Implications of the GRID on Networks (Random Thoughts)

Ian Leslie

University of Cambridge

Computer Laboratory

http://www.cl.cam.ac.uk/~iml/

Outline

- The Grid (my prejudices)
- Network Implications
- Virtual Networks
- Virtual Network Service

The Grid

- Driven originally by high energy physics community
- I personally don't find distributing a high volume of collider data to a small number of sites round the world terribly exciting.
- Globus is the current middleware used by most Grid projects (described by some as a bunch of shell scripts to do secure remote job entry)

The Grid

- Bit of a political animal, viewed with suspicion by much of the Computer Science community
- However, there are some interesting problems for computer science. (Just have to be careful to ensure that the other interesting problems don't suffer.)

The Grid: What's Interesting for CS?

- Large scale distributed computing.
  - Not big supercomputers (which are boring).
  - Lots of computers.
  - Lots of computers where "involvement" is dynamic but where resources must be allocated intelligently
- Discovery of information, understanding of information, assurance of information
- Virtual organisations, conflicting security models (don't need the Grid for this!)
- Oh, and some applications…

The Grid: What's Challenging for "Infrastructure"?

- Higher capacity demands (sort of obvious, these really relate to funding of science projects with communication component).
- Higher reliability
- No longer just the elephants and the mice?
  - Need for better resource management
  - Need for better performance monitoring
  - Need for self-provisioned SLAs
  - Need to open up control?
- Provision of Computational Platforms

Network Resource Management

- Need a better handle on what's going on
- Need better handle on user utility (No longer just the elephants and the mice.)
- Eg DiffServ with automatic SLA provisioning
- Eg Resource Assured Virtual Networks with automatic SLA provisioning
- Or just congestion pricing….

Grid and Dynamic Virtual Networks

- Want the ability on various timescales (the smaller ones are more interesting) to configure a subset of network resources to support a particular application.
- Solving the interesting problem solves quite a bit of the boring one (self provisioning circuits)
- Virtual organisation maps to virtual network

Virtual Networks

- Rather over abused term.
- Here I mean virtual in the sense that it “looks like a network” to those who are members.
- “Looks like a network” means it has real resources and can give guarantees
- Not just partitioning address space!
- Probably should say QoS VPNs to hit the jargon.

Virtual Networks: One I did earlier

- What follows is a quick run through some work on resource partitioning of networks.
- Work was based on ATM and then MPLS.
- One of the aims was for safe network programability (control your own virtual network but no one elses) but that's not really the point here.

