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].