

P4343—

OPENUH: AN OPEN SOURCE OPENACC COMPILER

Xiaonan (Daniel) Tian, Rengan Xu and Barbara
Chapman

HPCTools Group

Computer Science Department

University of **Houston**

GTC2014, San Jose, CA; 03/26 /2014

**O
U
T
L
I
N
E**

I. Motivation

II. Introduction to OpenUH

III. Loop Scheduling

IV. Data Movement

V. Performance

VI. Future and Conclusion

I. Motivation

Motivation


WHY do we implement OpenACC support in OpenUH?

- ✓ Performance gap between OpenACC and CUDA → more research on OpenACC compiler optimization
- ✓ Open Source OpenACC compiler is required for research purposes.

WHY is this talk important?

- ✓ BETTER understand OpenACC implementation, BETTER knowledge on application optimization.

II. Introduction to OpenUH



Website: <http://web.cs.uh.edu/~openuh/>

Source: <https://github.com/pumpkin83/OpenUH-OpenACC>

Email: openuh@cs.uh.edu

Introduction to OpenUH

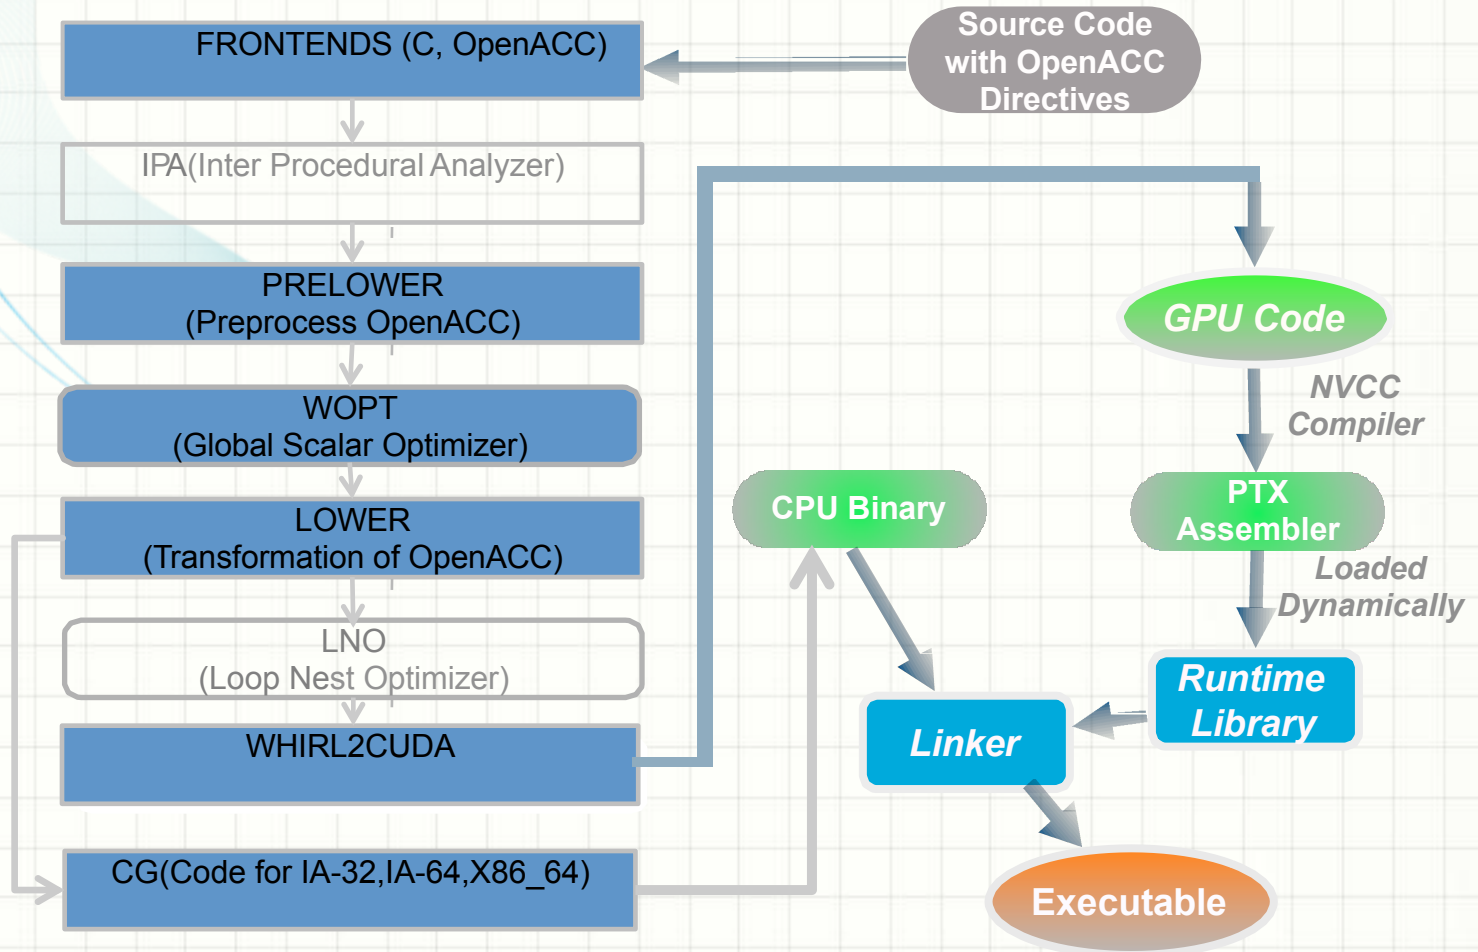
- ***Open Source Research Compiler***
 - ✓ Open64 based
 - ✓ Support C/C++/Fortran/Coarray
- ***Parallel Programming model***
 - ✓ OpenMP
 - ✓ OpenACC
 - ✓ COARRAY

Introduction to OpenUH OpenACC

- ***OpenACC 1.0 implementation***
 - ✓ **Directives:** Parallel, kernels, Data, Loop, Wait
 - ✓ **Data Clause:** copy/copyin/copyout/create/update
 - ✓ **Loop Scheduling Clauses:** gang/worker/vector
 - ✓ **Async clause:** async/wait
 - ✓ **Unsupported:** host_data/declare/cache

Introduction to OpenUH OpenACC

OpenUH OpenACC Compiler Infrastructure





III. LOOP SCHEDULING

Loop Scheduling

1

- What's Loop Scheduling?

