Platform Choices for LS-DYNA

Manfred Willem and Lee Fisher
High Performance Computing Division, HP
lee.fisher@hp.com October, 2004

Public Benchmarks for LS-DYNA

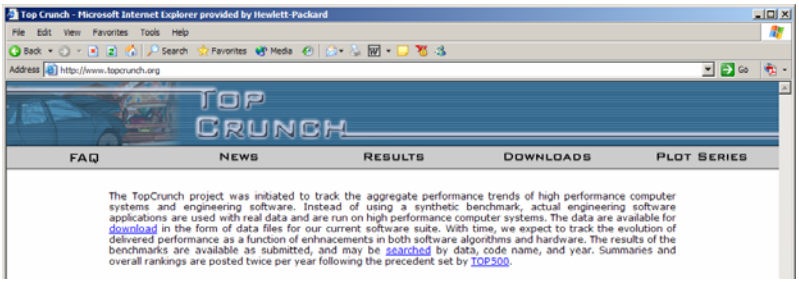

- www.TopCrunch.org
- administered by University of California, San Diego (Prof. David Benson)
- Refined Neon and 3 Vehicle Collision models for LS-DYNA

2

LS-DYNA Neon – from PA-RISC to IPF

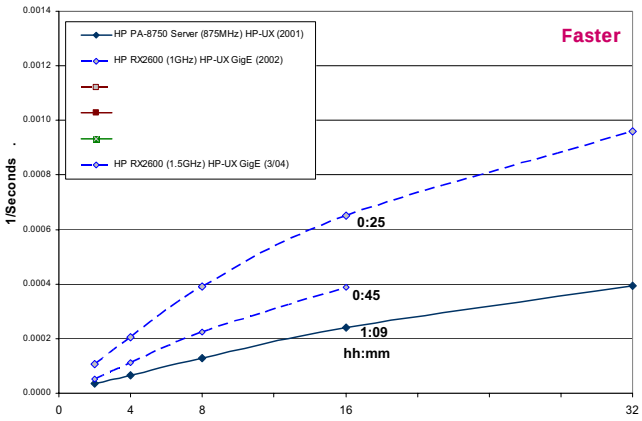

Observation:

• 1.5GHz Itanium is 70% faster than old Itanium (fall, 2002)

Reference: LS-DYNA v970.3858. Refined NCA Neon model for public benchmarks, 535K elements, posted as of 1-Oct-04 at: *www.topcrunch.org/benchmark_results.sfe*

3

LS-DYNA Neon – Power4 SP2 cluster

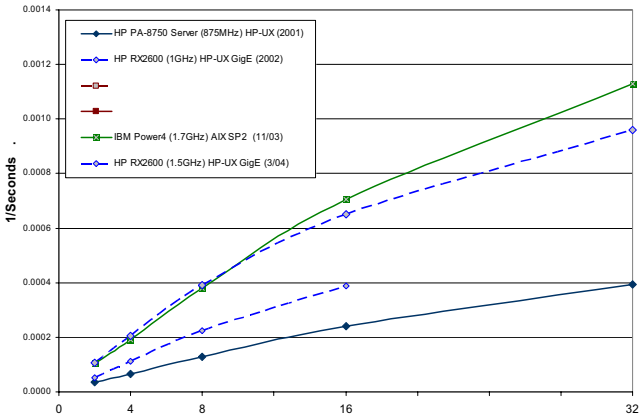

Observation:

• Itanium faster than Power4 at 2-8 CPUs

• Power4 faster above 8 due to low latency switch vs. GigE

Reference: LS-DYNA v970.3858. Refined NCA Neon model for public benchmarks, 535K elements. IBM data posted 13-Nov-03 at: *www.topcrunch.org/benchmark_results.sfe*

4

LS-DYNA Neon – Xeon GigE & Myrinet

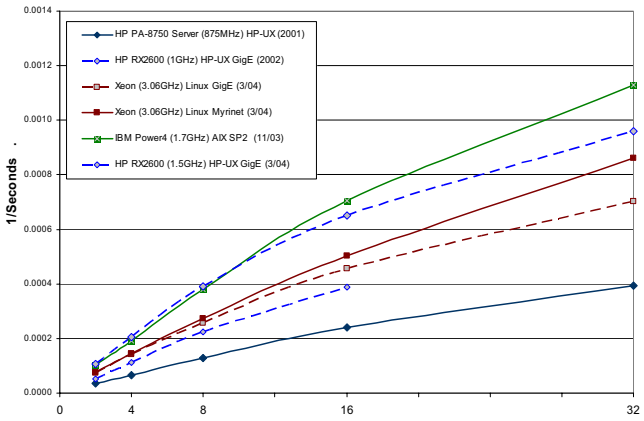

Observations:

- at 16 CPUs, Xeon adds half again as much time as Itanium
- scalability of GigE Xeon & Itanium track each other; scalability of Myrinet & SP2 track each other

