CAE 分野での GPU 活用のご紹介

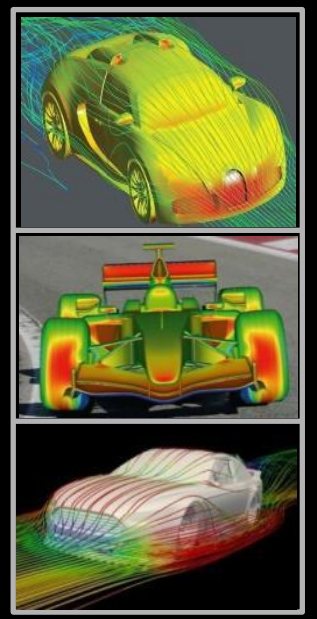

エヌビディア ジャパン
マーケティング本部 部長　林 憲一

NVIDIAについて

1993年に設立

設立以来、半導体企業の中で最速で10億ドルの収益を達成

創業者：Jen-Hsun Huang

従業員：20ヵ国に約8,500人

本社：カリフォルニア州サンタクララ


GeForce
Quadro
NVIDIA
Tegra
Tesla

CPU

+

GPU

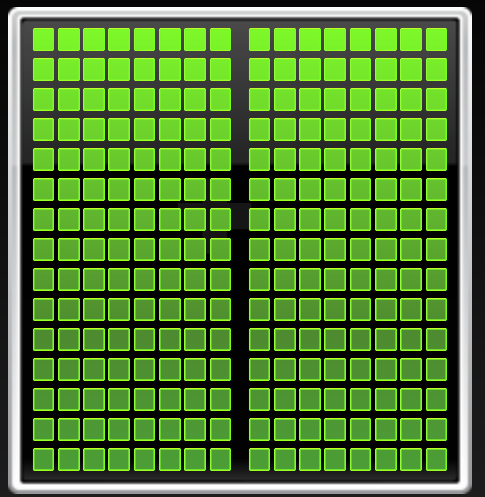

= スピードアップ

コンパニオン プロセッサ

GPUをCPUに追加することで、アプリケーションが高速化。ハイパフォーマンス・コンピューティングを実現。

GPU アクセラレーションの仕組み

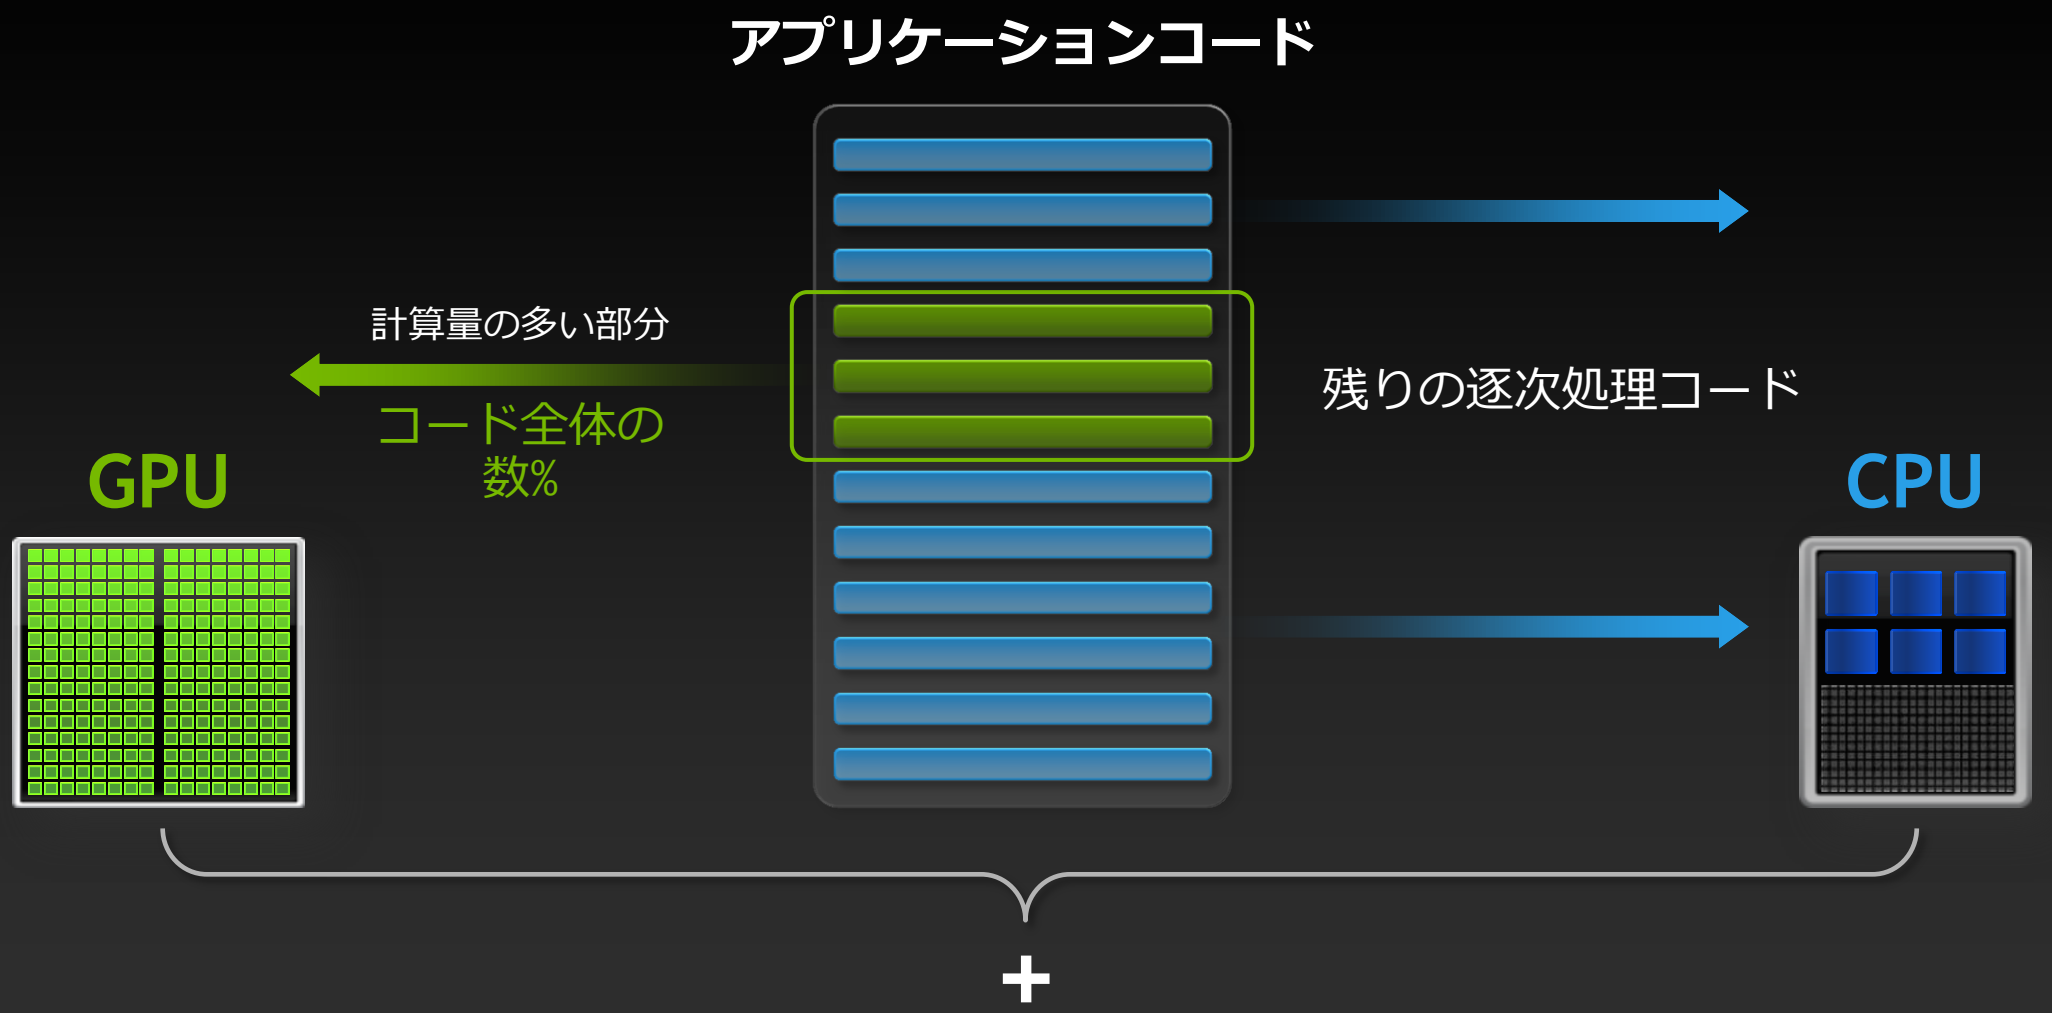

Titan: 世界最速のオープンスーパーコンピュータ

