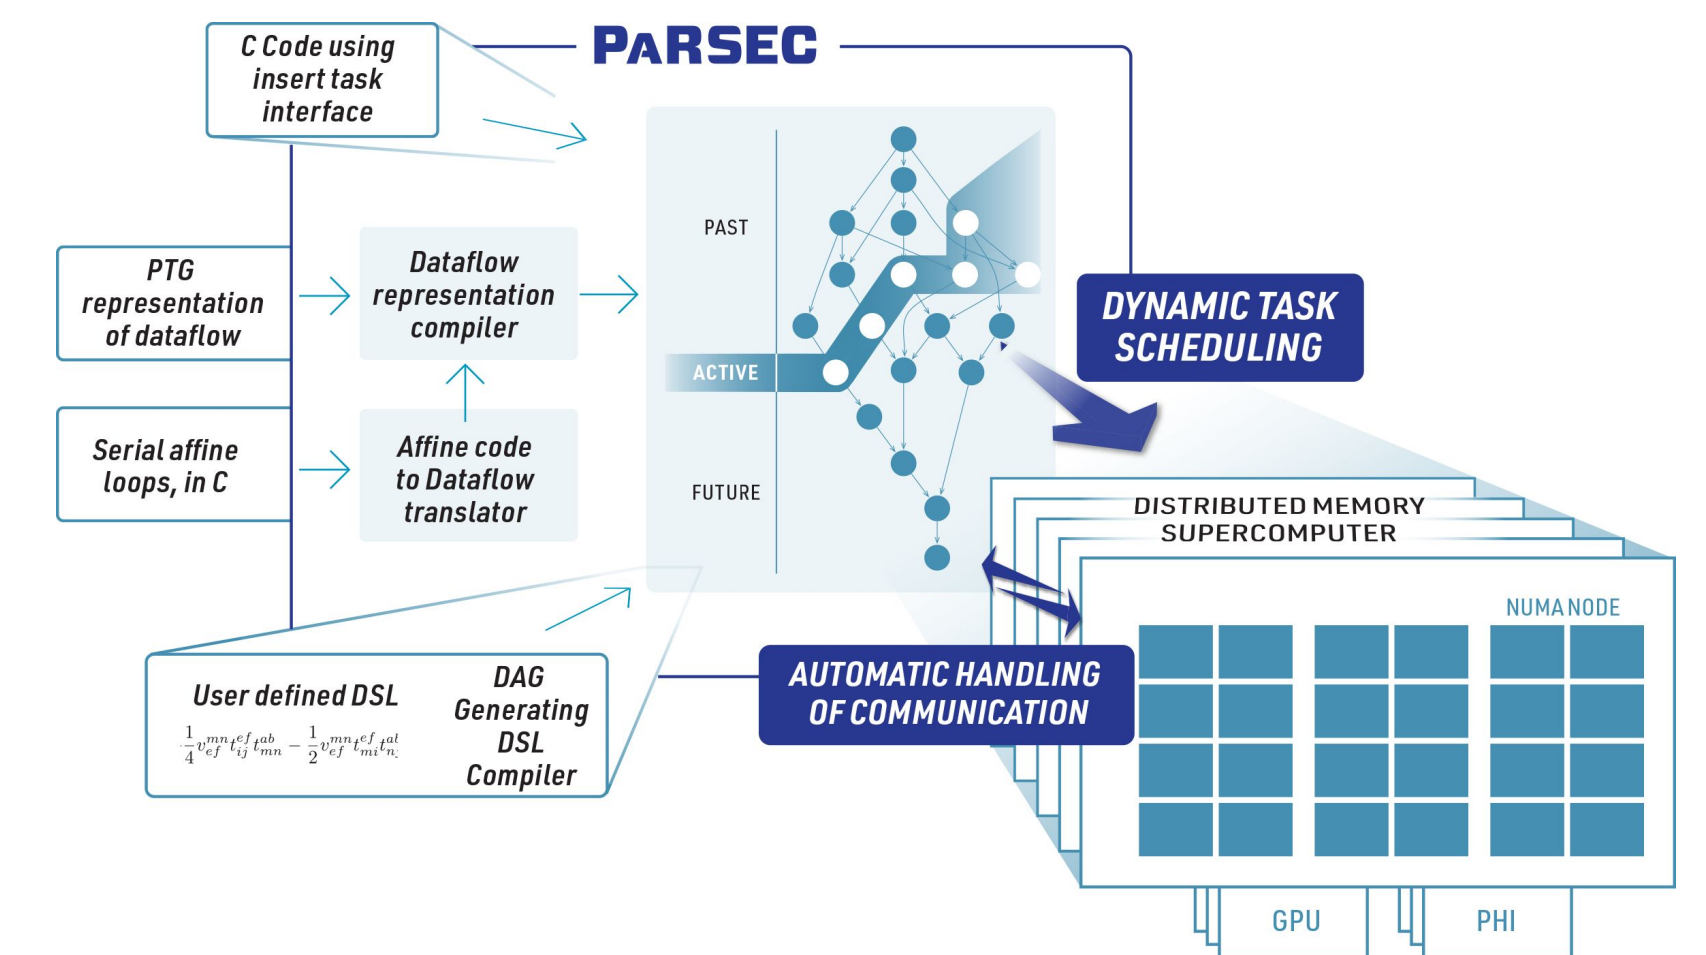


PaRSEC ENABLED LIBRARIES AND APPLICATIONS

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