2

- Parallel Loop Scheduling

3


- Kernels Loop Scheduling

Loop Scheduling

- What is Loop Scheduling?
 - Solutions to distribute sequential loop iterations across a large number of threads
- Why we have two different Loop Scheduling strategies?
 - Explore multi-dimensional topology of NVIDIA GPGPU architecture

Loop Scheduling


- `#pragma acc loop gang(4)`
- `For(i=0; i<11; i++){...}`



Gangs

Loop Scheduling

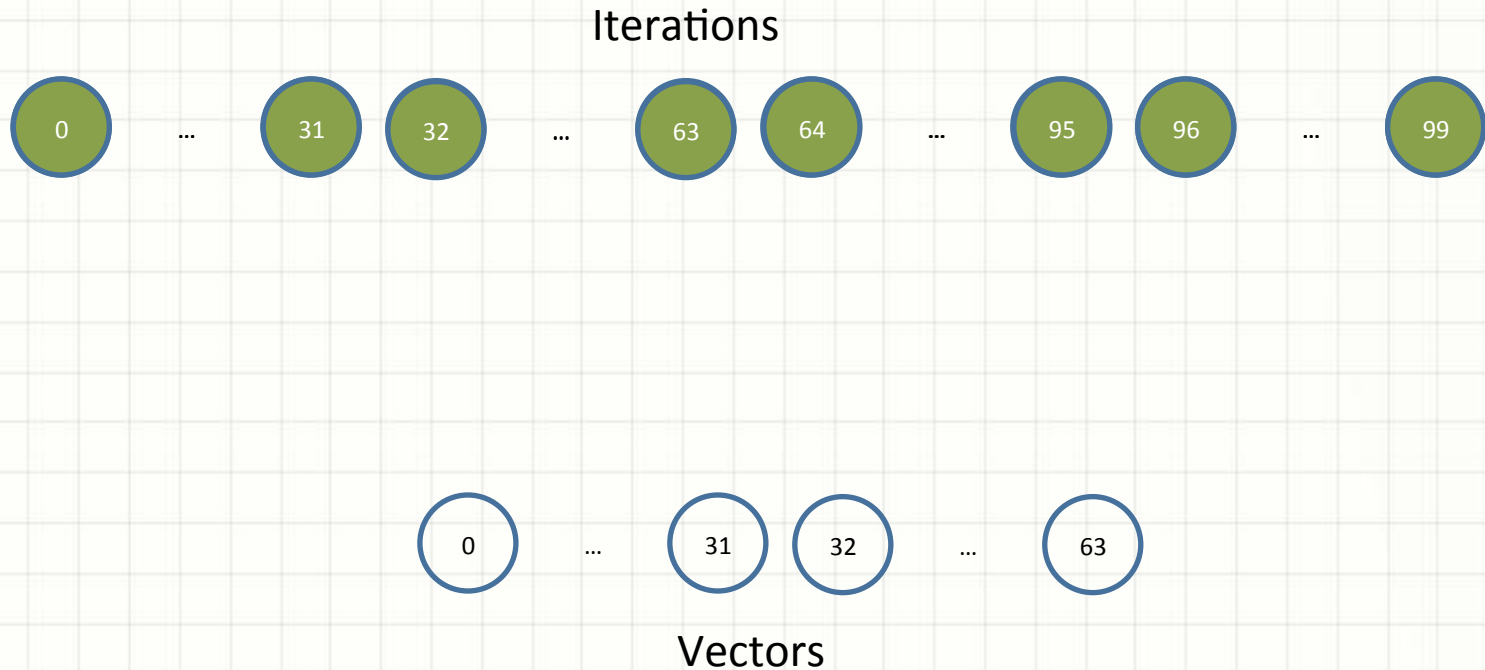
- `#pragma acc loop gang(4)`
- `For(i=0; i<11; i++){...}`



Gangs


Loop Scheduling

- `#pragma acc loop vector(64)`
- `For(i=0; i<99; i++){...}`



Loop Scheduling

- `#pragma acc loop gang(3) vector(32)`
- `For(i=0; i<130; i++){...}`



Parallel Loop Scheduling

- λ **Gang** \rightarrow (CUDA) thread-block
- λ **Worker** \rightarrow (CUDA) y dimensional threads in a thread block
- λ **Vector** \rightarrow (CUDA) x dimensional threads in a thread block
 - 1D Grid, and 1D/2D thread-block.
 - # of Worker * # of Vector \leq 1024
 - Requires minimal lower-level knowledge.
 - Follows OpenACC 2.0: gang contains worker and vector; worker can only include vector.

Parallel Loop Scheduling

λ 1. Single Loop

- #pragma acc loop **gang worker vector**
- for(...){}

λ 2. Two-level Nested Loop

λ 2.1. loop **gang** / loop **worker vector**

- #pragma acc loop **gang**
- for(...){
- #pragma acc loop **worker vector**
- for(...){
- }
- }

λ 2.2. loop **gang worker** / loop **vector**

- #pragma acc loop **gang worker**
- for(...){
- #pragma acc loop **vector**
- for(...){
- }
- }


2.3. loop **gang** / loop **vector**

Parallel Loop Scheduling:example

- `#pragma acc loop gang(2) worker(4) vector(64)`
- `For(i=istart; i<iend; i++){...}`

Parallel Loop Scheduling:example

- #pragma acc loop gang(2) worker(4) vector(64)
- For(i=istart; i<iend; i++){...}



Block 0


CUDA
Architecture



Block 1

Parallel Loop Scheduling

- `#pragma acc loop gang(2) worker(4) vector(64)`
- `For(i=istart; i<iend; i++){...}`



Block 0

CUDA
Architecture



Block 1

Parallel Loop Scheduling

- #pragma acc loop gang(2)
for(i=istart; i<iend; i++){
 #pragma acc loop worker(4) vector(64)
 for(j=jstart; j<jend; j++){...}
}

Parallel Loop Scheduling

- #pragma acc loop gang(2)

