FUJITSU

shaping tomorrow with you

14. LS-DYNA Forum, Oktober 2016, Bamberg

Erkenntnisse aus aktuellen Performance-Messungen mit LS-DYNA

Eric Schnepf[1], Dr. Eckardt Kehl[1], Chih-Song Kuo[2], Dymitrios Kyranas[2]

[1]Fujitsu Technology Solutions GmbH, München

[2]ict information communication technology GmbH, Aachen

FUJITSU

Agenda

- Fujitsu HPC Portfolio
- New PRIMERGY CX600 M1 with Intel® Xeon Phi™ 7200
- Performance Measurements with LS-DYNA
- Early Experiences with LS-DYNA on Intel® Xeon Phi™ 7200
- Best Practices

PRIMEFLEX for HPC
Integrated HPC Solutions

Application Appliance	Reference Configuration
Integration & Support services	
Assembly, Test & Delivery	
User workplace	HPC Gateway HPC Gateway Add-ons
Management software	Batch / Operation / Administration
System design	Head node / Compute nodes / Storage Interconnect / Rack & Power / Graphics

HPC Simplicity & Expertise

Select Your Preferred Hardware Platform

CELSIUS

Workgroup
Departmental

PRIMERGY
Rack Server

Capacity

Scalability,
Compute density

PRIMERGY
Scale-out Server

Data Center

PRIMERGY
Blade Server

Data Center

Scalability,
Infrastructure
density

Capability

Flexibility to address all kinds of customer requirements

- PRIMERGY CX400 / CX600 skinless server
  - Massive scale-out due ultra dense server
  - GPU coprocessor support
- PRIMERGY blade server
  - Industry leading blade server density
- PRIMERGY rack server
- CELSIUS workstations

Fujitsu PRIMERGY CX2550 M2

Feature Overview

- Highest performance & density
  - Condensed half-width-1U server
  - Up to 4x CX2550 M2 into a CX400 M1 2U chassis
  - Intel Xeon E5-2600 v4 product family, 16 DIMMs per server node with up to 1,024 GB DDR4 memory
- High reliability & low complexity
  - Variable local storage: 6x 2.5” drives per node, 24 in total
  - Support for up to 2x PCIe SSD per node for fast caching
  - Hot-plug functionality for server nodes, power supplies and disk drives enables enhanced availability and easy serviceability

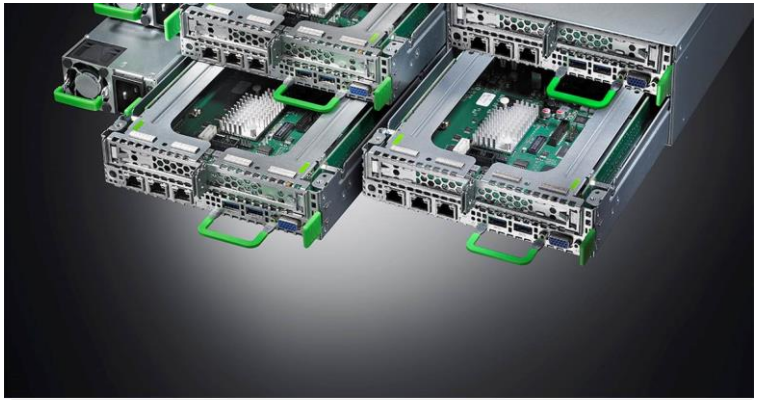

Usage Scenarios

Scale-out cloud infrastructure

HPC server farms

Intel® Xeon® E5-2600 v4 processors

Best combination of performance, built-in capabilities and cost-effectiveness

1.5 TB RAM with 2,400 MHz speed (DDR4)	»	Accelerate ‘in-memory’ solutions
14 nm manufacturing process	»	Improved energy efficiency
Up to 22 cores and 55 MB cache	»	Boost your general computing performance by about 20%
Intel® Turbo Boost Technology 2.0	»	Performance that adapts to spikes in your workload
Intel® Advanced Vector Extensions 2 (AVX2)	»	Achieve higher throughput in technical computing