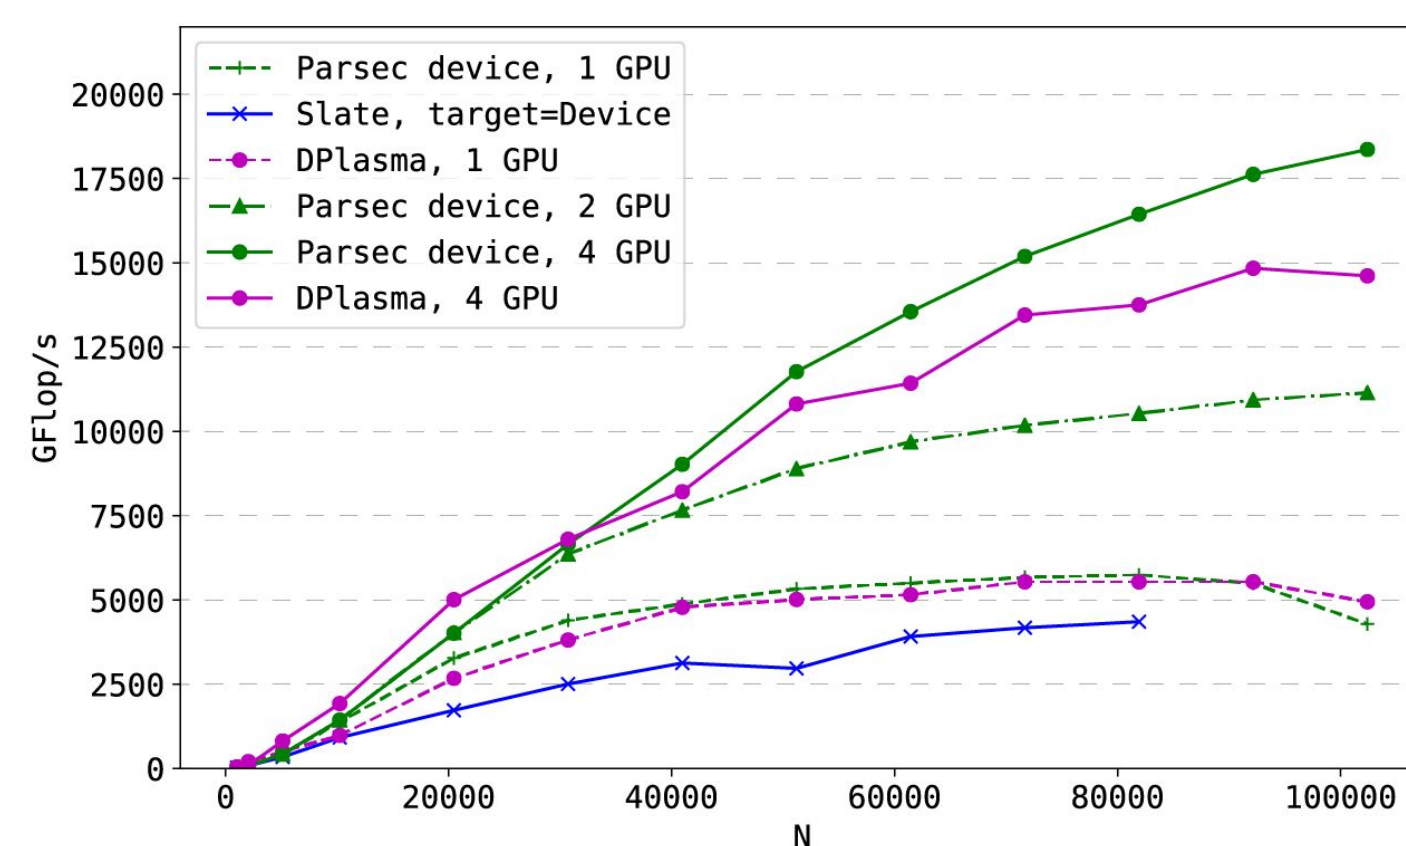
Dependencies are expressed between block columns of the matrix

