CCIT Report No. 311

Optimal Transport Strategies for Best-Effort Traffic over Priced Connections

Lavy Libman and Ariel Orda

ABSTRACT

In recent years, economic research has led to the development of several usage-based pricing schemes, designed to replace the flat-rate pricing common in today's communication networks and overcome its deficiencies. The majority of works, however, assume that the network infrastructure and communication protocols are not affected by the pricing scheme, and leave the pricing-related decisions to be made by (human) users or, at most, by top-level applications. We contend that lower-layer networking protocols, i.e. in the contexts of routing and end-to-end transport control, must also take pricing into account and attempt to be not only technically efficient, but also *cheap*. In this paper we provide an in-depth analysis of the typical transport-layer problem of choosing the optimal timeout to wait for an acknowledgment after sending a single packet over an unreliable connection that charges a linear tariff of a per unit time plus b per packet retransmission. Specifically, we study the dependence of the timeout on the time-volume price ratio, propagation and queuing delay, and loss probability (the latter two indirectly affected by network load). We show that there is a finite threshold of $\frac{a}{b}$ above which the first timeout is zero, i.e. an immediate retransmission is worthwhile. For subsequent timeouts, we prove several general properties that hold for any delay and loss probability distributions. Finally, we demonstrate that, strikingly, when all network users employ the optimal strategy, the network resources can be overloaded in the resulting operating equilibrium.