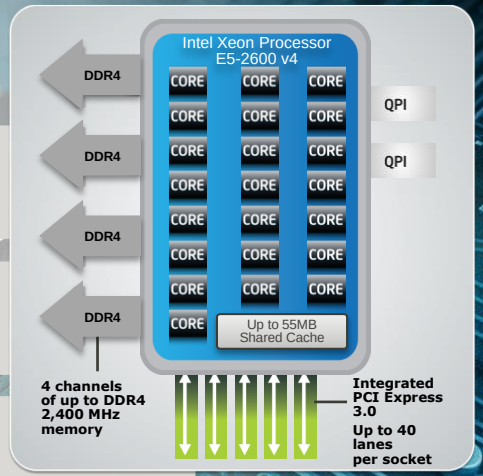

Intel® Xeon Phi™ 7200 Processor
A Highly-Parallel CPU

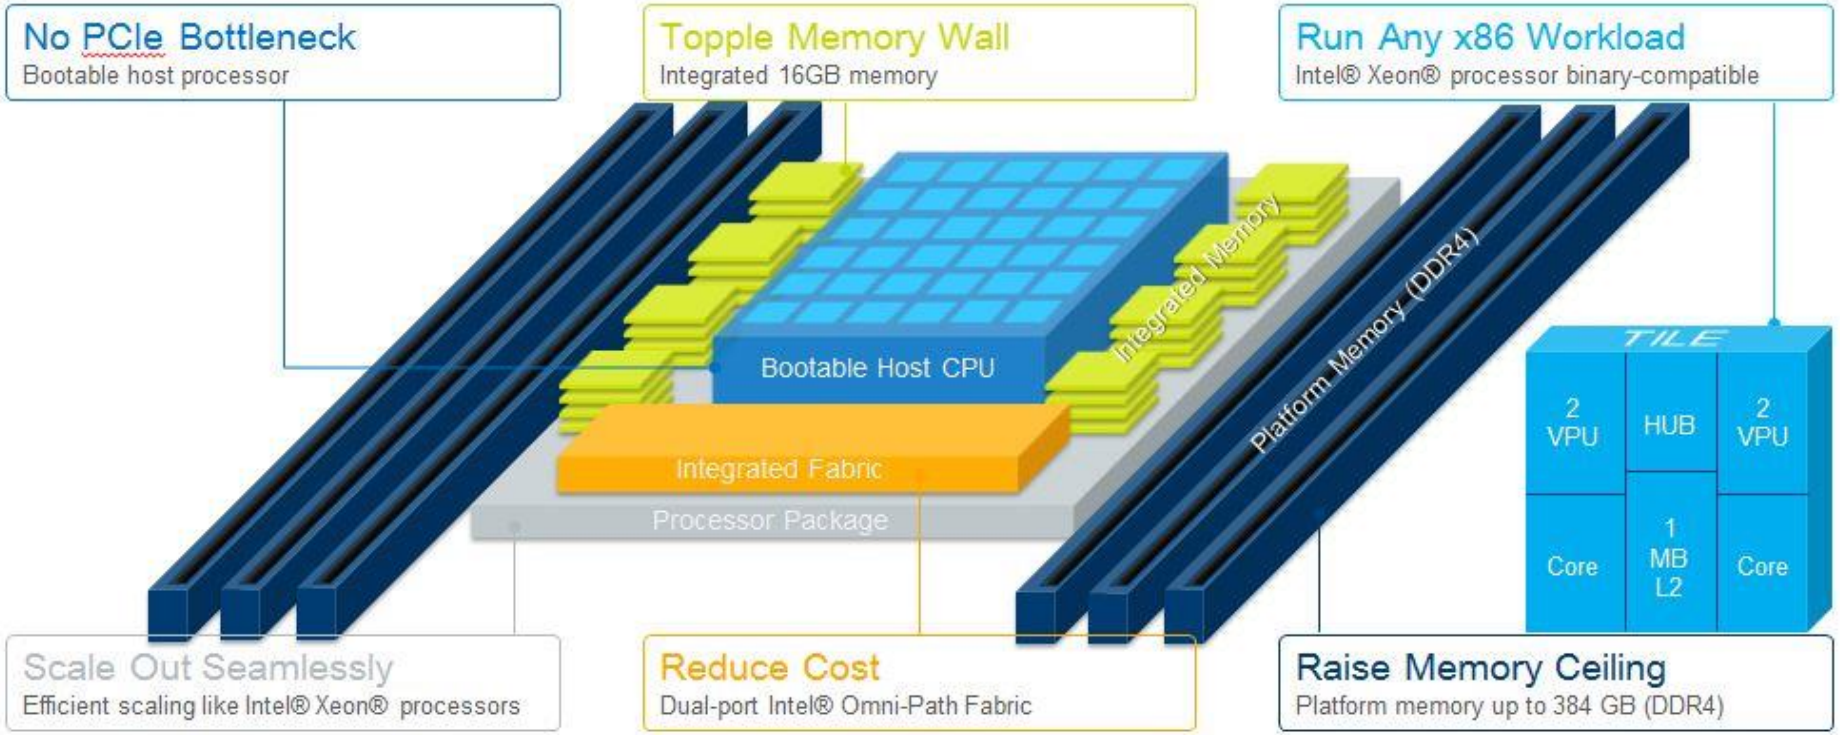

[1]Reduced cost based on Intel internal estimate comparing cost of discrete networking components with the integrated fabric solution

Fujitsu PRIMERGY CX600 M1

The new platform for highly parallel computing

The FUJITSU Server PRIMERGY CX600 M1 is the perfect choice for highly parallel applications in the area of scientific research, product development and business intelligence.

Up to eight server nodes per 2U and Intel® Xeon Phi™ Processors make for new levels of compute density.

