Endpoint Control techniques & timescales

Panos Gevros, Jon Crowcroft
University of Cambridge
panos.gevros@cl.cam.ac.uk

TAPAS Project Meeting, Bologna, Italy
25-26 September 2003

TAPAS

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

Connection with previous work

- we have identified the similarities between ASP and ISP but focused on the ISP
- presented a high-level discussion of the available options to the ASP (common access ISPs, strategic co-location, issuing recommendations)
- ISP use of traffic engineering (TE) (capacity planning & delay analysis) for meeting SLAs

Problem statement

Optimize performance to/from destinations of interest

(s.t. cost constraints)

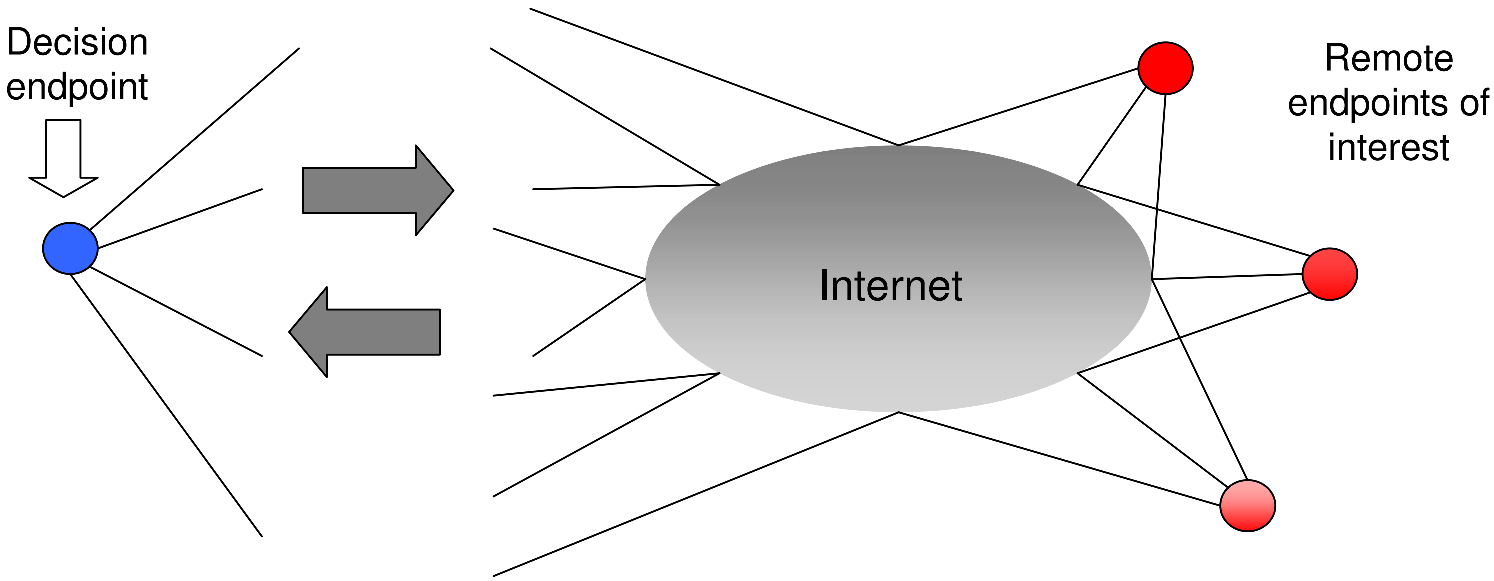

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

Assumptions

- 1 there is a market for access to global Internet with several access ISPs offering internet connectivity services in the same area (in certain areas can be >50 ISPs)
- 2 performance (or QoS) measured by (delay, loss, throughput) most often used metric delay
- 3 performance to different parts of the network varies with upstream provider - reason: path diversity across providers in a given location.
- 4 :endpoint is a stub network in a metropolitan area and has no control over the end-to-end path, can only choose 1st hop provider(s).

