This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2010 proceedings
This paper was presented as part of the main Technical Program at IEEE INFOCOM 2010.

Characterizing the Spread of Correlated Failures in Large Wireless Networks

Yi Xu and Wenye Wang
Department of Electrical and Computer Engineering
North Carolina State University, Raleigh, NC, USA
Email: {yxu2, wwang}@ncsu.edu

***Abstract*—Correlated failures pose a great challenge for the normal functioning of large wireless networks, because an initial local failure may trigger a global sequence of related failures. Given their potentially devastating impact, we characterize the spread of correlated failures in this paper, which lays the foundation for evaluating and improving the failure resilience of existing wireless networks. We model the failure contagiousness as two generic functions: the failure impact radius distribution function $f_r(x)$ and the failure connection function $g(x)$. By using the percolation theory, we determine the respective characteristic regimes of $f_r(x)$ and $g(x)$ in which correlated failures will and will not percolate in the network. As our model represents various failure scenarios, the results are generally applicable in understanding the spread of a wide range of correlated failures.**

I. Introduction

In large wireless networks, failures happen unavoidably as a result of the network structure complexity and the communication medium openness. With large network size, it is difficult to guarantee every node to work normally and hence failures have an increased chance to occur. Besides, the nodes in an open wireless environment are exposed to various adversary factors, such as unexpected environmental changes and malicious network attacks, which threaten the correct functioning of the nodes. When failures happen, a network cannot achieve its designed performance specifications.

Given the fact of failure inevitability and its detrimental consequences, many studies have been dedicated to the design of effective solutions that prevent [1], [2], detect [3], [4], mitigate [5], [6] and repair [7], [8] failures to retain normal network functioning. The majority of these studies deal with specific types of failures and their countermeasures, for example, planning traffic paths away from the failure-prone regions [2], providing continuous surveillance coverage when a subset of sensors fail [7], finding alternative routes when nodes or links become unavailable [8], [9], and using small-sized packets [10] and error control codes [11] to communicate in networks with low-quality links. These studies have greatly contributed to the failure resilience of large networks.

Different from these existing studies, we are interested in this paper in quantifying the impact of failures. Instead of inspecting the details of a particular failure, we intend to find out the extent of network changes in its composition and structure at the occurrence of failures. In the literature, the impact of failures has been studied with the assumption that failures are random and independent. In [12], Xing and Wang analyzed the critical phase transition time of a large wireless network in which random failures gradually break down an initially connected network into small pieces of disjoint components. Their results present an insightful observation on the impact of failures from the network connectivity perspective. However, the work in [12] does not characterize all the failure possibilities. In many network scenarios, causal relations exist among failures, i.e., some failures happen as a result of other earlier failures. One example of correlated failures is traffic overloading and energy depletion [13]. When a node fails, its traffic is redistributed to the neighboring nodes. Some neighbors may work under heavy traffic loads, resulting in early energy depletion and node failure. Another commonly observed case of failure correlation is virus propagation. Computer viruses not only sabotage the infected nodes, but also reproduce themselves to propagate the infection. In comparison to the random and independent network failures, correlated failures pose an even greater challenge for keeping a large wireless network healthy.

In view of the potentially devastating impact caused by correlated failures, we focus in this paper on characterizing the spread of correlated failures in large wireless networks. In particular, we attempt to determine the conditions under which a single initial failure will and will not spread to the entire network. In an effort to gain a generally applicable understanding on the spread of correlated failures, we model the failure correlations as general functions and determine their characteristic regimes in terms of the ability of an initial failure to permeate the network. The correlation functions model the geometric constraints in failure propagation, i.e., the distance and probability for a failure to spread in one hop. Based on the correlation functions, we then use the percolation theory [14], [15] to tackle the failure spreading problem. Percolation theory provides a mighty tool for analyzing a wide range of contact-and-relay problems observed in reality. We determine analytically the pervasiveness of failure spreading conditioned on the failure correlation functions by using the concept of percolation. Intuitively, stronger correlations drive an initial failure to spread into a larger area than weaker correlations. Our results provide a quantified answer on the relation between failure correlations and failure pervasiveness in large wireless networks. The results hence help network designers and operators evaluate the resilience of large wireless networks to the

978-1-4244-5837-0/10/$26.00 ©2010 IEEE

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2010 proceedings
This paper was presented as part of the main Technical Program at IEEE INFOCOM 2010.

correlated node failures.

The rest of this paper is organized as follows. We describe the network model and formulate the failure percolation problem in Section II. Section III analyzes one type of failure correlations in which a node is subject to failure only at the first contact from other failed nodes. Sections IV and V present the analysis of another type of failure correlations in which a node may fail each time when contacted by other failed nodes. Finally, Section VI concludes this paper.

II. Problem Formulation

Failure spreading can be modeled as a *bond percolation* process. In its most primitive form, bond percolation [15] depicts the phenomenon of liquid transfusion in a porous solid. The solid contains many tiny interstices (called sites) and neighbor interstices are connected by open paths (called bonds) with some probability. Bond percolation states that there exists a critical bond open probability that demarcates two different phases of the liquid distribution in the solid. If bonds are open with sufficiently high probability, the liquid can reach a giant cluster of sites. Otherwise, liquid transfusion is confined in a negligibly small space. Given the similarity between failure spreading and liquid transfusion, we use the analytical techniques of percolation theory to characterize failure spreading in large wireless networks. To proceed, we first define the network model.

A. Network Model

We consider a large wireless network consisting of n nodes in a region $\mathcal{B} = [-\frac{L}{2}, \frac{L}{2}]^2$ ($L \to \infty$). The number of nodes n is assumed to be a Poisson distributed random variable with constant density λ everywhere in the network. Let X_i ($1 \leq i \leq n$) denote the random location of node v_i that is uniformly distributed in the network, independent of n and any X_j ($i \neq j$). By definition [16], $\mathcal{H}_\lambda = \{X_1, \cdots, X_n\}$ is a homogeneous Poisson point process.

The nodes are correlated in their working status. When a node fails, it may trigger other failures in its neighbors. We define two probabilistic functions to model failure correlations. Let $\|X_i - X_j\|$ denote the Euclidean distance between v_i and v_j. We define the *failure impact radius* r to be the farthest distance between the location of a triggering failure and the location of an immediate follow-up failure, i.e., failure may propagate from v_i to v_j in one hop only if $\|X_i - X_j\| \leq r_i$, where r_i is the r of v_i. Considering the difference of the nodes in their respective failure impacts, for example a failed node with higher load may cause more neighbors to overload and a virus-infected node with higher transmission power may pass the virus to others at longer distance, the impact radius r is modeled as a random variable with probability density function $f_r(x)$ ($0 \leq x \leq R$). By the scaling property [14], we choose $R = 1$ without loss of generality. For different nodes v_{i_1} and v_{i_2}, r_{i_1} and r_{i_2} are independent. Besides, we define the *failure connection function* $g(x)$ to model the likelihood of failure propagation from v_i to v_j. If v_j is located within the impact radius of v_i, failure spreads to v_j with a probability $g(\|X_i - X_j\|)$ that depends on their distance but not on their respective locations. If v_j is beyond the impact radius of v_i, failure cannot spread from v_i to v_j. For any two nodes v_{j_1} and v_{j_2}, failure propagation is independent from each other.

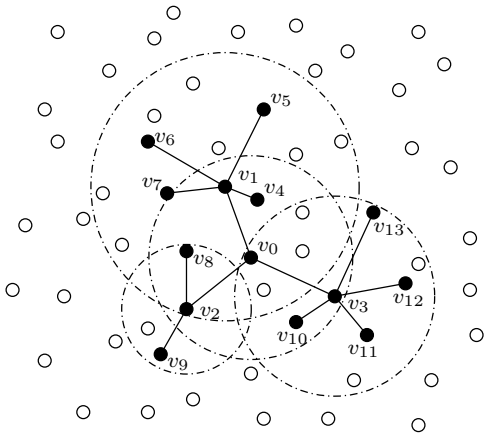

Fig. 1. Failure spreading in large wireless networks.

