Congestion balancing using re-ECN

João Taveira Araújo Miguel Rio George Pavlou

Networks and Services Research Laboratory
University College London

MSN 2009

UCL

Outline

1. **Introduction**
   - Congestion Pricing
   - Congestion Exposure
2. **Traffic Engineering**
3. **Multipath routing**

Motivation for congestion pricing

- Internet made to be shared, but expects users to be polite
  - TCP "fairness" prevails, first use of sarcasm in IETF.
  - heavier users cause most congestion. . .
  - . . . but operators must pass costs onto all.
- Operators reciprocate with available means:
  - Throw bandwidth at the problem.
  - Cap volume allowance.
  - Throttle bandwidth.
  - Deep packet inspection.

Motivation for congestion pricing

- Internet made to be shared, but expects users to be polite
  - TCP "fairness" prevails, first use of sarcasm in IETF.
  - heavier users cause most congestion. . .
  - . . . but operators must pass costs onto all.
- Operators reciprocate with available means:
  - Throw bandwidth at the problem.
  - Cap volume allowance.
  - Throttle bandwidth.
  - Deep packet inspection.

Congestion as a cost

- "Taxing" wrong metrics leads to (un)predictable results

Congestion as a cost

- "Taxing" wrong metrics leads to (un)predictable results
- Congestion is the marginal cost for provisioning network
  - no extra cost for carrying traffic until congestion occurs
  - congestion is a cost inflicted on all - externality

"Social welfare is maximized if sources are charged a shadow price proportionate to congestion they cause" **Frank Kelly**

- "You can't do this scalably/efficiently/incrementally"
- By exposing congestion at network layer, maybe you can

Congestion as a cost

- "Taxing" wrong metrics leads to (un)predictable results
- Congestion is the marginal cost for provisioning network
  - no extra cost for carrying traffic until congestion occurs
  - congestion is a cost inflicted on all - externality

"Social welfare is maximized if sources are charged a shadow price proportionate to congestion they cause" **Frank Kelly**

- "You can't do this scalably/efficiently/incrementally"
- By exposing congestion at network layer, maybe you can

Congestion as a cost

- "Taxing" wrong metrics leads to (un)predictable results
- Congestion is the marginal cost for provisioning network
  - no extra cost for carrying traffic until congestion occurs
  - congestion is a cost inflicted on all - externality

"Social welfare is maximized if sources are charged a shadow price proportionate to congestion they cause" **Frank Kelly**

- "You can't do this scalably/efficiently/incrementally"
- By exposing congestion at network layer, maybe you can

Congestion as a cost

- "Taxing" wrong metrics leads to (un)predictable results
- Congestion is the marginal cost for provisioning network
  - no extra cost for carrying traffic until congestion occurs
  - congestion is a cost inflicted on all - externality

"Social welfare is maximized if sources are charged a shadow price proportionate to congestion they cause" **Frank Kelly**

- "You can’t do this scalably/efficiently/incrementally"
- By exposing congestion at network layer, maybe you can

Why should congestion be visible at network layer?

- Congestion is viewed purely end-to-end:
  - network sends explicit/implicit signals
  - end host adjusts sending rate
- Ideally, network could view congestion too,
  - enforce congestion pricing
  - impose fairness between sources
  - hold neighbouring domains accountable for costs

Why should congestion be visible at network layer?

- Congestion is viewed purely end-to-end:
  - network sends explicit/implicit signals
  - end host adjusts sending rate
- Ideally, network could view congestion too,
  - enforce congestion pricing
  - impose fairness between sources
  - hold neighbouring domains accountable for costs

Re-ECN: an elevator pitch

- Re-ECN inserts congestion info in IP header
  - protocol in draft status in IETF
  - incrementally deployable (only 1 bit needed)
- Empowers networks by revealing congestion
  - each packet contains congestion seen 1 RTT ago
  - routers can estimate upstream congestion
- No obligation on how protocol should be used
  - "implement policy at runtime, not design time"
  - philosophical shift from TCP

Re-ECN: an elevator pitch

- Re-ECN inserts congestion info in IP header
  - protocol in draft status in IETF
  - incrementally deployable (only 1 bit needed)
- Empowers networks by revealing congestion
  - each packet contains congestion seen 1 RTT ago
  - routers can estimate upstream congestion
- No obligation on how protocol should be used
  - "implement policy at runtime, not design time"
  - philosophical shift from TCP

Re-ECN: an elevator pitch

- Re-ECN inserts congestion info in IP header
  - protocol in draft status in IETF
  - incrementally deployable (only 1 bit needed)
- Empowers networks by revealing congestion
  - each packet contains congestion seen 1 RTT ago
  - routers can estimate upstream congestion
- No obligation on how protocol should be used
  - "implement policy at runtime, not design time"
  - philosophical shift from TCP

Summary so far

- Congestion pricing aligns incentives
- Exposing congestion proposed as solution

If we can view congestion at the network layer, what else can we improve?

