
Blooming on c.mmp

Kwame

John

Quick Recap

- C.mmp has 16 PDP 11 processors.
 - 32 MB of shared memory in 16 banks
 - Crosspoint switch to link each processor to every memory bank.
 - SAP
 - Hydra quick-and-dirty OS which runs c.mmp.
-

IXP vs C.mmp

- Similarities

- Multiprocessor
- Shared Memory
- Synchronization mechanisms

- Differences

- No multithreading support
 - Switch interconnect
 - 16-bit address space
-

Bloom Requirements

- Memory for hash tables.
 - Fast memory access time
 - Computation /processing
 - Means by which data can be accessed easily by each Bloom implementation
-

C.mmp Implementation

- Bloom filter can be implemented on each processor.
 - Dedicate one processor to “receive” and one to “transmit”.
 - Assuming a network I/O device on c.mmp, Receive process can implement a Ring type data structure that is shared among all the Bloom filters.
 - Circumvent the “paging” limitation in c.mmp to improve performance.
 - Deploy hash tables in two ways
 - Having multiple hash tables
 - Having one shared hash table among all processors (Bloom filters)
-

C.mmp Performance

- Contention for shared data structure on c.mmp. IXP scratch rings for hardware support for reducing it.
 - Local memory on IXP significantly faster than accessing shared memory on c.mmp
 - No dedicated process to fetch packets into memory. IXP MSF is capable of placing packets directly in memory
 - Can use multithreading on IXP that can hide memory latencies, but not on c.mmp
-