This paper was presented as part of the Workshop on Network Science for Communication Networks (NetSciCom)

Non-Binary Information Propagation: Modeling BGP Routing Churn

Nicholas C. Valler*, Michael Butkiewicz*, B. Aditya Prakash†, Michalis Faloutsos* and Christos Faloutsos†
*Department of Computer Science and Engineering, University of California - Riverside
Email: {nvaller, butkiewm, michalis}@cs.ucr.edu
†Computer Science Department, Carnegie Mellon University
Email: {badityap,christos}@cs.cmu.edu

***Abstract*—In this work, we generalize the commonly-used "binary" (or categorical) information propagation model to describe the propagation of a continuous-value node-property in a network. Most efforts so far focus on discrete states for nodes (i.e. *healthy, sick*). Here, we extend the above model to describe the propagation of a node property that is characterized by a real value. As a case study, we focus on routing messages at the Internet backbone (BGP level), which we refer to as routing instability or churn. Our goal is to develop the simplest possible model that can characterize the propagation of routing instability. To capture an important routing property (routing policies), we enrich the model in a non-trivial way. Varying our small set of model parameters, we show that our model can exhibit a wide range of behaviors, from fast "die-out" to non-zero steady-state and oscillations. To the best of our knowledge, this is the first work that casts routing as a network-wide propagation problem and sets the stage for a theoretical analysis of routing instability, and the propagation of non-binary node properties in general.**

I. Introduction

How can we extend the binary/categorical epidemic propagation model to describe propagation of a non-binary entity in a network? So far most work focuses on binary epidemic models, where a node is either healthy or infected, though some models provide additional discrete states, such as incubation, removed, or immune. However, many phenomena are non-binary and they cannot be adequately described by the above models (e.g., models describing product marketing, product adoption, rumors or public opinion). The key difference is that now the state is defined by a strength and can take many different values: one's excitement about the Android phone could be a cumulative number of the excitement one receives from one's friends.

To ground our work, we focus on a case-study: the propagation of routing updates in Internet routing at the backbone. The Border Gateway Protocol (BGP) [1] is the Internet's *de facto* inter-domain routing protocol, facilitating the sharing of routing information between independent administrative domains (i.e., ISPs or other large enterprise networks), called Autonomous Systems (or ASes). Each AS has one or more BGP routers that maintain routing information describing how to reach groups of IP addresses (IP prefixes). BGP routers maintain routing information dynamically by exchanging **routing updates or messages** with their neighbors. Upon receiving an update message, a BGP-router: (1) evaluates the update against its routing policies; (2) modifies, as necessary, its routing table according to routing policies; and (3) generates and propagates a new update to all (or a subset) of its neighbors. Which neighbors receive the update depends, again, on specific routing policies formed according to business relationships that exist between the neighboring ASes (we discuss this more in §III). Routing updates can be seen as a measure of **routing instability**, and are often referred to as **churn**, and we use these terms interchangeably here.

In this paper, we focus on developing the simplest possible model that can characterize the propagation of routing instability, and incorporate in the model essential routing features (policy awareness). To the best of our knowledge, this is the first work that casts routing as a network-wide propagation problem and sets the stage for a theoretical analysis of routing instability. From a theoretical point of view, our work can be seen as a first attempt to explore the outcome of two competing phenomena: the "attrition of churn in every propagation hop" that wants to bring the system to an equilibrium (zero instability), and the multiplication of churn by the propagation of the instability to all adjacent nodes (one incoming update can generate multiple outgoing updates). Our contributions can be summarized in the following points.

a. A Model of Non-Binary Information Propagation: We develop an analytical framework, focusing on the simplest possible non-linear dynamic system, capable of characterizing the propagation of non-binary information, as exemplified in such systems as BGP instability propagation. We extend the power of this model through a simple, yet effective, graph transformation that incorporates BGP policies in the propagation of updates, which is a critical aspect for a theoretical BGP model.

