Computer Communications 112 (2017) 93–108

Contents lists available at ScienceDirect

Computer Communications

journal homepage: www.elsevier.com/locate/comcom

Using realistic factors to simulate catastrophic congestion events in a network

Christopher Dabrowski*, Kevin Mills

National Institute of Standards and Technology (NIST), 100 Bureau Drive, Mail Stop 8900, Gaithersburg, MD 20899-8900, United States

ARTICLE INFO

Article history:
Received 1 February 2016
Revised 27 June 2017
Accepted 10 August 2017
Available online 19 August 2017

Keywords:
Internet
Networks
Percolation
Phase transition
Simulation

ABSTRACT

With the rapid growth of the Internet, there has been increased interest in the research literature in the use of computer models to study the dynamics of communication networks. 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 representation of such important network factors as variable router speeds and packet buffer size limits. Here, we modify this typical approach, adding realistic network factors to a graph model of a single Internet Service Provider (ISP) network that can have more than a quarter million nodes. The realistic factors in our model, including router classes, variable router speeds, flows, the transmission control protocol (TCP), sources and receivers, and packet dropping, can be enabled and disabled in combinations. For each combination of realistic factors, we gradually increase network load, and gauge spread of congestion throughout the network. While there are realistic computer communications 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, does not appear. Yet, as load increases, 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.

Published by Elsevier B.V.

1. Introduction

With the rapid growth of the Internet, there has been increased interest in the research literature in the use of models, based on graphs, to study the dynamics of computer communication networks. 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 [10,13,22,23,27,30]. In these studies, 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 packets,[1] 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 computer communications networks, as reported in previous papers, have been abstract graphs with simplified behavioral factors, e.g., simplified forms of routing and queue management [2,10,15,23,25,27,29,30]. Few papers have explored the effects of including more realistic factors.

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 nodes. In our model, we simulate configurations of both (a) abstract

* Corresponding author.
E-mail addresses: cdabrowski@nist.gov (C. Dabrowski), kmills@nist.gov (K. Mills).

[1] A packet is a well-known atomic unit of information transferred in a network.

http://dx.doi.org/10.1016/j.comcom.2017.08.006
0140-3664/Published by Elsevier B.V.

communication networks with little realism and (b) increasingly realistic networks, in which six realistic network factors are gradually added. We then compare and analyze congestion spread for combinations of realistic factors.

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. 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. Catastrophic congestion events might, or might not, occur in less abstract models, depending on which realistic network factors are included. Generally, a catastrophic collapse or phase transition can occur only if the graph model omits key realistic factors. We study this collapse in terms of a *percolated state*[2] [3], preceded by a phase transition–in which all routers congest. We show that the presence of two realistic factors—variable router speeds and the transmission control protocol (TCP)–mitigate catastrophic congestion and the occurrence of phase transitions. We find that the spread of congestion, which leads to a catastrophic failure in a communications network, should be modeled using realistic factors. This is because congestion behaves differently in an abstract network model than in a network modeled with 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 for realistic factors should be considered for use in computer models of other types of networks, such as electrical power grids [6] or models of virus and disease spread [18,20].

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

To preview what follows below, we describe and present simulations for 18 combinations of six realistic network factors. For each combination, network load is elevated gradually by increasing the number of packets injected into the system at discrete intervals. By increasing the rate of packet injection, we hope to trigger a congestion collapse within our model. With the addition of variable router speeds or TCP in 10 of the 18 configurations, a complete collapse and the preceding 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 any sources and receivers attached below these routers cannot communicate in the network. This is effectively evidence of a percolated state [3]. In our case, a percolated state occurs when all routers of the ISP network are occupied by congestion, as will be explained below. Including more realistic network factors in the model 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 and a realistic topology in modeling a communications network, which were raised previously by Alderson and Willinger [1].

However, as congestion spreads, we find that ultimately the operation of the network, and the ability of member nodes to communicate declines. This decline is evident even when a total congestion collapse and a phase transition are not observed. Sources and receivers cannot communicate if they are attached to a congested router. In this paper, Section 2 discusses previous work, Section 3 covers definitions and the experiment plan, Section 4 presents results, Section 5 discusses the results, and Section 6 concludes. An Appendix presents information that supports our conclusions. A technical note [9] 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) to a critical point, after which increased load altered the behavior of the system [2,13,15,19,22–25,27,30]. 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 in two studies [15,22], while others [2,19,22,23] modeled a two-dimensional lattice, but differentiated host and router sites (with [2] also studying triangular and hexagonal lattices). One study [21] used ring and toroidal topologies. Other studies focused on *scale-free*[3] networks [13,30], with a realistic topology and simplified behavior [13].

The routing algorithms used in these efforts appear to be motivated by Internet processes. For example, some studies [2,23] used *shortest path first* (SPF) routing to determine which site to forward a packet to. Alternatively, the packet could be forwarded to the least congested site [23] or forwarded along the least-used link [2]. These models also sometimes used a combination of criteria to determine where to route packets, including shortest path, availability of buffer space at the destination site, and congestion awareness [13,15,25]. Some researchers used routing procedures based, at least in part, on randomly determining where to forward a packet [19,22,24,30]. Infinite-sized buffers were assumed, except for some examples [24,27,29,30], where finite-buffer sizes with packet dropping procedures for overflow were assumed. In one paper [30], models were varied to use both finite and infinite 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 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, we also increased packet injection, but we used a realistic topology adapted from a single ISP. In addition, our routing algorithms and network factors were more realistic than in previous studies.

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

[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 most sites have far fewer links. See also [11].

Among the more realistic topologies in previous studies, was that used by Echenique, Gómez-Gardeñes, and Moreno [13], which featured a 11 174-router snapshot taken from the Internet Autonomous System (AS) topology circa 2001, collected by the Oregon Route Views Project. The work of Echenique, Gómez-Gardeñes, and Moreno was the baseline for our work (as it was in Wang et al. [26], who studied congestion in city automobile traffic), on top of which we introduced realistic network factors. For this reason, we repeated their experiments and obtained similar results. The work of Echenique, Gómez-Gardeñes, and Moreno also showed that congestion occurred in two ways, depending on the method of packet forwarding. A gradual phase transition to a congested state was observed if packets were forwarded based on a SPF algorithm, as occurs in the Internet. However, a steeper phase transition occurred when the routing algorithm was changed to forward packets through routers that were less congested. Some have likened this difference in steepness to the difference in the order of phase transitions. Such difference occurs in our work as well, as explained below.

3. Definitions and experiment plan

This section consists of four subsections, in which Section 3.1 gives key percolation definitions; Section 3.2 describes our experiment topology, method of packet injection, and the simulation platform that we used; Section 3.3 explains the realistic network factors that we used, and how we derived router speeds and buffer sizes; and Section 3.4 discusses how we combined these realistic factors into network configurations.

3.1. 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 $N \equiv \{n_1, n_2, \ldots, n_N\}$ are the *nodes* (or *vertices*) of the graph G, while the elements of $L \equiv \{l_1, l_2, \ldots, l_K\}$ are its *links* (or *edges*), where each l_i consists of a pair of elements of N. Percolation processes have been studied in graphs [3–5,7]. In our ISP model, a site is considered a node (or vertex), where a site is any router, source, or receiver.

In site percolation, a spreading property progressively occupies all sites within an infinite network [3]. 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 router i in direction d (*up*, or *down*) reaches or exceeds Q_T, or $Q_{i,d} \geq Q_T$. (Note that, in our model, core routers have only one queue, while access and PoP routers have two queues (d = down toward the edge and up toward the core). 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 have no queues. In the ISP topology that we study, the network is declared percolated only if *all* routers are occupied with congestion. Sources and receivers cannot communicate across the network if they are attached to a congested access router. In our model, p will stand for the rate (per time step) of packet injection.

In principal, if p increases so that it exceeds some critical probability, known as the *percolation threshold* or *critical point*, p_c, then all routers congest. If $p > p_c$, a *percolation transition* occurs in an infinite network. An infinite *Giant Connected Component* (*GCC*) then emerges, such that all nodes in the GCC exhibit congestion. The proportion of nodes included in the GCC is represented by P_∞, and the GCC grows until $P_\infty = 1$. The symbol for infinity is present, because theoretically a GCC forms only in an infinite-sized network [3].

In our model, which of course is finite, all routers congest in eight configurations, in which case their attached sources and receivers, if any, cannot communicate across the network. Therefore, in these eight configurations, a GCC begins to emerge at p_c and continues to grow as more routers are occupied by congestion. In our finite-sized network, even though percolation cannot occur in the theoretical sense, we consider percolation to be *observed* in a practical sense when all routers are occupied. For these eight configurations, we consider a percolation transition to be *observed* when $p > p_c$, and we provide an estimate of p_c. Above p_c, in our model, the GCC is represented by the largest cluster of occupied sites. A phase transition is considered to take place when the percolation transition occurs. This process can also be inverted, with a second GCC of unoccupied, or uncongested, sites declining, as congestion increases, until there are no unoccupied, or uncongested, sites. As described in the next section, the decline of a GCC, or largest cluster, of uncongested sites represents the breakdown of network connectivity.

3.2. Experimental topology, packet injection, and simulation platform

In our model, in place of the 11 174-router topology used by Echenique, Gómez-Gardeñes, and Moreno [13], we substituted a 218-router topology encompassing a single-ISP network spanning the continental United States. The topology we used is depicted in Fig. 1.

Our 218-router single-ISP topology has three hierarchical router tiers: (1) 16 core routers (A–P in Fig. 1); (2) 32 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., the eight red nodes in Fig. 1, which connect directly to core routers), **F**-class (e.g., the 40 green nodes) and **N**-class (e.g., the 122 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 below the third tier. If sources and receivers are added (as explained below), there are a total of 258 158 sites in the topology. When sources and receivers are not present, but routers are classified, packet injection takes place only at the access router sites on the edge of the model. If the fourth tier is present, packet injection takes place only at source sites and not at the attached access routers. The three tiers, all classes, as well as all sources and receivers, can be eliminated. Eliminating these allows for a uniform, flat 218-router topology. In this case, routers are not classified, and packet injection can occur at any router in the topology.

As with Echenique, Gómez-Gardeñes, and Moreno [13], packets are injected 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 (ts) using the equation:

$$injectionChance = InjectionRateP / numberInjectionSites, \quad (1)$$

where *InjectionRateP* is the current packet-injection rate, given in packets / *ts*. The number of sites where the packet can be injected is given by *numberInjectionSites*. Applied repeatedly at each *ts*, Eq. (1) determines the next *ts* when a source will inject a packet. Eq. (1) is repeated at each site where packets can be injected, until a defined number of *ts* is reached and the simulation ends. Packets flowing between a source-receiver pair follow a single ingress/egress path between an access router and a top-tier core router. When flows are present, Eq. (1) is altered so that the average flow size, measured in packets, is also in the denominator.

