



中山大學
SUN YAT-SEN UNIVERSITY



国家超级计算广州中心
NATIONAL SUPERCOMPUTER CENTER IN GUANGZHOU

Advanced Computer Architecture

高级计算机体系结构

第8讲：DLP and GPU (3)

GPU在计算流体力学（CFD）中的应用

张曦

DCS5367, 11/23/2021



中山大學
SUN YAT-SEN UNIVERSITY



Self-introduction





- XI ZHANG (张曦)
- Work: Engineer in NSCC-GZ
- Interests: High Performance Computing, Computational Fluid Dynamics
- Contact: xi.zhang@nscg-gz.cn
zhangx299@mail.sysu.edu.cn

Detail

- Background
- Developing Stage
- Optimizations Stage
- Conclusion
- Future work
- Thinking more

Background

- Scientific computing (科学计算) in Science and Engineering
 - Scientific computing is regarded as the third methodology in science and engineering. (theory 理论 and experiment 实验)
 - Scientific computing is widely used in aerospace science and technology, ocean engineering, nuclear industry, etc.



Background

- Scientific computing in Science and Engineering
 - GPU plays more and more important role in scientific computing.
 - Many High Performance Computing Systems are built with GPU.



天河-1:
4.7 PFLOPs
CPU (飞腾1000)+GPU
(Nvidia Tesla M2050)



Summit:
200 PFLOPs
CPU (Power 9)+GPU
(Nvidia Tesla V100)



Frontier:
CPU (AMD
Zen 3)+GPU
(AMD MI200)

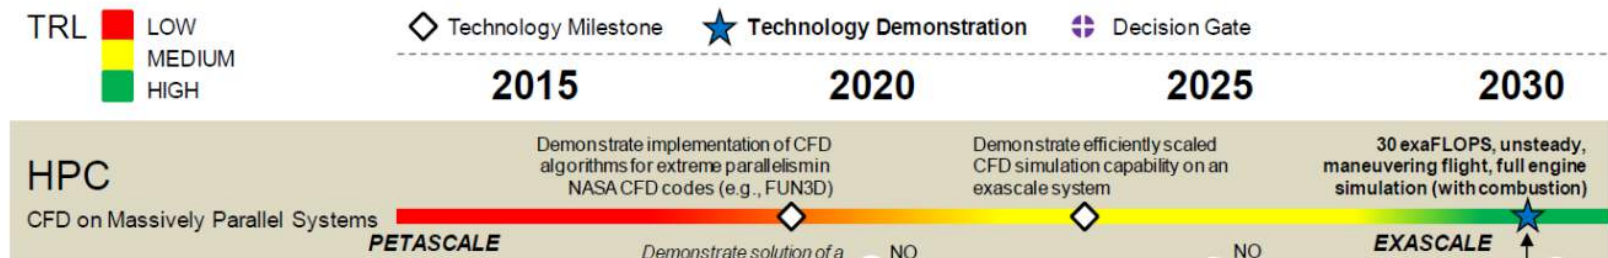
Background

- Computational Fluid Dynamics (CFD)

- A process of mathematically modeling (数学建模) a physical phenomenon (物理现象) involving fluid flow and solving it numerically (数值求解) using the computational prowess.

$$\frac{\partial u}{\partial t} + \nabla \cdot \mathbf{f}(u, \nabla u) = S(\mathbf{x}, t),$$

- A cross-discipline subject including mathematics (数学), fluid dynamics (流体力学), and computer science (计算机科学).
- CFD requires large computing source (E级计算需求)



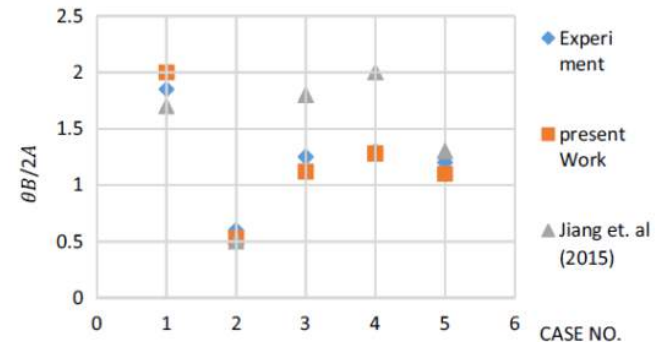

Background

- What is CFD?
 - Quite like Physical Rendering in Computer Graphics

CG




CFD



Background

- The process of CFD
 - Domain discretization (计算域离散)



- Equation discretization (方程离散)

$$\frac{\partial u}{\partial t} + \nabla \cdot \mathbf{f}(u, \nabla u) = S(\mathbf{x}, t), \quad \longrightarrow \quad \begin{aligned} \frac{V}{\Delta \tau} \Delta \mathbf{q} + \frac{\partial \hat{\mathbf{R}}}{\partial \mathbf{q}} \Delta \mathbf{q} &= -\mathbf{R}(\mathbf{q}^n) \\ \mathbf{q}^{n+1} &= \mathbf{q}^n + \Delta \mathbf{q} \end{aligned}$$

Background



- Difficulties of CFD by GPU

- Precision (精确): code from CPU to GPU, data dependence (数据依赖)

Algorithm 1 Flux Summation (FS) by FVM face-loop

```

1: for faceID = 0 to numFaces-1 do
2:   ownVolID ← owner[faceID]
3:   ngbVolID ← neighbor[faceID]
4:   res[ownVolID] ← res[ownVolID]+flux[faceID]
5:   res[ngbVolID] ← res[ngbVolID]-flux[faceID]
6: end for
  