b. Exploring the descriptive power of the model: The key observation is that the model has excellent descriptive power. It can generate a large number of behaviors, many of which have been observed in practice. We vary the key parameters of the model and show the operational regime to which the system is driven. We show that the model can exhibit: (a) quick die-out of the churn, (b) stabilization to non-zero **steady-state churn**, and (c) oscillating behavior around a non-zero steady-state. In all these cases, we start with only one initial excitation to the system, so the fact the system is stabilized to a non-zero value indicates that the system is unable to "get rid" of the churn,

978-1-4244-9920-5/11/$26.00 ©2011 IEEE

but falls into a perpetual self-sustaining mode.

Conjectures and future work: As a first step in a new area, our work creates more questions than answers.

Conjecture 1: We conjecture that the non-zero churn stabilization is an intrinsic topological property, for a given set of model parameters. The conjecture is based on our initial experiments, where churn seems independent of the initial excitation. For example, a subsequent excitation adds churns momentarily but the system returns to the same non-zero churn level as before.

Conjecture 2: We also conjecture that oscillations (period and amplitude) are intimately related to topological properties, and are revealed for particular values of the parameters of our model.

Future work in this direction can focus on: (a) a more extensive and detailed understanding of the observed behaviors via simulations, and (b) theoretical work that could relate eigenvalues of the adjacency matrix and the model parameters. Note that eigenvalues have already been shown to summarize effectively topological properties in other propagation problems[2].

The rest of this paper is organized as follows. We present related works in Section §II, followed by the description of our model in Section §III. A simulation-based evaluation of our model and discussion of results are presented in Sections §IV and §V. Finally, we conclude and present future work in Section §VI.

II. Related Works

In this section, we briefly cover works related to our own.

Modeling and Measurement of BGP Instability. Measurement studies dominate the academic landscape of BGP instability research. Earlier work by Prakash et.al. [3] provides motivation for this work. The authors identify a number of interesting features in BGP routing updates seen as a time series. Specifically, they find that churn exhibits chaotic features, and this has guided our selection of a non-linear dynamic system as central to our model (described in §III). Other BGP measurement studies informing our model include the topics of: (a) general BGP dynamics (e.g., [4], [5]), (b) catastrophic BGP events (e.g., [6], [7]) and (c) robustness and convergence properties of BGP (e.g., [8])

Modeling studies of BGP churn instability are few and far between, but excellent work has been done by [9], [10], [11]. In particular, Coffman et. al. [11] propose a model to study the cascading features of BGP instability. Cascading features are often seen in information propagation systems.

Information Propagation. Information propagation is analogous to disease or epidemic spreading processes described in epidemiological literature. For a gentle introduction to the subject, refer to [12]. Early propagation models assumed a *homogeneous* population; that is, a population with no social or spatial structure. Newer models apply underlying structure to the population, creating a *heterogeneous* population [13],[14]. We too apply a network structure to our model, deriving topological structure from general models of the Internet as detailed in [15]. These topologies retain the general statistical properties of the Internet, yet are of small enough size to make simulation studies feasible.

Our work differs from the referenced material in one distinct manner. We do not view information as a discrete quanta of information, rather, we model information as a continuous range of values, which in aggregate, represent the whole of the propagating information.

III. Our Model

In this section, we describe our model of propagation and show how it can describe BGP routing.

Intuition and Motivation. BGP routers exchange routing information via routing updates, and at any given time, a router can be thought of as having a number of such messages that are to be sent to its neighbors. Upon receiving a message a router decides if the message provides information that changes its routing table. This leads to two cases.

Case 1: If the message does not change the routing table, the message is not propagated further, i.e., we can say that the update stops there.

Case 2: If the message changes the routing information of the router, then the router will propagate the message to all or a subset of its neighbors depending on the relationships between the ASes that the routers belong to (more on this later in this section). In general, we can think of the message being propagated to multiple neighbors.

