Fair Background Data Transfers of Minimal Delay Impact

Costas Courcoubetis
Antonis Dimakis
Department of Computer Science
Athens University of Economics and Business
{courcou,dimakis}@aueb.gr

30. July 2011

Abstract

In this work we present a methodology for the design of congestion control protocols for background data transfers, that addresses the trade-off between the fair sharing of excess capacity and the delay externalities on short flows, when background transfers compete with other coexisting background flows that use TCP. First, we analytically solve the basic problem of finding the bandwidth sharing policy with the minimal delay impact on short TCP flows, if a target share of the excess capacity is to be achieved, the rest of the excess capacity being consumed by competing background TCP transfers. As an alternative to the optimal policy, which turns to be not practical to implement, we consider a weighted TCP policy, which maintains a target proportion of TCPs bandwidth at all times in order to achieve the same share of excess capacity. We establish that this much simpler policy achieves a delay always within 17% of the optimal. Moreover, the relative optimality error rapidly decreases to zero as the number of coexisting background TCP flows increases.

Next, we consider a general utility-based fairness criterion with a delay penalty term for the delay caused to short flows, jointly optimized over all allocations of excess capacity to background flows (including TCP ones) on long timescales and all bandwidth sharing policies on short timescales that influence the delays of short flows. Even though the delay optimal sharing policy that solves the above optimization problem does not lead to distributed congestion control algorithms and more significantly, requires the number of competing background TCP flows, both problems are solved under the weighted TCP policy. A distributed weight adjustment policy is considered where, at equilibrium, the overall performance is nearly optimal, with a quickly vanishing relative optimality error, as the number of background TCP flows increases.

We illustrate the methodology by giving two examples of con- gestion control algorithms for background transfers. Both achieve low delay for short flows relative to TCP, but at the same time they present strong incentives for adoption against incumbent low priority solutions in public environments.