18,688 個の Tesla K20X GPU

ピーク性能: 27 ペタフロップス（性能の90%はGPU）

Linpack 性能： 17.59 ペタフロップス

500 システムのうち、50 システムが NVIDIA Tesla を利用

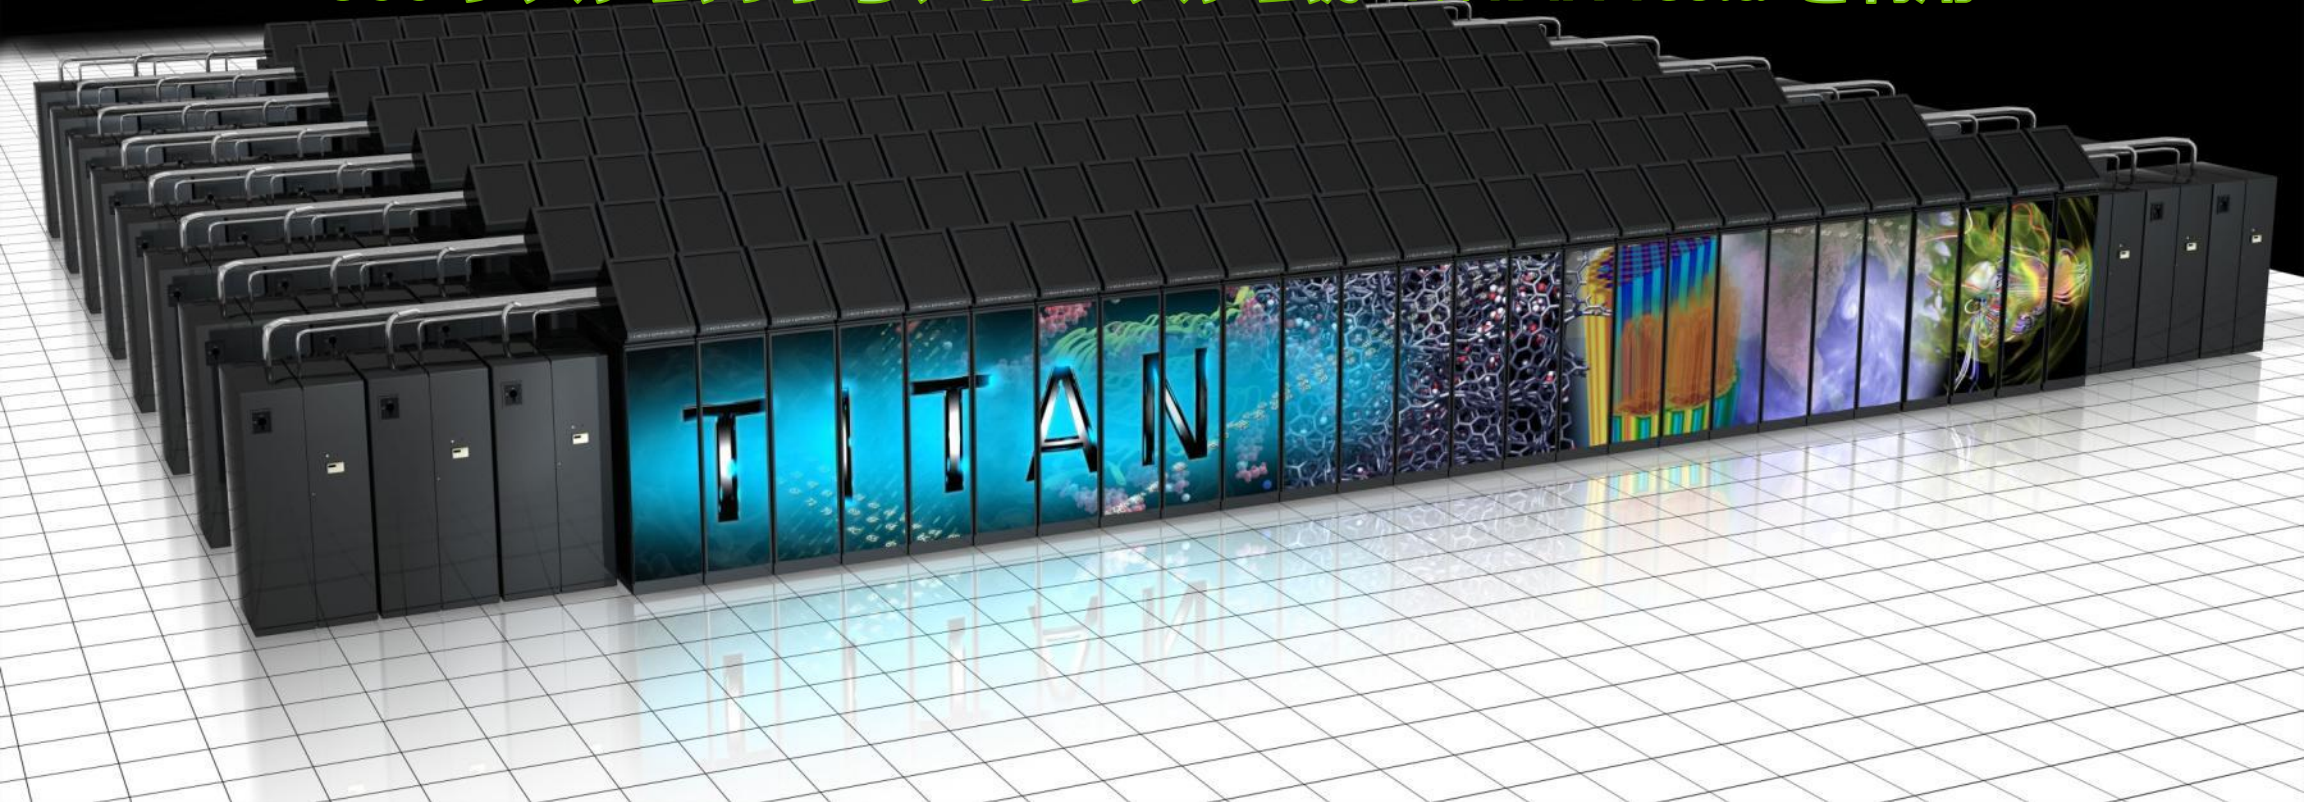

世界第21位：東工大 TSUBAME 2.0

TSUBAME 2.0
1,408 ノード
- 4,224 GPU = 2,175 TFlops
- 2,816 CPU = 216 TFlops

メモリ = 80.55 TB
SSD = 173.88 TB

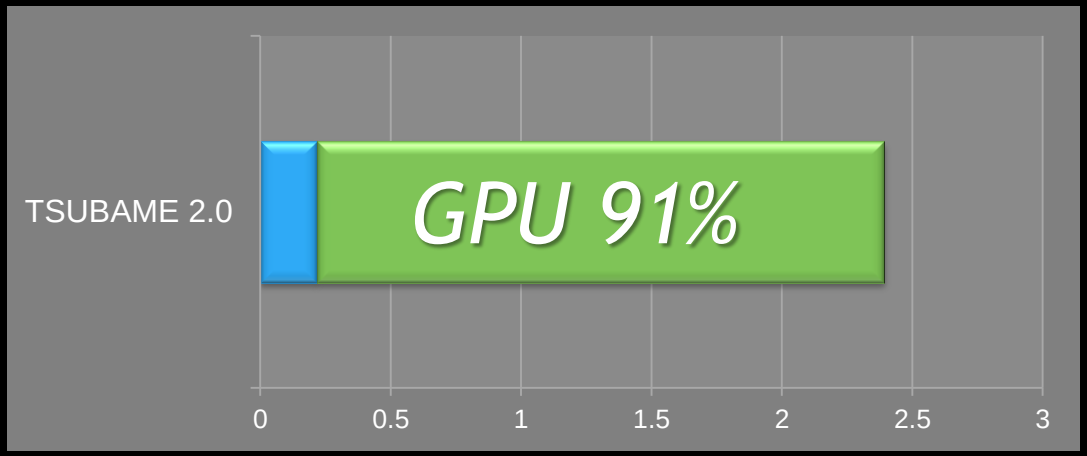

HP SL390 サーバー
3x NVIDIA Tesla M2050 GPU
2x Intel Westmere-EP CPU
52 GB DDR3 メモリ
2x 60 GB SSD
2x QDR InfiniBand