Using re-ECN to improve interdomain routing?

- Interdomain routing driven by commercial interests
- But we still evaluate paths based on "performance"
  - AS hops
  - load factor
- Not optimal, requires periodically computing costs

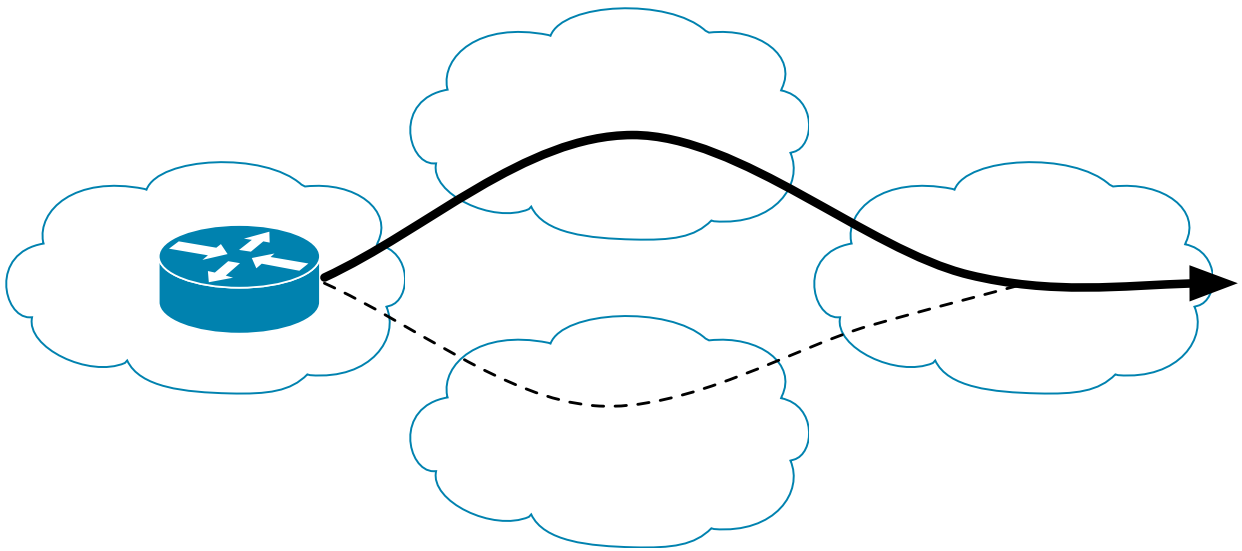

Congestion balancing

- Use expected congestion to select egress
- balance congestion, not load (less myopic)
- stub networks can split prefixes

