# Some questions about multipath

Damon Wischik, UCL

Trilogy UCL

# 1. Cut constraints and routing

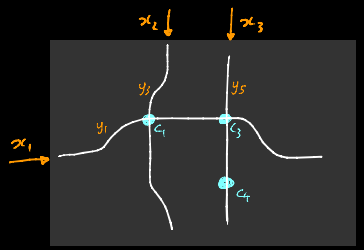


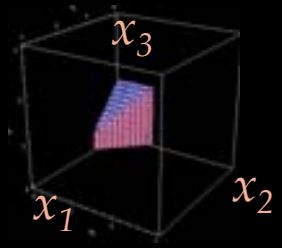


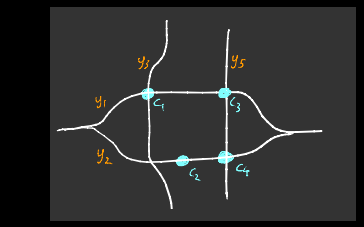


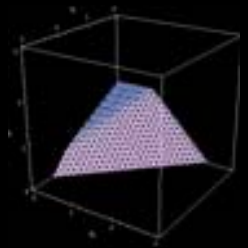

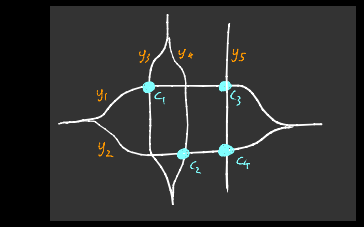


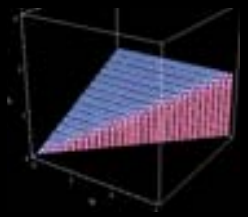

- QUESTION

  If the routing system can provide multiple paths to each flow, what should they be? Or, if users are multihomed, what choice of ISPs is good?

- THEORY

  Given a set of multipath allocations, we can work out the capacity region, i.e. the set of acceptable traffic matrices. The more path choice, the bigger the capacity region.

# 1. Cut constraints and routing

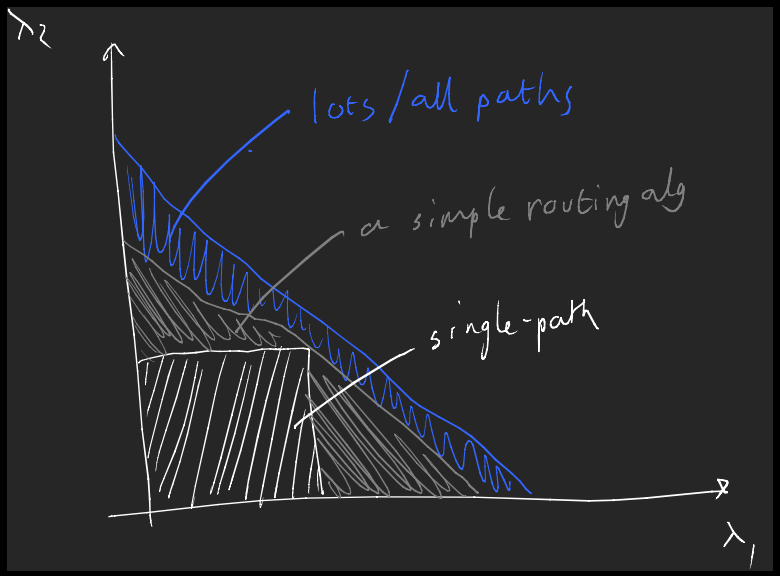


- METRICS

  Devise a multipath routing algorithm. Measure its effectiveness by
  - the volume of the capacity region achieved, as a fraction of the best possible capacity region (when flows have access to all possible paths)
  - the maximum sum of user utilities, according to the TCP utility function, assuming some given demands/weights

- ISSUES
  - Compact routing. Is there some sort of heterogeneity score for paths which would mean we don't have to propagate the entire path?
  - Should we always try to find shortest paths, or do we get more path diversity by choosing paths which aren't quite shortest?