The Data Path

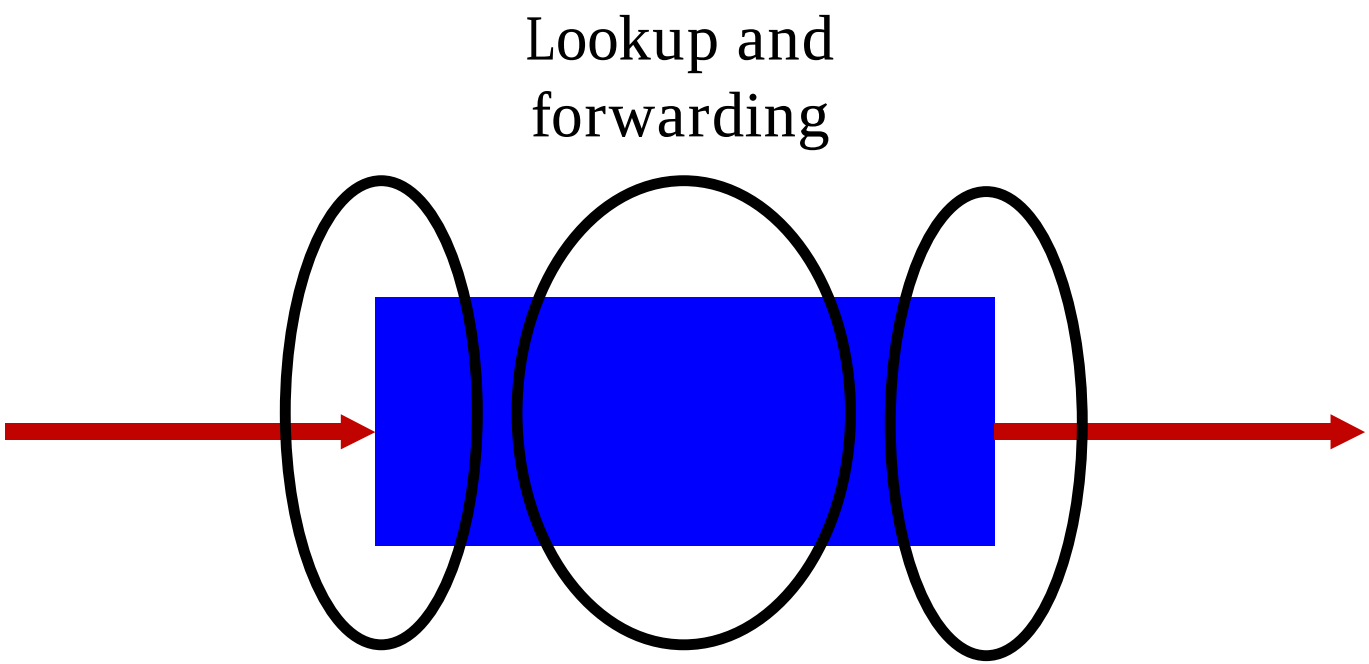

Data Path

- Forwarding path can be a variety of technologies:
  - IP forwarding, particularly MPLS !!!
  - ATM switching
  - Circuit switching
  - WDM switching

