

Using Realistic Factors to Simulate Catastrophic Congestion Events in a Network

Christopher Dabrowski and Kevin Mills
{cdabrowski, kmills}@nist.gov

Abstract:

With the rapid growth of the Internet, there has been increased interest in the use of computer models to study the dynamics of communication networks in the research literature. An important example of this has been the study of dramatic, but relatively infrequent, events that result in abrupt, and often catastrophic, failures in the network due to congestion. These events are sometimes known as phase transitions. With few exceptions, the models of such computer communications networks used in previous studies have been abstract graphs that include simplified models of such important network factors as router speeds and congestion control procedures. Here, we modify this typical approach, adding realistic network factors which can be varied, including router classes, variable router speeds, flows, the transmission control protocol (TCP), sources and receivers, and packet dropping to a graph model of a single Internet Service Provider (ISP) network that can have more than a quarter million members. For each valid configuration of realistic factors, congestion is then gradually increased. While there are realistic network models reported in the literature, to our knowledge none of these have been used to study catastrophic failures in computer networks. We show that the addition of realistic network factors to our model of an ISP network can mitigate catastrophic events. With the addition of variable router speeds or TCP, a phase transition to a congested state, where all routers are congested in a single ISP network is not observed. Yet, as congestion spreads, ultimately the operation of the ISP network appears to decline, along with the ability of its nodes to communicate. The results of this study should be cautionary for other domains, such as electrical power grids, and the spread of viruses or diseases, where abstract graph models are often used to study phase transitions.

Keywords: networks; percolation; phase transition; Internet; simulation.

1. Introduction

With the rapid growth of the Internet, there has been increased interest in the use of models, based on graphs, to study the dynamics of computer communication networks in the research literature. An important example of modeling dynamics has been the study of dramatic, but relatively infrequent, events that result in abrupt and often catastrophic failures in a graph model of a network. These events are sometimes referred to as phase transitions or as resembling phase transitions [Solé and Valverde, 2001, Echenique, Gomez-Gardenes, and Moreno, 2005, Wu, Wang and Yeung, 2008; Sarkar et al., 2012]. In these models, a network model may go from a state in which communications flow freely to a state where the network is severely degraded and effectively ceases to operate. Often, these failures are due to congestion in the form of an increased number of packets¹ in a system, though other factors, most notably computer viruses, may come into play. With few exceptions, the models used to study the spread of congestion and catastrophic failures in communications networks reported in previous papers have been abstract graphs using simplified network factors for, for example, simplified forms of routing and congestion control (Cohen et al., 2000; Solé and Valverde, 2001; Woolf et al., 2004; Arrowsmith et al., 2004; Lawniczak et al., 2007; Wu, Wang and Yeung, 2008; de Martino et al., 2009; WangD et al., 2009; and WangW et al., 2009). Few papers have explored the effects of making such factors more realistic.

This paper is motivated by the abstract nature of the network models used in such studies in which catastrophic failures occur and in which evidence of phase transitions is observed. In this paper, the effects of congestion, and its spread, are studied over a graph model of a single Internet Service Provider (ISP) network, where our ISP model contains 218 routers, which can be expanded to include source and receiver sites for a total of 258 158 sites. In our model, we simulate configurations of both (a) abstract communication networks with little realism and (b) increasingly realistic networks, in which six realistic network factors are gradually added. We then compare and analyze the results for combinations of realistic factors. In this paper, the effects of congestion, and its spread, are studied over a graph model of a single Internet Service Provider (ISP) network. In our model, we simulate configurations of both (a) abstract communication networks with little realism, and (b) increasingly realistic networks. We then compare and analyze the results. The contributions of our paper are as follows. First, we show that realistic factors can be grouped coherently and added to a graph model of a communications network as it collapses. Second, we show that catastrophic congestion events or observation of phase transitions, in which *all* routers congest, occur only in more abstract models of a network and might or might not occur in less abstract models, depending on which realistic network factors are included. Generally catastrophic failures or phase transitions, in which all routers congest, can occur only if the graph model of a network has few realistic factors. We study this collapse in terms of a *percolation state*² [ben-Avraham and Havlin, 2000], seen in phase transitions of graphs. Closely related to the second contribution is a third, where we show that the presence of two realistic factors—variable router speeds and the transmission control protocol (TCP)—mitigates catastrophic congestion and phase transitions. We find that the spread of congestion, which leads to catastrophic failures on the Internet,

¹ A packet is a well-known unit of information in a network.

² A percolated state occurs only in an infinite graph, as is explained below.

should be modeled using realistic factors. This is because congestion behaves differently in an abstract network model than in a network modeled using realistic factors. Conclusions about catastrophic congestion on the Internet should not be drawn from simpler, more abstract models that do not use sufficient realism. Further, the need to use realistic factors should be considered in computer models of other types of networks, such as electrical power grids [Carreras et al., 2002], or models of virus and disease spread [Pastor-Satorras and Vespignani, 2001; Moreno, Pastor-Satorras and Vespignani, 2002].

Here, we incorporated six realistic factors into our model of a computer communications network (described more fully below): (1) *classes* of sites, including router classes for core routers, point-of-presence (PoP) routers, and access routers, (2) *variable speeds* that differentiate between speeds of fast core routers and slower PoP and access routers, (3) *sources and receivers*, which have speeds when present, placed under access routers, (4) *flows*, in which packets are organized into correlated streams, which represent the transmission of a piece of information, such as a Web page, from a source to a receiver, (5) TCP, which detects and adapts to network congestion, (6) and *packet dropping*, which occurs when the number of packets exceeds a finite-buffer size within routers. Of these, we find that two, variable router speeds and TCP, are essential in regulating the spread of congestion and preventing catastrophic collapse. The presence of packet dropping affects the number and proportion of packets reaching their intended destinations.

To summarize what follows below, we describe and present simulations for 18 valid combinations of six realistic network factors. For each combination, congestion is gradually elevated by increasing the number of packets injected into the system at each time step in an attempt to trigger a congestion collapse within our model. With the addition of variable router speeds or TCP in 10 of the 18 configurations, a phase transition to a congested failed state is never observed. That is, we never observe a state in which all routers are congested. However, in eight cases, which have neither variable speeds nor TCP, all routers become congested and therefore any sources and receivers attached below cannot communicate with anyone in the network. This is effectively evidence of a percolated state [ben-Avraham and Havlin, 2000], in which, in our case, all routers of the ISP network are occupied by congestion. Including more realistic network factors leads generally to lower congestion, and all routers are not occupied. In more abstract configurations with fewer realistic factors, congestion spreads completely, and all routers are occupied. These results are consistent with important questions about the necessity and use of realistic network factors in modeling a communications network, which were raised previously by Alderson and Willinger [Alderson and Willinger, 2005; Willinger et al., 2002]. Alderson and Willinger focused more on the use of realistic topology, which we have incorporated as well. In our experiments, as congestion spreads, ultimately the operation of the network, and the ability of its member nodes to communicate declines, even when a total congestion collapse is not observed. That is, part of the network will congest even if the entire network does not. In this case, sources and receivers cannot communicate if they are attached to a congested router. These portions can be readily identified and constitute potentially vulnerable parts of the network. In this paper, section 2 discusses previous work, section 3 covers definitions and the experiment plan, section 4 present results and section 5 concludes. A companion technical report [Dabrowski and Mills, 2015] provides further details.

2. Previous Work

In the realm of computer communication systems, a number of researchers have studied how increasing a quantity, such as load, causes a phase transition from a network-wide operational state to a catastrophic congested state. Researchers have developed simulation models, in which they studied the effects of increasing load (i.e., the number of packets in a system) as load approached a critical point, after which increased load altered the behavior of the system [Solé and Valverde, 2001; Arrowsmith et al., 2004; Echenique, Gomez-Gardenes and Moreno, 2005; Lawniczak et al., 2007; Mukherjee and Manna, 2005; Tadic, Rodgers, and Thurner, 2007; Lei and Yu, 2008; Wu, Wang and Yeung, 2008; WangD et al., 2009; WangW et al., 2009, Rykalova, Levitan, and Brower, 2010; Sarkar et al., 2012]. The approach in most of these studies was primarily empirical, relying strongly on observation of simulations and models, which were most often abstract and simplified. For instance, a square lattice topology, in which sites acted as both hosts and routers, was used by [Lawniczak et al., 2007], while [Solé and Valverde, 2001; Arrowsmith et al., 2004; Mukherjee and Manna, 2005] modeled a two-dimensional lattice but differentiated host and router sites (with Arrowsmith et al., 2004 also studying triangular and hexagonal lattices). Other studies focused on *scale-free*³ networks, in which hosts and routers not were differentiated [Echenique, Gomez-Gardenes and Moreno, 2005].

The routing algorithms used in these efforts appear to be motivated by Internet processes. For example, [Solé and Valverde, 2001; Arrowsmith et al., 2004] used shortest-path first routing (SPF) to determine which site to forward a packet to. If SPF could not be used, the packet was forwarded to the least congested site [Solé and Valverde, 2001] or forwarded along the least used link [Arrowsmith et al., 2004]. A combination of criteria (depending on the researchers) was also sometimes used to determine where to route packets in these models, including SPF, availability of buffer space at the destination site, and congestion awareness [Echenique, Gomez-Gardenes and Moreno, 2005; Lawniczak et al., 2007; WangD et al., 2009]. Some researchers used routing procedures based, at least in part, on randomly determining where to forward a packet [Mukherjee and Manna, 2005; Tadic, Rodgers, and Thurner, 2007; Wu, Wang and Yeung, 2008; Sarkar et al., 2012]. Infinite-sized-buffers were assumed, except for example [Woolf et al., 2004; Tadic, Rodgers, and Thurner, 2007; Wu, Wang, and Yeung, 2008, Wang et al., 2009] where finite-buffers sizes with packet dropping procedures for overflow were assumed. In [Wu, Wang, and Yeung, 2008] models were varied to use both infinite and finite buffers.

In models used in many of these previous papers, the number of packets served as the control parameter. At higher packet-injection rates, packet forwarding was inhibited because queues formed at sites due to local congestion, leading to observable network-wide congestion. At some critical point, increased levels of congestion caused a change from an uncongested state in which packets regularly arrived at their destinations in a timely manner to a congested state, after which network throughput

³ A scale-free network is given by $P_k \sim k^{-\gamma}$, where P_k is the probability that a site has k links and γ is an exponent. In a scale-free network, $k^{-\gamma}$ is skewed so that only a few sites, or hubs, have many links incident upon them, while the vast majority of sites have far fewer links. See also [Dorogovtsev, Goltsev, and Mendes, 2008].

fell dramatically. In some of these studies, this change was likened to a phase transition and evidence of percolation was observed in the network.

In our study, a realistic topology is adapted from a single ISP. In addition, our routing algorithms, variable router speeds (VS), sources and receivers (SR), and other network factors are more realistic than in previous studies. Among the more realistic topologies in previous studies, is that used by Echenique, Gomez-Gardenes and Moreno [Echenique, Gomez-Gardenes and Moreno, 2005], which featured a 11 174 router snapshot taken from the Internet Autonomous System (AS) topology circa 2001, collected by the Oregon Router Views Project. Using the model of Echenique, Gomez-Gardenes and Moreno as a basis, we repeated their experiments and got similar results. This is discussed in [Dabrowski and Mills, 2015]. In the AS topology used by Echenique, Gomez-Gardenes and Moreno, the probability of a router having k links was given by $P_k \sim k^{-\gamma}$, where $\gamma = 2.2$. The distribution of router links in this topology was considered scale-free, in that there was a low probability of a router having many links, i.e., there were few hub routers. While the topology was large and taken from the real-world, the operation of each router had two variants: (1) either SPF was used (a packet used the fewest number of routers or “hops” in going from a sending to destination router), or (2) the level of congestion of a router’s neighbors was considered along with SPF. Each packet was first injected at a randomly selected router, which was to forward it to a randomly selected destination, at each time step. The rate of packet injection or p was increased. The measure of congestion in the network was given by ρ , computed by

$$\rho = \lim_{t \rightarrow \infty} \frac{A(t + \tau) - A(t)}{\tau p}, \quad (1)$$

where A represented the number of packets in the system, p is as before, and τ was the observation time. There were no classes, variable router speeds, flows, TCP, packet dropping, or network delays, as opposed to the real Internet, where these factors exist. However, a process similar to percolation (defined below) was observed, using both variants of router operation as seen in Figure 1 below:

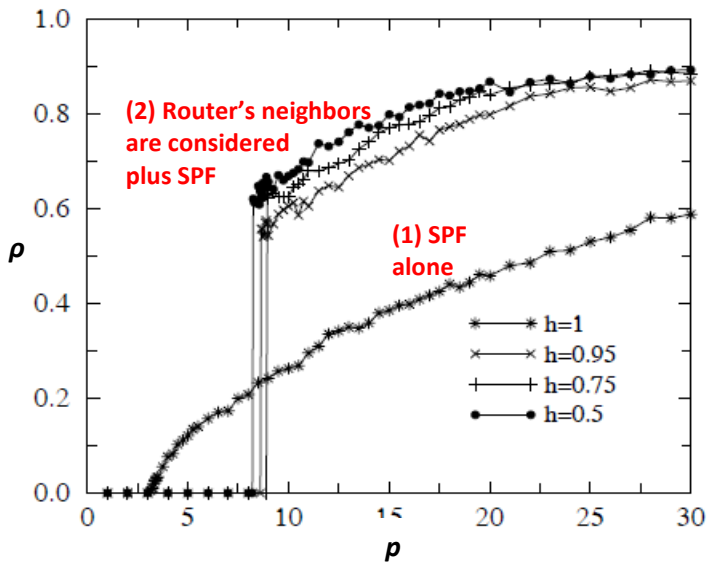


Figure 1. Response of ρ as p is increased, showing events which resemble phase transitions at p_c when variant (1) only SPF is used in lower curve, and variant (2) the level of congestion of a router's neighbors was considered along with SPF in upper three curves. Variant (2) is given by $\delta_i = h d_i + (1 - h) c_i$, where h is a traffic awareness parameter, whose value $0 \dots 1$, i is the index of a nodes neighbor, d_i is the minimum number of hops through neighbor i , and c_i is the queue length (or buffer size) of i . Note that when $h = 1$, only SPF is considered.