Optimisation techniques and timescales

Months/weeks	*Provider* selection (multihoming)
days/hours	*Route* selection
min/msec	*Control behaviour* selection

two distinct **perspectives** (from the endpoints) :

- send side (- content provider, ASP, web site, data center –)
- recv side (- enterprise, subscriber -)

Provider selection for multi-homing

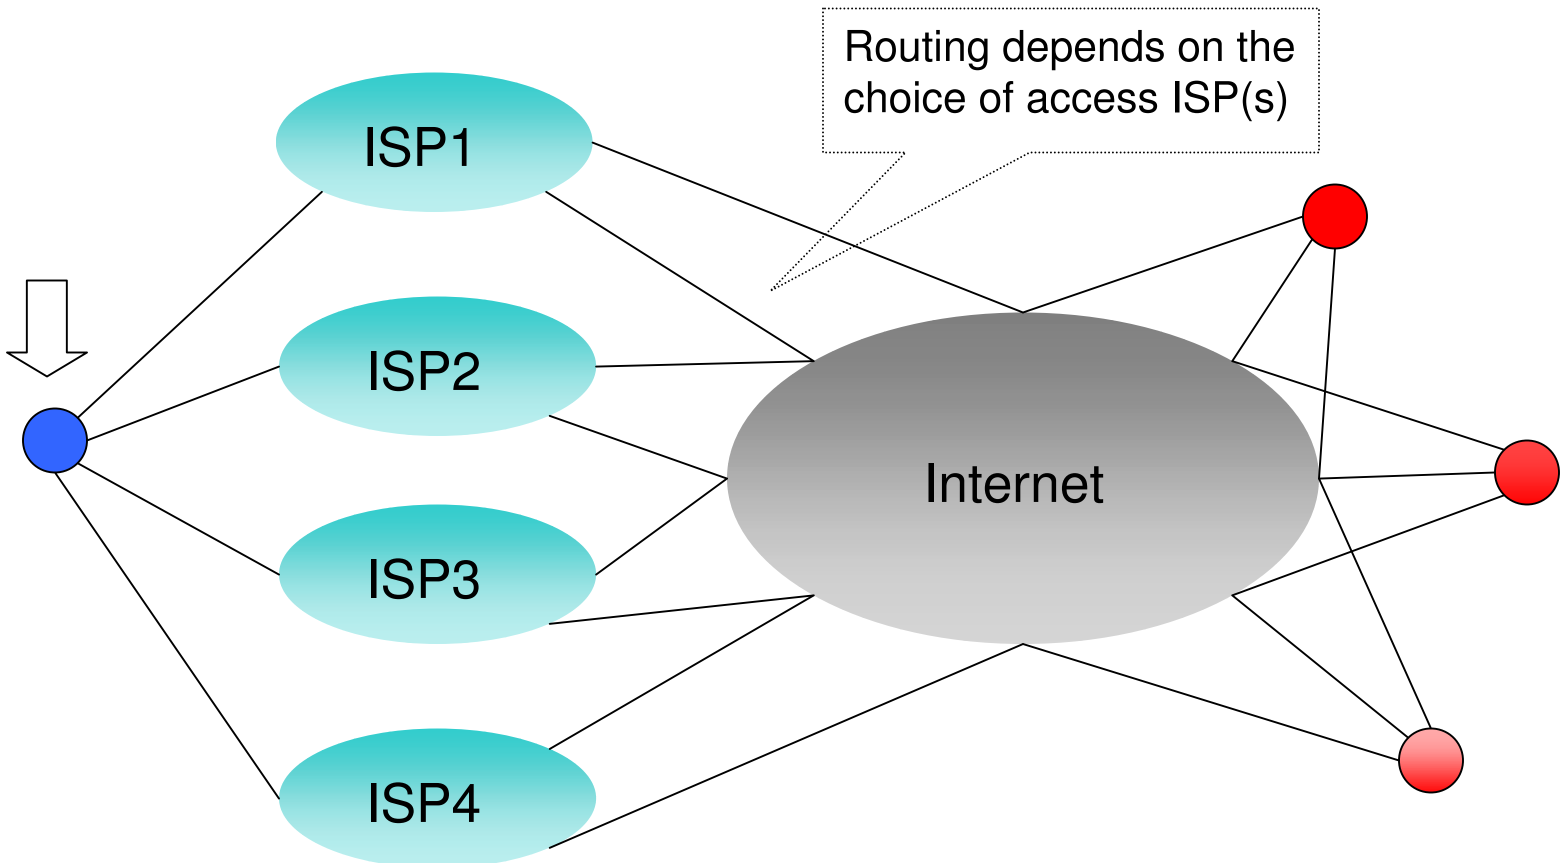

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

Stage 1: multi-homing

- Multi-homing : having > 2 external links to the same or different providers
- *k*-multihoming employ *k* ISPs
- k is usually small diminishing returns after a point (k>4)
- **reasons** for multi-homing :
  - fault tolerance
  - optimizing performance
  - reducing bandwidth costs
- requires having an AS# , running BGP,
- model the decision process for the selection of a subset of ISPs
- evaluate the relative benefits from provider selection from several available options

