

# Massively scalable computing method to tackle large eigenvalue problems for nanoelectronics modeling

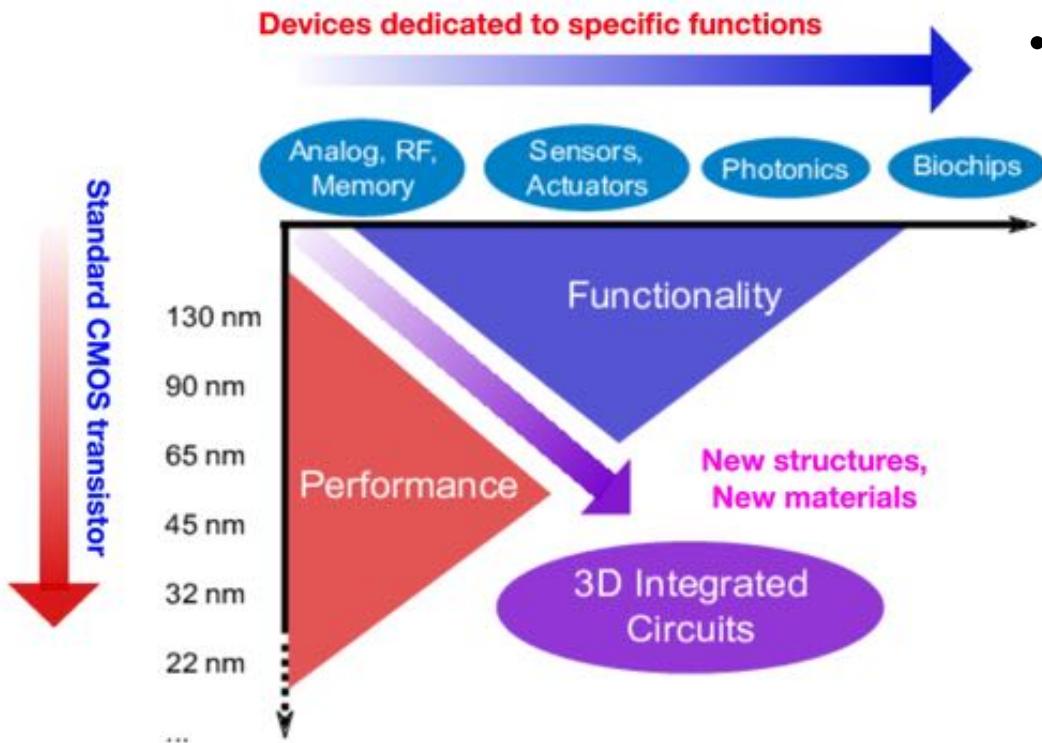
Hoon Ryu, Ph.D.

(E: [elec1020@kisti.re.kr](mailto:elec1020@kisti.re.kr))

Korea Institute of Science and Technology Information (KISTI)  
KISTI Intel® Parallel Computing Center

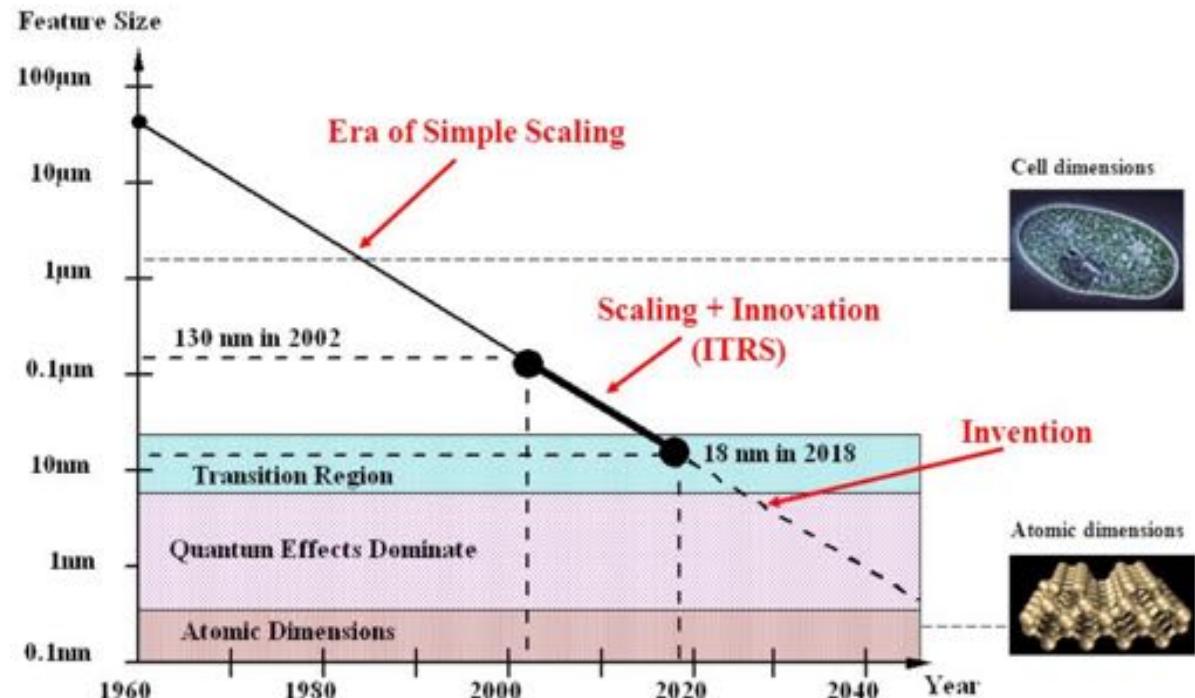
# Two Big Features of Advanced Device Designs

## Miniaturization + Functionalization



- **Size miniaturization of CMOS transistor**
  - Quantum / atomistic effects start to play!
  - New structures to increase the TR density!

- **Diversification in device functions**
  - Devices dedicated to specific functions: Sensors, LEDs
  - No CMOS transistors: Need to explore the feasibility of new materials and structures beyond CMOS

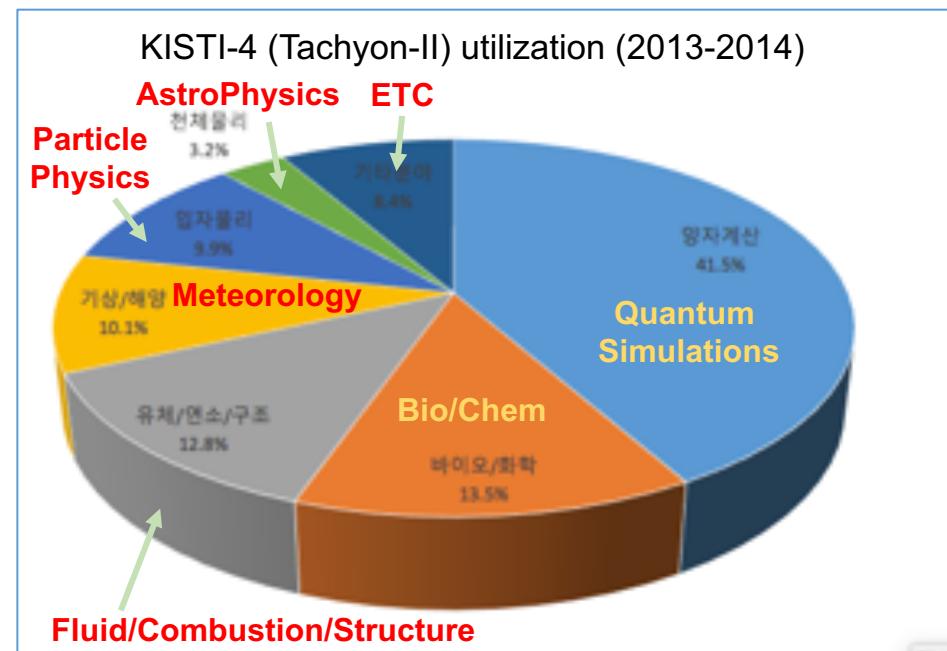
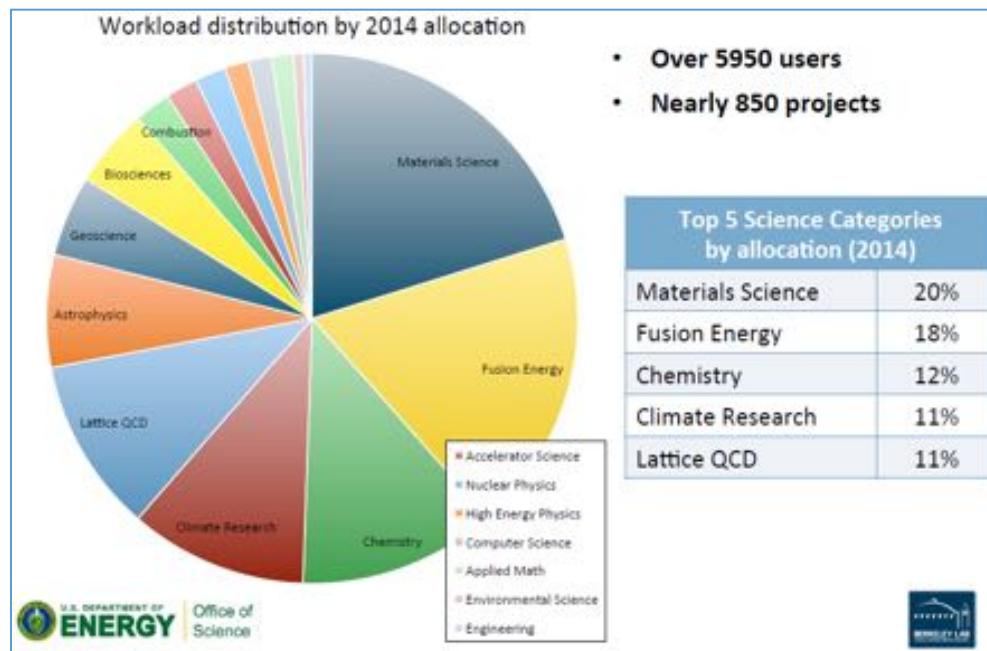


# Electronic Structure Calculations

## Prediction of electron motions in nanoscale materials



- The state of electron motions in an electrostatic field created by the stationary nuclei.
  - Prediction of electron motions in nanoscale materials and devices
- Physics, Chemistry, Materials Science, Electrical Engineering
  - Huge customers in the society of computational science

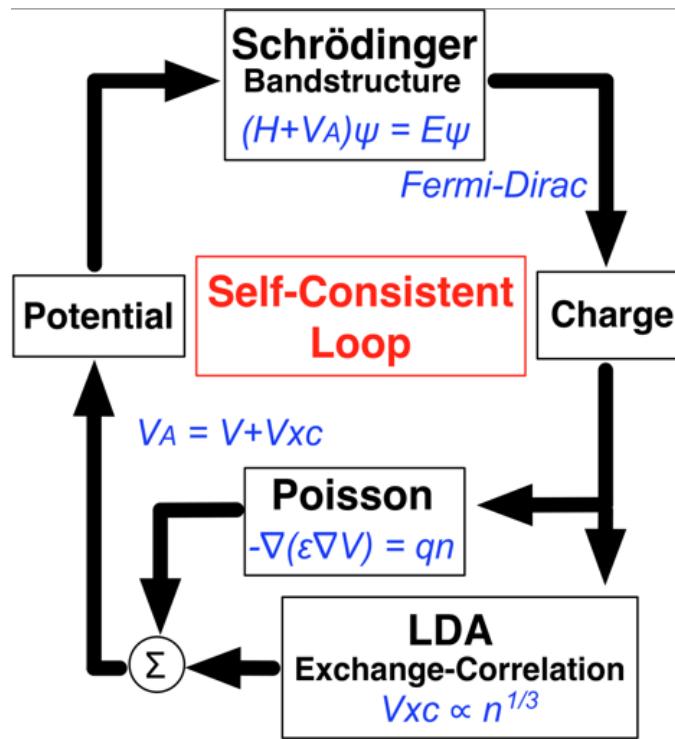


# Electronic Structure Calculations

In a perspective of “numerical analysis”



- Two PDE-coupled Loop: **Schrödinger Equation** and **Poisson Equation**
- Both equations involve **system matrices** (Hamiltonian and Poisson)  
→ DOFs of those matrices are proportional to the # of grids in the simulation domains



