

Intel Exascale Activities in Europe

Karl Solchenbach Director Intel European ExaScale Labs

EPoPPEA Workshop 7th HiPEAC Conference Paris Jan 2012

Intel Confidential

ExaScale is coming ...

... but it won't be easy to get there





1994 1996 1998 2000 2002 2004 2006 2008 2010 2012 2014 2016 2018 2020



US DOE's Exascale Expectations

- DOE driving for an ExaFLOP system by 2018
- Fundamental challenges on Parallelsim, Energy, Resilience
- Lower ratios of floating point speed vs
 - Memory size
 - Memory bandwidth
 - Communication bandwidth

Char.	2010	2018	Ratio
Peak Perf	2 PF	1 EF	x 500
Sys Memory	0.3 PB	32-64 PB	x 100-200
Node Perf	125 GF	1-2 TF	x 8-16
Node Mem BW	25 GB/s	200-400 GB/s	x 8-16
Node Concurrency	12	~1000	x 80
Interconnect BW	1.5 GB/s	50 GB/s	x 30
# Nodes	20K	~ 1 M	x 50
Total Concurrency	225K	~1 B	x 4000
Storage	15 PB	300 PB	x 20
I/O	0.2 TB/s	60 TB/s	x 300
MTTI	Days	1 Day	x 0.1
Power	6 MW	20 MW	х З





http://www.csm.ornl.gov/~engelman/publications/engelmann10facilitating.ppt.pdf

Technologies and Solutions That Got Us to Petascale Science...



The Challenges to Exascale





Ine

Moore's Law is Well and Alive



What are the right cores?

"large" cores	"medium" cores	"small" cores

There is no one model for parallelism



Intel® Many Integrated Core (MIC) Architecture





- Extending IA/x86 to flexible, fully programmable general purpose Many-Core computing as co-processor
- Common IA programming models, languages, techniques and software tools with Intel® Xeon® processors
- Industry leading Performance for highly parallel workloads
- Higher Efficiency than other attached solutions
- Optimized Efficiency for a solution in combination with Intel Xeon processors

Single-source approach to Multi- and Many-Core



Eliminates Need to Fork Application Code



Spectrum of Programming Models and Mindsets



Productive Programming Models Across the Spectrum



Exascale Challenges for HW and SW developers

- Exploiting massive parallelism
 - Mathematical models, numerical methods, and software implementations will all need new conceptual and programming paradigms to make effective use of unprecedented levels of concurrency.
- Reducing power requirements
 - Reducing the power requirement by a factor of at least 100 is a challenge for future hardware and software technologies.
- Coping with run-time errors
 - an exascale system will have approximately one billion processing elements. An immediate consequence is that the frequency of errors will increase while timely identification and correction of errors become much more difficult.

Source: Exascale Computing Summary Report of the Advanced Scientific Computing Advisory Committee (ASCAC) Subcommittee



Exascale Power Scaling Requirements

Petascale Machine of 2010: TFLOP of Compute



Compute 40x Memory 75X Comms 20x Disk/Storage 33x Other 900x



A Holistic Perspective of the Challenge



HW-SW Co-Design & Context Awareness

Future Systems Must Be Balanced, Dynamic, and Adaptive



Intel Exascale Labs - Europe

Strong commitment to advance computing leading edge: Intel collaborating with HPC community & European researchers 4 labs in Europe, Exascale computing is the central topic





Intel European Exascale Labs

Basics

- Started 2010/2011 as codesign centers
- With leading European HPC R&D organizations
- In total ~60-70 researchers
- Work on joint R&D program with partners
- Part of ILE network

Role

- Understand requirements for exascale apps
- Provide feedback to HW architects, educate developers
- Build exascale HW and SW prototypes
- Be involved in European and national projects





Leuven: Numerical Kernels, Resilient Load-Balance and Scheduling, Multicore Simulator



ExaScience Lab

Paris: Performance Characterization and Application Enabling





Barcelona: Scalable RT (StarSs), Tools (Paraver/Dimemas), New Algorithms

StarSs: ... taskified ...

```
#pragma css task input(A, B) output(C)
void vadd3 (float A[BS], float B[BS],
           float C[BS]);
#pragma css task input(sum, A) inout(B)
void scale add (float *sum, float A[BS],
              float B[BS]):
#pragma css task input(A) inout(sum)
void accum (float A[BS], float *sum);
for (i=0; i<N; i+=BS)
                                // C=A+B
  vadd3 ( &A[i], &B[i], &C[i]);
for (i=0; i<N; i+=BS)
                          // sum(C[i])
  accum (&C[i], &sum);
. . .
for (i=0; i<N; i+=BS) // B=sum*A
  scale add (&sum, &E[i], &B[i]);
for (i=0; i<N; i+=BS)
                               // A=C+D
  vadd3 (&C[i], &D[i], &A[i]);
. . .
for (i=0; i<N; i+=BS) // E=G+F
  vadd3 (&G[i], &F[i], &E[i]);
```

Compute dependences @ task instantiation time



Color/number: order of task instantiation Some antidependences covered by flow dependences not drawn



Juelich: DEEP KNC-Cluster HW+SW architecture





Thank You!