When failure propagates from v_i to v_j in one hop, we say that v_i and v_j are connected (or the connection is *open*). Otherwise, v_i and v_j are not connected (or the connection is *closed*). Note that the connections considered in this paper represent the failure correlations among nodes, which are completely different from the communication links.

In percolation terminology, each node in the wireless network is a site and failure connections define the bonds between neighboring sites. As node locations are randomly distributed, this percolation process is also known as *continuum percolation* [14]. In its basic form, continuum percolation assumes an open bond between any two neighboring sites. We have introduced here the probabilistic functions $f_r(x)$ and $g(x)$ for each bond, so our network model is also called the *random connection model* [14] in which a bond is open only probabilistically.

B. Failure Percolation

Given node locations and their bonds, our problem is to determine whether an initial failure will spread to the entire network. An example of failure spreading is illustrated in Fig. 1. In this example, the initial failure occurs at node v_0. As a result of this failure, nodes v_1–v_3 fail subsequently and spread the failure further away to nodes v_4–v_{13}. In each step of spreading, a node that has just failed in the previous step passes failure to a random subset of nodes in its neighborhood, as modeled by the impact radius distribution function $f_r(x)$ and the connection function $g(x)$. As time goes, there are two possible results regarding failure spreading: either the spreading continues for ever or it stops automatically. Percolation theory tells us that if each bond is open with sufficiently high probability, failure percolates in the network; otherwise, failure only reaches a negligibly small area. Therefore, to determine the trend of failure spreading, our task is to characterize the failure correlations via functions $f_r(x)$ and $g(x)$.

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2010 proceedings
This paper was presented as part of the main Technical Program at IEEE INFOCOM 2010.

Next, we formulate the problem mathematically. We represent a wireless network as a *random geometric graph* [16] and denote it as $G(\mathcal{H}_\lambda, f_r(\cdot), g(\cdot))$. Regarding node v_0 that initiates the failures, we define the *percolation probability* $p_\infty(\lambda, f_r(\cdot), g(\cdot))$ as

$$p_\infty(\lambda, f_r(\cdot), g(\cdot)) \triangleq \Pr[|C(v_0)| = \infty], \tag{1}$$

where $C(v_0)$ denotes the cluster of nodes that fail as a result of the initial failure at v_0 and $|C(v_0)|$ denotes the size of cluster $C(v_0)$. When $|C(v_0)| = \infty$, we call $C(v_0)$ a *giant component*. If $p_\infty(\lambda, f_r(\cdot), g(\cdot)) > 0$, failure percolates in the network with a positive probability. If $p_\infty(\lambda, f_r(\cdot), g(\cdot)) = 0$, failure does not percolate in the network almost surely.

Note that $p_\infty(\lambda, f_r(\cdot), g(\cdot)) > 0$ indicates only the possibility for percolation, in addition to which the actual occurrence of percolation depends on the initial failure location. To be more accurate, when $p_\infty(\lambda, f_r(\cdot), g(\cdot)) > 0$, there exists an inter-connected giant component of nodes almost surely. If the initial failure happens at a node inside this giant component, failure percolates. Otherwise, failure does not percolate. For the purpose of concise presentation, we say that failure percolates whenever $p_\infty(\lambda, f_r(\cdot), g(\cdot)) > 0$, by assuming that the initial failure location belongs to the giant component. Besides, when percolation happens, failure does not necessarily reach every node. Rather, it reaches a significant portion of the network (at least a constant percentage). We say that failure spreads to the entire network because it cannot be confined in any finite area when network size increases to infinity.

Our goal is to determine the regimes of functions $f_r(x)$ and $g(x)$ such that $p_\infty(\lambda, f_r(\cdot), g(\cdot)) > 0$ and $p_\infty(\lambda, f_r(\cdot), g(\cdot)) = 0$ with given λ, respectively. Obviously, $G(\mathcal{H}_\lambda, f_r(\cdot), g(\cdot))$ is a subgraph of $G(\mathcal{H}_\lambda, 1, 1)$ and $p_\infty(\lambda, f_r(\cdot), g(\cdot)) = 0$ whenever $p_\infty(\lambda, 1, 1) = 0$. To avoid triviality, we only consider the case $p_\infty(\lambda, 1, 1) > 0$ and determine the respective constraints on $f_r(x)$ and $g(x)$ for $p_\infty(\lambda, f_r(\cdot), g(\cdot)) > 0$ and $p_\infty(\lambda, f_r(\cdot), g(\cdot)) = 0$. It is well known that there exists a critical density λ_c in $G(\mathcal{H}_\lambda, 1, 1)$ defined as

$$\lambda_c \triangleq \inf\{\lambda > 0 : p_\infty(\lambda, 1, 1) > 0\} \tag{2}$$

that specifies the minimum λ for $p_\infty(\lambda, 1, 1) > 0$. So far, the best known rigorous bounds on λ_c are $0.7698 < \lambda_c < 3.372$ [17]. Hence, we assume $\lambda > 3.372$ in the rest of this paper. In addition, if r is a constant, we also assume $\lambda r^2 > 3.372$, which guarantees $p_\infty(\lambda, r, 1) > 0$ by the scaling property [14].

III. Percolation of One-Time Failures

In this section, we study the percolation of one possible type of failures. We say that the failures are *one-time* if each node is subject to the impact of other nodes only once. In reality, these failures happen due to a common security risk that is shared by a subset of nodes in the network. For example, a computer virus that exploits the security vulnerability of a commonly used software may infect the nodes that have installed this software. Once a node is infected, it continues to spread the virus to other nodes. However, each node can be infected only once. If a node has installed this software, it fails at the first contact from a virus-carrying node. If not, it is immune to the first and any onward attacks.

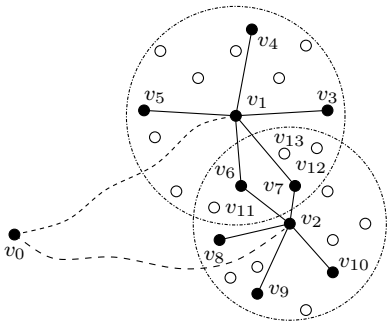

Fig. 2. Spread of one-time failures.

We illustrate an example of one-time failures in Fig. 2, in which failure has occurred at v_0 and propagated via some paths to v_1 and v_2. Node v_1 spreads the failure to nodes v_3–v_7, but not to v_{11}–v_{13}. Correspondingly, node v_2 must also spread the failure to v_6–v_7 but not to v_{11}–v_{13}, because v_6–v_7 are subject to this failure while v_{11}–v_{13} are immune.

For one-time failures, a node fails due to its own vulnerability, which is independent of the specific contact node and their distance. Therefore, we can simplify the failure connection function as $g(x) \equiv p$, where $0 < p < 1$ is a constant. Furthermore, $g(x)$ applies only to the first connection from a failed node. For subsequent connections, $g(x) = 0$ if the node did not fail at the first contact. For clearer presentation, we denote the percolation probability of one-time failures as $p_\infty^{(\text{one})}(\lambda, f_r(\cdot), p)$, on which we have the following lemma.

Lemma 1: The spread of one-time failures with $g(x) \equiv p$ can be modeled equivalently as a site percolation problem in which each site is open with probability p independently.

Proof: Before the initial failure occurrence, we predetermine the working status of each node independently with the failure probability p. Obviously, percolation in the resulting graph of predetermined failed nodes has the same property as in the original graph in which each node may fail at the first contact from a failed node. Thus, the original bond percolation problem is equivalent to a site percolation problem. ■

By modeling the spread of one-time failures as site percolation, we next characterize the regimes of $f_r(x)$ and p under which $p_\infty^{(\text{one})}(\lambda, f_r(\cdot), p) > 0$ and $p_\infty^{(\text{one})}(\lambda, f_r(\cdot), p) = 0$ respectively.