```



GPU multi-thread computing

$$A = L + D + U$$



$$(L + D)x^* = b - Ux^n$$

$$(L + U)x^{n+1} = b - Ux^n - Lx^*$$


$$\begin{bmatrix} a_{11} & a_{12} & a_{13} & \dots & a_{1n} \\ a_{21} & a_{22} & a_{23} & \dots & a_{2n} \\ a_{31} & a_{32} & a_{33} & \dots & a_{3n} \\ \dots & \dots & \dots & \dots & \dots \\ a_{n1} & a_{n2} & a_{n3} & \dots & a_{nn} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ \dots \\ x_n \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 \\ b_3 \\ \dots \\ b_n \end{bmatrix}$$

Background

- Difficulties of CFD by GPU
 - Performance (性能): non-coalescing (非对齐) memory access




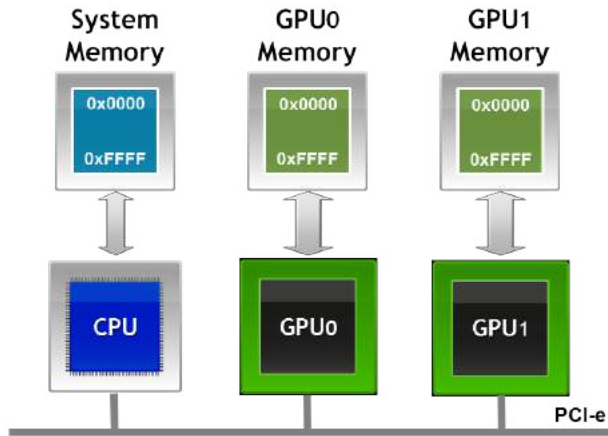
Irregular data storage (不规则存储)



Indirect data access (间接访问)

Background

- Difficulties of CFD by GPU
 - Data transfer between Host and GPU
 - Multi-GPU computing(多GPU计算): parallel framework



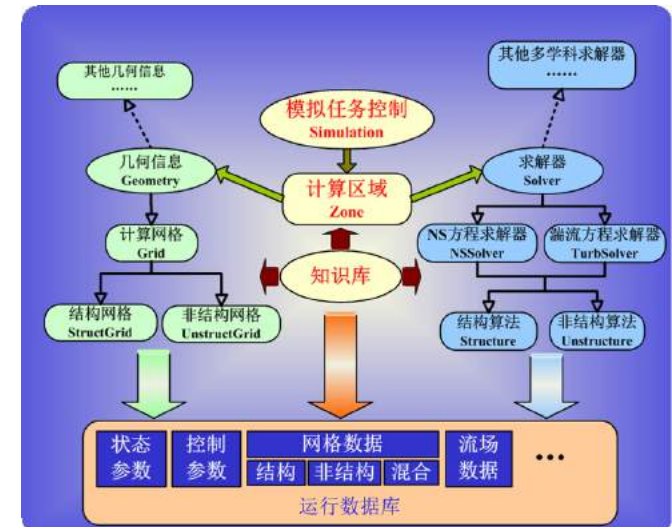
Background

- CFD application: NNW-PHengLEI (风雷)
 - High speed and compressible flow
 - Aeronautics (空气动力学) e.g. plane
 - CPU computing only
 - C++, Object-Oriented
 - Open source



(<https://forge.osredm.com/projects/p68217053/PHengLEI>)

- A similar CFD software Fun3D
 - Spend almost 10 years in R&D on GPU



Developing Stage

- CUDA C Programming
 - Close to loop induced computing (靠近f循环计算部分)
 - Interface functions for calling CUDA kernels (接口函数)
 - Test every CUDA kernel by comparing with CPU results (测试)

```
#ifdef CPURUN
    for ( int m = 0; m < n1; ++ m )
    {
        for ( int iCell = 0; iCell < nTotal; ++ iCell )
        {
            q[m][iCell] = qold[m][iCell];
        }
    }
#endif
#ifdef __CUDACC__
    CallGPULoadQ(nTotal, n1);
    #ifdef CUDAUNITTEST
        TestGPULoadQ(q);
    #endif
#endif
```

```
void CallGPULoadQ(...) {
    GPULoadQ<<<gridSize, blockSize>>> (...);
}


__global__ void GPULoadQ(...) {
    const int tidx= threadIdx.x;
    const int bidx= blockIdx.x;

    for ( int m = 0; m < n1; ++ m ){
        for ( int icell = bidx* blockDim.x+ tidx;
                icell < nTotal;
                icell += blockDim.x* blockDim.x){
            q[m*nTotal+ icell] = qold[m*nTotal+icell];
        }
    }
}
```

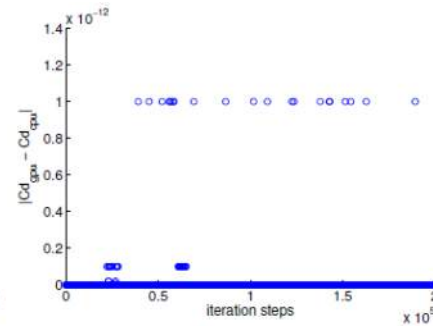
截图(Alt + A)

Developing Stage


- What can induce error on GPU computing?
 - Loop order on GPU (乱序执行循环)
 - CUDA supported optimizations such as MAD (乘加操作)
 - Some CUDA supported mathematical functions such as pow (幂次运算)




(a) C_d



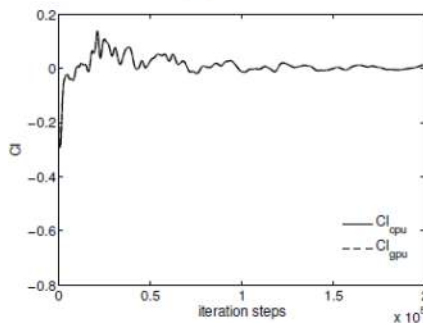
(c) $|C_{d_{GPU}} - C_{d_{CPU}}|$



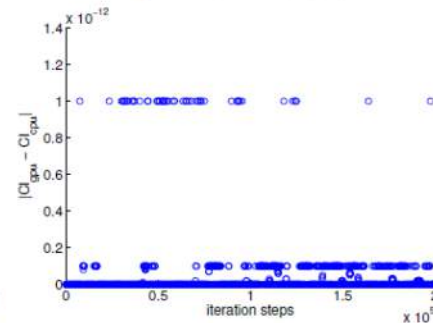
(a) C_p on top side of Chnt (CPU)




