Re'Arch 2008

Policing Freedom…

to use the Internet Resource Pool

Arnaud.Jacquet, Bob.Briscoe, Toby.Moncaster {@bt.com}

December 9 2008

Agenda

- Architectural choices for policing usage
- Design of a bulk congestion policer
- Impact on traffic
- Implications on congestion signals

Policing usage – state of affairs

- Distributed resource control
- Many parties involved in the outcome
- Fair usage policies on broadband services
Techniques	Assumptions
volume caps	each packet has the same packet
fair queuing	single access bottleneck
deep packet inspection	application type implies congestion
- They limit flexibility to shift usage (over links and time) around the Internet resource pool, and prevent evolution towards more efficient rate adaptation

Policing usage – what to change

- What matters is usage of scarce resources, as reflected by congestion
- Each packet is accountable for the congestion it causes on its path
- It is possible to monitor accountability of any collection of flows
- For any accountable party, monitor and control
  - Congestion volume (rather than volume)
  - Congestion bit rate (rather than throughput)
- Granularity of resource usage accountability
  - Not per flow (can open several in parallel)
  - Per customer, where there is a contractual relationship
- Congestion pricing leads to dynamic prices ☹
- Congestion policing is the rationing version
- To enforce such policies at the technical level, we need to consider control mechanisms (policing) and interfaces (signalling)

Architectural considerations

- Policing is located at the 'enforcement point' where a customer attaches, rather than at network resources
  - ➔ Need for suitable congestion signalling
- Only *the overall traffic* of each customer is policed: flow isolation would limit flexibility

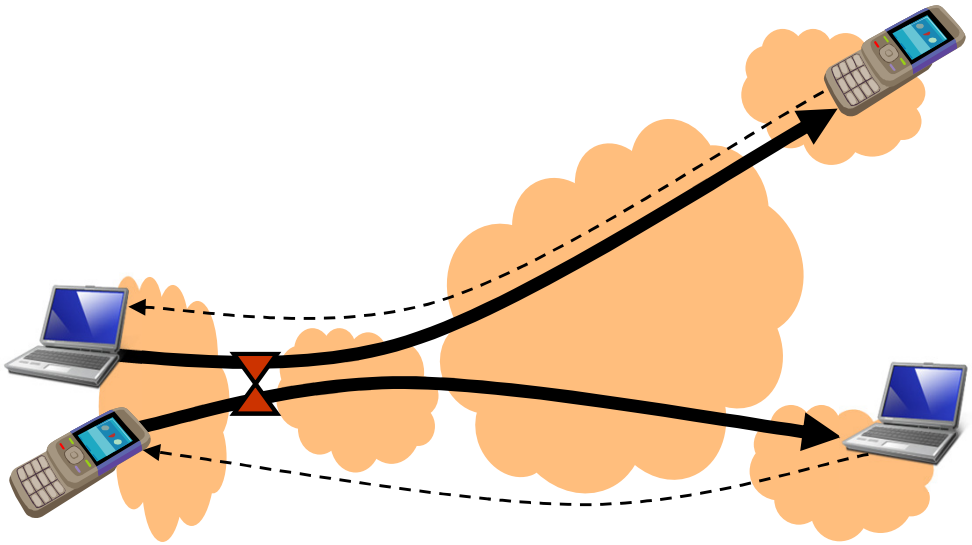

Bulk congestion policer

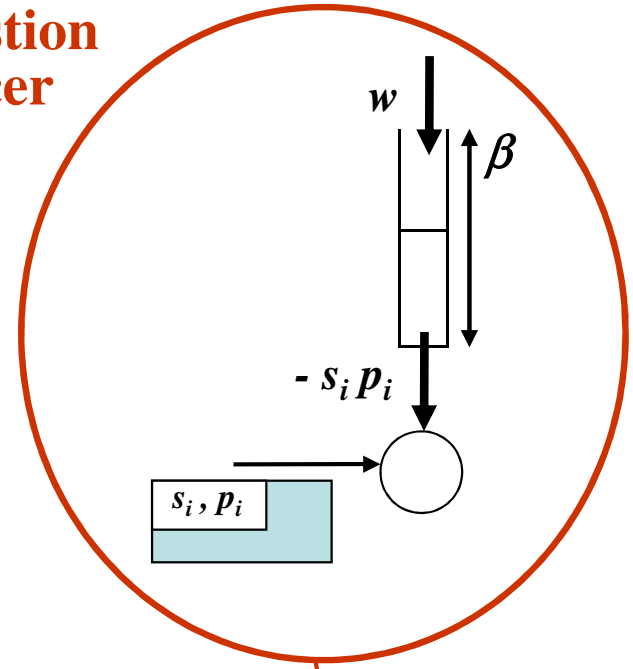

Design different to 'classic token bucket'

- still decides fate of each packet
- only congested bits consume tokens
- if not enough tokens, policer drops packet
  (alt. delay, charge..)
- sanction can be gradual
- for binary signals: policer sanction only marked packets