In Figure 1, the change in ρ shows that a process similar to a gradual phase transition to a congested state occurred when SPF alone was used with $p_c = 3$ (see variant (1)), but a steeper phase transition occurred when congested routers were avoided and $p_c = 8$ (see variant (2)). Some have likened the difference in this steepness to the difference in first- and second-order phase transitions. However, we find the greater steepness in variant (2) is due to space in uncongested router buffers being suddenly exhausted at the same time resulting also in a higher p_c . The work of Echenique, Gomez-Gardenes and Moreno [Echenique, Gomez-Gardenes and Moreno, 2005], was the baseline for our work, into which we introduced realistic network factors.

3. Definitions and Experiment Plan

This section consists of four subsections, which (a) give critical definitions and discuss (b) our experiment topology, (c) the realistic network factors that we used, and (d) how we combined these realistic factors into network configurations.

a. Percolation Definitions.

As used in this paper, percolation is the process by which some property of interest spreads through an infinite graph until each member of a graph is occupied, or possesses the property. In this paper, the ISP model we study can be considered a graph, where a *graph* $G = (N, L)$ consists of two sets N and L . The elements of $N \equiv \{n_1, n_2, \dots, n_N\}$ are the *nodes* (or *vertices*) of the graph G , while the elements of $L \equiv \{l_1, l_2, \dots, l_k\}$ are its *links* (or *edges*), where each l_i consists of a pair of elements of N . Percolation processes have been studied in graphs in [Barabási and Albert, 1999; Bollobás, 1984; Callaway et al., 2000; and Cohen et al., 2000]. A node or vertex is considered a site, where a site is any router, source, or receiver.

In site percolation, a spreading property progressively occupies all sites within an infinite network [ben Avraham and Havlin, 2000]. The property of interest, which spreads through our ISP model or graph, is congestion. We define a *congested router* if the count $Q_{i,d}$ of packets waiting for transmission in direction d (*up*, or *down*) exceeds Q_T , or $Q_{i,d} \geq Q_T$. (Note that core routers have only one queue ($d = \text{only}$), while access and PoP routers have two queues ($d = \text{down and up}$). In this experiment, $Q_T = 70\%$ of the space in a buffer and its related router queue. When $Q_{i,d} \geq Q_T$, a router is said to be *occupied*. In our model, sources and receivers (SR) have no queues. In the ISP topology we study, there is evidence of percolation only if all routers are occupied with congestion. In our model, SR are attached to one access router, therefore they cannot communicate across the network if that router is congested. In our model, p will stand for the rate of packet injection.

If p increases so that it exceeds some critical probability, p_c , then all routers congest. In a percolation process, p_c is known as the *percolation threshold* or *critical point*. If $p > p_c$, a *percolation transition* occurs in an infinite network, and an infinite *Giant Connected Component (GCC)* emerges. The GCC is represented by P_∞ , and grows until $P_\infty = 1$. The symbol for infinity is present because theoretically the GCC would only form in an infinite-sized network [ben Avraham and Havlin, 2000]. In a random graph, further proofs about the size and growth of a GCC in relation to p_c can be found in [Bollobas, 1984]. In our model, which of course is finite, all routers congest in eight configurations, in which case their attached sources and receivers (SR), if any, cannot communicate across the network. Therefore, a GCC emerges at a p_c and continues to grow as more routers in the network are occupied by congestion. In our finite network configurations, percolation can only be observed when all routers are occupied; and a percolation transition may be *observed* when $p > p_c$. Thus in our results, p_c can be observed in eight configurations. For each on these configurations, we provide an estimate of p_c . Above p_c , the GCC forms a largest cluster of occupied sites. This process can also be inverted, with a second GCC of unoccupied (uncongested) sites declining, as congestion increases, until there are no sites within the GCC. As will be described in more detail the next section, this second GCC of uncongested (unoccupied) sites represents the breakdown of network connectivity. Here, the largest cluster of uncongested sites (i.e., the GCC) declines to zero.

b. Experimental Topology.

In our model, we substituted a 218-router topology encompassing one ISP network, obtained from an Internet service provider, for the 11 174 router topology taken from [Echenique, Gomez-Gardenes and Moreno, 2005]. However, we wanted to see first what would happen within this smaller network. The topology we used is depicted in Figure 2.

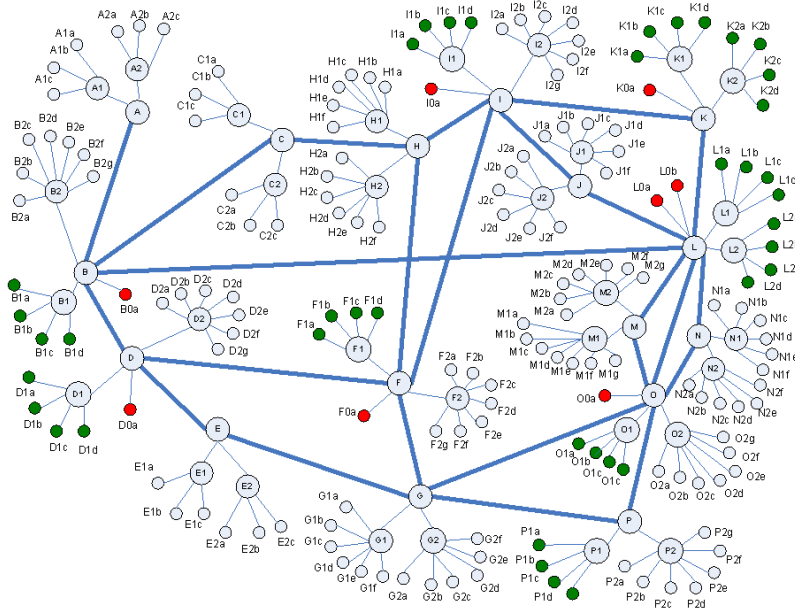


Figure 2. Three-Tier Topology with 16 Core Routers (A-P), 32 Point of Presence Routers (A1-P2) and 170 Access Routers (A1a-P2g) – 8 red and 40 green Access Routers. From [Mills et al., 2010].

Our 218 router single ISP topology has three hierarchical router tiers: (1) 16 core routers (A-P in Figure 2), (2) 32 point of presence (PoP) routers (A1-P2), and (3) 170 access routers (A1a-P2g). To model heterogeneity in network access, our topology has three different types of access routers: **D**-class (e.g., eight red nodes in Figure 2, which connect directly to core routers), **F**-class (e.g., 40 green nodes) and **N**-class (e.g., 122 small gray nodes). Classifying access routers enables different speeds to be assigned to each class. A fourth tier of sources and receivers may optionally be connected to the third tier, so that sources and receivers can be attached to access routers. If sources and receivers are added (as explained below), there are a total of 258 158 sites in a configuration. When access routers are classified, and there are no sources and receivers, packet injection takes place only at the access router sites. If the fourth tier is present, packet injection takes place only at sources and receiver sites and not at the routers. In both cases, packet injection takes place at the edges of the topology we study. The three tiers and all classes within them, as well as all sources and receivers, can be eliminated. Eliminating these allows for a uniform 218-router topology similar to that of the uniform topology of [Echenique, Gomez-Gardenes and Moreno, 2005]. In such cases, when routers are not classified, packet injection can occur at any router in the topology we study.

As with [Echenique, Gomez-Gardenes and Moreno, 2005], packets are injected into the model at randomly-chosen sites and are sent to randomly-chosen destinations. In our 218-router single ISP topology, each source determines whether it will inject a packet at the next time step on the basis of the formula

$$injectionChance = InjectionRateP / numberInjectionSites, \quad (2)$$

where $InjectionRateP$ is the current packet injection rate, given in packets per time step (ts). The number of sites where the packet can be injected is given by $numberInjectionSites$. When formula (2) is applied repeatedly at each time step (ts), the result is the time when a source will inject its next packet. Formula (2) is repeated at each site where packets can be injected until the number of ts in the run is exhausted. Packets flowing between a source-receiver pair follow a single ingress/egress path between an access router and a top-tier backbone router. When flows are present, formula (2) is altered so that the average flow size (in packets) is also in the denominator. This change allows computation of when the next packet in a flow will be injected.

c. Realistic network Factors.

In what follows, we describe six realistic factors derived from [Mills et al., 2010]. First is packet dropping (PD). When PD is present, incoming packets to a router whose buffer is full are simply dropped and never reach their destinations. (The size of the buffer is explained below.) If PD is not present, incoming packets are added to the end of a queue and not dropped. Of course with no PD present, queues can become quite long. However packets, once processed by a router, are forwarded to their destinations. Second are classes (CL). When CL is present, there are core, PoP, and access routers, with three sub-types of access routers as described above. If sources and receivers (SR) are also present, they are classified separately (see below) from the routers. With CL, packets are always injected only at network edges: at access routers when SR is not present and at SR otherwise. Also with CL, packets are required to transit at least one core router.

Third, are variable speeds (VS), as shown in Table 1 below, which was obtained from [Mills et al., 2010]. When VS is present, core routers, PoP routers, and different classes of access routers process packets at different speeds. Unlike real networks, where links have transmission speeds and associated buffers, transmission speeds are assigned directly to routers. Because packets have no size, router speeds are assigned in units of packets/time step (packets / ts). Here we scale down router speeds by $1 / 40$ from the speeds shown in Table 1. Therefore in our model, core routers process 80 packets / ts and PoP routers 10 packets / ts , **D**-class access routers process 10 packets / ts , **F**-class access routers process 2 packets / ts , while **N**-class routers process 1 packet / ts . If sources and receivers are also present together with VS, then 50 % of sources and receivers (randomly selected) operate at 2.0 packets / ts , while 50 % process 0.2 packets / ts . With or without VS, buffer sizes can be determined based on speeds of routers, using the formula $\text{ceil}(250 \times \text{routerProcessingSpeed})^4$. Without VS, routers, sources, and receivers (if any) operate at the same speed: 9.0 packets / ts , which is the average of all routers, computed when VS is present. As indicated in Table 1, our goal was to establish reasonable engineering relationships among the speeds of the various router classes.

⁴ In this formula, buffer size is calculated by a variant of the Bush-Meyer algorithm or $(250 \text{ } ts \times \text{router packet forwarding rate})$. Therefore with VS, core, PoP, and different classes of access routers will have different buffer sizes. If the router has two buffers instead of one, this speed is halved for each of the two buffers.

Table 1. Speed Relationships among Router Classes. From [Mills et al., 2010].