These two cases represent two competing mechanisms, (a) the termination of an update in case 1, and (b) the multiplication of the routing messages of updates in case 2. Our work can be seen as the study of these two competing mechanisms.

Determining the Propagation Function. A critical part of our model is how the updates residing in one router are propagated to a neighbor. We use a function that we call the **propagation function** to quantify this. In [3], the authors observe burstiness, periodicity and self-similar behavior in real traces of BGP updates, and this led us to the following choice of function inspired by the so-called logistic equation, which has been used widely in dynamic systems [16]. Let $F(x_t)$ be the amount of churn that a node with current churn x_t will send to a neighbor node:

$$F(x_t) = \frac{r \times x_t \times (M - x_t)}{M} \tag{1}$$

where, t is time, x_t the number of updates, and r and M are parameters of the system that we will discuss more later.

Framing a Dynamic Control System. We view our network as a dynamic system with feedback: routers send updates to neighbors, which then potentially cause the neighbors to send them back more updates. This feedback depends on the topology of the network.

To explore a connection with dynamic systems theory, we consider the following dynamic system which is often referred to us **logistic map**:

$$x_{t+1} = \frac{r \times x_t \times (M - x_t)}{M} \tag{2}$$

which corresponds to the trivial case of one node feeding churn to itself. Although this is an unrealistic scenario, it gives us some very interesting intuition and provides connections to analyzing our system using control theory. As shown in Figure 1(a,b,c), the logistic map displays a variety of behaviors, depending on the values of r. Specifically, we observe three interesting behaviors: (a) for $0 < r < 1$, the system goes to zero. (b) for $1 < r < 3$, the system converges to a non-zero fixed point, and (c) for $3 < r < 4$, the system oscillates around a non-zero fixed point. Furthermore, M is a dampening factor, such that as $x_t \to M$, $(M - x_t)/M$ approaches 0 and $x_t \to 0$, $(M - x_t)/M$ approaches 1. Intuitively, M is effectively a bound of how much churn can be propagated in one time step, and we will discuss this issue and the importance of these parameters later in this paper.

The simplicity and wide variety of behaviors makes this function a reasonable and promising choice, since, as we can see in Figure 1(d, e, f), our simulated behavior of large systems can mirror those behaviors, as we discuss later.

One nice feature of our propagation function is that, for appropriately selected values of r, x_t will map to a given range $(0, M)$. Note that this function can produce negative values, if $x_t > M$, which are not meaningful in our case, and thus we eliminate them, as we discuss in Section IV.

Our Model. The topology of our network can be described by a graph $\mathcal{N} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{v_1, v_2, \ldots, v_N\}$ is a set of $|\mathcal{V}| = N$ vertices and $\mathcal{E} = \{(i_0, j_0), (i_2, j_2), \ldots, (i_M, j_M)\}$. Without additional labeling, we can see that $\mathcal{N}$ strictly describes topological, or connectivity, information. The intuition behind our model is that, at time $t+1$, each node $i \in \mathcal{V}$ receives an aggregated quantity of churn $S_{i,t+1}$, accumulated from each of i's neighbors $j|(i,j) \in \mathcal{E}$ and a "retained" quantity $\delta S_{i,t}$ contributed by i itself. The retained churn can be thought of as the amount of churn i that did not transmit at the previous time step time t.

Given $\mathcal{N}$, we define our network propagation model as follows:

$$S_{i,t+1} = \delta_i S_{i,t} + \sum_{j|(i,j)\in\mathcal{E}} \frac{(1-\delta_j) h_j S_{j,t} (M_j - S_{j,t})}{M_j} \quad (3)$$

Explanation of symbols: $S_{i,t}$ is the quantity of churn at vertex i at time t. The parameters δ and h are the **retention** and **transmission parameters**, respectively. $\delta_i S_{i,t}$ is the amount of churn that is not sent to its neighbors, but is kept at vertex i for time $t+1$. The propagation function, as discussed above, appears within the summation and defines, how much churn neighbor j contributes to node i.