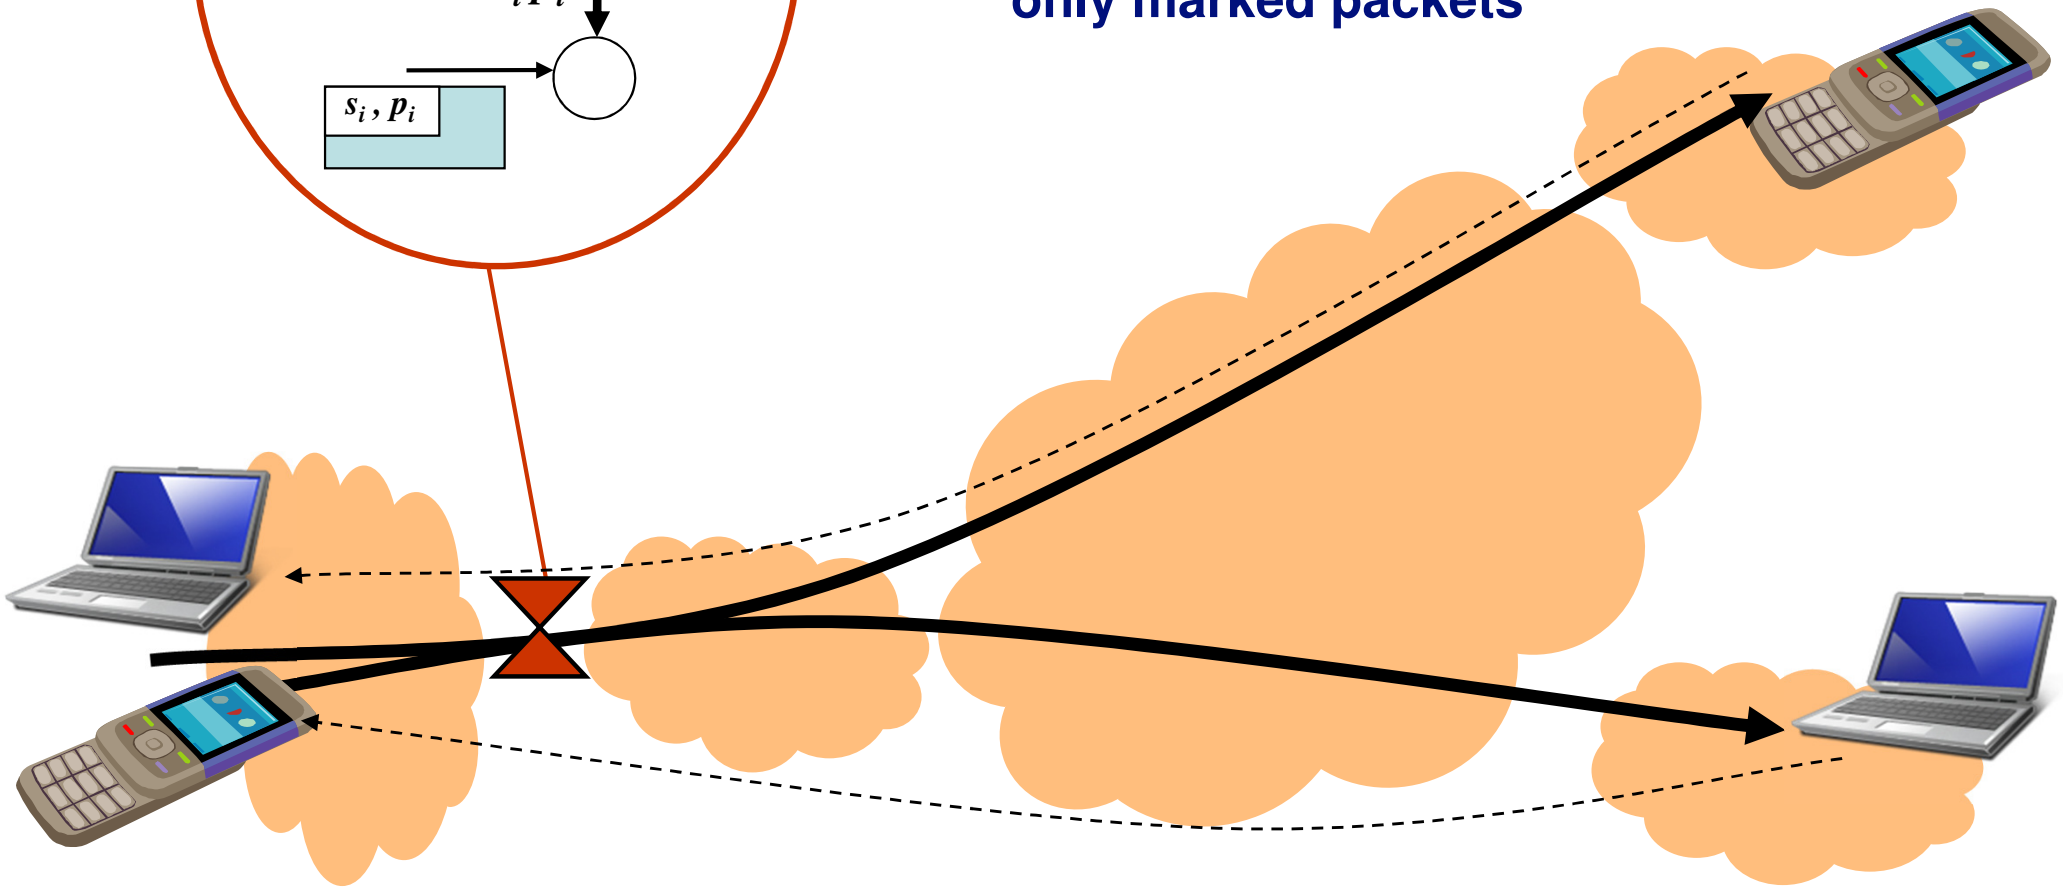

Acceptable Use Policy

Fair usage is defined by a 'congestion volume' allowance: of 1GB per month

That is equivalent to a constant congestion bit-rate of about 3kb/s

Bit-rate is otherwise unlimited

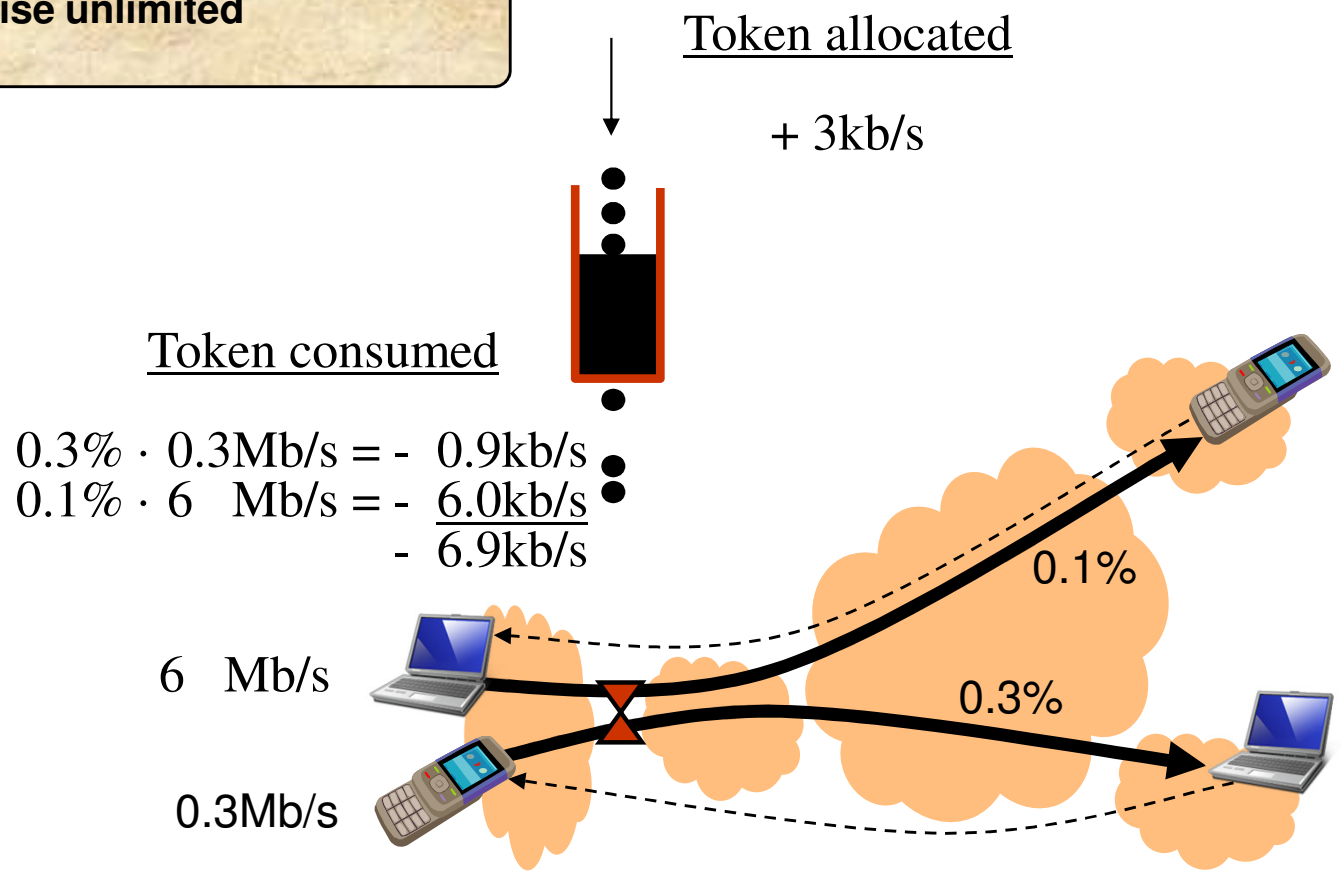

Cross-effects

