

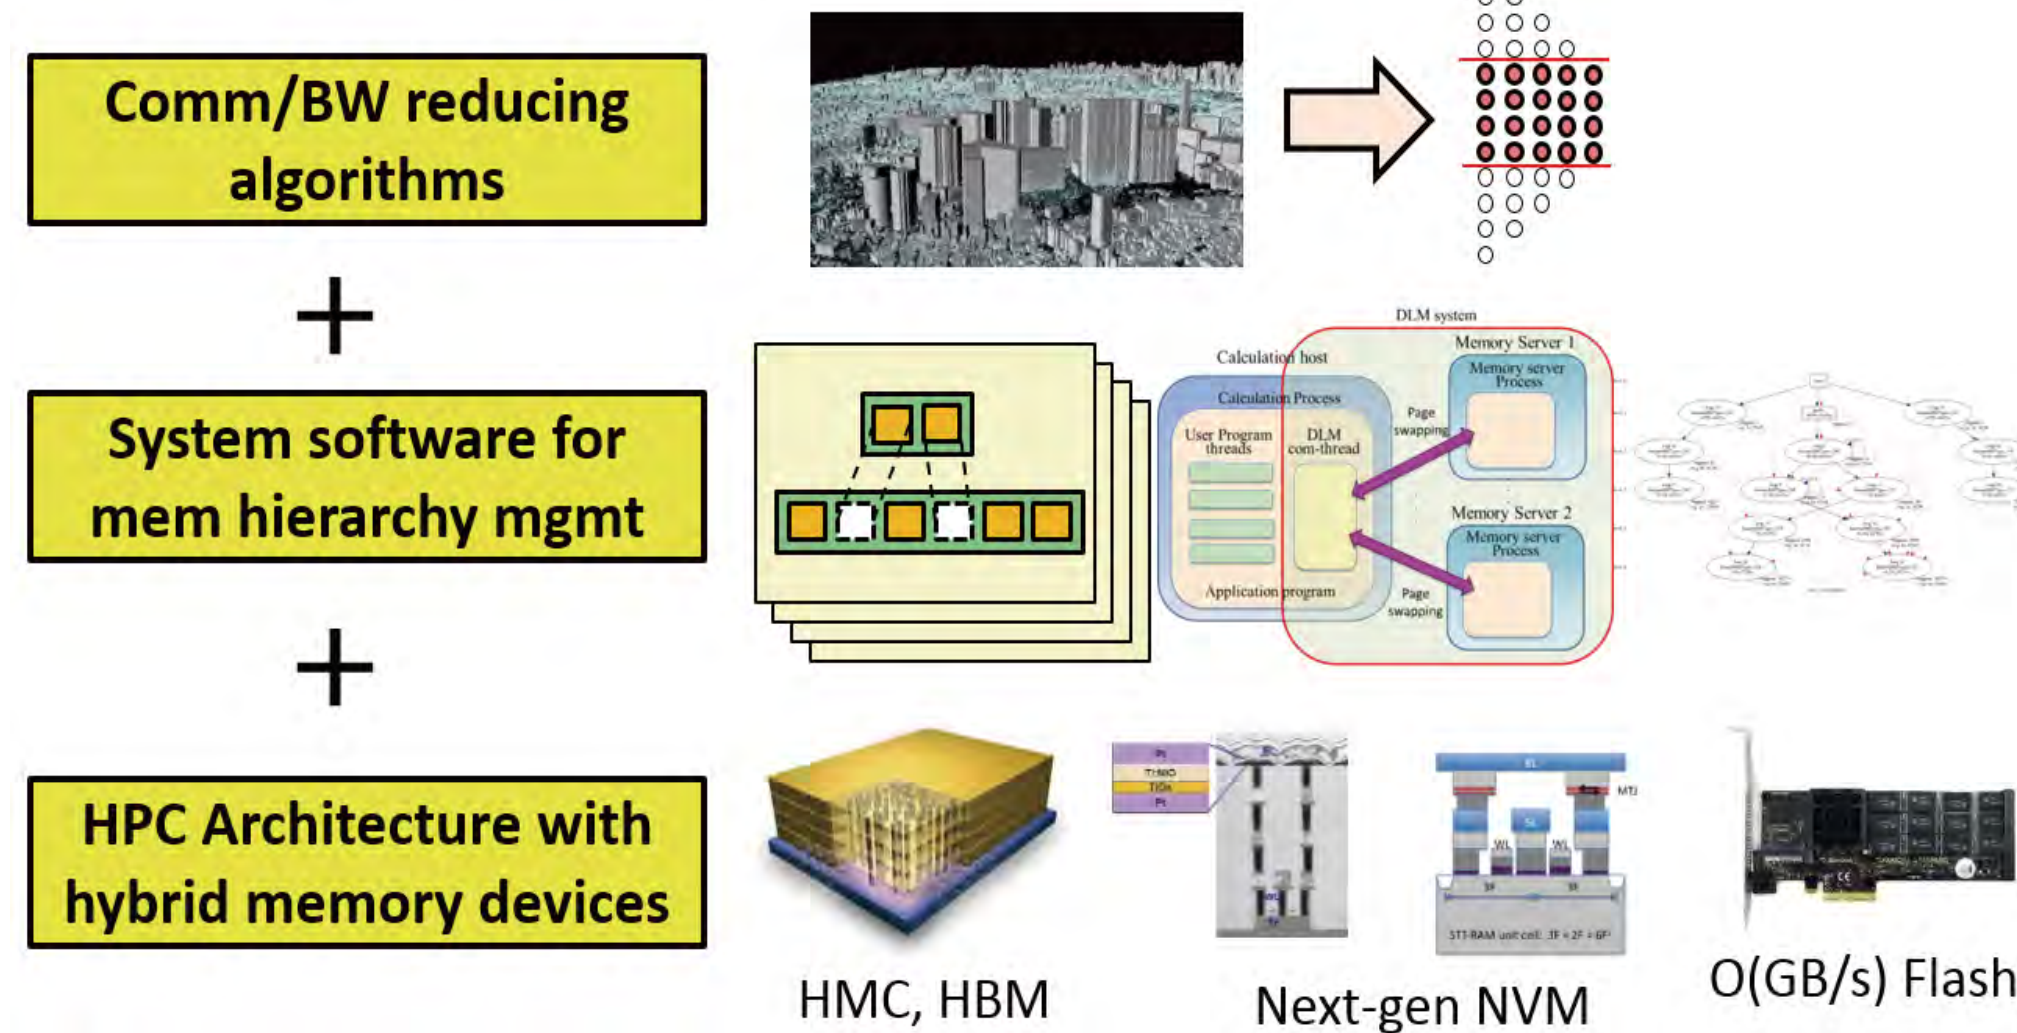
