Fair Background Data Transfers of Minimal Delay Impact

C. Courcoubetis

A. Dimakis

AUEB

Workshop on Capacity Sharing, Oct 13, 2011, Stuttgart

- Propose a framework for the design of protocols for background transfers
- File sizes differ by >10 orders of magnitude, connection rates by few orders of magnitude
- Main concern: how to obtain reasonable throughput with minimal delays on short flows
- Current approaches: TCP-nice, LEDBAT,...
  - behave as second priority traffic (low impact on short flows)
  - no consideration of fairness relative to other long flows
  - no adoption incentives
- Related work: Key, Massoulie, etc.
  - substitution of all long TCPs by on-off senders based on threshold price
  - prove that there is *some* delay improvement
  - mostly a different traffic model, assumes all flows to convert to new protocol

The competition environment

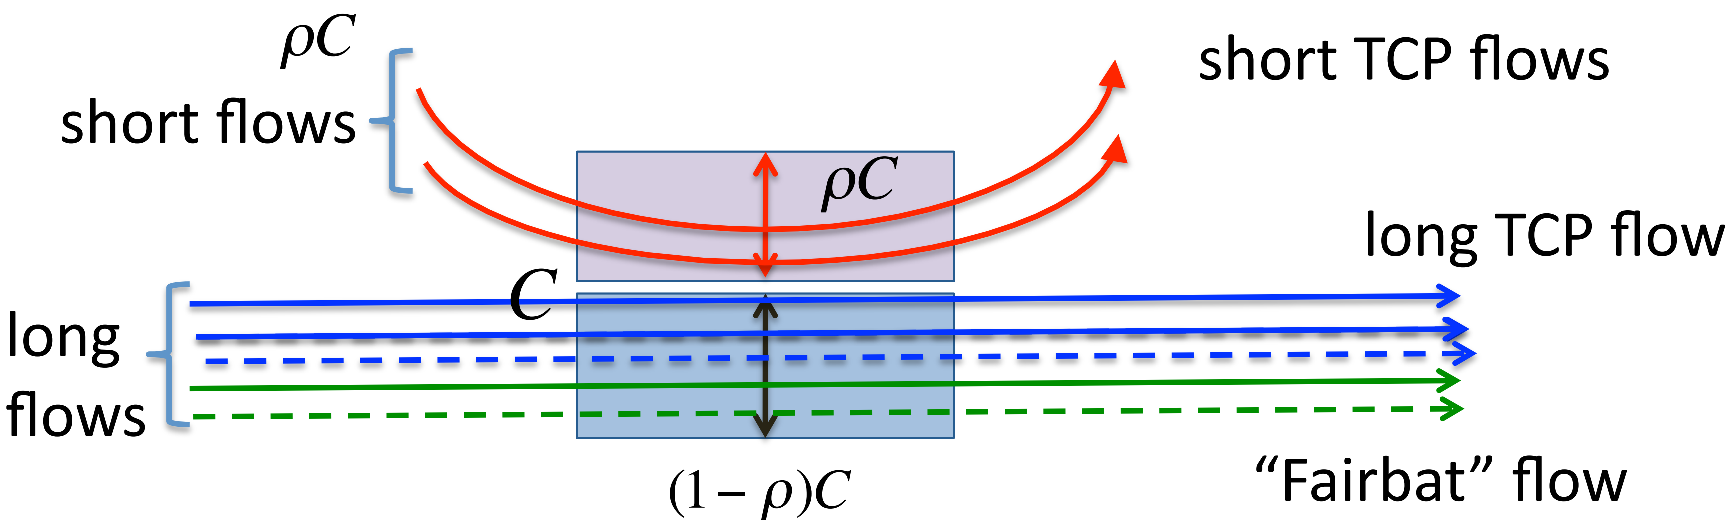

- No competing long TCP: easy case!! FB: 2^{nd} priority
- 1 long TCP: FB 2^{nd} priority => zero throughput
- existing solutions: > 2^{nd} priority, unspecified throughput