- (Stationary) Schrödinger Equations

→ **Normal Eigenvalue Problem**

$$H\Psi = E\Psi$$

- Poisson Equations

→ **Linear System Problem**

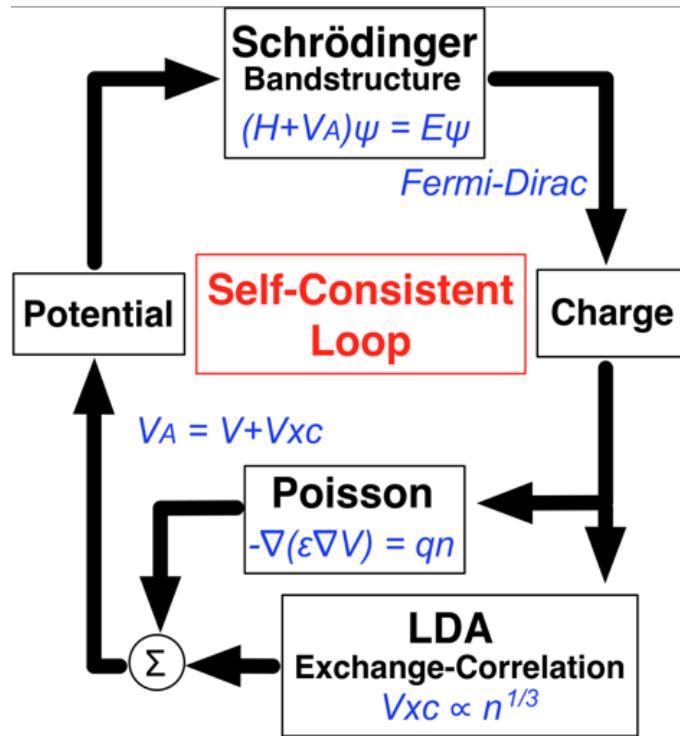
$$-\nabla(\epsilon\nabla V) = \rho \rightarrow Ax = b$$

# Electronic Structure Calculations

In a perspective of “numerical analysis”



- Two PDE-coupled Loop: **Schrödinger Equation** and **Poisson Equation**
- Both equation involve **system matrices** (Hamiltonian and Poisson)  
→ DOFs of those matrices are proportional to the # of grids in the simulation domains



- Schrödinger Equations  
→ **Normal Eigenvalue Problem**
- Poisson Equations  
→ **Linear System Problem**

$$H\Psi = E\Psi$$

$$-\nabla(\epsilon\nabla V) = \rho \rightarrow Ax = b$$

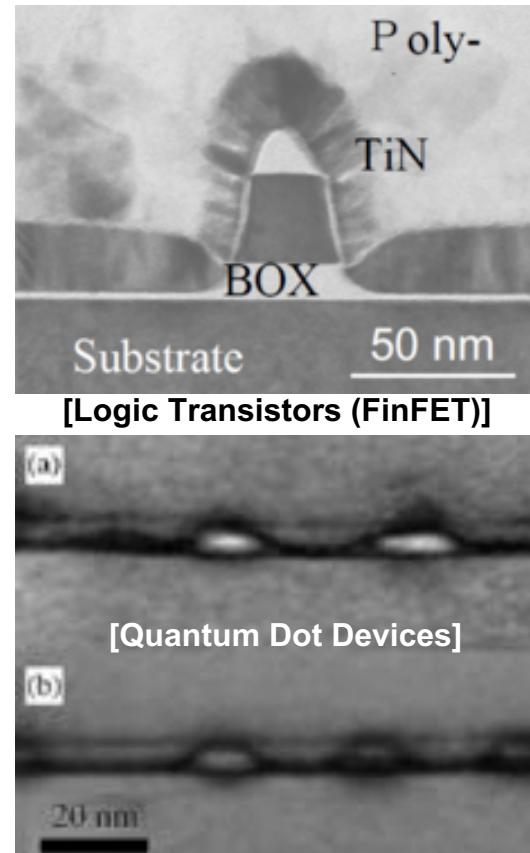
**How large are these system matrices?**  
**Why do we need to handle those?**

# Needs for “Large” Electronic Structures

Electron motion happens in cores, but we need more



1. Quantum Simulations of “Realizable” Nanoscale Materials and Devices  
→ Needs to handle large-scale atomic systems (~ A few tens of nms)



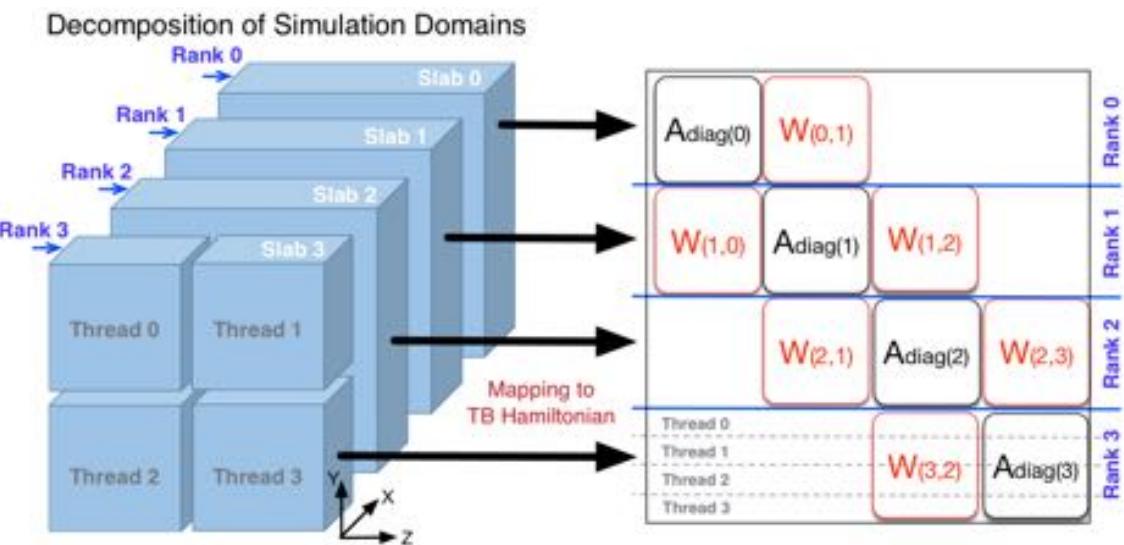
30nm<sup>3</sup> Silicon Box? → About million atoms

2. DOF of Matrices of Governing Equations  
→ Linearly proportional to # of atoms (w/ some weight)

### 3. Parallel Computing

$$Ax = b$$

$$H\Psi = E\Psi$$



# Development Strategy: DD, Matrix Handling

## System matrices for Schrödinger and Poisson equations



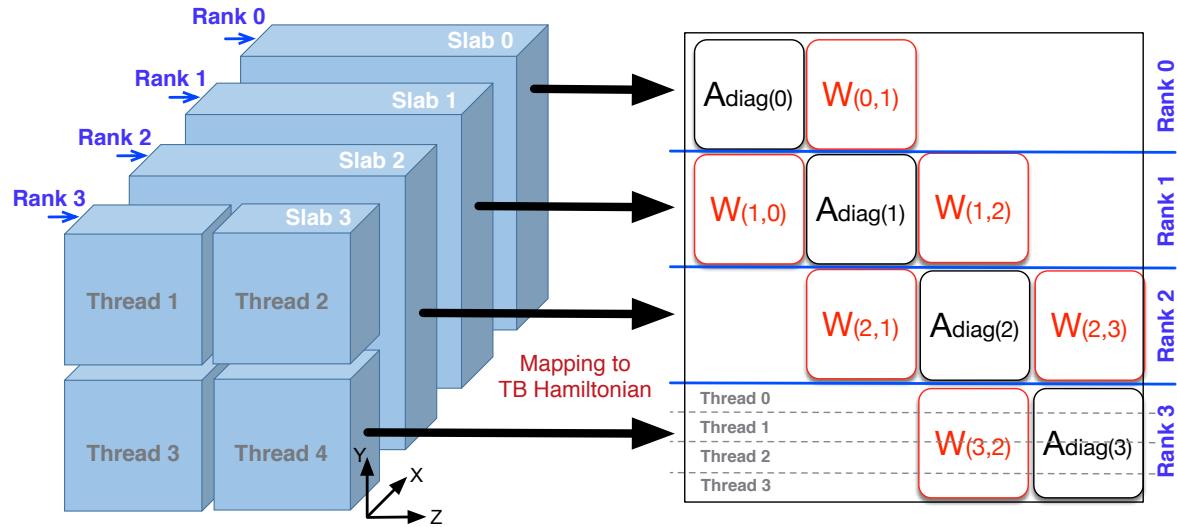
### Schrödinger Equation

- Normal Eigenvalue Problem (Electronic Structure)
- Hamiltonian is always symmetric

$$H\Psi = E\Psi$$

### Domain Decomposition

- MPI + OpenMP
- Effectively multi-dimensional decomposition



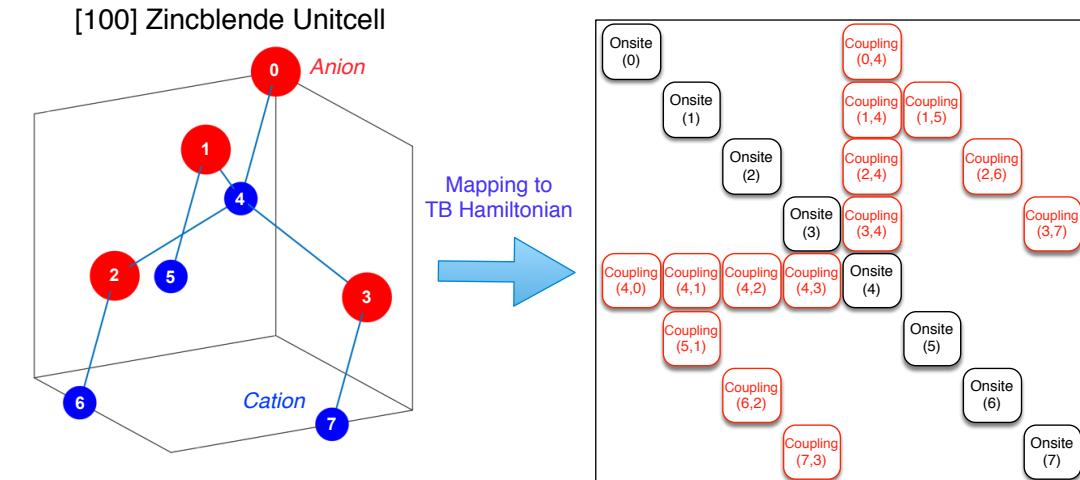
### Poisson Equation

- Linear System Problem (Electrostatics: Q-V)
- Poisson matrix is always symmetric

$$-\nabla(\epsilon \nabla V) = \rho$$

### Matrix Handling

- Tight-binding Hamiltonian (Schrödinger Eq.)
- Finite Difference Method (Poisson Eq.)
- Nearest Neighbor Coupling: Highly Sparse → CSR

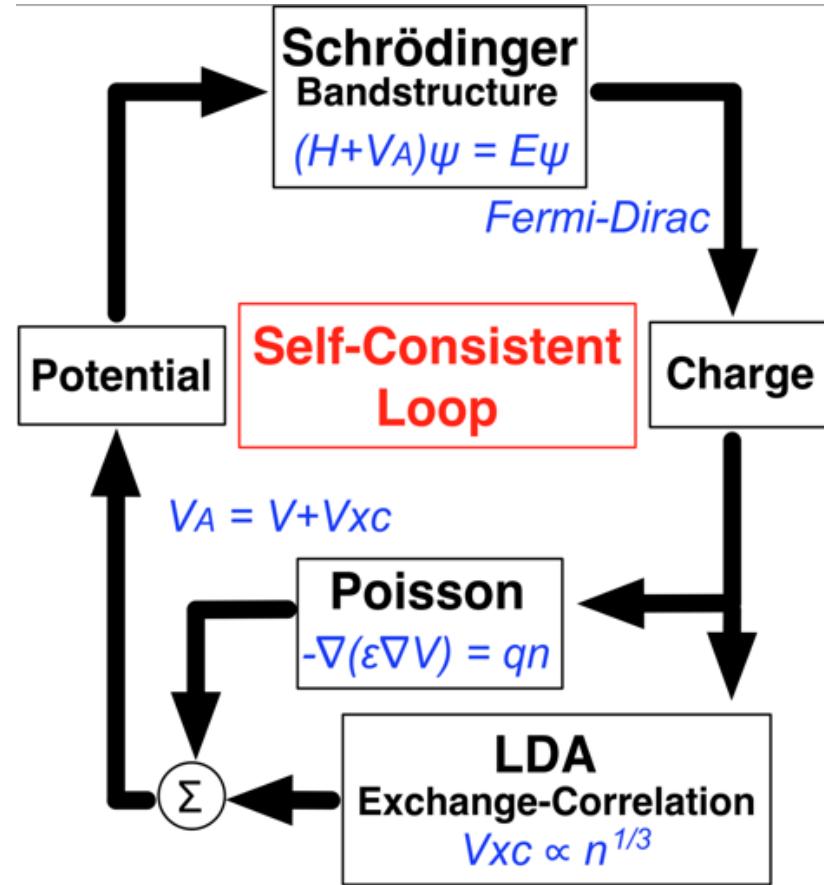


# Development Strategy: Numerical Algorithms

## Schrödinger equations



### Self-consistent Loop for Device Simulations



### Schrödinger Eqs. w/ LANCZOS Algorithm

- C. Lanczos, *J. Res. Natl. Bur. Stand.* 45, 255
- Normal Eigenvalue Problem (Electronic Structure)
  - Hamiltonian is always symmetric
  - Original Matrix → T matrix-reduction 
$$H\Psi = E\Psi$$
  - Steps for Iteration: Purely Scalable Algebraic Ops.

