

Dynamically Maximizing the Performance of Large Data Transfer over Dedicated Network Links

Benjamin Eckart

Department of Electrical and Computer Engineering, Tennessee Technological University

EXTENDED ABSTRACT

New networks are emerging for the purpose of transmitting large amounts of scientific data among research institutions quickly and reliably. These characteristics of these networks only marginally resemble those of the Internet and as a result render the established Internet protocols ineffective. Recent methods have been developed to circumvent these problems, including new protocols which implement both reliable (TCP) and unreliable (UDP) data transfer algorithms. The construction of faster networks and more efficient protocols, however, doesn't necessarily result in better data transfer performance. At speeds of 10 Gigabits per second or more, the end-systems involved may be unable to support such speeds due to hardware limitations of the CPU and hard drive. It is therefore necessary to build a high-speed protocol adaptive the performance of each system. This research develops such a protocol, Performance Adaptive UDP (PA-UDP), which aims to dynamically maximize data transfer performance under many system environments.

A mathematical model is proposed to describe the theoretical basis behind effective buffer and CPU management. Based on the formulas derived from this model, a dynamic flow control algorithm is developed that uses a feedback system to monitor the hosts during the data reception period. By monitoring the performance of the systems, the transfer rate between hosts can be adjusted dynamically to match the optimal sending rate predicted by our equations, achieving stability and high performance. A prototype based on the PA-UDP architecture is implemented as an application level protocol and the experimental results demonstrate that it outperforms other high-speed protocols on commodity hardware. Experiments also show that the prototype closely matches the theoretically optimum throughputs predicted by our formulas. We can conclude that PA-UDP is an effective protocol for reliable high-speed bulk data transfer.