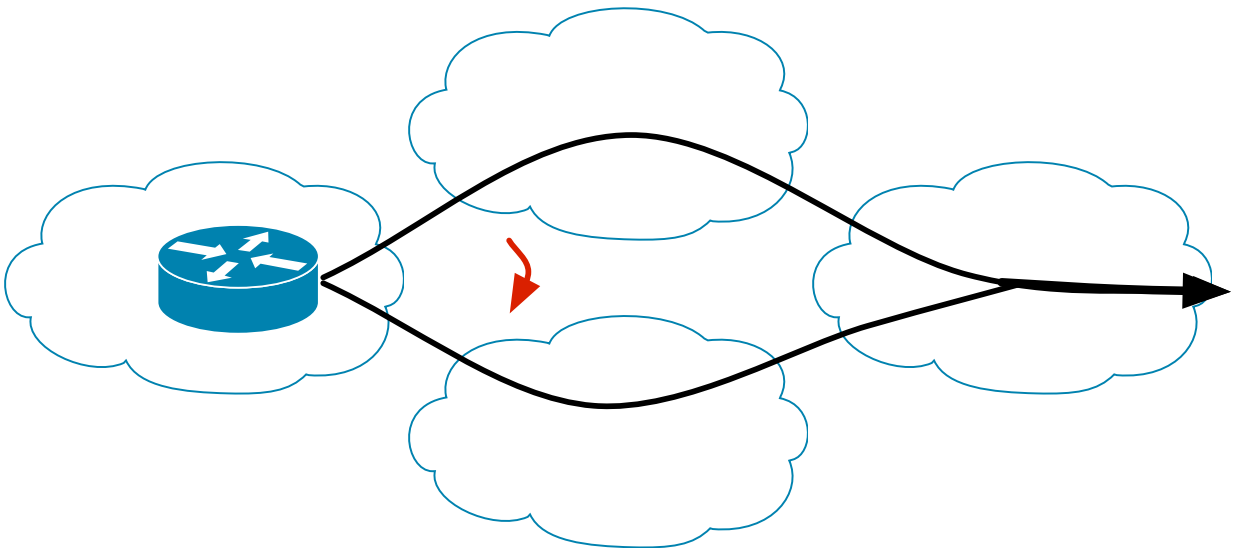

Open issues

- Significant improvement?
- Oscillations?
- Is re-ECN fine-grained enough?

Extending TE approach across domains

- Multipath routing has recognized advantages
  - reliability
  - higher throughput
  - optimize traffic according to requirements
- No incentive for domains to provide multiple paths
  - shifting bits a commodity, provide as cheaply as possible
  - no incentive for transit domains to carry more traffic
- Congestion pricing aligns incentives
- Congestion exposure provides path evaluation

Extending TE approach across domains

- Multipath routing has recognized advantages
  - reliability
  - higher throughput
  - optimize traffic according to requirements
- No incentive for domains to provide multiple paths
  - shifting bits a commodity, provide as cheaply as possible
  - no incentive for transit domains to carry more traffic
- Congestion pricing aligns incentives
- Congestion exposure provides path evaluation

Extending TE approach across domains

- Multipath routing has recognized advantages
  - reliability
  - higher throughput
  - optimize traffic according to requirements
- No incentive for domains to provide multiple paths
  - shifting bits a commodity, provide as cheaply as possible
  - no incentive for transit domains to carry more traffic
- Congestion pricing aligns incentives
- Congestion exposure provides path evaluation

Extending TE approach across domains

- Multipath routing has recognized advantages
  - reliability
  - higher throughput
  - optimize traffic according to requirements
- No incentive for domains to provide multiple paths
  - shifting bits a commodity, provide as cheaply as possible
  - no incentive for transit domains to carry more traffic
- Congestion pricing aligns incentives
- Congestion exposure provides path evaluation

Propogating congestion information

- Domains exchange per-prefix congestion information
  - still apply policy
  - choose small subset of best-performing paths
  - offer paths to end hosts

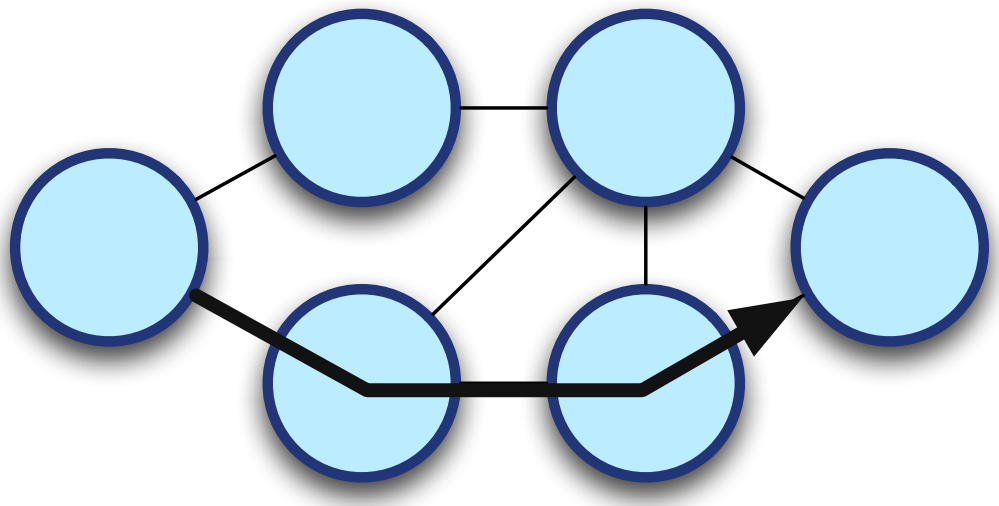

Propogating congestion information

- Domains exchange per-prefix congestion information
  - still apply policy
  - choose small subset of best-performing paths
  - offer paths to end hosts