$\mathbf{v}_i$ : ( $N \times 1$ ) vectors ( $i = 0, \dots, K$ );  $a_i$  and  $b_i$ : scalars ( $i = 1, \dots, K$ )

$\mathbf{v}_0 \leftarrow \mathbf{0}$ ,  $\mathbf{v}_1$  = random vector with norm 1 ;

$\mathbf{b}_1 \leftarrow \mathbf{0}$  ;

loop for ( $j=1; j \leq K; j++$ )

$\mathbf{w}_j \leftarrow \mathbf{A}\mathbf{v}_j$  ;

$\mathbf{a}_j \leftarrow \mathbf{w}_j^* \mathbf{v}_j$  ;

$\mathbf{w}_j \leftarrow \mathbf{w}_j - \mathbf{a}_j \mathbf{v}_j - \mathbf{b}_j \mathbf{v}_{j-1}$  ;

$\mathbf{b}_{j+1} \leftarrow \|\mathbf{w}_j\|$  ;

$\mathbf{v}_{j+1} \leftarrow \mathbf{w}_j / \mathbf{b}_{j+1}$  ;

construct T matrix;

end loop

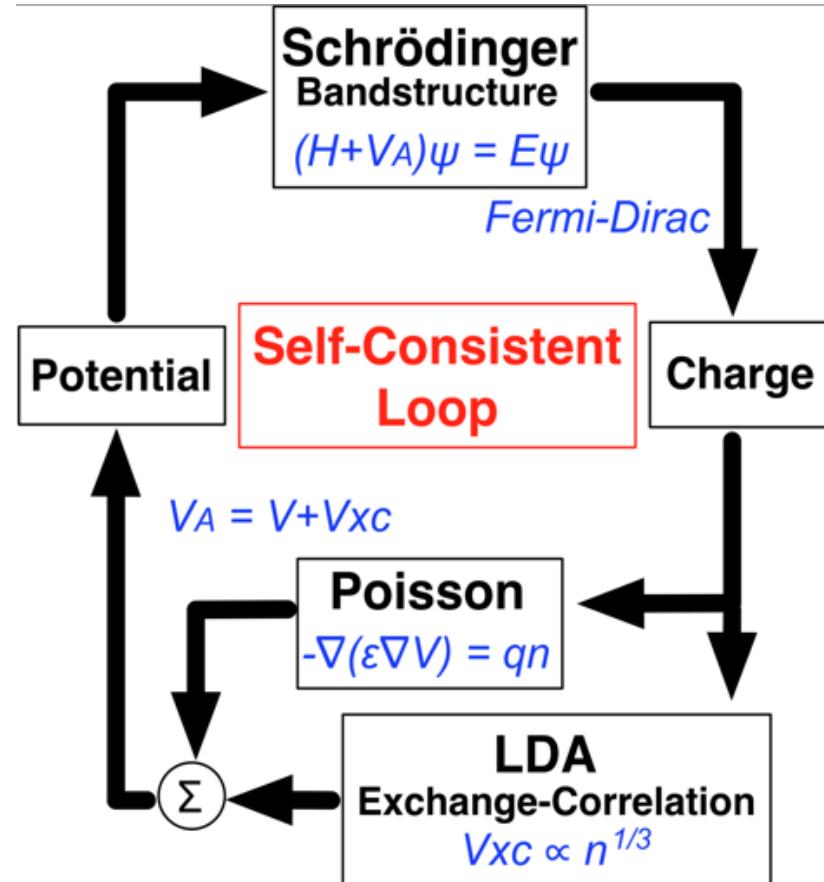
$$T = \begin{pmatrix} a_1 & b_2 & 0 & \cdots & \cdots & 0 \\ b_2 & a_2 & b_3 & & & \vdots \\ 0 & b_3 & \ddots & \ddots & & \vdots \\ \vdots & \ddots & \ddots & \ddots & b_{k-1} & 0 \\ \vdots & & \ddots & \ddots & b_{k-1} & a_{k-1} & b_k \\ 0 & \cdots & \cdots & 0 & b_k & a_k \end{pmatrix}$$

# Development Strategy: Numerical Algorithms

## Poisson equations



### Self-consistent Loop for Device Simulations



### Poisson Eqs. w/ CG Algorithm

- A Problem of Solving Linear Systems
- Conv. Guaranteed: Symmetric & Positive Definite
- Poisson is always S & PD.
- Steps for Iteration: Purely Scalable

$$-\nabla(\epsilon\nabla V) = \rho$$

Algebraic Ops.

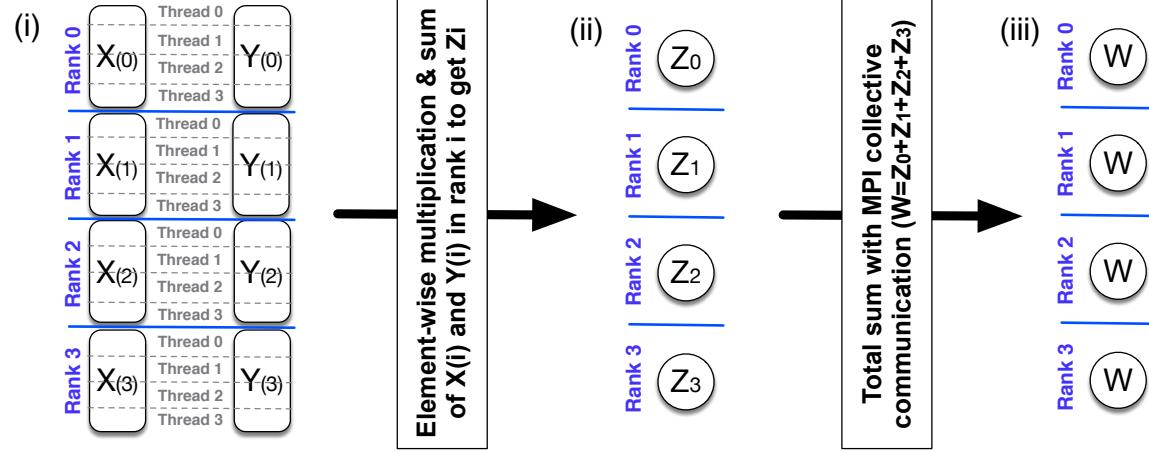
We want to solve  $Ax = b$ . First compute  $r_0 = b - Ax_0$ ,  $p_0 = r_0$   
loop for ( $j=1$ ;  $j \leq K$ ;  $j++$ )  
 $a_j \leftarrow \langle r_j \cdot r_j \rangle / \langle A p_j \cdot p_j \rangle$ ;  
 $x_{j+1} \leftarrow x_j + a_j p_j$ ;  
 $r_{j+1} \leftarrow r_j - a_j A p_j$ ;  
if ( $\|r_{j+1}\| / \|r_0\| < \epsilon$ )  
    declare  $r_{j+1}$  is the solution of  $Ax = b$  and break the loop  
 $c_j \leftarrow \langle r_{j+1} \cdot r_{j+1} \rangle / \langle r_j \cdot r_j \rangle$ ;  
 $p_{j+1} \leftarrow r_{j+1} + c_j p_j$ ;  
end loop

# Performance Bottleneck?

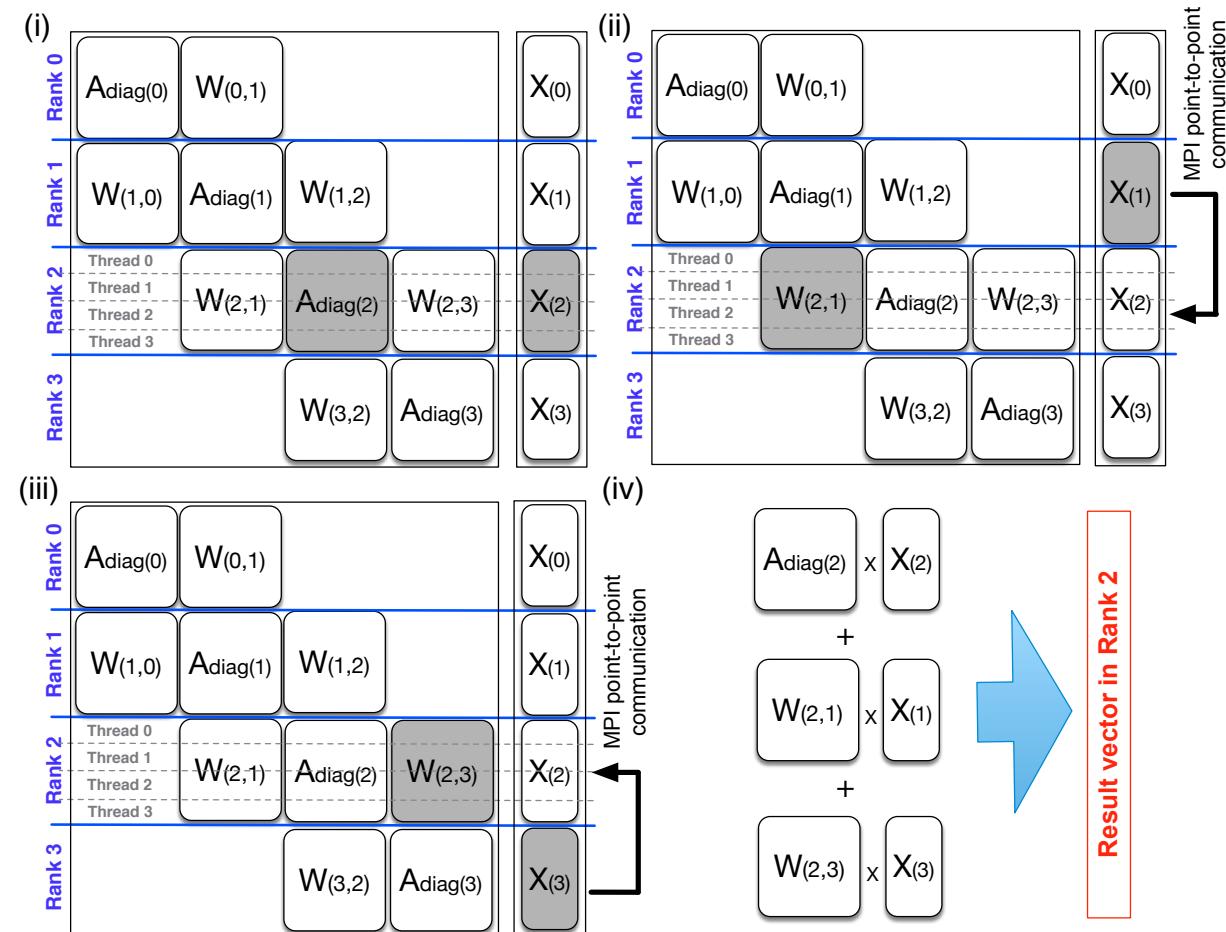
## Matrix-vector multiplier: Sparse matrices



### Vector Dot-Product (VVDot)



### (Sparse) Matrix-vector Multiplier (MVMul)



### Main Concerns for Performance

- Collective Communication
  - May not be the main bottleneck as we only need to collect a single value from a single MPI process
- Matrix-vector Multiplier
  - Communication happens, but would not be a critical problem as it only happens between adjacent ranks
  - Data locality affects vectorization efficiency

# Performance w/ Intel® KNL Processors

## Single-node Performance: Whole domain in MCDRAM



### Description of BMT Target and Test Mode

- 10 CB states in  $16 \times 43 \times 43 (\text{nm}^3)$  [100] Si:P quantum dot
  - Material candidates for Si Quantum Info. Processors (Nature Nanotech. **9**, 430)
  - $15.36 \times 15.36 \text{M}$  Hamiltonian Matrix ( $\sim 11 \text{GB}$ )
- Xeon Phi 7210: 64 cores
- MCDRAM control w/ [numactrl](#); Quad Mode

