

Routing

Joachim Malmquist, KTHNOC

jocke@sunset.se

February 4, 2004

What is a router?

Host Can have several network interfaces, but cannot forward packets between them.

Router Can forward packets between its interfaces.

The router's job consists of . . .

1. Packet Forwarding, and . . .
2. . . building a routing table. Often routers come with other functions: packet filtering, routing protocols, etc.

Under the hood

Most routers are built using ASICs (Application Specific IC). This is what distinguishes a heavy-duty router from a PC-router.

- Network interfaces

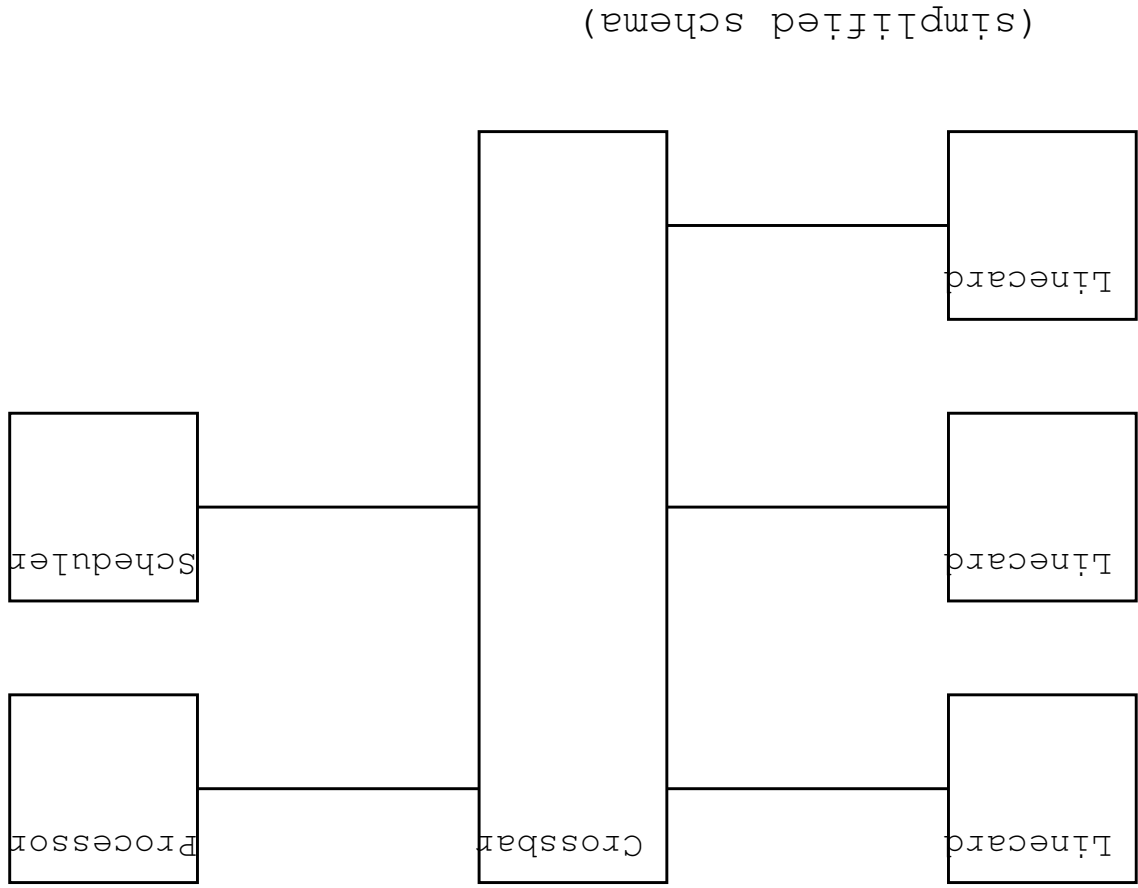
- PROCESSOR

- Cross Bar

- Scheduler

- The case with its fans, temperature measuring equipment, etc.

