ABSTRACT

The notion of *timeout* (namely, the maximal time to wait before retrying an action) turns up in many networking con-texts, such as packet transmission, connection establishment, etc. Usage of timeouts is encountered especially in large-scale networks, where negative acknowledgments (NACKs) on failures have significantly higher delays than positive acknowledgments (ACKs) and frequently are not employed at all. Selection of a proper timeout involves a tradeoff between waiting too long and loading the net-work needlessly by waiting too little. The common approach is to set the timeout to a large value, such that, unless the action fails, it is acknowledged within the timeout duration with a high probability. This approach is conservative and leads to overly long, far from optimal, timeouts.

It will take a quantitative approach with the purpose of computing and studying the optimal timeout strategy. The above tradeoff is modeled by introducing a "cost" per unit time (until success) and a "cost" per repeated attempt. The optimal timeout strategy is then defined as one that a selfish user would follow to minimize its expected cost. Then discuss the various practical interpretations that these costs may have, and derive the formulas for the optimal timeout values and study some of their fundamental properties. In particular, it identify the conditions for making parallel attempts from the outset to be worthwhile. In addition, will demonstrate a striking property of *positive feedback*. This motivates us to study the interaction resulting when many users selfishly apply the optimal timeout strategy; specifically, will use a noncooperative game model and show that it suffers from an inherent instability problem. Some implications of these results on network design are discussed