On Static and Dynamic Partitioning Behavior of Large-Scale Networks

Derek Leonard

Department of Computer Science
Texas A&M University
College Station, TX 77843

27th October 2005

Outline

1. Background
    - Motivation
    - Related Work
2. Static Resilience
    - Verification of Classical Result
    - Node Isolation Model
3. Dynamic Resilience
    - Lifetime Model
    - Verification of Classical Result
    - Disconnection Model

Outline

Motivation I

Questions about Peer-to-Peer Networks

- The recent explosion of Peer-to-Peer networks has sparked interest in several ancillary fields
  - routing efficiency, decentralized security, etc.
- The decentralized nature of P2P systems has brought their resilience into question
  - How many neighbors must a node maintain to remain connected to the system?
  - Is it likely that a P2P graph will partition into more than one subgraph?
  - How long can we expect the system to maintain desirable conditions for operation?

Motivation II

Static and Dynamic Resilience

- Peer-to-Peer network resilience has been explored in two directions
  - Static failure: analysis of fully-populated networks after simultaneous node failures with independent probability p
  - Dynamic failure: analysis of systems in which users join and leave according to some arrival/departure process

Our Contribution

We show that the problem of graph partitioning under *both* types of failure can be reduced to the probability that a P2P network develops at least one isolated node

Outline

1. Background
   - Motivation
   - Related Work
2. Static Resilience
   - Verification of Classical Result
   - Node Isolation Model
3. Dynamic Resilience
   - Lifetime Model
   - Verification of Classical Result
   - Disconnection Model

Related Work I

Random Graph Connectivity

- Erdös and Rényi in the 1960s demonstrated that *almost every* (i.e., with probability $1 - o(1)$ as $n \to \infty$) random graph is connected if and only if it has no isolated vertices.
- Define $\Phi(G)$ to be the probability that graph G remains connected under node or edge failure

$$\Phi(G) = P(G \text{ has no isolated nodes})$$

- It can be shown after some technical manipulation that this holds for certain *deterministic* networks as well
  - Burtin and then later Bollobas prove that this holds under independent uniform node failure for the hypercube

Related Work II

Deterministic Graph Connectivity

- Connectivity of *generic* deterministic graphs is called *residual node connectivity*:

$$\Phi(G) = \sum_{i=1}^{n} S_i(G)p^{n-i}(1-p)^i,$$

  where $S_i(G)$ is the number of connected induced subgraphs of G with exactly i nodes
  - $S_i(G)$ is an NP-complete metric, whose expression is unknown even for the hypercube
- Najjar and Gaudiot proposed a combinatorial model for the probability of disconnection of k-regular graphs, a result we compare to our model

Outline

Verification of Classical Result I

Node Isolation

- Recall that $\Phi(G)$ is defined as the probability that G remains connected under node or edge failure:

$$\Phi(G) = P(G \text{ has no isolated nodes})$$

- We verified this by simulating many well-known P2P networks under different values of p, with one example below

Table: Chord with $n = 16384$ under p-percent failure

p	$\Phi(G)$	P(no isolated nodes)
0.5	0.99996	0.99996
0.6	0.99354	0.99354
0.7	0.72619	0.72650
0.8	0.00040	0.00043

Verification of Classical Result II

Percentage of Disconnections Due to Isolation

- To enhance understanding of the classical result, we introduce metric $q(G)$ that captures the percentage of graph disconnections that contain at least one isolated node:

$$q(G) = \frac{P(X > 0)}{1 - \Phi(G)},$$

  where X is the number of isolated nodes.
- We calculated $q(G)$ for all tested graphs and all values of p
- Metric $q(G)$ ranged from 0.9966 and 1, further verifying that almost every disconnection occurs with at least one isolated node

Outline

Node Isolation Model I

Preliminaries

- We now develop a simple closed-form model for the probability that there are no isolated nodes under static node failure
- Assume that each node i has k_i neighbors in some graph G
- Define X_i to be a Bernoulli indicator variable measuring whether node i is isolated or not after each node has been removed from the system with independent probability p:

$$X_i = \begin{cases} 1 & \text{isolated and alive} \\ 0 & \text{otherwise} \end{cases}.$$

- Denote by $p_i = P(X_i = 1) = (1-p)p^{k_i}$ the probability that i is isolated and alive after the failure

Node Isolation Model II

General Model of Isolation

- Define $X = \sum_{i=1}^{n} X_i$ to be the total number of isolated nodes in G
- Notice that X is a sum of a large number of Bernoulli random variables
- Due to the diminishing dependency between $\{X_i\}$ as $n \to \infty$ we apply the Chen-Stein method to X