Relating Parameters to Real BGP Behavior. Parameter δ represents "memory" or the "stickiness of updates" in the system: in every time step, a node i defers for later the propagation of δ_i of its current churn. Similarly, the amount of churn contributed by node j to node i is $(1 - \delta_j)$.

Parameter h_j quantifies what percentage of the arriving churn will create outgoing churn. In other words, it represents **Case 1** that we saw in the beginning of this section, which is a realistic behavior in BGP.

Algorithm 1 $\mathcal{N}(\mathcal{V}, \mathcal{E}, \mathcal{L}) \to \mathcal{Q}(\mathcal{V}', \mathcal{E}')$ Transform

for edge $e = (i, j) \in \mathcal{E}$ **do**
  if $q_n \notin \mathcal{V}'$, where $n = \{i|j\}$, **then**
    $add\_node(q_n)$ to $\mathcal{V}'$
  end if
  $add\_node(q_{i\to j})$ to $\mathcal{V}'$
  $add\_edge(q_i, q_{i\to j})$ to $\mathcal{E}'$
  $add\_node(q_{j\to i})$ to $\mathcal{V}'$
  $add\_edge(q_j, q_{j\to i})$ to $\mathcal{E}'$
end for
for queue $q_{i\to j} \in \mathcal{V}'$ **do**
  $add\_edge((q_{i\to j}, q_j))$ to $\mathcal{E}'$
  if $l_{i\to j} \in \mathcal{L}$ is 'P→C' **then**
    for $q_{j\to x}$ such that $l_{j\to x}$ = 'P→C' **do**
      $add\_edge((q_{i\to j}, q_{j\to x}))$ to $\mathcal{E}'$
    end for
  end if
  if $l_{i\to j} \in \mathcal{L}$ is 'C→P' **then**
    for $q_{j\to x}$ such that $x \neq i$ **do**
      $add\_edge((q_{i\to j}, q_{j\to x}))$ to $\mathcal{E}'$
    end for
  end if
  if $l_{i\to j} \in \mathcal{L}$ is 'P↔P' **then**
    for $q_{j\to x}$ such that $l_{j\to x}$ = 'P→C' **do**
      $add\_edge((q_{i\to j}, q_{j\to x}))$ to $\mathcal{E}'$
    end for
  end if
end for

Considering policy-aware routing. We explain how, given an initial AS-level connectivity graph $\mathcal{N}$, we create a new, directed, graph, $\mathcal{Q}$, that incorporates BGP policies in its topology. To achieve this, we transform every router in $\mathcal{N}$ into a group of queues (nodes in $\mathcal{Q}$) which each queue forwarding churn to a specific neighbor.

In a nutshell, BGP update propagation is governed by peering policies, the effect of which is that not all received updates are propagated to all neighbors. The distinction is made based on the business relationship between nodes: customer-provider (C→P), provider-customer (P→C) or peer-to-peer (Peer↔Peer) [1].

The policy rules are simple:

1) C→P: an update message received from a customer is propagated to all the neighbors of a provider.
2) P→C or Peer↔Peer: an update message received from a Provider or a Peer is propagated only to customers of the receiving node.

Each vertex $i \in \mathcal{N}$ corresponds to a group of vertices in $\mathcal{Q}$, and we refer to the vertices in $\mathcal{Q}$ as queues. For each neighbor of node i we create one outgoing queue from i to that node. In this manner, policies can be encoded directly into the structure of the directed graph and used to determine in which queues the incoming churn should added to.

In addition, $\mathcal{Q}$ has an **internal queue** q_n for each original

vertex n in $\mathcal{N}$. The internal queue accounts for the ability of ASes to generate churn from within itself (i.e., from, say, the local network, local changes for routing, new IP addresses etc).