Parameter	Value	Speed	Relationships
s1	1600	Router Class	Speed
s2	4	Core	s1 X BBspeedup
s3	10	PoP	s1 / s2
BBspeedup	2	N-Class	s1 / s2 / s3
Bfast	2	F-Class	s1 / s2 / s3 X Bfast
Bdirect	10	D-Class	s1 / s2 / s3 X Bdirect

Without VS, all routers have the same buffer size—2250 packets, a size which is determined by 250 times the average speed, or $(250 \text{ } ts \times 9.0 \text{ packets} / ts)$. Each core router multiplexes packet forwarding from a single buffer, while PoP and access routers have two buffers each, one heading toward the core and one heading from the core. PoP and access routers alternate forwarding between each of the two buffers.

Fourth are sources and receivers (SR), which constitute an optional fourth tier. When present, they are distributed under the access routers. Sources equate to computers that have information that receivers wish to download. Bypassing some of the details contained in [Mills et al., 2010; Dabrowski and Mills, 2015], there are four times the number of receivers to the number of sources, with different probabilities that a source can be under an **N**-, **D**-, or **F**-class router. Briefly in our model, there are 297 sources under each **D**- and **F**-class router and 306 sources under each **N**-class router, while there are 1188 receivers under each **D**- and **F**-class router and 1224 receivers under each **N**-class router. This gives a total of 51 588 sources and 206 352 receivers, which adds up to 257 940 sources and receivers, creating an overall configuration with 258 158 sites when the 218 routers are included.

Fifth are flows (FL). Each flow consists of a number of packets and represents an email, Web page, or document, as in the actual Internet. Each source will periodically transfer a flow, after randomly selecting a receiver from under a parent core router that differs from the source’s parent core router. Since access routers of different classes have differing speeds, the locations of a source-receiver pair influence the characteristics of the path for each flow of packets. When FL is present, if sources and receivers (SR) are also present, a source first needs to establish a connection with a randomly selected receiver. Once the source receives a Connection *Ack* from the receiver within three tries, the source sends at average of 350 packets without acknowledgement or retransmission. The number of packets in a flow (if FL is present) is selected randomly

$$\text{Number of packets} = \text{ceil}(rv_pareto(\text{random number stream}, 350 \times (\text{Shape}-1) / \text{Shape}, \text{Shape})), \quad (3)$$

where the formula and Shape = 1.5 are obtained from [Mills et al., 2010], while the average number of packets in a flow (350) is chosen by us. This effectively means, the source operates as a UDP Source, and can potentially complete a flow without transmitting a packet that reaches its destination. However with FL, a source has to organize its sent packets within a correlated stream and cannot send packets independently. Each source and receiver is limited to one flow at a time and cannot start (or receive) a second flow until completing a flow that it has begun.

Finally, there is TCP, which in our model, can operate only when flows (FL) are present. Our TCP model is taken from [Mills et al., 2010]. TCP seeks to detect and adapt to congestion, where the rate at which packets are sent by a source is controlled by the size of a *congestion window*. This window is effectively the number of packets that can be sent by a source before receiving corresponding acknowledgements (or *acks*). An acknowledgement is sent by a receiver for every packet received from the source. At first, the size of the congestion window is only two packets, and TCP begins in an *initial slow start* phase. When *acks* are received by the source, the size of the congestion window is increased exponentially until an initial threshold of 100 packets is reached. Once this threshold is exceeded, the size is increased logarithmically to a second threshold ($2^{30} / 2$ packets). If no packet losses are encountered, this has the effect of speeding up packet transmission. After reaching the second threshold, the source enters *congestion avoidance*, and the size of the congestion window is increased linearly.

If, at any time during this process, a packet is received out-of-order by the receiver, a negative acknowledgement (*nack*) is sent to the source. Upon receiving a *nack*, a source decreases its congestion window by one half, and enters a congestion avoidance phase—in effect slowing down the transmission. The out-of-order (*nacked* packet) sent originally by the source must be retransmitted. A *timeout* is declared if a source fails to receive an *ack* or *nack* within a specified time. When a timeout is declared, congestion avoidance is again entered, the congestion window is cut down to two, and the threshold is cut by one half. In no case is the size of congestion window allowed to fall below two. The process repeats for each *nack* or *timeout*. Once *acks* are regularly received by the source, probably because congestion is lessened along the path, the source begins again to transmit packets more quickly. The above process is repeated, if a packet is received out-of-order or lost. This process of slowing down and speeding up in response to varying levels of congestion is called *TCP rate adaptation*. When repeated, the resulting behavior of rate adaptation exhibits a saw-tooth pattern of increasing and decreasing transmission rate on a flow. Verification that the model described here adheres to the model of [Mills et al., 2010] is documented and discussed in [Dabrowski and Mills, 2015].

d. Configurations: Combinations of Realistic network Factors.

In our experiments, realistic network factors were added individually or in combination to produce 18 valid configurations. Since there were six realistic factors, this yielded 2^6 (64) possible combinations, but all were not valid. Dependencies among factors (with their abbreviations) are shown in Table 2.

Table 2: Dependencies among Realistic network Factors.

Dependencies
Variable Speeds (VS), Sources and Receivers (SR), Flows (FL) and TCP require Classes (CL)
FL and TCP require SR
TCP requires FL

We eliminated any combination of realistic network factors that violate the dependencies. Violations identified as illogical combinations were (for example) TCP without flows (FL) or Variable (router) speeds without (router) classes. There were 46 such illogical combinations. The 18 valid combinations appear in Table 3. In this table for each valid combination, a decimal configuration is provided, followed by the

abbreviated factors present in the related configuration. Valid combinations represent realistic network factors which can reasonably be grouped together and therefore may not be in strict numeric order. In a related technical note [Dabrowski and Mills, 2015], there is an additional realistic network factor, which is omitted here—delays (DE). This brings the total number of possible configurations to 128^5 . So, we retain the (0-128) numbering of our configurations to maintain consistency with the technical notes.

Table 3: Valid Configurations Simulated. The table shows the 18 valid configurations in which packet dropping (PD), classes (CL), variable speeds (VS), sources and receivers (SR), flows (FL), and TCP are present.

Decimal Configuration Code	Factors in Valid Configuration
c0	
c1	PD
c2	CL
c3	PD + CL
c6	CL + VS
c7	PD + CL + VS
c18	CL + SR
c19	PD + CL + SR
c22	CL + VS + SR
c23	PD + CL + VS + SR
c50	CL + SR + FL
c51	PD + CL + SR + FL
c54	CL + VS + SR + FL
c55	PD + CL + VS + SR + FL
c114	CL + SR + FL + TCP
c115	PD + CL + SR + FL + TCP
c118	CL + VS + SR + FL + TCP
c119	PD + CL + VS + SR + FL + TCP

Each of the 18 configurations in Table 3 was run once, where a single run constituted 200 000 *ts*. For each run, p , the packet-injection rate, was advanced by 10, from one to 2500. In eight configurations, percolation was observed within 200 000 *ts*. When all router sites congested, evidence of percolation was declared in the configuration and the run was stopped after all routers congested for three increasing values of p . Packets were injected at a router (if sources and receivers (SR) were not present) or at a source (if SR were present). If classes (CL) are present, all packets needed to be forwarded by at least one core router. Note that configuration c0 is most similar to the model of Echenique, Gomez-Gardenes and Moreno [Echenique, Gomez-Gardenes and Moreno, 2005], but c0 uses a smaller 218-node topology from a single ISP network in our experiments. Configuration c0 uses the SPF ($h = 1$), has no packet dropping (PD) and has no other realistic factors from in Table 2.

For each of the 18 configurations, the growth of the largest cluster of congested sites was tracked as p increased. In configurations in which all routers became congested, the growth of the largest cluster of congested sites may occur rapidly so that there is evidence of percolation within the configuration. The total of congested sites was also tracked, whether or not it was part of the largest cluster. For some

⁵ With delays (DE), the number of valid combinations rises to 34. DE is omitted because it has the least influence on the results.

configurations, this largest cluster was small, yet the totals were large. For each of the 18 configurations, the decline of a second largest cluster of uncongested sites (which in some cases, ultimately consisted of no sites) was also tracked. This second process represented the breakup of network connectivity. The decline of this second largest cluster of uncongested sites was contrasted with the growth of the largest cluster of congested sites and the total number of congested sites. *The decline of this second largest cluster of uncongested sites, i.e., would only be observed if the total congested sites were large.* In our experiments, the decline of this second largest cluster of uncongested sites more closely followed the inverse of the total congested sites. This second largest cluster could not be exactly derived from the inverse of the largest cluster of congested sites.

4. Results

In what follows, results are presented in a series of four figures which display one chart for each configuration with our ISP topology of a computer communications network. The decimal configuration code from Table 3 is marked across the top of each chart. The vertical axis tracks the proportion of sites that are occupied by congestion (0 to 1), while the horizontal axis tracks the rate of injection of packets: $p = 1$ to $p = 2500$ in increments of 10 (note: our simulation will stop if all routers become congested). Configurations in which all routers become congested have their decimal configuration codes in marked in red. In such cases, a phase transition as defined in Section 3a can be said to be observed in a finite network. If congested routers have attached sources and receivers, these sources and receivers cannot communicate, and for practical purposes these sources and receivers are occupied.

Subsection (a) discusses the growth of the largest cluster of congested sites. Figure 3 in subsection (a) shows the results graphically. In eight configurations, all routers become congested. We may consider percolation does occur in these configurations, and a largest cluster of congested sites grows until it encompasses all routers. Thus, evidence of a phase transition can be said to be observed. Correspondingly, Figure 4 in subsection (b) shows the growth in totals of congested sites as congestion spreads, and subsection (b) discusses the related causes. In contrast, Figure 5 in subsection (c) shows decline of the largest cluster of uncongested sites and the breakdown of network connectivity. In eight configurations (the same eight configurations identified in subsection (a)), this largest cluster of uncongested sites ultimately drops to zero, and thus percolation occurs in these configurations—this is an inverted process, that is not necessarily the exact inverse of the processes discussed in subsections (a) or (b). Again, subsection (c) discusses the related causes for these plots. In all three graphs (Figs. 3-5), we see that if variable speeds (VS) or TCP are present, the largest cluster of congested sites never encompasses all sites; i.e., a catastrophic event or phase transition to a global failed state, does not occur. This conclusion is supported in the appendix, where a clustering analysis of Figs. 3-5 shows the influence of VS or TCP and of other realistic network factors. Packet dropping (PD) may affect the number of packets successfully delivered. However, other factors—classes (CL), sources and receivers (SR), and flows (FL)—may or may not be present in configurations where all clusters ultimately become congested (or the number of uncongested clusters drops to zero). Figure 6 in subsection (d) shows a decline in packets reaching their destinations, and subsection (d) discusses the underlying causes.

In all figures below, classes (CL) are present in all configurations but c0 and c1, and packet dropping (PD) is present in odd numbered configurations. Sources and receivers (SR) are present for configurations c18 and higher. Flows (FL) are present for 50 and above. Below c50, only individual packets are transmitted from a source to destination. TCP is present in configurations c114 and higher. Variable speeds (VS) are present in configurations c6, c7, c22, c23, c54, c55, c118, c119.

a. Growth of the Largest Cluster of Congested Sites

In Figure 3, the curve shows the growth of the largest cluster of sites with respect to congestion for 18 configurations of our ISP topology for $p = 1$ to $p = 2500$ in increments of 10 (unless all routers become congested before $p = 2500$ is reached).