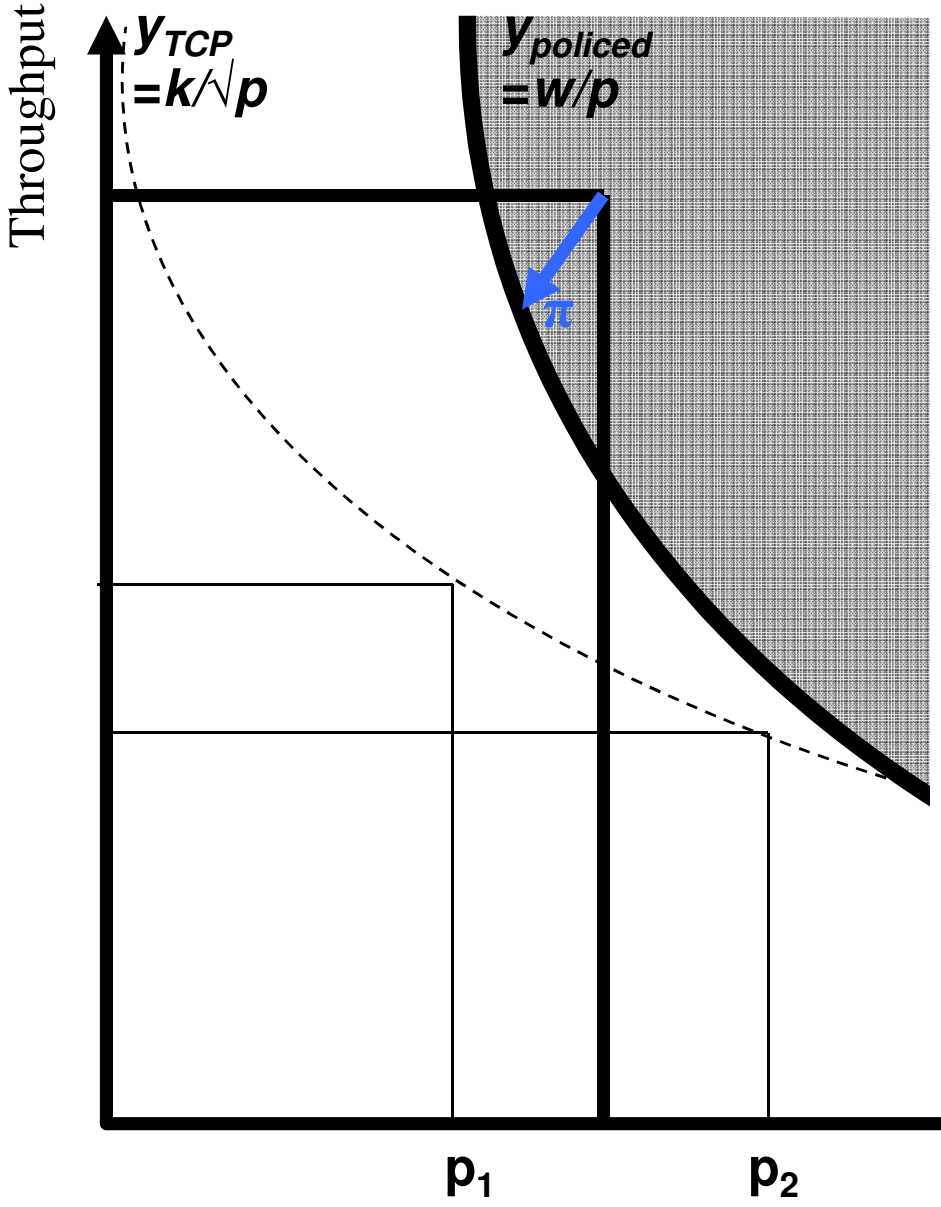

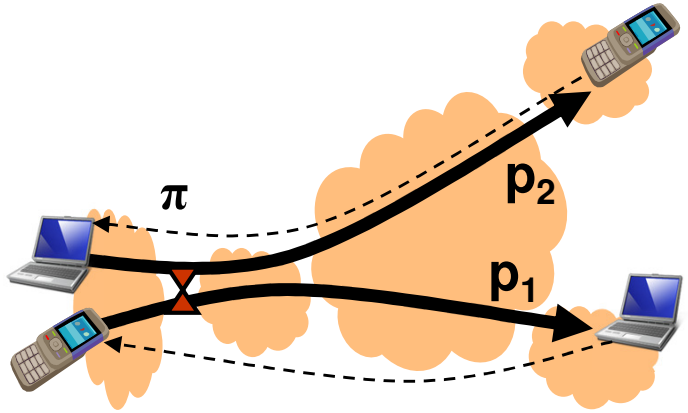

- The policer operates as a congestible resource
- When congestion volume exceeds allowance, it introduces its own congestion signal π
- … based on the congestion bit rate of the aggregate traffic

Cross-effect on responsive flow

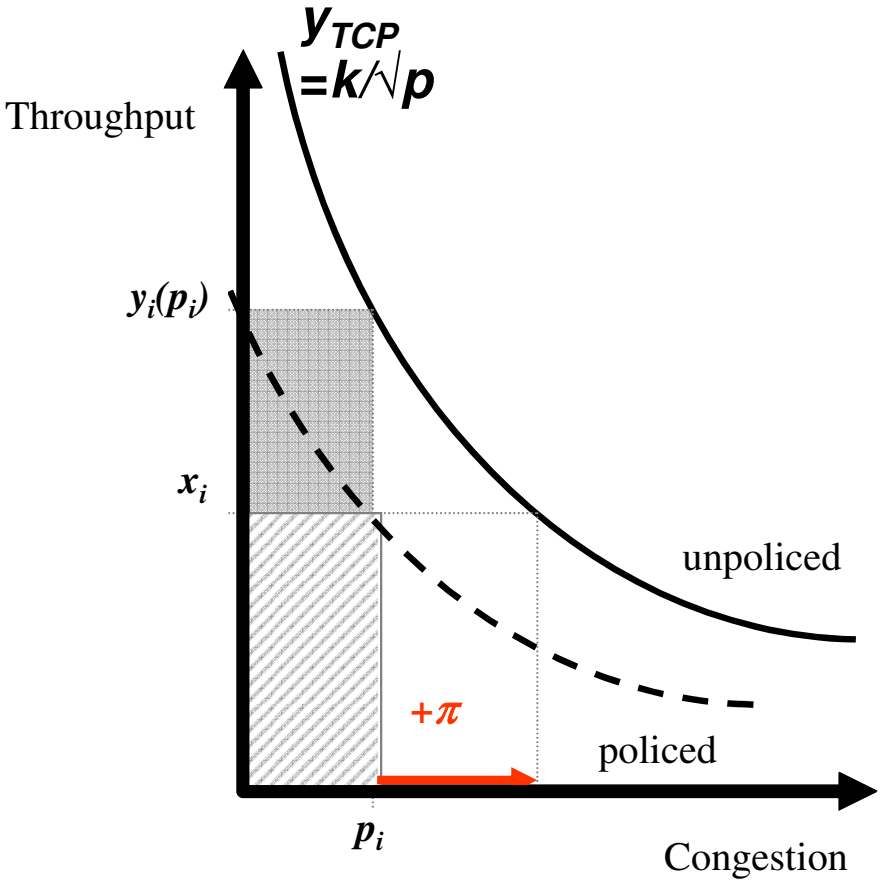

- Flow i experiences congestion p_i
- Other flows through same policer experience congestion forcing the policer to be active
- The bulk policer acts as a congestible resource with apparent congestion level π
- The figure shows how the congestion response of the flow changes from unpoliced to policed

