

*Bring Many-core Thinking to
High End Research and Education*



RANS CFD Solver on Fermi

Peng Wang, NVIDIA

On Behalf of

James Lin

HPC Center, Shanghai Jiao Tong University

GTC @San Jose, May 2012



Outline

- Related Work
- Background
- Design and Implementation
- Results analysis
- Conclusion and Future work

Related Work



Navier-Stokes Stanford University Solver(NSSUS)

15-20x Speedup On NVIDIA 8800GTX

Stanford

E.Elsen etc

2008.12

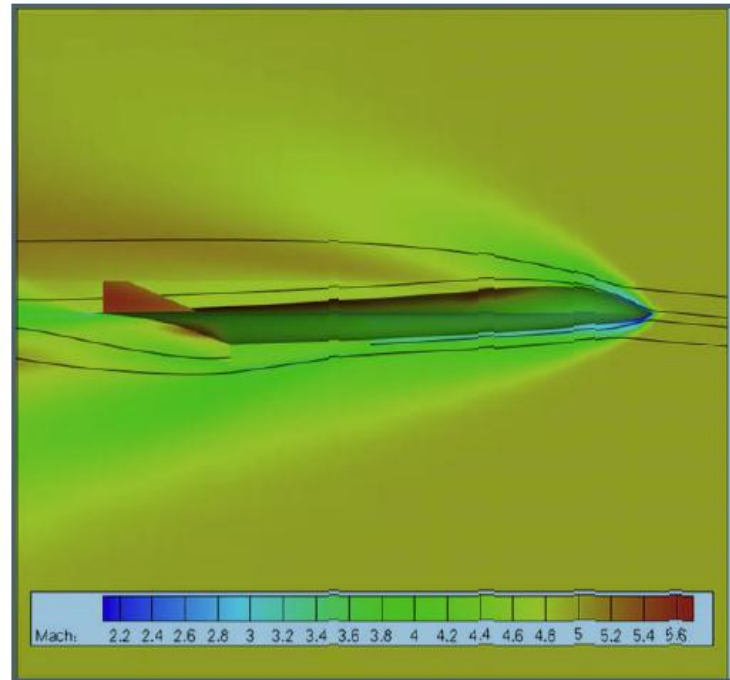


Table 2
Speed-ups for the hypersonic vehicle computation

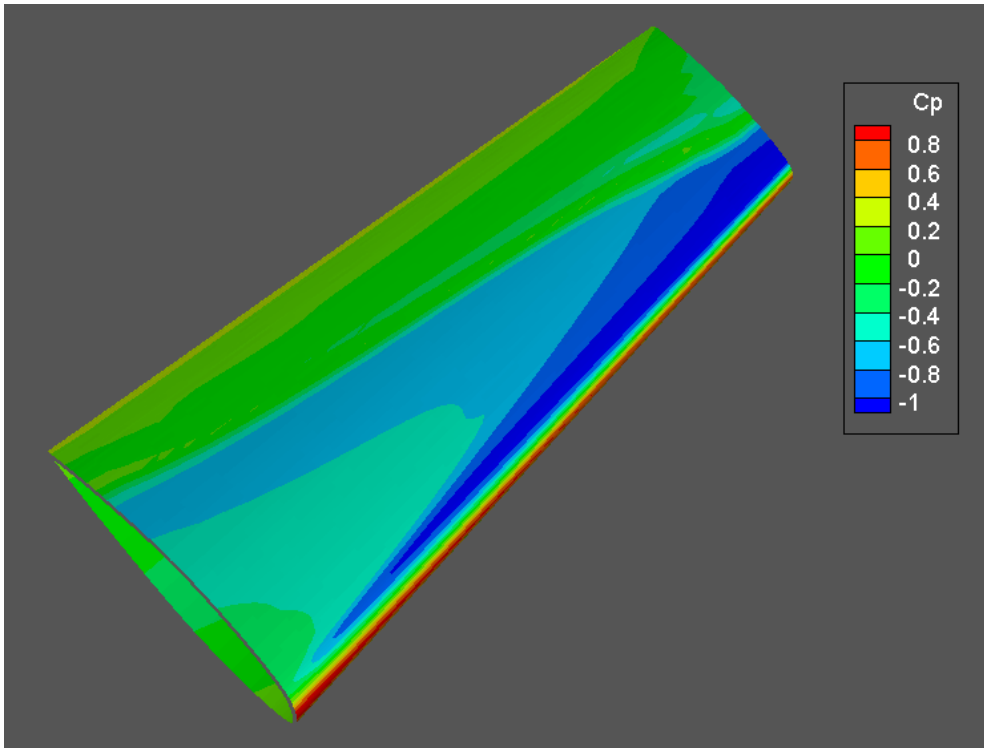
Mesh size	Multi-grid cycle	Speed-up
720 k	Single grid	15.4
720 k	2 Grids	11.2
1.5 M	Single grid	20.2
1.5 M	2 Grids	15.8



