

Router Bottlenecks we will study

- Exact match: this lecture
- Longest prefix Match
- Packet Classification
- QoS
- Switching

Context 1

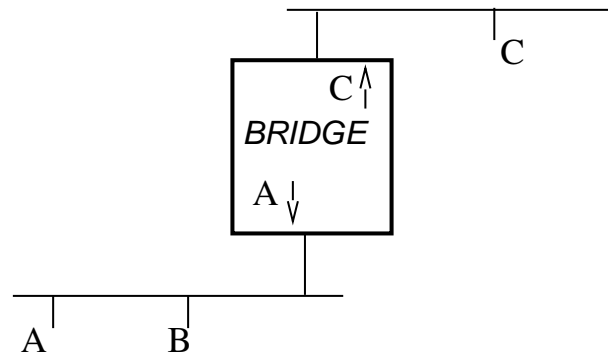
‘Challenge-and-response’ is a formula describing the free play of forces that provokes new departures in individual and social life.

An effective challenge stimulates men to creative action . . .

— *Arnold Toynbee*

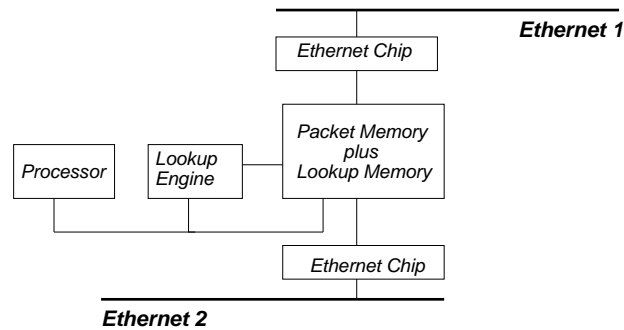
- Query specifies key K . Retrieve state associate with K . Hash table or binary search but memory references crucial and possible hardware speedups.
- Simplest lookup problem, yet crucial for a bridge. Birthed by 3 challenges

Challenges that led to bridges



- *Challenge 1:* Ethernets under fire, routers too slow.
- From packet repeater, to selective repeater, to learning bridge

Challenge 2: Wire Speed Forwarding



- Without wire-speed, can lose packets to be forwarded in a stream of minimum size packets to be filtered.
- Finessed issue by doing 2 lookups on either port in 51.2 usec


Attention to Algorithmics

- **Architecture:** 4-port cheap DRAM with cycle time of 100 nsec for packet buffers and lookup memory. Bus parallelism, memort bandwidth, page mode.
- **Data Copying:** Ethernet chips used DMA, packets copied from one port to other by flipping pointers.
- **Control Overhead:** Interrupt overhead minimized by processor polling, staying in a loop after a packet interrupt.
- **Lookups:** Used caveats. Wrote software to verify lookup bottleneck (**Q2**)

Eight Cautionary Questions in Action

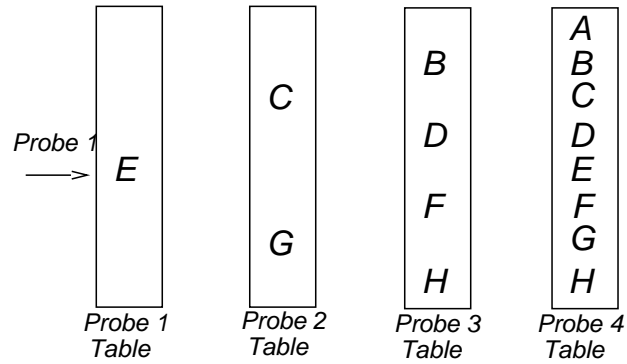
- **Q1, Worth improving performance?** Yes.
- **Q2, Really a bottleneck?** 68,000 code timing did not meet 25.6 usec budget.
- **Q3 Impact on system?** Allows wire-speed.
- **Q4, Initial analysis shows benefits?**
Compare and branch instructions in H/W. DRAM references = $\log_2 8000$, so 1.3 usec.
- **Q5, Worth adding custom hardware?**
Cheap using a PAL.
- **Q7, Do Prototypes confirm promise?** Lab prototype tested before product was built.
- **Q8: Sensitive to environment changes?**
Worst-case design

Challenge 3: Scaling to Higher Speeds



- Scaling via hashing for 32 port FDDI (100 mbps) bridge. Avoids worst case via perfect hashing.
- $A(X) * M(X) \bmod G(X)$, where $G(X)$ is $X^{48} + X^{36} + X^{25} + X^{10} + 1$, and A is picked randomly

Using Hardware Parallelism



- Seemingly arcane and clever idea. Can be derived readily by theory!

Lessons

- *Challenge-Response*: Bridges invented to solve a problem (slow multi-protocol routers) not present today.
- *Cost-Performance*: is what justifies bridges today in so-called switches together with a small number of extra features like VLANs.
- *Introduced techniques*: Broke wire-speed barrier, spread to routers, ideas like fast lookups, trading updates for lookup speed, minimal hardware all became classical.