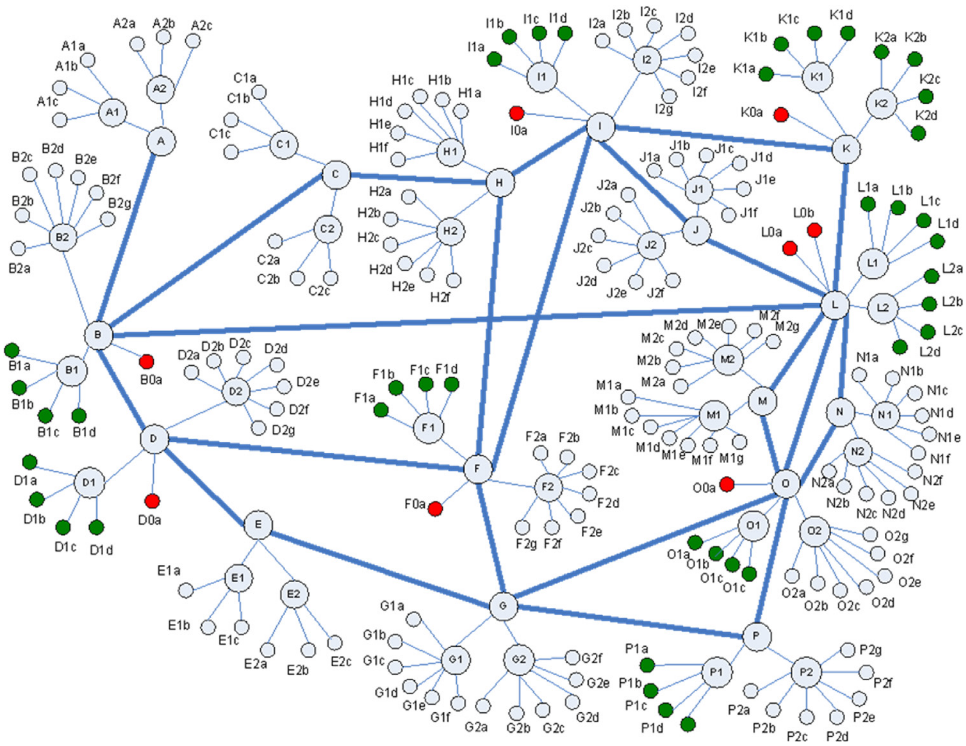

Fig. 1. The three-tier topology with 16 core routers (A–P), 32 PoP routers (A1–P2), and 170 access routers (A1a–P2g), including 8 red D-class, 40 green F-class, and 122 gray N-class access routers. From [17]. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

This change allows computation of when the next flow of packets will be injected at a site.

The simulation was conducted on a Windows High Performance Computing (HPC) cluster, where each node in the cluster contained dual Xeon X5670 2,93 GHz processors (some with 4 cores and some with 6 cores). The simulation was developed and executed in SLX [14], a commercial, highly reliable, high-speed, simulation environment, written in the C language.

Table 1
Speed relationships among router classes. From [17].

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

3.3. Realistic network factors, router speeds, and buffer sizes

In what follows, we describe six realistic factors derived from the study by Mills et al. [17]. Realistic factors also were chosen partly based on the paper by Alderson and Willinger [1]. Both works discussed the need for realism in models of the Internet. Mills et al. modeled in detail several realistic factors present in the Internet, including packet dropping (PD), classes (CL), sources and receivers (SR), flows (FL), variable speeds (VS), and TCP. When PD is present, incoming packets to a router whose buffer is full are simply dropped and never reach their destinations. (Buffer sizing 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, once processed by a router, packets are forwarded on toward their destinations. In our model, only routers can drop packets.

When CL is present, there are core, PoP, and access routers, with three subtypes 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 the network periphery: at sources when SR is present, and at access routers otherwise. With CL, packet receivers are also network peripheral nodes: receivers when SR is present, and access routers otherwise. Also with CL, packets are required to transit at least one core router. Without CL, packets may be injected and received at any node; hence network paths can be shorter then when CL is present.

When present, variable speeds (VS), as shown in Table 1 [17], allow core routers, PoP routers, and different classes of access routers to forward packets at different speeds. Unlike real networks, where links have transmission speeds and associated buffers, transmission speeds are subsumed directly into our router forwarding speeds. In our model, packets have no size, and so forwarding speeds are assigned in units of packets / *ts*. Here, we scale down speeds by 1 / 40 from the speeds shown in Table 1. Therefore, in our model, core routers forward 80 packets / *ts* and PoP routers 10 packets / *ts*, **D**-class access routers process 10 packets / *ts*, **F**-class access routers process two packets / *ts*, while **N**-class routers process 1 packet / *ts*. If sources and receives are also present together with VS, then 50% of sources and receivers are randomly selected to operate at 2.0 packets / *ts*, while 50% process 0.2 packets / *ts*. Mirroring Table 1, our goal was to establish reasonable engineering relationships among the speeds of the various router classes. Without VS, routers, sources, and receivers, if any, operate at the same speed: 9 packets / *ts*, which is computed from the weighted average of all routers when VS is present.

With, or without VS, buffer sizes are determined based on router speeds, using **ceil** ($250 \times$ *routerProcessingSpeed*).[4] Without VS, all routers have the same buffer size—2250 packets, a size which is determined by 250 times the average speed or 250 $ts \times 9$ packets / ts. Each core router multiplexes packet forwarding from a single buffer, while PoP and access routers have two buffers each, one heading in the direction d toward the core ($d =$ up), and one heading from the core ($d =$ down). PoP and access routers alternate forwarding between each of the two buffers.

Sources and receivers (SR) constitute an optional fourth tier. When present, sources and receivers are distributed under access routers. Sources equate to computers that have information that receivers wish to download. Bypassing some of the details contained in previous papers [9,17], 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, 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 topology with 258 158 sites, including the 218 routers.

When flows (FL) are present, packets are injected as correlated streams. Each flow consists of many packets, representing an email, Web page, or document, as in the actual Internet. Each source will periodically transfer a flow, after randomly selecting a receiver 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, a source first needs to establish a connection with the randomly selected receiver. Once the source receives a connection *acknowledgement*, or *ack*, from a receiver within three tries, the source sends the flow. The number of packets in a flow is selected randomly, using a Pareto distribution with a Shape parameter of 1.5 [9,17]. The average number of packets is 350, which we selected for these simulations. When TCP is not present, the source sends flow packets at its assigned transmission speed, without waiting for an ack and without engaging in any retransmissions. In this case, the source operates as a UDP Source, and the source can potentially complete transmitting a flow without any packet in the flow reaching its destination. Each source is limited to one flow at a time and cannot start a second flow until completing a previous flow. Similarly, a receiver cannot receive two flows simultaneously.

A final realistic factor, TCP congestion control, is used on over 90% of all traffic on the Internet. In our model, TCP can operate only when flows (FL) are present, i.e., there is no congestion control applied to injection of individual packets. Our TCP model, as taken from Mills et al. [17], monitors acknowledgments for packets in a flow, and detects and adapts to congestion. TCP modulates the rate at which a source sends packets. The sending rate is controlled by the size of a *congestion window*. This window represents the number of packets that can be sent by a source before receiving a corresponding ack. With TCP, an ack 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, or 2^{30} / 2 packets. If no packets are lost, 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 with each received *ack*.

If an incoming packet is received out-of-order (as for instance, when packet dropping causes packets to arrive out of order), the receiver sends a negative acknowledgement, or *nack*, back to the source. Upon receiving a nack, the source decreases its congestion window by one half, and enters a congestion avoidance phase—in effect slowing down transmission. The out-of-order, or nacked, packet, sent originally by the source, must be retransmitted. If a source fails to receive an ack, or nack, within a specified time, a *timeout* is declared, the congestion window is cut down to two packets, and the slow-start threshold is cut by one half. In no case is the size of congestion window allowed to fall below two packets. This 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 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 our model adhered to Mills et al. [17] was documented and discussed elsewhere [9].

Table 2
Existing dependencies among realistic network factors.

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

3.4. Configurations: combinations of realistic network factors

In our experiments, realistic network factors were added individually or in combination to produce 18 configurations. Since there were six realistic factors, this yielded 2^6, or 64, possible combinations, but all were not implemented, due to various dependencies, shown in Table 2. Dependencies occur when one realistic factor needs another to function coherently. For example, variable speeds (VS) need classes (CL), so that VS can be assigned to appropriate router types. In Table 1, core routers have one speed (3200 packets / ts), while PoP routers have another (400 packets / ts). Without CL, VS cannot be assigned. Similarly, sources and receivers (SR) need CL to be assigned.[5] Flows (FL) are started by sources, and for each flow, a receiver also must be identified. Without the presence of SR and with the presence of FL, flows would have to be implemented between routers alone, resulting in unnecessary redundancy.[6] Hence, FL is dependent upon SR. TCP needs FL, since TCP operates only on flows.[7]

We eliminated any combination of realistic network factors that did not observe these dependencies. For example, variable speeds (VS) without classes (CL) and TCP without flows (FL), were elimi-

[4] Buffer size is calculated by a variant of the Bush-Meyer criterion, or (250 $ts \times$ router packet forwarding rate). Therefore, with variable speeds (VS), core, PoP, and different classes of access routers will have different buffer sizes. If the router has two buffers instead of one, each buffer has one half of the total buffer size.

[5] While sources and receivers (SR) might be included as a second tier under a flat topology (i.e., without classes (CL)), SR is restricted to be a fourth tier under access routers: (1) to emulate the common case on the Internet and (2) because additional configurations would be needed. Therefore, our dependencies stipulate that SR requires CL.

[6] Implementing FL between access routers would replicate cases where flows were assigned to sources and receivers. This would provide redundant results and again increase the number of configurations.

[7] TCP regulates the rate of packet transmission within a stream of related packets, retransmitting those that are not received, and deciding when all packets in the stream have been delivered successfully. These steps cannot be taken without a flow.

Table 3
The configurations that were simulated. The table shows 18 configurations with various combinations of packet dropping (PD), classes (CL), variable speeds (VS), sources and receivers (SR), flows (FL), and TCP.

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

nated. The 18 remaining combinations that we simulated appear in Table 3 (see below), i.e., these were all the combinations allowed by the dependencies shown in Table 2.

In Table 3, each combination is a configuration that has a decimal number, followed by a list of abbreviations for realistic factors that are present in the configuration, ranging from c0 (no realistic factors present) to c119 (all realistic factors present). In a related technical note [9], there is an additional realistic network factor, propagation delay (DE), which is omitted here.[8] With DE, this brings the total number of possible configurations to 128. We retain the c0–c128 numbering of our configurations to maintain consistency with the technical note.