Proposition

The number of isolated vertices X tends to a Poisson distribution with mean $\lambda = \sum_i p_i$ and the probability $\Phi(G)$ of having a connected graph converges to $e^{-\lambda}$ with probability 1 as $n \to \infty$

Node Isolation Model III

Degree-Regular Model

For degree-regular networks, the proposition simplifies to:

$$\Phi(G) = e^{-n(1-p)p^k}$$

Degree-Irregular Model

For degree-irregular networks, the following holds:

$$\Phi(G) = e^{-(1-p)\sum_i p^{k_i}} \approx e^{-n(1-p)E[p^{k_i}]},$$

where $\sum_i p^{k_i}$ is approximated by $nE[p^{k_i}]$ treating k_i as a random variable.

Najjar and Gaudiot Contribution

Najjar and Gaudiot Model

Their model for the probability of a connected graph is as follows:

$$\Phi(G) = \sum_{i=0}^{n} Q_i \binom{n}{i} p^i (1-p)^{n-i},$$

where

$$Q_i = \prod_{j=1}^{i} \left[1 - \frac{k(n-k-1)!(j-1)!(n-j)!}{(n-1)!(j-k)!} \right].$$

Problem

Q_i is not a simple calculation

Node Isolation Model IV

Table: Simulation results and model for two regular graphs.

p	Chord $n = 16384$, $k = 27$			de Bruijn $n = 20736$, $k = 24$		
	$\Phi(G)$	Model	Najjar	$\Phi(G)$	Model	Najjar
.4	.9999	1	.9986	.9999	.9999	.9955
.45	.9999	1	.9984	.9999	.9999	.9948
.5	.9999	.9999	.9982	.9993	.9994	.9940
.55	.9992	.9993	.9976	.9944	.9945	.9892
.6	.9935	.9933	.9916	.9618	.9615	.9550
.65	.9500	.9503	.9463	.7954	.7907	.7750
.7	.7262	.7239	.7055	.3199	.3037	.2737
.75	.1788	.1766	.1501	.0079	.0055	.0033
.8	.0004	.0004	.0002	0	10^{-9}	10^{-10}

Node Isolation Model IV

Observations

- Note that our model is very accurate for all values of disconnection probability p
- Further note that the significantly more complex result of Najjar and Gaudiot is less accurate than our model
- There have been several other attempts at solving this model, with none to our knowledge being both simple and accurate simultaneously
- We now show that our model also applies to *irregular* graphs, which has not been attempted previously in the literature

Node Isolation Model V

Table: Simulation results and model for three irregular graphs.

p	Symphony		Gnutella		Randomized Chord	
	$\Phi(G)$	Model	$\Phi(G)$	Model	$\Phi(G)$	Model
.4	.9999	.9999	.9932	.9934	.9999	.9999
.45	.9998	.9996	.9661	.9666	.9999	.9999
.5	.9977	.9977	.8626	.8646	.9997	.9997
.55	.9875	.9875	.5804	.5829	.9975	.9976
.6	.9391	.9394	.1708	.1700	.9844	.9845
.65	.7552	.7535	.0055	.0053	.9162	.9151
.7	.3115	.3107	0	10^{-7}	.6375	.6372
.75	.0127	.0122	0	10^{-15}	.1299	.1282
.8	0	10^{-7}	0	10^{-34}	.0003	.0002

Static Resilience Wrap-up

Static Resilience Wrap-up

- We confirmed through simulations that $\Phi(G) = P(X = 0)$ for most P2P graphs under static node failure
- We derived a simple closed-form model for $P(X = 0)$ for both k-regular and irregular graphs with large n and verified the model with simulations
- The connectivity of P2P and other large networks under uniform independent node failure has been reduced to a simple expression
- We now consider the case of dynamic node failure

Outline

Overview of Lifetime Model I

Lifetime-based Node Failure

- What can be said about node-failure in real-world P2P systems?
  - The p-percent model may be useful in some cases; however, there is no evidence that such failure patterns occur in real P2P networks
  - Nodes arrive/depart dynamically instead of remaining static
- Model: we assign each user a random lifetime L_i from a distribution $F(x)$ the reflects the behavior of the user and represents the duration of his/her service (e.g., sharing files) to the system