Other Work

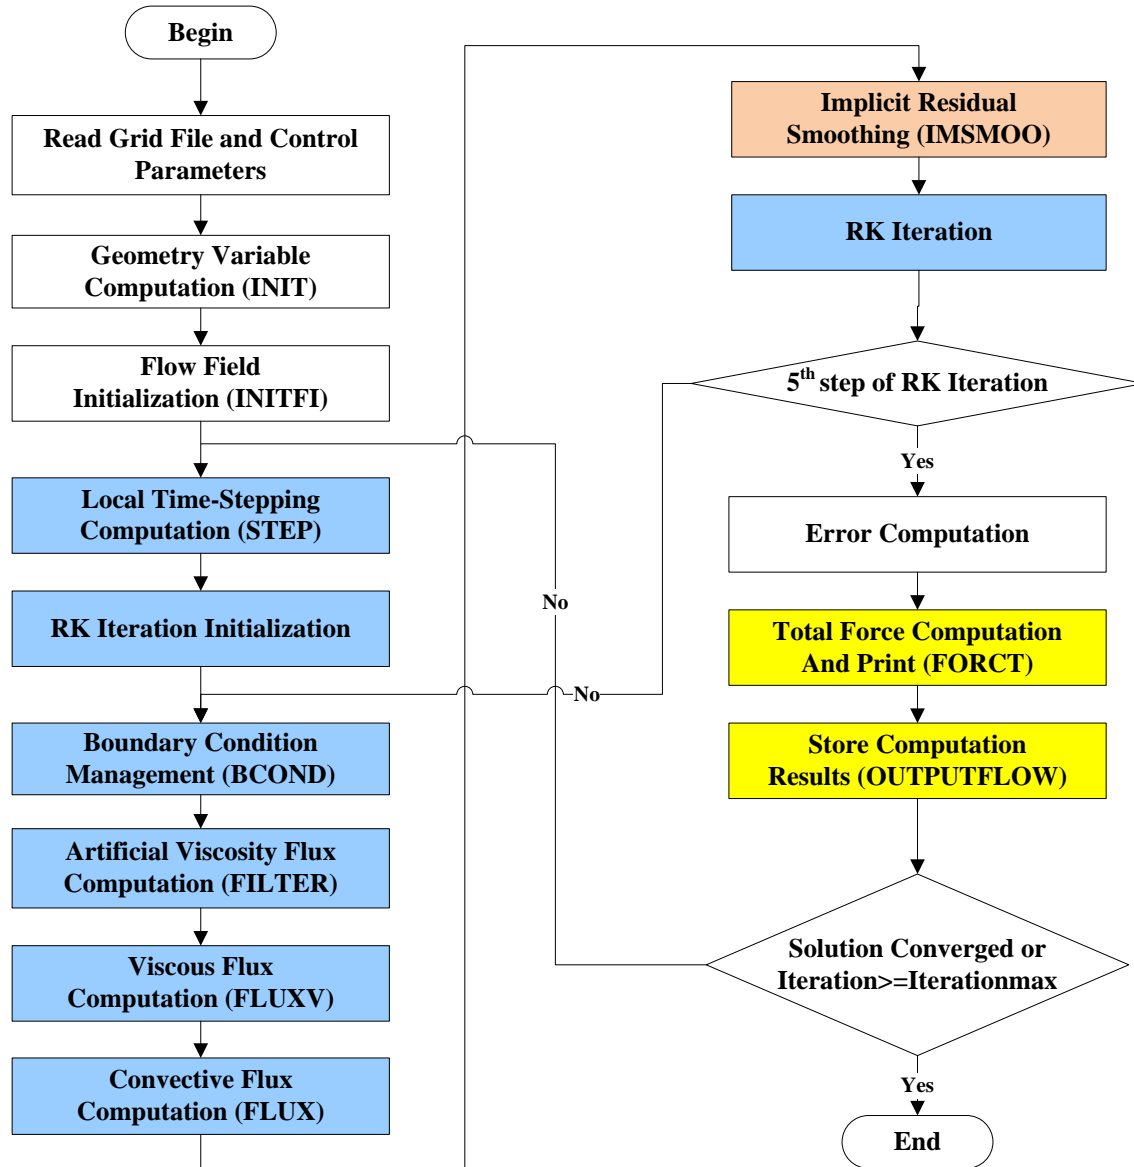
1. J. C. Thibault et al. implemented a Navier-Stokes(NS) Equation solver for incompressible flow on Tesla C870 boards, 13X on 1 GPU and 21X on 2 GPUs
2. Tobias Brandvik et al. ported Euler Equations onto GPU and a speedup of 29x in 2D and 16x in 3D
3. A. Corrigan et al. implemented 3D unstructured grid Euler solver for the inviscid compressible flow, and get 33x speedup on 1 C1060
4. I. C. Kampolis et al. use of the 2-D Euler solver for inviscid compressible flow to optimize the aerodynamic performance of the airfoil, get 15 X speedup with single precision and 21X with double on GTX 285