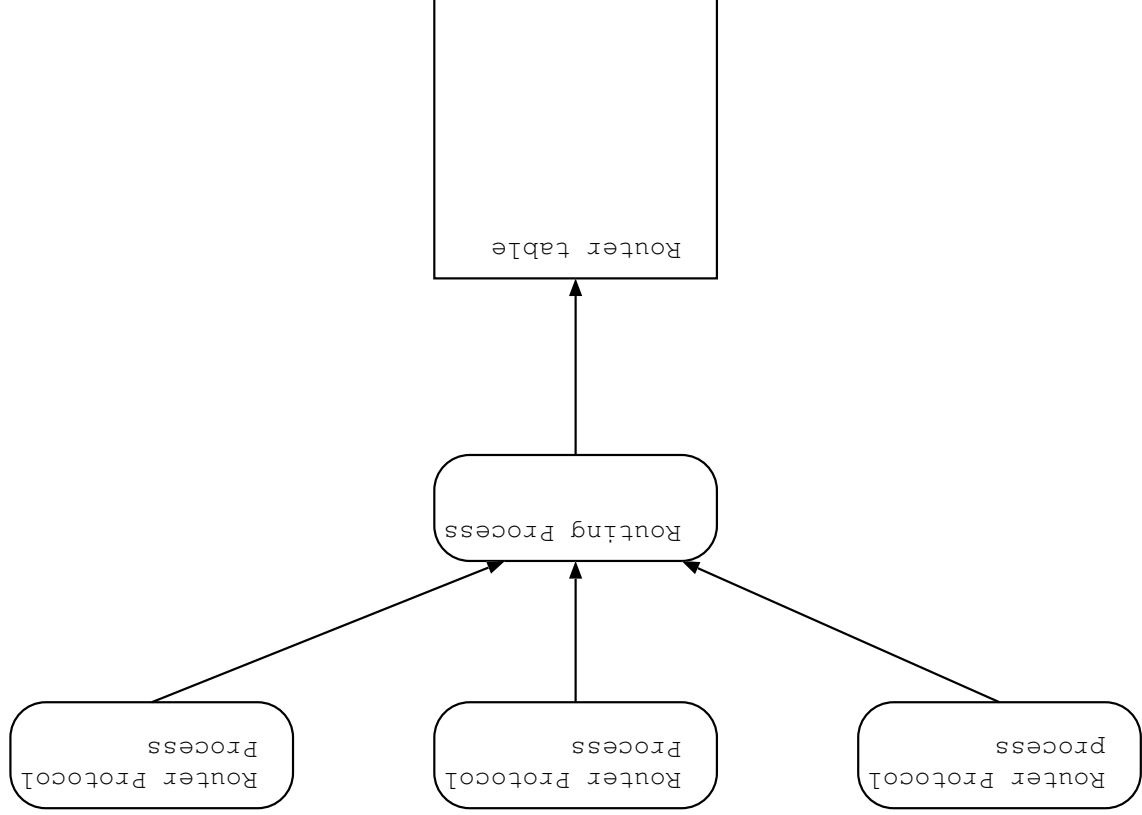
Example



Router Architecture problems

- Head of line blocking
- Synchronizing the forwarding table in the line card with the routing table

The Routing Process



The Routing Table

- How to handle large routing tables...
 - Choice of routing table datastructure and its consequences.
- Which routing protocol is allowed to write to the routing table?
 - Administrative distance

Routing Algorithms

How does a router find the optimal path to the destination?

Most solutions build on SPF (Shortest Path First) Algorithms, which are well known in graph theory.

- Bellman-Ford
- Dijkstra

Load Balancing

If there are more than one optimal path to a destination, what do you do?

It would be preferable to balance the load between the paths. The question is how to do this? What consequences do we get?

- Per IP packet

- Per routing table entry

- Per “flow”

What happens if network congestion arises at some point in the network?

Routing Protocol Classes

All routing protocols can be classified into one of two groups:

- Link State Algorithms (OSPF, IS-IS, NLSP)
- Distance Vector Algorithms (RIP, IGRP, BGP)

Routing table example (Cisco GSR)

```

stk-bb-1>sh ip ro
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
* - candidate default, U - per-user static route, o - ODR
Gateway of last resort is not set

B      208.221.13.0/24 [200/10] via 193.10.252.177, 2w2d
170.170.0.0/16 is variably subnetted, 5 subnets, 3 masks
B      170.170.0.0/20 [200/10] via 193.10.252.177, 2w2d
B      170.170.0.0/19 [200/10] via 193.10.252.177, 1w3d
S      130.242.48.0/25 [1/0] via 130.242.48.253
S      130.242.52.0/24 [1/0] via 130.242.205.46
i L2   130.242.204.176/30 [15/6] via 130.242.94.6, SRP4/0
i L2   130.242.203.180/30 [15/22] via 130.242.94.5, SRP4/0

```