Each of the 18 configurations in Table 3 was run once[9], where a single data point constituted 200 000 *ts*. A data point was stopped early if computer memory was exhausted, as for example could happen when packet dropping (PD) is absent. For each run, p, the packet-injection rate, was advanced from one to 2500 in increments of 10, yielding 250 data points per run (except in the case of configuration c51, which congested at $p = 3000$). In eight configurations, percolation was observed within 200 000 *ts*. When all router sites congested for three increasing values of p, evidence of percolation was declared in the configuration and related simulations stopped. Thus, some runs finished prior to, or after, p reached its limit. Packets were injected at a router (if sources and receivers (SR) were not present) or at a source (if SR was present). If classes (CL) was present, all packets were forwarded by at least two core routers. Note that configuration c0 is most like the model of Echenique, Gómez-Gardeñes, and Moreno [13], but c0 uses a smaller 218−node topology. Configuration c0 uses SPF, has no packet dropping (PD), and has no other realistic factors.

For each of the 18 configurations, the total number of congested sites and 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 could occur rapidly, so that there was evidence of percolation. For each of the 18 configurations, the decline of the largest cluster of uncongested sites (which in some cases, ultimately consisted of no sites) was also tracked. The decline in the largest cluster of uncongested sites represented the breakup of network connectivity. However, the largest cluster of uncongested sites could not be exactly derived from the inverse of the largest cluster of congested sites.

4. Results

The motivation and purpose of this paper is to study the spread of congestion through a network under realistic conditions. Therefore, we focus on the formation of the largest cluster of congested sites, because the largest cluster provides a distinctive view on the influence of realistic factors in spread of congestion. However, to give perspective and to show that congestion and congestion spread are not identical concepts, we first examine total congested sites, and then proceed to show the largest cluster of congested sites. Taking this approach to presenting the results allows readers to see that congestion can be quite substantial without necessarily spreading throughout a network.

Results are presented in a series of five figures, each of which displays one chart for each configuration we simulated in our topology. Each chart is labeled at the top with the decimal configuration code from Table 3. The vertical axis tracks the proportion (0 to 1) of sites that are occupied by congestion, while the horizontal axis tracks the rate of injection of packets: $p = 1$ to $p = 2500$ in increments of 10 (except in cases where all routers congest or where we extend p beyond 2500). Configurations in which all routers become congested have their decimal configuration codes given in red. In such cases, a phase transition as defined in Section 3.1 is observed. If congested routers have attached sources and receivers, these sources and receivers cannot communicate across the network, and for practical purposes are also occupied.

Fig. 2 shows the growth in the proportion of total congested sites as load increases, and Section 4 1 discusses related causes. Subsequently, Section 4 2 discusses growth in the largest cluster of congested sites. Fig. 3 shows these results graphically. In eight configurations, all routers become congested. We say percolation occurs in these configurations, and a largest cluster of congested sites grows to encompass all routers. Here, evidence of a phase transition is observed. Fig. 4 shows that the largest cluster of congested sites holds among the 18 configurations, even as we increase load up to $p = 5000$.

In contrast, Fig. 5 shows decline of the largest cluster of uncongested sites and thus the breakdown of network connectivity. Section 4 3 discusses related causes for these plots. In eight configurations (the same eight where all routers congest), the largest cluster of uncongested sites ultimately drops to zero.

In Figs. 3 and 4, we see that if variable speeds (VS) or TCP are present, the largest cluster of congested sites never encompasses all sites. That is, a catastrophic event or phase transition to a global failed state does not occur if VS or TCP is present. This conclusion is supported in the appendix, where a clustering analysis of Figs. 2, 3, 5, and 6 shows the influence of VS and TCP, relative to other realistic network factors. Fig. 6 shows a decline in packets reaching their destinations, and Section 4 4 discusses the underlying causes. Other factors, including packet dropping (PD), classes (CL), sources and receivers (SR), and flows (FL), do not exhibit the same influences on network congestion, as do VS and TCP.

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

[8] Delays (DE), or propagation delays on a link, did not influence spread of congestion in our topology. The effect of DE is documented in a technical note [9]. Propagation delays would be important in any cases where they exceed queuing delays, such as inter-planetary networks or perhaps networks with links spanning continents. We also recognize that propagation delays are important in cases where there are high transmission rates on long, uncongested paths (see [17]), which did not apply in our study.

[9] Each data point was measured after simulating 200 000 time steps (after reaching steady state). We found that there was little change in the curves for each configuration if we used multiple runs. Further, some configurations took multiple weeks to run to the maximum p.

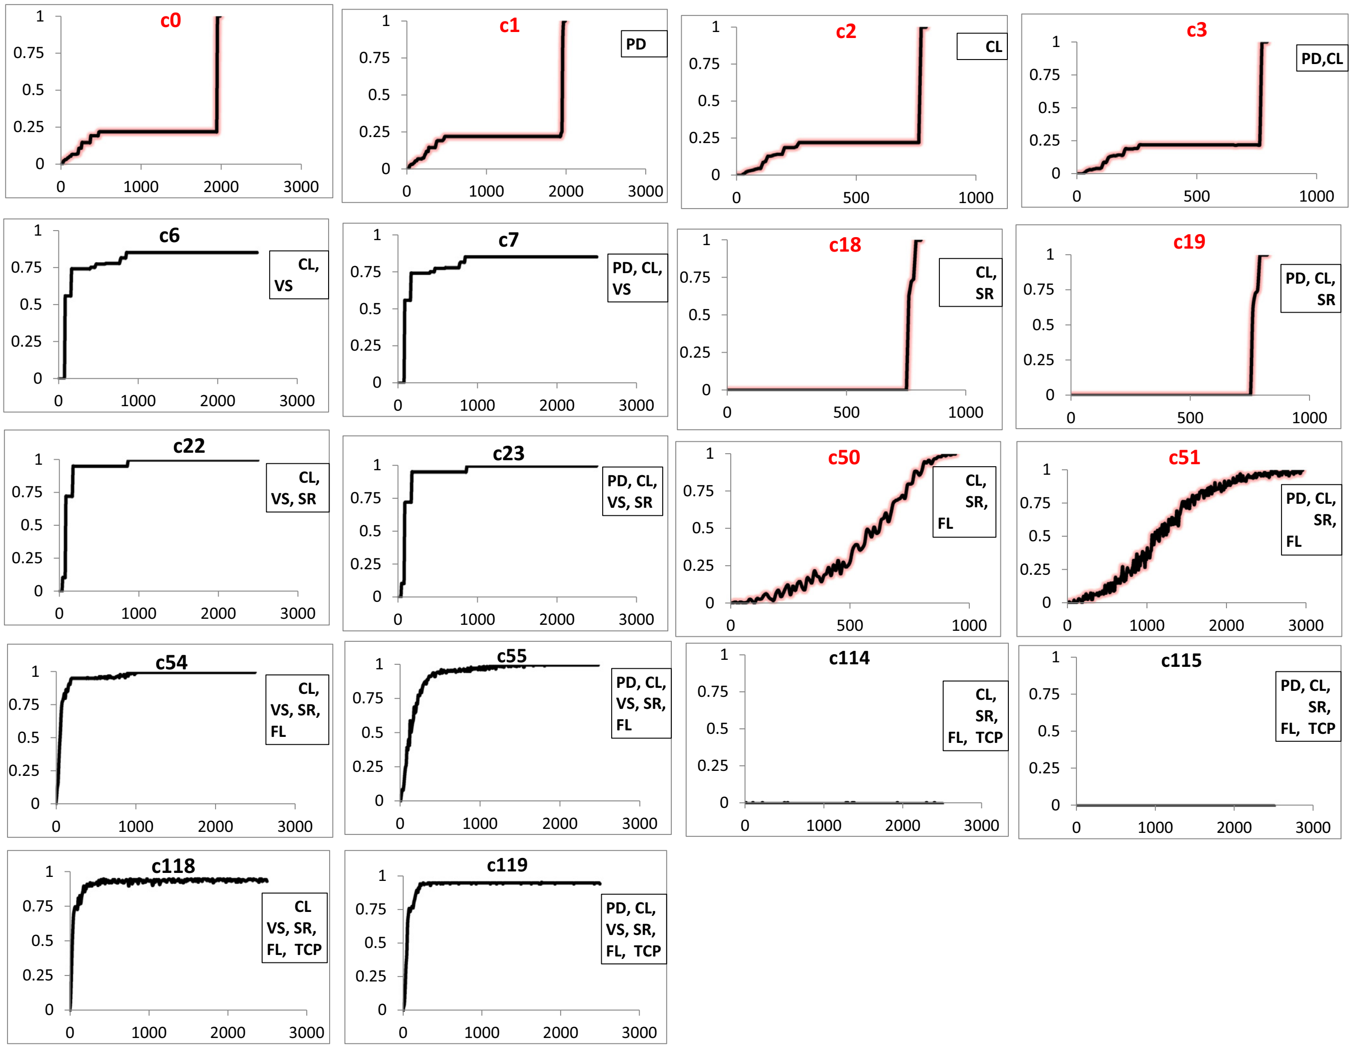

Fig. 2. Growth of the proportion of total congested sites for 18 configurations for $p = 1$ to $p = 2500$ in increments of 10 (unless all routers congest). Proportion of all sites occupied by congestion is the vertical axis (including all sources and receivers beneath a congested router, where applicable). 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 where all sites congest have their configuration numbers marked in red. In the other ten configurations, all sites do not congest. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

for configurations c18 and higher. Flows (FL) are present for c50 and above. Below c50, only individual packets are transmitted between sources and destinations. TCP is present in configurations c114 and higher. Variable speeds (VS) are present in configurations c6, c7, c22, c23, c54, c55, c118, c119.

4.1. Growth of the totals of congested sites

Fig. 2 shows the proportion of total sites that congest with respect to 18 configurations of our ISP topology for $p = 1$ to $p = 2500$ in increments of 10 (unless all routers congest). Total congested sites eventually include all routers in eight configurations, c0–c3, c18, c19, c50, c51. In these configurations, *all* 218 routers are full to a point of at least 70% of their buffer size (computed as if PD was enabled, even in cases where PD is disabled). These configurations have their configuration numbers marked in red. However, total congested sites are also high in other configurations, even where not all routers congest. When all routers do not congest, the growth in the curve for total congested sites can be sudden in configurations c6, c7, c22, c23, c54, c55, c118, and c119. Further, as explained below in Section 4 3, network connectivity breaks down in all cases.