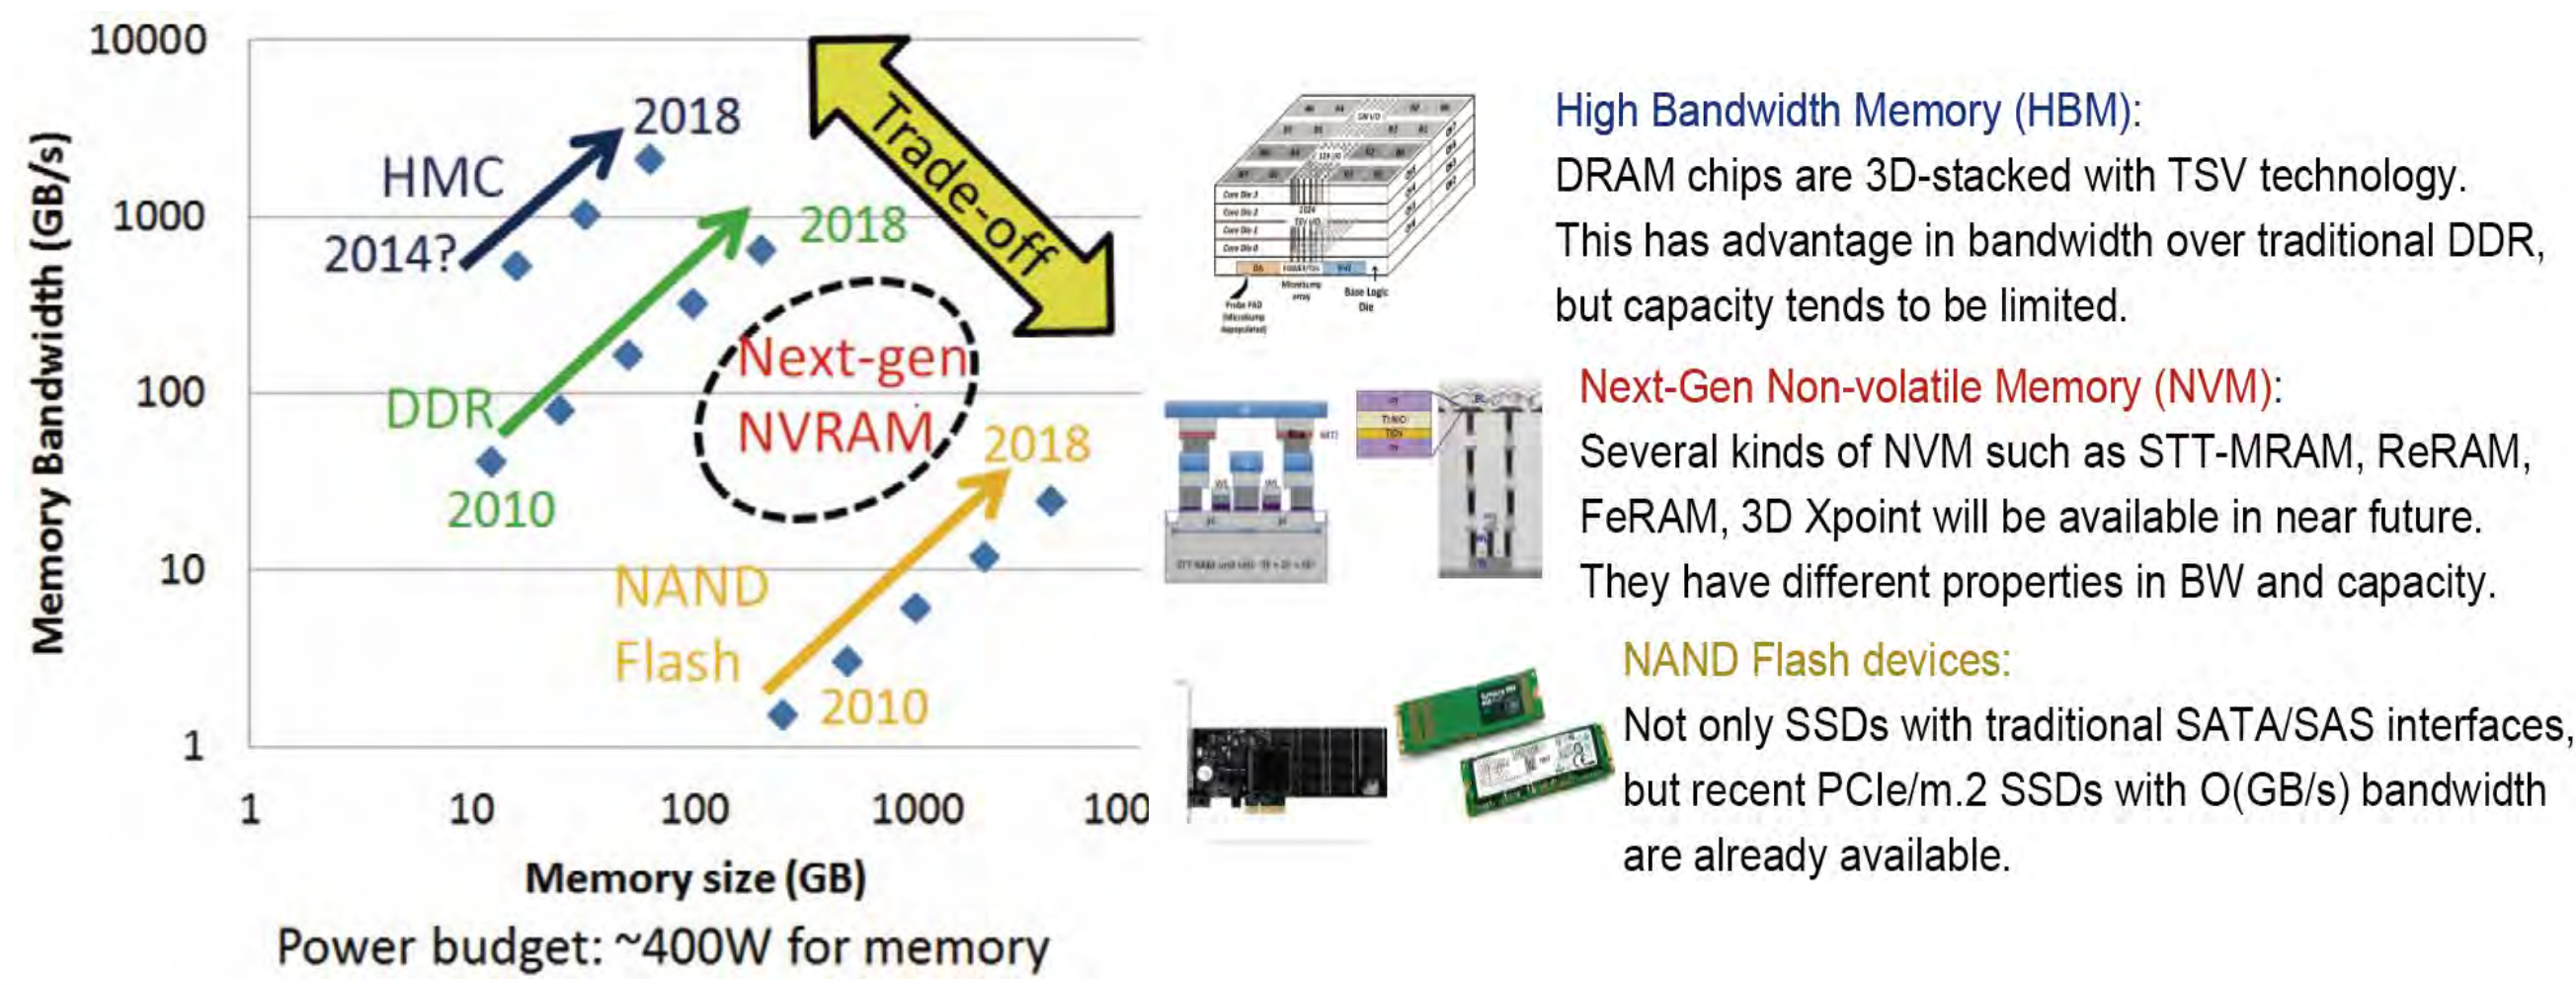