Stage 1: multi-homing

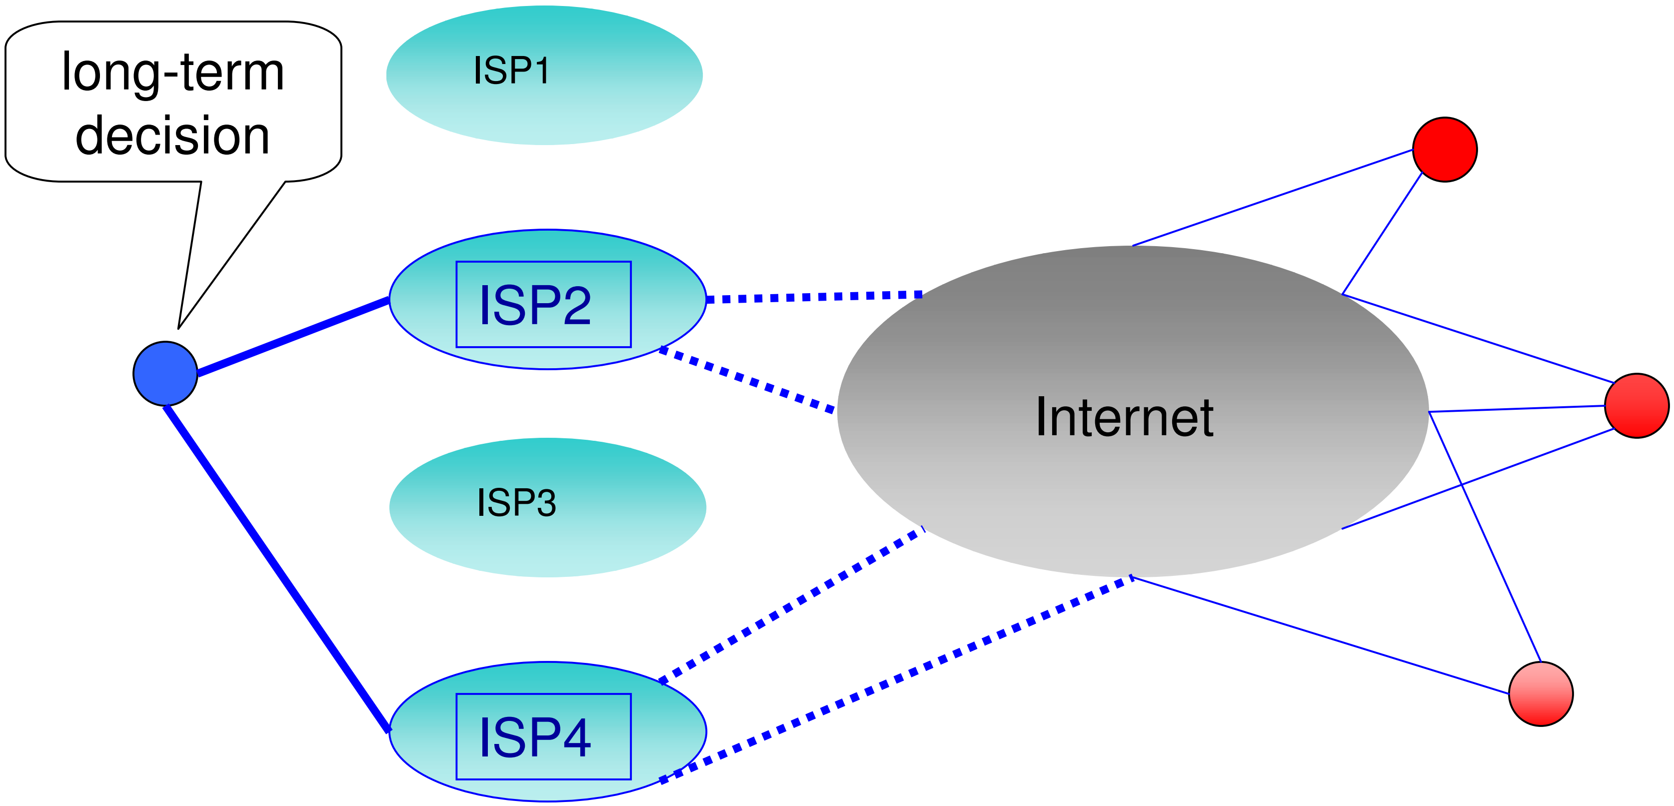

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

Modelling Provider Selection (1)

Available ISPs	j = 1, 2, … P
Destinations (nets or ASes)	i =1, 2, … N

Volume (MB)	Cost ($/MB)	Delay (msec)
v_i	c_i	d_i

- information available to the endpoint measured over all destinations (for all *I=1..N*) over a period of time.
- endpoint wants to select **k** (out of P) providers (k >=2)
- **delay vs cost** optimization
- we are interested in the properties of the delay (D) distribution.

TAPAS

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

Modelling Provider Selection (2)

Total Transit Costs

$$C = \sum_{i=1}^{N} v_i c_i$$

Total Traffic

$$V = \sum_{i=1}^{N} v_i$$

Weight per destination (probability)

$$w_i = \frac{v_i}{V}$$

Delay distribution (expected, stddev)

$$\mu = E[D] = \sum_{i=1}^{N} w_i d_i$$

$$\sigma = \sqrt{E[(D-\mu)^2]}$$

TAPAS

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

Modelling Provider Selection (3)

Optimisation Objectives:

I. Minimize Total Costs C s.t. <= E(D) stddev(D)

II. Minimize Delay Variance stddev(D) s.t. <= E(D) and C

III. Minimize Expected Delay E(D) s.t. <= stddev(D) and C

Modelling Provider Selection (3)