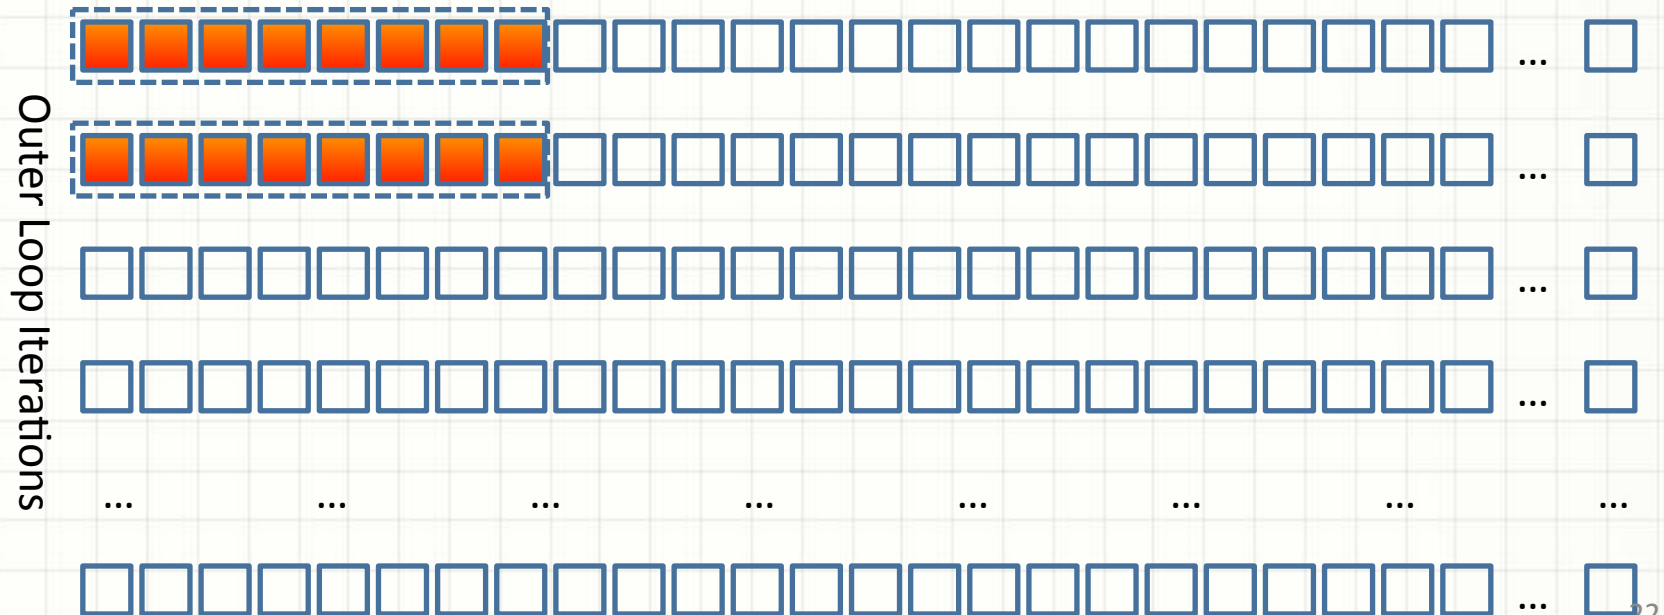
```
for(i=istart; i<iend; i++){
```

```
    #pragma acc loop worker(4) vector(64)
```

```
    for(j=jstart; j<jend; j++){...}
```

```
}
```

Inner Loop Iterations



Parallel Loop Scheduling

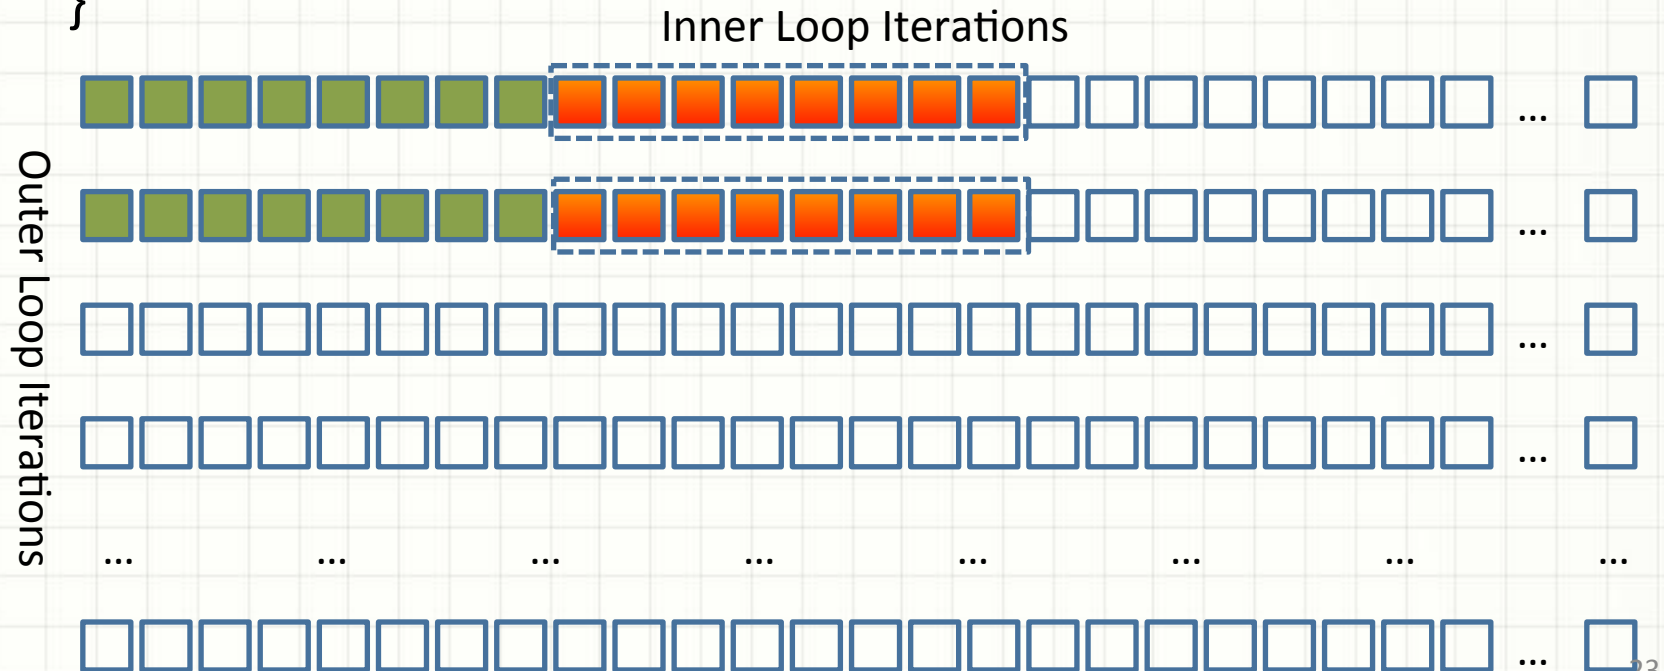
- #pragma acc loop gang(2)

```
for(i=istart; i<iend; i++){
```

```
    #pragma acc loop worker(4) vector(64)
```

```
    for(j=jstart; j<jend; j++){...}
```