 (π stays constant)

Cross-effect on unresponsive flow

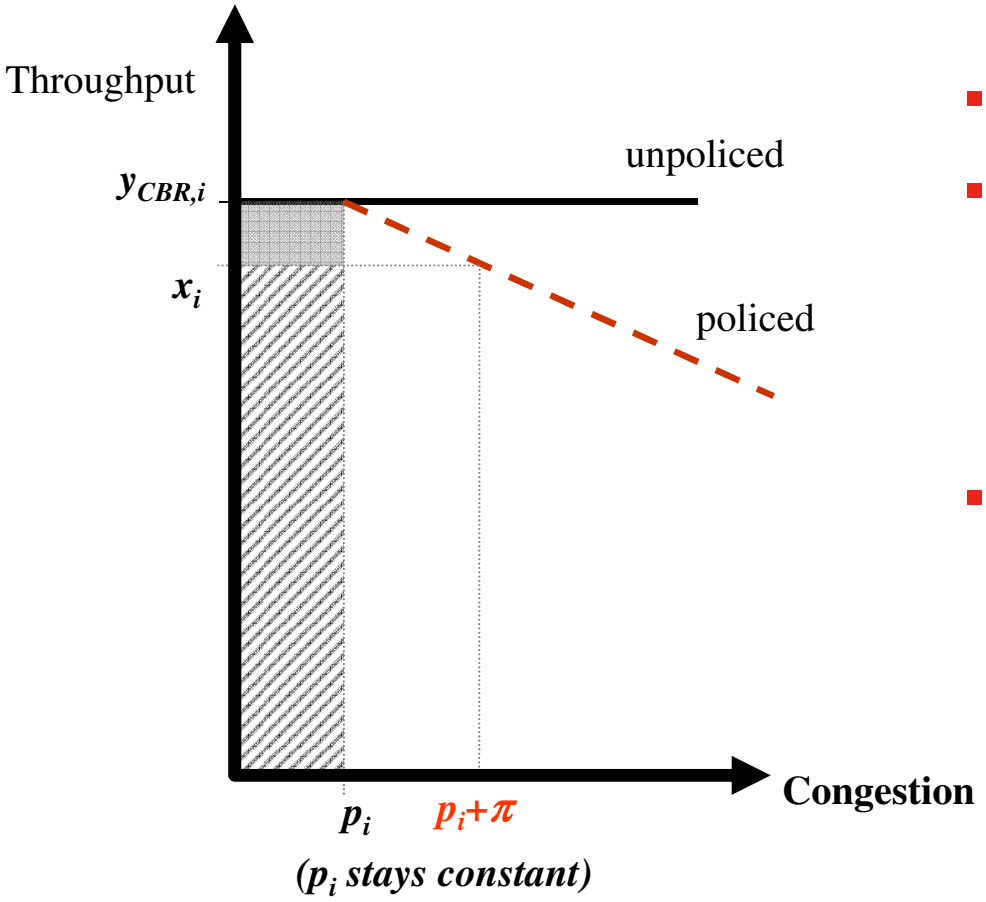

- The effect is similar with unresponsive flows
- Even an unresponsive application might be throttled on the basis of the congestion caused by other flows from the same customer
- However, responsive traffic remains more affected

Promoting self-policing

- The bulk policer imposes a joint constraint on all the traffic of a customer
- This can have disproportionate impact on some of the most valuable flows
- Thus encouraging customers to actively control the apportionment of their bit rate allowance:
  - → weighted congestion control
  - → protect foreground traffic
  - → shift background to less congested time-space