Adoption incentives vs "niceness"

- Why "long TCP" users adopt FB instead of TCP?
- Which are sensible properties of FBs?
  - When competing with long TCPs for $C(1-\rho)$:
    - obtain a given fraction f of $C(1-\rho)$,
    - cause minimum extra delays on short flows
  - Example: obtain same average throughput as TCP $f = \frac{l}{k+l}$
- Achieve all that with reasonable context information
  - public Internet context, competition with non-local flows

Our results

- Obtain optimal BW sharing policy under complete information
  - minimize delays on short flows while competing with *k* long TCPs and obtaining a share *f* of the leftover capacity
- Implementable approximation: weighted TCP
  - short time scales: use *w*-TCP $f = \frac{w}{k+w}$
  - delay deterioration $\leq 17.2\%$ for $k = 1$,

    $\downarrow 0$ as $k \to \infty$
- Delay impact $\delta(f)$

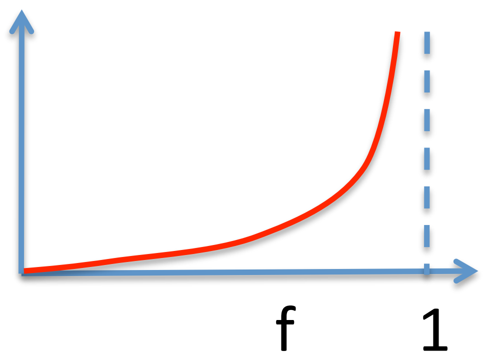

Our results (cont.)

- Adopt Kelly's approach for fairness, but for long-term
  - FBs don't get a fixed fraction *f* of leftover capacity
  - max some social welfare function
    - sum of utilities for average throughput of long flows TCPs, FBs
    - add a negative externality term (extra delay to short flows)
  - implementation: using w-TCPs
    - don't need to know $k, C, \rho, \ldots$
- Use this framework to design new protocols
  - examples: $y_0 / y_i = 1 + \gamma\, y_0^{-1}$

$$y_0 / y_i = 1 + \gamma\, y_0$$

A corollary

- If we substitute any subset of long TCP flows by "equivalent" optimal FBs, the max improvement of the delay of short flows is less than 17.2%
- The best improvement is achieved when there is competition of 1FB and 1 long TCP flow
- A negative result?
- The incentive compatibility constraint (obtain same average throughput as TCP) in larger systems implies small optimal delay improvements
- To get significant delay improvement we need to relax the IC condition (how?)

The general fairness framework

- Problem: “fair” share of excess capacity
- Express fairness on long-term rates “à la Kelly”
- Take into account delay spillovers to short flows
  - remember the tradeoff $f \leftrightarrow \text{delay}$
- Engineering: translate into flow control algorithms
  - decompose controls for short and long timescales
  - make reasonable assumptions on what is known locally
- Reverse engineering: translate existing algorithms into this model

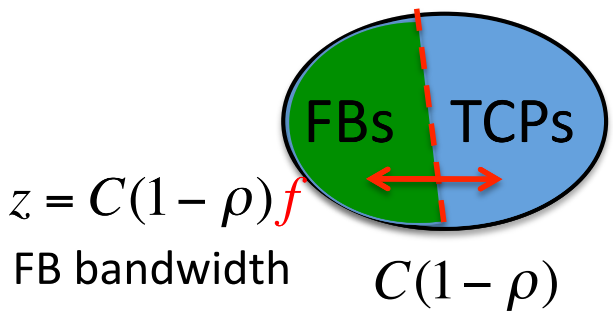

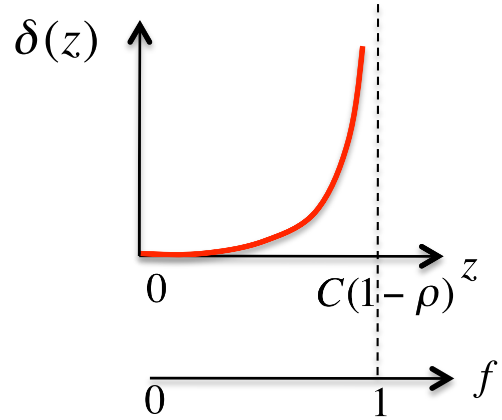

The optimization problem

$$\max \quad k u_0(y_0) + \sum_{i=1}^{l} u_i(y_i) - \int_0^{\sum_{i=1}^{l} y_i} F(\delta(z))\,dz$$