All routers do not congest when variable speeds (VS) are present, because in Table 1, the core routers (A–P) are fastest, the PoP routers (A1–P2) are slower than core routers, and the access routers are slower than the PoP routers (except for **D**-class access routers). VS enable routers to be hierarchically engineered so that higher tiers can accommodate the largest expected packet arrival rate from lower tiers, i.e., core routers can accommodate traffic from PoP routers, and PoP routers can accommodate traffic from access routers. Therefore, all core and some PoP routers never reach their threshold in buffer queue size and do not congest, when VS is present. Hence, when VS is present in eight configurations, c6, c7, c22, c23, c54, c55, c118, c119, congestion stays at the access routers on the edges of the network, and all routers never congest. However, when VS is not present, all sites have the same speed (9 packets / *ts*). In these cases, in configurations c0–c3, c18, c19, c50, and c51, core routers congest first, because they are not fast enough to keep pace with inputs from multiple neighboring routers. Ultimately, all sites congest in these configurations.

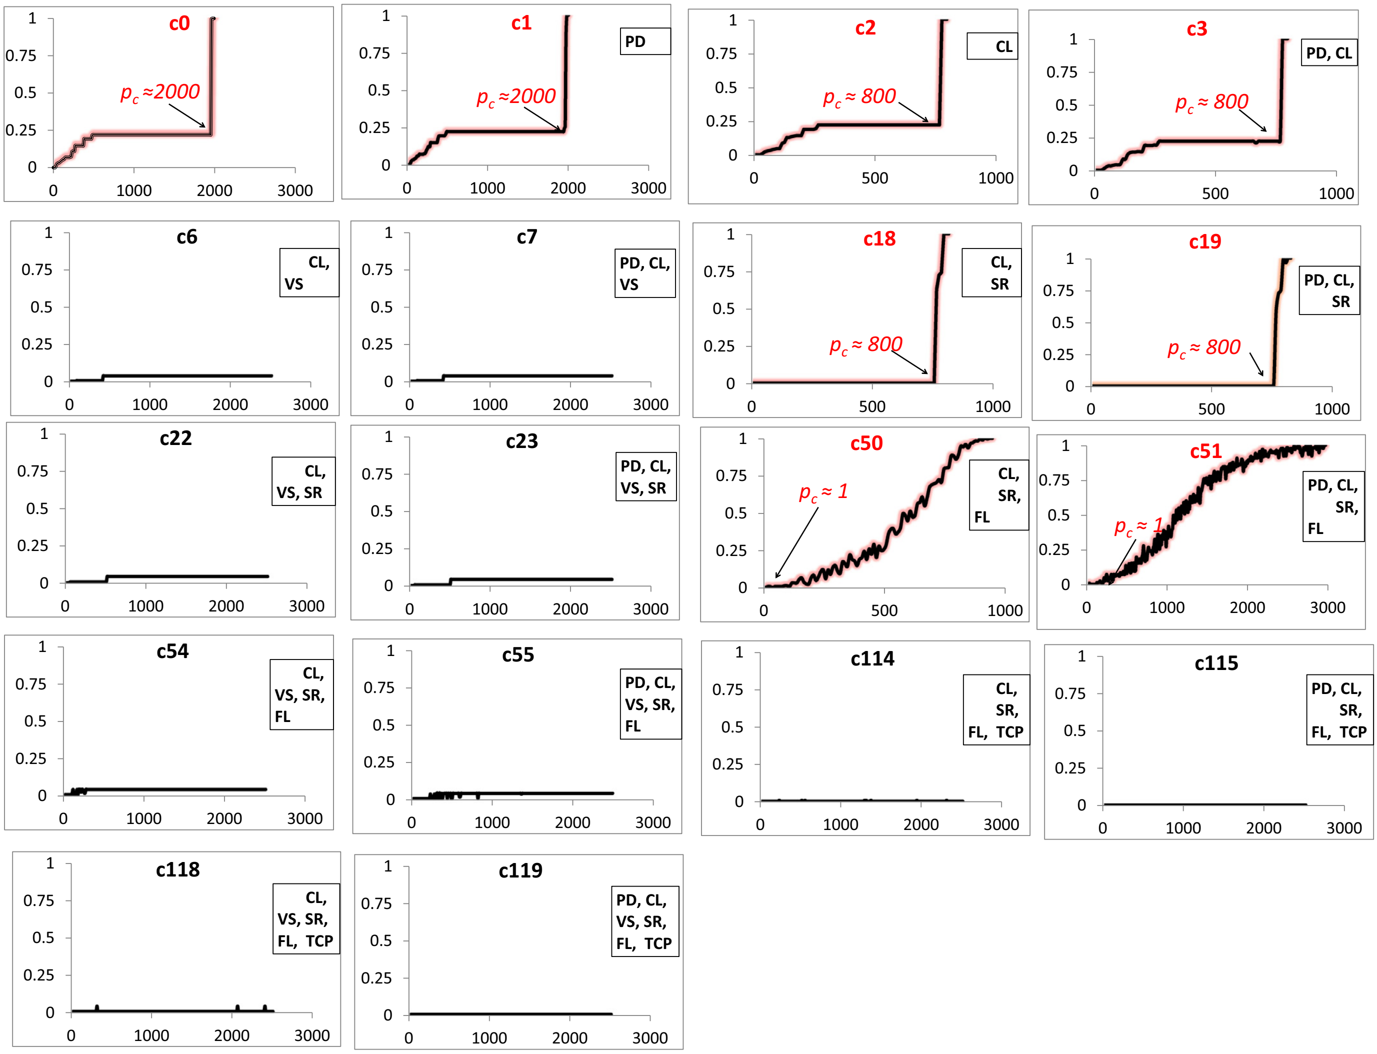

Fig. 3. Growth of the largest cluster of congested sites for 18 configurations for $p = 1$ to $p = 2500$ in increments of 10 (unless all routers congest). Proportion of all sites in the largest congested cluster is the vertical axis (including all sources and receivers beneath any congested access routers). 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 a phase transition is observed, an estimate is shown for p_c, and the configuration number is marked in red. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

In contrast, in two configurations, c114 and c115, the total number of congested sites remains small. In these configurations, TCP is present and variable speeds (VS) are not; hence without VS, all routers and all sources and receivers again have the same speed (9 packets / *ts*). However, TCP rate adaptation procedures mitigate congestion, both at core and peripheral routers. TCP lowers packet transmission rate in the face of congestion, and buffers fill more slowly. Consequently, the total number of congested sites is small. Meanwhile, in configurations c118 and c119, which have both TCP and VS, the total number of congested sites is large. In c118 and c119, because VS is present, the network is hierarchically engineered, Therefore, packets do not congest the fast core and can be forwarded to the periphery. Hence, the large periphery becomes congested, and the quarter million sources and receivers are cutoff from the network by congested access routers.

4.2. Growth of the largest cluster of congested sites

In Fig. 3, the curves show the growth of the largest cluster of sites with respect to congestion for 18 configurations for $p = 1$ to $p = 2500$ in increments of 10 (unless all routers congest). In eight of these configurations, configurations c0–c3, c18, c19, c50, and c51, with configuration numbers marked in red, evidence of a phase transition is observed, as the largest cluster of congested sites grows to encompass all 218 routers. These eight configurations have neither variable speeds (VS) nor TCP. In these configurations, p becomes greater than p_c. An estimate of p_c is provided. In these eight configurations, all sites have the same speeds. The largest cluster of congested components includes core routers, which congest first, because all traffic flows through them, followed by peripheral routers (thus isolating sources and receivers, if any). As with total congested routers described in Section 4 1, it is possible to say that congestion forms first at the center and flows from the inside out, that is, from core routers and PoP routers to access routers. Therefore, a largest cluster of congested sites grows to encompass all routers in these eight configurations. A phase transition then can be observed.

In the other ten configurations, c6, c7, c22, c23, c54, c55, c114, c115, c118, and c119, either VS or TCP, or both, are present. Hence, the largest congested cluster does not grow to include all routers, i.e., a phase transition is not observed in these configurations. In configurations, c114 and c115, TCP is present without variable speeds (VS). In c114 and c115, a phase transition is not observed,