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



Cholesky Factorization (POTRF)

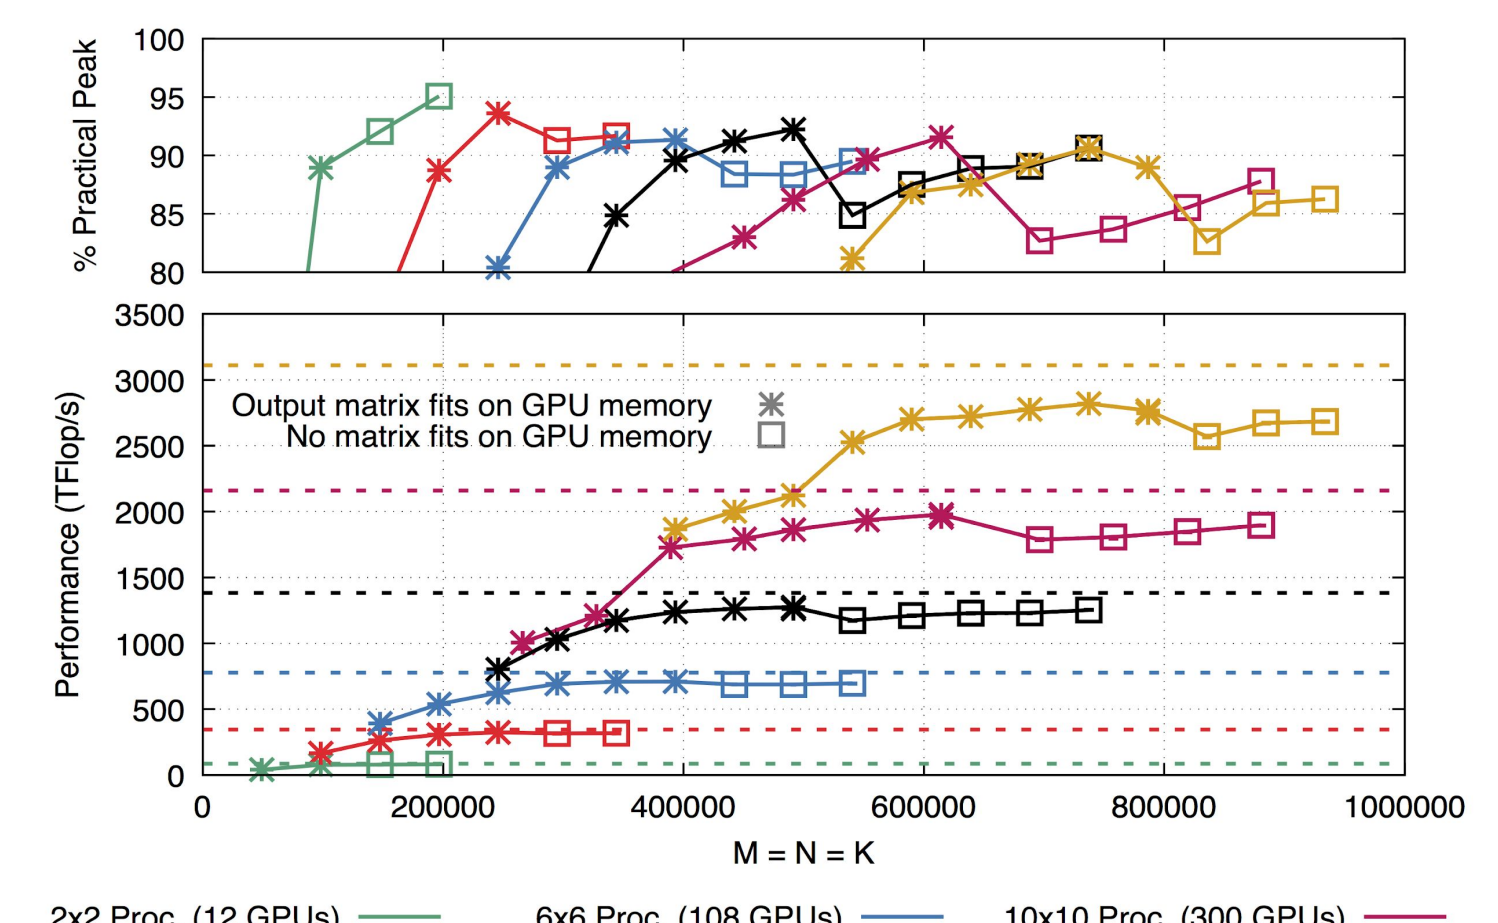
Double Precision • 64 cores, 1 to 4 V100 tiles of 1024x1024 doubles