A. Percolation with Constant Failure Impact Radius

First, we discuss the special case of a constant r for every node. In this case, we have the following theorem.

Theorem 1: For constant failure impact radius r,

i) $p_\infty^{(\text{one})}(\lambda, r, p) > 0$ if $pr^2 > \frac{\lambda_c}{\lambda}$,
ii) $p_\infty^{(\text{one})}(\lambda, r, p) = 0$ if $pr^2 < \frac{\lambda_c}{\lambda}$.

Proof: By Lemma 1 and Thinning Theorem [16], the predetermined failed nodes have density $p\lambda$ and $p_\infty^{(\text{one})}(\lambda, r, p) = p_\infty(p\lambda, r, 1)$. By the definition of λ_c in Eq. (2) and the scaling property [14], the critical density for $p_\infty(\lambda, r, 1) > 0$ is $\frac{\lambda_c}{r^2}$.

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2010 proceedings
This paper was presented as part of the main Technical Program at IEEE INFOCOM 2010.

Thus, $p_\infty(p\lambda, r, 1) > 0$ if $p\lambda > \frac{\lambda_c}{r^2}$ and $p_\infty(p\lambda, r, 1) = 0$ if $p\lambda < \frac{\lambda_c}{r^2}$. The results follow immediately. ■

B. Percolation with Random Failure Impact Radius

When r is a random variable, the results in Theorem 1 can be generalized, as stated below in Theorem 2.

Theorem 2: For random failure impact radius r with probability density function $f_r(x)$ $(0 \le x \le 1)$,

i) $p_\infty^{(\text{one})}(\lambda, f_r(\cdot), p) > 0$ if $ph_1(f_r(x)) > \frac{\lambda_c}{\lambda}$,
ii) $p_\infty^{(\text{one})}(\lambda, f_r(\cdot), p) = 0$ if $p(1 - h_2(f_r(x))) < \frac{\lambda_c}{\lambda}$,

where $h_1(f_r(x)) = \max_{\{0\le a\le 1\}}\{a^2 \int_a^1 f_r(x)\mathrm{d}x\}$, $h_2(f_r(x)) = \max_{\{0\le a\le 1\}}\{(1-a^2)\int_0^a f_r(x)\mathrm{d}x\}$.

Proof: Similar to the proof of Theorem 1, we have equation $p_\infty^{(\text{one})}(\lambda, f_r(\cdot), p) = p_\infty(p\lambda, f_r(\cdot), 1)$. We next focus on $p_\infty(p\lambda, f_r(\cdot), 1)$.

We first prove statement (i). When r has probability density function $f_r(x)$, given any $0 \le a \le 1$, $r \ge a$ with probability $\int_a^1 f_r(x)\mathrm{d}x$. We decompose graph $G(\mathcal{H}_{p\lambda}, f_r(\cdot), 1)$ into two subgraphs: one consisting of the nodes with $r \ge a$, denoted as $G(\mathcal{H}_{p\lambda\int_a^1 f_r(x)\mathrm{d}x}, f_r(\cdot), 1)$, and the other one consisting of the nodes with $r < a$, denoted as $G(\mathcal{H}_{p\lambda\int_0^a f_r(x)\mathrm{d}x}, f_r(\cdot), 1)$. Obviously, if failure percolates in $G(\mathcal{H}_{p\lambda\int_a^1 f_r(x)\mathrm{d}x}, f_r(\cdot), 1)$, it also percolates in $G(\mathcal{H}_{p\lambda}, f_r(\cdot), 1)$. Furthermore, every node in $G(\mathcal{H}_{p\lambda\int_a^1 f_r(x)\mathrm{d}x}, f_r(\cdot), 1)$ has $r \ge a$. If we reduce their r to the constant a and failure still percolates, then percolation must also happen in $G(\mathcal{H}_{p\lambda\int_a^1 f_r(x)\mathrm{d}x}, f_r(\cdot), 1)$. Failure percolates in $G(\mathcal{H}_{p\lambda\int_a^1 f_r(x)\mathrm{d}x}, a, 1)$ (where $r = a$) when $p\lambda\int_a^1 f_r(x)\mathrm{d}x$ exceeds the critical density $\frac{\lambda_c}{a^2}$, i.e.,

$$pa^2 \int_a^1 f_r(x)\mathrm{d}x > \frac{\lambda_c}{\lambda}. \tag{3}$$

Finally, we note that percolation happens as long as we find any a that satisfies Eq. (3). In other words, condition

$$p \cdot \max_{\{0\le a\le 1\}} \left\{a^2 \int_a^1 f_r(x)\mathrm{d}x\right\} > \frac{\lambda_c}{\lambda} \tag{4}$$

suffices for $p_\infty(p\lambda, f_r(\cdot), 1) > 0$.

We next prove statement (ii). Similar to the graph decomposition method used above, we divide the original graph $G(\mathcal{H}_{p\lambda}, f_r(\cdot), 1)$ into the subgraph $G(\mathcal{H}_{p\lambda\int_a^1 f_r(x)\mathrm{d}x}, f_r(\cdot), 1)$ and the subgraph $G(\mathcal{H}_{p\lambda\int_0^a f_r(x)\mathrm{d}x}, f_r(\cdot), 1)$. For each node in $G(\mathcal{H}_{p\lambda\int_a^1 f_r(x)\mathrm{d}x}, f_r(\cdot), 1)$, we scale its r to ar. As a compensation for reducing r, we increase the node density from $p\lambda\int_a^1 f_r(x)\mathrm{d}x$ to $\frac{p\lambda}{a^2}\int_a^1 f_r(x)\mathrm{d}x$. The resulting subgraph $G(\mathcal{H}_{\frac{p\lambda}{a^2}\int_a^1 f_r(x)\mathrm{d}x}, f_{ar}(\cdot), 1)$ has the same connectivity as subgraph $G(\mathcal{H}_{p\lambda\int_a^1 f_r(x)\mathrm{d}x}, f_r(\cdot), 1)$. We then merge $G(\mathcal{H}_{\frac{p\lambda}{a^2}\int_a^1 f_r(x)\mathrm{d}x}, f_{ar}(\cdot), 1)$ and $G(\mathcal{H}_{p\lambda\int_0^a f_r(x)\mathrm{d}x}, f_r(\cdot), 1)$ to obtain graph $G(\mathcal{H}_{\lambda'}, f_{r'}(\cdot), 1)$ with $\lambda' = p\lambda(\int_0^a f_r(x)\mathrm{d}x + \frac{1}{a^2}\int_a^1 f_r(x)\mathrm{d}x)$ and $r' \le a$, which has the same percolation property as $G(\mathcal{H}_{p\lambda}, f_r(\cdot), 1)$. When $\lambda' < \frac{\lambda_c}{a^2}$, i.e.,

$$p\lambda\Big(\int_0^a f_r(x)\mathrm{d}x + \frac{1}{a^2}\int_a^1 f_r(x)\mathrm{d}x\Big) < \frac{\lambda_c}{a^2}, \tag{5}$$

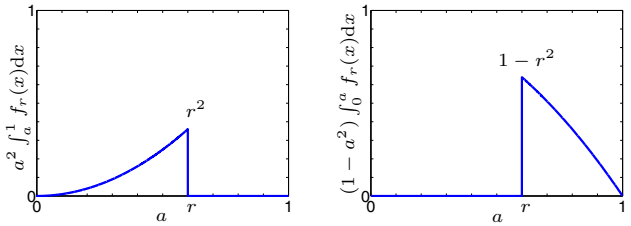

(a) $a^2\int_a^1 f_r(x)\mathrm{d}x$ (b) $(1-a^2)\int_0^a f_r(x)\mathrm{d}x$

Fig. 3. Functions $a^2\int_a^1 f_r(x)\mathrm{d}x$ and $(1-a^2)\int_0^a f_r(x)\mathrm{d}x$ when r is a constant, i.e., $f_r(x) = \delta(x-r)$.