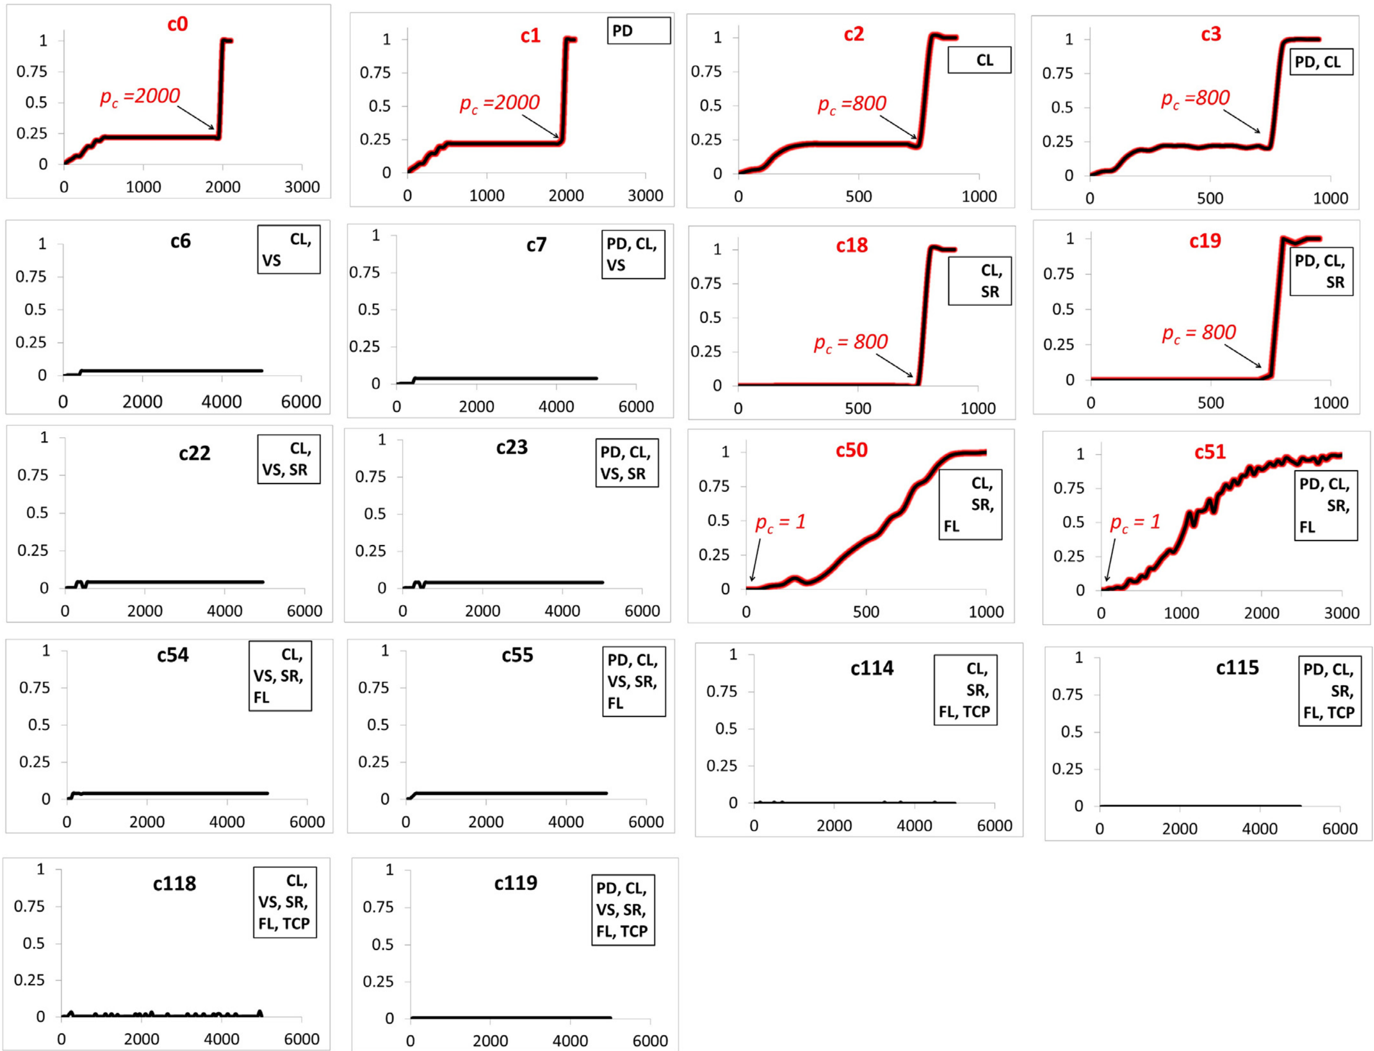

Fig. 4. Growth of the largest cluster of congested sites for 18 configurations for $p = 1$ to $p = 5000$ in increments of 50 (unless all routers congest). Proportion of all sites in the largest congested cluster is the vertical axis (including all sources and receivers beneath a congested router). The injection rate (p) is the horizontal axis. Classes (CL), 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, an estimate is shown for p_c, and the configuration number is marked in red. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

and the largest cluster of congested sites remains small. Since VS is absent in configurations c114 and c115, some core routers do congest, while the access routers at the periphery remain largely 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 prevent a phase transition. A phase transition can be observed in configurations c1, c3, c19, and c51 when PD is present, in configurations c2, c3, c18, c19, c50, and c51 when CL is present, in configurations c18, c19, c50, and c51 when SR is present, and in configurations c50 and c51 when FL is present.

In Fig. 4, p is increased up to 5000, in increments of 50. Congestion spreads as described above and as shown in Fig. 3. At increased levels of p, the largest cluster of congested sites becomes evident at the same thresholds and grows to encompass all routers. As discussed above, because of variable speeds (VS), or because of TCP, the largest cluster of congested sites does not encompass all sites when p is increased in 10 of the 18 configurations. However, individual core routers could congest if 217 of 218 routers send all their packets to a single access router, as might occur in a distributed denial of service attack. If additional ISPs are added, and they send packets at a high rate, the connecting routers between ISPs could also congest.

4.3. Decline of network connectivity—the largest cluster of uncongested sites

In Fig. 5, the curves for all 18 configurations show a precipitous decline in network connectivity as p increases toward 2500, i.e., one sees the decline of uncongested routers (and the isolation of attached sources and receivers (SR), if any). This decline in network connectivity is demarked by the decline of the largest cluster of uncongested sites. In eight configurations, c0–c3, c18, c19, c50, c51, this largest cluster of uncongested sites becomes nonexistent. These configurations have their configuration numbers marked in red in Fig. 5, and they are the same configurations in which the largest component of congested sites grows to encompass all sites in Figs. 3 and 4. Again, the presence or absence of packet dropping (PD), classes (CL), SR, and flows (FL), or combinations of them, does not determine whether the largest cluster of uncongested sites goes to zero in the eight configurations.

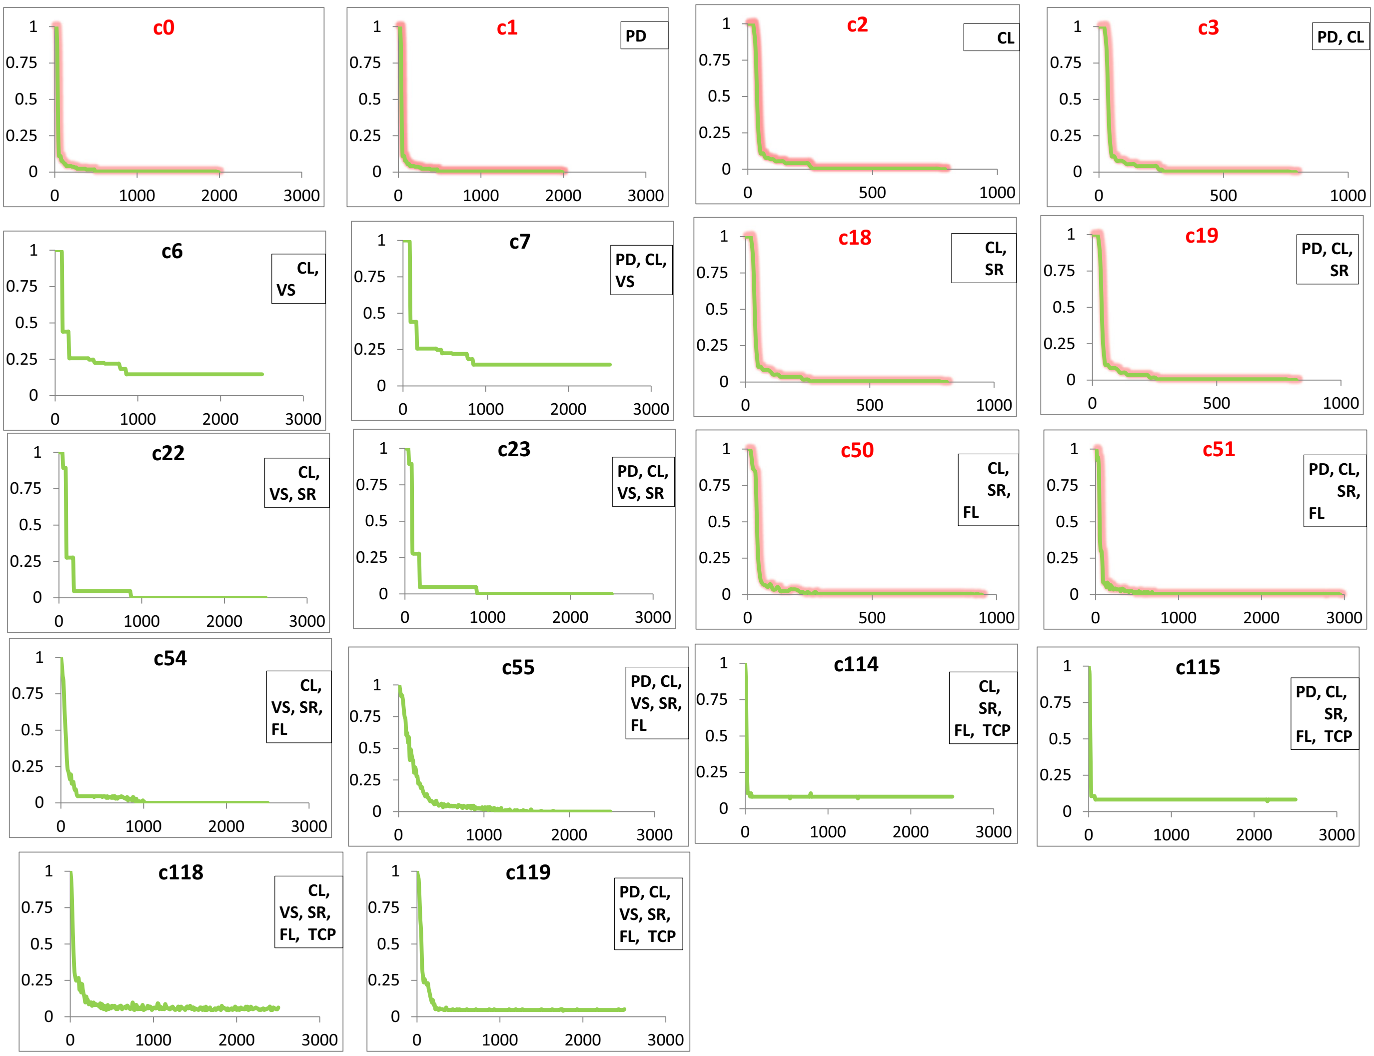

Fig. 5. Decline in uncongested sites, i.e., the breakdown of network connectivity for 18 configurations from $p = 1$ to $p = 2500$ in increments of 10 (unless all routers congest). Proportion of all sites not occupied by congestion is the vertical axis (including all sources and receiver). The injection 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. Configurations where the largest cluster of uncongested sites declines to zero have configuration numbers marked in red. Four configurations, c22, c23, c54, and c55, appear to decline to zero, but do not. In these four configurations, VS is present, and therefore the core routers remain uncongested, while the peripheral access routers can congest. In these four configurations, SR is also present, and over a quarter-million sources and receivers are isolated and are included in breakdown of network connectivity. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

As mentioned above, neither configurations c114 nor c115 have variable speeds (VS), but they do have TCP. Therefore, with this 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 fragmented clusters of uncongested sites represent diminished network connectivity. However, in all 18 configurations, Fig. 5 shows that even if variable speeds (VS) are present, network connectivity declines as the total number of congested routers rises. This means network connectivity ultimately suffers if p is raised sufficiently.

4.4. Decline in packets reaching destinations

Fig. 6, shows the proportion of injected packets reaching their destinations in 18 configurations for $p = 1$ to $p = 2500$ (unless all routers congest). The proportion of packets reaching their destinations is also shown in Table 4 (see below). In eight configurations, c0–c3, c18, c19, c50, c51, the number of packets reaching their destinations goes to zero. They are the same configurations in which a phase transition is observed in Figs. 3 and 4. These configurations have their configuration numbers marked in red.

When both TCP and VS are present together without PD, over 97% of injected packets reach their destinations in configuration c118. In configuration c119 where TCP and VS are both present with PD, over 83% of the injected packets reach their destinations. In four configurations, c114, c115, c118, and c119, which have flows (FL) and TCP, a higher proportion of packets reach their destinations than in other configurations. TCP rate adaptation has the effect of slowing flows in the face of congestion. Thus, while fewer packets are injected, a higher proportion of packets reach their destinations, as shown both in Table 4 and Fig. 6. In configurations c114 and c115, where TCP is present and VS is not, over 94% of packets reach their destinations in configuration c114, and over 75% of packets reach their destinations in c115. In effect, TCP rate adaptation improves the probability that injected packets will be delivered.