(c) C_p on bottom side of Chnt (CPU)




(b) C_l



(d) $|C_{l_{GPU}} - C_{l_{CPU}}|$



(b) C_p on top side of Chnt (GPU)



(d) C_p on bottom side of Chnt (GPU)

Optimization Stage

- Atomic operations or graph coloring for resolving data dependence
 - Atomic operations: hardware supported method
 - Graph coloring: software supported method


```

1: for faceID = nBoundFace to nTotalFace-1 do
2:   Le ← leftCellOfFace[faceID]
3:   Re ← rightCellOfFace[faceID]
4:   for eqnID = 0 to numEqn - 1 do
5:     setAdd(res[eqnID][Le], flux[eqnID][faceID])
6:     setAdd(res[eqnID][Re], flux[eqnID][faceID])
7:   end for
8: end for
    
```

Algorithm 4 Summation of Flux by atomic operation (SF-AT)

```

1: <GPU kernel Begin>
2: threadID ← threadIdx.x + blockIdx.x * blockDim.x
3: for faceID = nBoundFace + threadID to nTotalFace-1 do
4:   Le ← leftCellOfFace[faceID]
5:   Re ← rightCellOfFace[faceID]
6:   for eqnID = 0 to numEqn - 1 do
7:     atomicAdd(res[eqnID*nTotalCell+Le],
8:               flux[eqnID*nTotalFace+faceID])
9:     atomicAdd(res[eqnID*nTotalCell+Re],
10:              flux[eqnID*nTotalFace+faceID])
11:   end for
12: setAdd(faceID, blockDim.x * gridDim.x)
13: end for
14: <GPU kernel End>
    
```



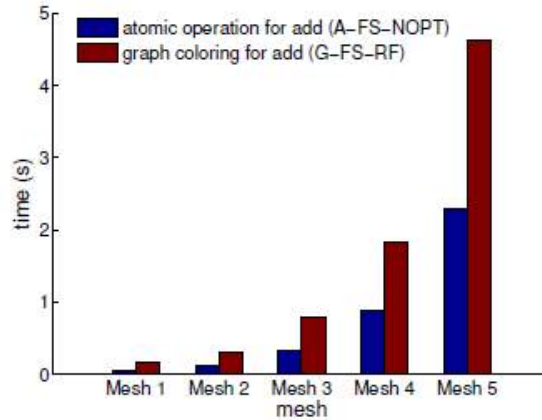
Algorithm 3 Summation of Flux by graph coloring(SF-GC)

```

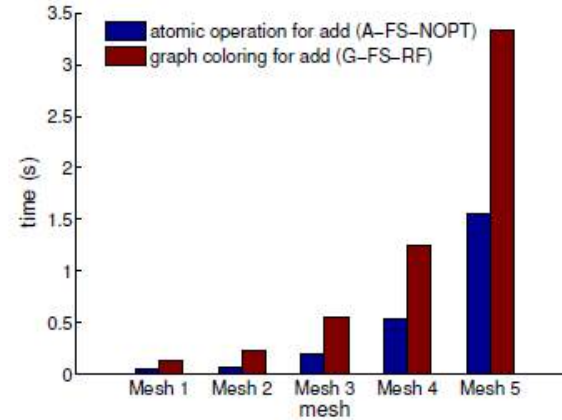
1: Reorder flux according to faceGroup
2: for groupID = 0 to nGroups-1 do
3:   numFacesInGroup ← numFaceOfGroup[groupID]
4:   groupStart ← offsetFaceGroup[groupID]
5:   <GPU kernel Begin>
6:   threadID ← threadIdx.x + blockIdx.x * blockDim.x
7:   for faceGroupID = threadID to numFacesInGroup-1 do
8:     groupFaceID ← groupStart + faceGroupID
9:     faceID ← faceGroup[groupFaceID]
10:    Le ← leftCellOfFace[faceID]
11:    Re ← rightCellOfFace[faceID]
12:    for eqnID = 0 to numEqn - 1 do
13:      setAdd(res[eqnID*nTotalCell+Le],
14:            flux[eqnID*nTotalFace+faceID])
15:      setAdd(res[eqnID*nTotalCell+Re],
16:            flux[eqnID*nTotalFace+faceID])
17:    end for
18:    setAdd(faceGroupID, blockDim.x * gridDim.x)
19:  end for
20: <GPU kernel End>
21: end for
    
```