failure does not percolate in $G(\mathcal{H}_{\lambda'}, a, 1)$, and hence not in $G(\mathcal{H}_{\lambda'}, f_{r'}(\cdot), 1)$ and $G(\mathcal{H}_{p\lambda}, f_r(\cdot), 1)$. Rewriting Eq. (5) to

$$p\Big(1 - (1-a^2)\int_0^a f_r(x)\mathrm{d}x\Big) < \frac{\lambda_c}{\lambda}, \tag{6}$$

and noting that percolation does not happen as long as there exists any a that satisfies Eq. (6), we arrive at

$$p\Big(1 - \max_{\{0\le a\le 1\}}\Big\{(1-a^2)\int_0^a f_r(x)\mathrm{d}x\Big\}\Big) < \frac{\lambda_c}{\lambda}, \tag{7}$$

which suffices for $p_\infty(p\lambda, f_r(\cdot), 1) = 0$. ■

Note that Theorem 1 is in fact a special case of Theorem 2. When r is a constant instead of a random variable, $f_r(x) = \delta(x-r)$, where $\delta(x-r)$ is called the *Dirac delta function* with the properties: (i) $\delta(x-r) = \infty$ when $x = r$ and $\delta(x-r) = 0$ when $x \ne r$, and (ii) $\int_{r_-}^{r_+}\delta(x-r)\mathrm{d}x = 1$. We depict functions $a^2\int_a^1 f_r(x)\mathrm{d}x$ and $(1-a^2)\int_0^a f_r(x)\mathrm{d}x$ in Fig. 3 with $f_r(x) = \delta(x-r)$. It is observed that these two functions achieve their maxima at $a = r$. Therefore, we have $h_1(f_r(x)) = r^2$ and $h_2(f_r(x)) = 1-r^2$. Plugging them into Theorem 2, we reach the same results as in Theorem 1.

Intuitively, r and p represent the degree of failure contagiousness. When r and p increase, failure tends to percolate. When they decrease, percolation becomes unlikely to happen. Our results in Theorems 1 and 2 present a quantified measure on r and p. Moreover, in the general case of a random r, Theorem 2 indicates that the chance of failure percolation increases if the probability distribution $f_r(x)$ shifts toward large r (such that $h_1(f_r(x))$ increases) and decreases if $f_r(x)$ shifts toward small r (such that $h_2(f_r(x))$ increases).

IV. Percolation of Persistent Failures: Constant Failure Correlations

In real network scenarios, another widely observed case of failures is that a node may be impacted by every failure that occurs in its neighborhood. We say that the failures are *persistent* if a node may fail each time one of its neighbors fails. Different from one-time failures, persistent failures happen not due to node vulnerability, but because of the interconnection between nodes. One example of persistent failures is traffic overloading and energy depletion. When a node fails, its traffic will be redistributed to the neighbors via dynamic

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2010 proceedings
This paper was presented as part of the main Technical Program at IEEE INFOCOM 2010.

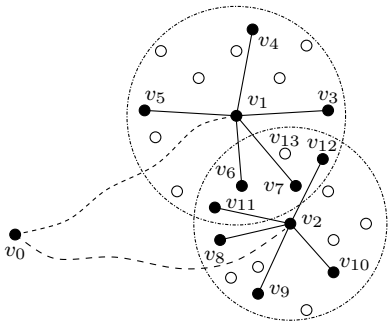

Fig. 4. Spread of persistent failures.

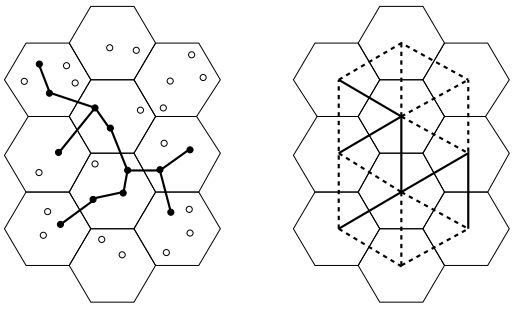

(a) continuous (b) discrete

Fig. 5. Mapping from continuous percolation to discrete percolation.

routing. Some neighbors may fail subsequently after a short time when their energy is depleted by the excessive traffic loads. Obviously, a node has the risk of failure whenever it receives extra load from neighbors.

Compared to one-time failures, persistent failures are easier to percolate in the network. Intuitively, a node subject to failure at each contact from failed nodes has an increased chance to join the cluster of failed nodes and, therefore, the giant component becomes likely to emerge. An example of persistent failures is given in Fig. 4, in which the initial failure at v_0 has spread to v_1 and v_2. Nodes v_1 and v_2 cause further failures at nodes v_3–v_7 and v_8–v_{12} respectively. Note that v_2 does not need to fail v_6–v_7 again (though this could happen) since v_6–v_7 are already part of the failed component connected through v_1. However, v_2 may fail v_{11}–v_{12}, which have withstood the impact of v_1. In comparison to Fig. 2, the component of failed nodes is larger in Fig. 4.

We explore the percolation of persistent failures in this and next sections. Specifically, we focus on constant failure impact radius r and constant failure connection probability p in this section. We will study the generalized functions $f_r(x)$ and $g(x)$ in the following section. Next, we discuss separately the sufficient conditions for percolation and non-percolation of persistent failures with constant r and p. We use $p_\infty(\lambda, r, p)$ in this section to denote the percolation probability of persistent failures with constant r and p.

A. Sufficient Condition of Percolation: Constant r and p

1) A Primitive Condition: Our discussion on the difference between persistent and one-time failures shows that persistent failures are easier to percolate in the network. Therefore, percolation of persistent failures should happen under the same sufficient condition for one-time failures, as stated below.

Theorem 3: For constant impact radius r and connection probability p, $p_\infty(\lambda, r, p) > 0$ if $pr^2 > \frac{\lambda_c}{\lambda}$.

Proof: For each node in the network, we first require that it fail at most once. We let the initial failure to spread in this modified network and obtain a graph G_1 of failed nodes, which is the case of one-time failures. Then, we start a second round of failure spreading, in which every failed node contacts its neighbors and passes the failure with probability p to each of them if the contacted neighbor is in the normal working status. This second round of failure spreading is repeated until there are no more failures. The resulting graph G_2 is the case of persistent failures. Under condition $pr^2 > \frac{\lambda_c}{\lambda}$, percolation happens in G_1. Since G_1 is a subgraph of G_2, percolation must also happen in G_2. ■

Theorem 3 provides a judgment regarding when persistent failures can percolate in a large network. As the current best known rigorous bounds on λ_c are $0.7698 < \lambda_c < 3.372$, we infer that percolation occurs when $pr^2 > \frac{3.372}{\lambda}$. In regime $pr^2 < \frac{3.372}{\lambda}$, Theorem 3 gives us no idea on whether failure percolates or not. Further judgment depends on the improved accuracy of the bounds on λ_c.

2) A Tighter Condition: Since persistent failures are easier to percolate than one-time failures, we expect a tighter sufficient condition that allows percolation with smaller r and p than the one given in Theorem 3. Next, we determine a new condition by using the technique of continuous-to-discrete percolation mapping.

We divide the network area into many small hexagonal cells, as illustrated in Fig. 5(a). The graph of the failed nodes and their connections now appears on the background of these cells. As failure spreads, it travels through a cluster of continuous cells. We define a cell as *open* if it contains at least one failed node and *closed* otherwise. Let $C_{\text{cell}}(v_0)$ denote the cluster of open cells and $|C_{\text{cell}}(v_0)|$ denote its size. It is obviously true that if $|C(v_0)| = \infty$, then $|C_{\text{cell}}(v_0)| = \infty$, and vice versa. The mapping between the cluster of failed nodes and the cluster of open cells thus allows us to find the sufficient condition for $|C_{\text{cell}}(v_0)| = \infty$ and use it for $|C(v_0)| = \infty$.