Kepler: 最速かつ電力効率の良いアーキテクチャ

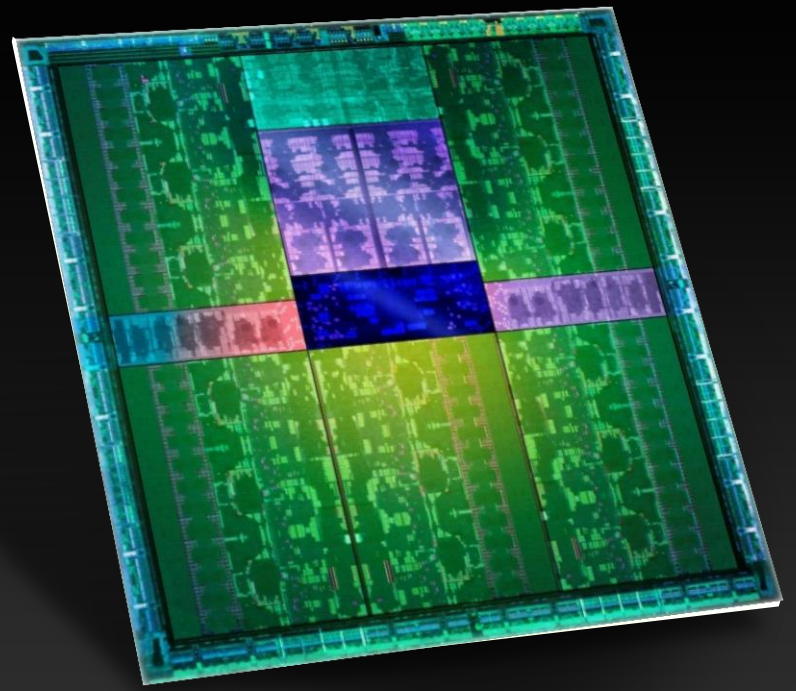

Kepler GK110 ブロックダイアグラム

アーキテクチャ

- 71億トランジスタ
- 最大 15 SMX ユニット
- 1 TFLOP以上の倍精度演算性能
- 1.5 MB L2 Cache
- 384-bit GDDR5

