Problemas y ampliación de temas: Arquitectura de routers

Jorge García Vidal

Problema

- Supongamos que en los algoritmos de lookup, los tiempos vienen dados por los tiempos de acceso a memoria.
- Supongamos:
  - SRAM, tamaño < 100 KB, t= 5ns
  - DRAM, t = 50 ns
- Queremos saber hasta qué punto escalarán los sistemas que estemos considerando

Caso 1: Switch ethernet

Supongamos que un único elemento de proceso es responsable del L2@ lookup. El switch tiene N puertos fast-ethernet.
El método de lookup es el mostrado en la figura. Podemos soportar hasta 1024 direcciones MAC distintas. Queremos tener p{Colisión} < 1%. ¿Cuánto valdrá h?

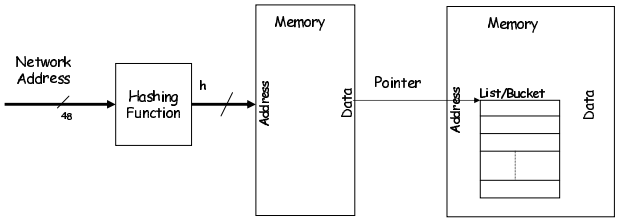

Caso 2: IP lookup

Supongamos que un único elemento de proceso es responsable del IP lookup. El router tiene N puertos fast-ethernet.
El método de IP lookup es el mostrado en la figura.

Binary Trie

a : 0*
b : 01000*
c : 011*
d : 1*
e : 100*
f : 1100*
g : 1101*
h : 1110*
i : 1111*

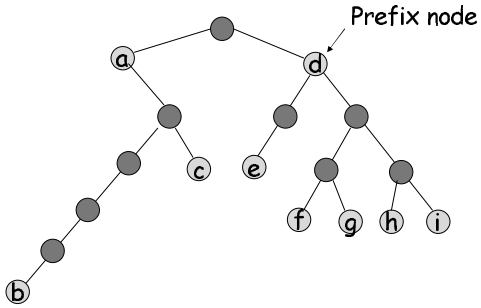

Caso 3: IP lookup optimizado

Supongamos que un único elemento de proceso es responsable del IP lookup. El router tiene N puertos fast-ethernet.
El método de IP lookup es el mostrado en la figura. Stride = 8 (CEF).

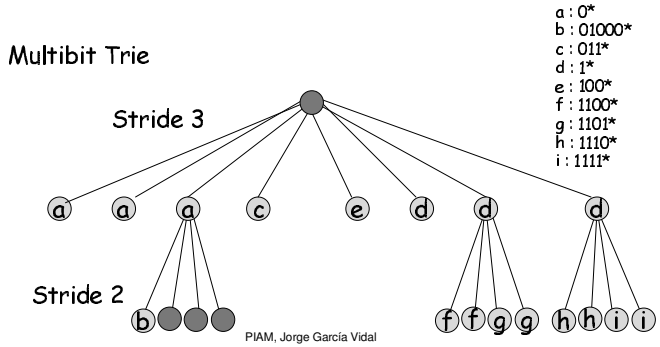

IP lookup

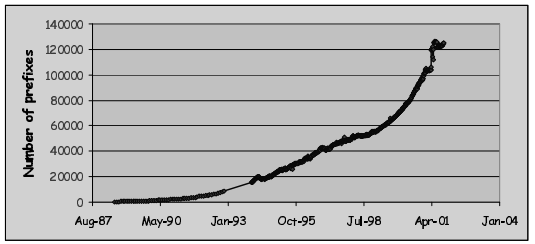

(De N. McKeown)