The cluster size of open cells is studied via bond percolation on a discrete grid, as shown in Fig. 5(b). In this abstract form, we represent a cell by a site located at the center of the cell. Two neighboring sites are connected by a bond, which represents the connectivity between the two corresponding cells. We define the bond to be *open* if, given one or more failed nodes in one of the end cells, at least one failed node connects to some nodes in the other end cell. In other words, an open bond propagates failure from one cell to another. In Fig. 5(b), open bonds are denoted as solid lines that are consistent with the continuous percolation in Fig. 5(a).

We observe in Fig. 5(b) that when the cells have hexagonal shape, the discrete grid is triangular. In the literature, square

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2010 proceedings
This paper was presented as part of the main Technical Program at IEEE INFOCOM 2010.

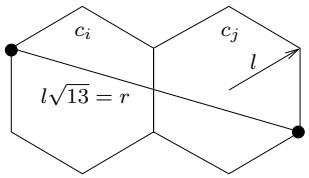

Fig. 6. Determination of cell size.

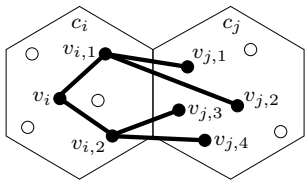

Fig. 7. Failure spread via stepstones.

cells are usually used [18], [19] as they generate square grid that is simple and easy for analysis. We use hexagonal cells in this paper for two reasons. First, hexagonal cells yield a tighter bound on r and p for failure percolation than square cells. Second, they render the study of failure non-percolation possible under this mapping framework. We will discuss these two reasons in details later. With triangular grid, there exists a critical probability $p_c^{\triangle} = 2\sin(\frac{\pi}{18}) = 0.3473$ [15] such that percolation occurs, i.e., $|C_{\text{cell}}(v_0)| = \infty$, if each bond is open with a probability higher than $p_c^{\triangle}$ and not occurs, i.e., $|C_{\text{cell}}(v_0)| < \infty$, otherwise. The discrete percolation in triangular grid allows us to reach the following theorem.

Theorem 4: For constant impact radius r and connection probability p, $p_\infty(\lambda, r, p) > 0$ if $pr^2 > \frac{2.1348}{\lambda}$.

Proof: We choose the size of each cell small enough such that given two arbitrary locations in two neighboring cells, one in each, their distance is at most r. This small cell size guarantees that a failed node has a positive chance to contact and infect every node in the neighboring cells. Two nodes are separated farthest as shown in Fig. 6, in which their distance is $l\sqrt{13}$ and l is the cell radius. Letting $l\sqrt{13} = r$, we obtain $l = \frac{r}{\sqrt{13}}$ and $s = \frac{3\sqrt{3}}{2}l^2 = \frac{3\sqrt{3}}{26}r^2$, where s is the area of each cell. Note that with this small cell size, it is possible for a failed node to connect in one hop to the nodes in non-neighboring cells. However, we can safely ignore these "jumping" connections, because they increase the probability of failure percolation and therefore do not change the validity of our result in this theorem.

We use p^{bond} to denote the bond open probability in the triangular grid. Given a failed node v_i in cell c_i, it connects to each node in the neighboring cell c_j with probability p independently. As the number of nodes in c_j follows a Poisson distribution with density λ, the number of nodes in c_j that fail after v_i is also a Poisson random variable, with density $p\lambda$. Hence, the probability of at least one node in c_j being connected by v_i is $1 - e^{-p\lambda s}$. Considering the possibility that other failed nodes in c_i besides v_i can also connect to the nodes in c_j, failure propagates from c_i to c_j with a probability higher than $1 - e^{-p\lambda s}$ and we have

$$p^{\text{bond}} \geq 1 - e^{-p\lambda s}. \tag{8}$$

Plugging $s = \frac{3\sqrt{3}}{26}r^2$ and solving

$$1 - e^{-p\lambda s} > p_c^{\triangle}, \tag{9}$$

we arrive at

$$pr^2 > \frac{-\ln(1 - p_c^{\triangle})}{\frac{3\sqrt{3}}{26}\lambda} = \frac{2.1348}{\lambda}, \tag{10}$$

under which $\Pr[|C_{\text{cell}}(v_0)| = \infty] > 0$ and $p_\infty(\lambda, r, p) > 0$. ■

We have mentioned earlier that hexagonal cells yield a tighter condition on r and p than square cells. To see this, we substitute $p_c^{\square} = 0.5$ (the critical bond open probability in square grid [15]) for $p_c^{\triangle} = 0.3473$ and $s = \frac{1}{5}r^2$ (the area of a square cell) for $s = \frac{3\sqrt{3}}{26}r^2$. We obtain $pr^2 > \frac{3.4657}{\lambda}$, which is not as tight as $pr^2 > \frac{2.1348}{\lambda}$.

3) An Even Tighter Condition: Though Theorem 4 improves the bound on r and p for failure percolation, it has only considered the case of one-hop failure propagation from node v_i to the nodes in cell c_j. There exists the possibility that v_i may fail some other nodes in c_i first and then spread failure to c_j via these stepstone nodes, as illustrated in Fig. 7. With these stepstones considered, the sufficient condition on r and p can be improved further. Before deriving this improved result, we first present two useful lemmas.

Lemma 2: Given a failed node v_i in cell c_i and a neighboring cell c_j, if none of the nodes in c_j is connected to v_i, then the probability of existing k nodes in c_j is $P(k) = e^{-(1-p)\lambda s}\frac{((1-p)\lambda s)^k}{k!}$, where p is the failure connection probability, λ is the node density, s is the area of a cell.

Proof: Let E_1 denote the event that there exist k nodes in c_j and E_2 denote the event that none of the nodes in c_j is connected to v_i. Then the conditional probability is written as $P(k) = \frac{\Pr[E_1 \cap E_2]}{\Pr[E_2]}$, from which we have

$$\begin{aligned} P(k) &= \frac{e^{-\lambda s}\frac{(\lambda s)^k}{k!}(1-p)^k}{e^{-p\lambda s}} \\ &= e^{-(1-p)\lambda s}\frac{((1-p)\lambda s)^k}{k!}, \end{aligned} \tag{11}$$

which shows that the conditional node distribution is Poisson with density $(1-p)\lambda$. ■

Lemma 3: Given neighboring cells c_i and c_j, assume that the number of failed nodes in c_i is Poisson distributed with density λ_i and the number of nodes in c_j is Poisson distributed with density λ_j. The probability of failure spread from c_i to c_j is bounded as $P \geq (1 - e^{-\lambda_i s})(1 - e^{-p\lambda_j s})$, where p is the failure connection probability, s is the area of a cell.