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

Massively Parallel Quantum Chemistry (MPQC)

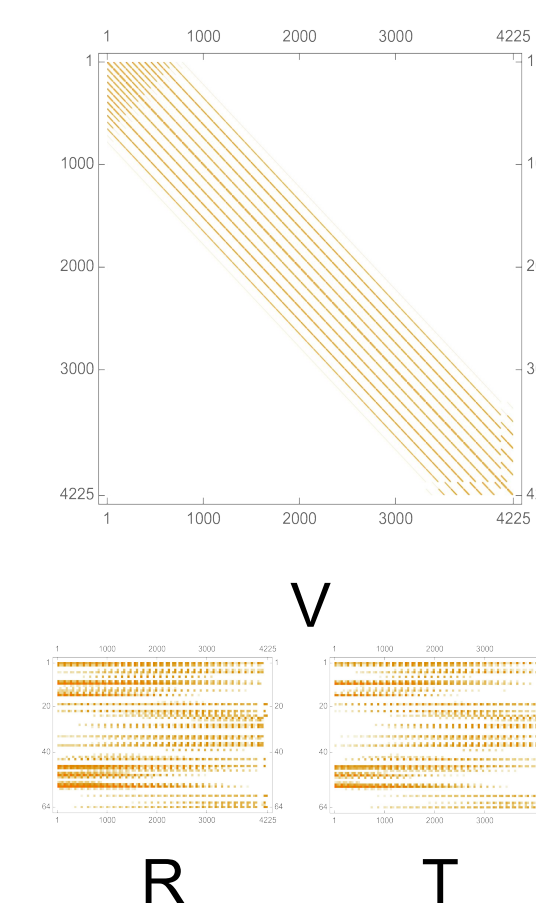
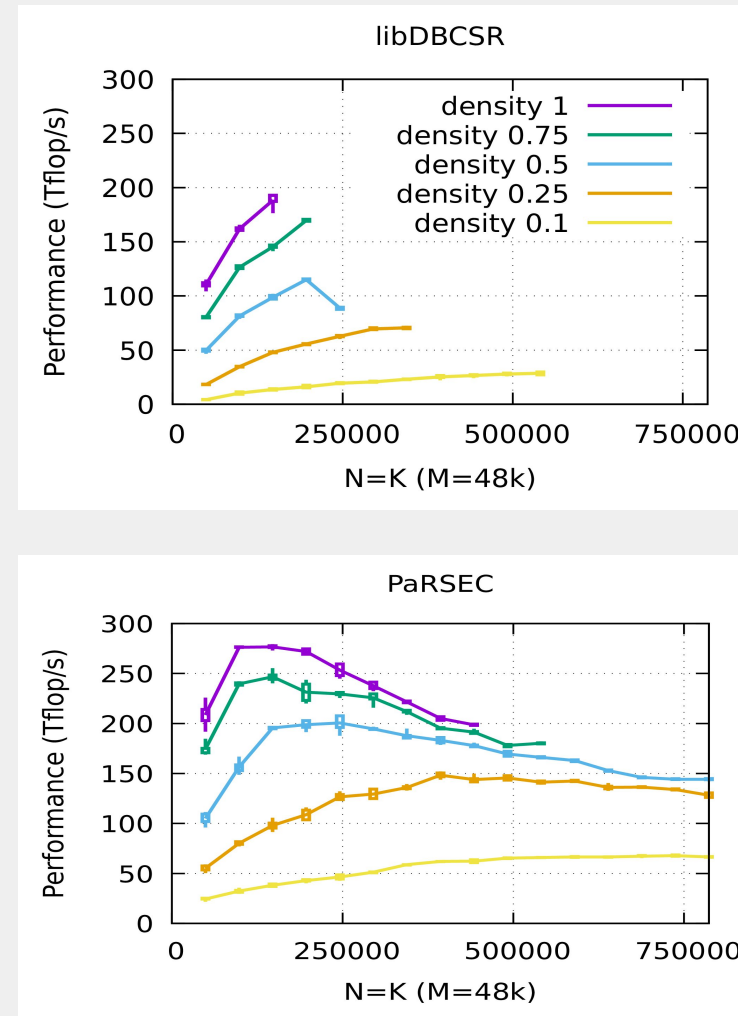
- Application part of NWChemEx
- Base implementation on top of TiledArray, itself programmed on top of MADNESS
- Replace tensor 'ABCD' tensor contraction in TA with a PaRSEC native implementation
- Use MADNESS over PaRSEC to simplify sharing of resources between MPQC/TA and PaRSEC native code