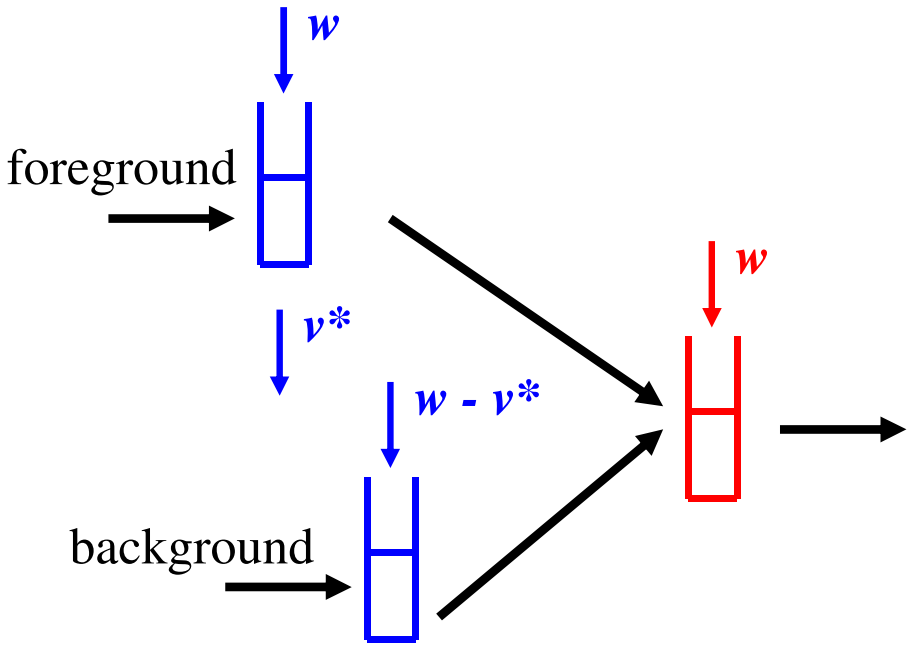

Example of resource control evolution

Requirement on signalling

- Each packet needs to signal what congestion is expected on its path
- This means each resource needs to signal congestion back to the source
  - ➔ ECN
- One way for the source to decide what to signal is to reinsert the congestion signal
  - ➔ re-feedback