Endpoint wants a multihoming contract for the next period - submits its contract requirements, revealing only that the *average delay to* **destination-i** *should not exceed a delay threshold.*

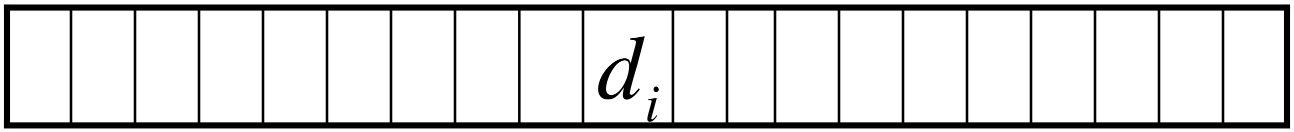

ISP- j responds with a bid for the contract specifying a maximum for the average delay for each destination, the volume that applies and its tariff for carrying the traffic. Normally there are **P** bids for the contract

The volume is spread over the period of the contract – burst rate limits may apply.

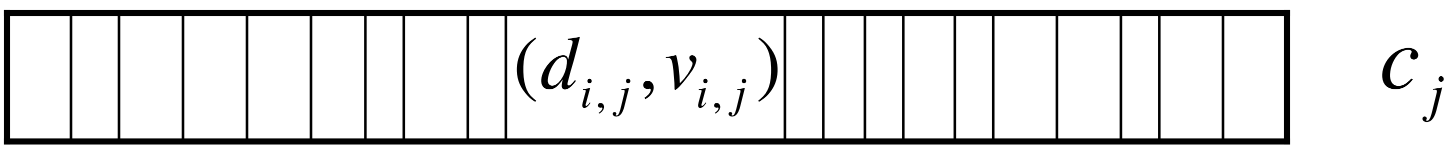

TAPAS

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

Modelling Provider Selection (4)

Search through the (k out of P) combinations

Compute the *new* cost and delay distribution parameters using the *minimum of the union* of the parameters in the current combination

Test if any of the Optimisation Objectives (I,II, III) is being met (compare with the current cost, delay and volume parameters)

For every objective keep a list of those combinations that satisfy it

Select the k-combination based on idiosyncratic preferences (i.e. a combination that minimizes cost may be preferable to one that maintains the same cost and minimizes the variance and the expected delay, will depend on higher level objectives).