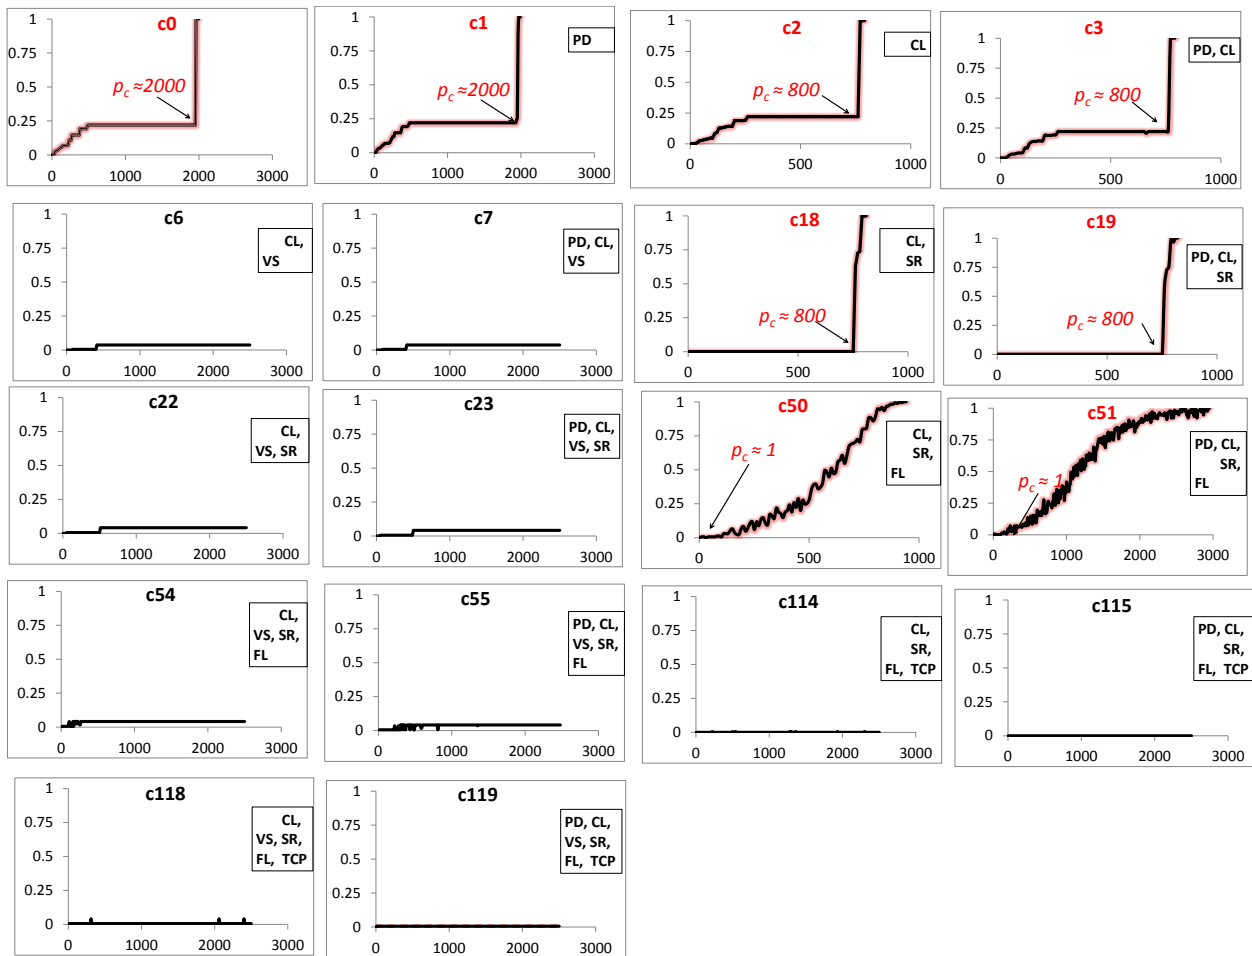


Figure 3. Growth of the largest cluster of congested sites for 18 configurations for $p = 1$ to $p = 2500$ in increments of 10 (unless all routers congest before $p = 2500$). Proportion of all sites in the largest congested cluster is the vertical axis (includes all sources and receivers beneath a congested router). The injection rate (p) is the horizontal axis. Classes (CL) are enabled in all configurations, except c0 and c1. Packet dropping (PD), sources and receivers (SR), variable speeds (VS), flows (FL), and TCP are enabled as shown. For configurations in which all routers congest and a phase transition is observed, p_c is shown.

In a subset of eight of these 18 configurations, evidence of percolation and a phase transition is observed when the buffers of *all* 218 routers are full to a point of at least 70 % of their buffer size. In these configurations, p is greater than p_c . These configurations have their configuration numbers marked in red, and an estimate of p_c is given. In none of the ten configurations with variable speeds (VS) or TCP present, does the largest congested cluster include all 218 routers (i.e., percolation is not observed in these configurations).

All routers do not congest when variable speeds (VS) are present, because the core routers (A-P) are fastest, the PoP routers (A1-P2) are slower than core routers, and the access routers are slower (except for **D**-class) according to Table 1. Variable speeds (VS) enables routers to be hierarchically engineered so that higher tiers (core routers) accommodate the largest expected packet sending rate from lower tiers (PoP routers and access routers beneath PoP routers). Therefore Core and PoP routers never reach their threshold in buffer queue size and therefore do not congest when VS is present. Hence, when VS is present in eight configurations in Figure 3 (c6, c7, c22, c23, c54, c55, c118, c119), congestion stays at the access routers on the edges of our single ISP network. Therefore, all routers never congest in these configurations. When VS is not present, also in eight different configurations in Figure 3 (c0, c1, c2, c3, c18, c19, c50, c51), all sites have the same speeds (9 packets / ts) and core routers will congest first because all traffic flows through them. The cluster of congested routers then grows to encompass all routers (and sources and receives, if any, attached to congested routers cannot communicate). When VS is not present in these eight configurations in Figure 3, it is possible to say that congestion—and congested routers—form first at the center and flow from the inside out (core and PoP routers to access routers)—a largest cluster of congested sites forms at the center—core and PoP routers—and grows to encompass all routers in the related configuration.

TCP provides connection procedures (associated with flows (FL)), rate adaptation, acknowledgements, and retransmission procedures in the event of a nack or timeout. Due to these procedures, buffers fill more slowly. In two configurations, c114 and c115, TCP is present without variable speeds (VS). In these two configurations, all routers do not congest, and the largest cluster of congested sites remains small. Hence, TCP can still establish a path along which to transmit packets even at high values of p . Therefore, sites do not congest in these two configurations as in others. But since VS is also absent in configurations c114 and c115, core router queues congest, while access routers (at the edges) remain uncongested. The congestion of the core routers has the effect of fragmenting the PoP routers and their access routers.

The presence of packet dropping (PD), classes (CL), sources and receivers (SR), and flows (FL) seems not to forestall the congestion of all routers. Similarly, in configurations c2, c4, c19, and c51, all routers congest when PD is present. In configurations c2, c4, c18, c19, c50, c51, all routers congest when CL is present. In configurations c18, c19, c50, and c51, all routers congest when sources and SR is present. In configurations c50 and c51, all routers congest when FL is present. Therefore, percolation can occur when PD, CL, SR, and FL are present.

b. Growth of the Totals of Congested Sites

Figure 4 shows totals with respect to congestion for 18 configurations for $p = 1$ to $p = 2500$ in increments of 10. Total congested sites obviously are high in the eight configurations where all routers congest and a phase transition is observed (c0, c1, c2, c3, c18, c19, c50, c51). However, total congested sites are also high in other configurations, though not all routers congest. As we will see in discussing Figure 5, network connectivity breaks down in such cases.

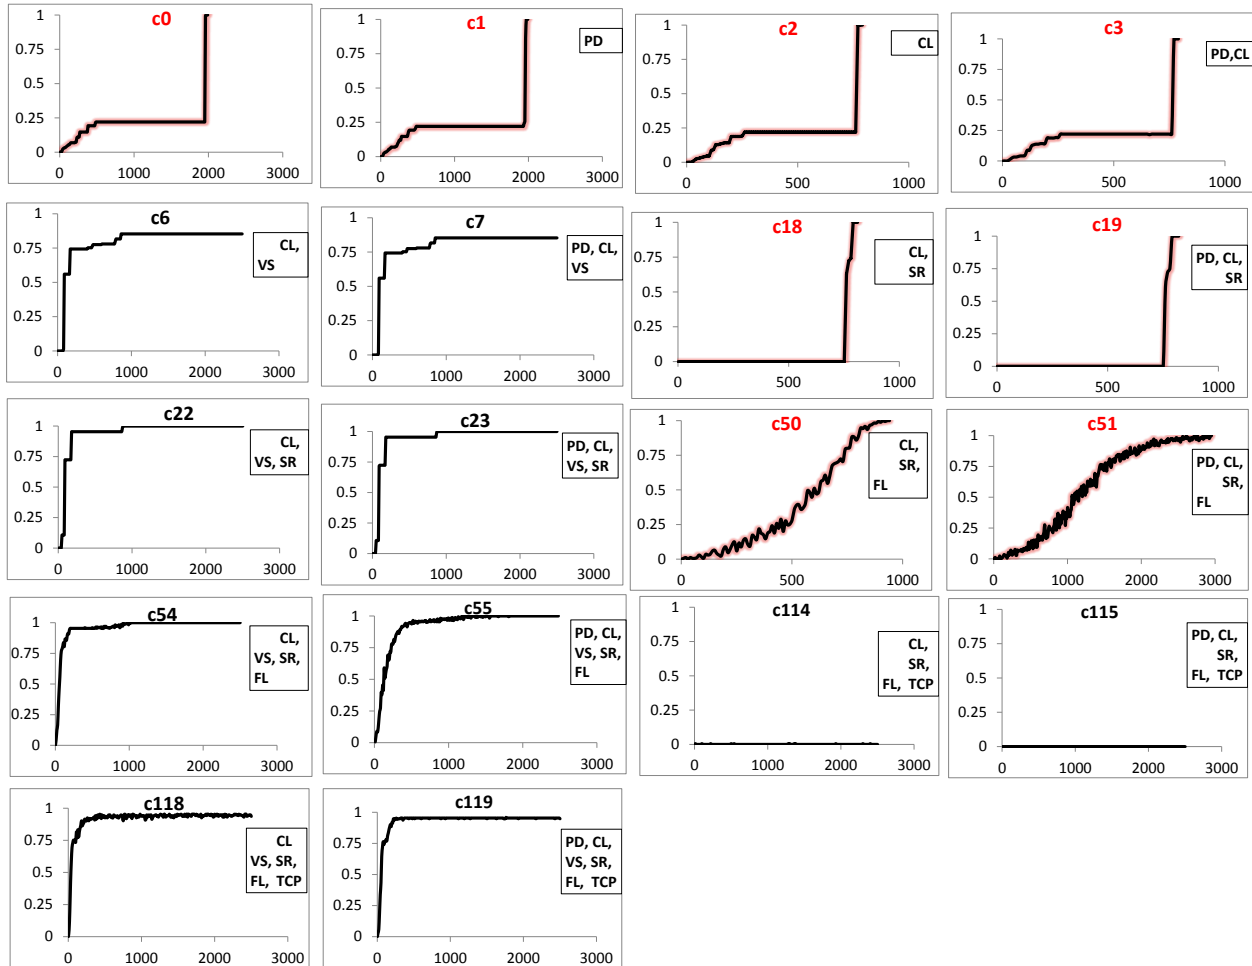


Figure 4. Totals of congested sites for 18 configurations for $p = 1$ to $p = 2500$ (unless all routers congest before $p = 2500$). Proportion of all sites occupied by congestion is the vertical axis (where applicable, includes all sources and receivers beneath a congested router). The injection rate (p) is the horizontal axis. Classes (CL) are enabled in all configurations, except c0 and c1. Packet dropping (PD), sources and receivers (SR), variable speeds (VS), flows (FL), and TCP are enabled as shown. Eight configurations are shown in red in which all routers congest and evidence of a phase transition is observed. In ten other configurations, all routers do not congest. Since these configurations have VS or TCP, the core routers remain uncongested, while the peripheral access routers may congest. Since some of these ten configurations have SR under access routers, over a quarter million source and receiver sites cannot communicate and are included in this figure.

In Figure 4, we see the growth in the total number of congested sites in eight configurations with variable speeds (VS), i.e., c6, c7, c22, c23, c54, c55, c118, and c119—even though in Figure 3, the largest cluster is small with respect to congestion for these configurations. This is because in these eight configurations, the congested routers are the slower, predominantly peripheral access routers, while the faster core routers, near the middle of the topology, remain uncongested as explained above. Hence, when VS is present in these eight configurations, the largest clusters are isolated at the periphery and small, as reflected in Figure 3. In six of these configurations, c22, c23, c54, c55, c118, and c119, the peripheral access routers also have sources and receivers (SR), which cannot communicate beyond themselves, and are therefore affected by their congested peripheral routers, hence the totals are almost one for these configurations. That is, these plots nearly show percolation, but the small number of uncongested core and PoP routers prevent a large cluster of congested nodes from forming.