Optimization Stage

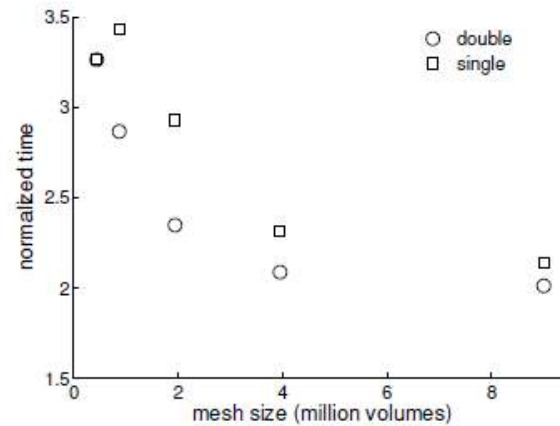
- Atomic operations V.S. graph coloring



(a) double



(b) single



(c) time ratio

Optimization Stage

- Data dependence resolving in LU-SGS scheme (数据依赖性)
 - Data dependence in LU-SGS
 - Multi-color LU-SGS

$$\begin{bmatrix} a_{11} & a_{12} & a_{13} & \dots & a_{1n} \\ a_{21} & a_{22} & a_{23} & \dots & a_{2n} \\ a_{31} & a_{32} & a_{33} & \dots & a_{3n} \\ \dots & \dots & \dots & \dots & \dots \\ a_{n1} & a_{n2} & a_{n3} & \dots & a_{nn} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ \dots \\ x_n \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 \\ b_3 \\ \dots \\ b_n \end{bmatrix}$$

$$A = L + D + U$$

$$\text{L-SGS} \quad (L + D)x^* = b - Ux^n$$

$$\text{U-SGS} \quad (L + U)x^{n+1} = b - Ux^n - Lx^*$$

```

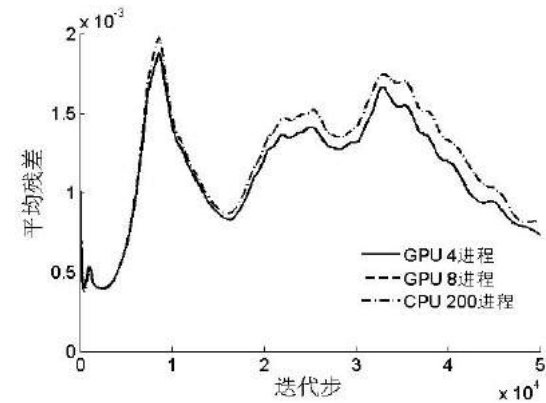
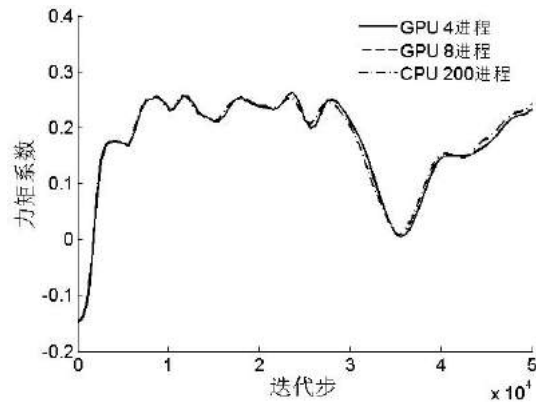

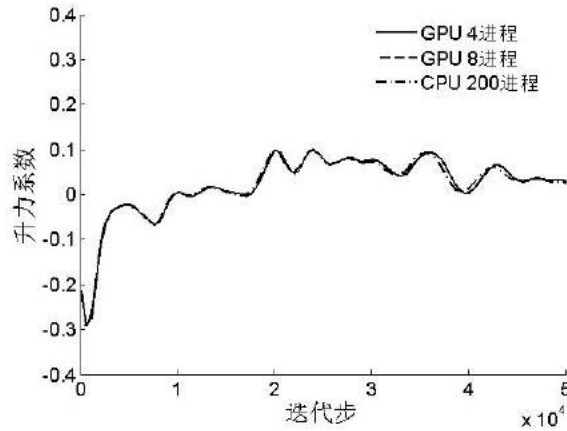
1: for colorID = 0 to numColors-1 do
2:   numColorGroup ← colorGroupNum[colorID]
3:   posiColorGroup ← colorGroupPosi[colorID]
4:   <GPU kernel Begin>
5:   for offset = 0 to numColorGroup do
6:     cellID ← colorGroup[posiColorGroup+offset]
7:     LowerSweepOnOneCell(cellID)
8:   end for
9:   <GPU kernel End>
10: end for
    
```

```

16: for colorID = numColors-1 to 0 do
17:   numColorGroup ← colorGroupNum[colorID]
18:   posiColorGroup ← colorGroupPosi[colorID]
19:   <GPU kernel Begin>
20:   for offset = 0 to numColorGroup do
21:     cellID ← colorGroup[posiColorGroup+offset]
22:     UpperSweepOnOneCell(cellID)
23:   end for
24:   <GPU kernel End>
25: end for
    
```




Optimization Stage

- Result of Multi-dolor LU-SGS




Optimization Stage

- Volume renumber (体编号重排序)
 - Reverse Cuthill-McKee (RCM)
 - Reduce adjacent matrix bandwidth
 - Easing non-coalescing data access



renumber




Optimization Stage

- Face renumber (面编号重排序)
 - Renumber face in a volume (以体为单位顺序重排面编号)
 - Optimizing data locality (提高数据局部性)


算法 1 面编号排序算法 \leftarrow

```
1: Reorder faces in cellFace by ascending order
2: mapFace  $\leftarrow$  -1
3: labelFace  $\leftarrow$  nBoundFace-1
4: for cellID = 0 to nTotalCell-1 do
5:   offset  $\leftarrow$  offsetCellFace[cellID]
6:   numFaces  $\leftarrow$  numFaceOfCell[cellID]
7:   for faceInCell = 0 to numFaces-1 do
8:     faceID  $\leftarrow$  cellFace[offset+faceInCell]
9:     if faceID > nBoundFace-1 then
10:      if mapFace[faceID] == -1 then
11:        labelFace  $\leftarrow$  labelFace+1
12:        mapFace[faceID]  $\leftarrow$  labelFace
13:      end if
14:    end if
15:  end for
16: end for
17: Update mesh connectivity information by mapFace  $\leftarrow$ 
```




Optimization Stage

- Results of volume and face renumber
 - Overall Performance (对GPU程序整体性能的影响)



- GPU kernels (对GPU kernels的影响)



Optimization Stage

- Loop mode adjust 1 (循环模式调整)

- Data interpolation
- Can a face loop be replaced?