Performance Measurement Environment

FUJITSU

	Intel® Xeon Platform	Intel® Xeon Phi™ Platform
Chassis	Fujitsu PRIMERGY CX400	Fujitsu PRIMERGY CX600
Server	Fujitsu PRIMERGY CX2550M2	Fujitsu PRIMERGY CX1640
CPU	2x Intel Xeon E5-2660v4/2690v4	1x Intel Xeon Phi 7210
Memory	192GB	192GB
Storage	SATA DOM 64GB	SATA DOM 64GB
Max Power	429 Watt	270 Watt

Models

FUJITSU

3 cars	0.8M elements	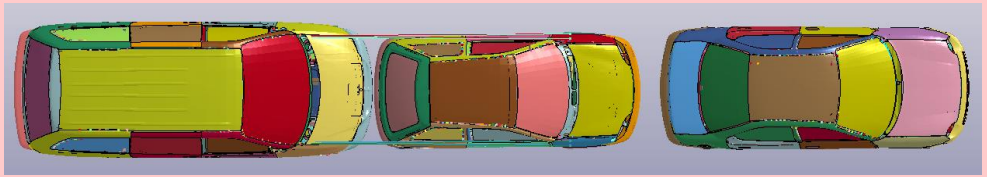
Car2car	2.5M elements	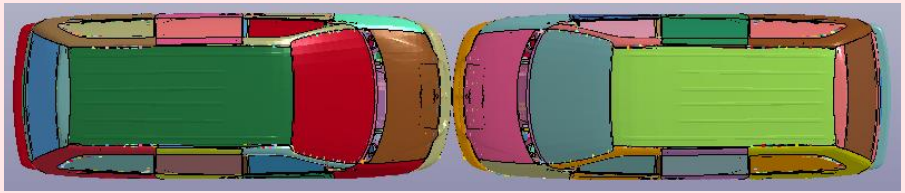
ODM-10M	10M elements	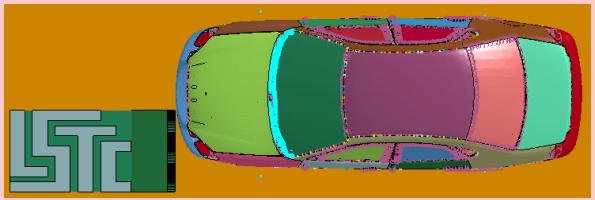