Proof: As long as there exists one open connection between a failed node in c_i and a node in c_j, failure propagates from c_i to c_j. The probability is then expressed as

$$\begin{aligned} P &= \sum_{k_1=0}^{\infty}\sum_{k_2=0}^{\infty} e^{-\lambda_i s}\frac{(\lambda_i s)^{k_1}}{k_1!}e^{-\lambda_j s}\frac{(\lambda_j s)^{k_2}}{k_2!}(1-(1-p)^{k_1 k_2}) \\ &= 1 - \sum_{k_1=0}^{\infty}\sum_{k_2=0}^{\infty} e^{-\lambda_i s}\frac{(\lambda_i s)^{k_1}}{k_1!}e^{-\lambda_j s}\frac{(\lambda_j s)^{k_2}}{k_2!}(1-p)^{k_1 k_2} \end{aligned}$$

$$
\begin{aligned}
&= 1 - \sum_{k_1=0}^{\infty} e^{-\lambda_i s} \frac{(\lambda_i s)^{k_1}}{k_1!} \sum_{k_2=0}^{\infty} e^{-\lambda_j s} \frac{(\lambda_j s)^{k_2}}{k_2!} (1-p)^{k_1 k_2} \\
&= 1 - \sum_{k_1=0}^{\infty} e^{-\lambda_i s} \frac{(\lambda_i s)^{k_1}}{k_1!} e^{-\lambda_j s (1-(1-p)^{k_1})} \\
&= 1 - e^{-\lambda_i s} - \sum_{k_1=1}^{\infty} e^{-\lambda_i s} \frac{(\lambda_i s)^{k_1}}{k_1!} e^{-\lambda_j s (1-(1-p)^{k_1})} \\
&\geq 1 - e^{-\lambda_i s} - \sum_{k_1=1}^{\infty} e^{-\lambda_i s} \frac{(\lambda_i s)^{k_1}}{k_1!} e^{-p\lambda_j s} \\
&= 1 - e^{-\lambda_i s} - e^{-p\lambda_j s}(1 - e^{-\lambda_i s}) \\
&= (1 - e^{-\lambda_i s})(1 - e^{-p\lambda_j s}), \qquad (12)
\end{aligned}
$$

where the inequality occurs because $(1-(1-p)^{k_1}) \geq p$ when $k_1 \geq 1$. ■

Lemmas 2 and 3 provide the necessary preparation for us to consider the failure spread via stepstones. Next, we present a new theorem that improves the bound on r and p further.

Theorem 5: For constant impact radius r and connection probability p, $p_\infty(\lambda, r, p) > 0$ if $pr^2 > \frac{1.8889}{\lambda}$.

Proof: We now consider two possible ways for failure to spread from node v_i to cell c_j: (i) direct connections from v_i to some nodes in c_j, and (ii) connections from v_i to stepstones in c_i and connections from these stepstones to some nodes in c_j. In the first case, we know from our proof of Theorem 4 that the probability of existing direct connections is $P_1 = 1 - e^{-p\lambda s}$. With the complementary probability $P_1^c = e^{-p\lambda s}$, there are no direct connections, in which we consider the stepstones.

As nodes are Poisson distributed with density λ, the number of nodes neighboring with v_i inside cell c_i is also Poisson with density λ. The appearance of v_i does not change the distribution of its neighbors. For each of these neighbors, v_i has a chance of p to connect to it independently from other neighbors. Therefore, the number of stepstones that v_i has is a Poisson random variable with density $p\lambda$. We next consider the node distribution in cell c_j. By Lemma 2, under the condition that none of the nodes in c_j is connected to v_i, the number of nodes in c_j is a Poisson random variable with density $(1-p)\lambda$. Letting $\lambda_i = p\lambda$ and $\lambda_j = (1-p)\lambda$, by Lemma 3, we obtain the probability of existing connections via stepstones when there are no direct connections from node v_i to cell c_j as $P_2 \geq (1 - e^{-p\lambda s})(1 - e^{-p(1-p)\lambda s})$.

Combining both cases and also considering the possibility of existing other failed nodes in c_i besides v_i, we bound the probability of failure propagation from c_i to c_j as

$$
\begin{aligned}
p^{\text{bond}} &\geq P_1 + P_1^c P_2 \\
&\geq 1 - e^{-p\lambda s} + e^{-p\lambda s}(1 - e^{-p\lambda s})(1 - e^{-p(1-p)\lambda s}) \\
&= (1 - e^{-p\lambda s})(1 + e^{-p\lambda s} - e^{-p(2-p)\lambda s}) \qquad (13)
\end{aligned}
$$

By defining $y = p\lambda s$ and noting $\lambda s = \frac{3\sqrt{3}}{26}\lambda r^2 > \frac{3\sqrt{3}}{26} \times 3.372 = 0.6739$, we rewrite Eq. (13) as

$$
\begin{aligned}
p^{\text{bond}} &\geq (1 - e^{-y})(1 + e^{-y} - e^{-y(2-\frac{y}{\lambda s})}) \\
&> (1 - e^{-y})(1 + e^{-y} - e^{-y(2-\frac{y}{0.6739})}). \qquad (14)
\end{aligned}
$$

Now we need to determine the condition on y such that

$$(1 - e^{-y})(1 + e^{-y} - e^{-y(2-\frac{y}{0.6739})}) > p_c^{\Delta} = 0.3473. \quad (15)$$

There is no closed-form solution for Eq. (15). Using numerical computation, we find that Eq. (15) holds when $y > 0.3775$. Since $y = p\lambda s$ and $s = \frac{3\sqrt{3}}{26} r^2$, we arrive at

$$\frac{3\sqrt{3}}{26} p\lambda r^2 > 0.3775, \qquad (16)$$

and finally

$$pr^2 > \frac{1.8889}{\lambda}, \qquad (17)$$

under which $\Pr[|C_{\text{cell}}(v_0)| = \infty] > 0$ and $p_\infty(\lambda, r, p) > 0$. ■

So far, Theorem 5 is our tightest result on the sufficient condition on r and p for failure percolation, where r and p are constants. Further improvement may be obtained by including the cases of multihop stepstones and multiple failed nodes in c_i other than v_i. However, an accurate modeling on the distribution of these nodes turns out to be extremely difficult. We will consider any further improvement in our future work.

B. Sufficient Condition of Non-Percolation: Constant r and p

Our results on the sufficient condition of percolation have quantified r and p that are large enough for an initial failure to spread to the entire network. Similarly, we are also interested in determination of r and p when failure cannot percolate in the network. From the continuous-to-discrete percolation mapping, we know that $|C(v_0)| < \infty$ if $|C_{\text{cell}}(v_0)| < \infty$. Thus, our task is to find the condition for $|C_{\text{cell}}(v_0)| < \infty$.

We still divide the network into many hexagonal cells. Compared to square cells, failure non-percolation is easier to study with hexagonal cells. When we choose the cell size sufficiently large, the failed nodes in one cell can only connect to the six neighboring cells and the connections are symmetric in probability. We can focus on the connections between any two neighboring cells to understand failure spreading. With square cells, however, a cell has eight neighbors: four horizontal or vertical ones and four diagonal ones. Its connections to the horizontal or vertical neighbors are not symmetric to those to the diagonal neighbors, rendering failure percolation difficult to analyze. By mapping into hexagonal cells, we have the following theorem to characterize the sufficient condition on r and p for failure non-percolation.

Theorem 6: For constant impact radius r and connection probability p, $p_\infty(\lambda, r, p) = 0$ if $pr^2 < \frac{-\ln(1-\frac{0.1642}{\lambda})}{2.5981\lambda}$.

Proof: We choose the cell radius $l = r$ such that a failed node in a cell can only connect to the six neighboring cells. In this way, the connections correspond to the triangular grid illustrated in Fig. 5(b) and there are no jumping connections, i.e., a failed node cannot connect to non-neighboring cells in one hop. With $l = r$, the area of each cell is $s = \frac{3\sqrt{3}}{2} r^2$.