1: <GPU kernel Begin>

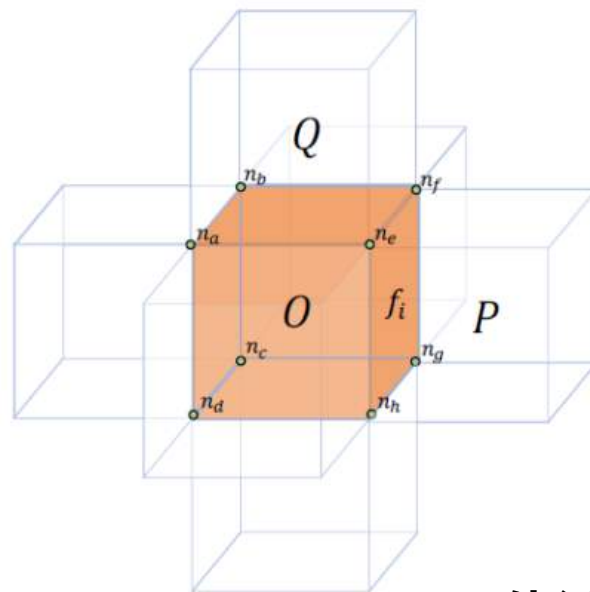
2: threadID ← threadIdx.x + blockIdx.x * blockDim.x
 3: **for** faceID = nBoundFace + threadID to nTotalFace - 1 **do**

4: Le ← leftCellOfFace[faceID]
 5: Re ← rightCellOfFace[faceID]
 6: faceNodeStart ← offsetFaceNode[faceID]
 7: numNodeInFace ← numNodeOfFace[faceID]
 8: **for** faceNodeID = 0 to numNodeInFace - 1 **do**

面循环

9: nodeID ← faceNodes[faceNodeStart + faceNodeID]
 10: **for** eqnID = 0 to numEqn - 1 **do**
 11: **atomicAdd**(qNode[eqnID * nTotalNode + nodeID],
 q[eqnID * nTotalCell + Le])
 12: **end for**
 13: **atomicAdd**(tNode[nodeID], t[Le])
 14: **atomicAdd**(nCount[nodeID], 1)
 15: **for** eqnID = 0 to numEqn - 1 **do**
 16: **atomicAdd**(qNode[eqnID * nTotalNode + nodeID],
 q[eqnID * nTotalCell + Re])
 17: **end for**
 18: **atomicAdd**(tNode[nodeID], t[Re])
 19: **atomicAdd**(nCount[nodeID], 1)

20: **end for**
 21: **setAdd**(faceID, blockDim.x * gridDim.x)
 22: **end for**
 23: <GPU kernel End>



体循环

2: <GPU kernel Begin>


3: threadID ← threadIdx.x + blockIdx.x * blockDim.x
 4: **for** cellID = threadID to nTotalCell - 1 **do**

5: cellNodePosi ← offsetCellNode[cellID]
 6: **for** offset = 0 to numNodeOfCell[faceID] - 1 **do**
 7: nodeID ← cellNodes[cellNodePosi + offset]
 8: accessFrequency ←
 cellNodeCount[cellNodePosi + offset]
 9: **atomicAdd**(qNode[eqnID * nTotalNode + nodeID],
 accessFrequency * q[eqnID * nTotalCell + cellID])


10: **end for**
 11: **setAdd**(cellID, blockDim.x * gridDim.x)
 12: **end for**
 13: <GPU kernel End>

Optimization Stage

- Results of loop mode adjust




(a) V100



(b) K80

Optimization Stage

- Loop mode adjust 2 (循环模式调整)
 - Data comparison
 - Can a face loop be replaced?



Algorithm 10 Local pressure comparing by face loop (LPC-F)

```

1: <GPU kernel Begin>
2: threadID ← threadIdx.x + blockIdx.x * blockDim.x
3: for faceID = threadID + nBoundFace to nTotalFace - 1 do
4:   Le ← leftCellOfFace[faceID]
5:   Re ← rightCellOfFace[faceID]
6:   atomicMin(pMin[Le], pressure[Re])
7:   atomicMax(pMax[Le], pressure[Re])
8:   atomicMin(pMin[Re], pressure[Le])
9:   atomicMax(pMax[Re], pressure[Le])
10:  setAdd(faceID, blockDim.x * gridDim.x)
11: end for
12: <GPU kernel End>
    
```

面循环

Algorithm 11 Local pressure comparing by cell loop (LPC-C)

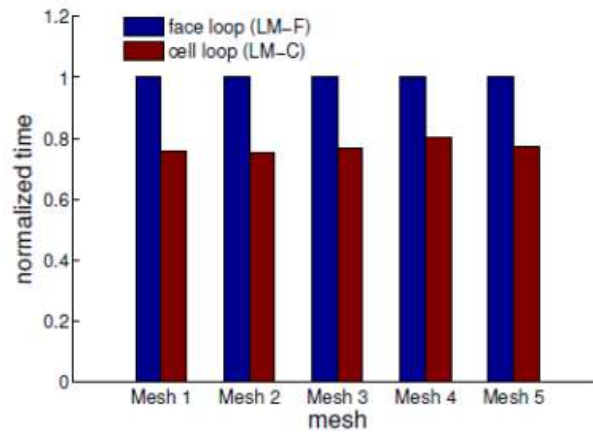
```

1: <GPU kernel Begin>
2: threadID ← threadIdx.x + blockIdx.x * blockDim.x
3: for cellID = threadID to nTotalCell - 1 do
4:   numCell ← numCellCell[cellID]
5:   cellStart ← offsetCellCell[cellID]
6:   for cellInCellID = 0 to numCell - 1 do
7:     cellCellID ← cellCell[cellStart + cellInCellID]
8:     setCompMin(pMin[cellID], pressure[cellCellID])
9:     setCompMax(pMax[cellID], pressure[cellCellID])
10:  end for
11:  setAdd(cellID, blockDim.x * gridDim.x)
12: end for
13: <GPU kernel End>
    
```