Tesla K20X / K20 GPU ファミリー製品

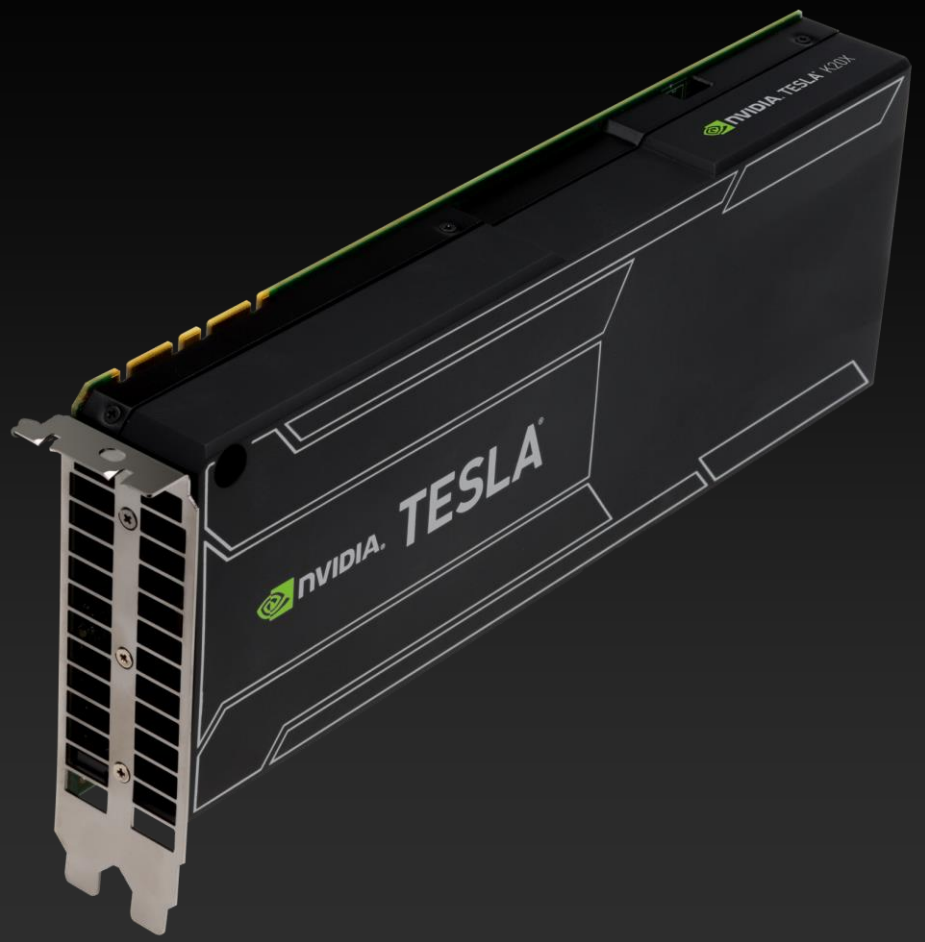

- 世界最速のアクセラレーター
- 世界最高効率
- 世界で最も普及した並列プログラミングモデル CUDA
- TITAN に 18,688 個搭載

前世代比 3倍の性能

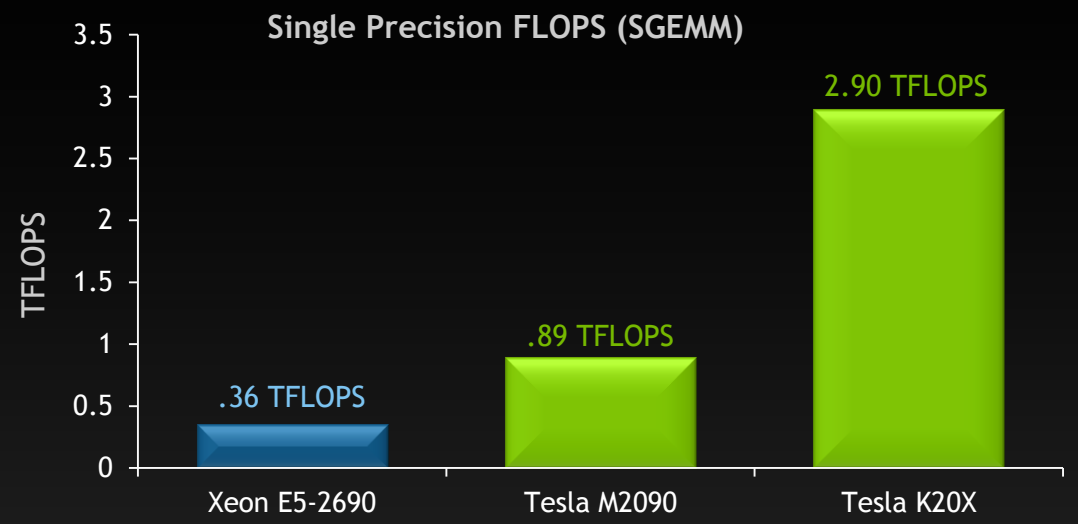

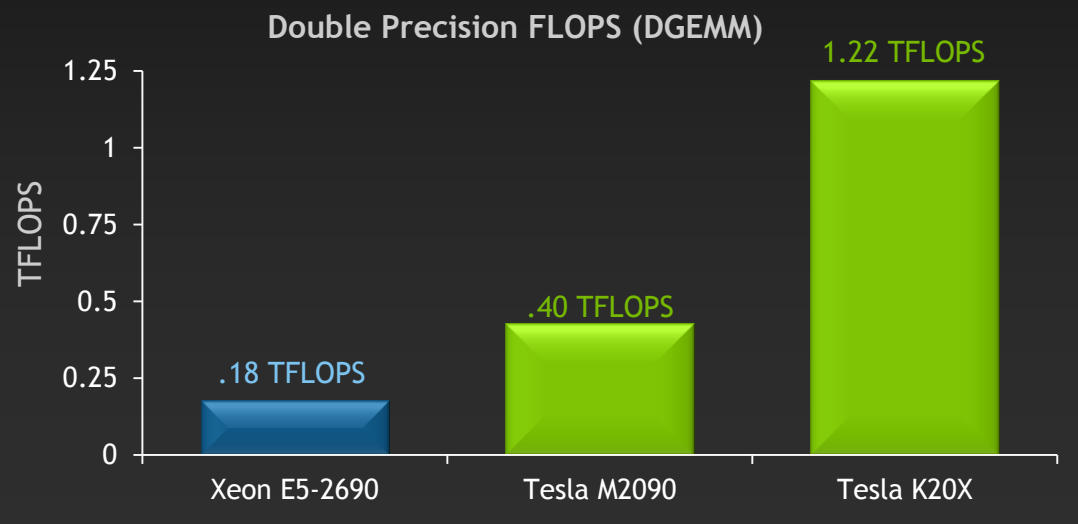

	Tesla K20X	Tesla K20
CUDA コア数	2688	2496
倍精度演算性能 DGEMM	1.32 TF 1.22 TF	1.17 TF 1.10 TF
単精度演算性能 SGEMM	3.95 TF 2.90 TF	3.52 TF 2.61 TF
メモリバンド幅	250 GB/s	208 GB/s
メモリサイズ	6 GB	5 GB
消費電力	235W	225W

Commercial CAE Software and GPU Progress

ISV	Primary Applications (Green color indicates CUDA-ready during 2013)
ANSYS	ANSYS Mechanical; ANSYS Fluent; ANSYS HFSS
DS SIMULIA	Abaqus/Standard; Abaqus/Explicit; Abaqus/CFD
MSC Software	MSC Nastran; Marc; Adams
Altair	RADIOSS; AcuSolve
CD-adapco	STAR-CD; STAR-CCM+
Autodesk	AS Mechanical, Moldflow, AS CFD
ESI Group	PAM-CRASH *imp*; CFD-ACE+
Siemens	NX Nastran
LSTC	LS-DYNA; LS-DYNA CFD
Mentor	FloEFD, FloTherm
Metacomp	CFD++

Other Commercial CAE and GPU Progress

ISV	Domain	Location	Primary Applications
FluiDyna	CFD	Germany	Culises for OpenFOAM; LBultra
Vratis	CFD	Poland	Speed-IT for OpenFOAM; ARAEL
Prometech	CFD	Japan	Particleworks
Turbostream	CFD	England, UK	Turbostream
IMPETUS	Explicit FEA	Sweden	AFEA
AVL	CFD	Austria	FIRE
CoreTech	CFD (molding)	Taiwan	Moldex3D
Intes	Implicit FEA	Germany	PERMAS
Next Limit	CFD	Spain	XFlow
CPFD	CFD	USA	BARRACUDA
Flow Science	CFD	USA	FLOW-3D

2013: Further Expansion of OF Community

- **ESI acquisition of OpenCFD from SGI during Sep 2012**
- **IDAJ investment in ICON (migration from CD-adapco)**
- **This Year 3 Global OpenFOAM User Conferences:**
  - **APR 24 – 26, Frankfurt, DE: ESI OpenFOAM Users Conference (first ever)**
    http://www.esi-group.com/corporate/events/2013/OpenFOAM2013
    Concentration on OpenFOAM from OpenCFD
  - **JUN 11 – 14, Jeju, KR : 8th International OpenFOAM Workshop (first in Asia)**
    http://www.openfoamworkshop2013.org/
    Concentration on OpenFOAM-extend and Wikki
  - **OCT 24 – 25, Hamburg, DE : 7th Open Source CFD International Conference (ICON)**
    http://www.opensourcecfd.com/conference2013/
    Concentration on both OpenFOAM and OpenFOAM-extend

NVIDIA Market Strategy for OpenFOAM

- **Provide technical support for commercial GPU solver developments**
  - FluiDyna Culises library with NVIDIA collaboration on AMG
  - Vratis Speed-IT library, development of CUSP-based AMG
- **Invest in alliances (but not development) with key OpenFOAM organizations**
  - ESI and OpenCFD Foundation (H. Weller, M. Salari)
  - Wikki and OpenFOAM-extend community (H. Jasak)
  - IDAJ in Japan and ICON in the UK – support for both OF and OF-ext
- **Conduct performance studies and customer benchmark evaluations**
  - Collaborations: developers, customers, OEMs (Dell, SGI, HP, etc.)

OpenFOAM Applied Use: Parameter Optimization

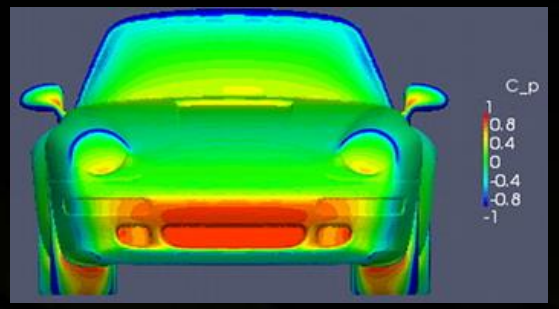

#1: Develop validated CFD model in ANSYS Fluent or other commercial CFD software in production

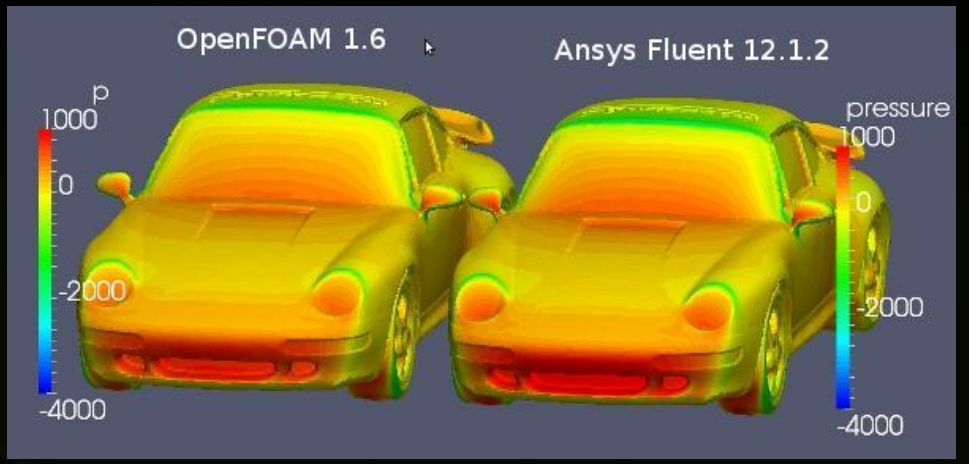

#2: Develop CFD model in OpenFOAM, validate against commercial CFD model

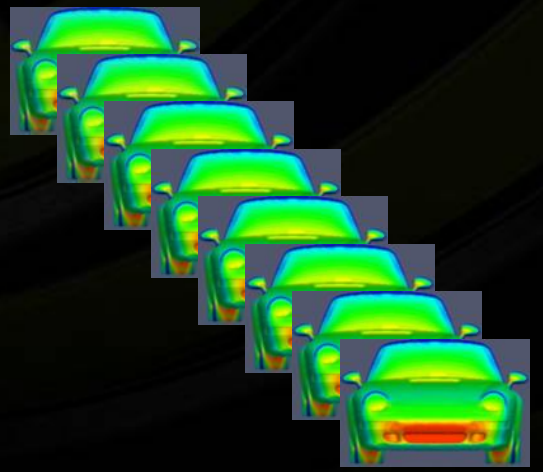

#3: Conduct parameter sweeps with OpenFOAM model to save on commercial CFD license costs

GPU Opportunity for Parameter Optimization

Problem Statement:

Demand for optimization – can existing CPU clusters manage 10x more jobs?