Downloaded from http://www.topcrunch.org/benchmark_problems.sfe

Pure Performance Comparison

FUJITSU

All results on Xeon Phi are preliminary and are expected to improve with the official V9.0 release

Please contact the author for actual information

∴ 2-socket Xeon node provides slightly higher performance than 1-socket Xeon-Phi node

*Performance is defined to be the reciprocal of simulation time, i.e. excluding preprocessing time

Price Performance Comparison

FUJITSU

All results on Xeon Phi are preliminary and are expected to improve with the official V9.0 release

Please contact the author for actual information

∴ Depending on the model, Xeon Phi can be more price-performance attractive than Xeon

*Performance is defined to be the reciprocal of simulation time; Fujitsu internal cost calculation is used.
*Price Performance = Performance / Price

Single-run Energy Consumption Comparison

FUJITSU

All results on Xeon Phi are preliminary and are expected to improve with the official V9.0 release

Please contact the author for actual information

∴ Xeon Phi is more energy efficient in almost all cases

*Taking only the simulation time into account

Model Preprocessing/Decomposition Comparison

FUJITSU

All results on Xeon Phi are preliminary and are expected to improve with the official V9.0 release

Please contact the author for actual information

∴ It is better to decompose the model on a Xeon node, then simulate it on Xeon-Phi.
This is because the I/O and decomposition phase do not scale on multiple cores.