体循环

Optimization Stage

- Results of loop mode adjust



(a) V100



(b) K80

Optimization Stage

- Nested loop split (嵌套循环拆分)
 - Loop on geometry
 - Loop on dimensions

```
1: for faceID = nBoundFace to nTotalFace-1 do
```

```
2:   Le ← leftCellOfFace[faceID]
```

```
3:   Re ← rightCellOfFace[faceID]
```

```
4:   for eqnID = 0 to numEqn - 1 do
```

```
5:     setAdd(res[eqnID][Le], flux[eqnID][faceID])
```

```
6:     setAdd(res[eqnID][Re], flux[eqnID][faceID])
```

```
7:   end for
```

```
8: end for
```

面循环

方程数量循环

```
1: <GPU kernel Begin>
```

```
2: threadID ← threadIdx.x + blockIdx.x * blockDim.x
```

```
3: for faceID = nBoundFace + threadID to nTotalFace - 1 do
```

```
4:   Le ← leftCellOfFace[faceID]
```

```
5:   Re ← rightCellOfFace[faceID]
```

```
6:   for eqnID = 0 to numEqn - 1 do
```

```
7:     atomicAdd(res[eqnID * nTotalCell + Le],  
              flux[eqnID * nTotalFace + faceID])
```

```
8:     atomicAdd(res[eqnID * nTotalCell + Re],  
              flux[eqnID * nTotalFace + faceID])
```

```
9:   end for
```

```
10:  setAdd(faceID, blockDim.x * gridDim.x)
```

```
11: end for
```

```
12: <GPU kernel End>
```



```
1: for eqnID = 0 to numEqn - 1 do
```

```
2:   <GPU kernel Begin>
```

```
3:   threadID ← threadIdx.x + blockIdx.x * blockDim.x
```

```
4:   for faceID = nBoundFace + threadID to nTotalFace - 1 do
```

```
5:     Le ← leftCellOfFace[faceID]
```

```
6:     Re ← rightCellOfFace[faceID]
```

```
7:     atomicAdd(res[eqnID * nTotalCell + Le],  
              flux[eqnID * nTotalFace + faceID])
```

```
8:     atomicAdd(res[eqnID * nTotalCell + Re],  
              flux[eqnID * nTotalFace + faceID])
```

```
9:     setAdd(faceID, blockDim.x * gridDim.x)
```


```
10:   end for
```

```
11:   <GPU kernel End>
```

```
12: end for
```


Optimization Stage

- Results of Nested loop split



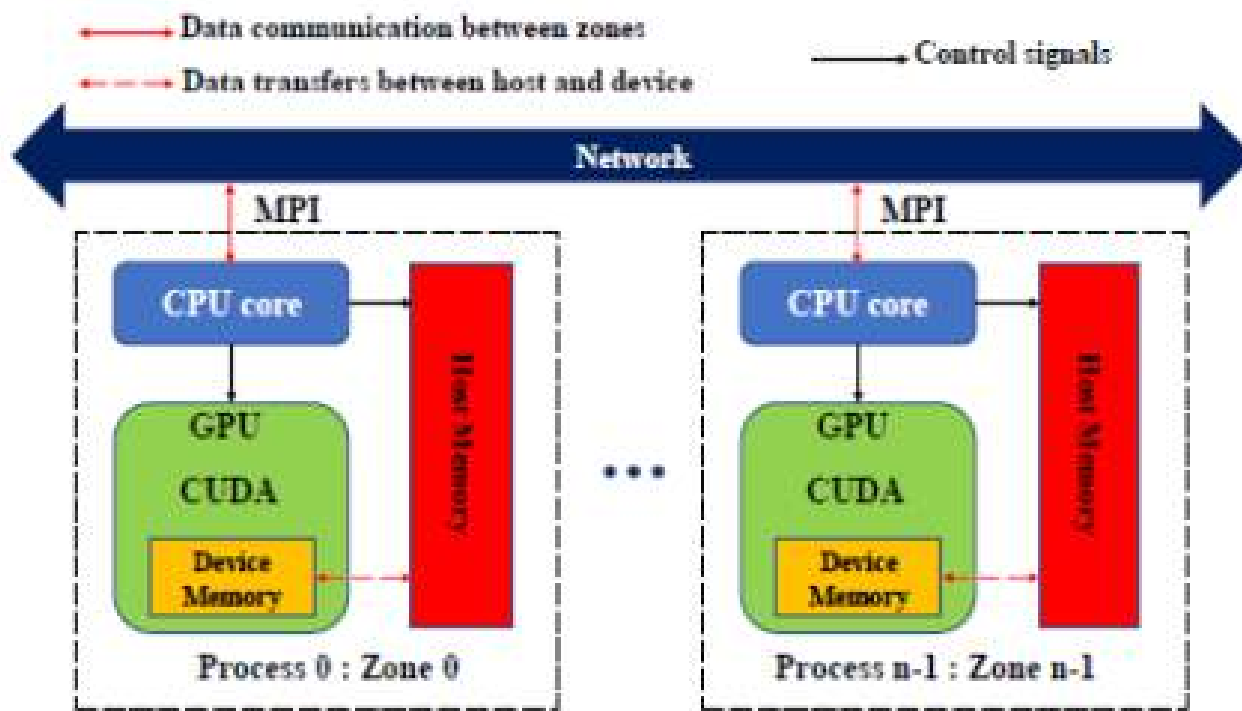
Optimization Stage

- Performance comparison of single GPU and 28 CPU cores
 - GPU: 4 x Nvidia Tesla V100
 - CPU: 2 x Intel Xeon Gold 6132