# 2. Cut constraints and topology planning

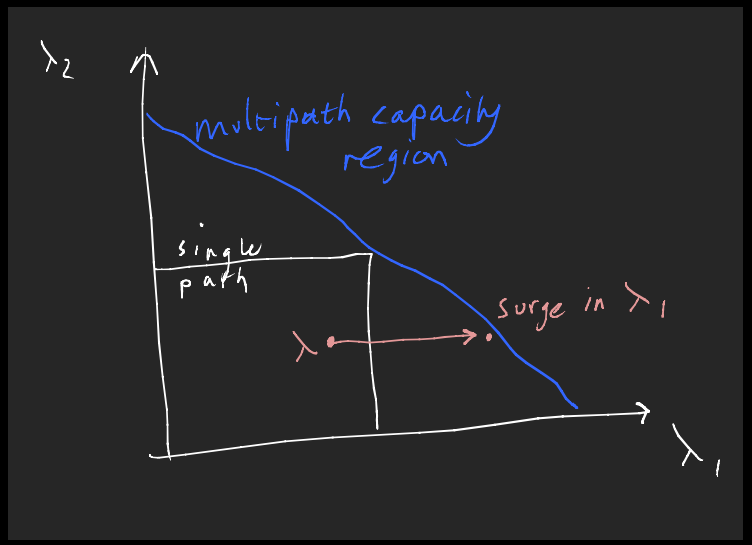


- QUESTION

  Which parts of the network are the most serious bottleneck, even with best possible multipath routing? Where should new links be added?

- MOTIVE

  Multipath means that the network copes automatically with surges in traffic and link failures. How much? How does it compare to single-path routing?

# 2. Cut constraints and topology planning

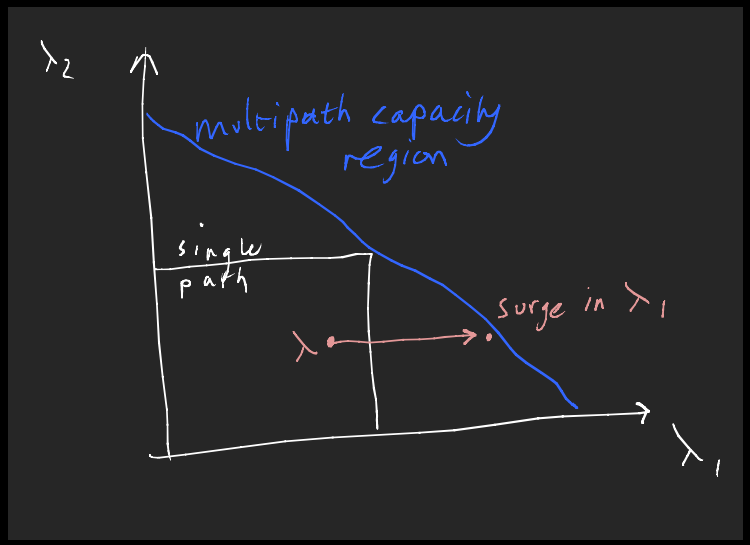


- METRIC
  Assume a traffic matrix $\lambda$, and measure the *surge factor*
  - If $\lambda$ for some source-destination pair is multiplied by $\alpha$, suppose we need to multiply the rest of the traffic matrix by $1-\beta$ to keep within the capacity region
  - Define the surge factor $s(\alpha)$ to be the worst-case $\beta$ across all source-destination pairs

- ISSUES
  The generalized cut constraints may span large parts of the network, and traffic routes itself around problems, so no single network provider knows what's going on

# 3. Power of two choices

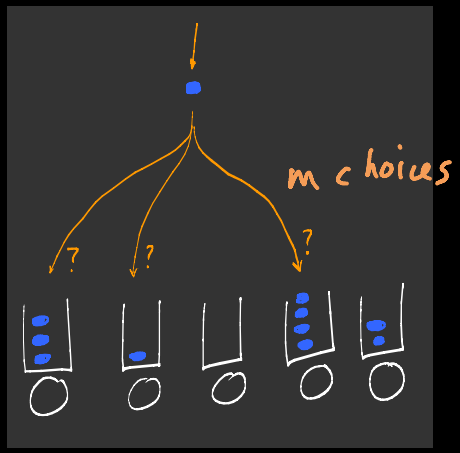


$m$=1 $m$=2 $m$=5

- QUESTION

  Does the network topology allow us to get the benefit of `power of two choices'?