Overview of Lifetime Model II

Model Assumptions

- Arrival: nodes arrive randomly according to any process; however, their arrival times are uncorrelated with lifetimes of existing nodes
- Departure: nodes deterministically die (fail) after spending L_i time units in the system
- Neighbor selection: neighbors are picked from among the existing nodes using any rules that do not involve node lifetimes or age (e.g., based on random walks, DHT space assignment, topological locality, content interests, etc.)
- Neighbor replacement: once a failed neighbor is detected, a replacement search is performed

Overview of Lifetime Model III

Definition

A node becomes isolated when all of the neighbors in its table are in the failed state

Node Departure

- All departures are considered to be abrupt, requiring each node to search for a replacement upon failure of its neighbor

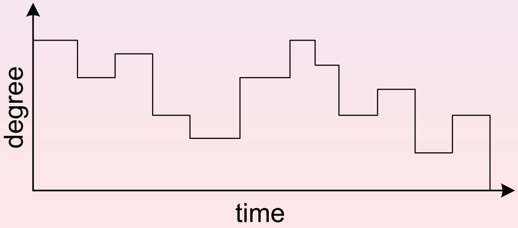

Overview of Lifetime Model IV

Lifetimes of Neighbors

- Node v enters at time t_v with random lifetime L_v
- The k neighbors of v are represented by residual lifetimes

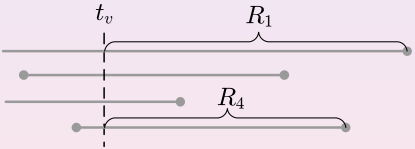

Definition

Let R_i be the remaining lifetime of neighbor i when v joined the system

Overview of Lifetime Model V

Formalizing Search Time

- How do nodes replace neighbors?
  - There is usually some mechanism for detecting that a neighbor has failed (e.g., periodic probing, etc.)
  - Systems often repair the failed zone of a DHT or find a random replacement neighbor in unstructured systems
- We allow this process to be arbitrary as the technique employed has no effect on our results

Definition

Let S_i be a random variable describing the total search time for the i-th replacement in the system

Outline

Verification of Classical Result I

Preliminaries

- We now confirm the classical result in the case of dynamic node failure
- Assume a graph G in which nodes join and leave the system according to the lifetime model
- Define Z to be the random time (in terms of user joins) when G disconnects for the first time
- Assign each joining node i a Bernoulli random variable X_i that determines whether the user is isolated from the network during its lifetime

Verification of Classical Result II

Model

- The probability that the graph stays connected for more than N user joins is almost surely:

$$P(Z > N) = P\left(\bigcap_{i=1}^{N}[X_i = 0]\right) = \prod_{i=1}^{N}(1 - E[X_i]).$$

- For k-regular graphs, each user has the same probability of isolation ϕ and the above reduces to:

$$P(Z > N) = (1 - \phi)^N$$

Verification of Classical Result III

Simulations

- We verify this using simulations where both $E[X_i]$ and ϕ are computed empirically

Table: Simulations for 12-regular CAN with $N = 10^6$ and degree-irregular Chord with $k \approx 13$ and $N = 50,000$

Search time	CAN	Model	Chord	Model
6	.9732	.9728	.6295	.6251
7.5	.8118	.8124	.3284	.3184
8.5	.5669	.5659	.2189	.2206
9	.4065	.4028	.1460	.1483
9.5	.2613	.2645	.1211	.1274
10.5	.0482	.0471	.0493	.0493

Verification of Classical Result IV

Implications

- As with the static case, under dynamic node failure the probability that the graph stays connected is equal to the probability of no isolated nodes
- Interestingly, in the case where disconnection does occur, the largest connected component of dynamic systems almost always contains *exactly* $n - 1$ nodes
- *For reasonably small search delays, network partitioning in lifetime-based systems almost surely effects only a single node in the system*
- We now derive the probability of isolation ϕ for the lifetime model

Outline

Probability of Isolation I

Preliminaries

- What is the probability ϕ that a node will become isolated from the network during its lifetime?
  - Denote by T the time at which a node is isolated when all of its neighbors are simultaneously in the failed state
  - Then $\phi = P(T < L_v)$ for node v with lifetime L_v
- We model the neighbor failure/replacement procedure as an on/off process $Y_j(t)$