Optimization Stage

- Multi-GPU computing (多GPU计算)
 - MPI-CUDA parallel framework
 - CPU is only used for controlling GPU



Optimization Stage

- Pack and Unpack MPI data on GPU

Algorithm 13 Pack MPI data on GPU



```
1: Get dataSend and dataIF by data name
2: for ngbZoneID = 0 to numNgbZone do
3:   startFace ← startFaceForSend[ngbZoneID]
4:   startSend ← startDataSend[ngbZoneID]
5:   numNgbFace ← nIFaceOfNgbZone[ngbZoneID]
6:   <GPU kernel Begin>
7:   threadID ← threadIdx.x + blockIdx.x * blockDim.x
8:   for faceID = threadID to numNgbFace do
9:     sendID ← faceForSend[startFace + faceID]
10:    for eqnID = 0 to numEqn do
11:      ngbZoneFaceID ← startSend +
        eqnID * numNgbFace + faceID
12:      interfaceID ← eqnID * nInterfaceTotal + sendID
13:      dataSend[ngbZoneFaceID] ← dataIF[interfaceID]
14:    end for
15:    setAdd(faceID, blockDim.x * gridDim.x)
16:  end for
17:  <GPU kernel End>
18: end for
```

Algorithm 14 Unpack MPI data on GPU

```
1: Get dataSend and dataIF by data name
2: for ngbZoneID = 0 to numNgbZone do
3:   startFace ← startFaceForRecv[ngbZoneID]
4:   startRecv ← startDataRecv[ngbZoneID]
5:   numNgbFace ← nIFaceOfNgbZone[ngbZoneID]
6:   <GPU kernel Begin>
7:   threadID ← threadIdx.x + blockIdx.x * blockDim.x
8:   for faceID = threadID to numNgbFace do
9:     recvID ← faceForRecv[startFace + faceID]
10:    for eqnID = 0 to numEqn do
11:      ngbZoneFaceID ← startRecv +
        eqnID * nIFaceOfNgbZone + faceID
12:      interfaceID ← eqnID * nInterfaceTotal + recvID
13:      dataIF[interfaceID] ← dataRecv[ngbZoneFaceID]
14:    end for
15:    setAdd(faceID, blockDim.x * gridDim.x)
16:  end for
17:  <GPU kernel End>
18: end for
```

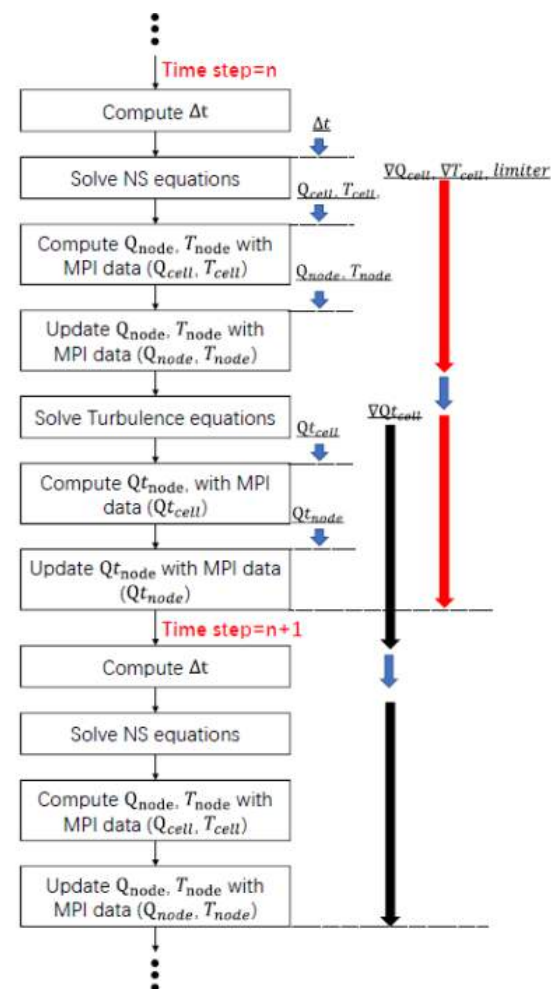
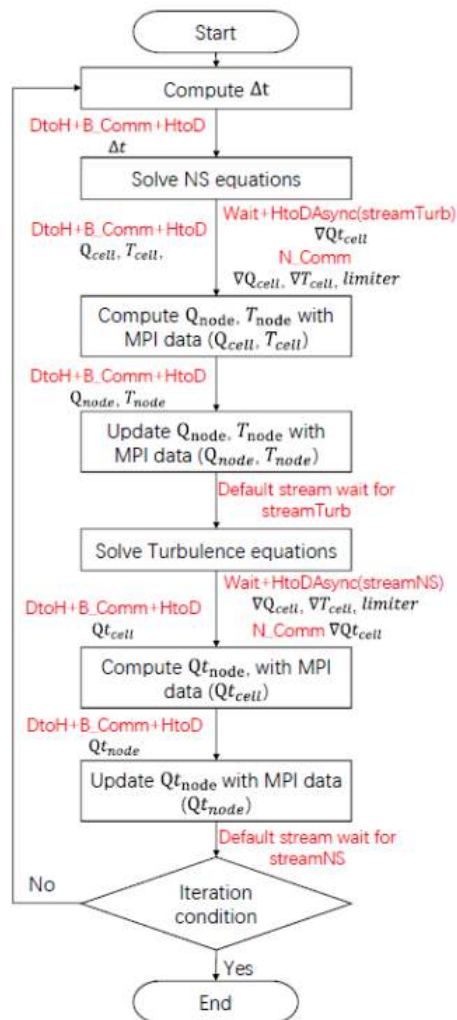
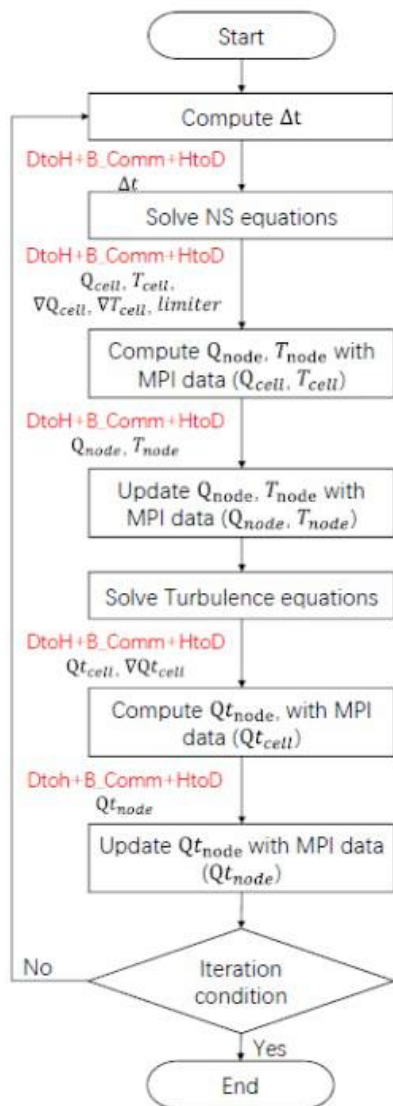
Optimization Stage