One option = re-ECN

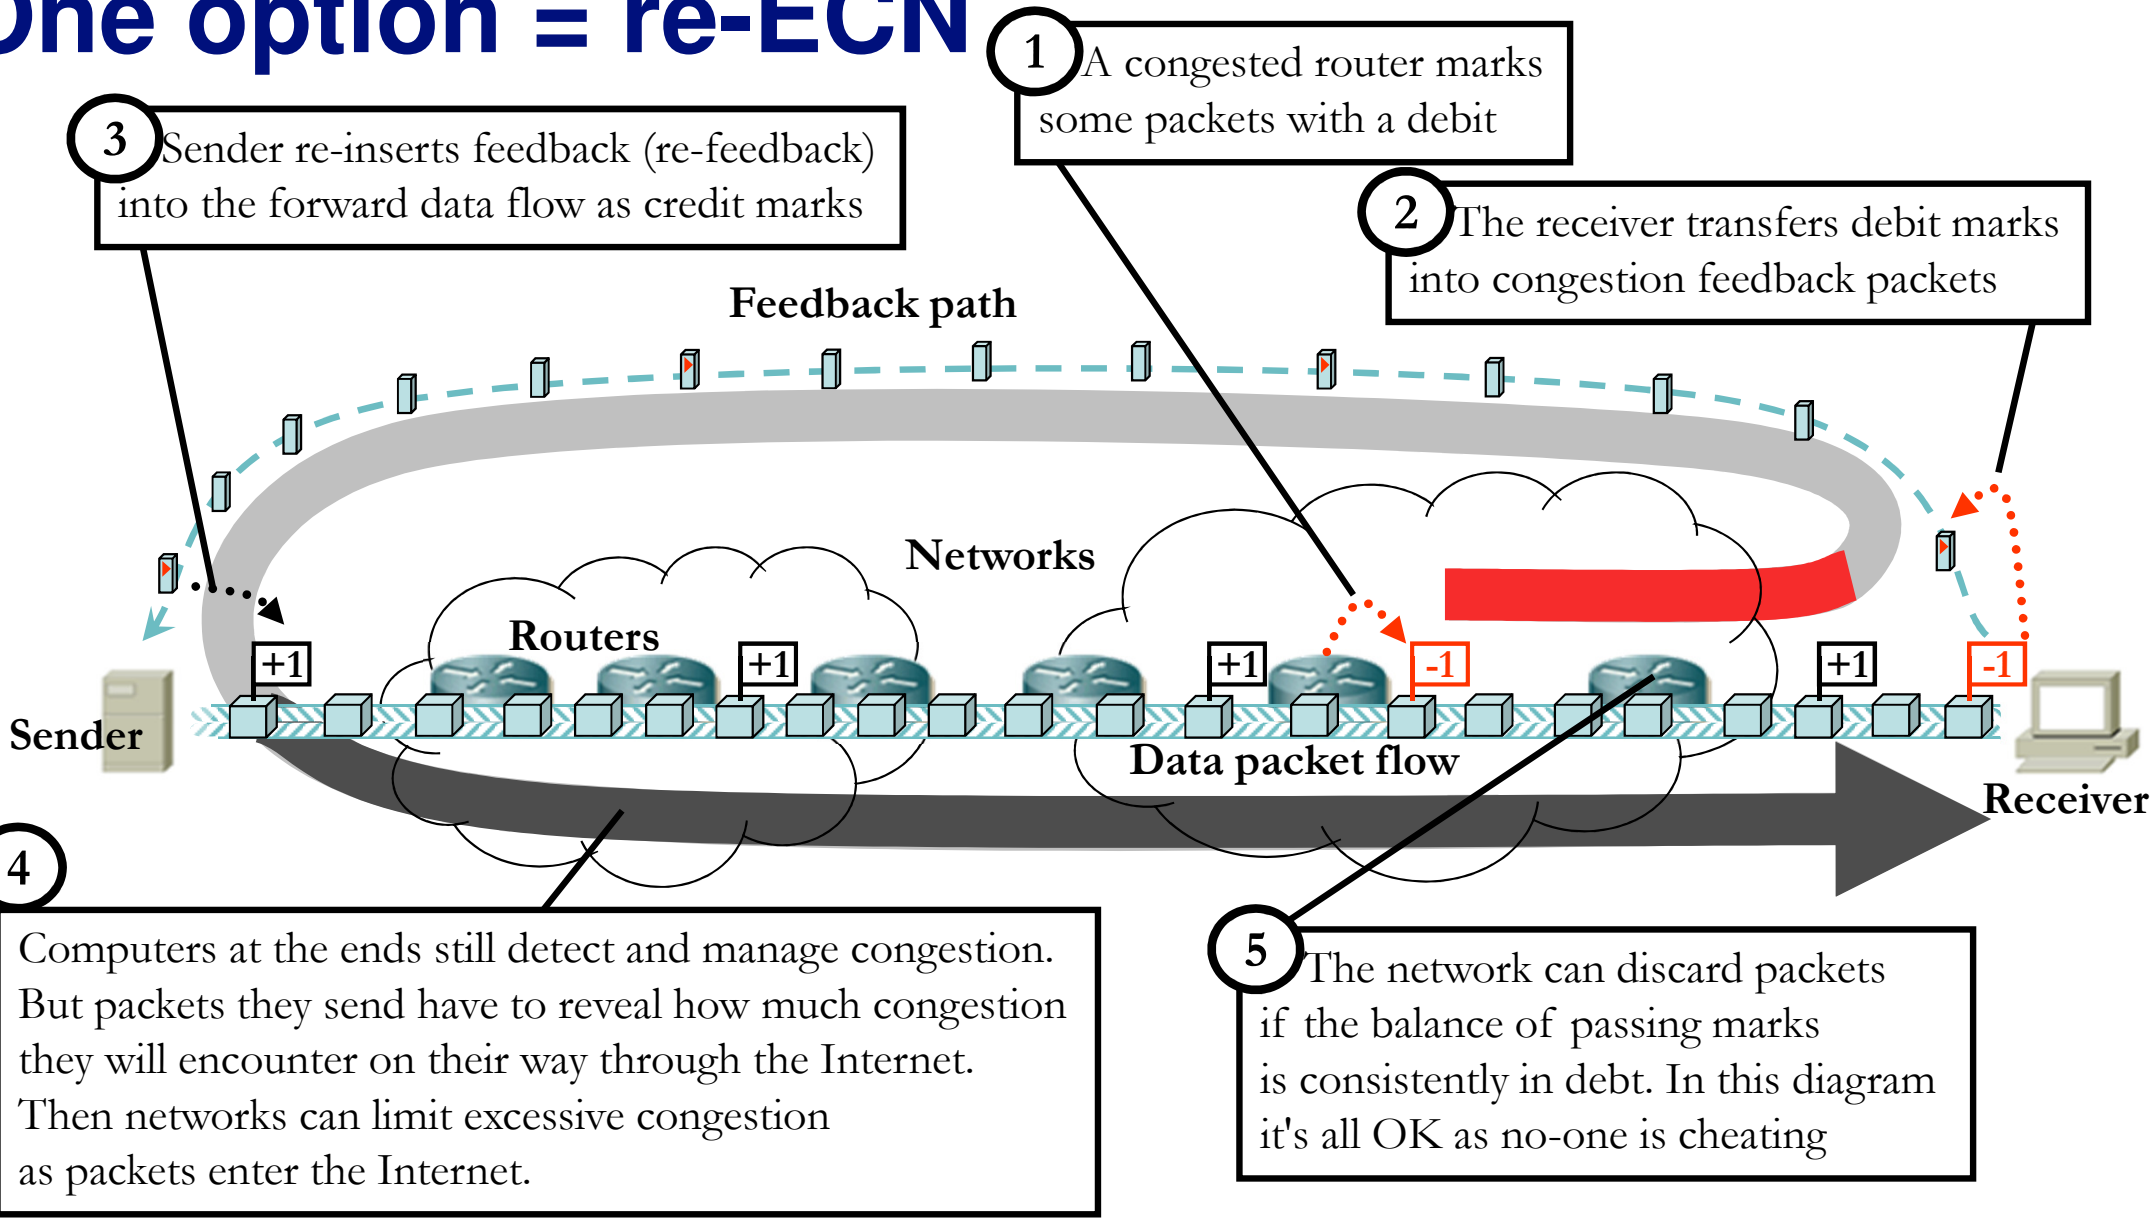

- Policing upload traffic (rather than download) requires end-of-path information validation but provides stronger protection against identity spoofing

Conclusions

- Make each packet accountable
- Control congestion volume rather than volume
- Don’t assume link between application type and congestion
- Enforce per customer, at contractual connectivity point
- Expose downstream congestion
- Bulk constraint forces the evolution of end-customer rate adaptation and encourages better use of shared resource pool

Re'Arch 2008

Policing Freedom…

to use the Internet Resource Pool

Arnaud.Jacquet, Bob.Briscoe, Toby.Moncaster {@bt.com}

December 9 2008

Direct effect on a single flow

(illustration purposes)