RANS Solver in SJTU

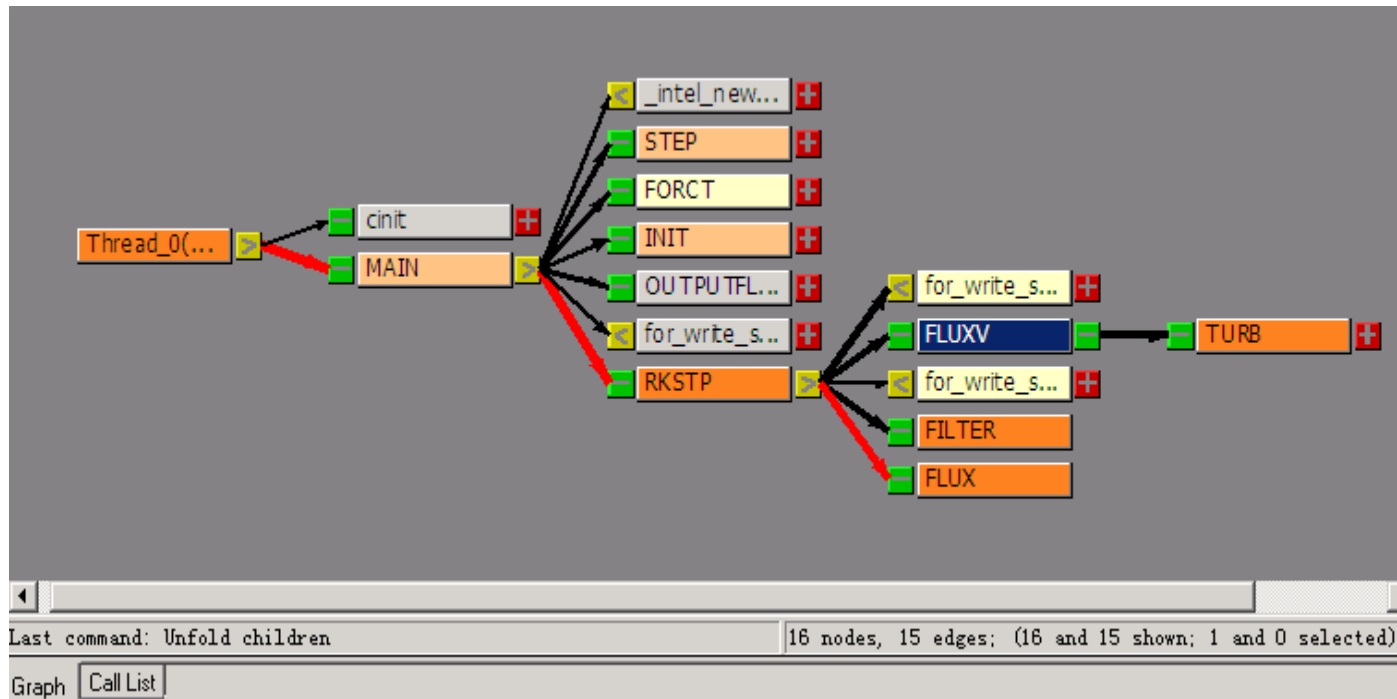


- CFD code by SJTU
- Used in COMAC
- Simu. For Wing design
- Based on RANS Equ.
- FVM on Struc. Grid

Flowchart of Sequential code



Finding hotspots of RANS Solver



Callee Function	Contribu... ▾	Edge Time	Edge Calls
RKSTP	95.7%	37,332,081	50
FORCT	1.1%	415,573	11
STEP	1.1%	415,028	10