Configuration

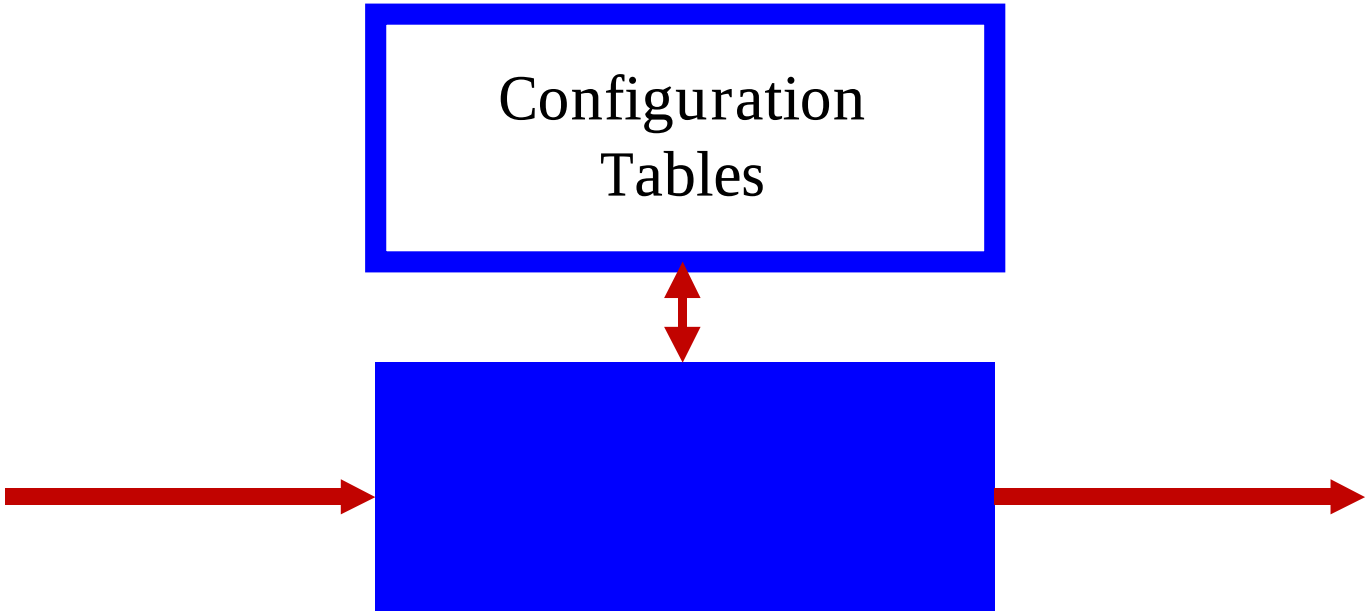

Control

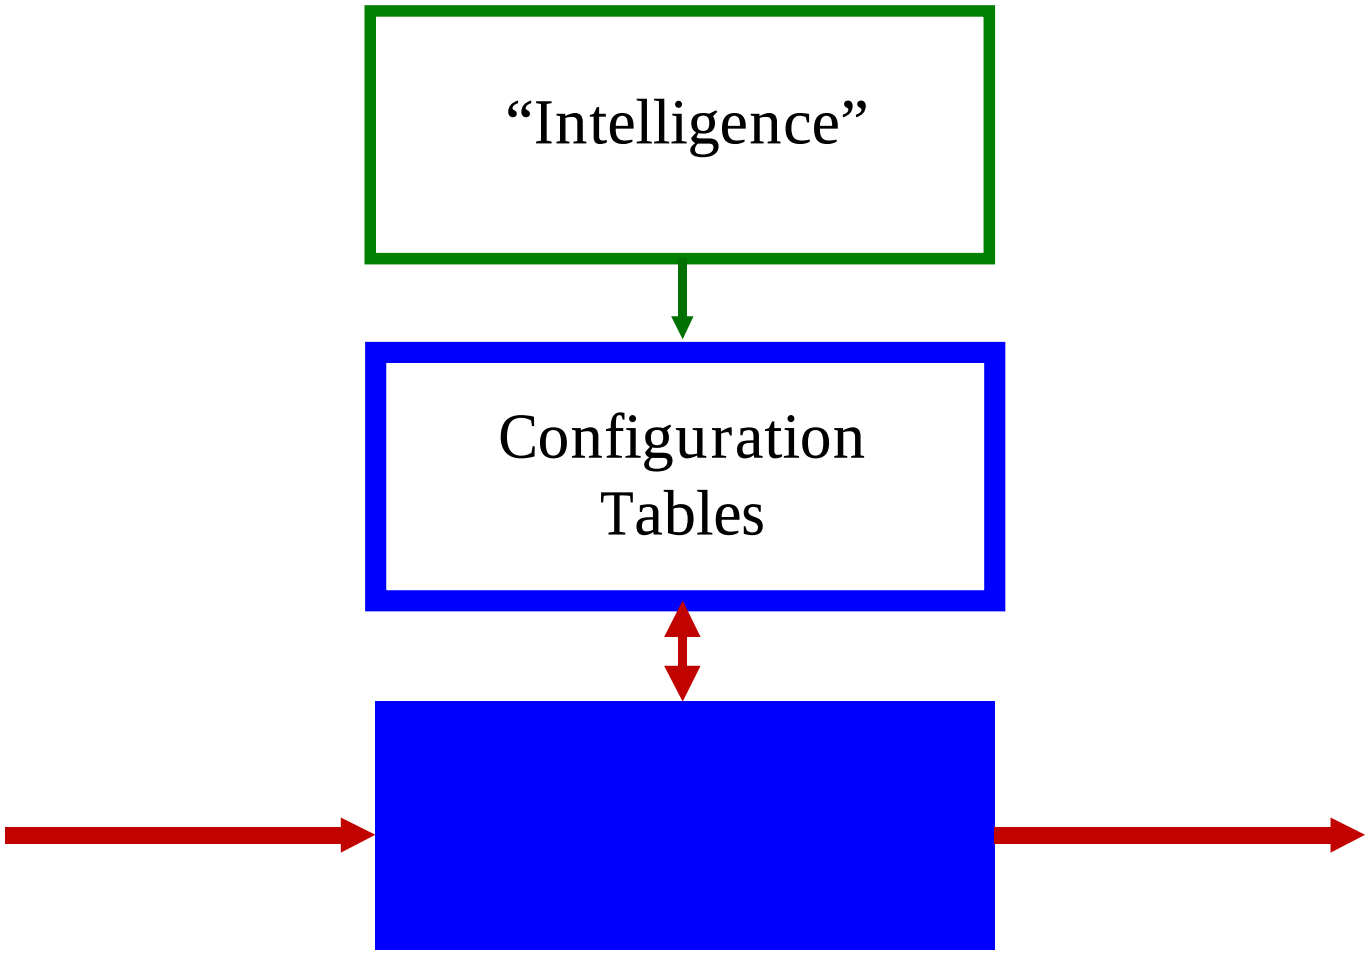

Control Path

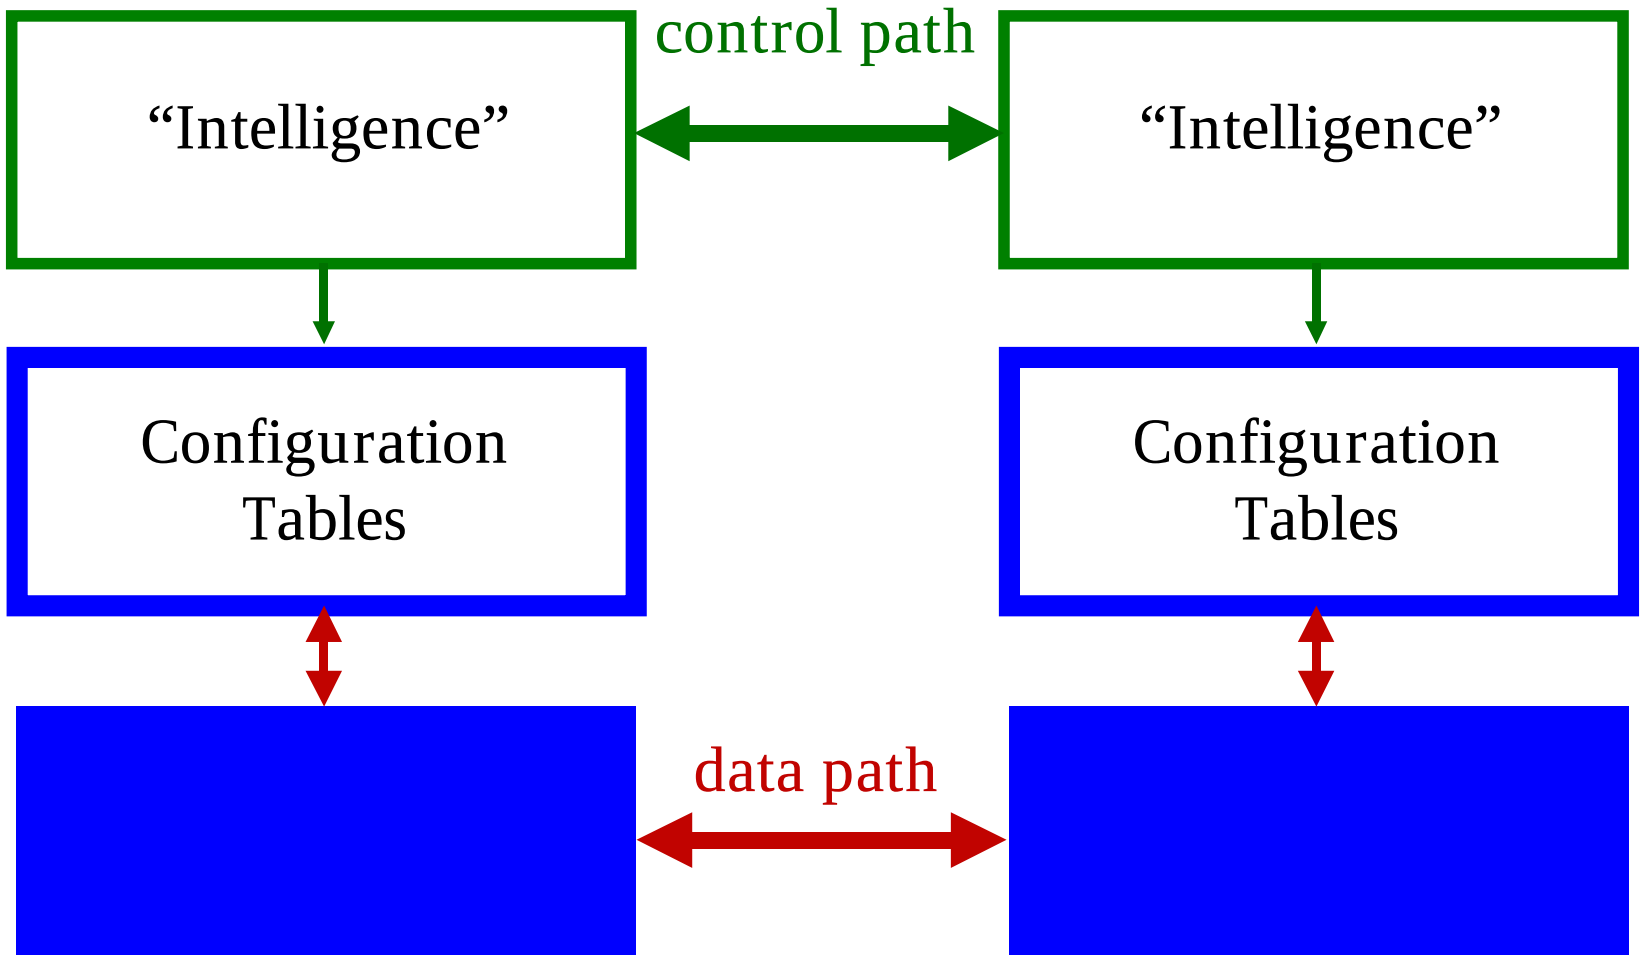