The propagation of churn between queues is still defined by equation 3 applied on the $\mathcal{Q}$.

In Algorithm 1, we explain how we do the transformation from $\mathcal{N}(\mathcal{V}, \mathcal{E}, \mathcal{L})$ where the policy set $\mathcal{L}$ is a type of the peering relationship between any two vertices of $\mathcal{N}$. These relationships are used in lines 10-27 to determine the connectivity between the internal queues and other nodes. The resulting "Queue" graph $\mathcal{Q}$ encodes all the adjacency information contained in $\mathcal{N}$, yet adds flexibility to send different churn to different neighbors.

IV. Simulation and Evaluation

In this section, we present the results of a simulation-based study of our model. All simulation experiments were conducted on 4× Xeon 2.53GHz quad-core processors with 72 GB of RAM, under the CentOS 5.5 operating system. The simulator was based in Python 2.6, using the popular mathematical libraries NumPy (v1.4.0), SciPy (v0.7.1), and NetworkX (v1.4).

Topology. We use a set of sampled versions of the Internet topology [15], which are state of the art topologies that are both realistic and have BGP policies. Clearly, given our model, we wanted topologies with associated policies. Unless otherwise noted, the results presented below were generated using a 600 node reduced AS topology, but simulations were selectively repeated on 1,500 and 2,500 node topologies as well with qualitatively similar results. For each initial topology, we used our transformation from Section III to obtain a transformed topology of queues.

Assumptions. We assume the network to be homogeneous: we use the same values for h and δ. We only differentiate between internal queues and non-internal queues. The h and δ parameters for all internal queues have been fixed at 0.05 and 0.0, respectively, as low values of internal queues allow us to focus on the effect of the network connectivity between ASes.

Our model can generate a wide range of behaviors with appropriately tuned parameters. Simulations have been conducted to $10,000$ time steps and all observed behaviors are present at the end of these runs. Figures 1(d), 1(e), 1(f) show the results of the total system churn over time with significantly different behaviors.

We also plot heat-maps to explore the effect of the different model parameters. Specifically, we examine the effect of δ (x-axis) and h (y-axis) on the magnitude of the steady-state churn in Figure 1(g), the amplitude of the period in Figure 1(h), the period in Figure 1(i) The darker values imply higher values of the observed metric. White values mean zero or non-existent such behavior.

Below we provide a more in depth discussion of our results.

The system reaches a steady-state for a wide range of parameters. We say that the system reaches a steady state when the system churn reaches a fixed point. This point can be either zero, as seen in Figure 1(d) when the h value is very small, or to a non-zero fixed point at higher h values as in Figure 1(e).

Exhibiting a periodic behavior for a wide range of parameters. In some cases, the system reaches a persistent periodic behavior as shown in Figure 1(f). This behavior is much less common, however, as can be seen by the small number of non-white boxes in Figure 1(i) that correspond to combinations that lead to periodic behavior.

Studying spikes and explosive behavior. One frequently observed behavior in BGP are spikes caused by router restarts. When a BGP speaking router reconnects to a neighbor, both will exchange routing tables. This causes BGP update spikes in both the restarted router and its neighbors. To initiate a queue restart in the simulation we inject a large amount of churn into the internal queue associated with the restarting queue's AS. The internal queue propagates this churn to all queues in the AS, as the graph transform has given it outgoing connects to each, simulating the queue reconnection to neighbors in the same AS. This churn continues to propagate outward, simulating the spreading of the restarted queue's default routing table, and subsequently back inwards, simulating the receiving of reconnected neighbor's routing tables. We see this desired result in Figure 2. Churn is injected into the restarting queue's AS at time $t = 50$ causing a small spike. At time $t = 51$ it leaves the AS and starts a spike in its neighboring ASes. This churn returns to the original AS at time $t = 52$ causing a second spike. We find it fascinating that **the model is able to significantly amplify the initial artificial spike, and then conversely reestablish it's original operating behavior so rapidly**.