for_write_seq_li...	1.0%	384,860	232
OUTPUTFLOW	.4%	173,576	1
INIT	.2%	61,507	1

Callee Function	Contribu... ▾	Edge Time	Edge Calls
FLUX	26.1%	9,726,010	250
FILTER	25.8%	9,616,121	100
FLUXV	23.1%	8,709,589	50
BCOND	6.2%	2,302,139	250

Parallel Approach



•Aggregation

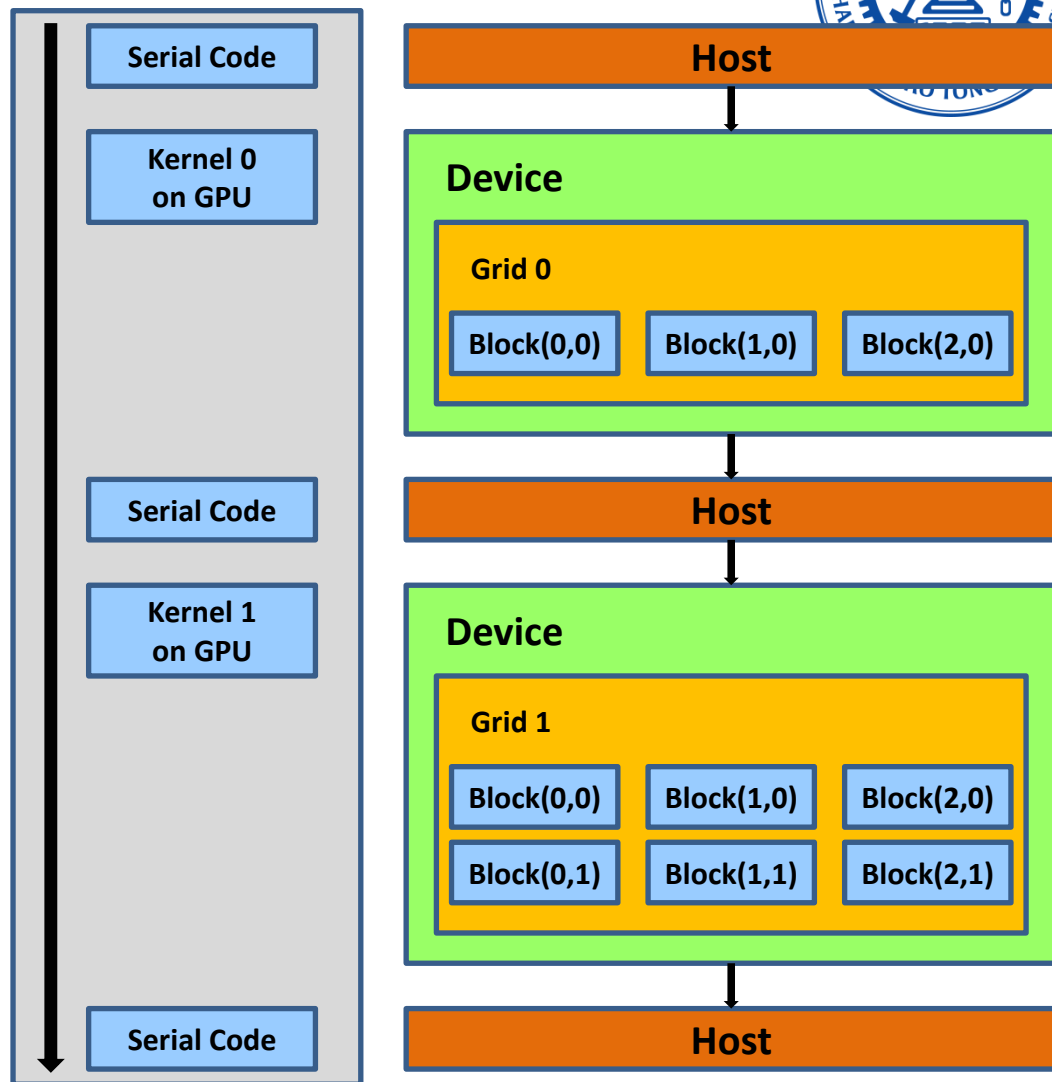
Compute: Rk inter.+
accer. Conv.