Stage 2: intelligent route control

- assumes multi-homing
- Intelligent Route Control : selecting which ISP to use for send/recv traffic to/from specific remote subnets (or ASes)
- provides more fine-grained control : based on feedback control loop
- obtain delay measurement (periodic ping), compute ISP delay ratio, alter BGP routing tables, (or NAT box or DNS)

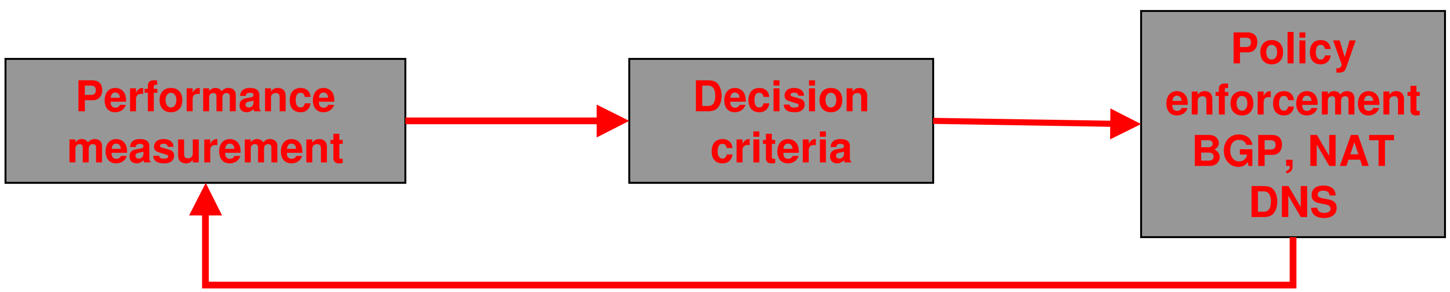

Stage 2: intelligent route control

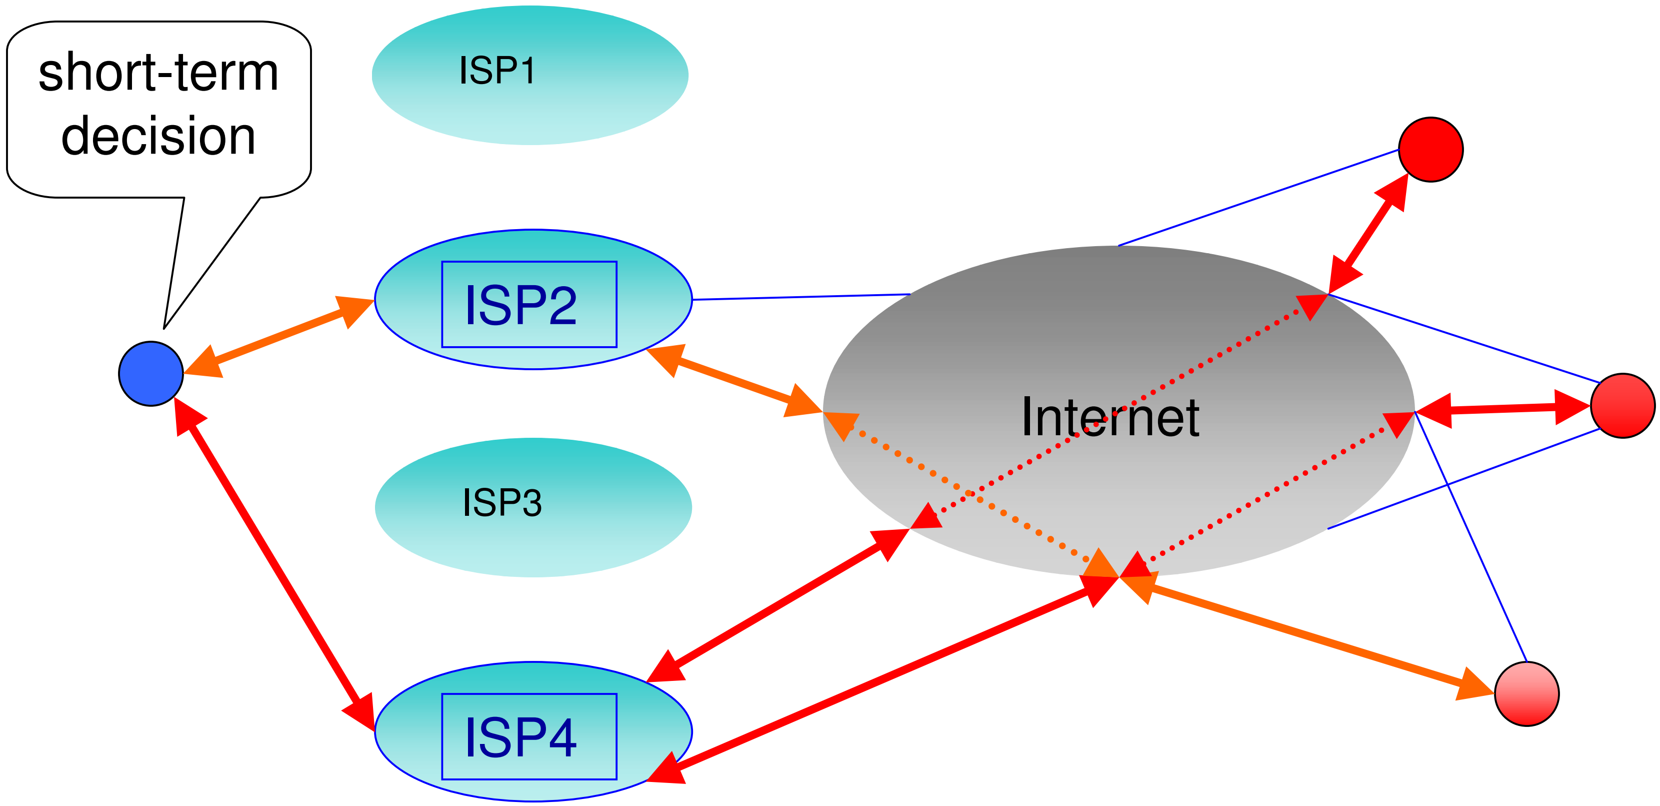

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

Stage 2: intelligent route control

two distinct **perspectives** (from the endpoints) :