We observe explosive behavior indirectly: when the churn of a queue becomes very large, larger than M, the propagation of churn is "capped" to zero (as explained earlier). However, when this capping function kicks in, it is an indication of a system reaching a highly active mode of operation. We see the artifact of this in the high initial peak shown in Figures 1(f) and 1(e). When the transmission of enough queues are reduced by our limiting of $S_{j,t}$ to M, the system forms the initial peak as it declines to find a stabilization point. Our study on when and how these "capping" events manifest themselves could not be incorporated in this paper due to space limitations.

The system is robust to large external or initial excitations. The behavior of the model is robust to the amount of initial churn. For example, the initial churn affects only the time it takes to reach the steady state in the scenarios we studied. Furthermore, external spikes destabilize the system temporarily, and then it returns to its previous steady-state churn level as shown in Figure 2.

Understanding the effect of M, h and δ. We attempt to understand how each parameter affects the performance of the system.

a. M has only a scaling effect on the system, but does not affect the type of the behavior otherwise. For this reason, and due to space limitations, we do not present any results

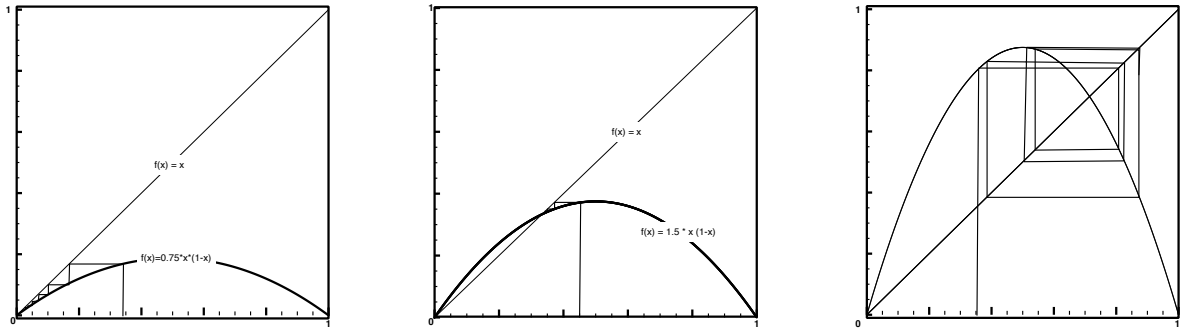

(a) Logistic Map, $r < 1$, convergence to the fixed point at 0.

(b) Logistic Map, $1 < r < 3$, convergence to the fixed point > 0.

(c) Logistic Map, $3 < r < 4$, oscillations behavior "around" the fixed point > 0.