IO ops: print and write
file

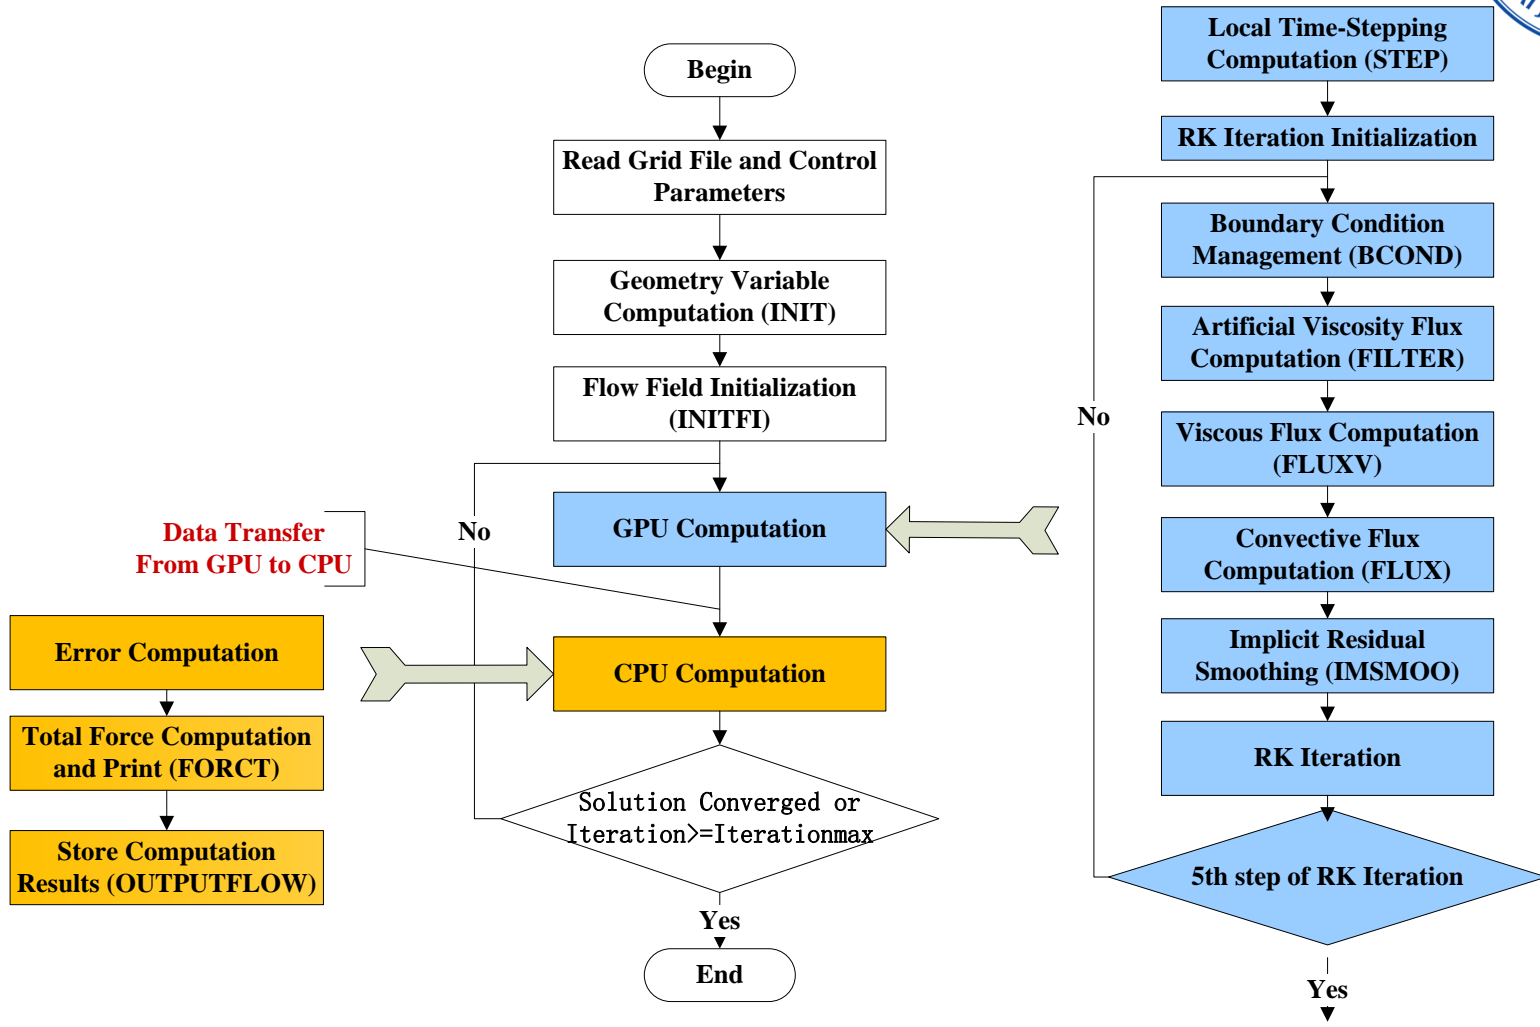
•Mapping

Compute:  GPU

IO ops:  CPU



CUDA version of Flowchart





Optimization Technique

- for better locality, transformed 1D matrices stored intermediate variables to 3D matrices
- to update data in neighbor vertexes, used multiple kernels to synchronize globally
- to eliminate data transfer between host and device memory, paralleled one direction each time in Implicit Residual Smoothing Algorithm
- for better access to global memory after matrices transformation: coalesce access.

Test Environment



Hardware

CPU: Intel(R) Core(TM) i7 CPU 920 @ 2.67GHz, Quad-Core

Cache: 32KB x 4 L1 Cache, 256 x 4 L2 Cache, 8 MB L3 Cache;

Memory: 6GB DDR3, 1033MHz;

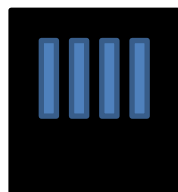
GPU: Tesla C2050 with 2.78GB global memory, 65.536KB constant memory, 49.152KB shared memory and 32768 32-bit registers per block, 14 multiprocessor, 448 CUDA cores, 1.15GHz。

Software

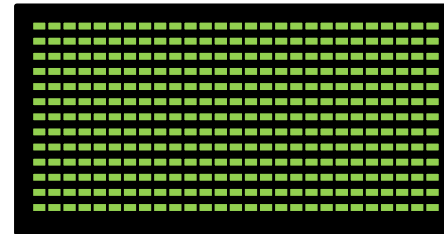
OS: Windows 7 Enterprise Edition;

Compiler: Microsoft Visual Studio 2008 C++ compiler;

CUDA: CUDA Toolkit 3.1 and SDK;