Switchlets 1: Mainframe Model

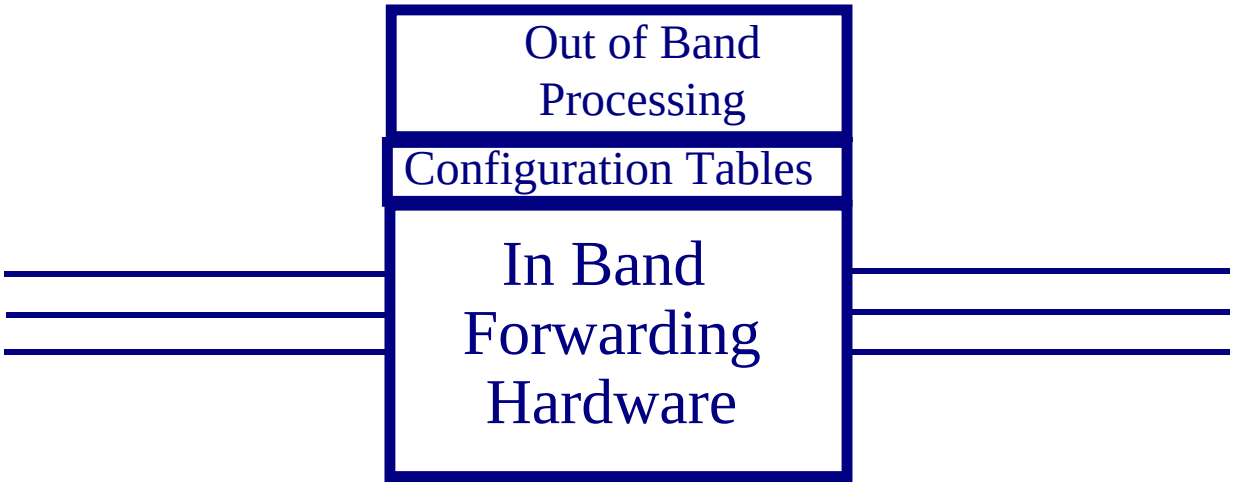

Switchlets 2: Open Control

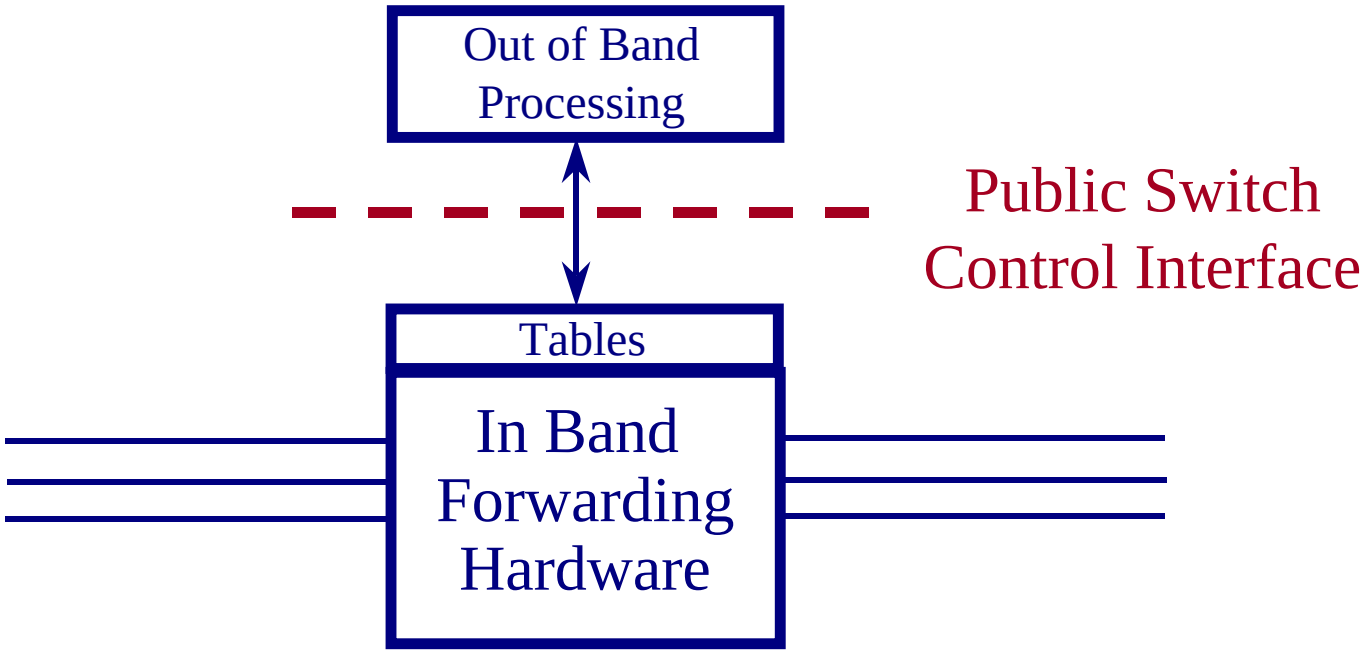