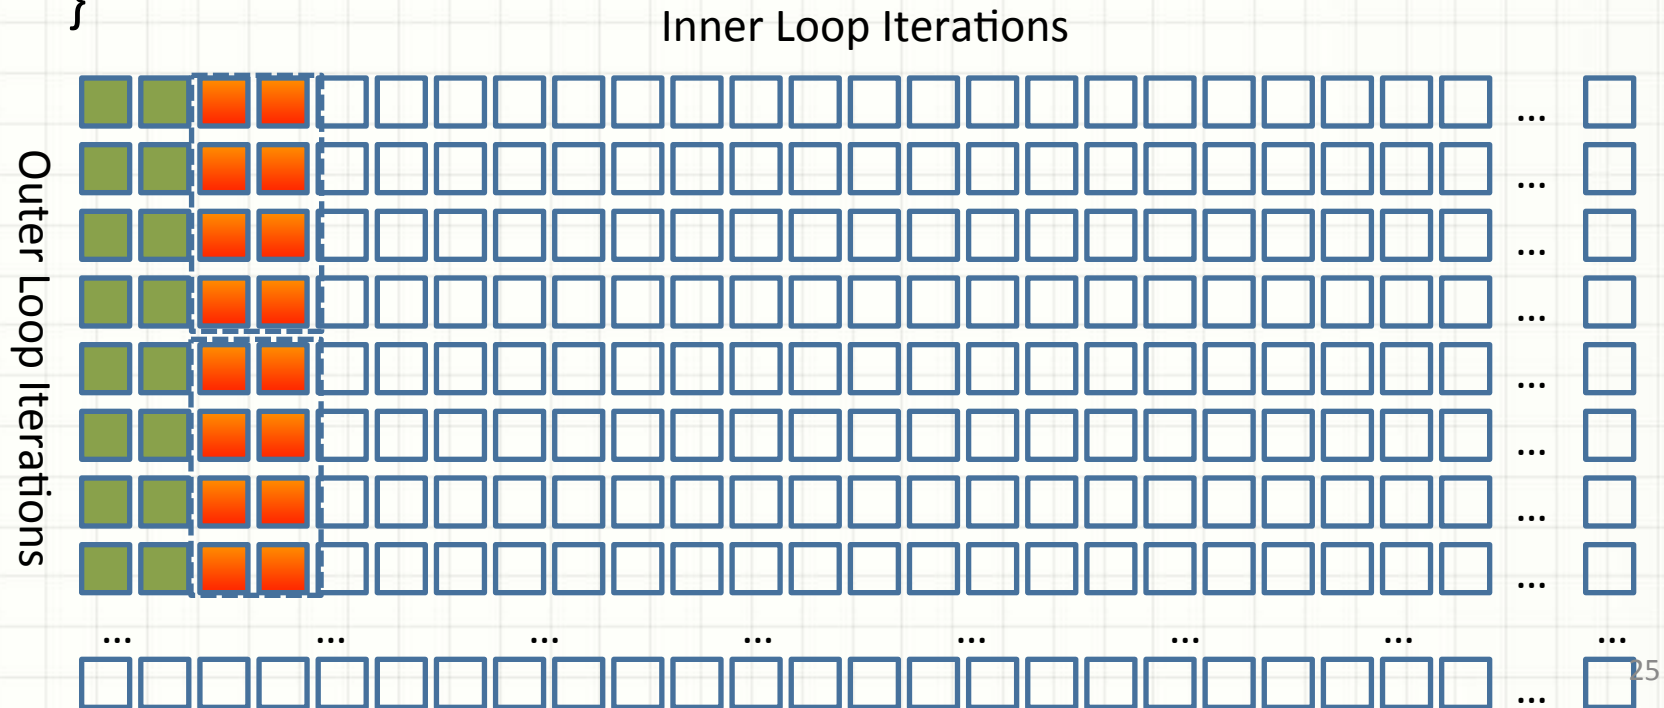
```
}
```



Parallel Loop Scheduling

```

• #pragma acc loop gang(2) worker(4)
  for(i=istart; i<iend; i++){
    #pragma acc loop vector(64)
    for(j=jstart; j<jend; j++){...}
  }
    
```



Parallel Loop Scheduling

λ 3. Three level Nested Loop

loop gang/loop worker/ loop vector

```
#pragma acc loop gang
```

```
for(...)
```

```
    #pragma acc loop worker
```

```
    for(...)
```

```
        #pragma acc loop vector
```

```
        for(...)
```

```
        {
```

```
        }
```

Why do we need different strategies for implementing loop scheduling?

```
#pragma acc loop gang(19)
```

```
for(i=0; i<19; i++)
```

```
    #pragma acc loop worker(32)
```

```
    for(j=0; j<1000000; j++)
```

```
        #pragma acc loop vector(32)
```

```
        For(k=0; k<100000; k++)
```

```
        {
```

```
        }
```

What is the maximum threads we have?

$19 * 32 * 32 = 19K$

Try this loop scheduling

- λ According the scheduling in the code, 2D grid and 2D thread-block in NVIDIA GPGPU are created.

```
#pragma acc loop gang(19)
for(i=0; i<19; i++)
    #pragma acc loop gang(32) vector(32)
    for(j=0; j<1000000; j++)
        #pragma acc loop vector(32)
        for(k=0; k<100000; k++)
        {
        }
```

-What is the maximum threads we have here?