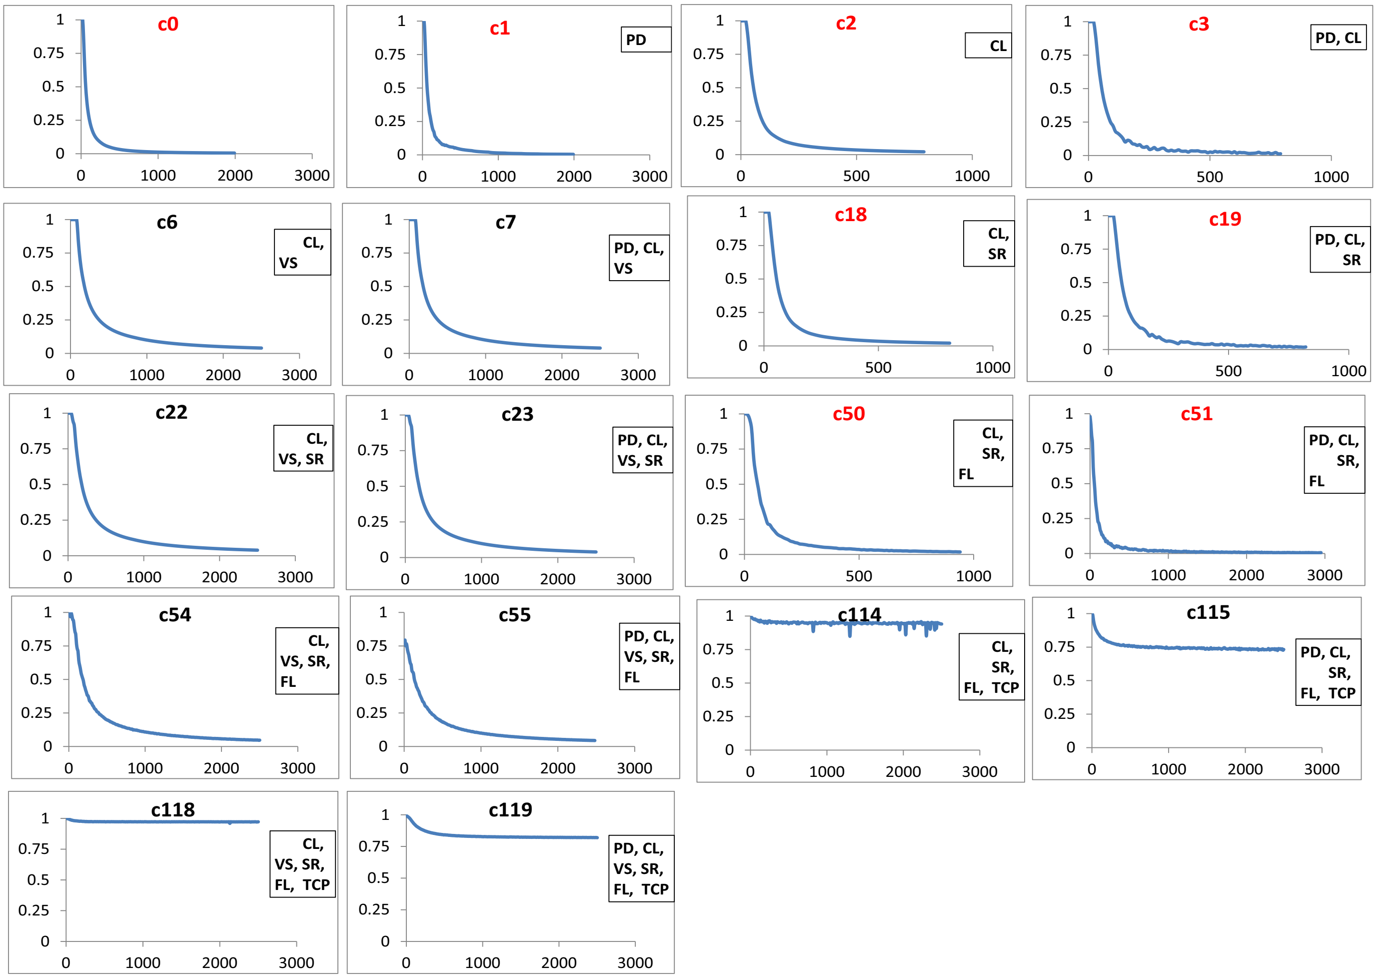

Fig. 6. Decline in packets reaching destinations for 18 configurations for $p = 1$ to $p = 2500$ (unless all routers congest). Proportion of packets reaching their destinations is the vertical axis includes all sources and receivers, if applicable. The injection rate (p) is the horizontal axis. Classes (CL), 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 marked in red. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

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

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 presence of packet dropping (PD) also makes a difference. We see that in many configurations including c114 and c118, a higher proportion of packets eventually reach their destinations without PD than with PD. 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 than with the presence of PD.

5. Discussion

The figures in Section 4 show that all sites congest, and a phase transition is observed, only when a configuration lacks both variable speeds (VS) and TCP in eight configurations. In these eight configurations, network connectivity also declines completely, i.e., the number of uncongested sites drops to zero. Yet, network connectivity declines significantly for all 18 configurations, even if a phase transition is not observed, and even if VS and TCP are both present.

The Growth in Congested Sites and the Phase Transition. In cases where all routers are operating at the same speed, core routers and then the rest of the network are overwhelmed with congestion, which spreads from the network core out to the edge. Hence, a phase transition and percolation, as both terms are defined in Section 3.1, can be seen. However, when either VS or TCP is present in ten configurations in Figs. 3 and 4, all routers do not congest, as congestion is concentrated in the network edge and not the core. When VS is present, routers are engineered with hierarchical speeds, so that the capacity of core routers in packets at least equals the packets received from PoP routers and access routers. When TCP is 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 the primary influence of these two factors. In eight configurations, which have neither TCP nor VS, the largest cluster of congested sites grows to encompass all routers. In these eight configurations, the largest cluster encompasses all routers, even these configurations have packet dropping (PD), classes (CL), sources and receivers (SR), flows (FL), or a combination of these, as Figs. 3 and 4 show.

Decline in Uncongested Sites: Connectivity Breakdown of the Network. Fig. 5 shows that network connectivity declines for all 18 configurations. These declines are complete in the eight configurations in which all routers congest; that is, in the eight configurations which have neither variable speeds (VS) nor TCP. Here, Fig. 5 shows that the largest cluster of uncongested sites, representing network connectivity, completely disintegrates precisely in the same eight configurations. In configurations c114 and c115, which have TCP and not VS, TCP rate adaptation prevents access routers at the periphery from congesting for long periods, because packet transmission slows. Thus, queues at the network periphery are also smaller. However, since all packets must flow through the core in c114 and c115, the core fragments, with some core routers congesting. Therefore, network connectivity is lost in c114 and c115. In all other configurations, the cause of the decline in network connectivity is the increase in congested routers. Proportion of total congested routers are seen in Fig. 2. Network connectivity declines in configurations c6, c7, c22, c23, c54, c55 in Fig. 5, which have VS and not TCP. In these six configurations, as total congested routers rise, the core remains uncongested. With VS, the core is faster than the periphery, and the core can transfer all packets to the periphery. Therefore, the periphery congests, and access routers at the periphery cannot communicate (nor can their sources and receivers, if any). In configurations c118 and c119, which have both VS and TCP, network connectivity also declines even though the core again remains uncongested. Congestion is again confined to the periphery, with sources and receivers cut off.

Total Packets Reaching Destinations. The causes of the decline in total packets reaching their destinations in Fig. 6 in 14 of 18 configurations are also ultimately due to growth of 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. Hence, fewer packets are injected, and the network outputs fewer packets in total. However, since TCP rate adaptation reduces congestion, a higher proportion of injected packets arrive at their destinations in these four configurations with TCP, than in the 14 configurations without TCP.

The Effect of variable speeds (VS). VS enables routers to conform to common Internet engineering principles, so that higher tiers accommodate the largest expected packet sending rate from lower tiers. In Table 1, with VS, core routers are faster than PoP routers, and PoP routers are faster than access routers (except for **D-**class access routers). Because of hierarchical engineering in configurations that have VS, fast core routers can forward all arriving packets. When heavy packet injection occurs, congestion arises at the network periphery, but the core cannot be overwhelmed and remains uncongested. Thus, with VS present, all routers do not congest, and no phase transition occurs. When VS is not present, all sites operate at the same speed. In such cases, core routers congest first, because packets arrive from multiple neighbors at a rate faster than the individual core routers can handle. Congestion then spreads toward the network edge. The presence of VS also affects queue sizes. With VS, queues are minimal for the 16 core routers and most of the PoP routers. Without VS, queues become longer in core and PoP routers.

The Effect of TCP. TCP rate adaptation modulates the injection rate of packets in response to detected congestion. By slowing packet injection, TCP reduces congestion on network paths, and thus mitigates congestion even when VS is not present. TCP interacts with PD. When PD is present, router queues are shorter than without PD. When PD is present, TCP receives feedback in the form of nacks, and first reduces packet injection rate and then increases packet injection as congestion dissipates. Thus, TCP injects more packets per time step when PD is present. When PD is absent, large queues build up in routers, and TCP receives feedback in the form of timeouts (see Section 3 .3). The large queues and timeouts cause the packet injection rate to be reduced significantly. Thus, absent PD, TCP rate adaptation provides lower network-wide throughput, even though a larger percentage of injected packets arrive at their destinations. With PD present, TCP rate adaptation provides higher network-wide throughput, but a lower percentage of injected packets reach their destinations. The congestion-mitigating effects of TCP persist even when p reaches 5000 packets / ts in Fig. 4. Even absent VS, all routers do not congest in configurations c114 and c115, which have only TCP. In these configurations, the largest cluster of congested sites remains small and never grows to encompass all routers. Yet, in both c114 and c115, network connectivity collapses, because a significant portion of core routers congest and fragment the entire network. Some core routers congest without VS, because packets from access routers cross the core, which cause queues to build in affected core routers. However, TCP then detects the congestion in the core and rate adapts, so that congestion at the periphery is reduced in configurations c114 and C115.

The Effect of Packet Dropping (PD). For the even-numbered configurations, which have no PD, a higher proportion of packets reach their destinations in Table 4. This is because without PD, packets stay in a queue which becomes very long, but packets can eventually reach their destinations. Further, in our simulations, users cannot exercise the realistic option of ending flows (FL) if flows take too long or progress too slowly. With PD in force in odd-numbered configurations, packets are dropped and must be resent. Therefore, with PD, congestion is less and queues are shorter, but a lower proportion of injected packets reach their destinations than in configurations without PD. As indicated above, PD allows TCP to receive quicker feedback and to inject more packets, and thus PD improves overall network throughput. As we saw in Section 4, phase transitions are observed with, and without, the presence of PD.