Achieving Extremely Fast&Big Computations on Post-Petascale Supercomputers

Project Overview

On Exa-scale supercomputers, the “Memory Wall” problem will become even more severe, which prevents the realization of Extremely Fast&Big Simulations.

This project promotes research towards this problem via co-design approach among application algorithms, system software, architecture.



Target: Realizing extremely Fast&Big simulations of O(100PB/s) & O(100PB) in Exa-scale era

System Software for Deeper Memory Hierarchy

HHRT: System Software for GPU Memory Swap

Towards achieving fast&big simulations, we need to exploit high speed of upper memory layer (e.g. GDDR/HBM on GPUs) and large capacity of lower memory layer (e.g. NAND Flash).


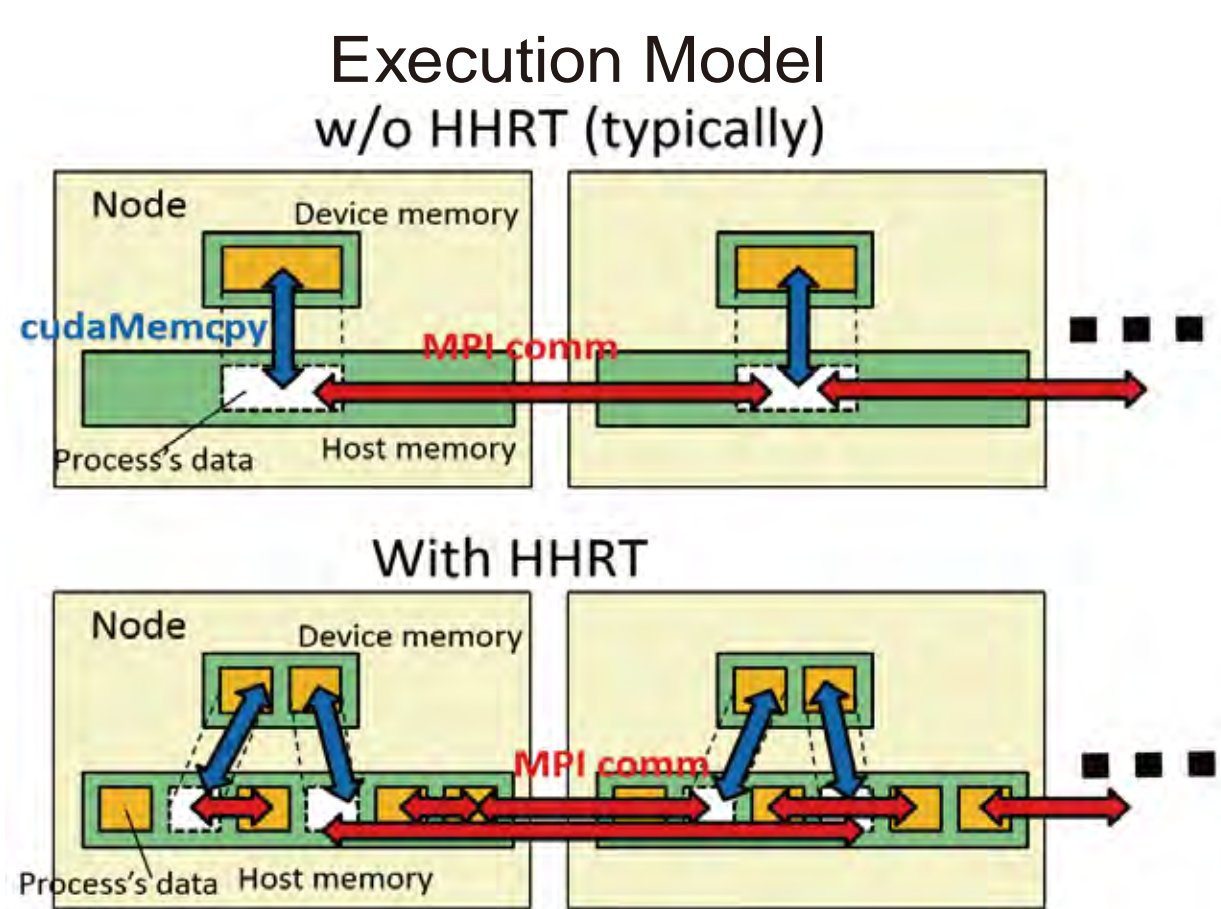
To make memory hierarchy programming easier, we implemented system software, named HHRT (hybrid hierarchical runtime).