- $19 * 32 * 32 * 32 = 32 * 19K$

Kernels Loop Scheduling

- λ **Gang** \rightarrow (CUDA) thread-block, can be in x, y, z dimension
- λ **Worker** \rightarrow Ignored
- λ **Vector** \rightarrow (CUDA) thread, can be in x, y, z dimension
 - \triangleright Multi-dimensional grid/thread-block, both of them can be extended into 3 dimensional topology.
 - \triangleright Fine tuning: provide more scheduling options for users.
 - \triangleright Users need to have more knowledge about compiler and hardware information(currently, no autotuning)
 - \triangleright Provided more choices to loop scheduling.
 - \triangleright In some cases, it does help improve performance

Kernels Loop Scheduling

λ 1. Single Loop

```
#pragma acc loop gang vector  
for(...){}
```

λ 2. Double Nested Loop

λ 2.1. *loop gang / loop vector*

λ 2.2. *loop gang vector / loop vector*

λ 2.3. *loop gang / loop gang vector*

λ 2.4. *loop gang vector / loop gang vector*

Kernels Loop Scheduling

3. Triple Nested Loop

3.1 loop gang / loop gang vector / loop vector

3.2 loop vector / loop gang vector / loop gang

3.3 loop gang vector / loop gang vector / loop vector

3.3 loop gang vector / loop gang vector / loop gang
vector

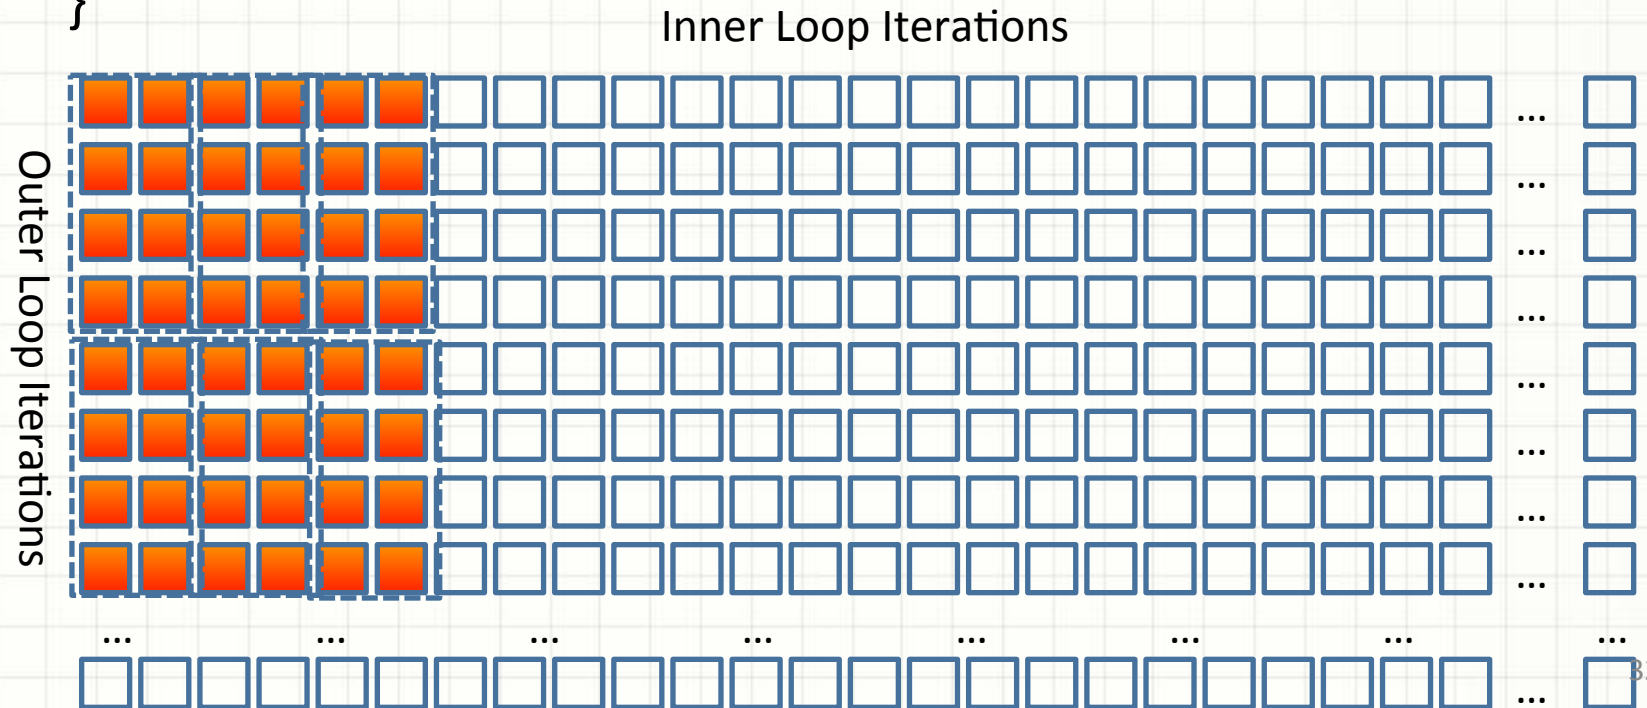
λ ...