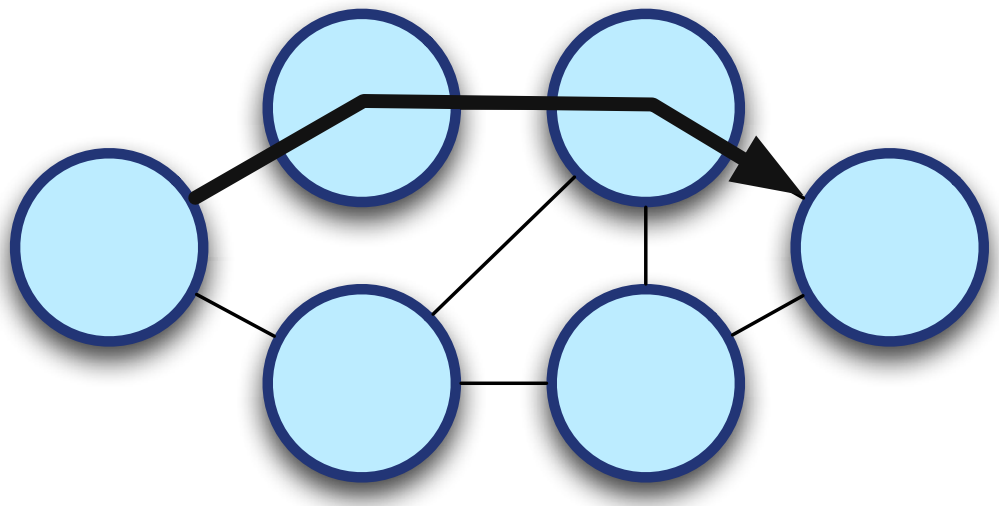

Propogating congestion information

- Domains exchange per-prefix congestion information
  - still apply policy
  - choose small subset of best-performing paths
  - offer paths to end hosts

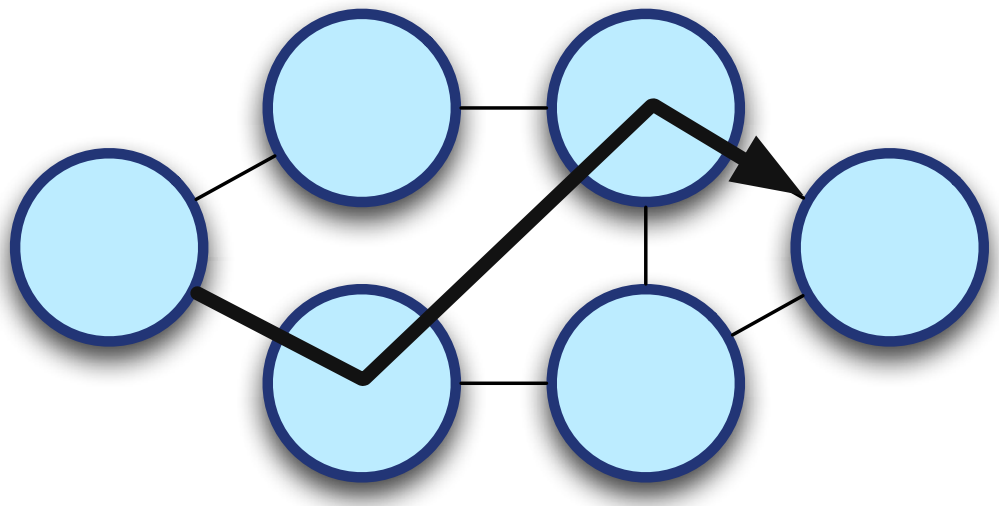

Open issues

- Scalability
- Loop-freeness
- Interfacing with end-host?
- Incentives in the absence of pricing

Conclusions

- Congestion exposure one solution to pricing problem
- We can use this to improve interdomain routing
- Will this require pricing, or is it complementary?
- Is re-ECN exposing enough for this purpose?

- Questions
- Comments
- Abuse