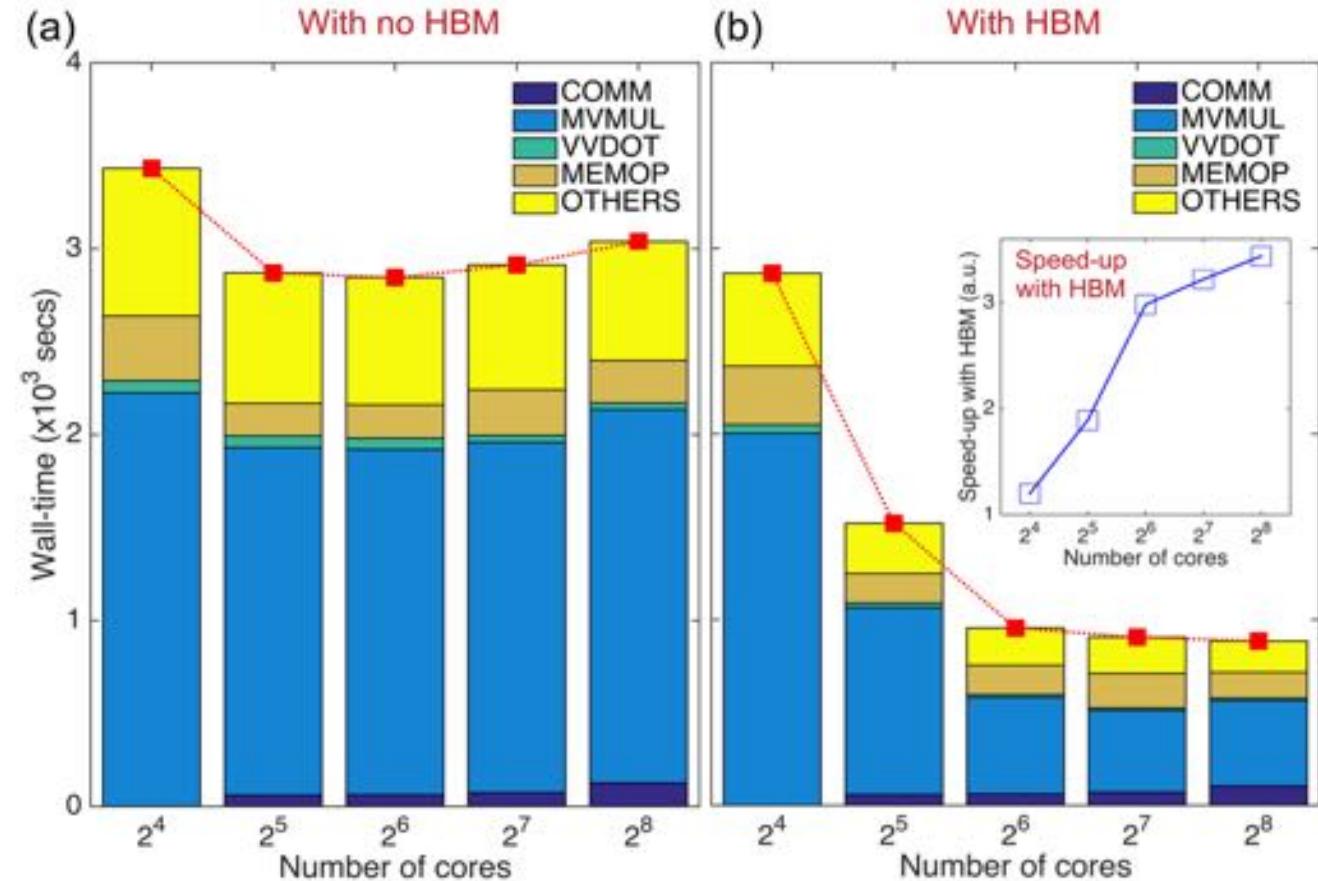
### Results

[H. Ryu, Intel® HPC Developer Conference \(2017\)](#)

- With no MCDRAM
  - No clear speed-up beyond 64 cores
- **With MCDRAM**
  - Up to  $\sim 4x$  speed-up w.r.t. the case w/ no MCDRAM
  - Intra-node scalability up to 256 cores

### Points of Questions

- How is the performance compared to the one under other computing environments? (GPU, CPU-only etc..)
  - In terms of speed and **energy consumption**



$2^4$  cores = (1 MPI proc(s), 16 threads),  $2^7$  cores = (2 MPI proc(s), 64 threads)  
 $2^5$  cores = (2 MPI proc(s), 16 threads),  $2^8$  cores = (4 MPI proc(s), 64 threads)  
 $2^6$  cores = (2 MPI proc(s), 32 threads)

# Performance w/ Intel® KNL Processors

## Speed in various computing platforms

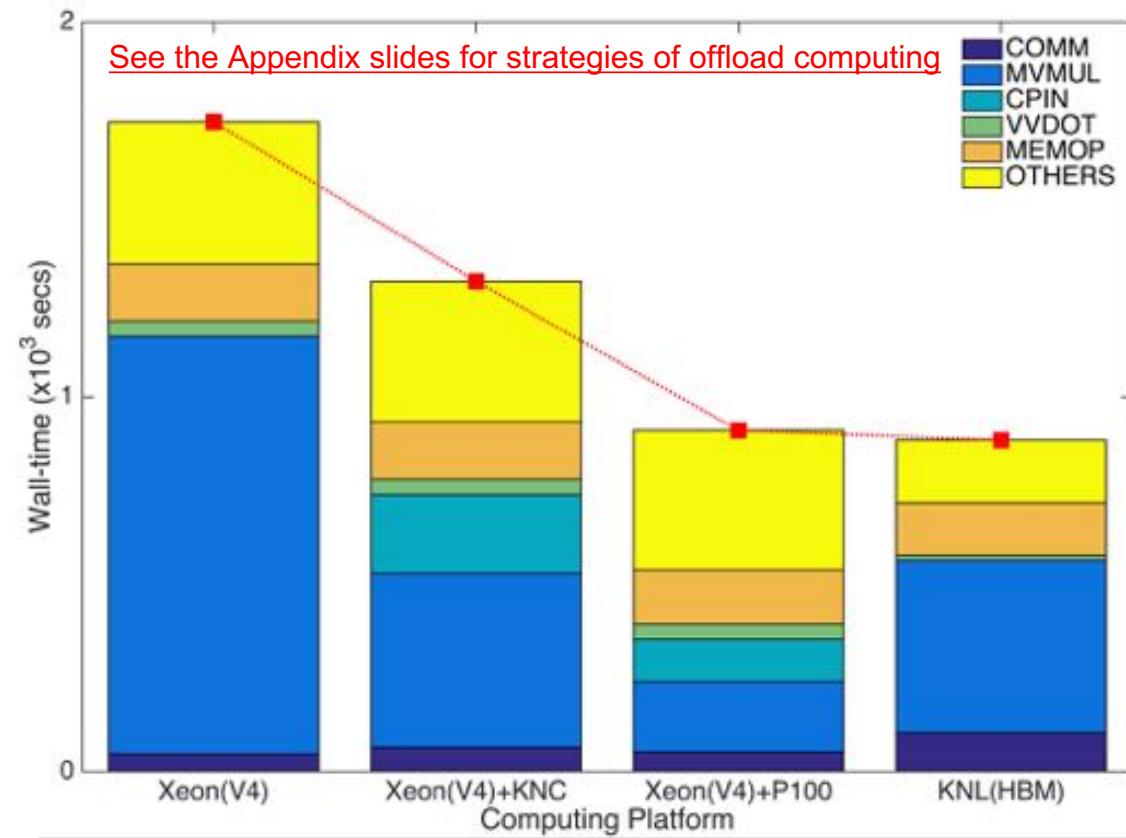


### Description of BMT Target and Test Mode

- 10 CB states in  $16 \times 43 \times 43 (\text{nm}^3)$  [100] Si:P quantum dot
  - Material candidates for Si Quantum Info. Processors (Nature Nanotech. **9**, 430)
  - $15.36 \times 15.36 \times 15.36$  M Hamiltonian Matrix ( $\sim 11$  GB)
- Specs of Other Platforms
  - Xeon(V4): 24 cores of Broadwell (BW) 2.50GHz
  - Xeon(V4)+KNC: 24 cores BW + 2 KNC 7120 cards
  - Xeon(V4)+P100: 24 cores BW + 2 P100 cards
  - KNL(HBM): the one described so far

### Results

- **KNL slightly beats Xeon(V4)+P100**
  - Copy-time (CPIN): a critical bottleneck of PCI-E devices
  - P100 shows better kernel speed, but the overall benefit reduces due to data-transfer between host and devices
  - CPIN would even increase if we consider periodic BCs
- **Another critical figure of merit: Energy-efficiency**



Xeon(V4) = (2 MPI proc(s), 12 threads)  
Xeon(V4) + KNC = (2 MPI proc(s), 12 threads) + 2 KNC 7120 cards  
Xeon(V4) + P100 = (2 MPI proc(s), 12 threads) + 2 P100 cards  
KNL (HBM) = (4 MPI proc(s), 64 threads) with HBM

# Performance w/ Intel® KNL Processors

## Energy consumption in various computing platform

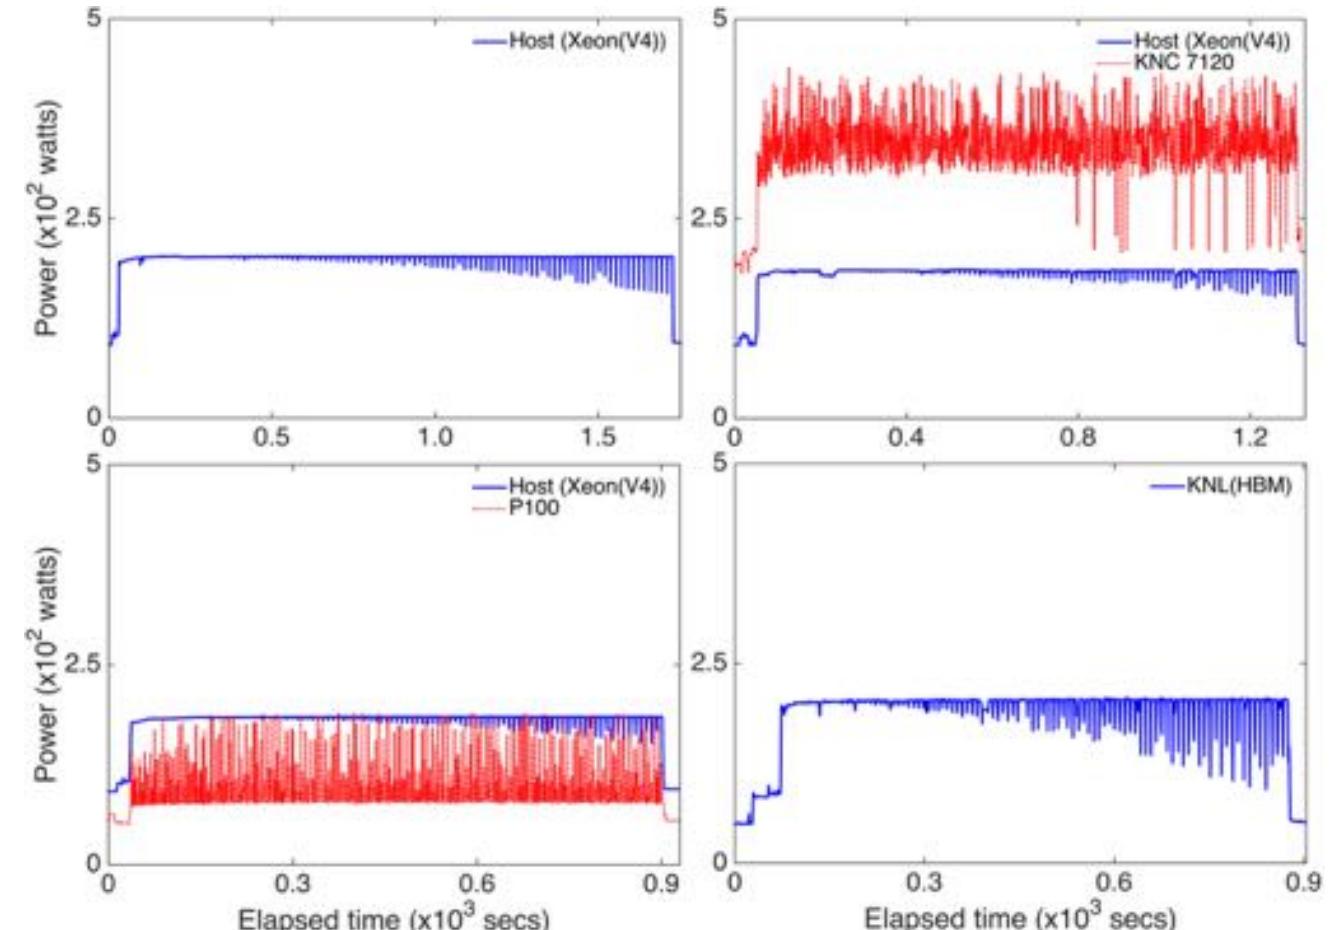
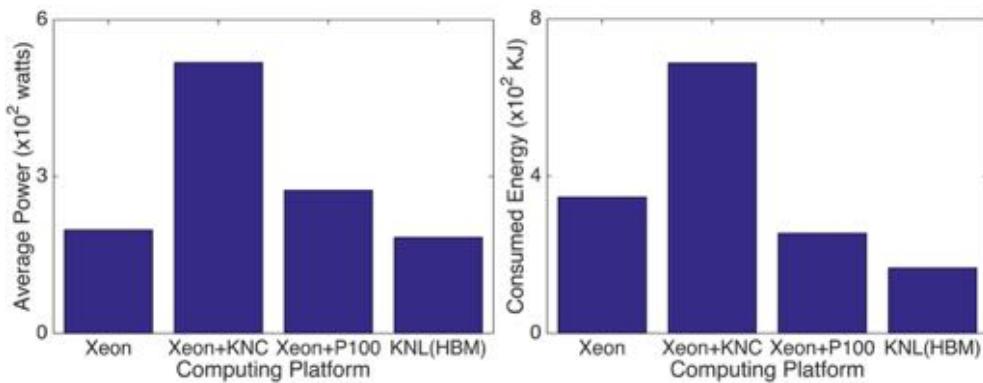