- THEORY

  Mitzenmacher (1996). If each job can choose the least loaded of $m$ paths, then $m$=2 is sufficient

  Key+Massoulié (2007). If each flow is split over $m$ paths chosen at random, we get the same benefit. If it does load balancing across these paths we get even more benefit

  Godfrey (2008). If these choices are not completely random, these benefits may not obtain: we may need $m$=log($n$) where $n$ is the number of physical links

# 3. Power of two choices

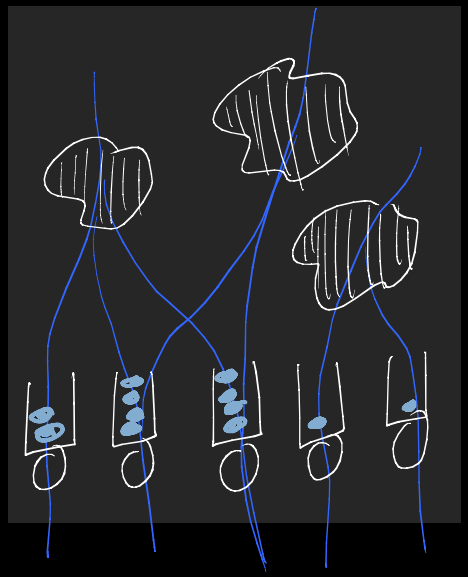

- ISSUE

  If your set of choices across the cut constraint are limited by the upstream topology, then Godfrey's result may limit the amount of resource pooling that happens

- METHOD
  - Combine the results of Key+Massoulié and Godfrey
  - Experiment with e.g. power-law distributions for the upstream topology

# 4. Behaviour of congestion control

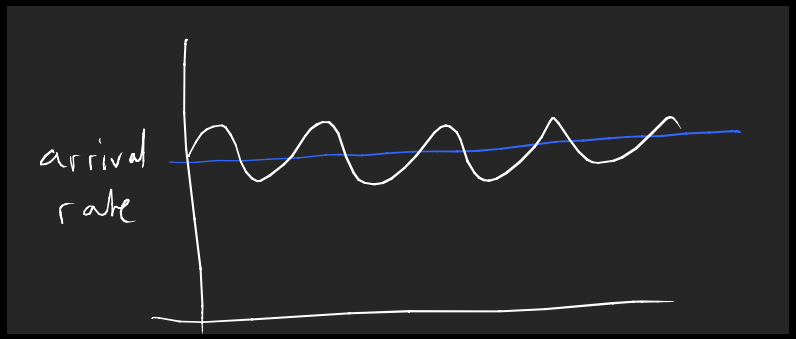


- QUESTION

  How well do the proposed algorithms for multipath congestion control actually work?

- THEORY

  TCP-like algorithms have been proposed by Han, Shakkotai, Hollot, Srikant & Towsley (2006), and by Kelly & Voice (2005). Their fluid-model stability has been analyzed.