The Effect of Classes (CL), Sources and Receivers (SR), and Flows (FL). One of these realistic factors, CL, is needed to support variable speeds (VS), and the others are needed to support TCP in Table 2. These factors are also present in the Internet. However, CL, SR, and FL do not prevent congestion by themselves in Fig. 2. Without VS or TCP, they do not prevent percolation, as seen in Figs. 3 and 4. CL, SR, and FL do not prevent the breakdown of network connectivity, as seen in Fig. 5. Further, these three factors seem also not to influence the proportion of packets which arrive at their destinations in Fig. 6 and Table 4.

However, CL and FL have other, more subtle effects in configurations where the largest cluster grows to encompass all routers.

For instance, in Figs. 3 and 4, a drop in the threshold is observed from $p_c = 2000$, in configurations c0 and c1, to $p_c = 800$, in configurations c2, c3, c18, and c19. The latter four configurations have classes (CL), while c0 and c1 do not. With CL, packet injection and destination selection takes place only at the network periphery, and so packets must cross the network core, as well as inbound and outbound routers. Without CL, as in c0 and c1, packet injection can take place in any router, and the packet destination can be any other router. Packets do not necessarily have to cross the core in c0 and c1. Therefore, packet paths are longer with CL than without CL. This enhances congestion and lowers p_c.

In Figs. 3 and 4, a further drop in the threshold is observed to $p_c = 1$ in configuration c50 and c51. In c50 and c51, CL remains but flows (FL) are added. With FL, packets are injected in correlated groups and packet injection therefore is clumpier. Access router packet queues are full only when many of their sources inject packet flows; otherwise access router packet queues may be far from full. Thus, overall congestion rises more gradually and less uniformly, until the rise in p causes all peripheral routers to receive many injected packet flows. Then all routers belong to the largest congested cluster. In contrast, packets are injected individually in configurations c2, c3, c18, and c19. As p rises, outward edge nodes tend to saturate more uniformly and in near synchrony. Hence, there is a sudden upsurge in the largest cluster, and a steep phase transition occurs in c2, c3, c18, and c19. In c50 and c51, with a more gradual phase transition, a low p_c is appropriate.

The Influence of Realistic Network Factors. In the two most realistic configurations, c118 and c119, most packets in a flow reach their destination. Configurations c118 and c119 have both variable speeds (VS) and TCP. Configuration c119 has all six realistic network factors, including packet dropping (PD). Hence, it is possible to generalize that more realistic network factors enable the Internet to operate in an intelligible fashion and mitigate congestion, at least in the network core. In configurations c114 and c115, which also have flows (FL) and TCP but not VS, most packets also reach their destinations. Yet, even with the presence of realistic network factors that mitigate congestion, as congestion spreads, network connectivity still fails in all configurations, when p is raised high enough. However, including more realistic factors reduces the likelihood that a phase transition will be observed.

6. Conclusions and future work

Based on the experiment, 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 and related 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 rapidly and reaches a critical point, after which the network can become entirely congested. Catastrophic failure is then said to occur. We see that two realistic factors—variable speeds (VS) and TCP–can prevent percolation from being observed with respect to congestion; that is, either of 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 some of these papers, percolation was observed and adverse conclusions about the Internet were drawn. Computer model studies of networks often do not include widely-used Internet congestion-control protocols, which would likely impact the spreading processes and change the nature of catastrophic collapses and observed phase transitions [8].

For computer communications networks, our conclusions are consistent with those of Alderson and Willinger [1], who pointed out differences between the network topologies used in research models and the topology of the Internet. They concluded that both abstract networks and the Internet had scale-free connectivity in a graph model, but only the Internet had hub sites primarily on the periphery. Consistent conclusions were also reached by Doyle et al. [12]. Alderson and Willinger [1] also concluded that realistic behavior was also necessary in a model of the Internet, without which accurate conclusions could not be drawn. We have included realistic behavior by incorporating our six realism factors.

Willinger et al. [28] 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 cases of Internet-related modeling. 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 provide an explanatory model with respect to catastrophic congestion rather than an evocative model. Further, we are showing that the assumptions underlying selected evocative models are incorrect, and that incorporating realistic network factors leads to observed behavior that is explainable and different from extant evocative models.

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. With the exclusion of these key realistic factors, the operation of our model also experiences congestion collapse, and a phase transition is observed. However, even with the inclusion of realistic factors, all 18 configurations show a severe decline in network connectivity as the number of injected packets gradually increases, and total congested routers rises. Network connectivity declines whether, or not, a phase transition is observed.

Even when the entire network does not experience a total collapse, parts of it congest, and sources and receivers within them cannot communicate across the network. In our simulations, these portions can be identified and constitute 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 [6], models of virus and disease spread [18,20], and city automobile traffic [26]. Accordingly, models such as these should be examined to see if inclusion of additional realistic factors, or realistic topologies, are necessary to obtain a more accurate picture of network collapse.

As for future work, we recommend for development of communications networks models that include more realistic network factors, such as user patience, improved models which include sources and receivers 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 having topologies of interconnected ISP net-

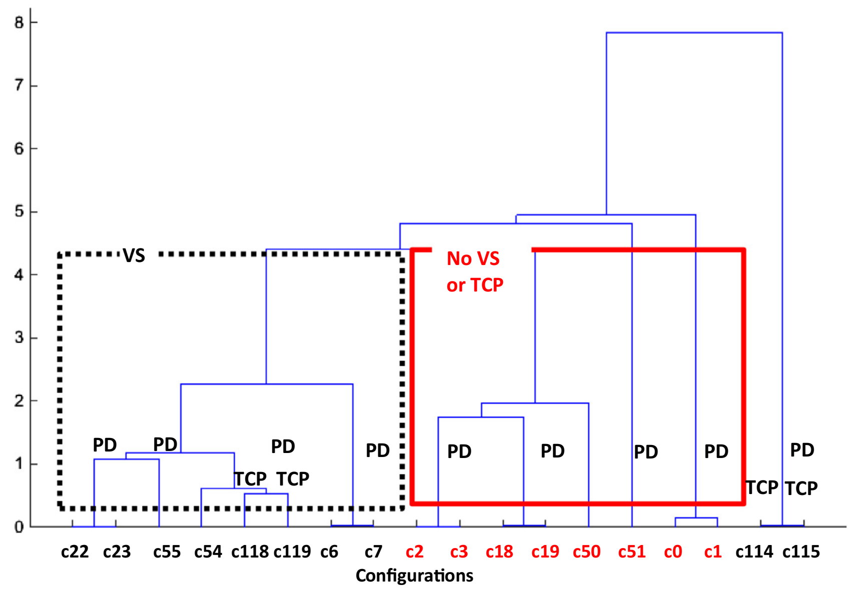

Fig. A1. Hierarchical Clustering for data in Fig. 2 shown as a dendrogram. Totals for congested sites for 18 configurations for $p = 1$ to $p = 2500$ (unless all routers congest). Configurations in which all routers congest are marked with their configuration number in red. These configurations marked have neither 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. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

works. 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.

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.

Appendix

This appendix contains four figures which show a hierarchical clustering executed with MATLAB[10] [16] for vectors of all 18 configurations in each of Figs. 2, 3, 5, and 6, from $p = 1$ to $p = 2500$ (unless all routers congest). In each figure in this appendix, 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 Figs. 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 packet dropping (PD), classes (CL), and receivers (SR), and flows (FL) are not grouped together in Figs. A1–A4. Moreover, when any combination of these four factors is present with both VS and TCP absent, then all routers 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. Yet in configuration c51, VS and TCP are both absent—and all routers congest.

In hierarchical clustering figures in this appendix, configurations which lack variable speeds (VS) are grouped together, while configurations that lack TCP are also grouped together. Configurations that lack both VS and TCP are also grouped together. Since these groupings appear in all four figures in this appendix, VS and TCP seem to be the most influential factors per hierarchical clustering. The hierarchical clustering represents a second method which supports the results given in the main body of the paper. Separate discussions of Figs. A1–A4 ensue.

Fig. A1 represents the clustering of plots from Fig. 2. In Fig. A1, 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, all routers congest. Configurations which have TCP, and not VS, 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.

Fig. A2 represents the clustering of plots from Fig. 3. In Fig. A2, 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 have neither VS nor TCP. Configurations that have neither VS nor TCP are also grouped together in a solid red box, and in these configurations, a phase transition is observed. Configurations which have only TCP are individually annotated as such. In contrast, configurations having packet

[10] Certain commercial products or company names are identified in this report to describe our study adequately. Such identification is not intended to imply recommendation or endorsement by the National Institute of Standards and Technology, nor is it intended to imply that the products or names identified are necessarily the best available for the purpose.

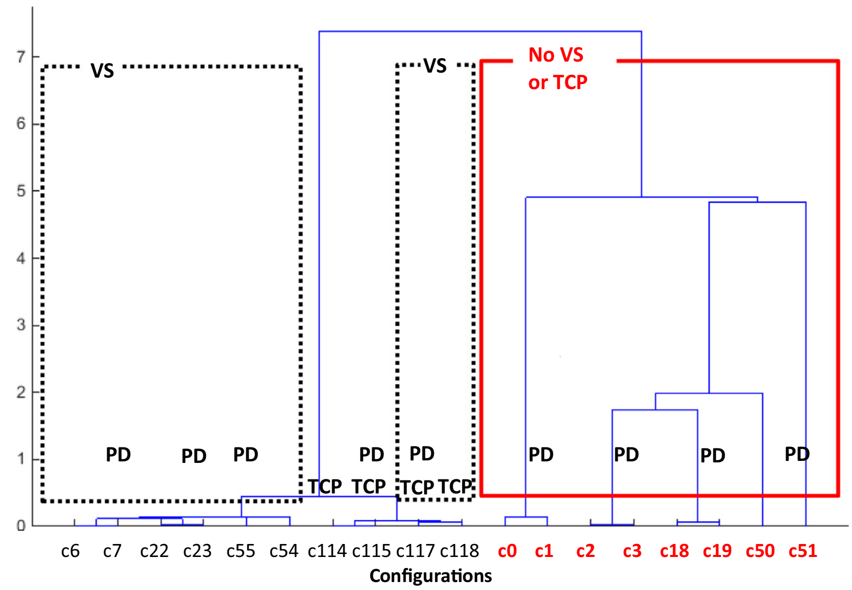

Fig. A2. Hierarchical Clustering for data in Fig. 3 shown as a dendrogram. Growth of the largest cluster of congested sites for 18 configurations for $p = 1$ to $p = 2500$ (unless all routers congest). Configurations in which a phase transition is observed are marked with their configuration number in red. These configurations have neither 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. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