### Description of BMT Target and Test Mode

- 10 CB states in  $16 \times 43 \times 43 (\text{nm}^3)$  [100] Si:P quantum dot  
→ Hamiltonian DOF:  $15.36 \text{M} \times 15.36 \text{M}$  ( $\sim 11 \text{GB}$ )
- Description of Device Categories
  - Xeon(V4): 24 cores of Broadwell (BW) 2.50GHz
  - Xeon(V4)+KNC: 24 cores BW + 2 KNC 7120 cards
  - Xeon(V4)+P100: 24 cores BW + 2 P100 cards
  - KNL(HBM): the one described so far

### Power Measurement

- w/ RAPL (Running Ave. Power Limit) API
- Host (CPU+Memory), PCI-E Devices



Results:

**KNL consumes 2x less energy than Xeon(V4)+P100**

# Performance w/ Intel® KNL Processors

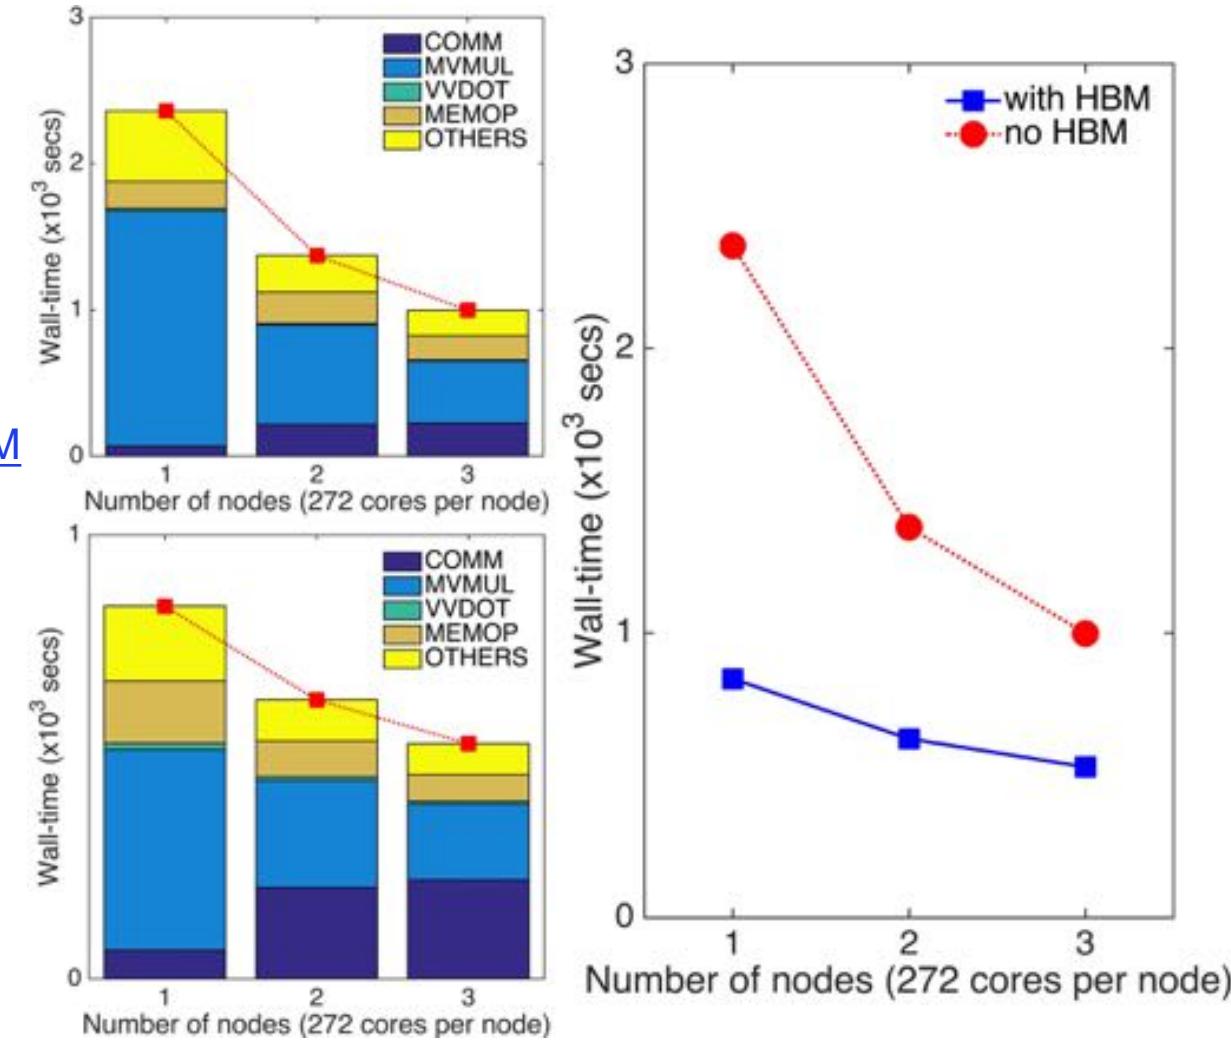
## Multi-node performance: Scalability

### Description of BMT Target and Test Mode

- 5 CB states in  $27 \times 33 \times 33 (\text{nm}^3)$  [100] Si:P quantum dot  
→ 14.4Mx14.4M Hamiltonian Matrix
- Xeon Phi 7250 nodes: 68 cores/node  
→ (4 MPI processes + 68 threads) per node  
→ Quad / Flat mode, No OPA (10G network)  
→ Strong scalability measured up to 3 nodes  
→ Problem size < 16GB; capable with a single node MCDRAM

### Results

- Speed-enhancement with HBM becomes larger as a single node takes larger workload
- Inter-node strong scalability is quite nice with no HBM  
→ **2.36x speed-up with 3x computing expense (no HBM)**  
→ ~78% scalability ( $2.36/3$ ) is what we usually get from multi-core base HPCs (Tachyon-II HPC in KISTI)  
→ 1.58x speed-up with HBM (prob. size is not large enough)



# Performance w/ Intel® KNL Processors

When problem sizes exceed > 16GB?



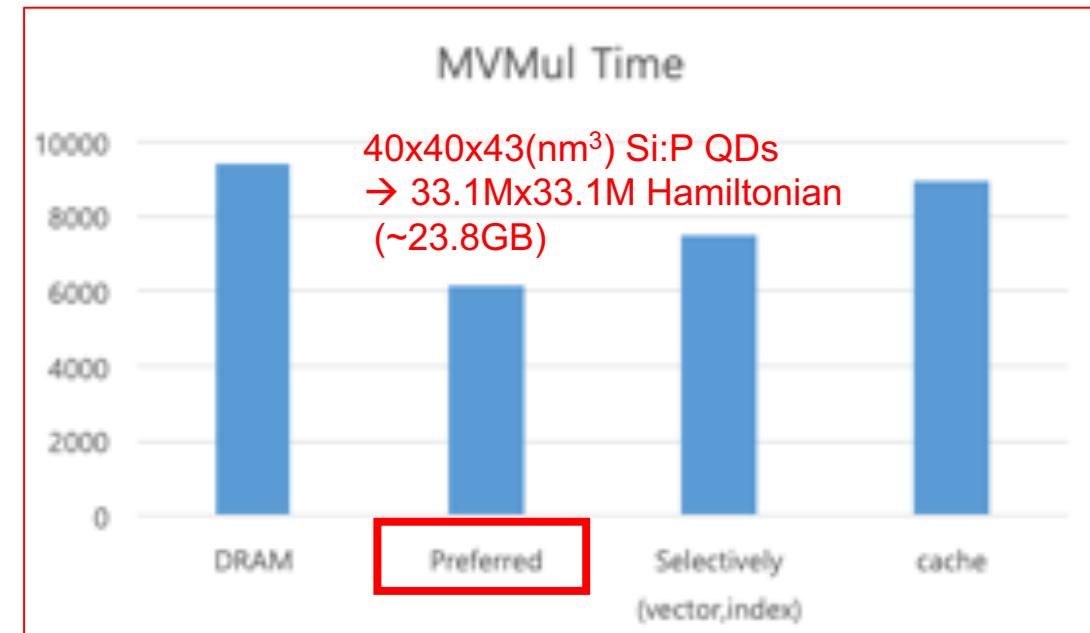
## Matrix-vector multiplier

```
for (unsigned int i = 0; i < nSize; i++) {  
    double real_sum = 0.0;  
    double imaginary_sum = 0.0;  
    const unsigned int nSubStart = pMatrixRow[i];  
    const unsigned int nSubEnd = pMatrixRow[i + 1];  
  
    for (unsigned int j = nSubStart; j < nSubEnd; j++) {  
        const unsigned int nColIndex = pMatrixColumn[i];  
        const double m_real = pMatrixReal[j];  
        const double m_imaginary = pMatrixImaginary[j];  
        const double v_real = pVectorReal[nColIndex];  
        const double v_imaginary = pVectorImaginary[nColIndex];  
  
        real_sum += m_real * v_real - m_imaginary * v_imaginary;  
        imaginary_sum += m_real * v_imaginary + m_imaginary * v_real;  
    }  
  
    pResultReal[i] = real_sum;  
    pResultImaginary[i] = imaginary_sum;  
}
```

1. index  
2. Matrix element  
3. Vector element

## Available options for MCDRAM utilization

- Cache mode: use MCDRAM like L3 cache
- Preferred mode: First fill MCDRAM then go to DRAM  
→ numactl --preferred=1 ...
- Library memkind: use dynamic allocations in code  
→ hbw\_malloc(), hbw\_free()