$$\text{such that } k y_0 + \sum_{i=1}^{l} y_i = C(1-\rho),$$

$$\text{over } y_0, \ldots, y_l \geq 0$$

Assume that w-TCP is used in the short t.s. Then magic!!!

$$\delta_{\mathrm{w}}\left(\sum_{i=1}^{l} y_i\right) == \frac{1}{y_0}$$

Optimality condition for long ts: $-u_0'(y_0) + u_i'(y_i) - \frac{1}{y_0^2} F\left(\frac{1}{y_0}\right) = 0, i = 1, \ldots, l$

Short ts: use w-TCP

Long ts: adapt the weights $\dot{w}_i = -u_0'\left(\frac{y_i}{w_i}\right) + u_i'(y_i) - \left(\frac{w_i}{y_i}\right)^2 F\left(\frac{w_i}{y_i}\right), i = 1, \ldots, l$

Engineering new protocols

$$\max \quad k u_0(y_0) + \sum_{i=1}^{l} u_i(y_i) - \int_0^{\sum_{i=1}^{l} y_i} \delta_w(z)^2 F(\delta(z))dz \qquad -u_0'(y_0) + u_i'(y_i) - \frac{1}{y_0^2} F\left(\frac{1}{y_0}\right) = 0, i = 1,\ldots,l$$

$$\text{such that } k y_0 + \sum_{i=1}^{l} y_i = C(1-\rho), \qquad \Leftrightarrow \dot{w}_i = -u_0'\left(\frac{y_i}{w_i}\right) + u_i'(y_i) - \left(\frac{w_i}{y_i}\right)^2 F\left(\frac{w_i}{y_i}\right), i = 1,\ldots,l$$

over $y_0,\ldots,y_l \geq 0$, w-TCP short ts controls

Case A:

$$u_i(y) = \log y, i = 0,\ldots,l, F(\delta) = \gamma$$

$$\frac{y_0}{y_i} = 1 + \gamma y_0^{-1}, i = 1,\ldots,l$$

Case B:

$$u_i(y) = \log y, i = 0,\ldots,l, F(\delta) = \gamma\delta^{-2}$$

$$\frac{y_0}{y_i} = 1 + \gamma y_0, i = 1,\ldots,l$$

Algorithm A

$$\frac{y_0}{y_i} = 1 + \gamma\, y_0^{-1}, i = 1,\ldots,l$$

- FBs get similar throughput as TCP when there is enough excess bandwidth, give away when it becomes scarce
- IC condition relaxed when resources are scarce (second priority when sensible)