Reference: LS-DYNA v970.3858. Refined NCA Neon model for public benchmarks, 535K elements. *www.topcrunch.org/benchmark_results.sfe*

5

LS-DYNA Neon – Itanium InfiniBand

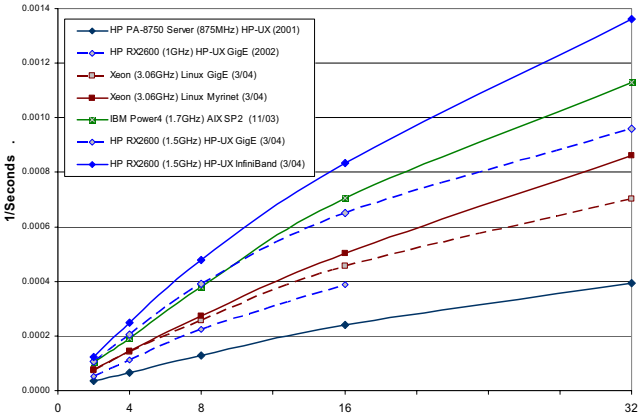

Observation:

- low latency InfiniBand switch puts Itanium/HP-UX in the lead

Reference: LS-DYNA v970.3858. Refined NCA Neon model for public benchmarks, 535K elements. *www.topcrunch.org/benchmark_results.sfe*

6

What about Opteron?

- All preceeding data is v970.3858
  - this version does NOT include native recompile for x86-64 (e.g. Opteron)
- New v970.5434 was released 3-Aug-04, supporting x86-64
  - no vendor has posted Opteron data at TopCrunch using this official release yet
  - LSTC has improved general performance by ~30% in this new release (contact algorithm)
  - therefore, BEWARE any comparison of hardware performance that mixes LS-DYNA versions!
  - also, insure runs do not leave idle CPUs in any nodes (classic benchmark trick)