CPU
4 cores



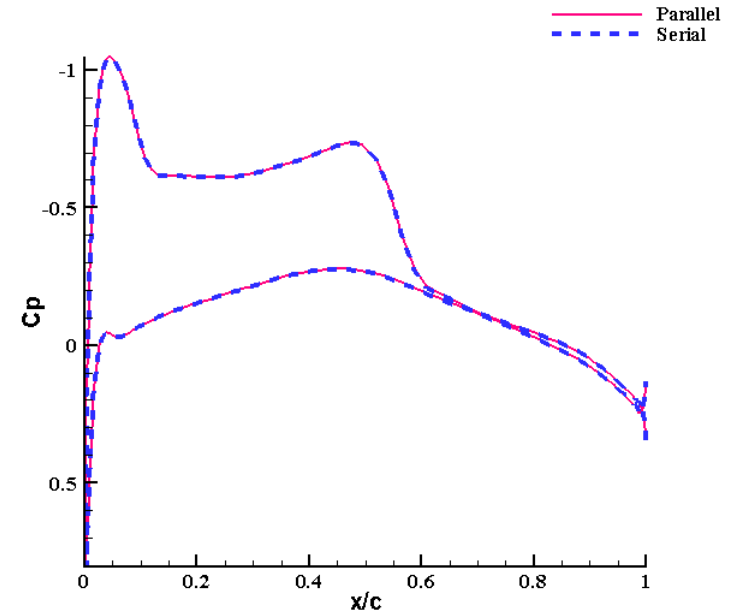
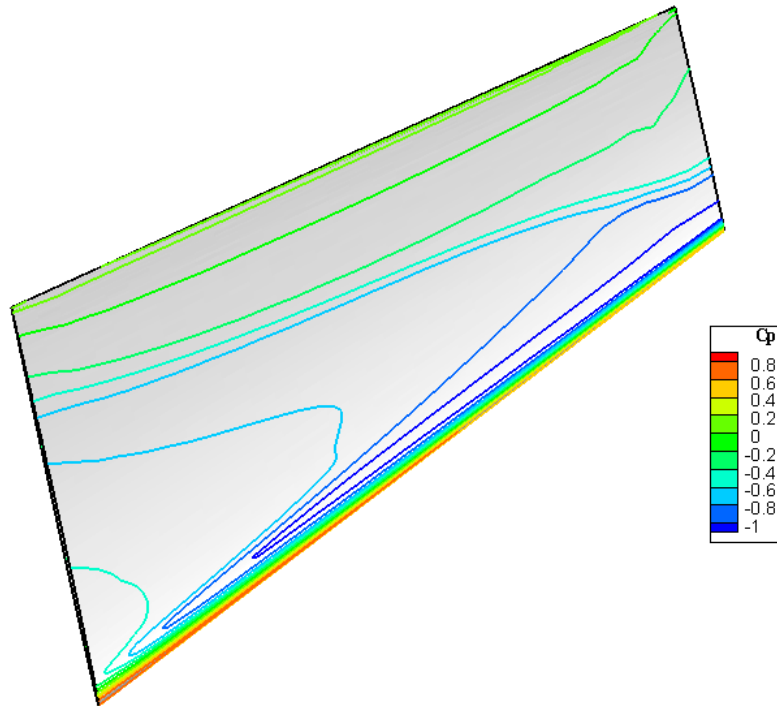
GPU
448 cores

ORENA M6 Wing 244800 (51x30x160) Mesh



Computation error

	CL	CD	CZ	CM
Serial	0.268945	0.022390	0.005948	-0.126988
CUDA	0.269663	0.022446	0.005969	-0.127724
Relative	0.3%	0.3%	0.4%	0.6%