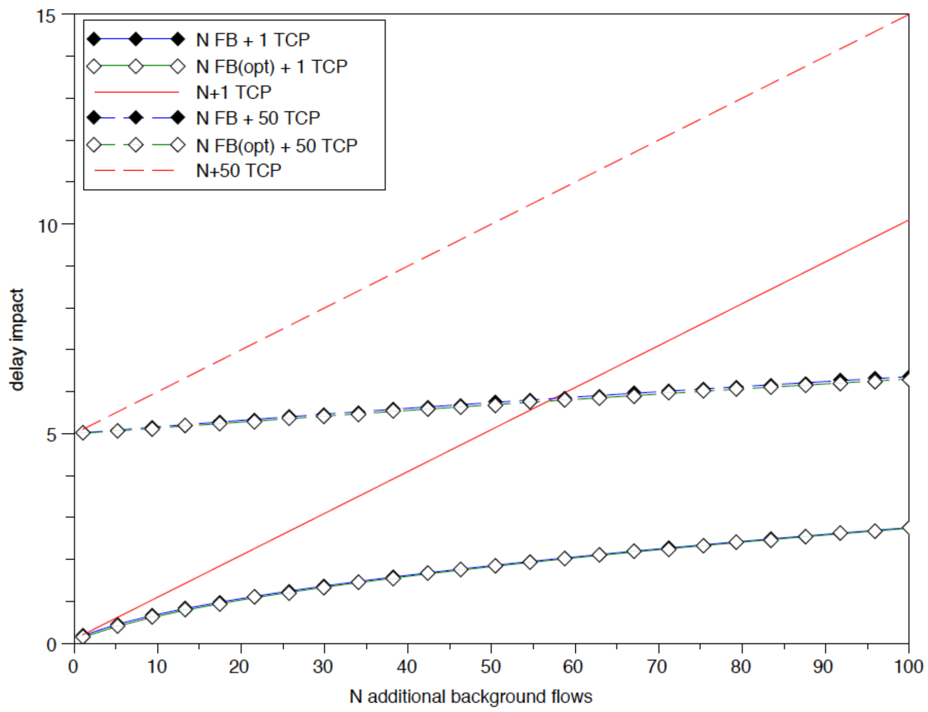

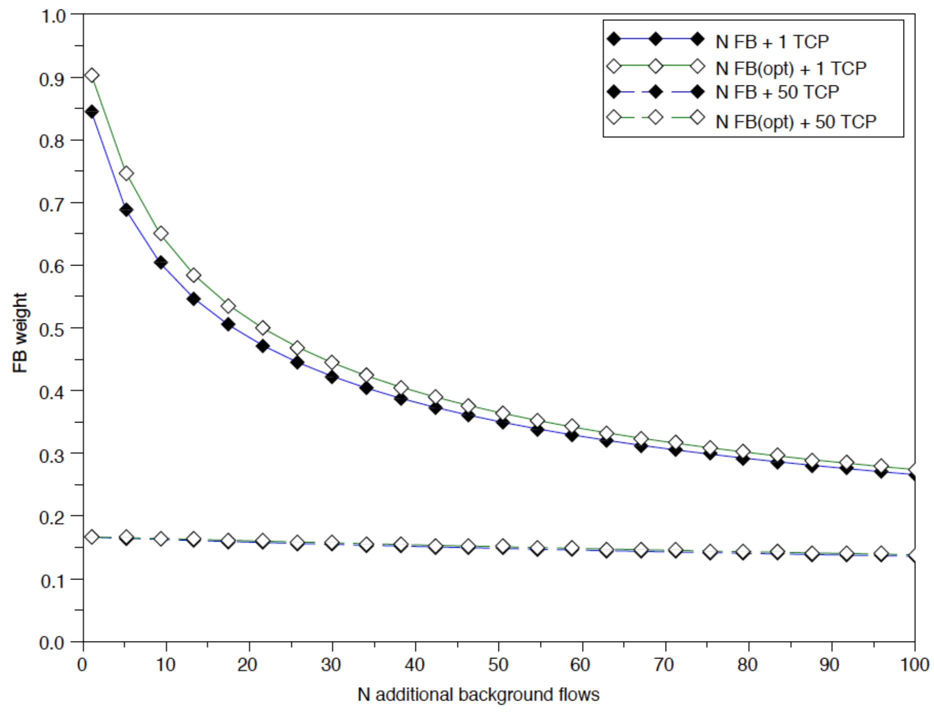

Conclusions

- Protocols for background transfers operate in the context of other long and short TCP flows
- TCP is the incumbent protocol, new protocols should compare to TCP
- We derived the optimal short time scale policy for achieving a given share of long term throughput, but has practical implementation issues
- w-TCP seems a reasonable practical alternative, provably small efficiency loss
- We provided a utility-based definition for fair sharing including a negative externality term for delay caused to short flows
- We derived two new interesting protocols for background transfers by relaxing the IC condition for adoption relative to TCP