Examples:

Automotive crashworthiness optimization

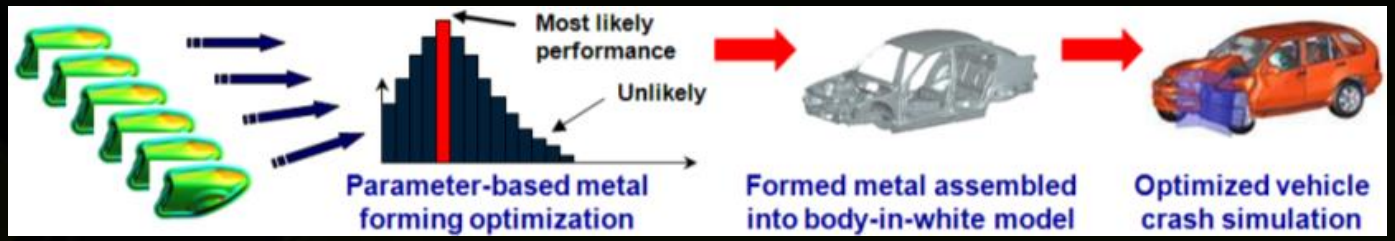

Jet engine CFD aerodynamics optimization

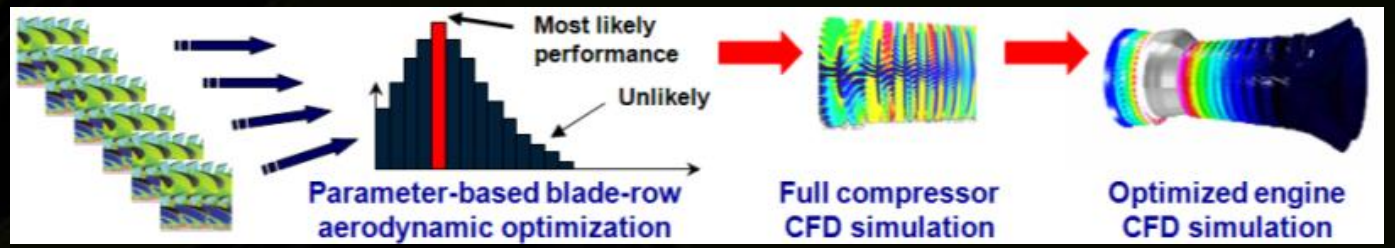

GPU Opportunity:

- Open source and proprietary not bounded by commercial CAE license costs
- ISV "optimization" licensing solved – ANSYS, Altair, SIMULIA, etc. – hardware problem next
  - GPUs: performance under smaller footprint with better power and cooling efficiency