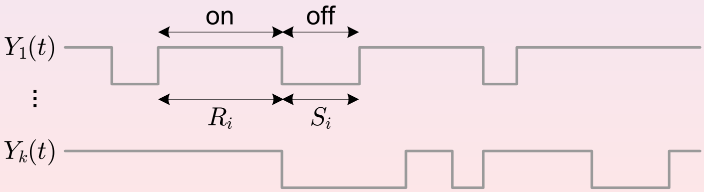

Probability of Isolation II

Degree Evolution

- The degree of node v at time t is:

$$W(t) = \sum_{j=1}^{k} Y_j(t)$$

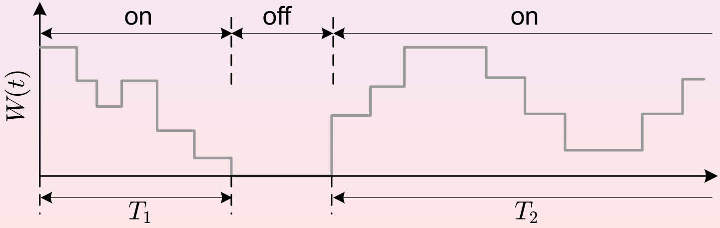

Probability of Isolation III

Result

- By deriving the expected time $E[T]$ before the first visit to state 0 by $W(t)$ and using an exponential approximation to the density of T, we obtain the following as $E[S_i] \to 0$:

$$\phi = \frac{\rho k}{(1+\rho)^k + \rho k - 1},$$

  where $\rho = E[L_i]/E[S_i]$ is the ratio of the mean user lifetime to the mean search delay
- We verify this result in simulations using four different distributions of search delay

Probability of Isolation IV

Simulations

We simulated a system with $E[L_i] = 0.5$ and $k = 8$ using four search distributions to verify the model

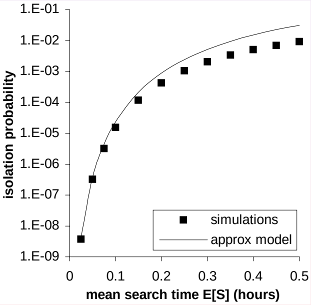

(a) exponential S_i

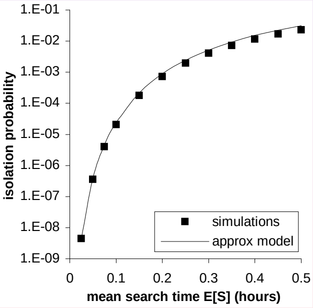

(b) constant S_i

Probability of Isolation V

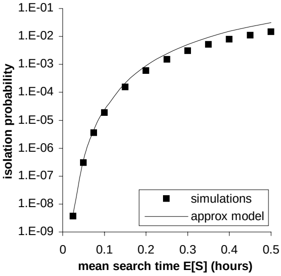

(c) uniform S_i

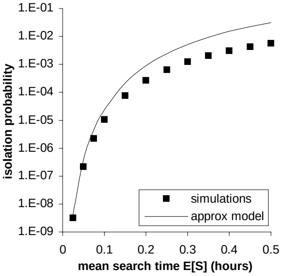

(d) Pareto S_i with $\alpha = 3$

Simulations

Note that as $E[S_i]$ becomes small the simulations converge to the model

Graph Disconnection I

Dynamic Disconnection Revisited

- We now apply the newly acquired model for ϕ to the dynamic graph disconnection model and verify its accuracy in simulations:

$$P(Z > N) = (1-\phi)^N \geq \left(1 - \frac{\rho k}{(1+\rho)^k + \rho k - 1}\right)^N,$$

  where Z is the number of user joins before the first disconnection of the system
- We now verify this result with simulations on 12-regular CAN with exponential lifetimes, $E[L_i] = 0.5$ hours, $n = 4096$, and $N = 10^6$ user joins

Graph Disconnection II

Table: Comparison of $P(Z > N)$ in CAN

Fixed search time (min)	Actual $P(Z > N)$	Model	Metric $q(G)$
6	.9732	.9728	1
7.5	.8118	.8215	1
8.5	.5669	.5666	1
9	.4065	.4016	1
9.5	.2613	.2419	1
10.5	.0482	.0424	1

Observation

Even with large search times, simulations follow the model well

Wrap-up

Wrap-up

- In the static case, we verified that most P2P graphs disconnect almost surely with at least one isolated node
- We then derived a model for graph connectivity using this property and verified its accuracy in simulations
- In the dynamic case, we also confirmed that graphs almost surely disconnect with *exactly* one isolated node
- We then derived the probability of isolation under lifetime-based node failure and verified the result via simulations
- *In both cases, global resilience has been effectively reduced to local resilience*