In contrast in two configurations, c114 and c115 (where all routers do not congest in Figure 3 for reasons given in subsection (a)), the total number of congested routers remains small. In these two configurations, TCP is present and variable speeds (VS) are not; hence without VS, most router speeds are slower and have the same speed. TCP rate adaptation procedures therefore cause congestion to be abated, both at core and peripheral routers. Consequently, the largest cluster of congested routers remains small, and totals are small also. Meanwhile, in configurations c118 and c119 which have both TCP and VS, the largest cluster of congested routers remains small and totals are large. In configurations c118 and c119, VS is present, routers speeds are faster, and packets accumulate more quickly than if VS is absent. Hence, the periphery becomes congested and totals are large, but the relatively small but fast core remains free. The largest cluster in the single ISP network is thus fragmented and percolation cannot take place for reasons given above.

c. Decline of Network Connectivity—the largest cluster of Uncongested Routers

In Figure 5, the curves for all 18 configurations of our ISP topology show a precipitous decline in network connectivity as p goes from one to 2500, i.e., one sees the decline of uncongested routers (and their attached sources and receivers (SR), if any). This decline in network connectivity is demarked by the decline of the largest cluster that contains all routers that remain uncongested. In eight configurations (c0, c1, c2, c3, c18, c19, c50, c51), this largest cluster consists of uncongested routers whose size goes to zero as network connectivity breaks up. These configurations are shown in red in Figure 5, and they are the same configurations in which all routers congest in Figure 3. In these eight configurations, $p_c \approx 1$ in Figure 5, and a phase transition is observed. Again, the presence or absence of packet dropping (PD) in all 18 configurations seems not to determine whether this precipitous decline in network connectivity occurs. Similarly, the presence or absence of PD does not determine whether the largest cluster of uncongested routers goes to zero in eight configurations. As also noted in discussing Figure 3, the presence or absence of classes (CL), SR, and flows (FL), or combinations of them, seems to have no effect either.

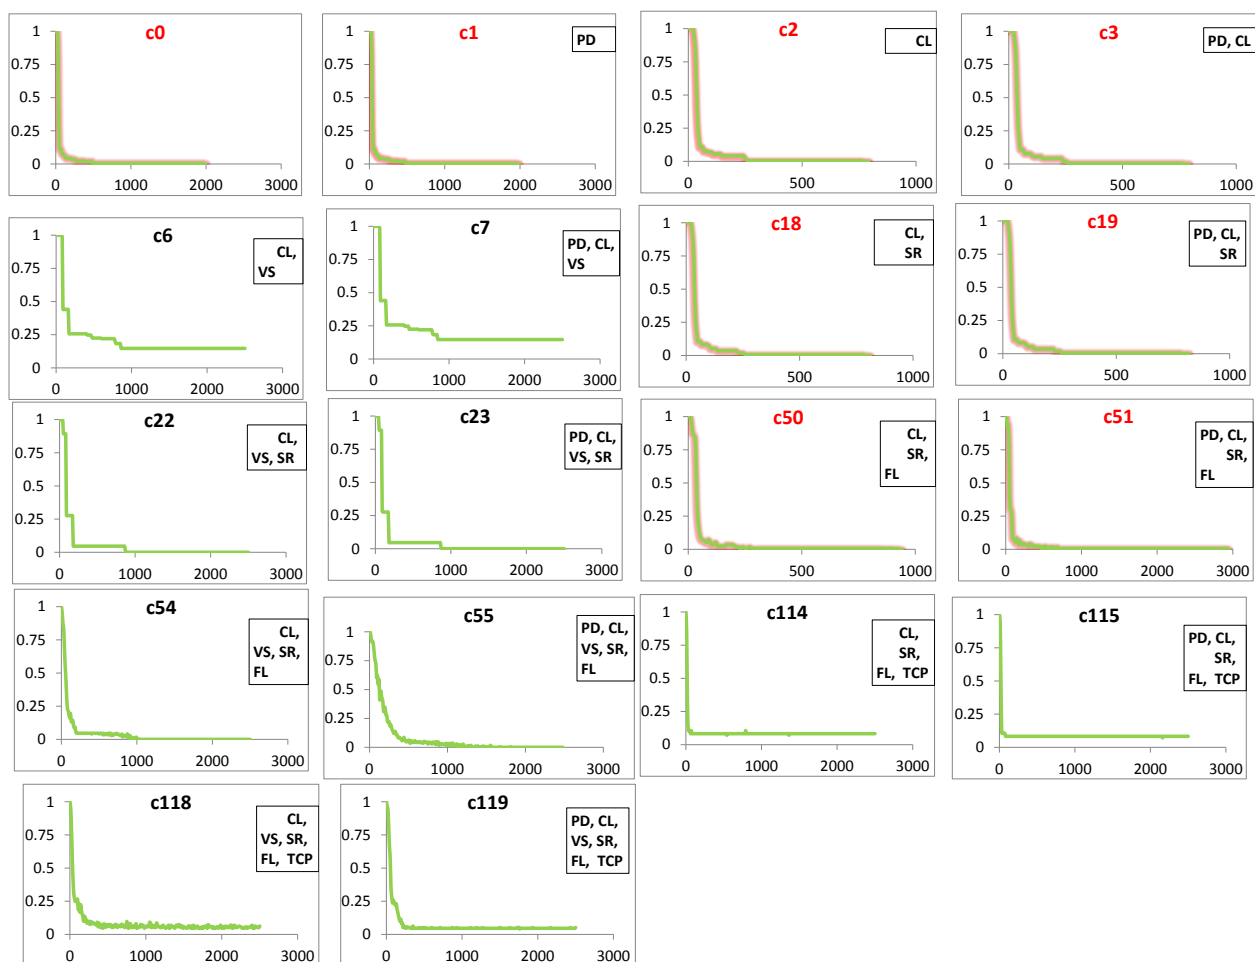


Figure 5. Decline in uncongested sites, i.e., the breakdown of network connectivity for 18 configurations from $p = 1$ to $p = 2500$ (unless all routers congest before $p = 2500$). Proportion of all sites not occupied by congestion is the vertical axis (includes all sources and receivers beneath a congested router). The injection rate (p) is the horizontal axis. Classes (CL) are enabled in all configurations, except c0 and c1. Packet dropping (PD), sources and receivers (SR), variable speeds (VS), flows (FL), and TCP are enabled as shown. For configurations where the largest cluster of uncongested sites declines to zero and evidence of a phase transition can be observed, $p_c \approx 1$ and their configuration numbers are marked in red. Four configurations, c22, c23, c54, and c55, appear to decline to zero, but actually do not. Since these four configurations have VS, the core routers remain uncongested, while the peripheral access routers congest. Since these four configurations also have SR, over a quarter million source and receiver sites cannot communicate and must be included in *breakdown of network connectivity*.

Neither configurations c114 nor c115 have variable speeds (VS), but they do have TCP. Because they do not have VS, some core routers congest. As mentioned above, with this particular combination, the relatively small clusters of congested sites formed by the core routers have the effect of fragmenting the PoP routers and their access routers—so the largest clusters of uncongested sites still decline, and therefore network connectivity also diminishes.

However in all 18 configurations, Figure 5 shows that even if variable speeds (VS) are present, network connectivity (i.e., the number of uncongested routers) declines due to a rise in the total congested routers. This means that at higher congestion levels shown here, the network connectivity ultimately suffers, if p is raised sufficiently.

d. Decline in Packets Reaching Destinations

In Figure 6, the curves show packets reaching destinations for 18 configurations of our ISP topology for $p = 1$ to $p = 2500$. The decline in proportion of packets reaching their destinations is also evident in Table 4. In eight configurations (c0, c1, c2, c3, c18, c19, c50, c51), the number of packets reaching their destinations goes to zero. They are the same configurations in which all routers congest and a phase transition is observed in Figure 3. These configurations are shown in red in Figure 5.

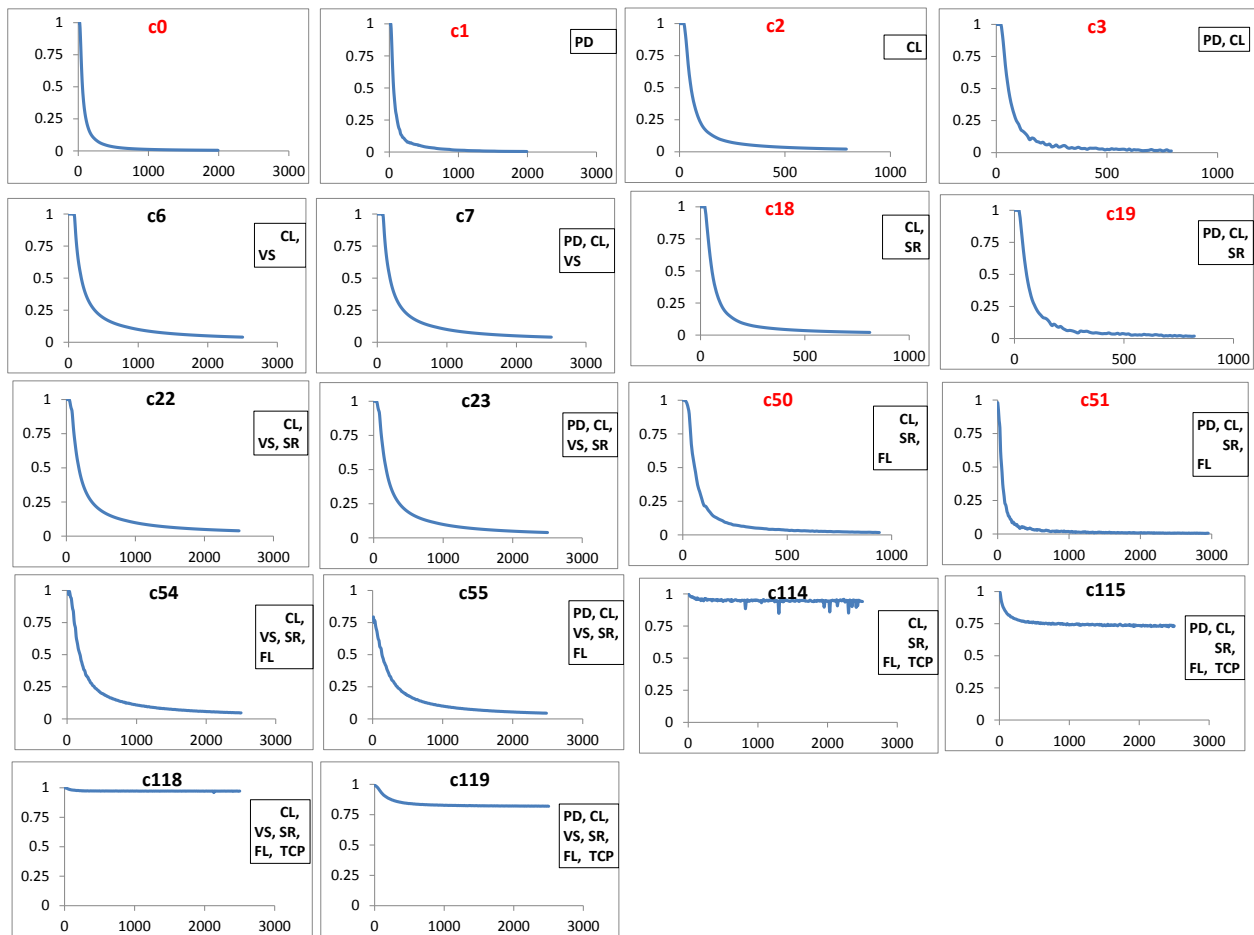


Figure 6. Decline in packets reaching destinations for 18 configurations for $p = 1$ to $p = 2500$ (unless all routers congest before $p = 2500$).` Proportion of packets reaching their destinations is the vertical axis includes all sources and receivers, if applicable. The injection rate (p) is the horizontal axis. Classes (CL) are enabled in all configurations except c0 and c1. Packet dropping (PD), variable speeds (VS), sources and receivers (SR), flows (FL), and TCP are enabled as shown. Configurations where all routers congest and evidence of a phase transition is observed have configuration numbers in red.

The impact of TCP and variable speeds (VS) in mitigating, and preventing, a phase transition, as shown in Fig. 6, can be explained further. When TCP is present, in configurations c114, c115, c118, and c119, this dramatic reduction does not take place—though the number of packets reaching their destinations is less than 100 %, as shown in Figure 6 and quantified in Table 4.

Table 4. Percentage packets reaching their destinations for packet injections from $p = 1$ to $p = 2500$ (unless all routers congest before $p = 2500$). Configurations where all routers congest have configuration numbers in red.