(1) HHRT automatically supports data swapping among three memory layers: GPU memory <=> Host memory <=> Flash SSD.

(2) HHRT supports “process-wise” swapping, not “page-wise” like OS.

(3) On the other hand, users still have responsibility to improve locality for better performance, such as Temporal blocking (TB) for stencils

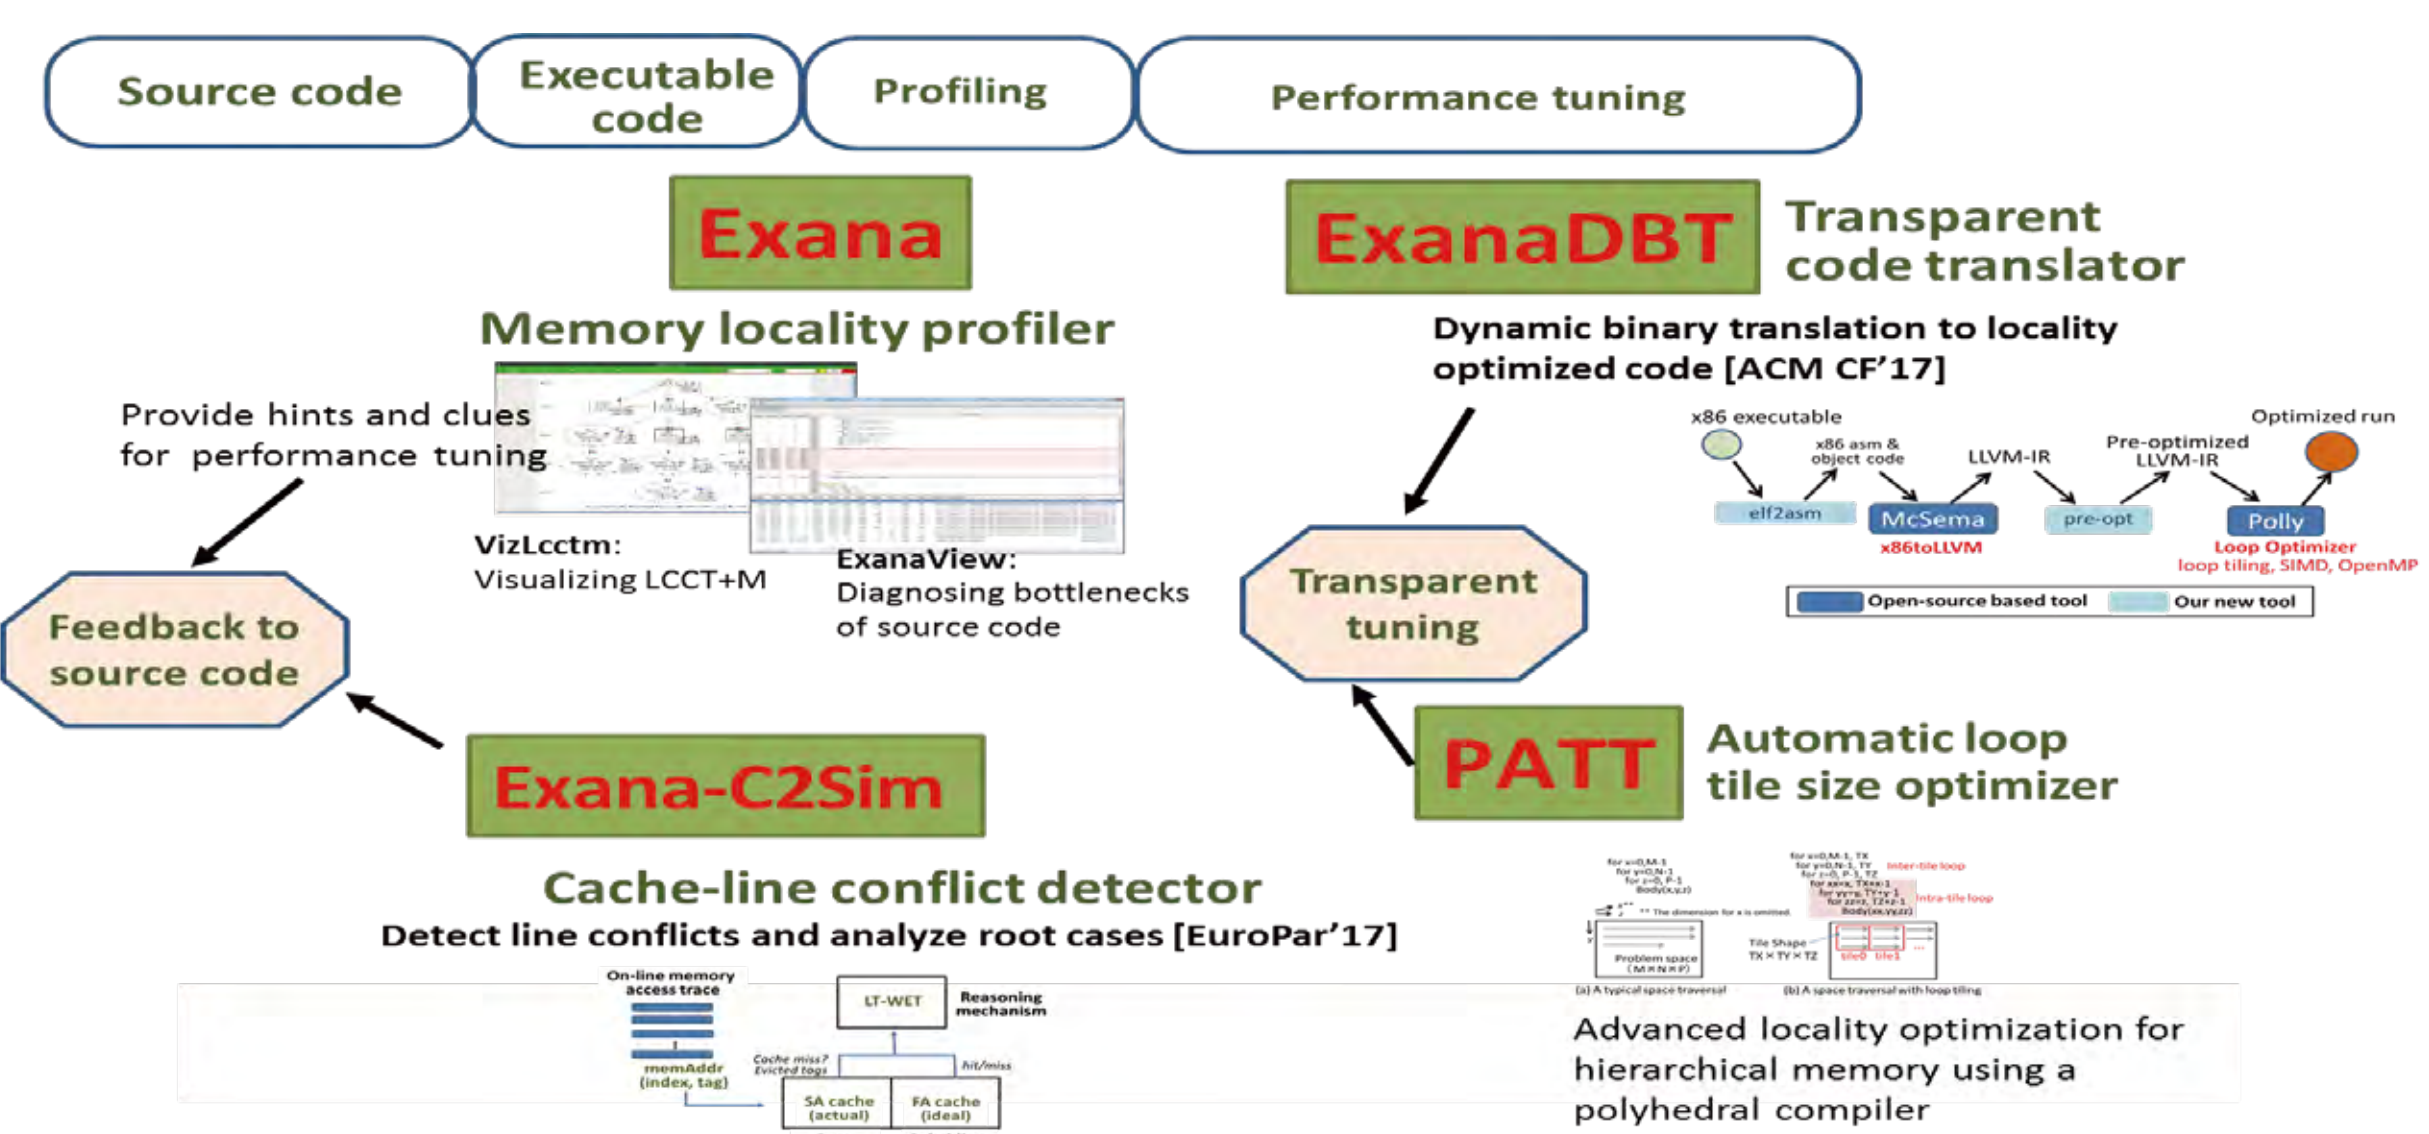
<https://github.com/toshioendo/hhrt>



Tool Chains for Memory Locality Profiling and Performance Tuning

We have been developing tool chains for accelerating system with deeply hierarchical memory. Starting with source code or its executable code, these tools profile and translate applications for the underlying memory subsystems.