Problem characteristics:


- Matrices are not all square: V is square and too large to fit in (accumulated memory), and T and R are short and wide
- Matrices are not dense, but block-sparse (density of 1% to 20%)
- Tiling is irregular (large variability)

$$R_{ab}^{ij} = \sum_{cd} T_{cd}^{ij} V_{ab}^{cd}$$

Synthetic Benchmark (random matrices)



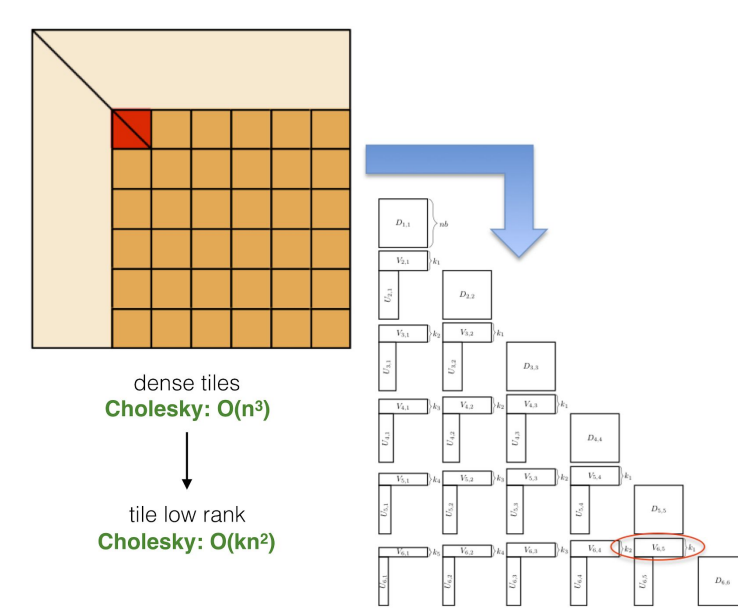

Applicative case (C65H132) on Hybrid System



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))



LEFT: problem size too large to possibly obtain a result with ScaLAPACK. Instead use a low-rank representation using HiCMA.

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

Tiles of the matrix are communicated under their existing representation (low-rank at most $2n \times n^2$ on the left and mixed precision $n^2 * \text{sizeof}(\text{representation})$ on the right).

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

RIGHT: time of Cholesky mixed precision factorization using multiple floating point representations using 128 hybrid nodes (6 GPU per node) on Summit. The percentage of each floating part data is indicated, as an example 10D:30S:60H represent a matrix with 10% of the band diagonal being double precision, 30% single precision and the remaining half-precision.

Mixed precision half/single/double

Summit: 128 nodes (6 GPU per node)

