


VIRAM-1 Floorplan – Tapeout June 01



Thanks to DARPA: funding
 IBM: donate masks, fab
 Avanti: donate CAD tools
 MIPS: donate MIPS core
 Cray: Compilers, MIT: FPU
 CalStan 3/201

- **Microprocessor**
 - 256-bit media processor
 - 12-14 MBytes DRAM
 - 2.5-3.2 Gops
 - 2W at 170-200 MHz
 - Industrial strength compiler
- **280 mm² die area**
 - 18.72 x 15 mm
 - ~200 mm² for memory/logic
 - DRAM: ~140 mm²
 - Vector lanes: ~50 mm²
- **Technology: IBM SA-27E**
 - 0.18μm CMOS
 - 6 metal layers (copper)
- **Transistor count: >100M**
- **Implemented by 6 graduate students**

1

Lessons learned from Past Projects

- **Maintenance of machines expensive**
 - ~10X cost of HW **per year**
 - System administration primarily keeps system available: System + clever human = uptime
 - Software upgrades necessary, dangerous
- **Everything has an error rate**
 - Well designed, manufactured HW: >1% fail/yr
 - Well designed, tested SW: > 1 bug / 1000 lines
 - Well trained, rested people: >1%??
 - Well run collocation site (e.g., Exodus): 1 power failure / year, 1 network outage / year
- **Can improve performance (and cost)**
 - Run on workload, measure, innovate, repeat
 - Benchmarks standardize workloads, lead to competition, turning debates into numbers

3

Goals, Assumptions of last 15 years

- **Goal #1: Improve performance**
- **Goal #2: Improve performance**
- **Goal #3: Improve cost-performance**
- **Assumptions**
 - Humans are perfect (they don't make mistakes during installation, wiring, upgrade, maintenance or repair)
 - Software will eventually be bug free (good programmers write bug-free code)
 - Hardware MTBF is already very large (~100 years between failures), and will continue to increase

CalStan 3/201

2

An Approach to Trouble-Free Systems

"If a problem has no solution, it may not be a problem, but a fact, not be solved, but to be coped with over time."


Shimon Peres, quoted in *Rumsfeld's Rules*

- **Rather than aim towards (or wait for) perfect hardware, software, & people, assume flaws**
- **Focus on Mean Time To Repair (MTTR), for whole system including people who maintain it**
 - Availability = MTTR / MTBF, so 1/10th MTTR just as valuable as 10X MTBF
- **Use techniques to make repair fast vs. programs fast: transactions for undo, benchmarks for competition, ...**

CalStan 3/201

4

Moore's Law vs. Common Sense?



- Scaled 32-bit, 5-stage RISC II
1/1000th of current MPU, die size or
transistors

CalStan 3/201

5



New view: ClusterOnaChip (CoC)

- 32-bit MPU as the new Nand Gate
 - “Cluster on a chip” with 100s of processors enable amazing MIPS/\$, MIPS/watt for cluster applications
 - MPUs combined with dense memory + system on a chip CAD + quick turn fab
- Inspiration: Google
 - Search engine for world: 100M/day
 - Economical, scalable build block: PC cluster today 6000 PCs, 12000 disks
 - Advantages in fault tolerance, scalability, cost/performance
- 30 years ago Intel 4004 used 800 gates: when an 800 processor CoC?

CalStan 3/201

6