Switchlets 3: Multiple Control Systems

Simultaneous Control Architectures

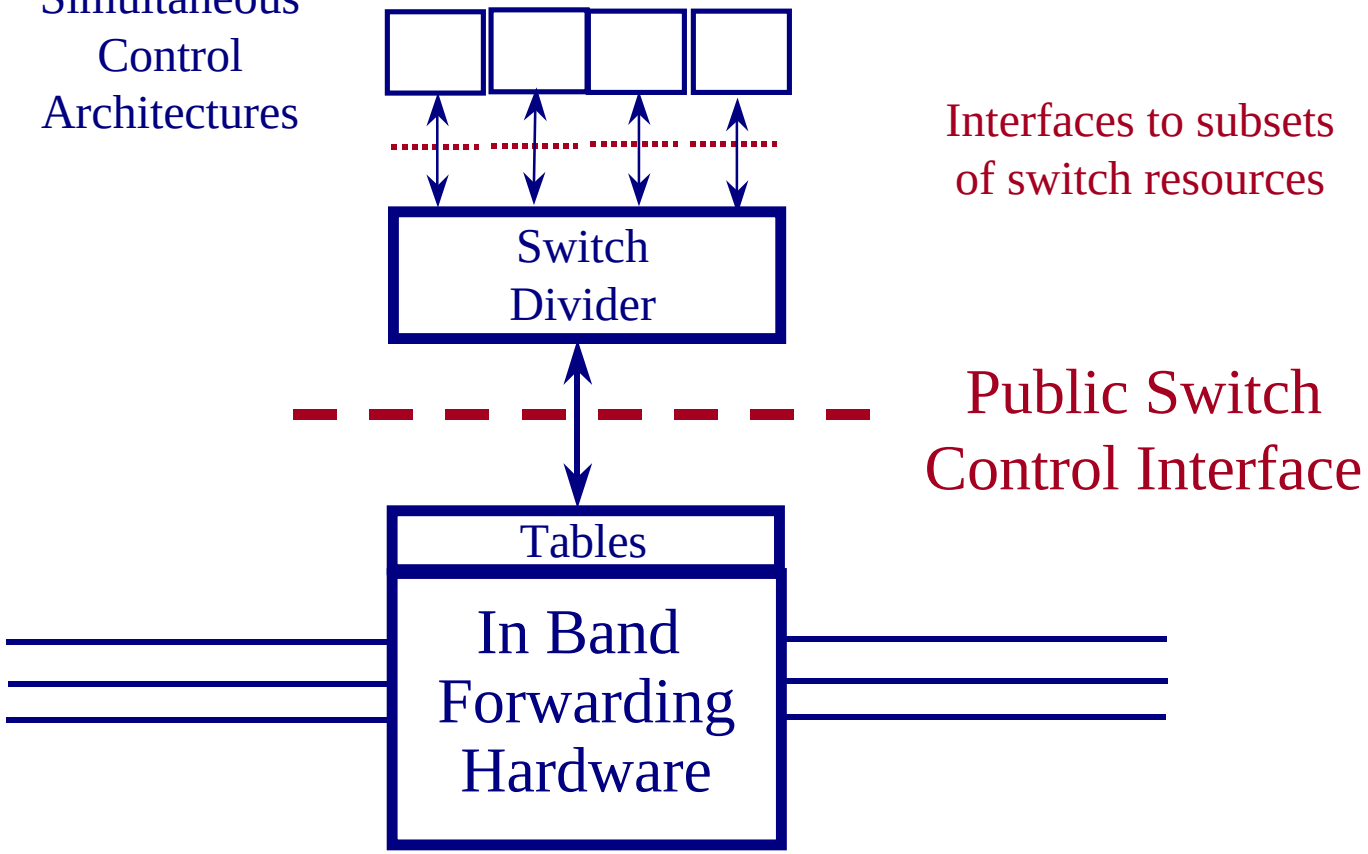

Simple WDM Example

- Consider a simple wavelength switch