

Questions for Embra: Fast and Flexible Machine Simulation

General

- Why is it useful to be able to switch the level of simulation detail while running?
- How is Embra simulation different than Shade simulation?
- How did the existence of SimOS help the development of Embra?

Similar to Shade

- Data structures: pc2tc hash table, translation cache
- Chaining to "improve performance by eliminating many of the lookups in the pc2tc data structure"

MMU issues

- MMU simulation performance is critical (say the authors); why? Why not an issue in Shade?
- TLBs are fully associative, yet Embra does not implement a fully-associative TLB; why?
- What is the MMU relocation array, and how is it used?
- How are ASIDs handled in Embra's TLB emulation?
- How would this approach scale as virtual address spaces grow?
- Why is instruction fetch easier to make efficient?
- What checks must be performed and when for instruction MMU simulation?
- Can a translated basic block cross page boundaries?
- Are uncached loads/stores handled efficiently by Embra?
- What about accesses to physical memory through KSEG0?

Chaining

- Why are chains in Embra harder to implement than chains in Shade?
- How does speculative chaining help overcome this?

Customized Translations

- vQC is a key data structure; what is it and how is it used?
How does the vQC interact with the MMU emulation described above?
- Why does Embra generally work better when emulating larger caches?

Evaluation (aka the Thing We Never Talk About)

- Table 5.1 gives a good sense of Embra performance; how does it compare to Shade?
- What types of applications make Embra run more slowly?
- How much does cache simulation impact Embra's performance?
- How do the authors study Embra performance?
- The study in Figure "5.3" changes the model of the CPU; what doesn't it change?
- How do the authors convince you that Embra is accurate?

Best line:

- Section 2: "Bootng and initializing the operating system and database with MXS would have taken several years."

Typos:

- 3.3.1: "an data structure", "Embra can modeling"
- 5.1: "a course multiprocessor interleaving"
- Figure 5.3 is called Figure 5.1, Table 5.2 is called Figure 5.2