- Result of pack and Unpack MPI data on GPU



Optimization Stage

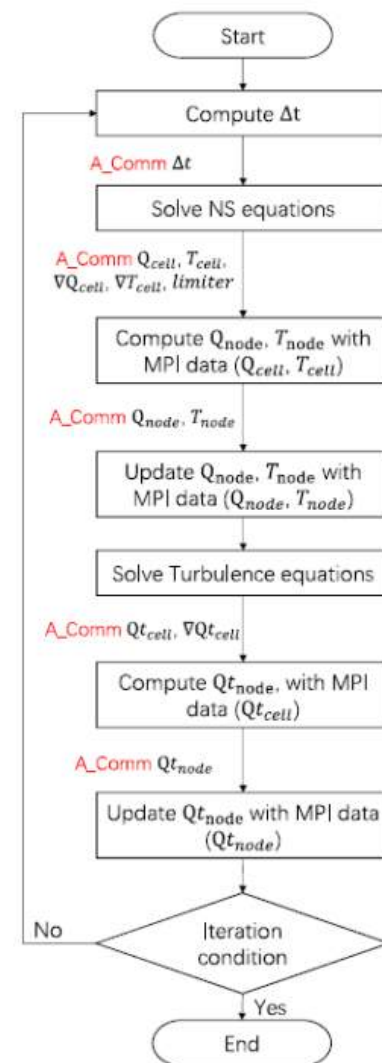


- Communication and computing overlap



Optimization Stage

- CUDA-AWARE-MPI


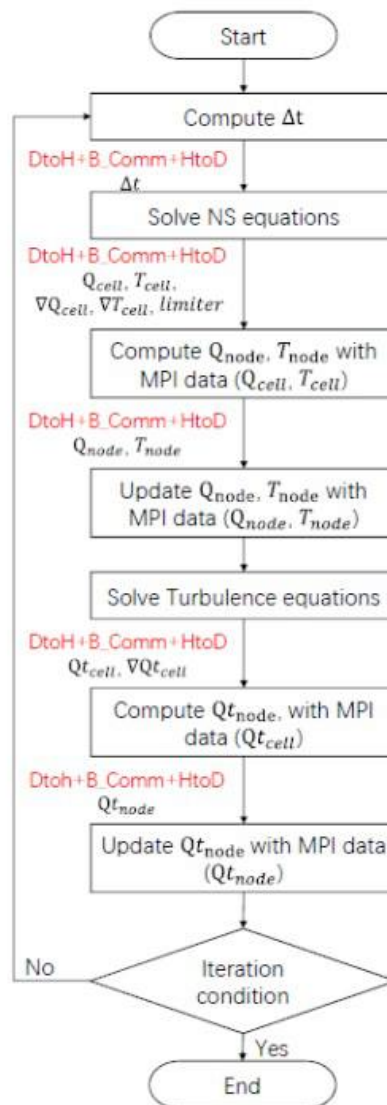

GPUDirect RDMA



Optimization Stage



- CUDA-AWARE-MPI

GPUDirect RDMA



Optimization Stage


- Result of 3 MPI-CUDA parallel framework




Optimization Stage

- Strong and weak scaling test (扩展性测试)

强扩展性测试



每个GPU计算50万网格弱扩展性



Conclusion

- 高保真异构程序架构确保计算精度
- Multi-color LU-SGS解决数据依赖性
- 提高数据局部性优化：网格编号重排、循环模式优化、循环嵌套拆分、原子操作解决资源竞争
- 多GPU计算：开发MPI-CUDA并行框架及其改进版本；实现数据在GPU上的打包、解包
- 实现了GPU的高精度计算；最多实现了200块GPU的并行计算，并行效率较高

Future Work

- Multi-color LU-SGS的进一步优化
- Mixed Precision computing
- Tensor core
- CFD+AI

Thinking more ...

- 硬件
 - 领域定制硬件 FPGA, RISC-V等
- 软件
 - 可移植性
 - 可维护性
 - 领域定制软件

Thank you for your listening

联系方式: zhangx299@mail.sysu.edu.cn
xi.zhang@nscg-gz.cn

