

# Making Integration our Friend (again...)

Christos Kozyrakis

Computer Systems Laboratory  
Stanford University

<http://csl.stanford.edu/~christos>

# Integration Used to Be Our Friend

---

- **Moore's Law used to be good for all of us**
  - Processors were getting faster at exponential rates
    - Larger caches, larger instruction windows, multiple ALUs, ...
    - Faster clocks
    - Better performance with 35 year old programming model!
  - ASIC-based systems were getting easier and cheaper
    - More functionality in one chip
    - Fewer components, smaller boards, ...
  - Power consumption was reduced
    - Reduced voltages, fewer board traces to drive, ...
- **Business as usual was enough to succeed**
  - My undergrad advisor's PhD (1984): a 40K-transistor processor
  - My PhD (2002): a 125M-transistor processor

# But Then It Got Ugly

---

- **We run out of instruction-level parallelism**
  - A bigger uni-processor is no longer faster
- **Power became an issue**
  - Cannot switch all transistors at the same time
  - Cannot scale the clock frequency
- **Design cost and complexity became an issue**
  - Difficult to verify and program that bigger ASIC
- **Off-chip bandwidth became an issue**
  - Cannot get data in to feed my transistors
- **The good news with all that?**
  - We've got some good problems to solve...

# Making Integration our Friend (1)

---

- (Re)configurable multiprocessor chips
  - General-purpose  $\Rightarrow$  volume
  - Modular & replicated  $\Rightarrow$  scalable, easy to verify
  - But can they be as efficient as an ASIC?
- Our approach
  - Start with simple processor and memory components
  - Add some FPGA-like reconfiguration flexibility
    - Is the memory a cache, a FIFO, a scratch-pad?
    - Is the processor reading operands from local storage or network?
  - Configure chip to most efficient design for each application
    - Vector processor for applications with data-level parallelism
    - Shared-memory multiprocessor for applications with task parallelism
    - Mix and match
- More details:
  - [http://www-vlsi.stanford.edu/smart\\_memories/](http://www-vlsi.stanford.edu/smart_memories/)

# Making Integration our Friend (2)

---

- **Practical parallel programming**
  - Currently, parallel programming is a black art
  - What good is a multiprocessor chip then?
- **Our approach**
  - Parallel programming with transactional memory semantics
  - Hardware supports light-weight transactions
    - Atomicity and performance using optimistic concurrency
  - Easy to write correct parallel programs
    - Coarse-grain & speculative parallelism
  - Hardware helps with performance tuning
    - Accurate feedback on performance loss events
    - Dynamic optimization for locality
- **More details:**
  - <http://tcc.stanford.edu>

# Other Technologies We Are Keeping an Eye On

---

- Adding “intelligence” into the system
  - Using AI techniques to learn best dynamic way to deal with memory latency, data placement, power usage, ...
- Power management in large data-centers
  - Room-level integration...
  - Statistical rigorous scheduling techniques, heterogeneous server architectures, ...
- Architecture opportunities with 3-D integration
  - More than area density
  - High-bandwidth connections between components
  - Free additional functionality for productivity, reliability, ...