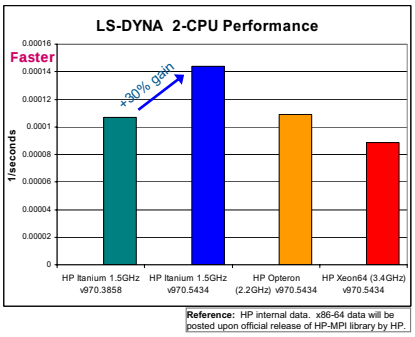

7

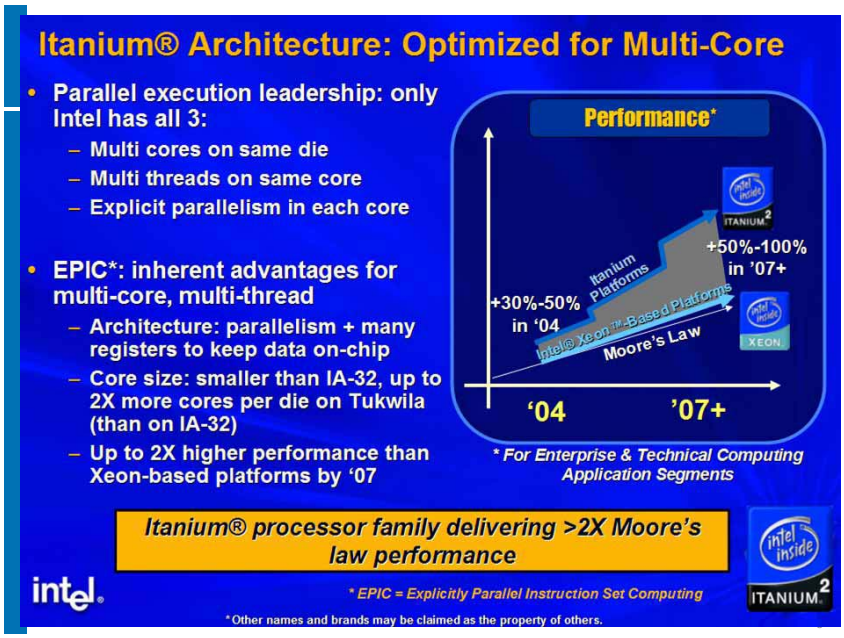

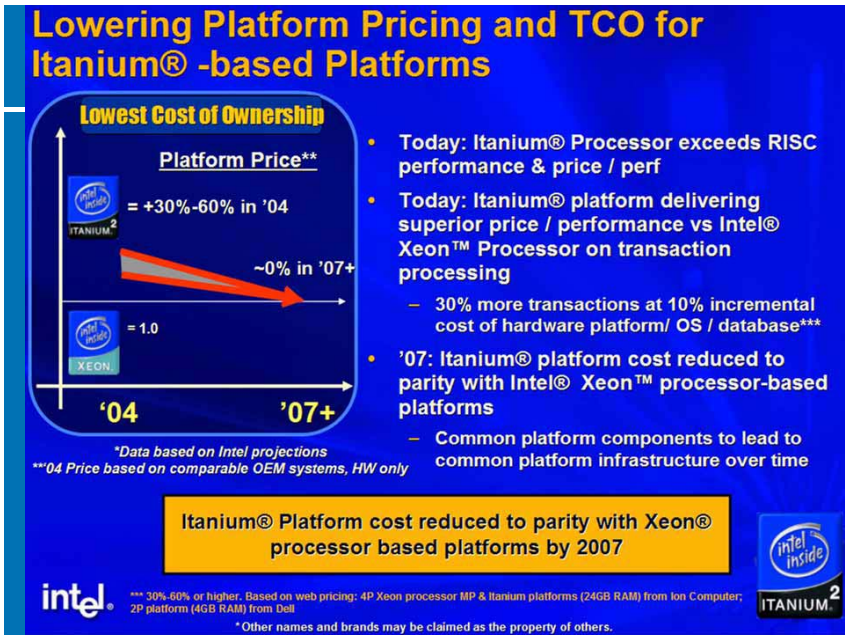

Value of HP Itanium platform increases as problems becoming more complex

- Bigger models
  - LS-DYNA 3 Vehicle Benchmark run 7.5 hours @ 16 CPUs
- Longer duration runs
  - Double Precision math is more efficient on 64-bit processors
- LS-DYNA for Implicit FEA
  - requires 64-bit address space and fast I/O

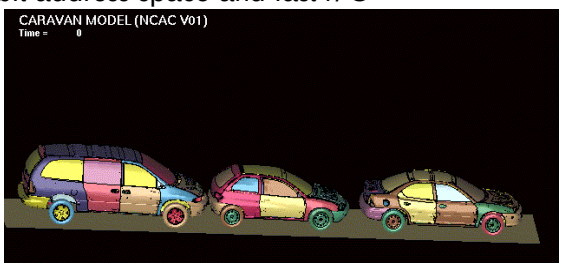

10

Benchmark Summary

- HP encourages hardware vendors and customers to use TopCrunch
- We offer choice: all leading architectures
  - Itanium / Xeon / Opteron clusters and SMP servers
- HP Itanium is the performance leader today
- x86 is price / performance leader today
- Technological changes will continue
- HP will consistently be a leader in performance & reliability

11

HP Platform Recommendations for LS-DYNA?

- it depends!

HP Servers & Clusters for CAE
Scalability of CAE Applications

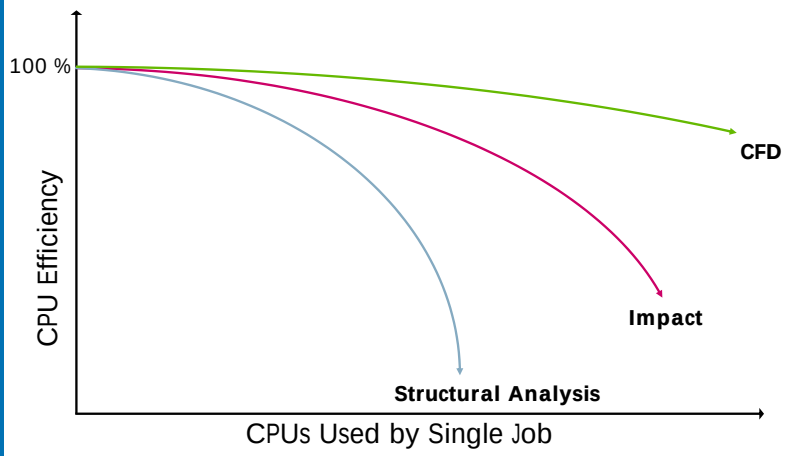

13

HP: Platform of Choice for LSTC™

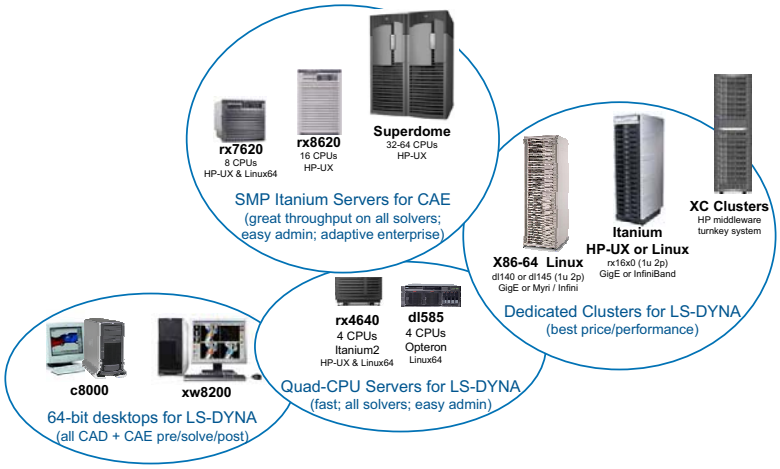

14