- send side (- content provider, ASP, web site, data center –)
- recv side (- enterprise, subscriber -)

Stage 2: intelligent route control techniques

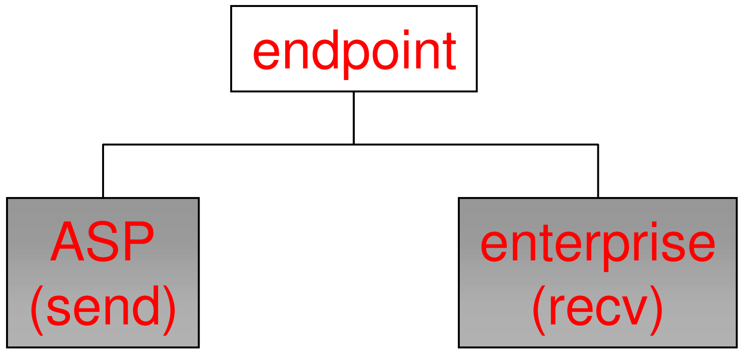

Dynamically direct traffic over the best performing ISP.

BGP filter : control which routes to accept from each ISP, based on performance measurements of different paths

Control routing announcements to different ISPs…(slow)

Outgoing connections: have address ranges from different ISPs, dynamically assign addresses from the best performing ISP (NAT).

Incoming connections: DNS returns address from the ISP that provides the best performance to the initiating host.