OpenFOAM GPU Focus on Implicit Sparse Solvers

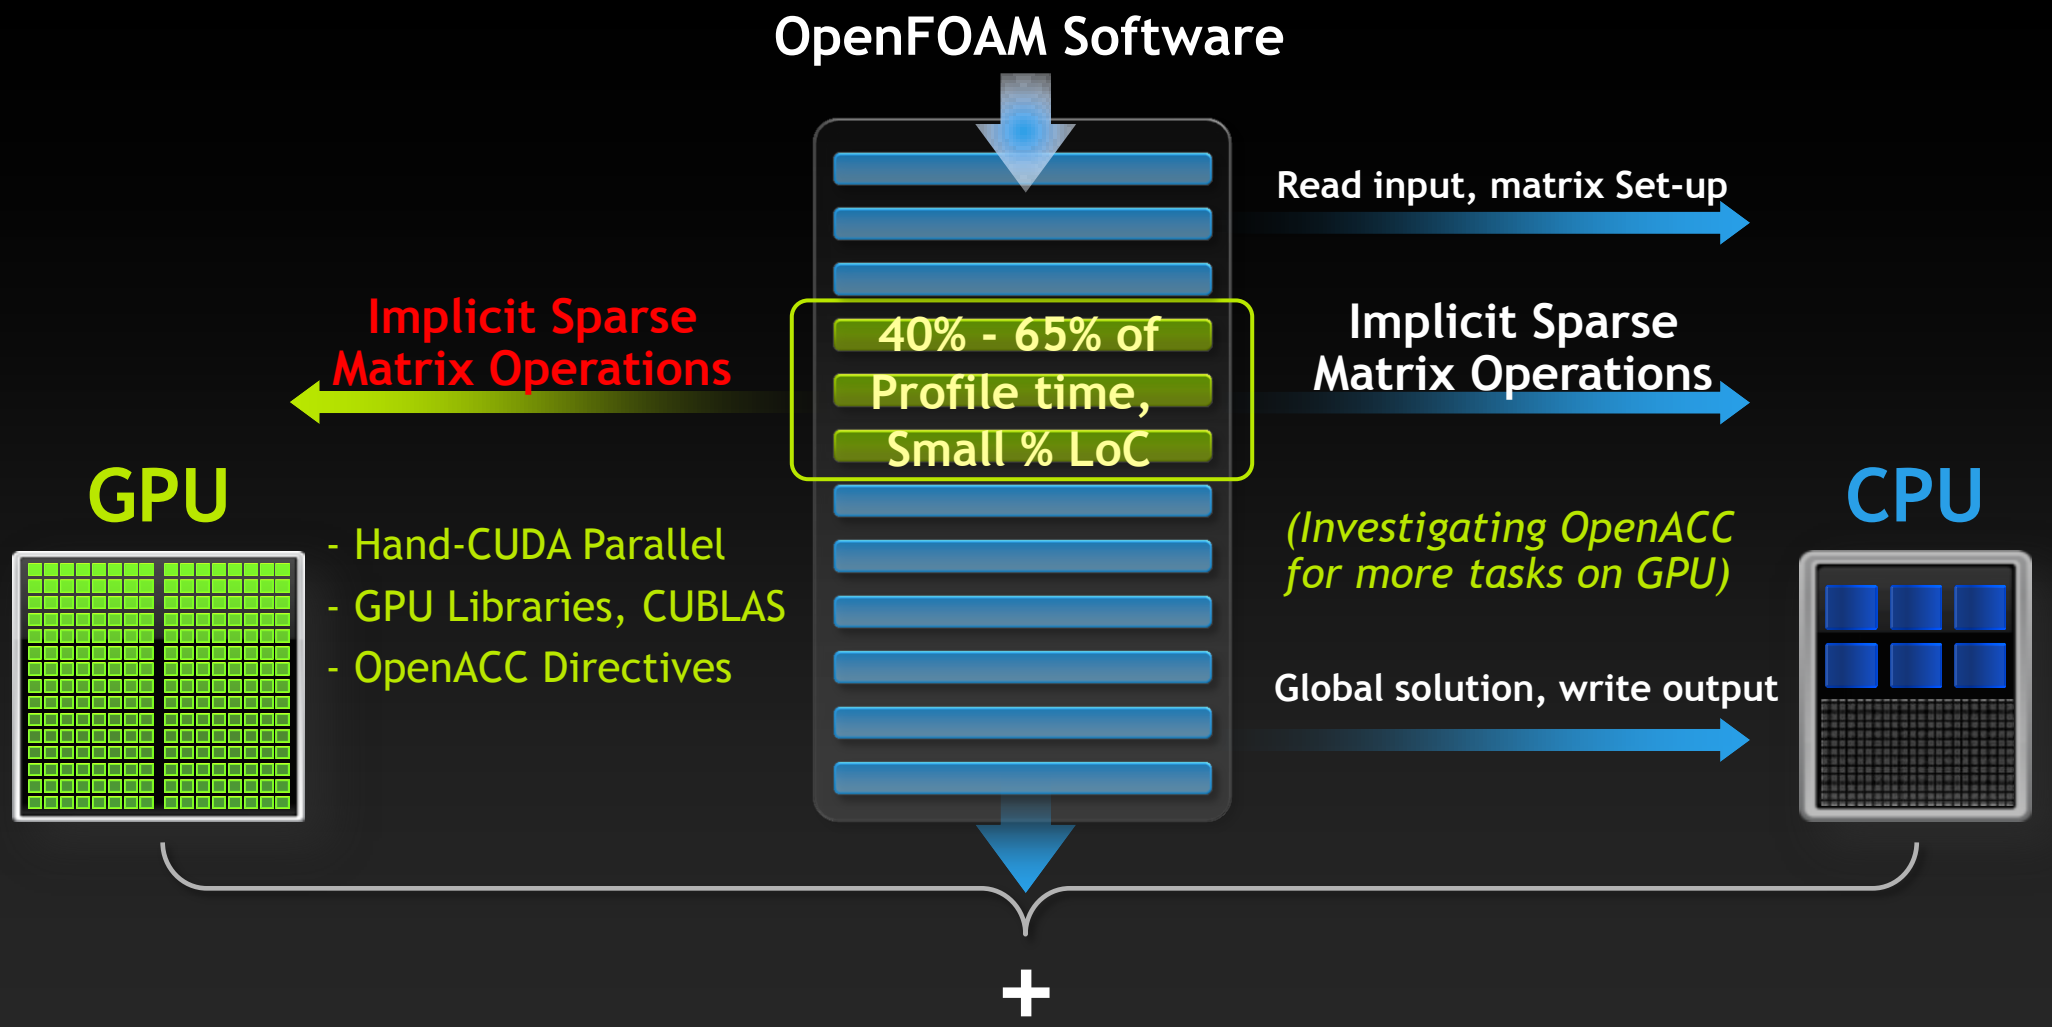

ANSYS Fluent Preview for 2 x CPU + 2 x Tesla K20X

ANSYS Fluent 15.0 Preview Performance – Results by NVIDIA, Feb 2013

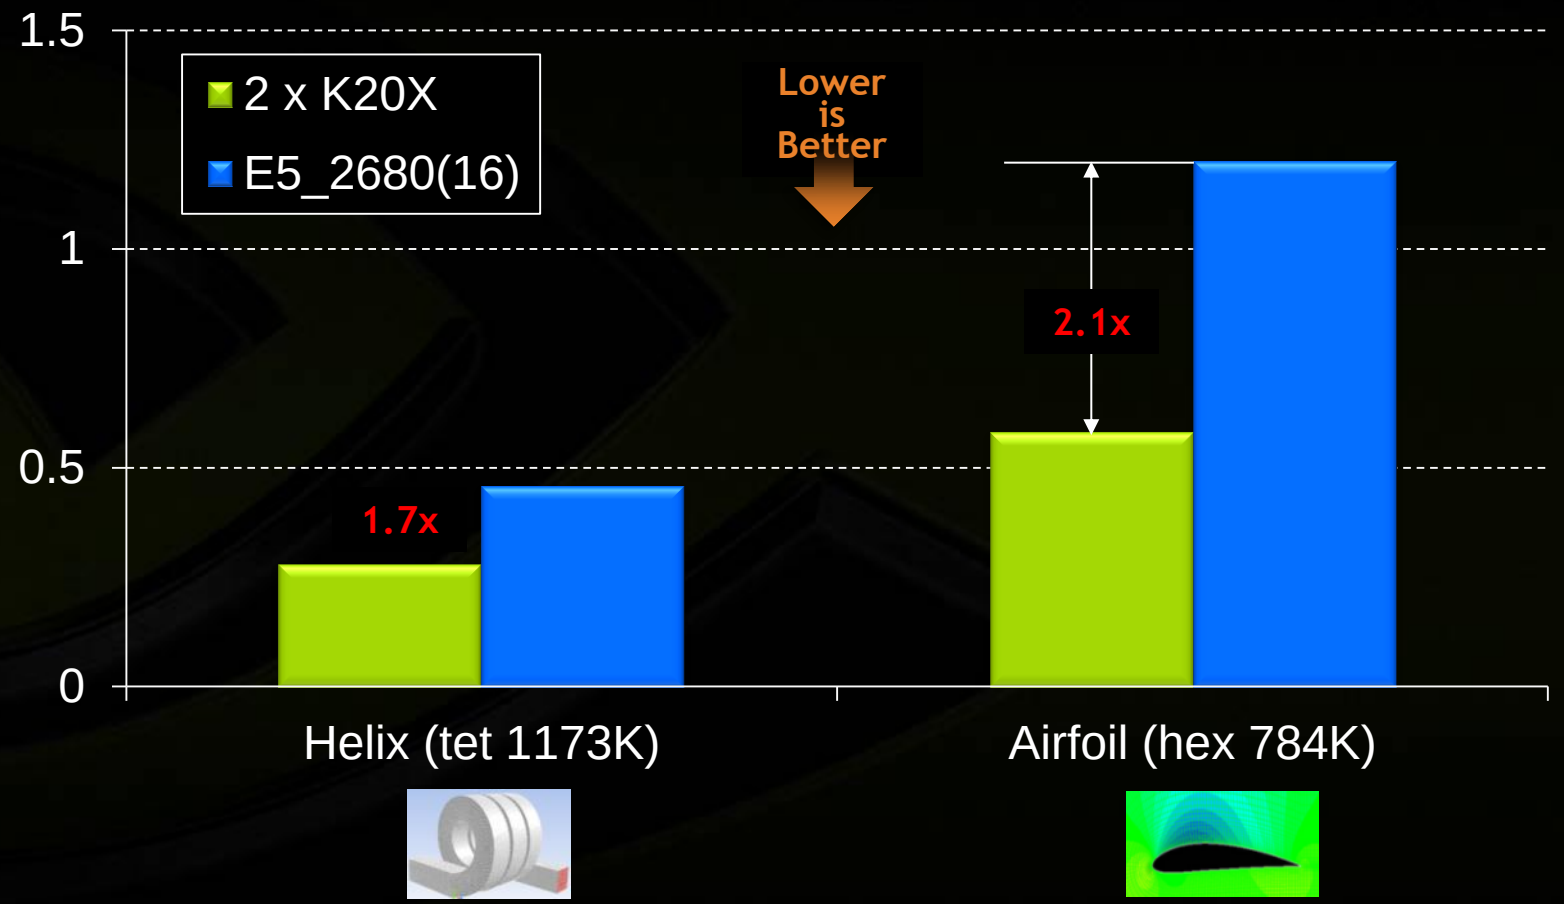

2 x E5_2680 SB CPUs, 16 cores total, only 2 cores used with GPUs

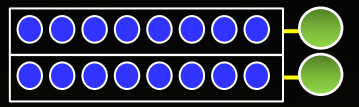

Solver settings:

CPU Fluent solver:
F-cycle, agg8, DILU, 0pre, 3post

GPU nvAMG solver:
V-cycle, agg8, MC-DILU, 0pre, 3post

NOTE: Times for solver only

Culises: New CFD Solver Library for OpenFOAM

www.fluidyna.de

Culises Features:

- State-of-the-art solvers for solution of linear systems
  - Multi-GPU capability enables HPC computing
  - Available solvers:
    - Krylov subspace methods
      - Conjugate or Bi-Conjugate Gradient method for symmetric or non-symmetric system matrices
      - Several preconditioning options for enhanced performance
    - Stand-alone Multigrid method
- Flexible interfaces for arbitrary applications
  e.g.: established coupling with OpenFOAM® CFD package

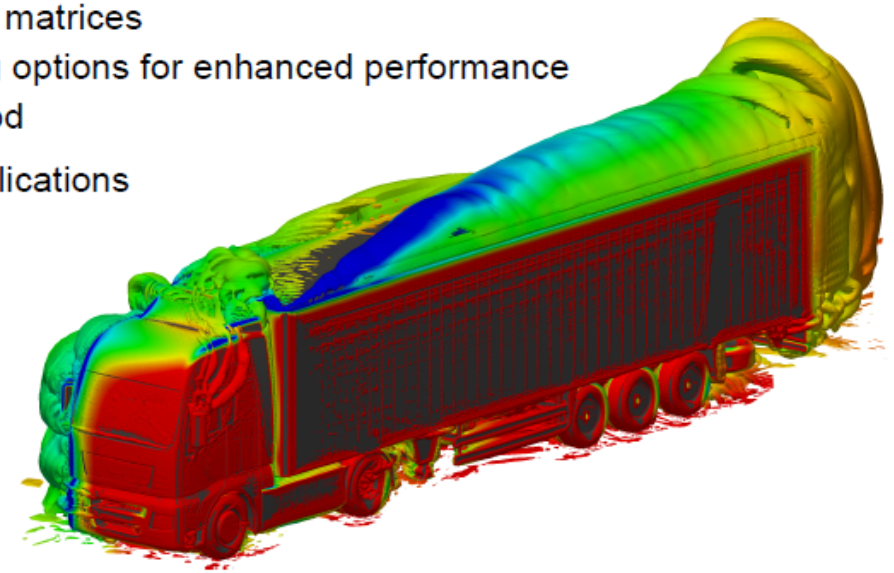

FluiDyna: TU Munich Spin-Off from 2006

Culises provides a linear solver library

Culises requires only two edits to control file of OpenFOAM

Multi-GPU ready

Contact FluiDyna for license details

Culises Coupling to OpenFOAM

Culises Coupling is User-Transparent:

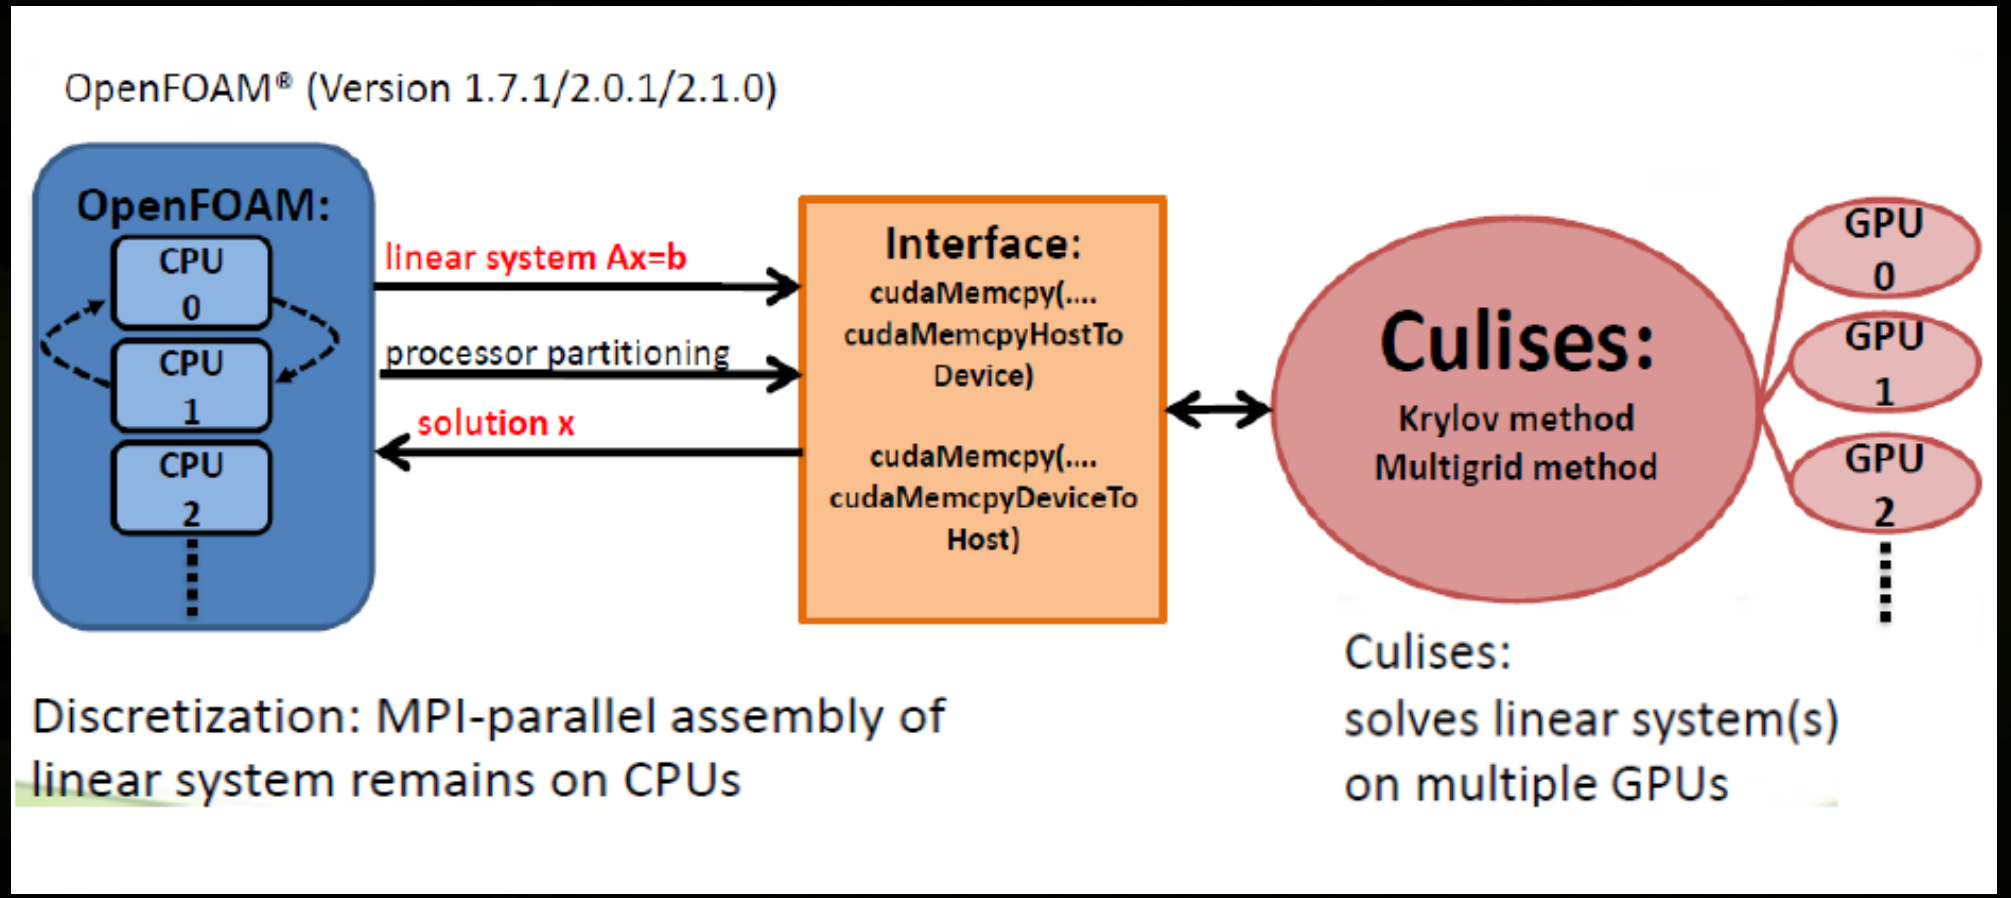

Culises: New CFD Solver Library for OpenFOAM

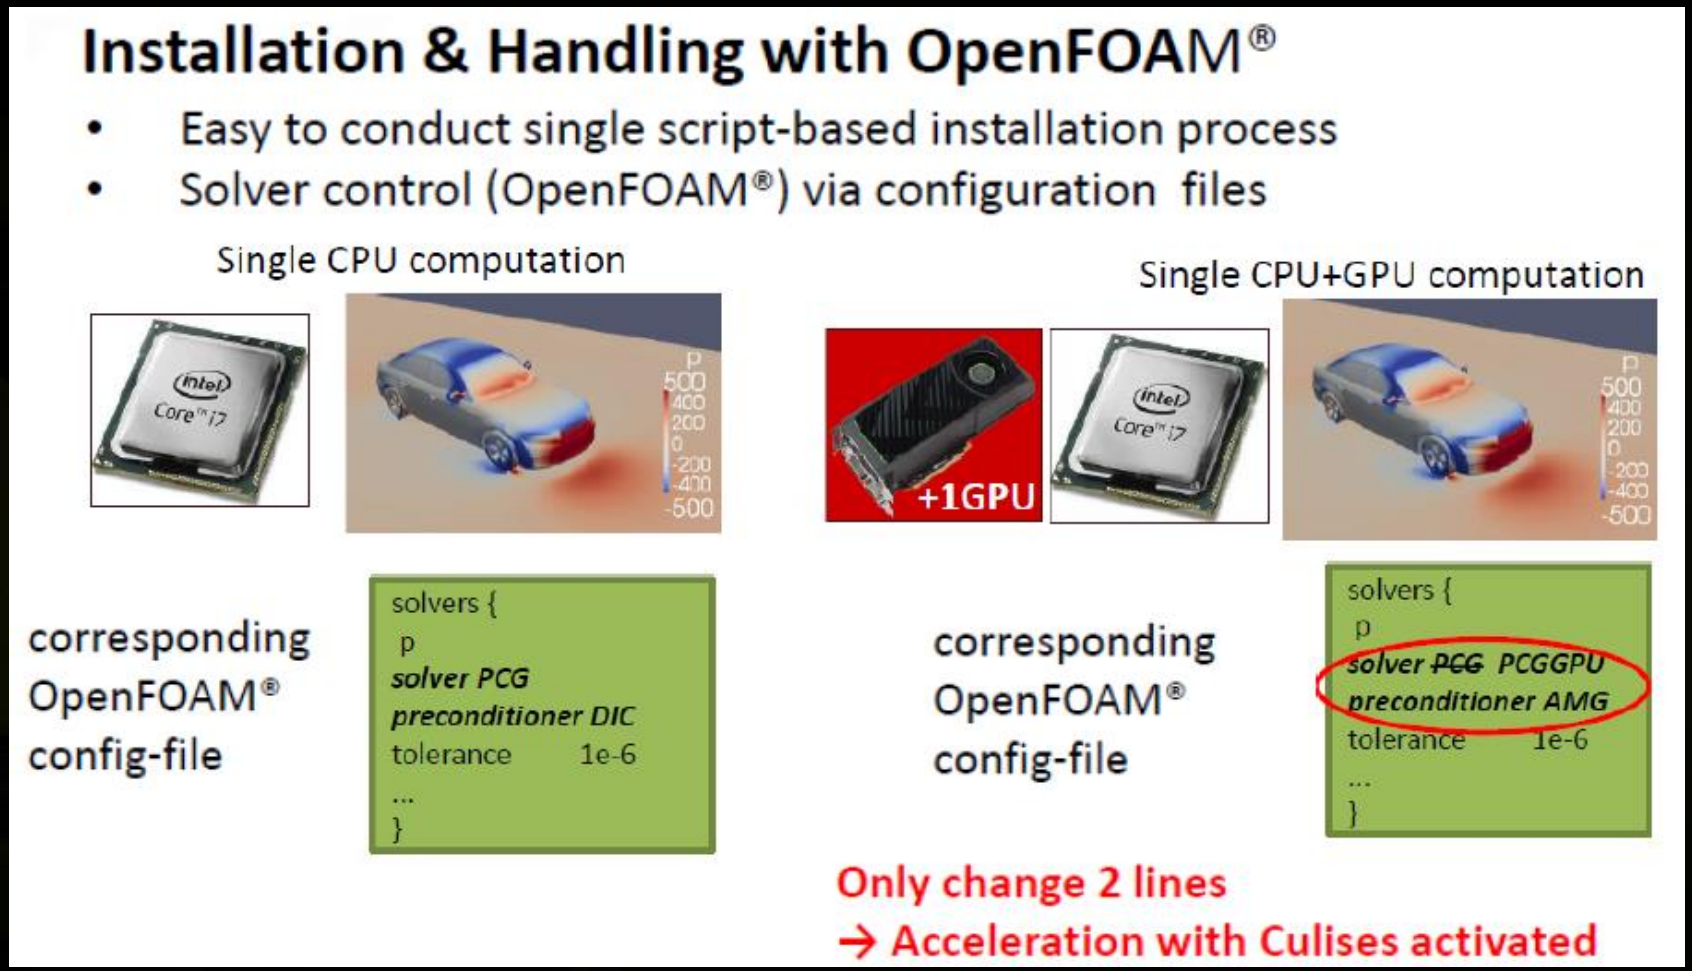

www.fluidyna.de

Easy-to-Use

#1. Download and license from www.FluiDyna.de

#2. Install with script provided by FluiDyna

#3. Activate Culises and use of GPUs with 2 simple changes to OF config-file

OpenFOAM GPU Speedups Based on Application

Speedups for a Range of Industrial Cases:

www.fluidyna.de

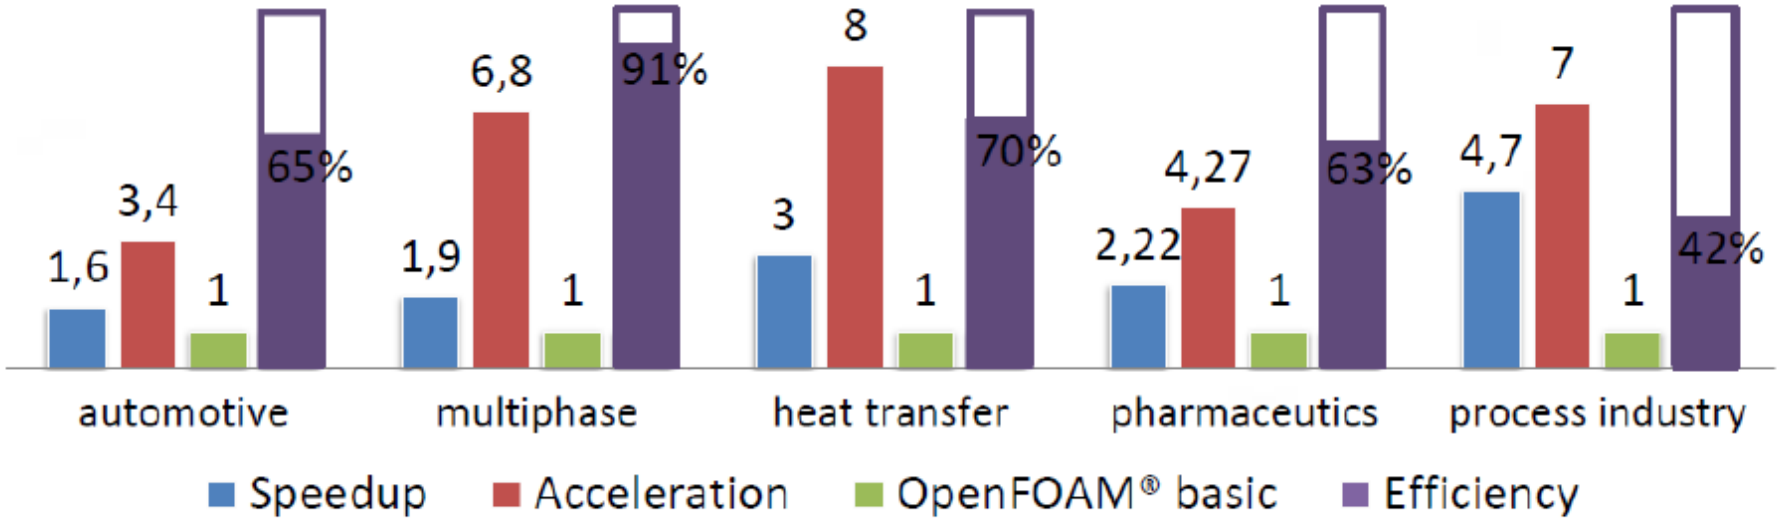

OpenFOAM® basic	CPU only reference run
Speedup	Speedup of total simulation compared to OpenFOAM basic run
Acceleration	Speedup of linear solver compared to OpenFOAM basic run
Efficiency	Ratio of obtained total speedup and theoretical maximum speedup

FluiDyna GmbH
Lichtenbergstraße 8
D-85748 Garching b. München
www.fluidyna.com

Dr. Bjoern Landmann

Accelerating the Numerical Simulation of Heavy-Vehicle Aerodynamics Using GPUs with Culises

