What Does it Take to Disconnect a P2P Network?*

Dmitri Loguinov

Department of Computer Science
Texas A&M University, College Station, TX 77843
dmitri@cs.tamu.edu

1 Abstract

We analyze the problem of network disconnection in the context of large-scale P2P networks and attempt to understand how both static and dynamic patterns of node failure affect the resilience of such graphs. We start by applying classical results from random graph theory to show that a large variety of deterministic and random P2P graphs almost surely (i.e., with probability $1-o(1)$) remain connected under random failure if and only if they have no isolated nodes. This simple, yet powerful, result subsequently allows us to derive in closed-form the probability that a P2P network develops isolated nodes, and therefore partitions, under both types of node failure. We finish with examples that demonstrate that dynamic P2P systems are extremely resilient under node churn as long as the neighbor replacement delay is much smaller than the average user lifetime.

2 Main Result

For the *static* disconnection pattern (i.e., all nodes fail with independent probability p), almost every existing P2P graph G survives node failure without partitioning with probability:

$$P(G \text{ is connected}) = e^{-n(1-p)p^k},$$

where n is the number of users and k is the fixed degree of each user.

In the *dynamic* case (i.e., users depart based on expiration of individual lifetimes), define L_i to be the random lifetime of user i, S_i to be the random search delay to find the replacement for the i-th failed user, T to be the random delay before a given joining user becomes isolated (i.e., time instance when all of its neighbors are in the failed state), and $\phi = P(T < L_i)$ to be the probability that the isolation event happens within the lifetime of a user. Then, for $E[S_i] \ll E[L_i]$, the following approximation holds:

$$E[T] \approx \frac{E[S_i]}{k}\left[\left(1+\frac{E[R_i]}{E[S_i]}\right)^k - 1\right]$$

Using the above result, the probability of isolation ϕ can be upper-bounded by $E[L_i]/E[T]$ for all distributions of lifetime *with exponential or heavier* tails. Translating this into global resilience, out main result is that almost every dynamic P2P graph can survive N user joins without disconnecting with probability *no less than* $(1-\phi)^N$.

*Supported by NSF grants CCR-0306246, ANI-0312461, CNS-0434940, CNS-0519442.