



chist-era



CHIST-ERA Projects Seminar
Day 2, Cross Topics
**Heterogeneous
Distributed Computing**

Hugh Leather

Bern, April 29th, 2016



FUNDING OPPORTUNITIES from the
FUTURE & EMERGING TECHNOLOGIES scheme






chist-era

Heterogeneity and Distribution In Action


This presentation was written by these
people:





Oh, no! No More Moore Anymore

Dennard Scaling





Power \propto area



Oh, no! No More Moore Anymore

Dennard Scaling



Power ~~\propto~~ area



chist-era

Oh, no! No More Moore Anymore

Moore's Law





Transistor count doubles every 2 years



Oh, no! No More Moore Anymore

Moore's Law




TSMC: Per transistor cost rises in 2015!

Heterogeneity to the rescue

Match the program to the hardware


Fat CPU




Heterogeneity to the rescue

Match the program to the hardware

Fat CPU




Small CPUs



Heterogeneity to the rescue

Match the program to the hardware


Fat CPU



GPU




Small CPUs



Heterogeneity to the rescue

Match the program to the hardware


Fat CPU



GPU



Small CPUs

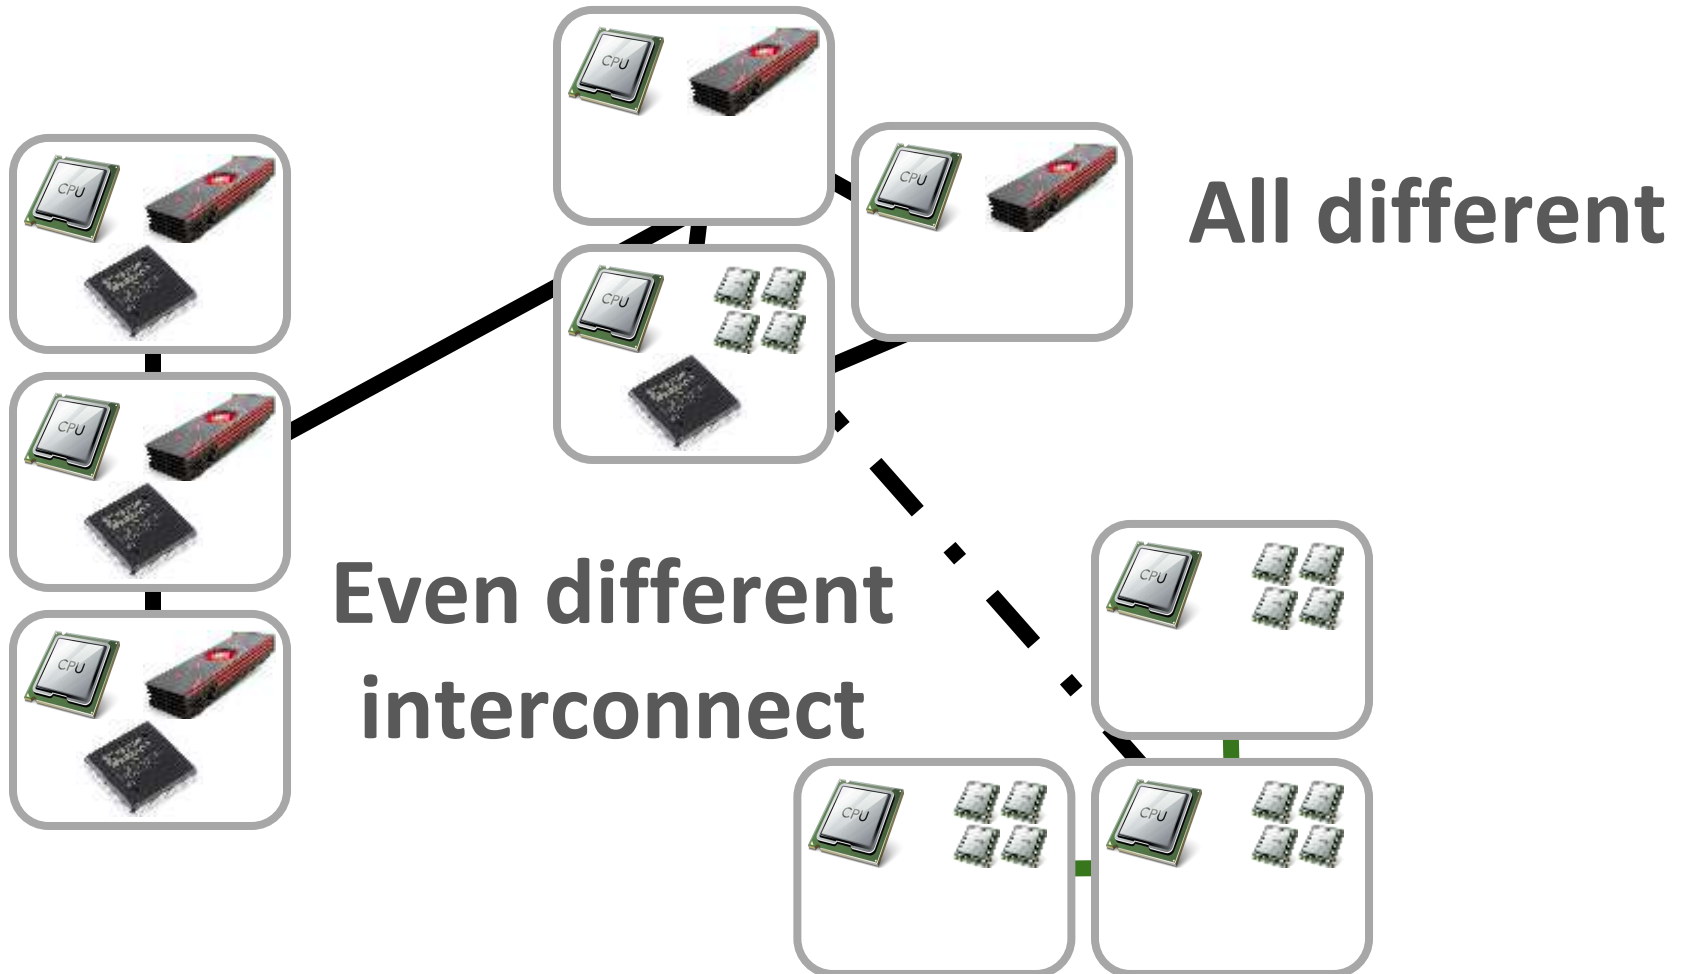


FPGA



Heterogeneity to the rescue

But doesn't fit on one machine






The **BIG PROBLEMS** Today

- Too hard to program
 - CPUs, GPUs, FPGAs - complex interactions
 - Massive distribution - complex network
 - Must optimise at multiple scales
 - Only experts can play

The **BIG PROBLEMS** Today

- Too much energy
 - ~20MW each, 1.5% globally, growth exponential





Heterogeneous Distributed Computing

- Machines internally heterogeneous
- Machines heterogeneous to each other
- Massive distributed networks
- Networks heterogeneous
- **Very hard to program**
- **If we don't get it right =>**
energy/performance disaster



The Projects

HPDCJ

Heterogenous Parallel Distributed Computing in Java

DIONASYS

Declarative and Interoperable Overlay Networks,
Applications to Systems of Systems

DIVIDEND

Distributed Heterogeneous Vertically Integrated
Energy Efficient Data centres

The Projects

	HPDCJ	DIVIDEND	DIONASYS
<i>Programming Model</i>	✓	✓	✓