Configuration Number	Percent Reached	Configuration Number	Percent Reached
c0	2.41 %	c23	7.57 %
c1	1.60 %	c50	4.84 %
c2	5.37 %	c51	1.17 %
c3	3.81 %	c54	11.56 %
c6	10.67 %	c55	7.88 %
c7	7.67 %	c114	94.75 %
c18	5.34 %	c115	75.09 %
c19	4.27 %	c118	97.33 %
c22	10.53 %	c119	83.46 %

The impact of TCP and variable speeds (VS) in mitigating, and preventing, a phase transition, as shown in Fig. 6, can be explained further. When both TCP and VS are present, over 97 % of packets reach their destinations in configuration c118 (see Table 4). In configuration c119 where TCP and VS are also both present, over 83 % of the packets reach their destinations. In those four configurations (c114, c115, c118, and c119) which have TCP, a higher proportion of packets reach their destinations (than in other configurations), according to Table 4. In these configurations, all packets are organized in flows (FL) which are subject to TCP rate adaptation. In configurations c114, c115, c118, and c119, TCP rate adaptation has the effect of slowing flows, so while fewer packets are transmitted (than in other configurations), a higher proportion of packets do reach their destinations in these four configurations, as shown both in Table 4 and Figure 6. In configurations c114 and c115, where TCP is present and not VS, over 94 % of packets reach their destinations in configuration c114, and over 75 % of packets reach their destinations in c115.

We see that in many configurations including c114 and c118, a higher proportion of packets eventually reach their destinations without packet dropping (PD) than with PD. As it is, when PD is absent, queues are longer but no packets are dropped. Therefore, a higher proportion of packets reach their destinations, though packets take longer to travel from source to destination. When PD is present, queues can be shorter. Some packets are lost due to the presence of PD, and the proportion of packets reaching their destinations will be lower with PD. Beyond this, it is possible to sum the number of packets that are sent and reach their destinations for all 18 configurations. It is then possible to see, by totaling the relevant configurations, that these total packets reaching destinations are higher (by

1.37E+10) when PD is present. In configurations where TCP is also present, TCP can operate more quickly with PD than without PD.

5. Discussion

The figures in Section 4 showed that all routers congested (and sources and routers, if any, beneath them could not communicate), only when a configuration of such a network lacked both variable speeds (VS) and TCP—as is the case for eight configurations which experience a catastrophic congestion collapse and transition to a global failed state. In these same eight configurations, network connectivity completely declined, i.e., the number of uncongested sites went to zero. Hence in these eight configurations, a phase-transition event can be seen, and percolation in a graph as defined in Section 3a takes place. Yet, network connectivity declined significantly for all 18 configurations, even if a phase transition event as defined in Section 3 was not observed. Figure 5 shows this most clearly. Even if VS was present, network connectivity declines severely as total congested routers rose.

The Growth in Congested Sites. All routers congested only when a configuration lacked both variable speeds (VS) and TCP in eight cases in Figure 3. In those eight cases, when all routers are operating at the same speed, a phase-transition event can be seen, and percolation of a graph takes place. Hence, the core routers and the rest of the network are overwhelmed with congestion. When either VS or TCP is present in the other 10 configurations, all routers did not congest. When VS is present, routers are engineered with hierarchical speeds, so that the capacity of core routers at least equals the packets received from PoP routers and access routers. When TCP was present, rate adaptation reduces congestion, preventing all routers from congesting. These conclusions about VS and TCP are supported in the Appendix, where a hierarchical clustering analysis shows that the influence of these two factors. In configurations which have other realistic network factors but have neither TCP nor VS—that is, configurations which have only packet dropping (PD), classes (CL), sources and receivers (SR), flows (FL) or some combination of all four—the largest cluster of congested sites grows to encompass all routers. In our finite network, Figure 3 shows a largest cluster of congested sites emerges and continues to grow in the eight configurations until all routers congest, that is in configurations which have neither VS nor TCP.

The Decline in Uncongested Sites: Connectivity Breakdown of the Network. Figure 5 shows that even though complete congestion spread to all routers is limited to eight configurations, network connectivity declines for all 18 configurations. These declines are complete in the eight configurations in which all routers congest in Figures 3; that is, in the eight configurations which have neither variable speeds (VS) nor TCP. Therefore, Figure 5 shows there is evidence that the largest cluster of uncongested sites, representing network connectivity, completely disintegrates precisely in the same eight configurations, which congest completely. Again, a phase-transition-like event can be seen, and percolation of a graph takes place. Nevertheless, if p is raised high enough, network connectivity suffers in all configurations. The cause of the decline in network connectivity (uncongested routers) in all configurations in Figure 5 is the growth of congested routers. These totals are seen clearly in Figure 4 for all configurations except c114 and c115. In configurations c114 and c115, which have TCP and not VS, rate adaptation prevents routers at the periphery from congesting because packet transmission slows as congestion is

encountered in the core, and thus queues at the network edge are smaller. Since all packets must flow through the core in configurations c114 and c115, the core fragments, and network connectivity is lost. Network connectivity also declined in configurations c6, c7, c22, c23, c54, c55 which have VS and not TCP, because of the growth in total congested routers in Figure 4 (though the presence of VS slightly delays the breakdown of network connectivity). In these cases, network connectivity declined, even though the core remained intact and congestion is confined to the periphery (while the largest cluster of congested sites remained small in Figure 3). Network connectivity also declined in configurations c118 and c119, which have both VS and TCP even though again the core again remained intact and congestion was confined to the periphery (while the largest cluster of congested sites again remained small). Hence, it is beneficial to see the graphs of total congested sites to fully understand how and why network connectivity breaks down.

Total Packets Reaching Destinations. The causes of the decline in total packets reaching their destinations in Figure 6 in 14 of 18 configurations, except c114, c115, c118, and c119, is also ultimately the growth in congestion. In configurations c114, c115, c118 and c119, TCP is present, and TCP rate adaption causes packet queues to be smaller. Packet transmission slows down because of rate adaptation: fewer packets are sent, and thus the network outputs fewer packets in total. However, TCP rate adaptation also reduces congestion, therefore a higher proportion of sent packets arrive at their destinations in these four configurations with TCP, than in the 14 configurations without TCP, as Table 4 shows. In fact, with TCP most packets reach their destinations in configurations c114, c115, c118, and c119.

The Effect of TCP. Even at high levels of p , all routers do not congest in configurations which have only TCP (c114 and c115) and not variable speeds (VS). In those configurations, the largest cluster of congested sites remains small and never grows to encompass all routers. Yet, in both c114 and c115, network connectivity with respect to congestion also collapses because a significant portion of core routers congest and fragment the entire network. Without VS, some core routers congest because many PoP and access routers send packets to the core, which cause queues to build in their associated core routers. In these configurations (c114 and C115), TCP then detects the congestion in the core and rate adapts, so that the congestion at the edge is reduced. Nevertheless, some core routers still congest, intermediate PoP routers have both congested and uncongested states, and so the network is fragmented and connectivity breaks down. With both VS and TCP (configurations c114, c115, c118, and c119), congestion remains at the periphery, while core routers are completely uncongested.

The Effect of variable speeds (VS). VS enables routers to conform to common Internet engineering principles, so that higher tiers (core routers) accommodate the largest expected packet sending rate from lower tiers (PoP routers and access routers beneath PoP routers). Therefore according to Table 1 with VS, core routers are much faster than PoP routers, and PoP routers are faster than access routers (except for D-class). All packet Injection takes place at the periphery in all configurations with VS. However, because of hierarchical engineering in these configurations, congestion remains at the periphery while the core remains uncongested. Thus with VS present, all routers do not congest, and observation of a phase transition event does not occur. When VS is not present, all sites were at the same packet speed (9 packets / s). In this case (unless TCP was present), core routers congested first

because all packet traffic flowed through the core and then spread outward. All packet Injection also takes place at the periphery (except configurations c0 and c1) when VS is not present. When VS and TCP are both not present, congestion also started at the core and spread to PoP and access routers—and all routers eventually congested (and sources and routers, if any, beneath them could not communicate) as Figures 3 and 5 show. The presence of VS also affects queue size. With VS, queues are smaller in some routers. With VS, queues are very small for the 16 core routers and most of the PoP routers. Without VS, queues are longer, and all routers congest as discussed above.

The Effect of Packet Dropping (PD). In Table 4, for the even-numbered configurations which have no PD, a higher proportion of packets actually reach their destinations, even though congestion is greater (for example, configurations c114 and c118). This is because without PD, packets stay in a queue which becomes very long, but will eventually reach their destinations. However, in our simulations users cannot exercise the realistic option of ending flows (FL) if they take too long or progress too slowly. With PD in force in odd-numbered configurations, packets are dropped and must be resent, even though congestion is less and queues are shorter than without PD (for example, configurations c115 and c119). With PD, total packets reaching destinations are higher than without PD. Yet without PD, in our simulations it turns out that a higher percentage of packets usually reach their destinations in configurations c114 and c118. As we see in Section 4, phase transition events are observed with and without PD.

The Influence of Realistic Network Factors. In the cases with most realistic network factors (configurations c118 and c119)—most packets in a flow reach their destination. In configurations c118 and c119, packets proceed to their destinations in a flow (FL), and therefore as part of a flow, they can more easily produce an intelligible web page or document at a destination site. Configurations c118 and c119 have variable speeds (VS) and TCP. To support these factors, it is necessary to have classes (CL), FL, and sources and receivers (SR). Configuration c119 has all six realistic network factors, including packet dropping (PD). In c118 and c119, the presence of VS causes congestion to collect on the periphery and the core to remain uncongested. Hence, it is possible to generalize that more realistic network factors enable the Internet to operate in an intelligible fashion and mitigate congestion, unless congestion is very high. In configurations c114 and c115, which also have FL and TCP but not VS, the majority of packets also reach their destinations, though fewer than in c118 and c119. Yet even with the presence of realistic network factors that mitigate congestion, network connectivity still fails in all configurations, when p is raised high enough. Hence, with more realistic factors, observation of a phase transition event is less likely.

6. Conclusions

Based on the experiment conducted above, it is possible to conclude that realistic factors in a computer communications network can be grouped coherently and added to a finite graph model. It is also possible to conclude that models of the Internet must include realistic network factors to provide an accurate picture of catastrophic failures or observation of phase transitions. Generally, there is evidence of percolation with respect to congestion if network models are abstract; percolation is not observed when more realistic network factors are included. In more abstracted models with fewer realistic

assumptions about network factors, congestion spreads more rapidly and reaches a critical point, after which the network becomes entirely congested and a catastrophic failure thus occurs. We see that two realistic factors—in particular variable speeds (VS) and TCP—can prevent percolation from being observed with respect to congestion; that is, these two factors can prevent a catastrophic event or phase transition to a global failed state. This is not the case with many abstract network models appearing in previous papers, which did not include VS or TCP, in which percolation was observed and in which adverse conclusions about the Internet were drawn. Computer model studies of networks often did not include widely-used Internet routing and congestion-control protocols, which would likely impact the spreading processes and change the nature of catastrophic collapses and observed phase transitions [Dabrowski, 2015].

For computer communications networks, our conclusions are consistent with Alderson and Willinger, who pointed out that the differences between the network topologies used in research models and the topology of the Internet [Alderson and Willinger, 2005]. They conducted a study that found both abstract scale-free networks and the Internet had scale-free connectivity in a graph model, but only the Internet had high-degree hub sites primarily on the periphery. Consistent conclusions were also reached by [Doyle et al., 2005]. Willinger et al. [Willinger et al., 2002] gave a method for determining whether a model explaining emergent behavior was evocative or explanatory. That is, their method contrasted models that showed how emergent behavior took place (i.e., were evocative) with models that showed why emergent behavior took place (i.e., were explanatory). They further suggested that models needed to be explanatory and gave examples of applying their method in the cases of Internet self-similarity and in the scale-free topology formed using preferential attachment. Specifically, Willinger et al. stated: “Because [evocative models] ignore important networking-specific details and fail to exploit the rich semantic content of the available measurements, they can lead to incorrect conclusions about the causes and origins of the emergent phenomena at hand”. Our paper provides a practical example of this point. In our case, we believe that we are providing an explanatory model with respect to catastrophic congestion rather than an evocative model. Further, we are showing that the assumptions underlying the evocative model are incorrect and that when realistic network factors are used, the observed behavior is both different and explainable.

We have presented six realistic network factors, whose inclusion in a graph model of an ISP network could forestall, or prevent, widespread congestion collapse and evidence of percolation. Widespread congestion collapse can occur as the level of congestion is increased in the sites of an ISP network. Realistic network factors can be included in our model individually or in combination to yield distinct configurations, which can be tested—as we have shown. The inclusion of these realistic network factors enables the Internet, as we have come to know it, to operate and communicate—and forestalls and prevents congestion collapse. The exclusion of key realistic factors yields widespread congestion collapse in our models, as many other papers have attested. With the exclusion of these realistic factors, the operation of our model also experienced congestion collapse and a phase transition event was observed. However, even with the inclusion of realistic factors, all 18 configurations experienced a severe decline in network connectivity as the number of packets gradually increased, and total

congested routers rose. Network connectivity declined whether or not a phase transition event was observed.

