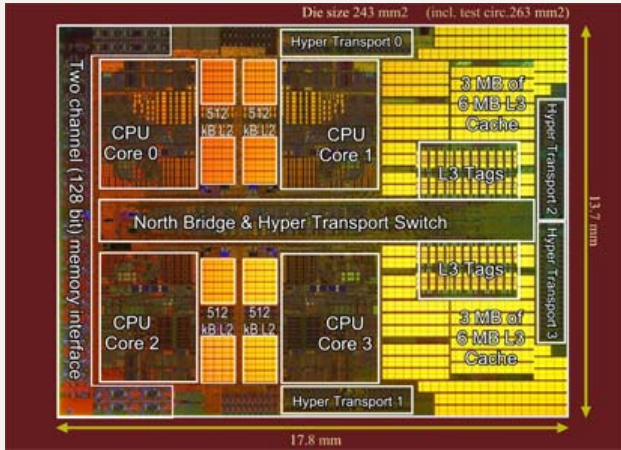


Programming Models for Accelerator Clouds


Nate Clark, Sudha Yalamanchili, Greg Damos,
Andrew Kerr, Haicheng Wu, Kirak Hong,
and many people at LogicBlox



What / Why Accelerator Clouds?



AMD Shanghai



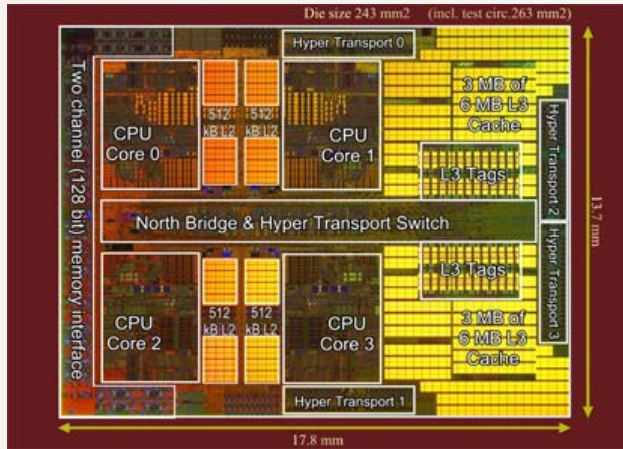
NVIDIA Fermi

- Power wall → heterogeneity
- Compute / \$

Problem: Programming




MIMD Coherent Shared Memory



AMD Shanghai

SIMD Data Parallel



NVIDIA Fermi

- Productivity
- Performance/system portability

A PARALLEL CHANNEL COMPUTING MACHINE


Lecture by
J. P. Eckert, Jr.
Electronic Control Company

. . . Again I wish to reiterate the point that all the arguments for parallel operation are only valid provided one applies them to the steps which the built in or wired in programming of the machine operates. Any steps which are programmed by the operator, who sets up the machine, should be set up only in a serial fashion. It has been shown over and over again that any departure from this procedure results in a system which is much too complicated to use.





Portability Solution

- Java, Fortran, MapReduce, etc.



- Side effect free kernels
- Explicit dependencies
- Enables aggressive optimization

Harmony Runtime






Productivity: LogicBlox's Datalog

- Declarative language (not imperative)
 - Think high-powered Excel
- Many examples of business uses
 - Risk analysis
 - Business analytics
- (Usually) maps well to kernel IR

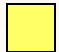
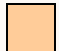

Current Status



Application Development




Legend


-  Open Source available
-  In progress
-  Commodity SW/GT



Translate to Harmony Kernel IR



Kernel IR




Harmony: Run-Time



Ocelot: Code Generation and Dynamic Optimization Environment



Risk Analysis Application





Conclusion

- Preliminary results look good
- Interaction benefits both LB and GT
- Open research
 - Scaling, scaling, scaling
 - Further optimizations
 - Other programming paradigms?