We examine the failure propagation from cell c_i to a neighboring cell c_j. Since there exist failed nodes in c_i, i.e., c_i is not empty, the number of nodes k_1 in c_i is a random variable slightly different from Poisson, given by $P(k_1) = \frac{e^{-\lambda s}(\lambda s)^{k_1}}{(1-e^{-\lambda s})k_1!}$ ($k_1 \geq 1$). The number of nodes k_2 in c_j is still Poisson with

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2010 proceedings
This paper was presented as part of the main Technical Program at IEEE INFOCOM 2010.

density λ. Because the failed nodes are a subset of k_1 and a failed node may not be able to contact every node in k_2 due to the limit of r, we bound the probability of failure propagation from c_i to c_j as

$$\begin{aligned} p^{\text{bond}} &\leq \sum_{k_1=1}^{\infty}\sum_{k_2=0}^{\infty} \frac{e^{-\lambda s}(\lambda s)^{k_1}}{(1-e^{-\lambda s})k_1!}\frac{e^{-\lambda s}(\lambda s)^{k_2}}{k_2!}(1-(1-p)^{k_1k_2}) \\ &= 1-\sum_{k_1=1}^{\infty}\frac{e^{-\lambda s}(\lambda s)^{k_1}}{(1-e^{-\lambda s})k_1!}e^{-\lambda s(1-(1-p)^{k_1})} \\ &\leq 1-\sum_{k_1=1}^{\infty}\frac{e^{-\lambda s}(\lambda s)^{k_1}}{(1-e^{-\lambda s})k_1!}e^{-k_1p\lambda s} \\ &= 1-\frac{e^{-\lambda s}}{(1-e^{-\lambda s})}(e^{\lambda s e^{-p\lambda s}}-1) \\ &= \frac{1-e^{-\lambda s(1-e^{-p\lambda s})}}{1-e^{-\lambda s}}. \end{aligned} \tag{18}$$

Noting that $\lambda s = \frac{3\sqrt{3}}{2}\lambda r^2 > \frac{3\sqrt{3}}{2} \times 3.372 = 8.7607$ and $s = \frac{3\sqrt{3}}{2}r^2 \leq \frac{3\sqrt{3}}{2} = 2.5981$, we rewrite Eq. (18) as

$$p^{\text{bond}} < \frac{1-e^{-2.5981\lambda(1-e^{-p\lambda s})}}{1-e^{-8.7607}}. \tag{19}$$

By solving

$$\frac{1-e^{-2.5981\lambda(1-e^{-p\lambda s})}}{1-e^{-8.7607}} < p_c^{\Delta} = 0.3473, \tag{20}$$

we finally obtain

$$pr^2 < \frac{-\ln(1-\frac{0.1642}{\lambda})}{2.5981\lambda} \tag{21}$$

for $\Pr[|C_{\text{cell}}(v_0)| = \infty] = 0$ and $p_\infty(\lambda, r, p) = 0$. ■

As λ increases, $\lim_{\lambda\to\infty}\frac{-\ln(1-\frac{0.1642}{\lambda})}{2.5981\lambda} = \frac{0.0632}{\lambda^2}$. Thus, Theorem 6 also indicates that failure does not spread to the entire network if $pr^2 = o(\frac{1}{\lambda^2})$. In addition, Theorem 5 implies that failure percolates if $pr^2 = \omega(\frac{1}{\lambda})$. Combination of them characterizes the asymptotic regimes of r and p for failure percolation and non-percolation. Currently, the percolation property in the regime $pr^2 = \omega(\frac{1}{\lambda^2})$ and $pr^2 = o(\frac{1}{\lambda})$ is unknown yet. In order to improve the result of Theorem 6, we need to model accurately the distribution of the failed nodes in cell c_i and all their possible connections to cell c_j, which is the same challenge for improving the result of Theorem 5. We will consider any further improvement in our future work.

V. Percolation of Persistent Failures: General Failure Correlations

When the failure impact radius r is random and the failure connection function $g(x)$ is in a general form, modeling of failure spread becomes even complicated. We present two theorems in this section that characterize the general failure correlations such that an initial failure will and will not percolate in the network. To facilitate our study, we first combine the functions $f_r(x)$ and $g(x)$ into a unified failure connection function, as stated in the lemma below.

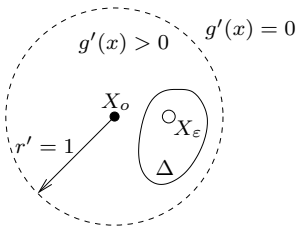

Fig. 8. Function $\Psi(g'(x), X_o, \Delta) \triangleq \int_\Delta g'(\|X_o - X_\varepsilon\|)\mathrm{d}\varepsilon$.

Lemma 4: For random failure impact radius r with probability density function $f_r(x)$ $(0 \leq x \leq 1)$ and general failure connection function $g(x)$, the failure correlations can be modeled equivalently by a constant $r' = 1$ and a unified connection function $g'(x) = g(x)\int_x^1 f_r(z)\mathrm{d}z$ if $0 \leq x \leq 1$ and $g'(x) = 0$ if $x > 1$.

Proof: When r is random and $g(x)$ is general, the connection between a failed node v_i and another node v_j is open with probability $g(x)$ conditional on $r_i \geq x$ and closed otherwise, where $x = \|X_i - X_j\|$ and r_i is the impact radius of v_i. Hence, the unconditional probability of an open connection is $g'(x) = g(x)\int_x^1 f_r(z)\mathrm{d}z$ if $0 \leq x \leq 1$. If $x > 1$, the connection is closed because $r_i \leq 1 < x$, so $g'(x) = 0$. Since $g'(x)$ has considered both $f_r(x)$ and $g(x)$, we can always assume the failure impact radius as a constant $r' = 1$. ■

Given the equivalence in modeling failure correlations, we consider r' and $g'(x)$ instead of $f_r(x)$ and $g(x)$ in the rest of this section. For clearer presentation, we now define a few concepts that will be used in our following theorems. We define function

$$\Psi(g'(x), X_o, \Delta) \triangleq \int_\Delta g'(\|X_o - X_\varepsilon\|)\mathrm{d}\varepsilon, \tag{22}$$

which is the integration of the probability $g'(x)$ over region Δ with respect to location X_o, as illustrated in Fig. 8. In addition, with respect to all the possible X_o and Δ ($|\Delta| = \sigma$), we define $\Psi_{\min}(g'(x), \sigma) = \min_{\{X_o, |\Delta|=\sigma\}}\{\Psi(g'(x), X_o, \Delta)\}$ and $\Psi_{\max}(g'(x), \sigma) = \max_{\{X_o, |\Delta|=\sigma\}}\{\Psi(g'(x), X_o, \Delta)\}$, which denote the minimum and maximum of $\Psi(g'(x), X_o, \Delta)$ respectively when we change X_o and Δ arbitrarily, as long as the area of Δ is kept constantly as σ. With the help of these definitions, we present our results regarding failure percolation with general $f_r(x)$ and $g(x)$ as follows.

A. Sufficient Condition of Percolation: $f_r(x)$ and $g(x)$

Theorem 7: For general failure correlations $f_r(x)$ and $g(x)$, $p_\infty(\lambda, f_r(\cdot), g(\cdot)) > 0$ if $\Psi_{\min}(g'(x), 0.1999) > \frac{0.4266}{\lambda}$.

Proof: We follow the same line of reasoning as in the proof of Theorem 4. Substituting $r' = 1$ for r, we obtain the area of each cell as $s = \frac{3\sqrt{3}}{26} = 0.1999$. Given a failed node v_i in cell c_i, because each node v_j in cell c_j is connected to v_i independently with probability $g'(\|X_i - X_j\|)$, the number of open connections is a Poisson random variable with parameter $\lambda\Psi(g'(x), X_i, c_j)$. Since there may exist other failed nodes in c_i besides v_i that also connect to c_j, we bound the probability

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2010 proceedings
This paper was presented as part of the main Technical Program at IEEE INFOCOM 2010.

of failure propagation from c_i to c_j as

$$\begin{aligned} p^{\text{bond}} &\geq 1 - e^{-\lambda\Psi(g'(x),X_i,c_j)} \\ &\geq 1 - e^{-\lambda\Psi_{\min}(g'(x),0.1999)}. \end{aligned} \quad (23)$$

By solving $1-e^{-\lambda\Psi_{\min}(g'(x),0.1999)} > p_c^{\Delta} = 0.3473$, we obtain

$$\Psi_{\min}(g'(x), 0.1999) > \frac{0.4266}{\lambda} \quad (24)$$

for $\Pr[|C_{\text{cell}}(v_0)| = \infty] > 0$ and $p_\infty(\lambda, f_r(\cdot), g(\cdot)) > 0$. ■

B. Sufficient Condition of Non-Percolation: $f_r(x)$ and $g(x)$

Theorem 8: For general failure correlations $f_r(x)$ and $g(x)$, $p_\infty(\lambda,f_r(\cdot),g(\cdot))=0$ if $\Psi_{\max}(g'(x), 1.2535) < \frac{-\ln(1-\frac{0.1642}{\lambda})}{\lambda}$.