with no wavelength conversion and 8 wavelengths on each link $\lambda_0\lambda_1...\lambda_7$

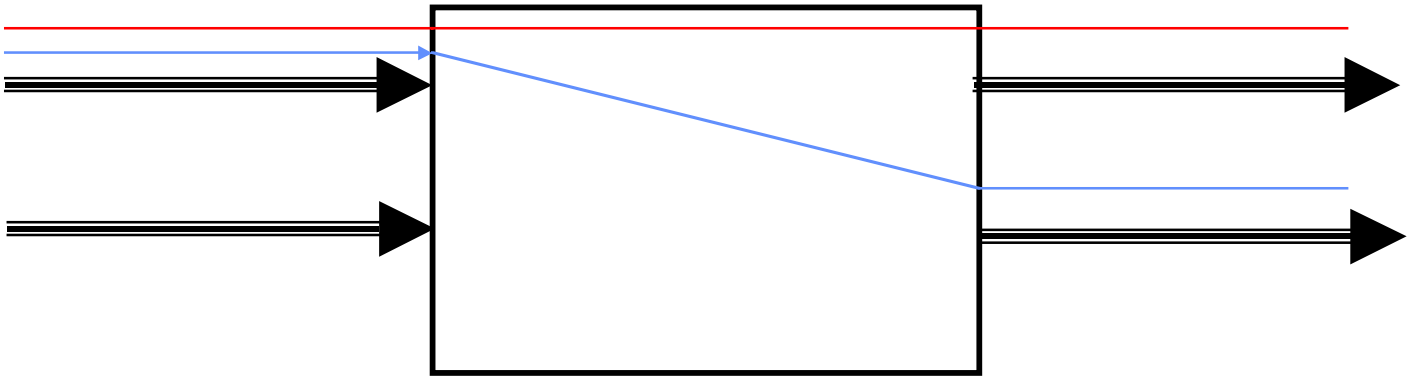

Simple WDM Example(2)