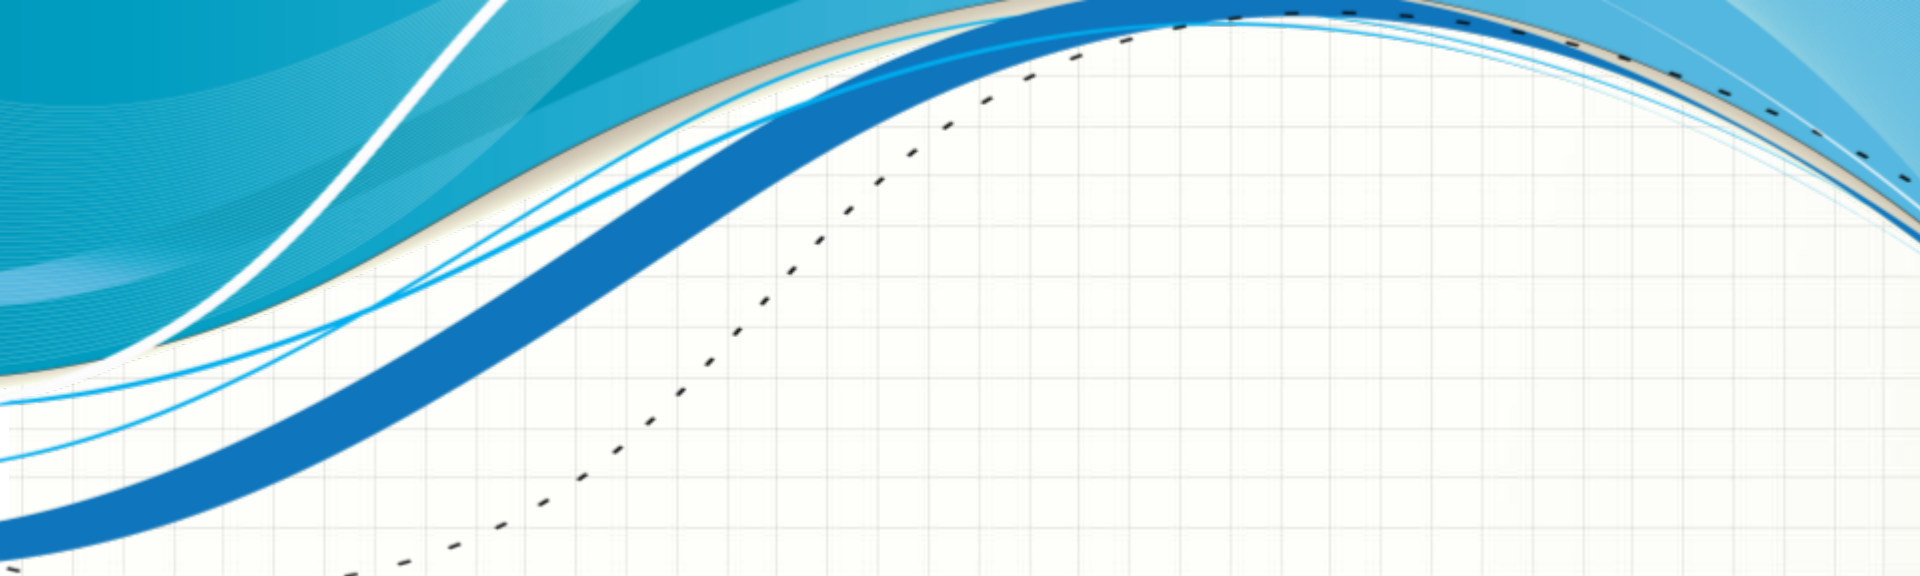
Kernels Loop Scheduling: Example

```

• #pragma acc loop gang(2) vector(4)
  for(i=istart; i<iend; i++){
    #pragma acc loop gang(3) vector(64)
    for(j=jstart; j<jend; j++){...}
  }

```





IV. DATA MOVEMENT

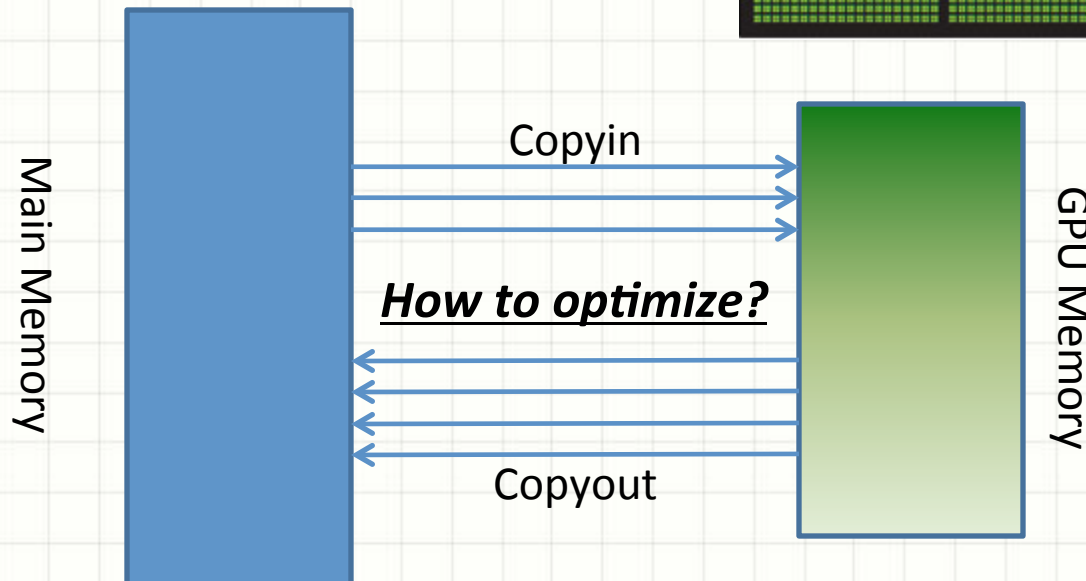

Data Movement

λ 1. Data transfer between CPU and GPU

Multi-core CPU



GPU Thousands of Cores



Data Movement

λ 2. Basic Implementation

λ *copy* → *pcopy*;

λ *copyin* → *pcopyin*

λ *copyout* → *pcopyout*

λ *create* → *pcreate*

- *Free buffer/variables when you exit the current region*

Goal: Avoid duplicate data traffic(malloc, copyin, copyout)

Data Movement

λ 2. Basic Implementation

λ #pragma acc data
data_clauses

λ {

λ #pragma acc data
data_clauses

λ {

λ #pragma acc kernels
data_clauses

λ {

λ ...

λ }

λ }

λ }



Data Movement

λ 2. Basic Implementation

λ #pragma acc data
data_clauses

λ {
λ #pragma acc data
data_clauses

{
λ #pragma acc kernels
data_clauses


λ {

λ ...

λ }

λ }

λ }



Data Movement

```

λ 3. Partial Array
λ #pragma acc data
  create(xx[0:N])
λ {
  λ Foo(&xx[start])
λ }
λ ...