# 4. Behaviour of congestion control

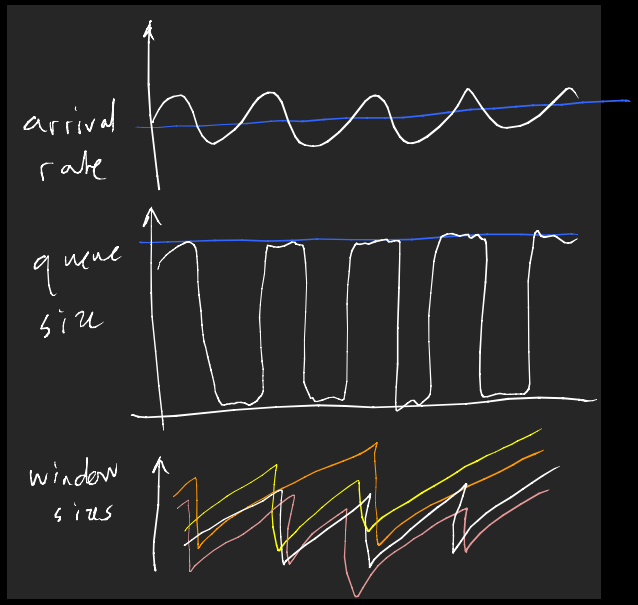


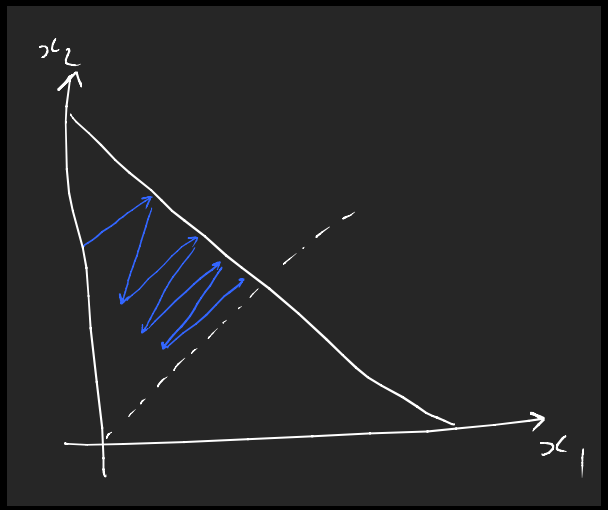


- EXPERIMENTAL ISSUES
  - In the single path case, instability corresponds to TCP synchronization and bursty losses. In multipath, does it correspond to synchronization, or to route oscillation?
  - In the single path case, TCP is fairly robust: oscillations aren't as bad as theory predicts. Might multipath have further issues?
  - How robust is stability to noise, heterogeneity etc.?
- THEORY ISSUES
  - How well do the proposed schemes work for small numbers of flows?
  - What are the appropriate modifications to other high-speed TCP algorithms?

# 5. Economics of multipath

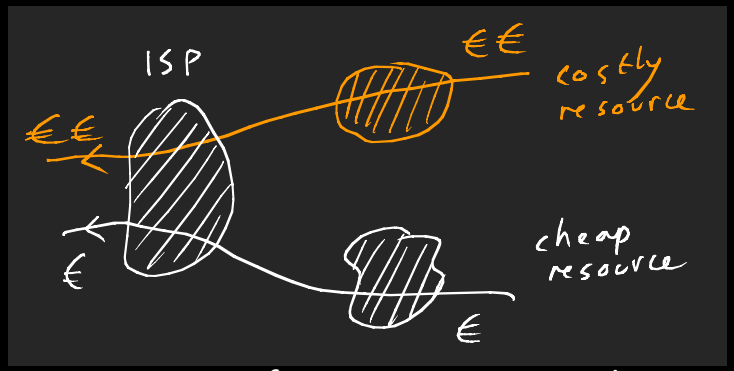


*Competition is forcing ISPs into tiered pricing: charge more to those customers who use more expensive resources.*

- **CLAIM**
  Multipath routing will increase competition. This will force ISPs into e.g. tiered pricing based on congestion costs, and it will force downstream networks into congestion pricing.

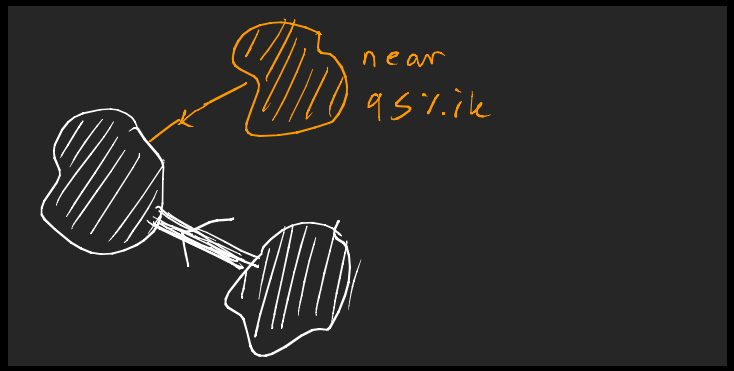


*A network provider can throttle its connection, when it is near its charging limit, so that load-balancing multipath users will back off. Thus prices should reflect congestion.*