<i>Dependability</i>	✓		
<i>Data Management</i>	✓	✓	✓
<i>Versatility</i>		✓	✓
<i>Optimisation Techniques</i>		✓	
<i>Distributed Techniques</i>	✓	✓	✓




High Performance Computing

- Parallel distributed computing in Java
 - PCJ library for parallel computing in Java
- Scalability up to 6000 cores
- CPU and GPGPU
- Fault Tolerance


- Easy for non expert programmers
 - New approach to teach students

System as a first-class component (holon)

- Generative programming
- Application in IoT
- System Composition



Self organising overlays
Prototypes Open Sourced

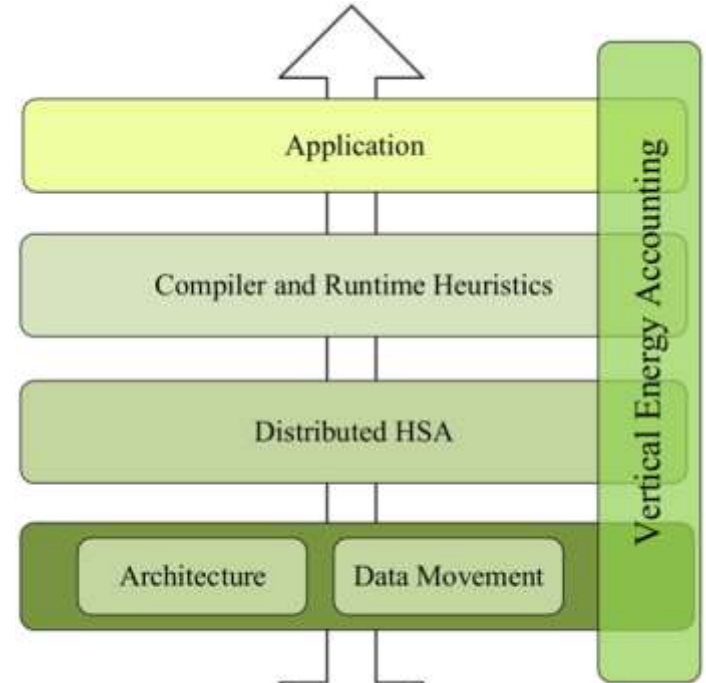




DIVIDEND

- Vertical integration
- Programming model
- Energy accounting
- Auto tuning
- More heterogeneity
- Fast networks

Prototypes Open Sourced
Already saving 22% energy



```
QModelIndex start;  
if (currentIndex().isValid())  
    start = currentIndex();  
else  
    start = d->model->index(0  
  
bool skipRow = false;  
bool keyboardTimeWasValid = d  
qint64 keyboardInputTimeElaps  
if (search.isEmpty() || !keyb  
|| keyboardInputTimeElaps  
d->keyboardInput = search  
skipRow = currentIndex().  
} else {  
    d->keyboardInput += searc  
}
```



Grand Challenge

Write a program, then system automatically

- Chooses the right hardware
 - Or creates new hardware
- Optimises everything
- It is easy for the programmer



Roadmap

In 5 Years we need

- Programming models for major domains
- DSLs to specialise to all devices (CPU, GPGPU, FPGA)
- Eliminate waste in computing
- SDN needs to be transparent the application



Roadmap

In 10 Years we need

- Universal languages for the masses
- Tool chains to co-design platforms and fabricate logic/network/memory blocks for services
- Programming without knowing what's out there





CHIST-ERA's Role

Provide new calls

- Parallel programming
- Energy optimisation
- Automatic hardware synthesis



Conclusion

- Energy/performance crisis looming
- Can't program and optimise HDC
- We are making progress on this
- Need more calls on this