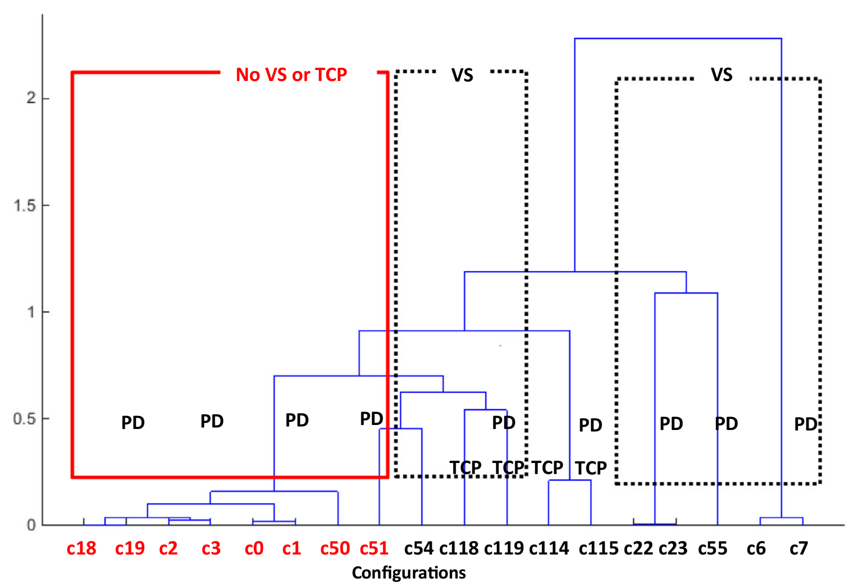

Fig. A3. Hierarchical Clustering for data in Fig. 5 shown as a dendrogram. Decline in the largest cluster of uncongested sites for 18 configurations for $p = 1$ to $p = 2500$ (unless all routers congest). Configurations in which a phase transition is observed are marked with their configuration number in red. These configurations have neither 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. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

dropping (PD), classes (CL), sources and receivers (SR), and flows (FL) seem interspersed among the above groupings, and therefore are not grouped together.

Fig. A3 represents the clustering of plots from Fig. 5. In Fig. A3, those configurations which have variable speeds (VS) can be grouped together, and a dotted black box encircles these. Those configurations that have neither VS nor TCP are also grouped together, and a solid red box encircles these configurations. In these eight configurations, a phase transition 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 Fig. A3 are generally smaller than those seen among groupings in Figs. 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 in all configurations—as Fig. 5 also shows, and as is discussed in Section. 5.

Fig. A4 represents the clustering of plots from Fig. 6. In Fig. 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 have neither VS nor TCP, and once again a solid red box surrounds these. In these eight configurations, a phase transition is observed. In contrast, configurations having packet dropping (PD), classes (CL), sources and receivers (SR), and flows (FL) are not grouped together. In Fig. A4, four configurations, c114, c115, c117,

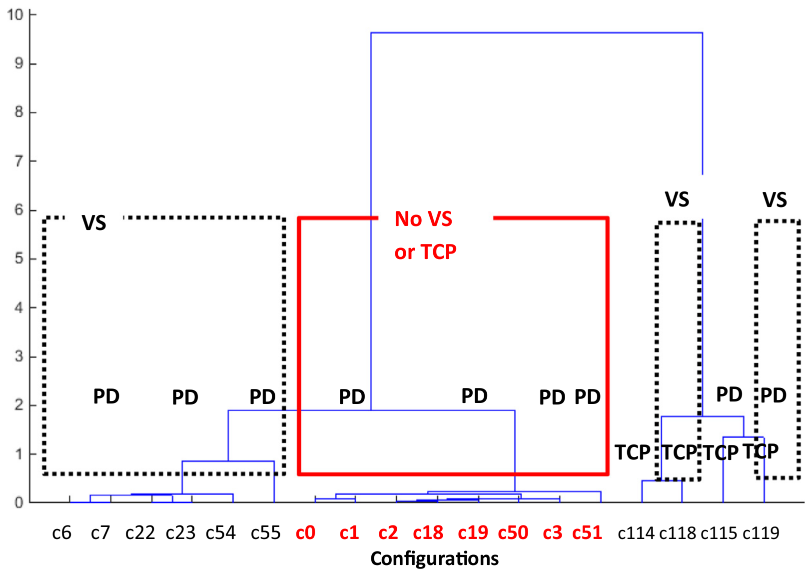

Fig. A4. Hierarchical Clustering for data in Fig. 6 shown as a dendrogram. Decline in packets reaching destinations for 18 configurations from $p=1$ to $p=2500$ (unless all routers congest). Configurations in which a phase transition is observed are marked with their configuration number in red. These configurations have neither 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. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

and c118, appear together. These four configurations all have TCP, and they also are the ones in which most packets reach their destinations in Fig. 6, as is discussed in Section. 5.

References

[1] D. Alderson, W. Willinger, A contrasting look at self-organization in the internet and next-generation communication networks, IEEE Commun. Mag. 43 (7) (2005) 94–100.
[2] D. Arrowsmith, R. Mondragón, J. Pitts, M. Woolf, in: Phase Transitions in Packet Traffic On Regular Networks: a Comparison of Source Types and Topologies, Institut Mittag-Leffler. Royal Swedish Academy of Science, 2004, p. 8.
[3] D. ben-Avraham, S. Havlin, Diffusion and Reaction in Fractals and Disordered Systems, Cambridge University Press, Cambridge, England, 2000.
[4] B. Bollobás, S. Janson, O. Riordan, The phase transition in the uniformly grown random graph has infinite order, Random Struct. Algorithms 26 (1-2) (2004) 1–36.
[5] D. Callaway, M. Newman, S. Strogatz, D. Watts, Network robustness and fragility: percolation on random graphs, Phys. Rev. Lett. 85 (25) (2000) 5468–5471.
[6] B. Carreras, V. Lynch, I. Dobson, D. Newman, Critical points and transitions in an electric power transmission model for cascading failure blackouts, Chaos 12 (4) (2002) 985–994.
[7] R. Cohen, K. Erez, D. ben-Avraham, S. Havlin, Resilience of the internet to random breakdowns, Phys. Rev. Lett. 85 (21) (2000) 4626–4628.
[8] C. Dabrowski, The study of catastrophic event phenomena in communication networks, Comput. Sci. Rev. 18 (2015) 10–45.
[9] C. Dabrowski, K. Mills, in: The Influence of Realism on Congestion in Network Simulations, National Institute of Standards and Technology (NIST), 2016, p. 62. Technical Note 1905.
[10] D. De Martino, L. Dall'Asta, G. Bianconi, M. Marsili, Congestion phenomena on complex networks, Phys. Rev. E 79 (2009) 015101-1–015101-4.
[11] S. Dorogovtsev, A. Goltsev, J. Mendes, Critical phenomena in complex networks, Rev. Mod. Phys. 80 (4) (2008) 1275–1445.
[12] J. Doyle, D. Alderson, L. Li, S. Low, M. Rougan, S. Shalunov, R. Tanaka, W. Willinger, The "robust yet fragile" nature of the internet, Proceedings of the National Academy of Sciences 102 (41) (2005) 14497–14502.
[13] P. Echenique, J. Gómez-Gardeñes, Y. Moreno, Dynamics of jamming transitions in complex networks, Europhys. Lett. 71 (2) (2005) 325–331.
[14] J. Henriksen, An introduction to SLX™, in: Proceedings of the 28th Winter Simulation Conference. Coronado, California, USA, 1996, p. 468–475. December 08–11.
[15] A. Lawniczak, P. Lio, S. Xie, J. Xu, Study of packet traffic fluctuations near phase transition point from free flow to congestion in data network model, in: Canadian Conference on Electrical and Computer Engineering, CCECE 2007. Vancouver, Canada, 2007, pp. 360–363. April 22–26.
[16] *MATLAB v.* 2013, *Mathworks*. https://www.mathworks.com/help/matlab/.
[17] K. Mills, J. Filliben, D. Cho, E. Schwartz, D. Genin, in: Study of Proposed Internet Congestion Control Mechanisms, National Institute of Standards and Technology (NIST), 2010 Special Publication 500-282, 534 pages.
[18] Y. Moreno, R. Pastor-Satorras, A. Vespignani, Epidemic outbreaks in complex heterogeneous networks, Eur. Phys. J. B 26 (2002) 521–529.
[19] G. Mukherjee, S. Manna, Phase transition in a directed traffic flow network, Phys. Rev. E 71 (2005) 055108-1–055108-7.
[20] R. Pastor-Satorras, A. Vespignani, Epidemic dynamics and endemic states in complex networks, Phys. Rev. E 6 (2001) 066117-1–066117-8.
[21] Y. Rykalova, L. Levitin, R. Brower, Critical phenomena in discrete-time interconnection networks, Phys. A 389 (22) (2010) 5259–5278.
[22] S. Sarkar, K. Mukherjee, A. Ray, A. Srivastav, T. Wettergren, Statistical mechanics-inspired modeling of heterogeneous packet transmission in communication networks, IEEE Trans. Syst., Man, and Cybern.—Part B 42 (4) (2012) 1083–1094.
[23] R. Solé, S. Valverde, Information transfer and phase transitions in a model of internet traffic, Phys. A 289 (3–4) (2001) 595–605.
[24] B. Tadić, G. Rodgers, S. Thurner, Transport on complex networks: flow, jamming and optimization, Int. J. Bifurcation Chaos 17 (7) (2007) 2363–2385.
[25] D. Wang, N. Cai, Y. Jing, S. Zhang, Phase transition in complex networks, in: American Control Conference. St. Louis, Missouri, USA, 2009, pp. 3310–3313. June 10–12.
[26] F. Wang, D. Li, X. Xu, R. Wu, S. Havlin, Percolation properties in a traffic model, Europhys. Lett. 112 (3) (2015) 38001-1–38001-6.
[27] W. Wang, Z. Wu, R. Jiang, G. Chen, Y. Lai, Abrupt transition to complete congestion on complex networks and control, Chaos 19 (3) (2009) 033106-1–033106-7.
[28] W. Willinger, R. Govindan, S. Jamin, V. Paxson, S. Shenker, Scaling phenomena in the internet: critically examining criticality, Proceedings of the National Academy of Sciences of the United States of America 99 (3) (2002) 2573–2580, Supplement 1.
[29] M. Woolf, D. Arrowsmith, R. Mondragón, J. Pitts, S. Zhou, Dynamical modeling of TCP packet traffic on scale-free networks, in: Institut Mittag-Leffler., Royal Swedish Academy of Science, 2004, p. 9.
[30] Z. Wu, W. Wang, K. Yeung, Traffic dynamics in scale-free networks with limited buffers and decongestion strategy, New J. Phys. 10 (2008) 023025-1–023025-12.