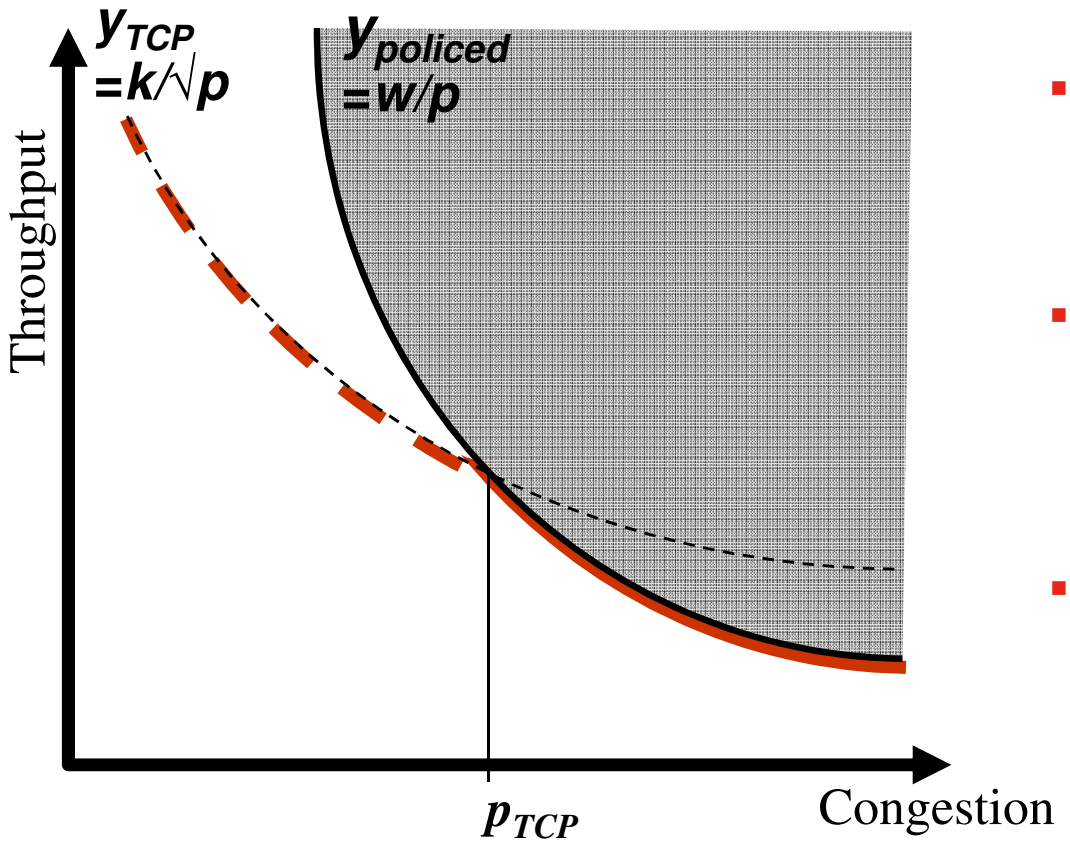

- Each flow has its natural congestion response, based on the application used

  ➔ eg. y_{TCP}
- The policer puts a constraint forcing the operational point of the application's throughput to remain out of the shaded area
- When congestion exceeds p_* , the policer takes over the congestion response

Direct effect on a single flow

(illustration purposes)

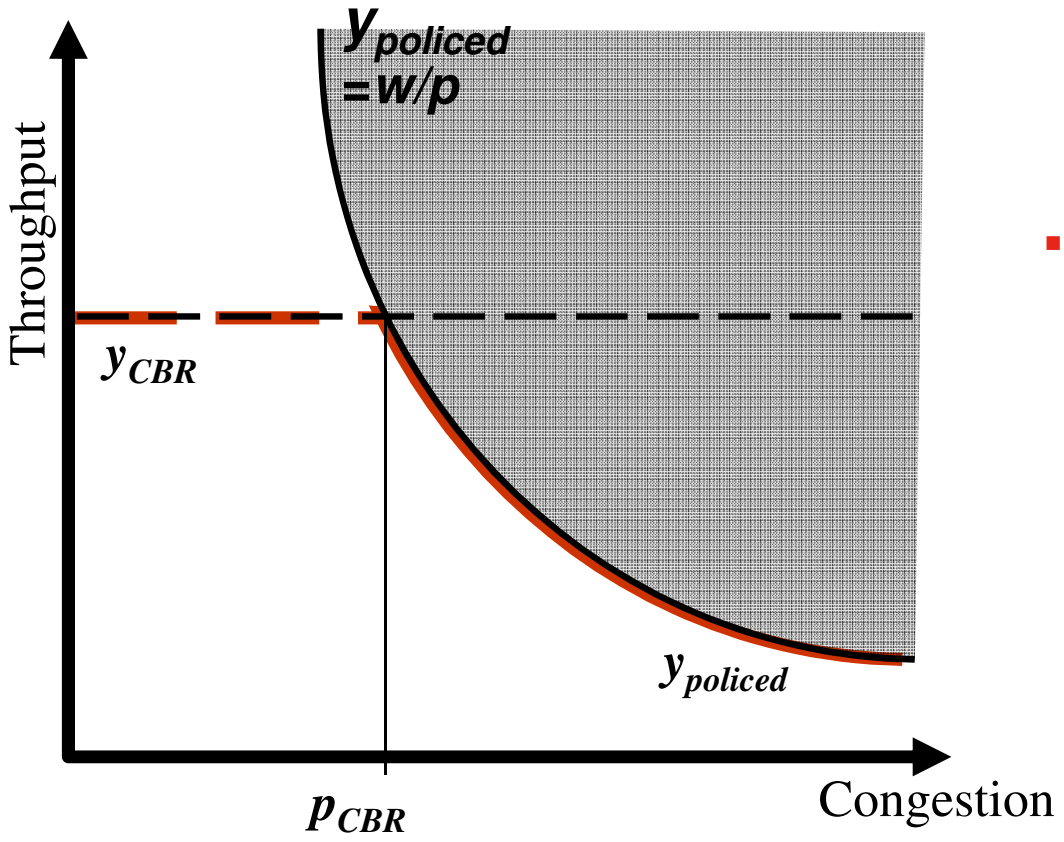

- This also applies for unresponive flows