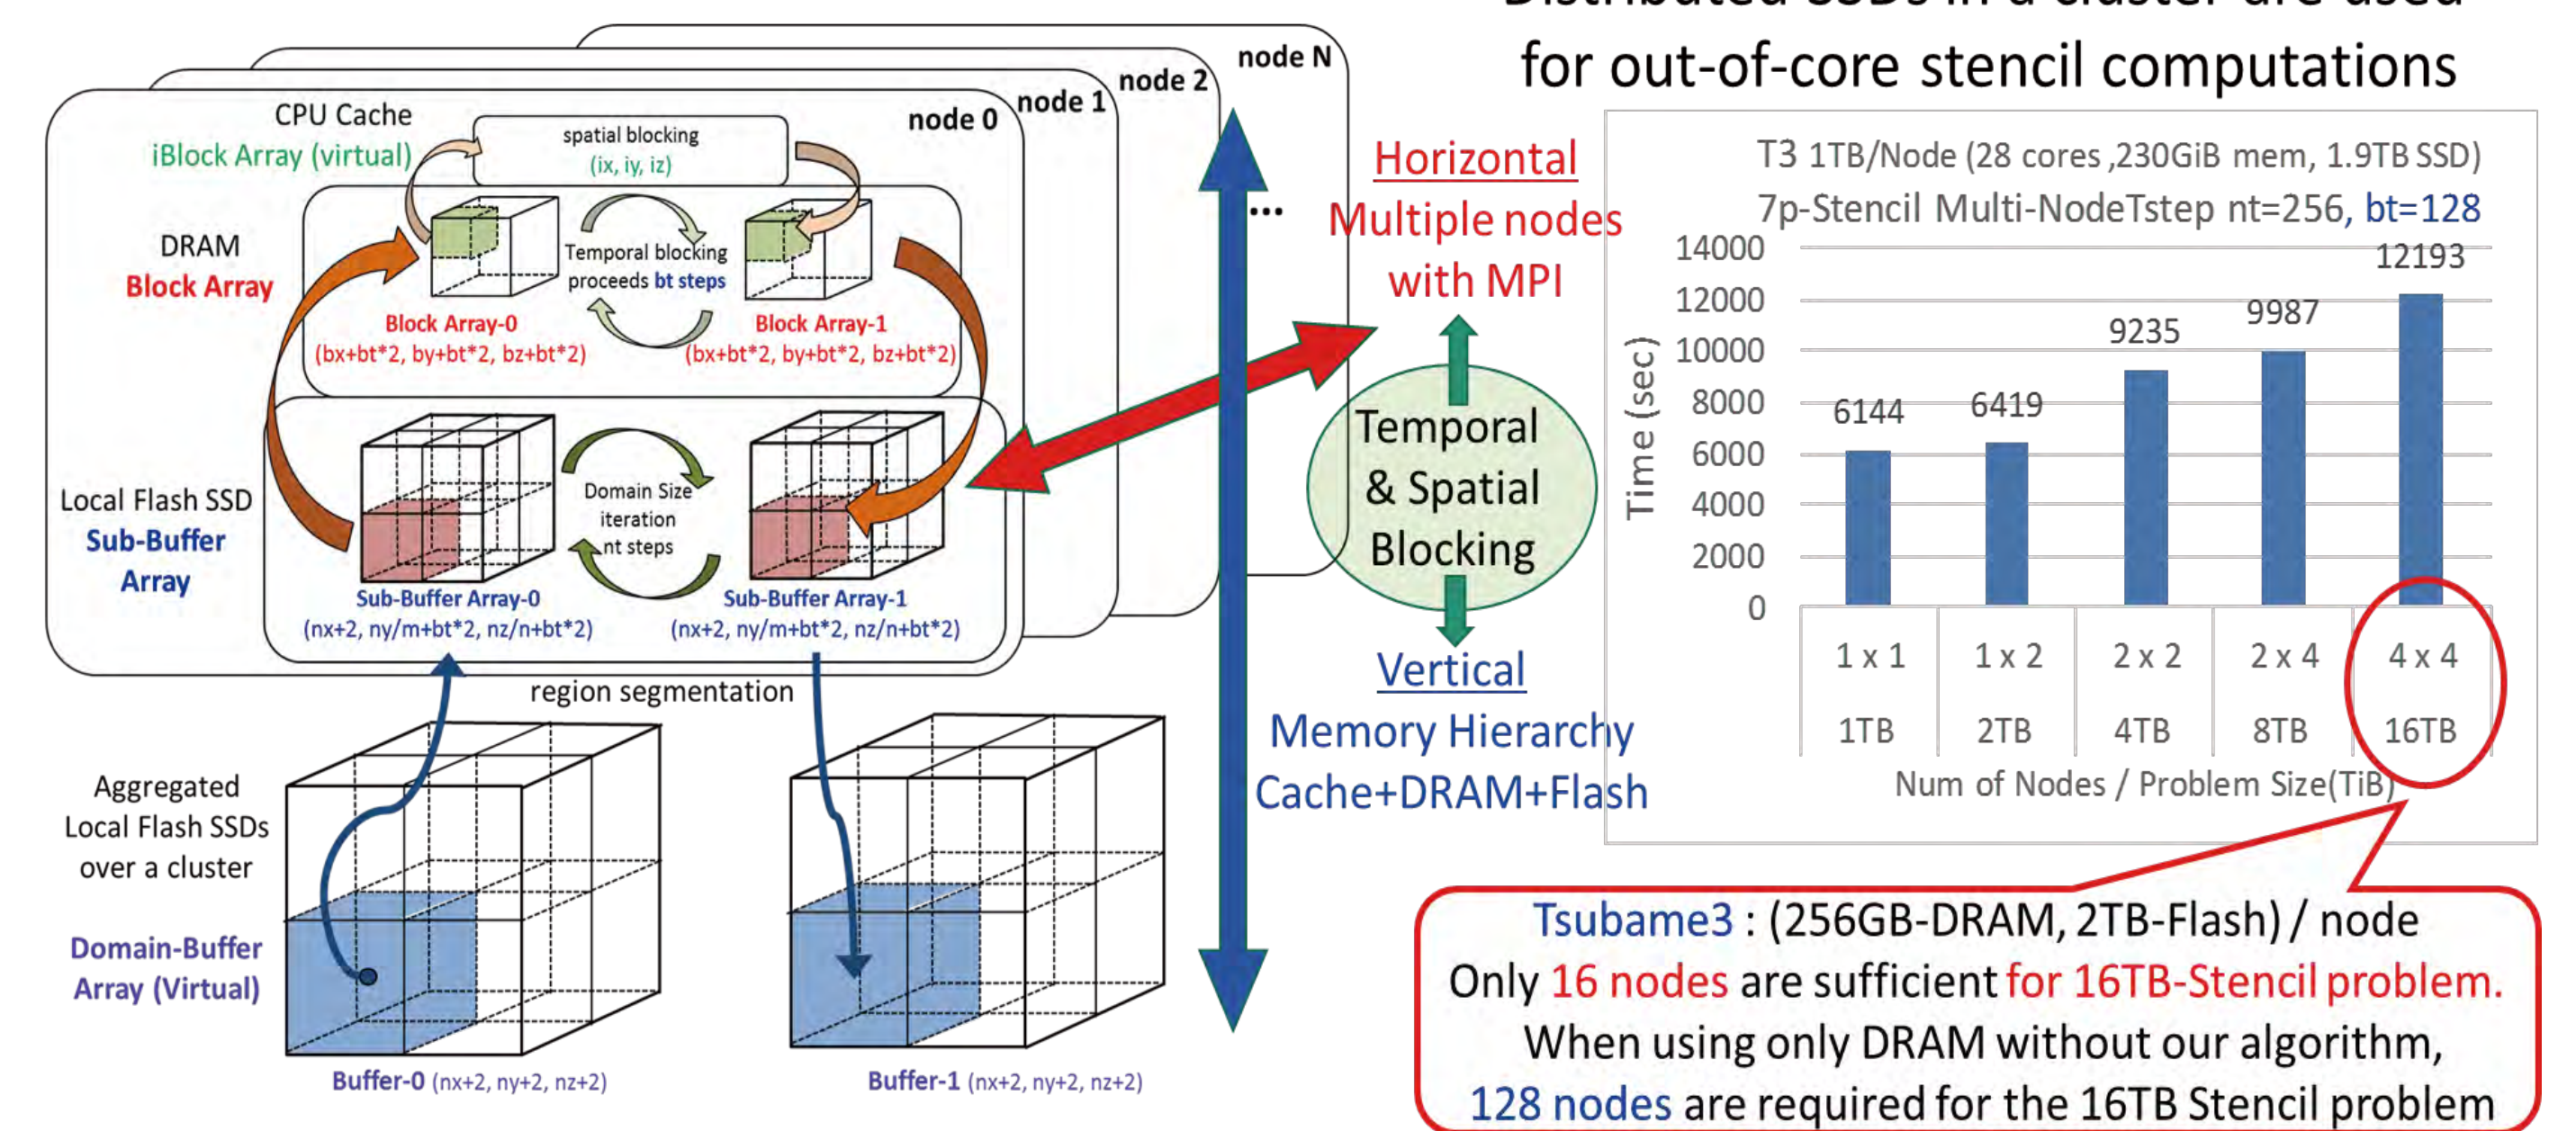
<https://github.com/YukinoriSato/ExanaPkgs>



