

CSC2231 – Internet Systems and Services

Paper Review – High Availability

Name: Alex Wun

Date: Nov. 13th, 05

Blake and Rodrigues make a convincing argument in this paper about the limitations of P2P networks as a high availability data store. They model the requirements of a high availability P2P system and show that the bandwidth needed to maintain availability is fundamentally impractical (in Gnutella-like systems). There are two reasons why this is the case. First, the capacity of hard disks is increasing at a much faster rate than the available bandwidth of home users. And second, the unstable uptimes of most P2P hosts results in high upkeep costs.

This paper concisely presents a very important result: that P2P systems are not suitable for high availability data storage. This task has always been more appropriate for high availability clusters. In my limited experience, it seems that there exists research topics that treat P2P systems as just another “decentralized” architecture. However, as these authors have shown, this is not the case. There are characteristics such as heterogeneity and unstable peer availability that differentiate P2P networks from traditional “decentralized” networks.

Although the early computations are based mostly on the unrealistic assertion that an arriving peer immediately downloads all data that it later serves, the majority of their arguments are convincing. Also, the authors do indeed discuss alternate schemes in which arriving nodes do not necessarily download data immediately.

The results of this paper suggest that P2P topologies are not practical for large-scale data distribution. Instead, it seems that applications such as sensor/ad-hoc networks or ubiquitous computing have characteristics that are more suited for P2P topologies. Perhaps P2P networks can be selectively applied to roaming technologies: i.e. using P2P routing in dense downtown areas (where everyone and his dog has a cellphone or PDA) and switching to base-station routing only in more sparse areas. This is just a spur of the moment idea ...