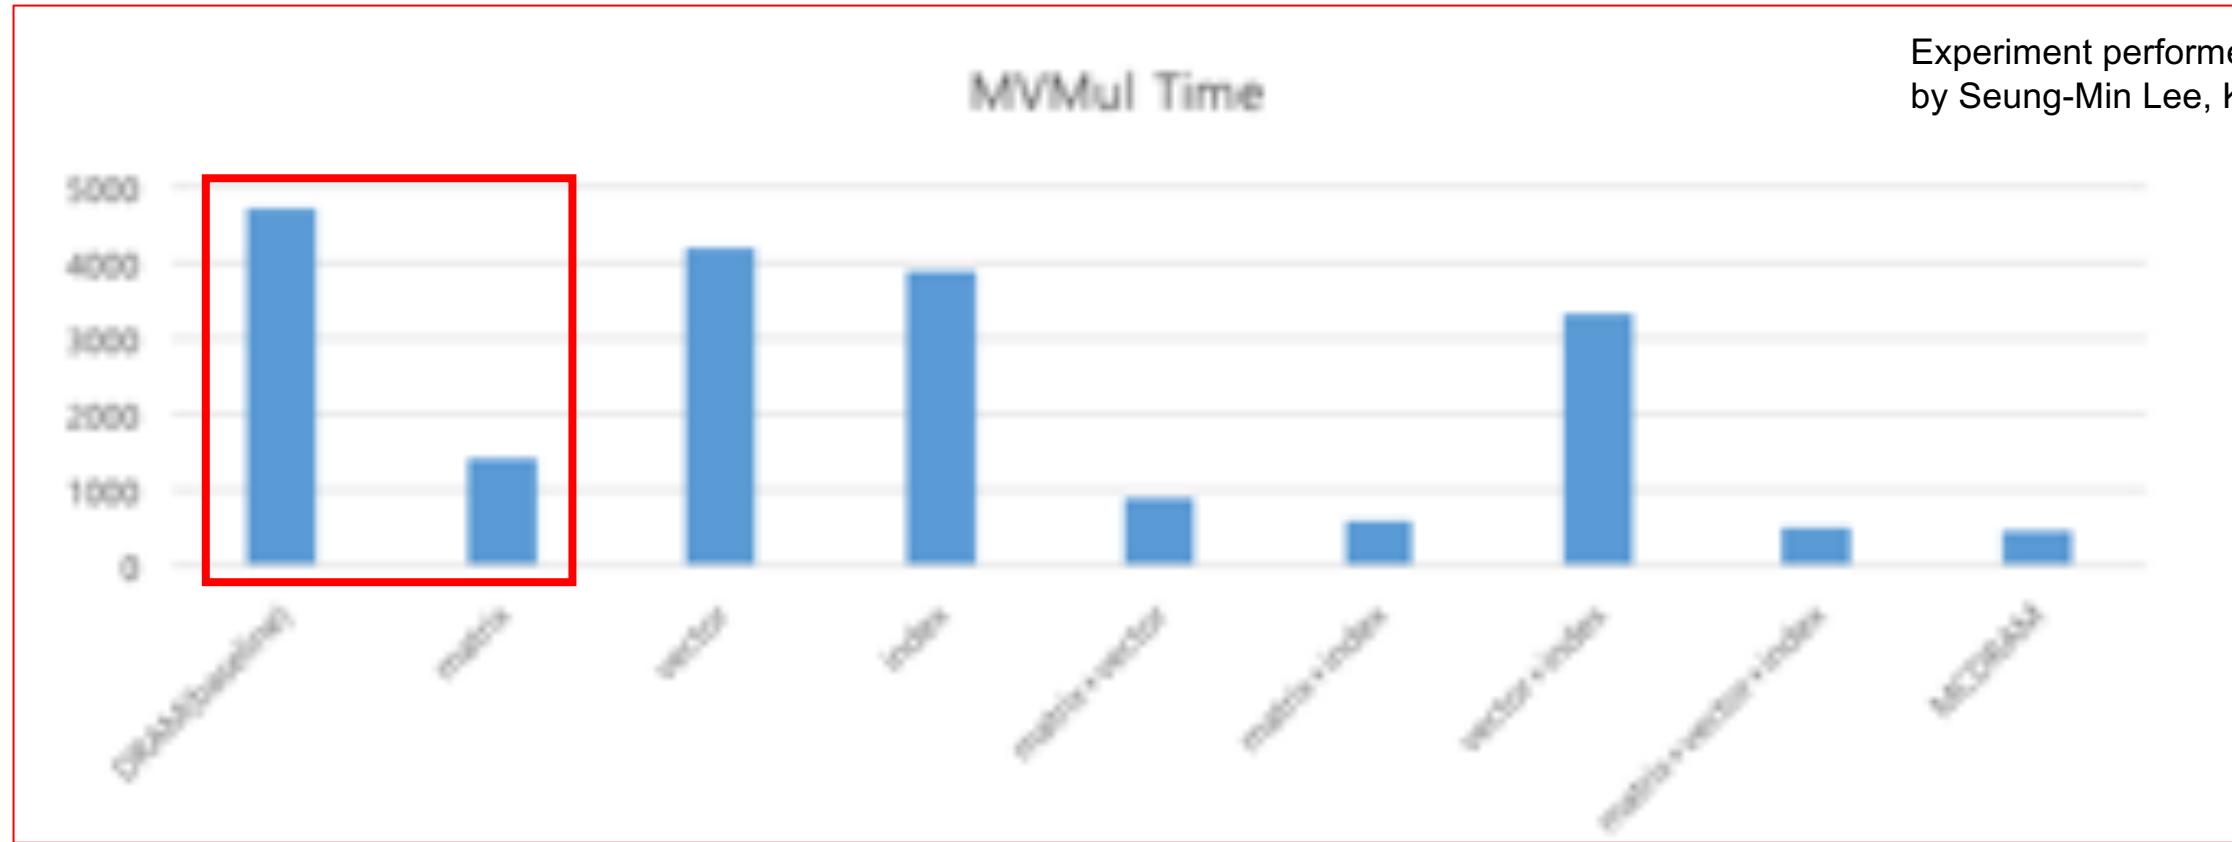


## Data to be saved in memory

- index: column index of matrix nonzero elements (indirect index)
- matrix and vector: matrix nonzero elements and vector elements

# Performance w/ Intel® KNL Processors

# When problem sizes exceed > 16GB?



**Which component would be most affected by the enhanced bandwidth of MCDRAM?**

- MVMul is tested with 11GB Hamiltonian matrix Good for us – matrix is built upon the definition of
  - Matrix nonzero elements drive the most remarkable performance improvement when combined w/ MCDRAM

Good for us – matrix is built upon the definition of geometry!

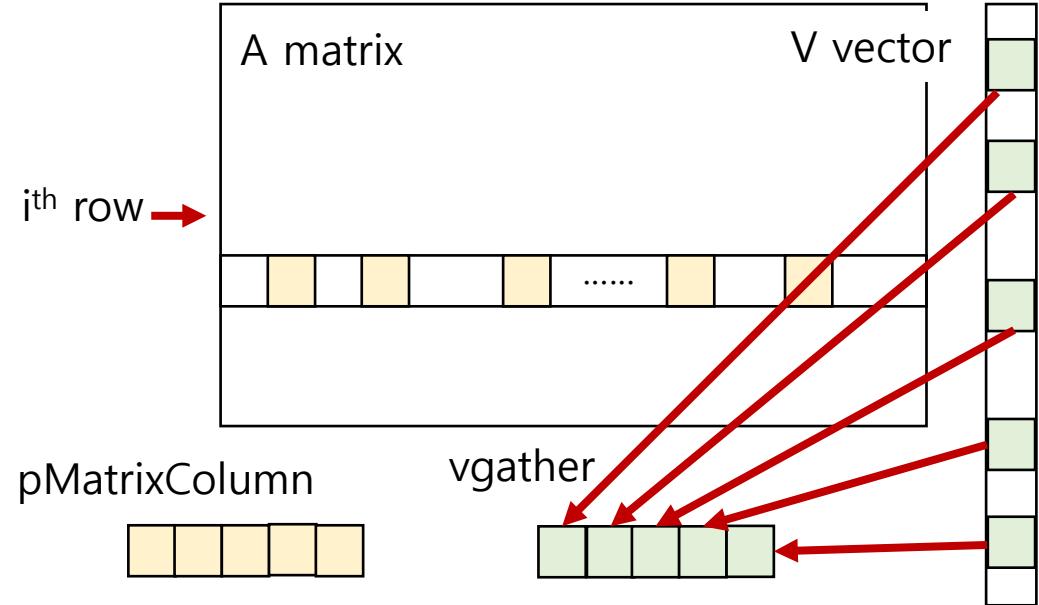
# Performance w/ Intel® KNL Processors

## Vectorization efficiency



### Matrix-vector multiplier: Revisit

```
for (unsigned int i = 0; i < nSize; i++) {  
    double real_sum = 0.0;  
    double imaginary_sum = 0.0;  
    const unsigned int nSubStart = pMatrixRow[i];  
    const unsigned int nSubEnd = pMatrixRow[i + 1];  
  
    for (unsigned int j = nSubStart; j < nSubEnd; j++) {  
        const unsigned int nColIndex = pMatrixColumn[j];  
        const double m_real = pMatrixReal[j];  
        const double m_imaginary = pMatrixImaginary[j];  
        const double v_real = pVectorReal[nColIndex];  
        const double v_imaginary = pVectorImaginary[nColIndex];  
  
        real_sum += m_real * v_real - m_imaginary * v_imaginary;  
        imaginary_sum += m_real * v_imaginary + m_imaginary * v_real;  
    }  
  
    pResultReal[i] = real_sum;  
    pResultImaginary[i] = imaginary_sum;  
}
```



**Matrix nonzeros are stored sequentially**  
→ But “necessary” vector elements are not, and need to be gathered in vector register first!!

- Efficiency of vectorization would not be super excellent  
→ Vector elements should be “gathered” onto register before processing vectorization for matrix-vector multiplier

# Performance w/ Intel® KNL Processors

## Vectorization efficiency



### Matrix-vector multiplier: Revisit

```
for (unsigned int i = 0; i < nSize; i++) {  
    double real_sum = 0.0;  
    double imaginary_sum = 0.0;  
    const unsigned int nSubStart = pMatrixRow[i];  
    const unsigned int nSubEnd = pMatrixRow[i + 1];  
  
    for (unsigned int j = nSubStart; j < nSubEnd; j++) {  
        const unsigned int nColIndex = pMatrixColumn[j];  
        const double m_real = pMatrixReal[j];  
        const double m_imaginary = pMatrixImaginary[j];  
        const double v_real = pVectorReal[nColIndex];  
        const double v_imaginary = pVectorImaginary[nColIndex]; }  
  
        real_sum += m_real * v_real - m_imaginary * v_imaginary;  
        imaginary_sum += m_real * v_imaginary + m_imaginary * v_real;  
    }  
  
    pResultReal[i] = real_sum;  
    pResultImaginary[i] = imaginary_sum;  
}
```

Assembly
block 2: leal (%rcx,%rdi,1), %r15d vpaddyy (%r11,%r15,4), %ymm0, %ymml kxnorw %k0, %k0, %k1 vpxord %zmm9, %zmm9, %zmm9 vpxord %zmm11, %zmm11, %zmm11 kxnorw %k0, %k0, %k2 vmovupsz (%rax,%r15,8), %zmm10 vmovupsz (%r10,%r15,8), %zmm12 add \$0x8, %edi vgatherdpdz (%r9,%ymml,8), %k2, %zmm11 vgatherdpdz (%r14,%ymml,8), %k1, %zmm9 vmulpd %zmm11, %zmm10, %zmm8 vmulpd %zmm10, %zmm9, %zmm13 vfmsub231pd %zmm12, %zmm9, %zmm8 vfmaadd231pd %zmm12, %zmm11, %zmm13 vaddpd %zmm4, %zmm8, %zmm4 vaddpd %zmm2, %zmm13, %zmm2 cmp %r8d, %edi jb 0x417920 <Block 2>

- Efficiency of vectorization would not be super excellent  
→ Vector elements should be “gathered” onto register before processing vectorization for matrix-vector multiplier

# Performance w/ Intel® KNL Processors

## Vectorization efficiency



### Matrix-vector multiplier: Revisit

```
for (unsigned int i = 0; i < nSize; i++) {  
    double real_sum = 0.0;  
    double imaginary_sum = 0.0;  
    const unsigned int nSubStart = pMatrixRow[i];  
    const unsigned int nSubEnd = pMatrixRow[i + 1];  
  
    for (unsigned int j = nSubStart; j < nSubEnd; j++) {  
        const unsigned int nColIndex = pMatrixColumn[j];  
        const double m_real = pMatrixReal[j];  
        const double m_imaginary = pMatrixImaginary[j];  
        const double v_real = pVectorReal[nColIndex];  
        const double v_imaginary = pVectorImaginary[nColIndex]; }  
  
        real_sum += m_real * v_real - m_imaginary * v_imaginary;  
        imaginary_sum += m_real * v_imaginary + m_imaginary * v_real;  
    }  
  
    pResultReal[i] = real_sum;  
    pResultImaginary[i] = imaginary_sum;  
}
```

#### Default (-O3)

Vector length 2  
Normalized vectorization overhead 1.020  
Vector cost : 26.0  
Estimated potential speedup: 1.70

#### AVX2 (-AVX2)

Vector length 2  
Normalized vectorization overhead 1.020  
Vector cost : 24.5  
Estimated potential speedup: 1.490

#### MIC-AVX512 (-xMIC-AVX512)

Vector length 8  
Normalized vectorization overhead 1.104  
Vector cost : 8.370  
Estimated potential speedup: 3.930

- Efficiency of vectorization would not be super excellent  
→ Vector elements should be “gathered” onto register before proc

# Extremely large-scale problems

## In NURION computing resource

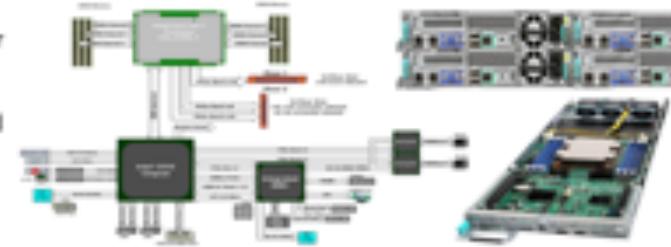


### NURION System Overview

#### Computing nodes

##### Cray 3112-AA000T(2U enclosure), 8,305 KNL Computing modules

- > 1x Intel Xeon Phi KNL 7250 processor
- > 96GB (6x 16GB) DDR4-2400 RAM
- > 1x Single-port 100Gbps OPA HFI card
- > 1x On-board GigE (RJ45) port