ISC 2013, Leipzig , June 2013

Multi-GPU runs

- Speedup by adding multiple GPUs:

(a) single-socket board

Mesh - # CPUs	**9M - 1 CPU**	18M - 1 CPU	27M - 1 CPU	36M - 1 CPU
# GPUs added	**+1 GPU**	**+2 GPUs**	**+3 GPUs**	**+4 GPUs**
Speedup linear solver a	**3.5**	5.7	7.8	10.6
Speedup total simulation	**1.45**	1.59	1.67	1.74
Theoretical max speedup s_{max}	**1.78**	1.82	1.85	1.89

(b) dual-socket board

Mesh - # CPUs	9M - 2 CPU	**18M - 2 CPU**	27M - 2 CPU	36M - 2 CPU
# GPUs added	**+1 GPU**	**+2 GPUs**	**+3 GPUs**	**+4 GPUs**
Speedup linear solver a	2.5	**4.2**	6.2	6.9
Speedup total simulation	1.36	**1.52**	1.63	1.67
Theoretical max speedup s_{max}	1.78	**1.82**	1.85	1.89

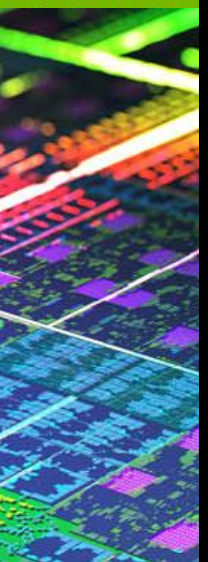

GTC Japan 2013

- イベント名: GTC Japan 2013
- 主催: エヌビディア ジャパン
- 日時: 2013年 7月30日（火）10:00 - 18:30
- 会場: 東京ミッドタウンホール
- 参加費: 無料
- イベントサイト:

http://www.gputechconf.jp


Thank you