- Tables are very simple:
  - F(input link, λ) -> output link
- With constraints that λ's on each output link are different

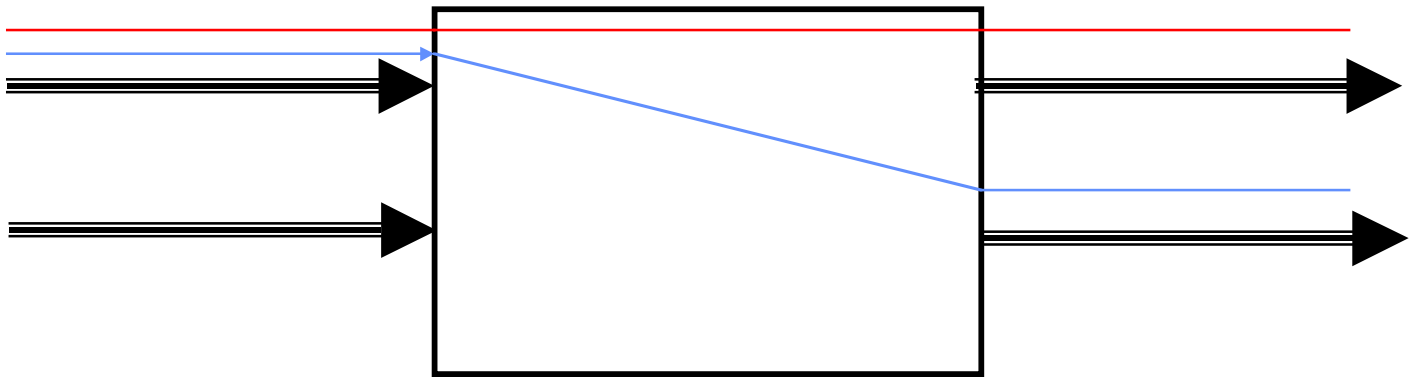

Simple WDM Example(3)

- Network programmability provides the ability to modify the table (subject to the constraints) to users of the network.