Pressure Distribution span section of 44%

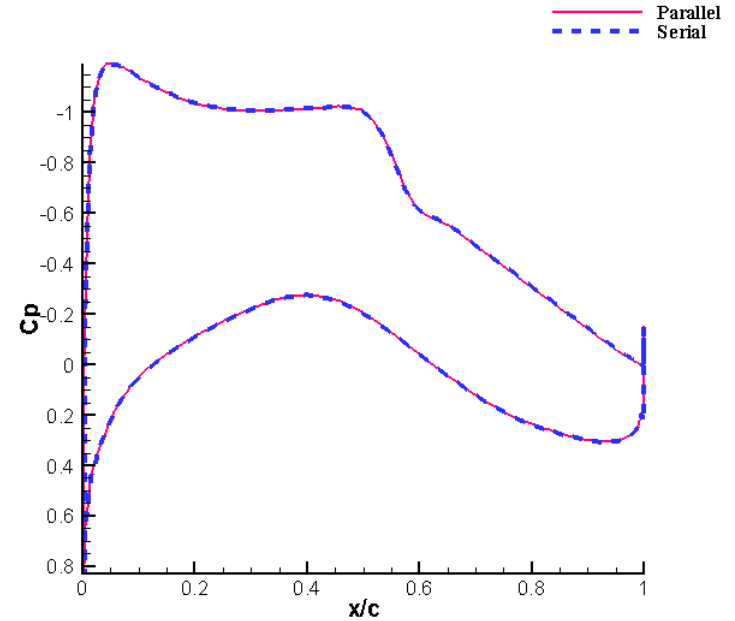
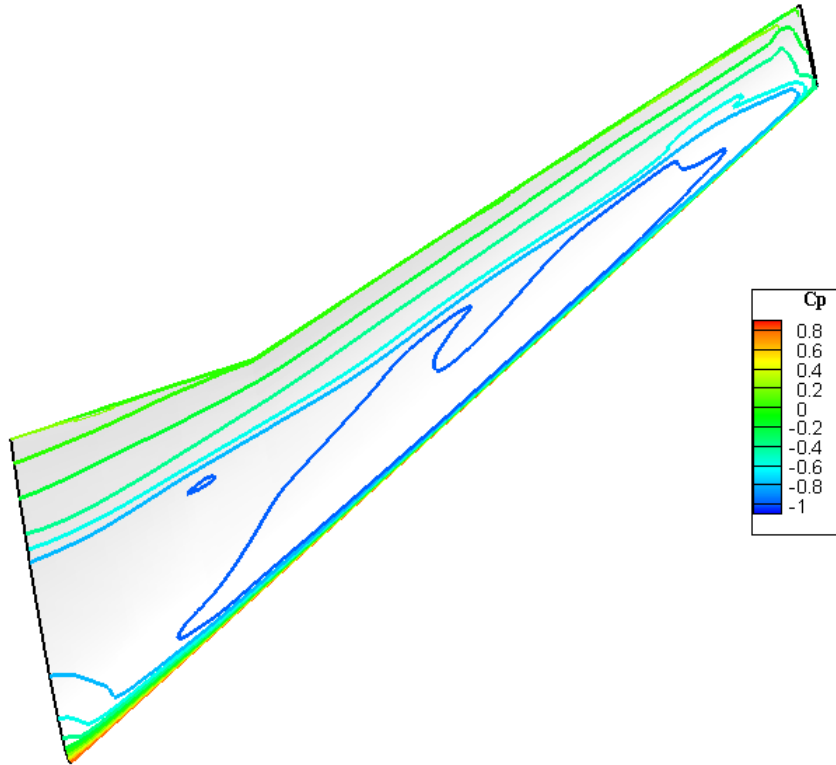
	Serial	Parallel	Speedup
OutPutFlow	15041.5s	740.0s	20.3
OutPutFlow*	14903.4s	601.9s	24.8

Another Wing 1676480 (208x65x124) Mesh, $Ma=0.785$, $AOA=2.13$, $Re=2.0E+07$



Computation error

	CL	CD	CZ	CM
Serial	0.519426	0.022840	-0.046544	1.349560
CUDA	0.518257	0.022661	-0.046310	1.346366
Relative	0.2%	0.8%	0.5%	0.2%



Pressure Distribution span section of 55%

	Serial	Parallel	Speedup
OutPutFlow	139751.5s	4652.4s	30.0
OutPutFlow*	138926.7s	3769.8s	36.9

Conclusion and Future Work



- RANS solver on GPU get 25X speedup for M6 Wing and 37X for another Wing without sacrifice in accuracy
- Next Step will scale up on GPU cluster

Acknowledgement



This research is financially supported by
Academic Partnership Program from NVIDIA



Contact Information

James Lin (林新华) lin-xh@sjtu.edu.cn

SJTU CUDA Center of Excellence
High Performance Computing Center, SJTU

***Bring Many-core Thinking to
High End Research and Education***