Advancing the Standard Xeon Line

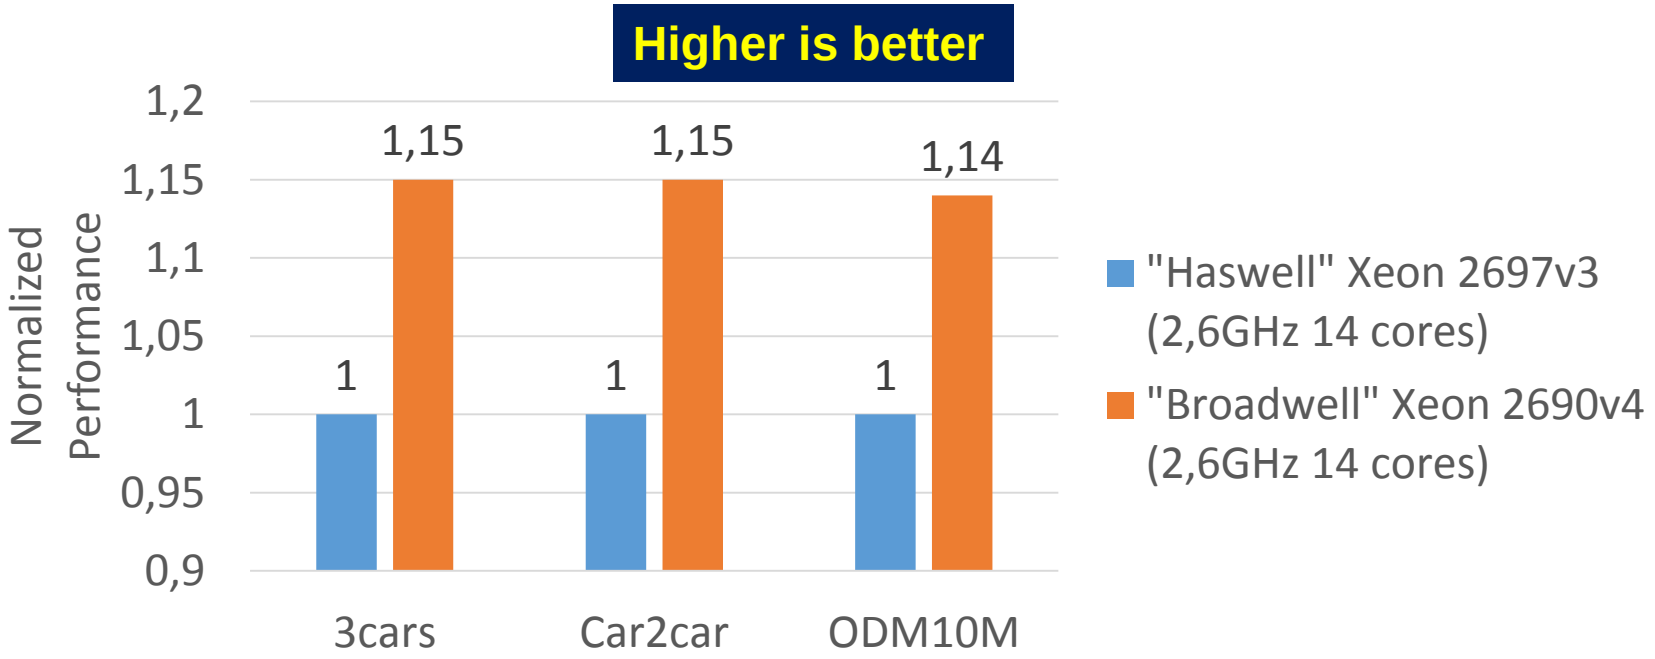

∴ Even with the same frequency and number of cores, the Broadwell architecture is still 15% faster than the previous Haswell generation

*Performance is based on the simulation time. The preprocessing time follows the same pattern.

Best practices: Process Pinning

- The default launching mode of Platform MPI might oversubscribe CPU cores and hence hinders the performance
- Use mpirun flags:
  -affcycle=core -affwidth=core -affblock=1 -affstep=1 -affopt=v,coreindex

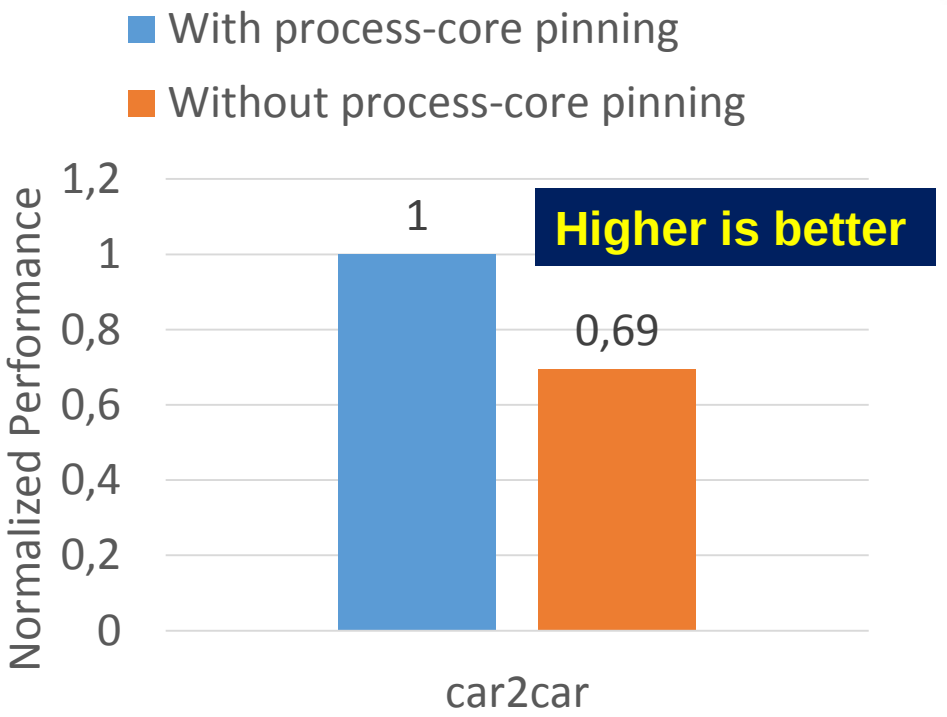

Best practices: Model decomposition

- Deformation occurs the most at the frontier of both cars
- Cut the y-axis so each process calculates the same amount of deformations
- P-file: decomp { sy <scaling factor> }