```

Memory Mapping Table

CPU	GPU

Foo(double *x)

```

λ #pragma acc parallel
  {
    Foo(&xx[start])
  }

```

CPU Memory



GPU Memory



Data Movement

```

λ 3. Partial Array
λ #pragma acc data
  create(xx[0:N])
λ {
λ   Foo(&xx[start])
λ }
λ ...
    
```

```

Foo(double *x)
    
```

```

#pragma omp parallel
{
    Foo(&xx[start]);
}
    
```

```


Foo(double *x)
    
```

```

#pragma omp parallel
{
    Foo(&xx[start]);
}
    
```

Memory Mapping Table

CPU	GPU
xx	xx'



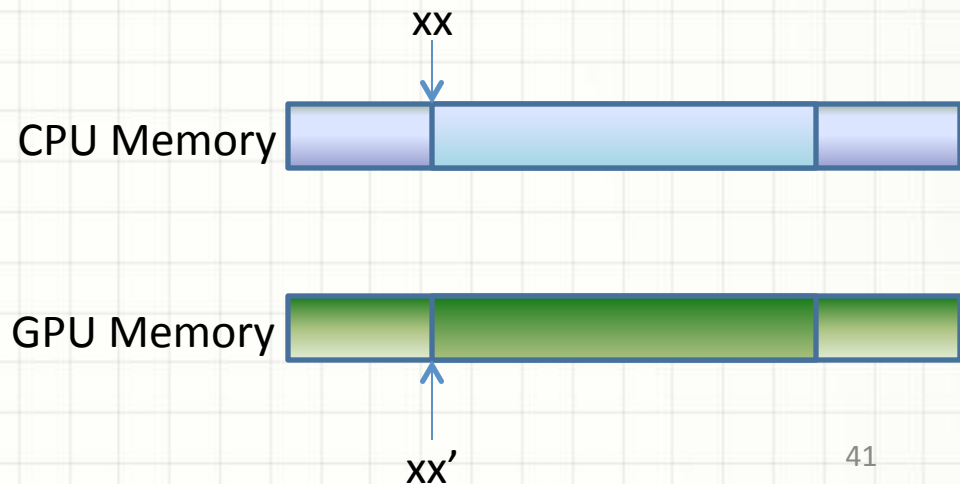
Data Movement

```

λ 3. Partial Array
λ #pragma acc data
  create(xx[0:N])
λ {
  λ Foo(&xx[start])
λ }
λ ...
λ Foo(double* x)
λ {
  λ #pragma acc parallel
    pcopy(x[n1:n2])
  λ {
  λ ...
  λ }
λ }
    
```

Memory Mapping Table

CPU	GPU
xx	xx'



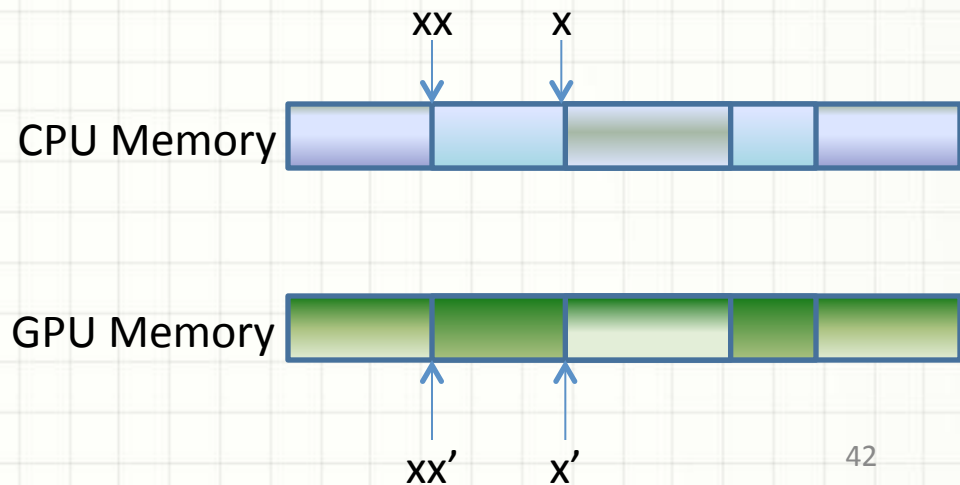
Data Movement

```

λ 3. Partial Array
λ #pragma acc data
  create(xx[0:N])
λ {
  λ Foo(&xx[start])
λ }
λ ...
λ Foo(double* x)
λ {
  λ #pragma acc parallel
    pcopy(x[n1:n2])
  λ {
  λ ...
  λ }
λ }
    
```

Memory Mapping Table

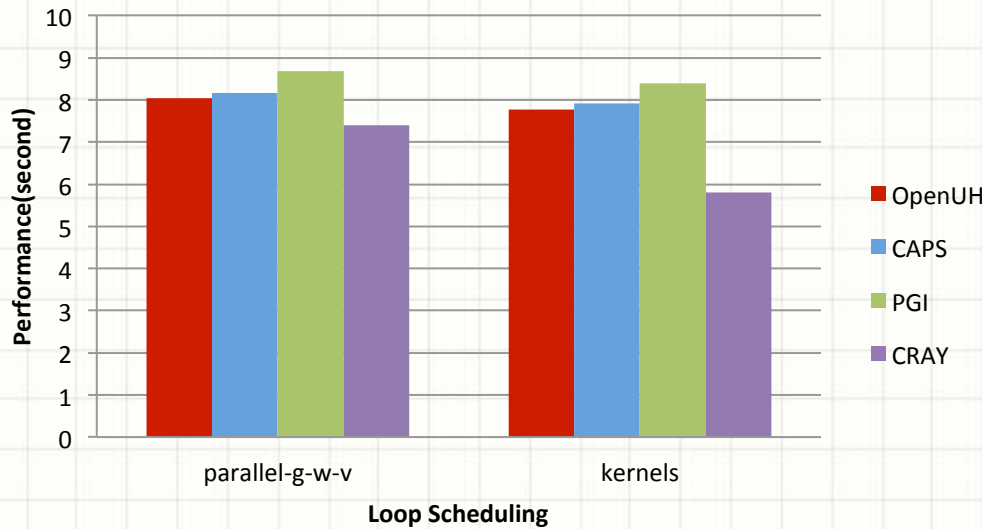
CPU	GPU
xx	xx'
x	x'