Even when the entire network did not experience such a collapse, parts of it congested, and sources and receivers within them could not communicate across the network. In our simulations, these portions could be readily identified and constituted potentially vulnerable parts of the network. The occurrence of a catastrophic collapse, or a phase transition, also has been seen in computer models of other types of networks, such as electrical power grids [Carreras et al., 2002], or models of virus and disease spread [Pastor-Satorras and Vespignani, 2001; Moreno, Pastor-Satorras and Vespignani, 2002]. Accordingly, models such as these should be examined to see if inclusion of additional realistic factors is necessary to obtain a more accurate picture of network collapse.

As for future work, we recommend for communications networks the development of further models that include more realistic network factors, such as user patience, improved models which include sources and receivers (SR) that have their own queues and can congest, expanded queuing schemes for core routers, and more elaborate linking among routers. Especially important for testing our conclusions, we recommend experimentation with different ISP topologies or with models with interconnected ISP topologies. We also recommend that models in other network domains be examined to see if inclusion of additional realistic factors is necessary to obtain an accurate picture of phase transitions.

7. Acknowledgments

We appreciate and thank those researchers who, over the past decade or so, studied the applicability of graph theory and percolation theory to congestion in communication networks. Their work inspired us to plan a program of research to explore the practicality of the ideas in real networks.

We thank our laboratory management and appreciate the financial support and encouragement provided by them. The technical staff benefit greatly from management's visionary thinking.

We thank those who reviewed this paper and provided suggestions for improving it, including Jim Filliben, who reviewed our experiment design. We also thank Phillip Gough of the Commonwealth Scientific and Industrial Research Organization of Australia for creating a multidimensional visualization that enabled exploration of congestion spread in our simulated topology.

8. References

- D. Alderson and W. Willinger, 2005. A Contrasting Look at Self-Organization in the Internet and Next-Generation Communication Networks. *IEEE Communications Magazine*. Volume 43, Issue 7, pp. 94–100.
- D. Arrowsmith, R. Mondragon, J. Pitts, and M. Woolf, 2004. *Phase Transitions in Packet Traffic on Regular Networks a comparison of source types and topologies*. Report 08, 2004/2005 Institut Mittag-Leffler.
- A. Barabási and R. Albert, 1999. Emergence of Scaling in Random Networks. *Science*. Volume 286, pp. 509–512.
- D. ben-Avraham and S. Havlin, 2000. *Diffusion and Reaction in Fractals and Disordered Systems*. Cambridge University Press, Cambridge England.
- B. Bollobás, 1984. The Evolution of Random Graphs. *Transactions of the American Mathematical Society*. Volume 286, Number 1, pp. 257–274.
- D. Callaway, M. Newman, S. Strogatz, and D. Watts, 2000. Network Robustness and Fragility: Percolation on Random Graphs. *Physical Review E*. Volume 85, pp. 5468–5471.
- B. Carreras, V. Lynch, I. Dobson, and D. Newman, 2002. Critical points and transitions in an electric power transmission model for cascading failure blackouts. *Chaos: An Interdisciplinary Journal of Nonlinear Science*. Volume 12, Number 4, pp. 985–994.
- R. Cohen, K. Erez, D. ben-Avraham and S. Havlin, 2000. Resilience of the Internet to Random Breakdowns. *Physical Review of Letters*. Volume 85, Number 21, pp.4626–4628.
- C. Dabrowski, 2015. The Study of Catastrophic Event Phenomena in Communication Networks. *Computer Science Review*. Volume 18, pp. 10–45.
- C. Dabrowski and K. Mills. *The Influence of Realism on Congestion in Network Simulations*. NIST Technical Note 1905, January 2015, 62 pages.
- S. Dorogovtsev, A. Goltsev, and J. Mendes, 2008. Critical phenomena in complex networks. *Review of Modern Physics*. Volume 80, Issue 4, pp. 1275–1445.
- J. Doyle, D. Alderson, L. Li, S. Low, M. Rougan, S. Shalunov, R. Tanaka and W. Willinger, 2005. The “robust yet fragile” nature of the internet. *Proceedings of the National Academy of Sciences*. Volume 102, Number 41, pp. 14497–14502.
- P. Echenique, J. Gomez-Gardenes and Y. Moreno, 2005. Dynamics of jamming transitions in complex networks. *Europhysics Letters*. Volume 71, Number 2, pp. 056105-1–056105-4.
- A. Lawniczak, P. Lio, S. Xie, and J. Xu, 2007. Study of Packet Traffic Fluctuations Near Phase Transition Point From Free Flow to Congestion in Data Network Model. *Canadian Conference on Electrical and Computer Engineering. CCECE 2007*. pp. 360–363, April 22–26.

- K. Mills, J. Filliben, D. Cho, E. Schwartz and D. Genin, 2010. *Study of Proposed Internet Congestion Control Mechanisms*, NIST Special Publication 500-282, May 2010, 534 pages.
- G. Mukherjee and S. Manna, 2005. Phase transition in a directed traffic flow network. *Physical Review E*. Volume 71, Issue 6.
- T. Lei and Z. Yu, 2008. Analytical Results of Phase Transitions Behaviors in Internet traffic on Complex Network. *4th International Conference on Wireless Communications, Networking and Mobile Computing, 2008. WICOM '08*. 12-14 Oct. 2008, pp. 1-4.
- R. Pastor-Satorras and A. Vespignani, 2001. Epidemic dynamics and endemic states in complex networks. *Physical Review E*. Volume 6, pp. 066117-1–066117-8.
- Y. Moreno, R. Pastor-Satorras and A. Vespignani, 2002. Epidemic outbreaks in complex heterogeneous networks. *European Physical Journal B*. Volume 26, pp. 521–529.
- Y. Rykalova, L. Levitan, R. Brower, 2010. Critical phenomena in discrete-time interconnection networks. *Physica A*. Volume 389. 2010. pp. 5259–5278.
- S. Sarkar, K. Mukherjee, A. Ray, A. Srivastav and T. Wettergren, 2012. Statistical Mechanics-Inspired Modeling of Heterogeneous Packet Transmission in Communication Networks. *IEEE Transactions on Systems, Man, and Cybernetics—Part B: Cybernetics*, Volume 42, Number 4, pp. 1083–1094.
- R. Solé and S. Valverde, 2001. Information Transfer and phase transitions in a model of internet traffic. *Physica A*. Volume 289, pp. 595–605.
- B. Tadic, G. Rodgers, and S. Thurner, 2007. Transport on Complex Networks: Flow, Jamming and Optimization. *International Journal of Bifurcation and Chaos*. Volume 17, Number 7, pp. 2363–2385.
- D. Wang, N. Cai, Y. Jing, and S. Zhang, 2009. Phase Transition in Complex Networks. *American Control Conference*. St. Louis, MO, USA. June 10-12, pp. 3310–3313.
- W. Wang, Z. Wu, R. Jiang, G. Chen, and Y. Lai, 2009. Abrupt transition to complete congestion on complex networks and control. *Chaos*. Volume 19, pp. 033106.
- W. Willinger, R. Govindan, S. Jamin, V. Paxson, and S. Shenker, 2002. Scaling phenomena in the Internet: Critically examining criticality. *Proceedings of the National Academy of Sciences of the United States of America*. Volume 99, Supplement 1, pp. 2573–2580.
- M. Woolf, D. Arrowsmith, R. Mondragón, J. Pitts, and S. Zhou, 2004. Dynamical modeling of TCP packet traffic on scale-free networks. *Institut Mittag-Leffler*. Volume 6. 9 pages.
- Z. Wu, W. Wang and K. Yeung, 2008. Traffic dynamics in scale-free networks with limited buffers and decongestion strategy. *New Journal of Physics*. Volume 10, pp. 023025.

APPENDIX

This appendix contains four figures which show a hierarchical clustering (MATLAB⁶) of 18 configurations for Figures 3-6, from $p = 1$ to $p = 2500$. In all four figures in this section, configuration numbers appear on the x -axis, while the y -axis shows squared Euclidean distance as computed by the Matlab program. That is, MATLAB computes the distance between vectors representing congestion of routers in the 18 configurations. Hierarchical clustering places data into a cluster tree or *dendrogram*. The tree is a multilevel hierarchy of clusters, where clusters at one level are joined to the level above. Configurations are grouped together based on their Euclidean distances in Figures A1-A4 by the hierarchical clustering program and through visual interpretation. For example, configurations that have variable speeds (VS) can be grouped together, while configurations that do not have VS or TCP can also be grouped together.

In contrast, configurations with other factors, such as classes (CL), packet dropping (PD), sources and receivers (SR), and flows (FL) are not grouped together in Figures A1-A4. Moreover, when any combination of these four factors is present with VS and TCP both absent, then all routers can congest (and therefore, all sources and receivers, if any, cannot communicate). For example, in one configuration (c51), CL, SR, FL, and PD are all present (while VS and TCP are both absent)—and all routers congest nevertheless.

According to the hierarchical clustering in the figures in this appendix, configurations which lack variable speeds (VS) are often grouped together, while configurations that lack TCP are also often grouped together. Configurations that lack both VS and TCP are grouped together. The lack of these realistic network factors also allows the congestion of all routers and the existence of a phase-transition-like event. Since these groupings are indicated in all four figures, VS and TCP seem to be the most influential factors according to hierarchical clustering. The hierarchical clustering represents a second method which supports the results given in the main body of the paper. Separate discussions of Figures A1-A4 ensue. For an analysis of all six realistic network factors, please see Sections 5 and 6.

⁶ The identification of any commercial product or trade name does not imply endorsement or recommendation by the National Institute of Standards and Technology, nor is it intended to imply that the materials or equipment identified are necessarily the best available for the purpose.

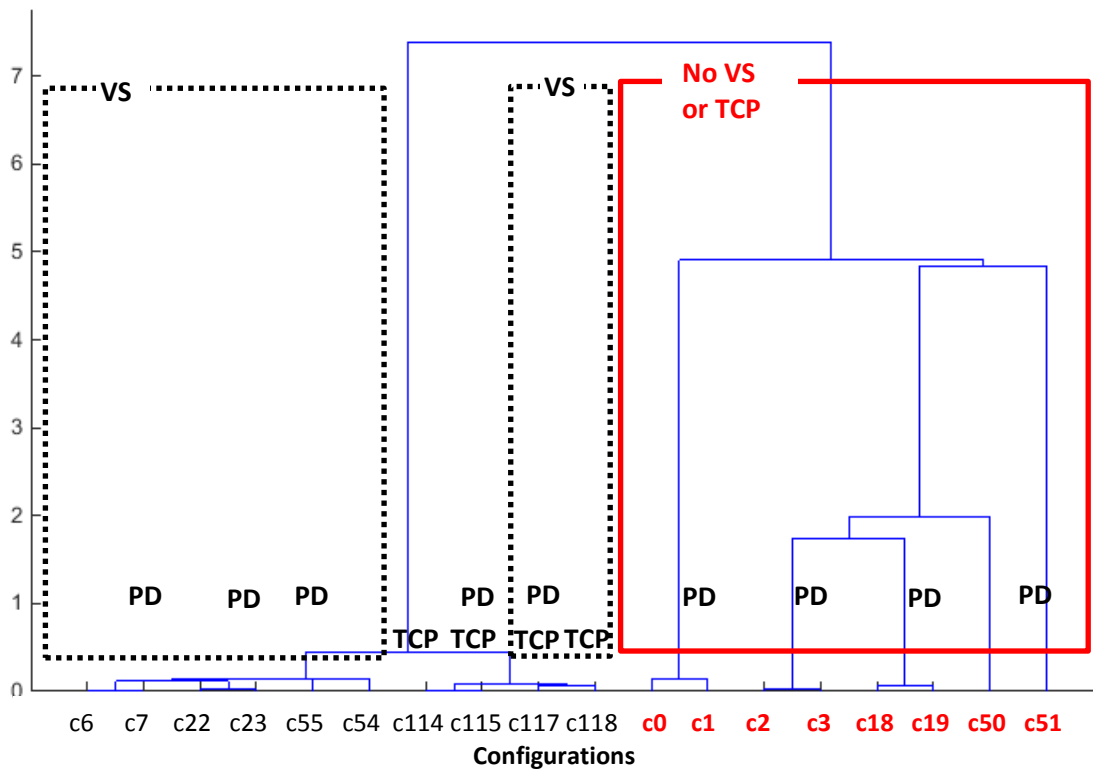


Figure A1. Hierarchical Clustering for data in Figure 3 shown as a dendrogram. *Growth of the largest cluster of congested sites for 18 configurations for $p = 1$ to $p = 2500$.* Configurations in which all routers congest (and their respective sources and receivers, if any, cannot communicate) are marked with their configuration number in red. These configurations do not have variable speeds (VS) nor TCP. A solid red box also denotes these configurations. Dotted black boxes denote configurations which have VS. Configurations which have packet dropping (PD) and TCP are individually marked.