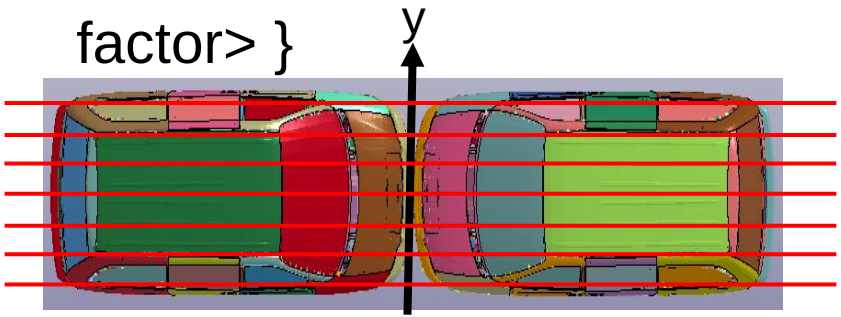

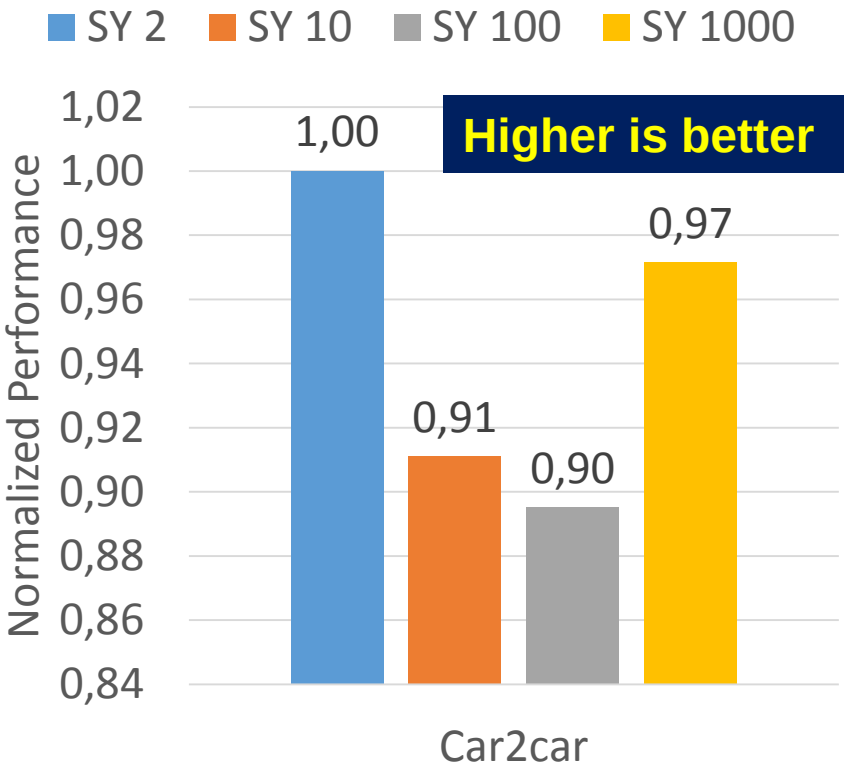

Best practices: Leverage AVX2/AVX512

- New hardware needs new code
- In particular wider vectorization
- AVX512 > AVX2 > SSE2