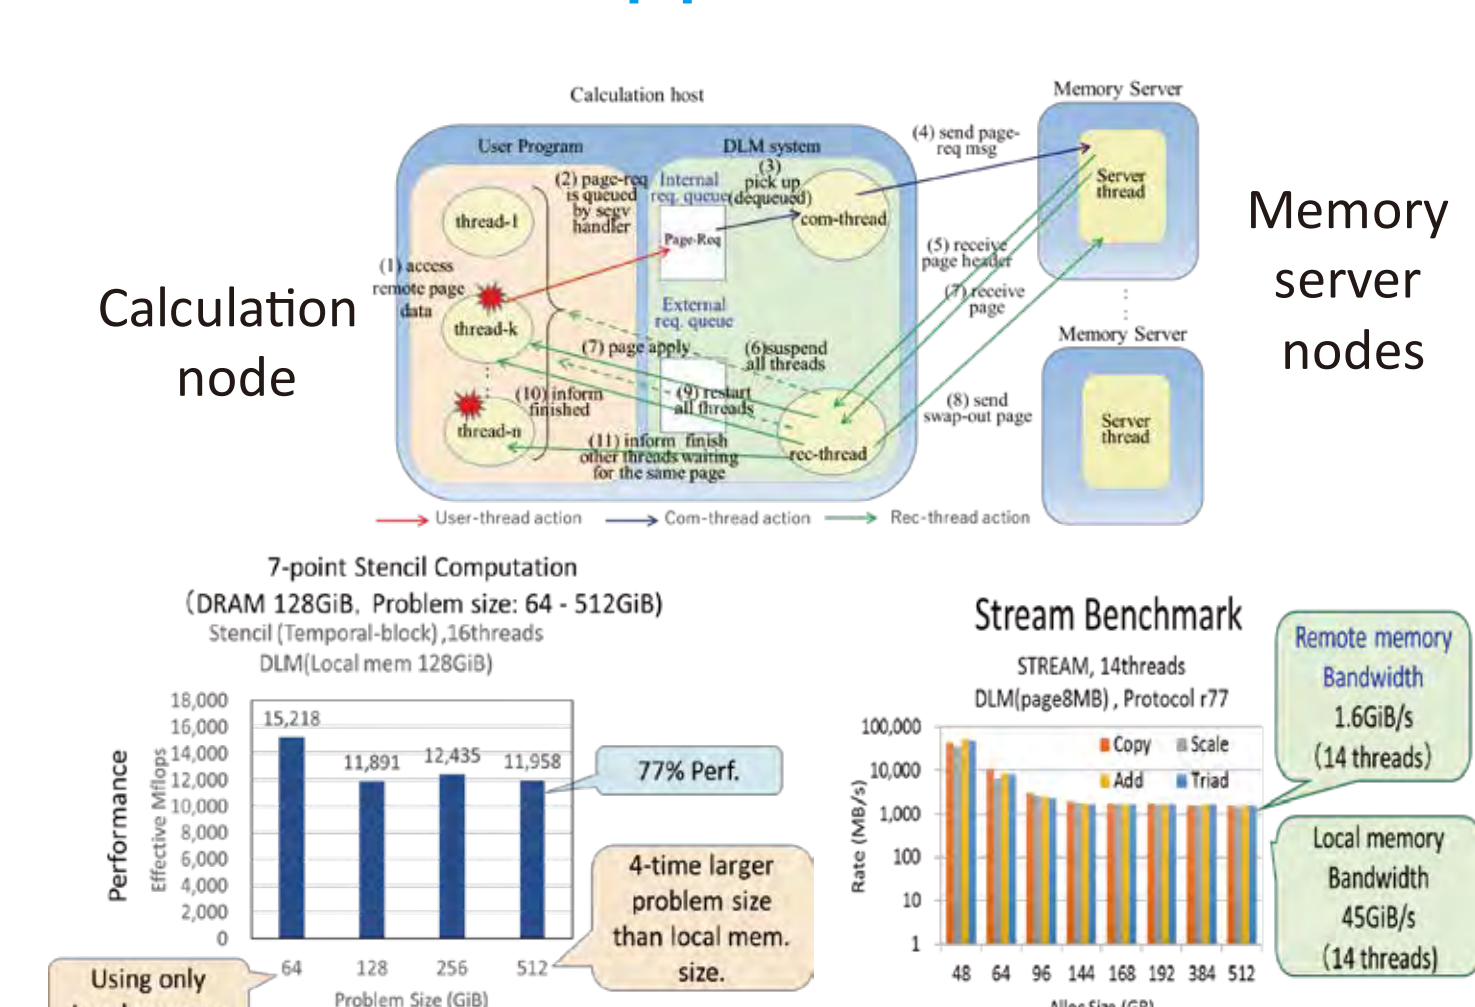
Project Leader: Toshio Endo (Tokyo Tech), Contact: endo@is.titech.ac.jp
Hiroko Midorikawa (Seikei Univ), Yukinori Sato, Shimpei Sato, Noboru Tanabe, Tomoya Yuki (Tokyo Tech)