VI. Performance

Three-Level Nested Loop Scheduling

Wave13pt



Kernels

OpenUH: g-gv-v scheduling

PGI: default

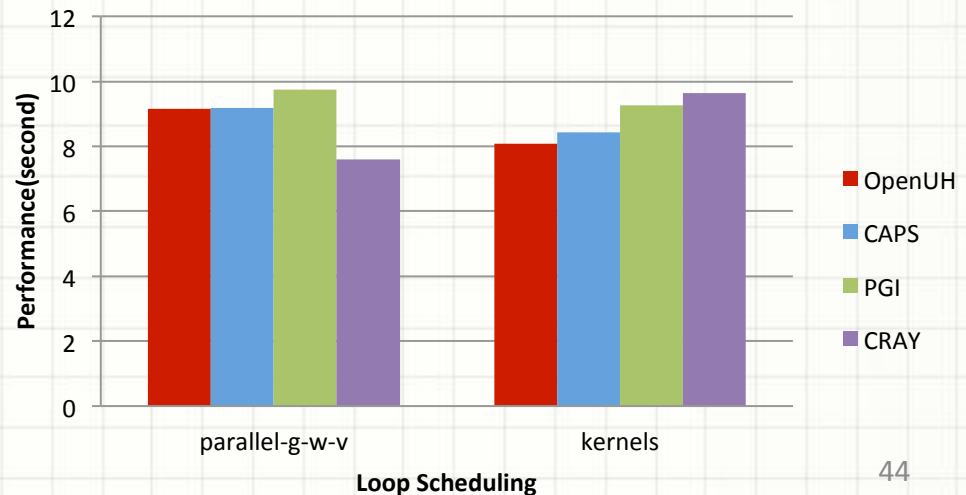
CAPS: default

CRAY: default

Same experimental platform used for OpenUH, CAPS and PGI

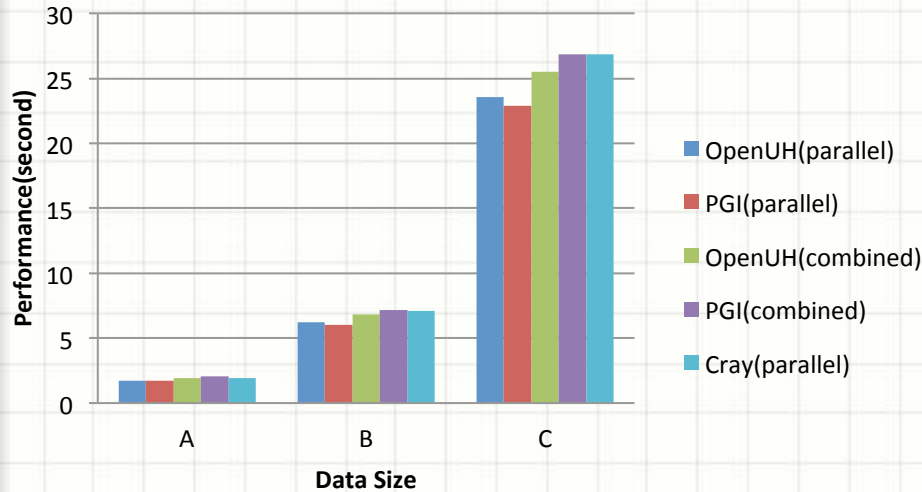
CRAY platform used for Cray machine

Laplacian



NAS Benchmark

NAS EP




Combined: parallel + kernels
Cray: use default loop scheduling, #pragma acc loop

Same experimental platform used for OpenUH, CAPS and PGI

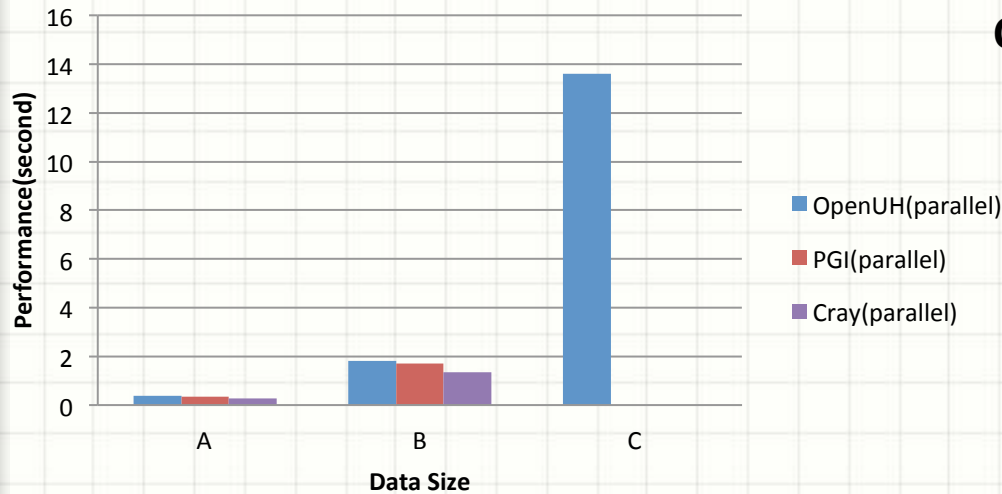
CRAY platform used for Cray machine

NAS CG




NAS Benchmark

NAS MG



Combined: parallel + kernels

NAS SP

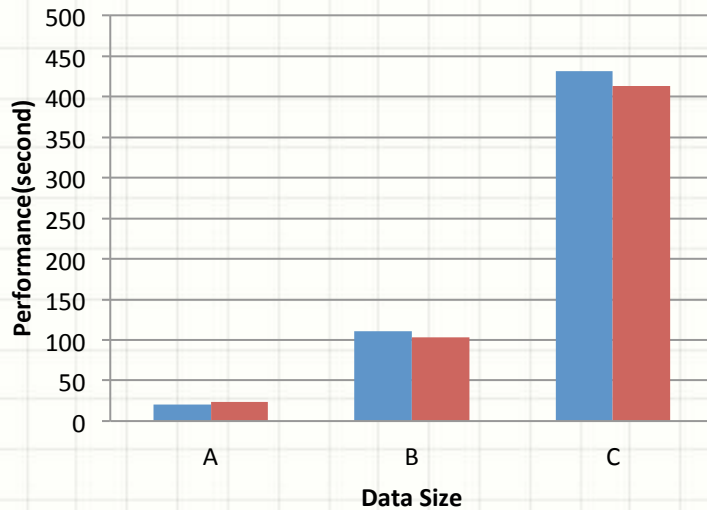


Same experimental platform used for OpenUH, CAPS and PGI

CRAY platform used for Cray machine


NAS Benchmark


NAS BT



Combined: parallel + kernels

NAS LU





V. Future and Conclusion

Future Work

- ❑ Support Fortran
- ❑ Support Xeon Phi/AMD GPGPUs and APU
- ❑ Perform more optimization: Irregular Memory access optimization
- ❑ Provide a more robust OpenACC implementation

Conclusion

- ❑ Open source OpenACC research compiler, based on Open64
- ❑ Competitive performance, compared to other commercial compilers
- ❑ Proposed regular loop scheduling for parallel region and non-standard loop scheduling for kernels region