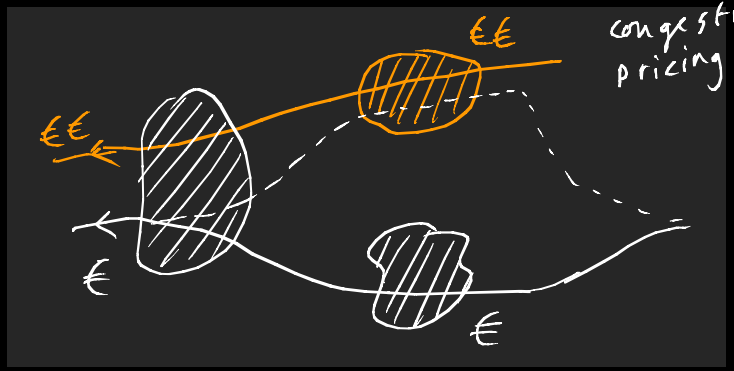


*Suppose the remote resources use congestion pricing. If multipath users adapt quickly to congestion, they end up costing less, so a competitive ISP will be forced to charge them less.*

# 5. Economics of multipath

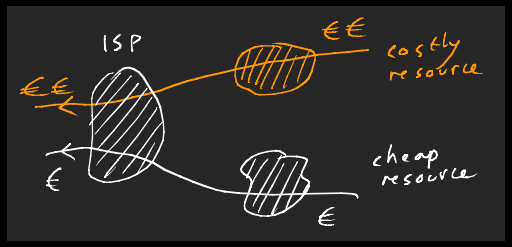


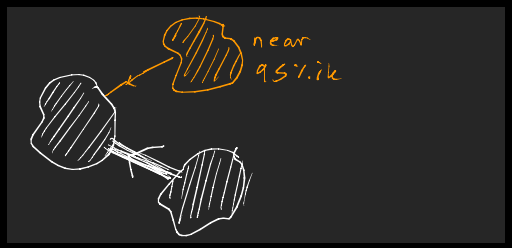


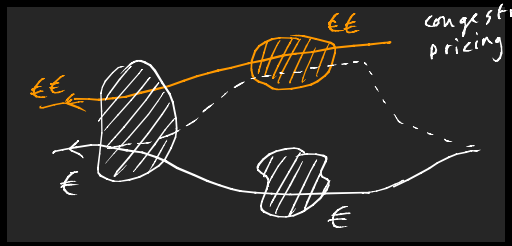


- QUESTION
  - Is this hand-waving plausible? Is it true? Can it be quantified?
  - Can we suggest a clean and logical migration path?
  - Multipath is a way to let users reveal their preferences (e.g. user chooses a path with his/her desired RTT, drop rate). This should increase the net economic efficiency

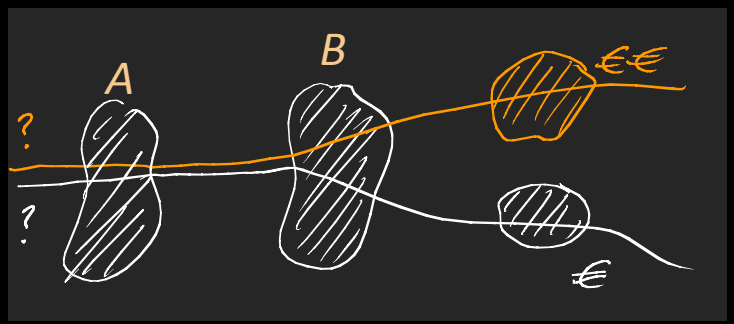


*Is there a cunning way for ISP A to work out how to do price discrimination?*

- *Maybe network provider B could provide two queues, one €€, one €, each able to reach the two resources but at different line rates.*
- *ISP A would then know the cost incurred by its different types of user*

# 6. Multipath and P2P


- P2P does multipath already, in some sense
- It also chooses peers
  - i.e. more degrees of freedom than single-flow multipath

- How should peers be chosen, compatibly with multipath rate allocation?