Figure A1 represents the clustering of plots from Figure 3. In Figure A1, those configurations which have variable speeds (VS) can be grouped together in a dotted box. They are separated by a large Euclidean distance from configurations that do not have VS or TCP. Configurations which do not have VS or TCP are also grouped together in a solid red box, and in these configurations a phase transition event is observed. Configurations which have only TCP are individually annotated as such. In contrast, configurations having packet dropping (PD), classes (CL), sources and receivers (SR), and flows (FL) seem interspersed among the above groupings, and therefore are not grouped together.

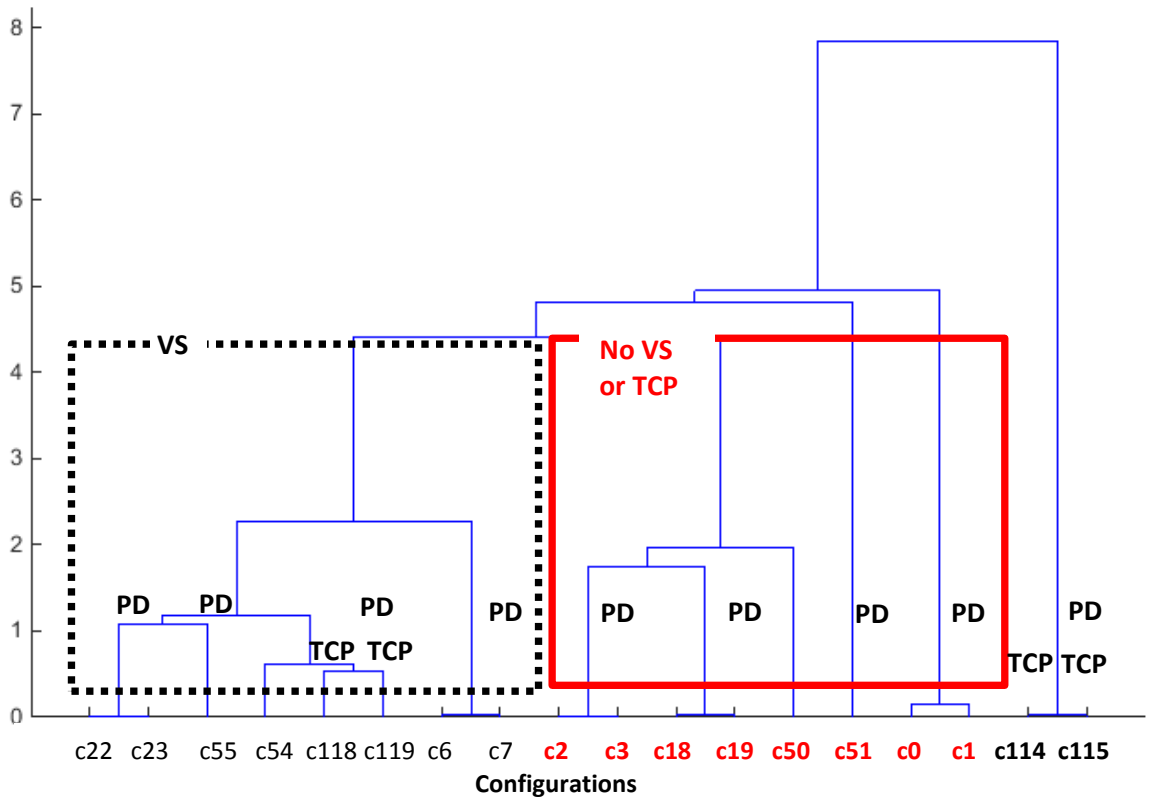


Figure A2. Hierarchical Clustering for data in Figure 4 shown as a dendrogram. *Totals for congested sites for 18 configurations for $p = 1$ to $p = 2500$.* Configurations in which all routers congest (and their respective sources and receivers, if any, cannot communicate) are marked with their configuration number in red. These configurations do not have variable speeds (VS) nor TCP. A solid red box also denotes these configurations. Dotted black boxes denote configurations which have VS. Configurations which have packet dropping (PD) and TCP are individually marked.

Figure A2 represents the clustering of plots from Figure 4. In Figure A2, those configurations which have variable speeds (VS) are grouped together under a dotted black box. Those configurations that do not have VS or TCP are also grouped together in a solid red box, and in these configurations a phase transition event is observed. Configurations which have only TCP are individually annotated as such. In contrast, configurations having packet dropping (PD), classes (CL), sources and receivers (SR), and flows (FL) are not grouped together.

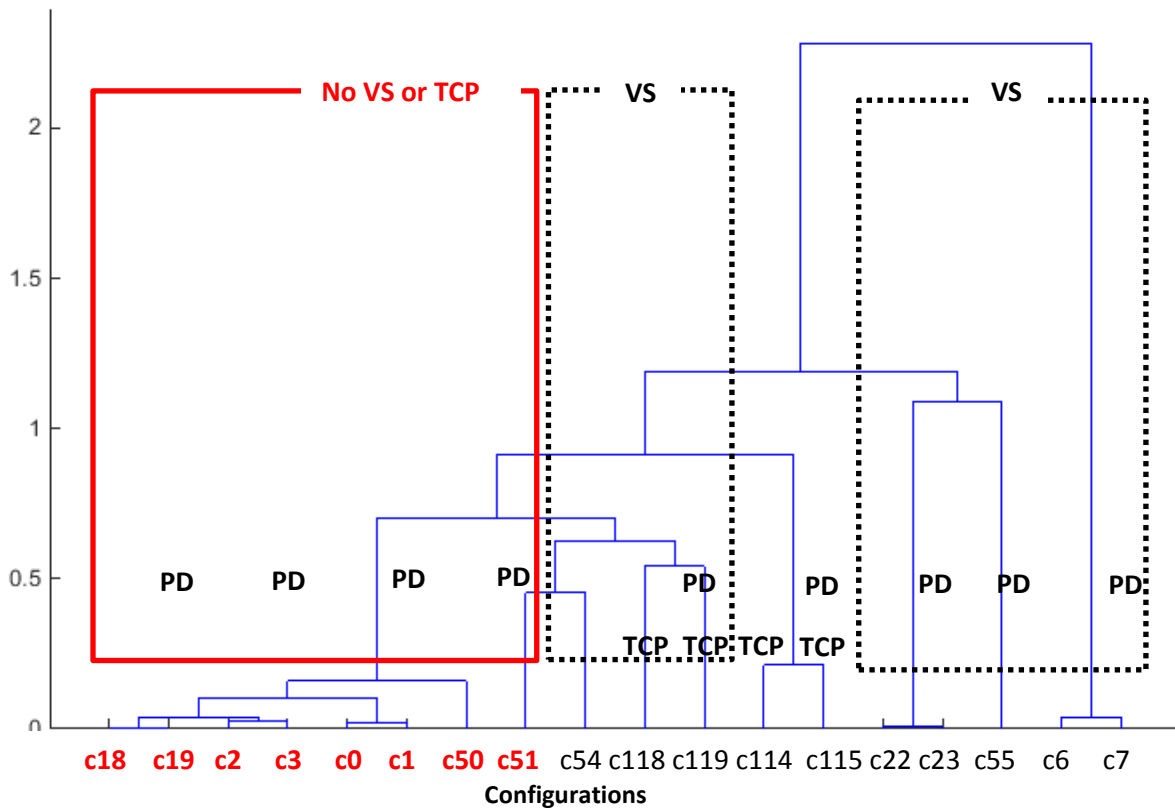


Figure A3. Hierarchical Clustering for data in Figure 5 shown as a dendrogram. *Decline in the largest cluster of uncongested sites* for 18 configurations for $p = 1$ to $p = 2500$. Configurations in which all routers congest (and their respective sources and receivers, if any, cannot communicate) are marked with their configuration number in red. These configurations do not have variable speeds (VS) nor TCP. A solid red box also denotes these configurations. Dotted black boxes denote configurations which have VS. Configurations which have packet dropping (PD) and TCP are individually marked.

Figure A3 represents the clustering of plots from Figure 5. In Figure A3, those configurations which have variable speeds (VS) can be grouped together, and a dotted black box encircles these. Those configurations that do not have VS or TCP are also grouped together, and a solid red box encircles these configurations. In these eight configurations, a phase transition event is observed. Configurations which have only TCP are individually annotated as such. In contrast, configurations having packet dropping (PD), classes (CL), sources and receivers (SR), and flows (FL) are not grouped together. Note that distances among the groupings in Figure A3 are generally smaller than those seen among groupings in Figures A1, A2, or A4. This reflects the fact that the decline in the largest cluster of uncongested sites, or the breakdown in network connectivity, is more similar among all configurations—as Figure 5 also shows.

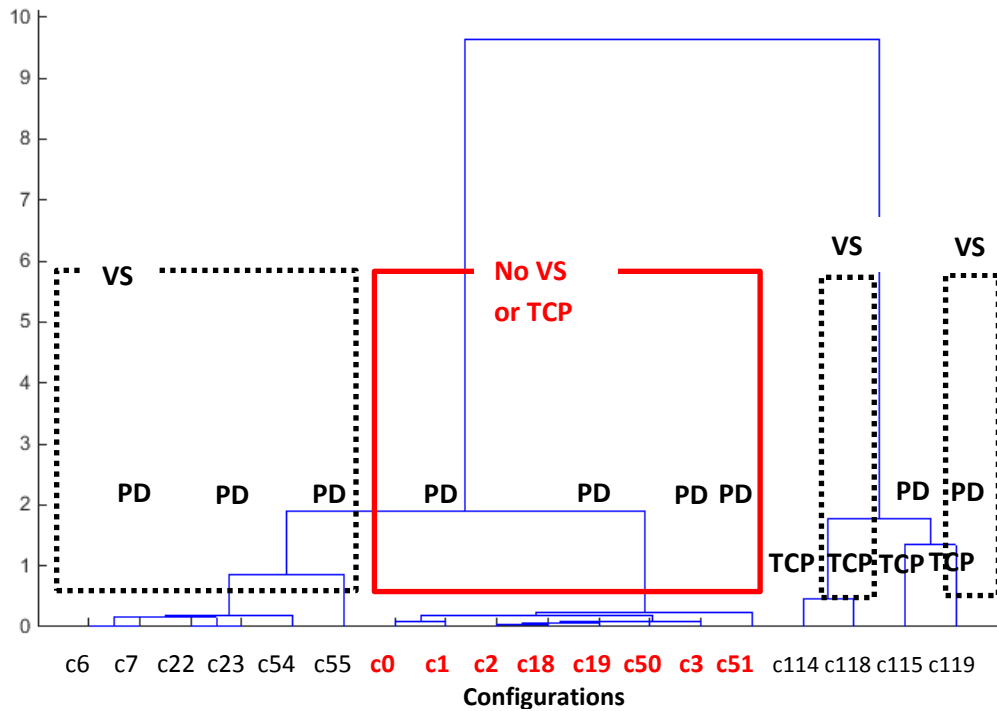


Figure A4. Hierarchical Clustering for data in Figure 6 shown as a dendrogram. *Decline in packets reaching destinations for 18 configurations from $p = 1$ to $p = 2500$.* Configurations in which all routers congest (and their respective sources and receivers, if any, cannot communicate) are marked with their configuration number in red. These configurations do not have variable speeds (VS) nor TCP. A solid red box also denotes these configurations. Dotted black boxes denote configurations which have VS. Configurations which have packet dropping (PD) and TCP are individually marked.

Figure A4 represents the clustering of plots from Figure 6. In Figure A4, those configurations that have variable speeds (VS) can be grouped together in a dotted black box. Configurations can also be grouped together, if they do not have VS or TCP, and once again a solid red box surrounds these. In these eight configurations, a phase transition event is observed. In contrast, configurations having packet dropping (PD), classes (CL), sources and receivers (SR), and flows (FL) are not grouped together. In this figure, configurations having TCP (c114, c115, c117, and c118) appear together. These four configurations with TCP also are the ones in which the majority of packets reach their destinations (recall Figure 6). In Figure A4, these four configurations could also be grouped together.