TAPAS

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

Stage 2: intelligent route control

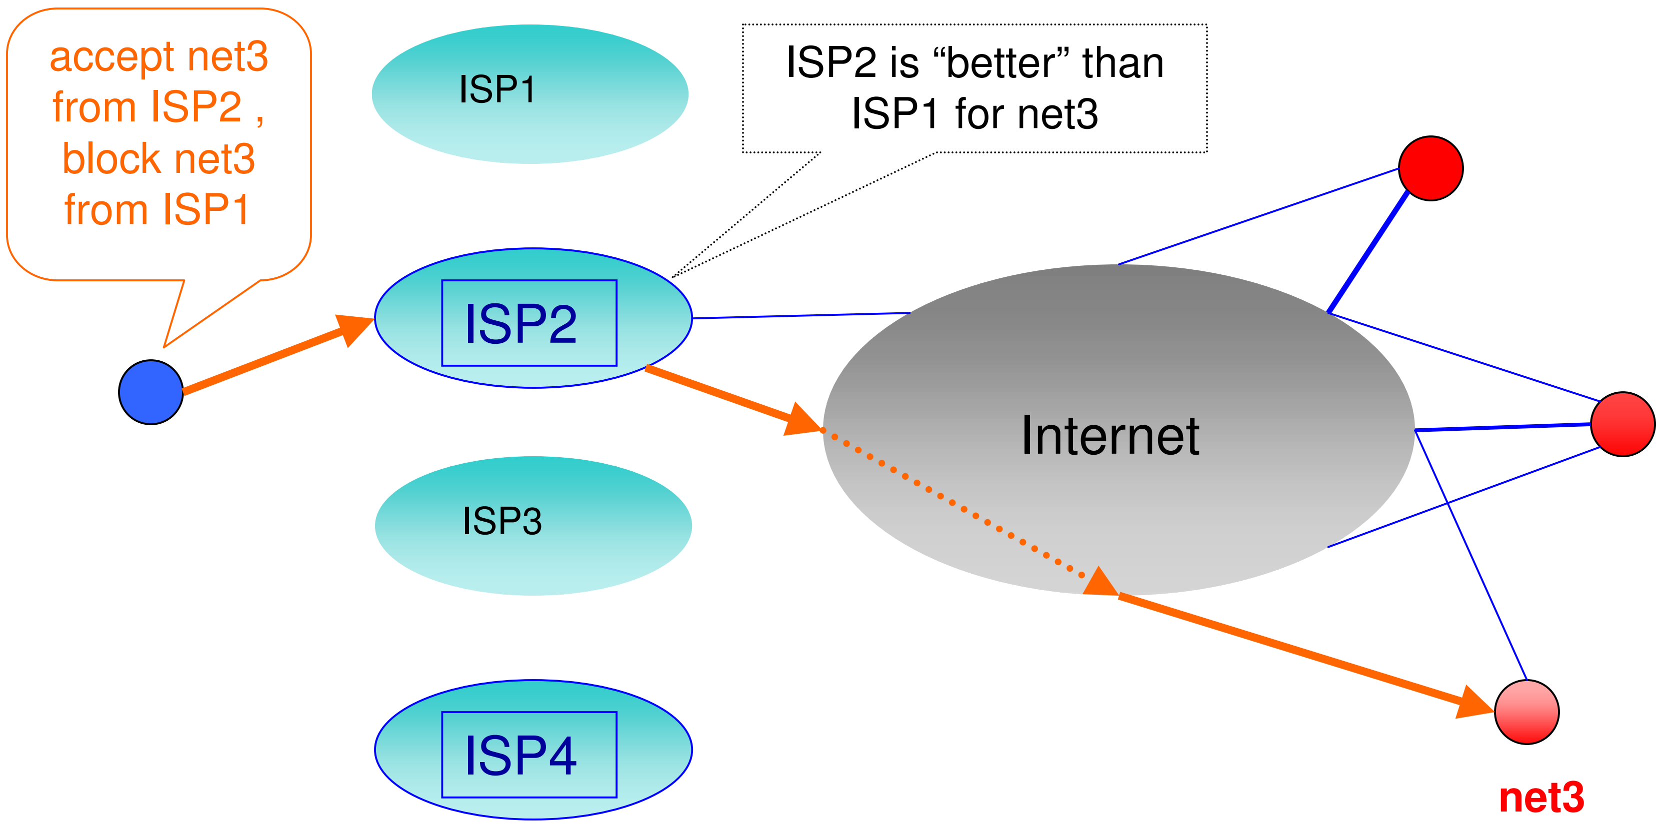

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

Summary

- presented a framework for the control options available to the Internet endpoints for influencing the performance of their end-to-end connections at different time-scales. In particular we discussed:
- *provider selection (or k-multihoming)* based on cost and performance (delay) criteria and a *formalism for the decision making process* (we have an early prototype in Matlab and an SQL database with traffic statistics for different networks/ASes, volume, delay as seen from UoCambridge, Computer Lab)
- *intelligent route control* in a multi-homing scenario (such products exist already by companies Sockeye, Routescience).
- transmission control behaviours Transport Options will be treated separately.

UNIVERSITY OF CAMBRIDGE
Computer Laboratory