#### CPU-only nodes

##### Cray 3111-BA000T(2U enclosure), 132 Skylake Computing modules

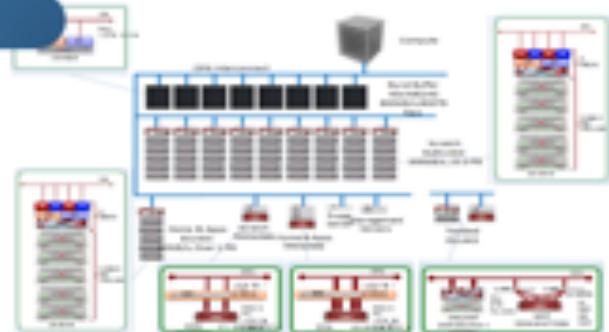
- > 2x Intel Xeon SKL 6148 processors
- > 192GB (12x 16GB) DDR4-2666 RAM
- > 1x Single-port 100Gbps OPA HFI card
- > 1x On-board GigE (RJ45) port



- 132 SKL (Xeon 6148) nodes / 8,305 KNL (Xeon Phi 7250) nodes
- Ranked at 13<sup>th</sup> in Top500.org as of 2018. Nov.  
→ Rpeak 25.7pFLOPS, Rmax, 13.9pFLOPS. <https://www.top500.org/system/179421>

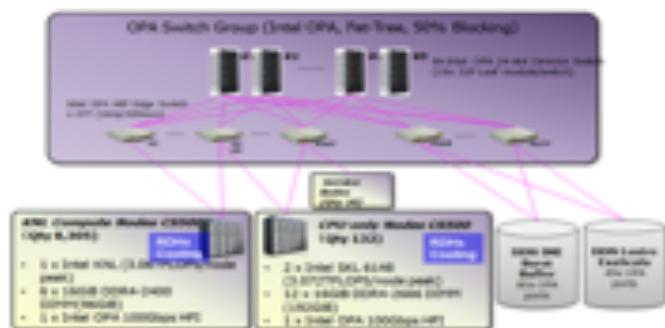
#### Storage 20PB SFS@300GB/s, 10PB Archiving

- > Global scratch: 20PB, 0.3TB/s (DDN ES14KX 9ea, 360 x 8TB disk each)
- > Home and application directory 1PB
- > NVMe Burst Buffer: 0.8PB, 0.8TB/s (IME240 40ea 19 NVMe SSD each)
- > Cray TSMSC and IBM TS4500



#### Interconnect OPA(Omni-Path Architecture), Fat-Tree, 50% Blocking

- > Intel OPA High-speed interconnect switch  
274x 48-port OPA edge switches  
8x 768-port OPA core switches
- > Bandwidth: 12.3 Gb/sec
- > Bisectional Bandwidth : 27 TB/sec
- > 10<sup>-14</sup> BER(Bit Error Rate), Adaptive routing



# Extremely large-scale problems

## In NURION computing resource



### Description of BMT Target

- Computed lowest 3 conduction sub-bands in 2715x54x54 (nm<sup>3</sup>) [100] Si:P square nanowire
  - contains 400 million (0.4 billion) atoms,
  - Hamiltonian matrix DOF = 4 billion x 4 billion

### Computing Environment

- Intel® Xeon Phi 7250 (NURION)
  - 1.4GHz/68 cores, 96GB DRAM, 16GB MCDRAM (/node)
  - OPA (100GB)

### Other Information for Code Compile and Runs

- Intel® Parallel Studio 17.0.5
- Instruction set for vectorization: MIC-AVX512
- MCDRAM allocation: numactl –preferred=1
- OPA fabric. 4 MPI processes / 17 threads per node
- Memory Placement Policy Control

→ export I\_MPI\_HBW\_POLICY = hbw\_bind,hbw\_preferred,hbw\_bind

(HBW memory for RMA operations and for Intel® MPI Library first. If HBW memory is not available, use local DDR)

RMA: Remote Memory Access (for MPI communications)

```
module purge
module load craype-network-opa
module load intel/17.0.5
module load impi/17.0.5

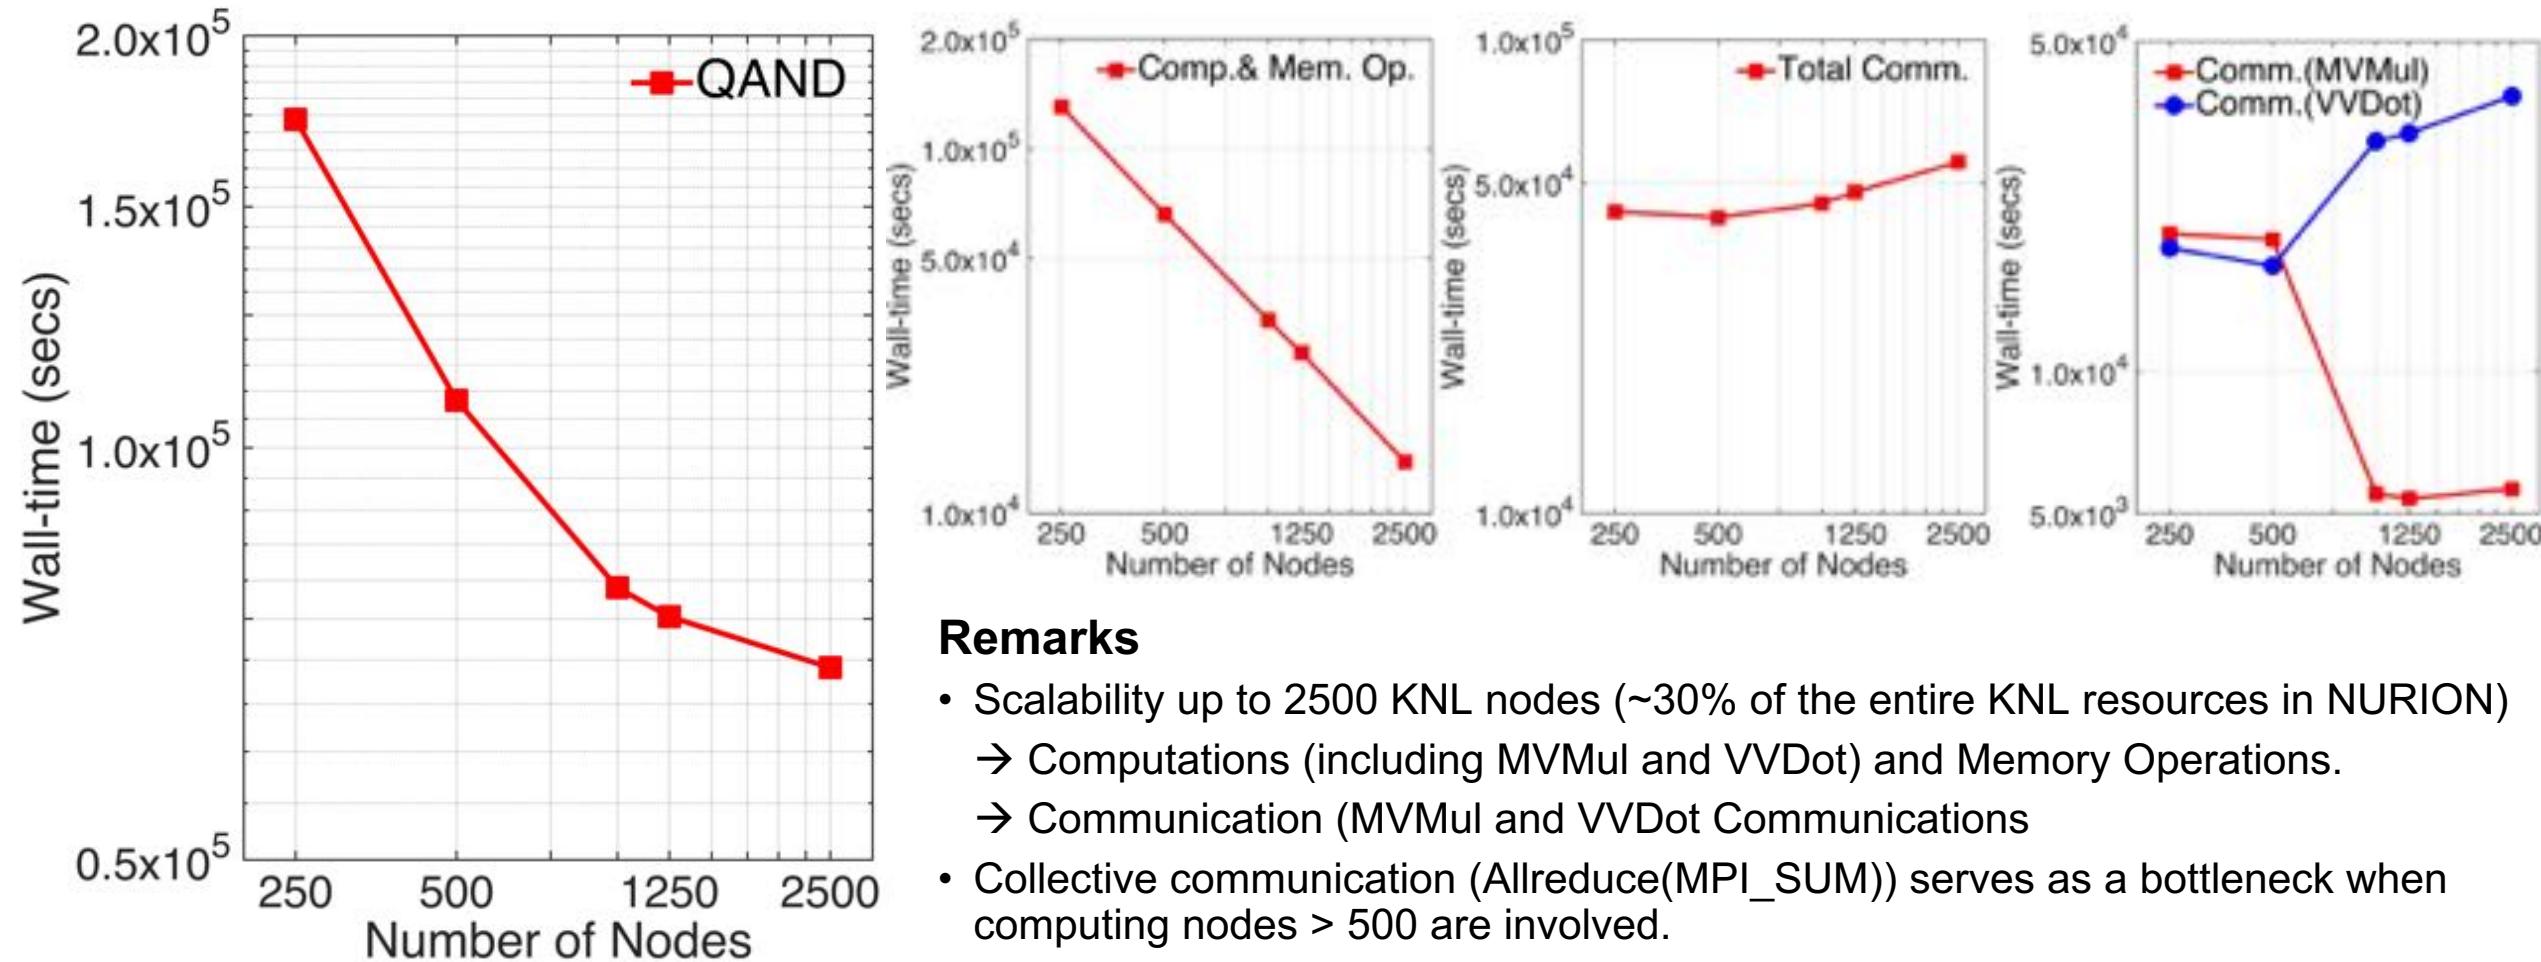
export NUMACTL="numactl --preferred=1"
export I_MPI_HBW_POLICY=hbw_bind,hbw_preferred,hbw_bind
export I_MPI_FABRICS=ofi

ulimit -s unlimited
cd $PBS_O_WORKDIR
cat $PBS_NODEFILE
time mpirun $NUMACTL ...
```

### Snapshot of a PBS script

# Extremely large-scale problems

## In NURION computing resource



### Remarks

- Scalability up to 2500 KNL nodes (~30% of the entire KNL resources in NURION)
  - Computations (including MVMul and VVDot) and Memory Operations.
  - Communication (MVMul and VVDot Communications)
- Collective communication (Allreduce(MPI\_SUM)) serves as a bottleneck when computing nodes > 500 are involved.