Horizontal and Vertical Memory Extensions

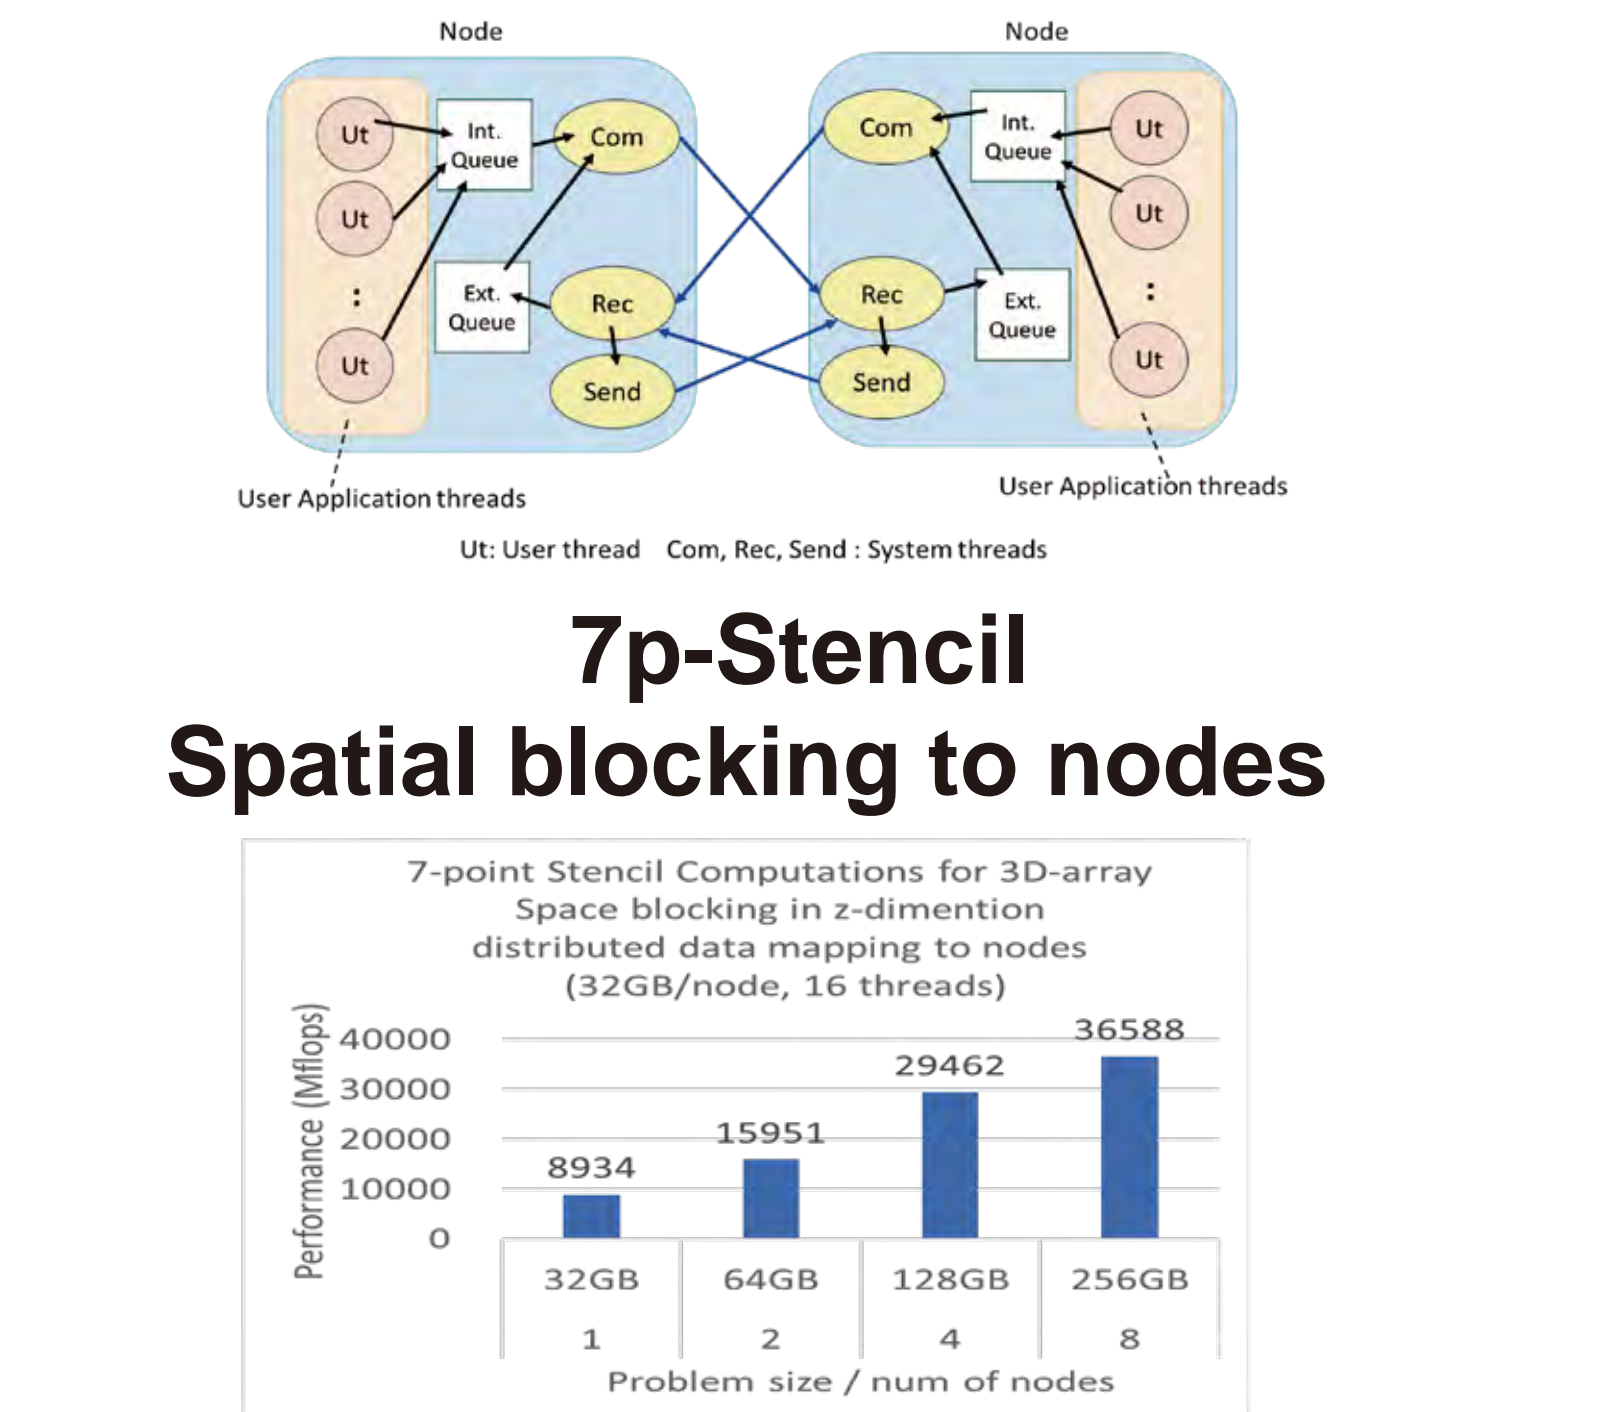
Out-of-Core Stencil Algorithm using Cache, DRAM, Flash SSDs



mDLM: Distributed Large Memory Remote memory paging for multi-thread applications



mSMS: Distributed Shared Memory



Integrating System Software and Real Application

We integrated HHRT system software with the dendrite simulation. This simulation is stencil-based, and originally written in MPI+CUDA, by T. Shimokawabe and T. Aoki.

The loop code has been modified for “temporal blocking”, and linked with HHRT. => High-performance and Larger problem size than GPU device memory