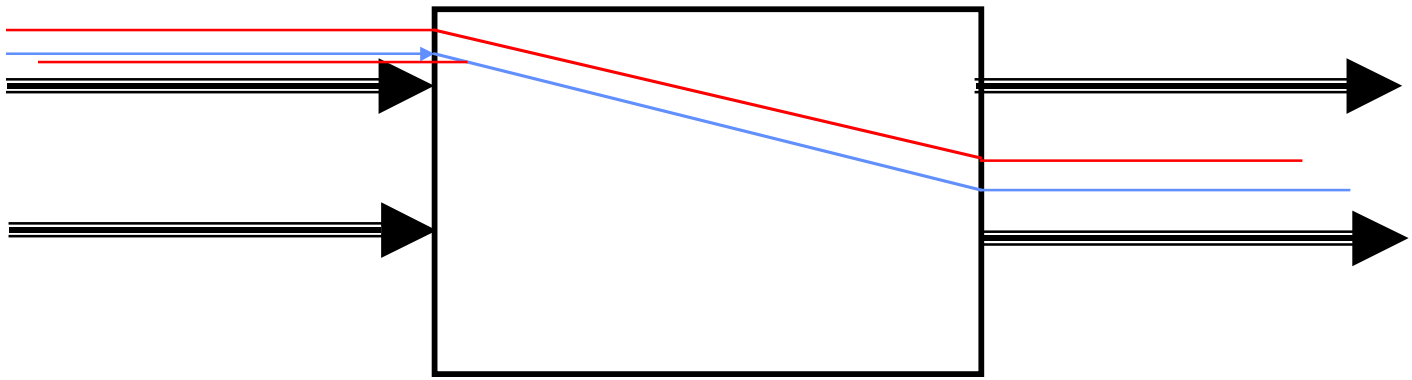

Switchlets™ 4: Partition Control

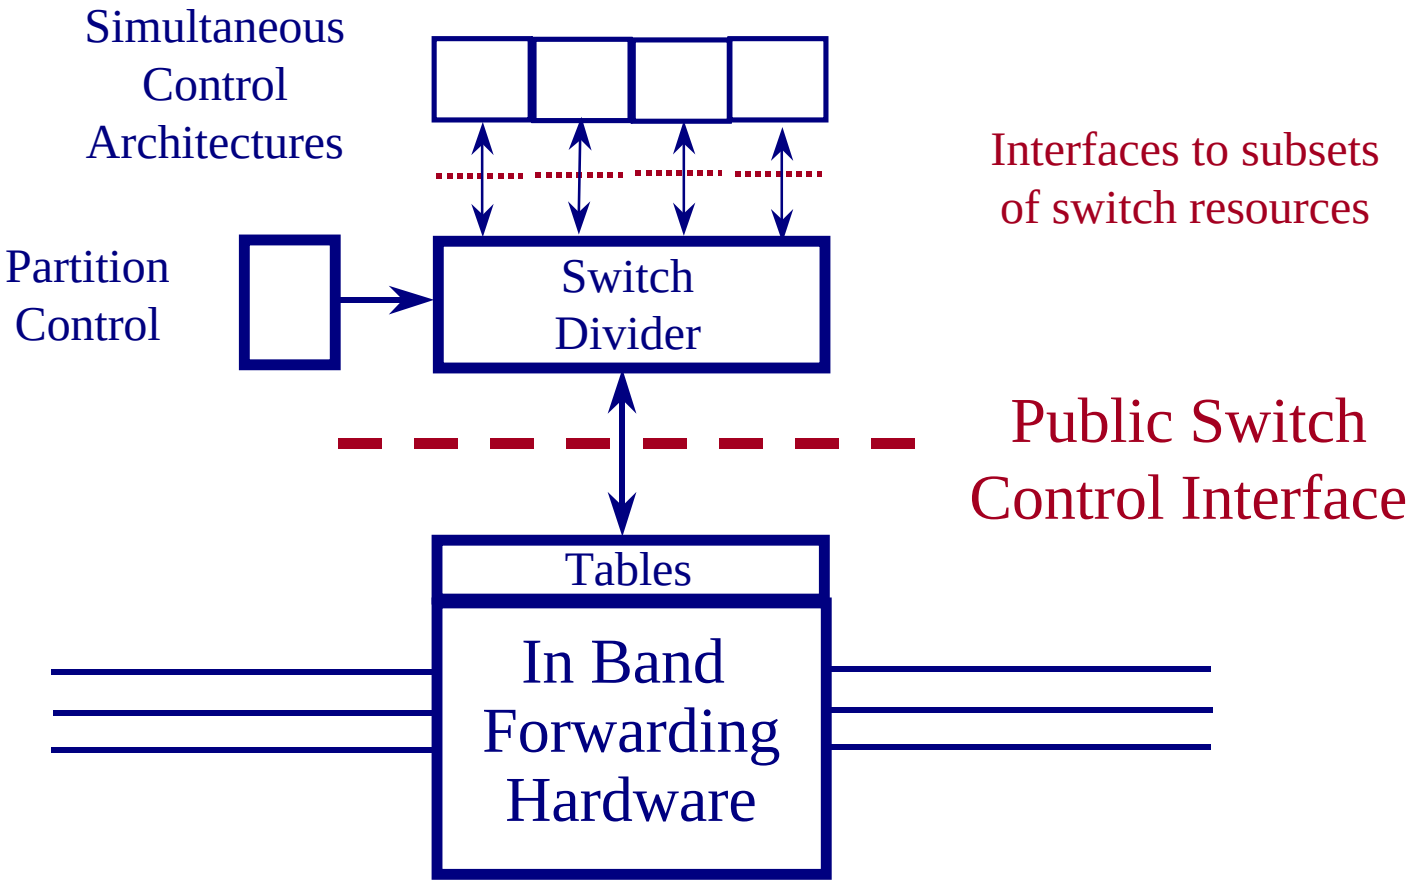

Simple WDM Example (4)

- A switchlet in the WDM example would be a partition of the wavelengths and links:

- switchlet 1 could be λ_0, λ_1 on links 2,3,4
- switchlet 2 could be λ_2, λ_3 on links 2,3,4
- switchlet 3 could be λ_0, λ_1, λ_2, λ_3 on links 1 and 5
- Requests to change mappings would check that request was on behalf of appropriate switchlet

Simple WDM Example (5)

- In the switchlet example:

  switchlet 1 : λ_0, λ_1 on links 2,3,4

  switchlet 2 : λ_2, λ_3 on links 2,3,4

  switchlet 3 : λ_0, λ_1 , λ_2, λ_3 on links 1 and 5

  Move λ_0, λ_1 on links 1 and 5 to switchlet 1

Using Dynamic VPNs

- Can move resource dynamically between virtual networks (eg time of day, disaster recovery)
- Can have relatively short lived virtual networks
- Even with identical control architectures, we have resource partitioning (VPN SLAs)

Building VPNs

- Take a specification of a VPN:

  endpoints, traffic matrices, reliability, etc, Control Policy

  Allocate resource and create switchlets dynamically, find computational resource and instantiate Control Policy

So what?

- Useful to think of four activities
  - Data path forwarding
  - Control path configuration
  - Component control for resource partitioning
  - Virtual Network Service Provision

Virtual Network Service

- Self provisioning of resource assured virtual networks
- Client provides description and authority (money??) in request to create
- We have needed this service since early 1990's when attempts to do *network* research on SuperJanet began (I personally gave up in ~1995)
- Grid will demand it (when they figure out what it is)