# Summary

## KISTI Intel® Parallel Computing Center



- Introduction to Code Functionality
- Main Numerical Problems and Strategy of Development
- Performance (speed and energy consumption) in a single KNL node
  - Benefits against the case of CPU + 2xP100 GPU devices
- Performance in extremely huge computing environment
  - Strong scalability up to 2,500 KNL nodes in NURION system
- (Appendix) Strategy of Performance Improvement towards PCI-E devices
- (Appendix) List of Related Publications

**Thanks for your attention!!**

# Appendix: Strategy for offload-computing

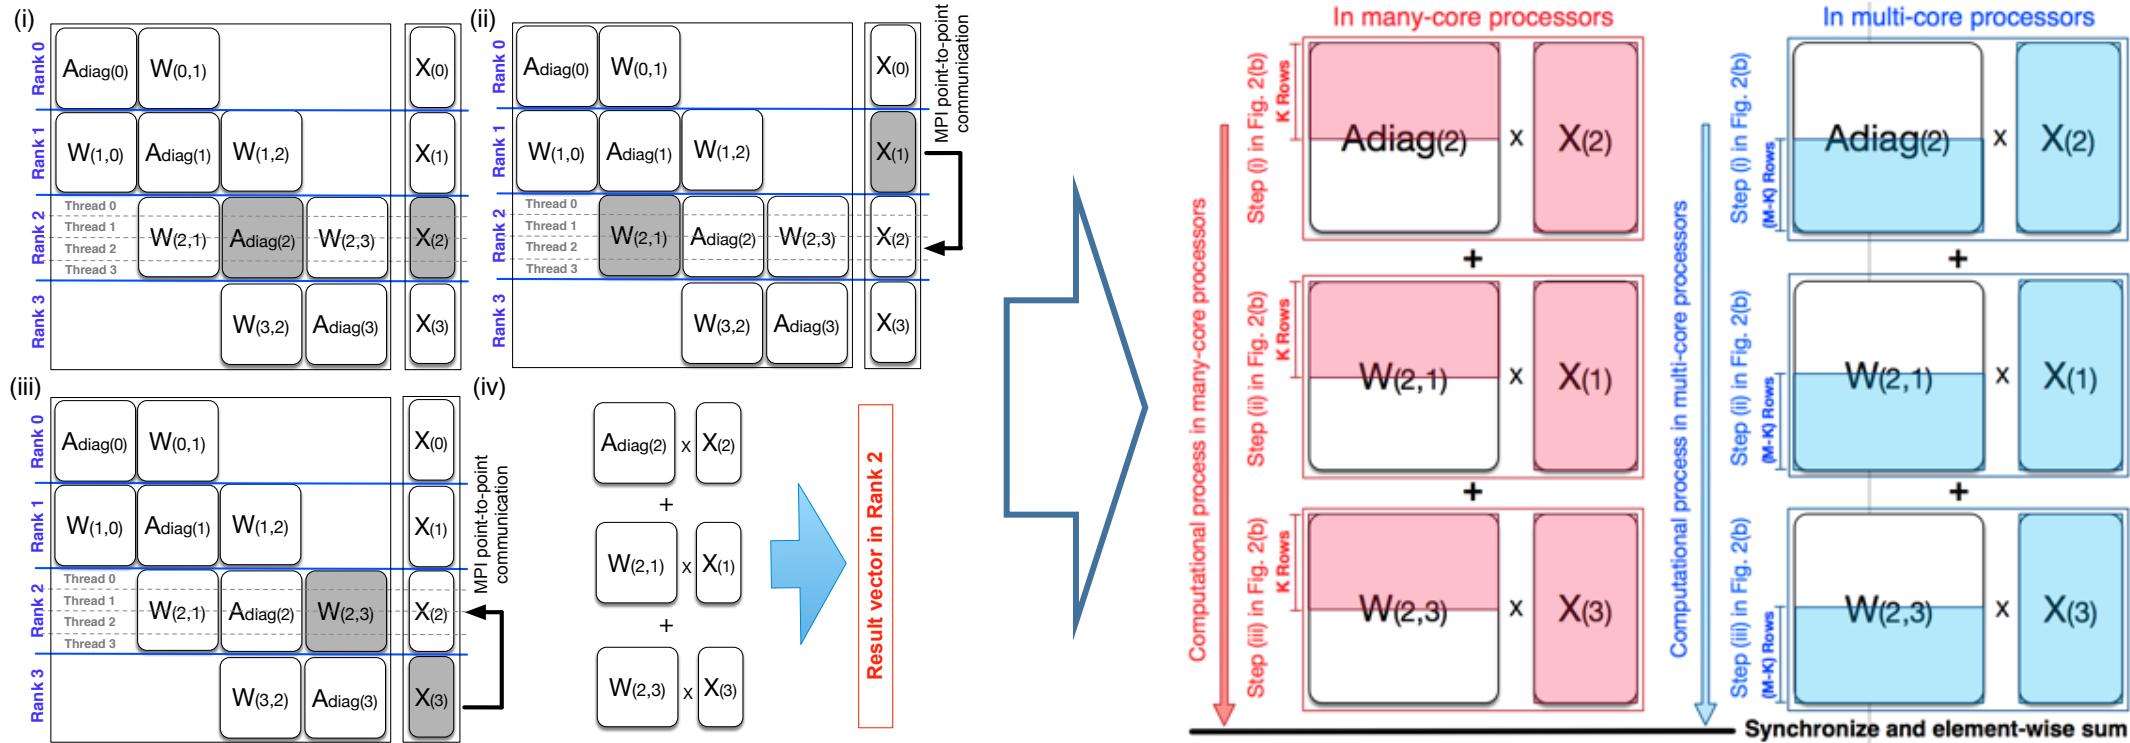
## Asynchronous Offload (for Xeon(V4) + KNC, Xeon(V4) + GPU)



### The real bottleneck of computing: Overcome with asynchronous offload

H. Ryu et al., Comp. Phys. Commun. (2016)  
(<http://dx.doi.org/10.1016/j.cpc.2016.08.015>)

- Vector dot-product is not expensive: All-reduce, but small communication loads
- Vector communication is not a big deal: only communicates between adjacent layers
- Sparse-matrix-vector multiplication is a big deal: Host and PCI-E device shares computing load



# Appendix: Strategy for offload-computing

## Data-transfer and Kernel Functions for GPU Computing

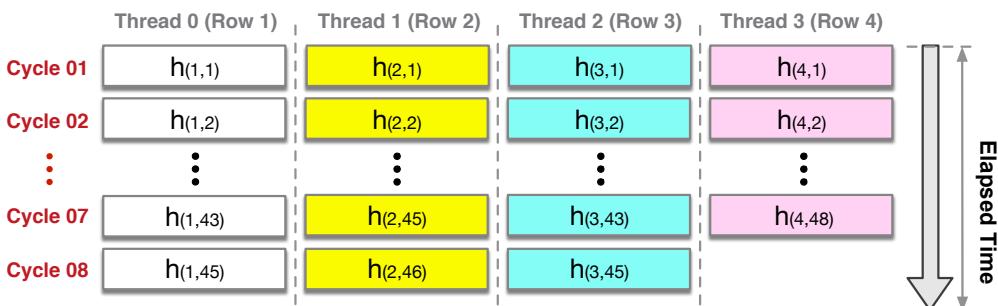
### Data-transfer between host and GPU Devices

- 3x increased bandwidth with **pinned memory**
- Overlap of computation and data-transfer with **asynchronous streams**

### Speed-up of GPU Kernel Function (MVMul)

- Treating several rows at one time with WARPs

Data-Access with a thread-base (no WARPs)



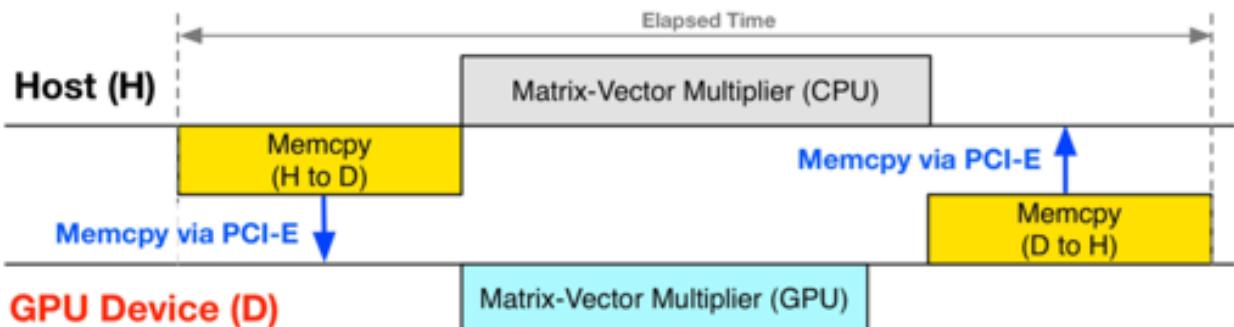
Data-Access with a WARP-base



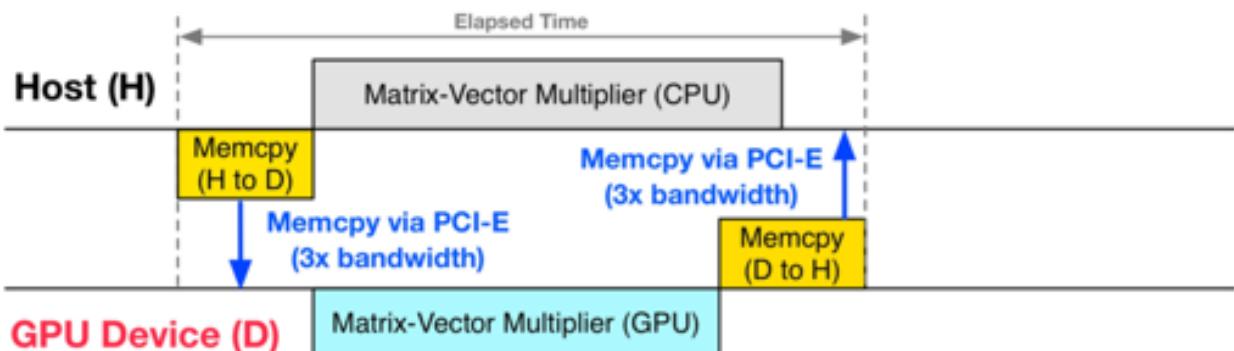
H. Ryu et al., J. Comp. Elec. (2018)

(<http://dx.doi.org/10.1007/s10825-018-1138-4>)

#### [Synchronous Data Transfer with Pageable Memory]



#### [Asynchronous Data Transfer with Pinned Memory]



# List of Related Publications

## Journals and Conference Proceedings



### Journal Articles / Book Chapters

- [1] H. Ryu, O. Kwon, Journal of Computational Electronics (2018) <https://doi.org/10.1007/s10825-018-1138-4>  
→ "Fast, Energy-efficient Electronic Structure Simulations for Multi-million Atomic Systems with GPU Devices",
- [2] S. Choi, W. Kim, M. Yeam, H. Ryu, International Journal of Quantum Chemistry (2018) <https://doi.org/10.1002/qua.25622>  
→ "On the achievement of high fidelity and scalability for large-scale diagonalizations in grid-based DFT simulations"
- [3] O. Kwon, H. Ryu, A Book Chapter in "High Performance Parallel Computing", InTechOpen (2018) <https://doi.org/10.5772/intechopen.80997>  
→ "Acceleration of Large-scale Electronic Structure Simulations with Heterogeneous Parallel Computing"
- [4] H. Ryu, Y. Jeong, J. Kang, K. Cho, Computer Physics Communications (2016) <https://doi.org/10.1016/j.cpc.2016.08.015>  
→ "Time-efficient simulations of tight-binding electronic structures with Intel Xeon Phi™ many-core processors"

### Conference Proceedings / Presentations

- [1] H. Ryu, K.-H. Hong (2018), Proceedings of IEEE SISPAD, <https://doi.org/10.1109/SISPAD.2018.8551719>  
→ "Optical Properties of Organic Perovskite Materials for Finite Nanostructures"
- [2] J. Kang, O. Kwon, J. Jeong, K. Lim, H. Ryu (2018), Proceedings of HPCS, <https://doi.org/10.1109/HPCS.2018.00063>  
→ "Performance Evaluation of Scientific Applications on Intel Xeon Phi Knights Landing Clusters"
- [3] H. Ryu, O. Kwon (2017), Top 20 Research Posters in GPU Technology Conference,  
<http://www.gputechconf.com/resources/poster-gallery/2017/computational-physics>  
→ "Q-AND: Fast, Energy-efficient Computing of Electronic Structures for Multi-million Atomic Structures with GPGPU Devices"
- [4] H. Ryu, Y. Jeong (2016), Proceedings of IEEE CLUSTER, <https://doi.org/10.1109/CLUSTER.2016.76>  
→ "Enhancing Performance of Large-scale Electronic Structure Calculations with Many-core Computing"