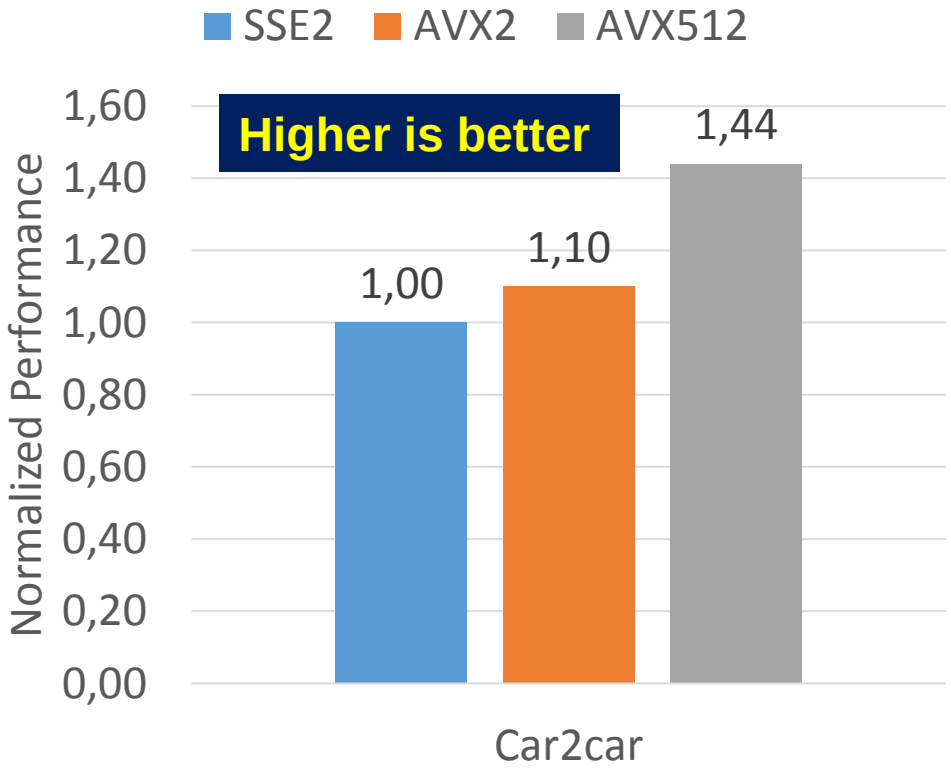

Conclusion

FUJITSU

- When simulating model deformations, Xeon-Phi offers
  - Superior price-performance (for simulation)
  - Lower energy consumption
  - Higher space density
- For model decomposition, Xeon is still the better choice
- Broadwell offers 15% higher performance than Haswell with the same core count and frequency, yet is offered at a lower price
- Pay attention to
  - Process pinning
  - Model decomposition methods
  - LS-DYNA version with the most advanced instruction set

High Performance Computing with Fujitsu

FUJITSU

Why PRIMEFLEX for HPC from Fujitsu?

- Benefit from
  - Supercomputing experience also in commercial deployments
  - established collaboration with leading ISVs
  - global HPC Competency Network

40

- Increase **innovation & productivity**
  with integrated HPC cluster solutions from Fujitsu

- **Compete effectively**
  in your global market place
  with Fujitsu PRIMEFLEX for HPC

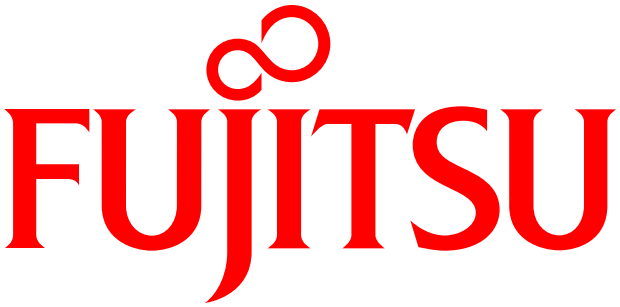

shaping tomorrow with you