IP lookup

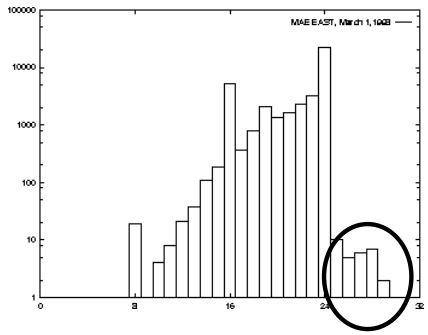

Most prefixes are 24-bits or shorter

Caso 4: IP lookup

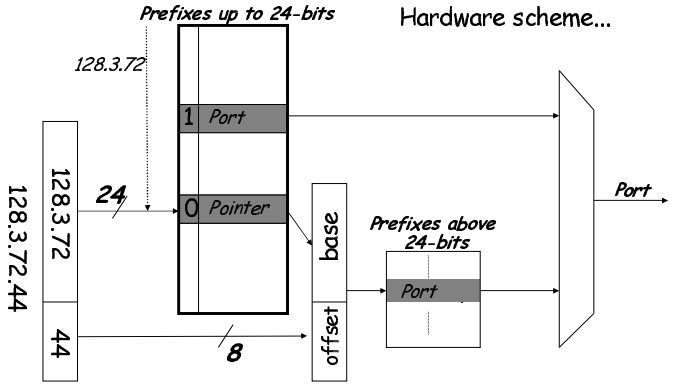

(De N. McKeown)

Shared Mem Routers (Cisco terminology)

- Cisco 1600, 2500 , 4000, 4500 and 4700
- Architecture: CPU, Main Memory and Interfaces
- Use System buffers for packet buffering

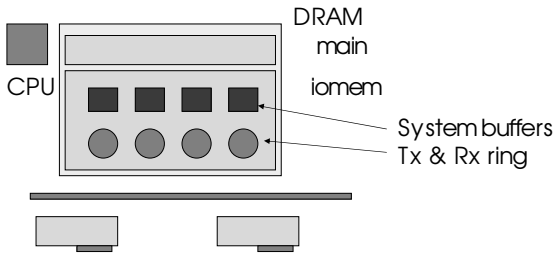

Shared Memory Routers

- The CPU is responsible for executing all the switching methods; there are no offload processors involved
- Cisco 1600, 2500: Motorola 68000
- Cisco 4500, 4700: RISC MIP
- "Shared Memory" is not a convenient name.

Shared Memory Routers

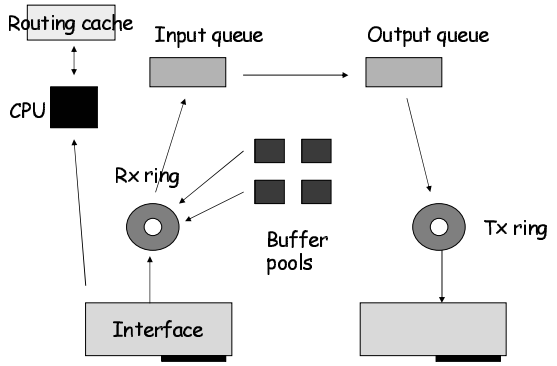

SMR: Receiving the pckt

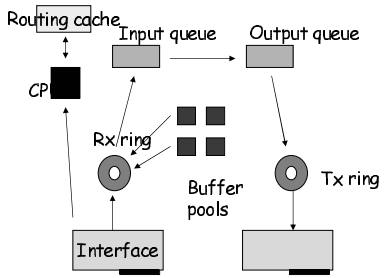

The interface copies the packet into a buffer in the receive ring using DMA

The media controller changes ownership of the pckt buffer to the processor and issues an interrupt to the CPU

The CPU removes the incoming pckt and removes the pckt buffer of the receive ring

SMR: Switching the pckt

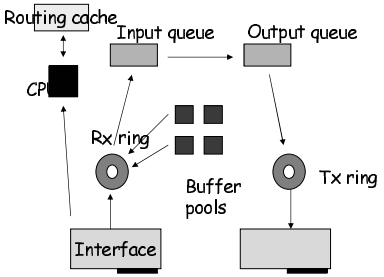

After the receive ring is replenished, the CPU begins switching the pckt. IOS attempts to switch using the fastest method configured on the interface: CEF, FS or PS

SMR: Transmitting the pckt

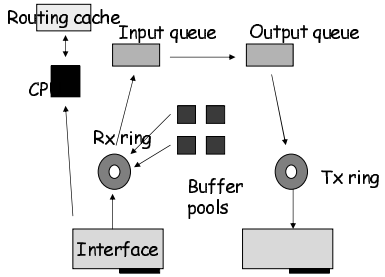

The pckt is placed in the Tx ring.

The Output interface periodically polls its tx ring. As soon as the media controller detects a pckt, it copies the pckt onto the network media and raises a Tx interrupt to the processor.

The CPU unlinks the pckt buffer from the transmit ring and returns the buffer to the pool of buffers from which it came