Proof: We follow the same line of reasoning as in the proof of Theorem 6. Substituting $r' = 1$ for r, we obtain the area of each cell as $s = \frac{3\sqrt{3}}{2} = 2.5981$. Given $r' = 1$, a failed node v_i in cell c_i cannot reach all the locations in c_j, but only a fraction of area 1.2535 at maximum. (We skip this geometric computation due to space limitation.) When there exist k_1 number of nodes in c_i, denoted as $\{v_{i_1}, \cdots, v_{i_{k_1}}\}$, because the failed nodes are a subset of $\{v_{i_1}, \cdots, v_{i_{k_1}}\}$ and each failed node propagates failure to a node v_j in cell c_j independently, v_j fails after $\{v_{i_1}, \cdots, v_{i_{k_1}}\}$ with a probability at most $1-\prod_{m=1}^{k_1}(1-g'(\|X_{i_m}-X_j\|))$. Hence, the number of nodes that fail in c_j after $\{v_{i_1}, \cdots, v_{i_{k_1}}\}$ is a Poisson random variable with parameter Λ that satisfies

$$\begin{aligned} \Lambda &\leq \lambda \int_{c_j} 1 - \prod_{m=1}^{k_1} (1 - g'(\|X_{i_m} - X_\varepsilon\|))\mathrm{d}\varepsilon \\ &\leq \lambda \int_{c_j} \sum_{m=1}^{k_1} g'(\|X_{i_m} - X_\varepsilon\|)\mathrm{d}\varepsilon \\ &= \lambda \sum_{m=1}^{k_1} \Psi(g'(x), X_{i_m}, c_j) \\ &\leq k_1\lambda\Psi_{\max}(g'(x), 1.2535). \end{aligned} \quad (25)$$

Similar to Eq. (18), by considering the distribution of k_1,

$$\begin{aligned} p^{\text{bond}} &= \sum_{k_1=1}^{\infty} \frac{e^{-\lambda s}(\lambda s)^{k_1}}{(1-e^{-\lambda s})k_1!}(1 - e^{-\Lambda}) \\ &\leq 1 - \sum_{k_1=1}^{\infty} \frac{e^{-\lambda s}(\lambda s)^{k_1}}{(1-e^{-\lambda s})k_1!} e^{-k_1\lambda\Psi_{\max}(g'(x),1.2535)} \\ &= \frac{1 - e^{-\lambda s(1-e^{-\lambda\Psi_{\max}(g'(x),1.2535)})}}{1 - e^{-\lambda s}}. \end{aligned} \quad (26)$$

Furthermore, by following Eq. (19), (20) and (21), we obtain

$$\Psi_{\max}(g'(x), 1.2535) < \frac{-\ln(1 - \frac{0.1642}{\lambda})}{\lambda} \quad (27)$$

for $\Pr[|C_{\text{cell}}(v_0)| = \infty] = 0$ and $p_\infty(\lambda, f_r(\cdot), g(\cdot)) = 0$. ■

VI. Conclusion

In this paper, we have studied the spread of correlated failures in large wireless networks. When there exist causal relations, a single failure may initiate a cascading sequence of related failures that severely impair the correct network functioning. In order to understand the occurrence of large-scale failure spreading, we model the contagiousness of various types of failures into two generic correlation functions and determine their characteristic regimes in which failures will and will not spread to the entire network. Our results provide a quantified answer on the relation between failure correlations and failure pervasiveness. They match the intuition that failures spread to a large area when they have strong correlations. By using a generalized model, our results are applicable to the understanding of a wide range of failure scenarios. The work in this paper helps network designers and operators evaluate the resilience of large wireless networks to correlated failures.

References

[1] Z. Yu, Y. Wei, B. Ramkumar, and Y. Guan, "An Efficient Signature-Based Scheme for Securing Network Coding against Pollution Attacks", *Proc. of IEEE INFOCOM*, 2008.
[2] S. Stefanakos, "Reliable Routings in Networks with Generalized Link Failure Events", *IEEE/ACM Transactions on Networking*, vol. 16, no. 6, pp. 1331–1339, December 2008.
[3] B. Awerbuch, D. Holmer, C. Nita-Rotaru, and H. Rubens, "An On-Demand Secure Routing Protocol Resilient to Byzantine Failures", *Proc. of ACM WiSe*, 2002.
[4] T. Bu, A. Chen, S. Wiel, and T. Woo, "Design and Evaluation of a Fast and Robust Worm Detection Algorithm", *Proc. of IEEE INFOCOM*, 2006.
[5] V. Bhandari and N.H. Vaidya, "Reliable Broadcast in Wireless Networks with Probabilistic Failures", *Proc. of IEEE INFOCOM*, 2007.
[6] Y. Kim, W.C. Lau, M.C. Chuah, and H.J. Chao, "Packetscore: Statistics-Based Overload Control against Distributed Denial-of-Service Attacks", *Proc. of IEEE INFOCOM*, 2004.
[7] S. Yu and Y. Zhang, "R-Sentry: Providing Continuous Sensor Services against Random Node Failures", *Proc. of IEEE DSN*, 2007.
[8] S. Nelakuditi, S. Lee, Y. Yu, Z.L. Zhang, and C.N. Chuah, "Fast Local Rerouting for Handling Transient Link Failures", *IEEE/ACM Transactions on Networking*, vol. 15, no. 2, pp. 359–372, April 2007.
[9] R.S. Bhatia, M. Kodialam, T.V. Lakshman, and S. Sengupta, "Bandwidth Guaranteed Routing with Fast Restoration against Link and Node Failures", *IEEE/ACM Transactions on Networking*, vol. 16, no. 6, pp. 1321–1330, December 2008.
[10] P.R. Jelenkovic and J. Tan, "Dynamic Packet Fragmentation for Wireless Channels with Failures", *Proc. of ACM MobiHoc*, 2008.
[11] S. Jain, M. Demmer, R. Patra, and K. Fall, "Using Redundancy to Cope with Failures in a Delay Tolerant Network", *Proc. of ACM SIGCOMM*, 2005.
[12] F. Xing and W. Wang, "On the Critical Phase Transition Time of Wireless Multi-hop Networks with Random Failures", *Proc. of ACM MobiCom*, 2008.
[13] Z. Kong and E.M. Yeh, "Distributed Energy Management Algorithm for Large-Scale Wireless Sensor Networks", *Proc. of ACM MobiHoc*, 2007.
[14] R. Meester and R. Roy, *Continuum Percolation*, Cambridge University Press, 1996.
[15] G. Grimmett, *Percolation*, 2nd ed., Springer, 1999.
[16] M. Penrose, *Random Geometric Graphs*, Oxford University Press, 2003.
[17] Z. Kong and E.M. Yeh, "Connectivity and Latency in Large-Scale Wireless Networks with Unreliable Links", *Proc. of IEEE INFOCOM*, 2008.
[18] Z. Kong and E.M. Yeh, "On the Latency for Information Dissemination in Mobile Wireless Networks", *Proc. of ACM MobiHoc*, 2008.
[19] O. Dousse, P. Mannersalo, and P. Thiran, "Latency of Wireless Sensor Networks with Uncoordinated Power Saving Mechanisms", *Proc. of ACM MobiHoc*, 2004.