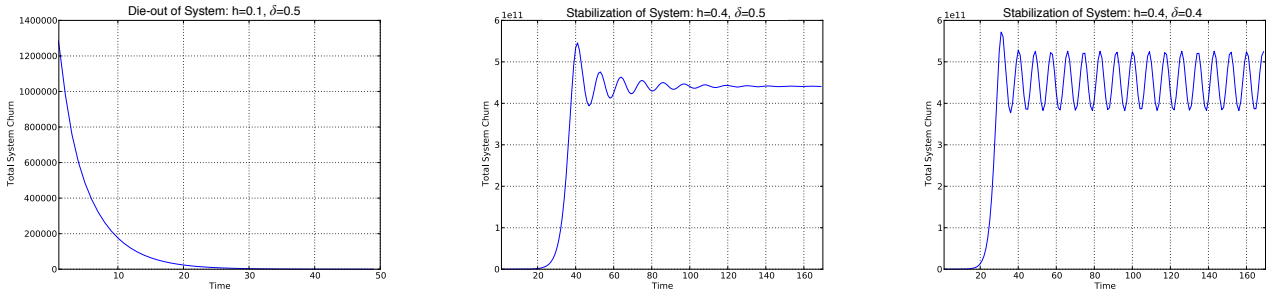

(d) System churn goes to zero

(e) System churn reaches non-zero steady state

(f) System oscillates

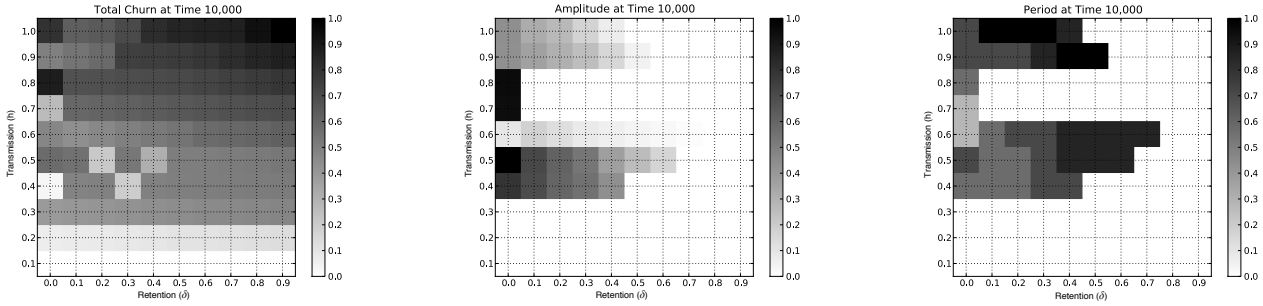

(g) Steady-state of System Churn

(h) Amplitude of System Churn: white means zero amplitude (no periodicity)

(i) Period of System Churn (white implies no periodicity)

Fig. 1. (a-c) Examples of Logistic Map behaviors, (d-f) Total System Churn over time t, (g-i) Heat Maps of Simulated System Churn as functions of δ (x-axis) and h (y-axis). The darker values imply higher values of the observed metric.

with varying M.

b. Increasing h greatly increases the general churn of the system. This rising of the churn stabilization point for the system as can be seen in Figure 1(g) as h and δ increase. We see that h contributes heavily in that a 10% increase in h will cause an average increase of 9.42% in system churn.

c. Increasing δ "kills" the periodicity. Increasing δ by 10% will, on average, decrease the amplitude by 9.14% and, inversely, increase the period by approximately 7.79%. 1(i), and 1(h) show this strong tendency. In contrast to h, δ's effect on the total level of system churn in general only 1.98% per 10% increase, as opposed to h's 9.42%. These general trends are not without exception. We see that for the $h = 0.7$ and $h = 0.8$ region, fixed point stabilization is almost ubiquitous in 1(i). We conjecture this is due to a shifting of the operational regime of the model, partly by our topology, which will be further explained in the next section.

V. Discussion

As show in Figures 1(a), 1(b) and 1(c), we note the presence of fixed points of operation for various r values of the logistic function. By basing our information propagation model on the logistic function (see Eq. 3) and producing a series of simulations, we note behaviors in Figures 1(d), 1(e) and 1(f) that would be expected from the logistic function, namely, stabilization to the fixed point (either 0 or non-zero, as shown in Figures 1(a) and 1(b)) or an induced periodic behavior (as shown in Figure 1(c)). We suspect that for our model, the r parameter is a function of δ, h and some metric of connectivity.

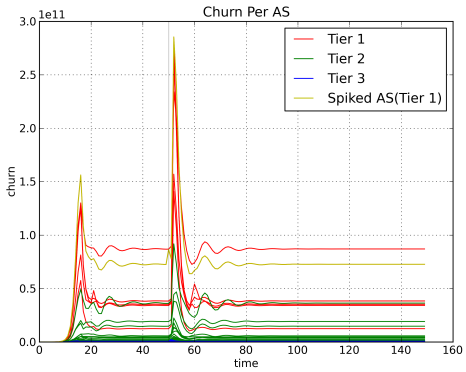

Fig. 2. The lines on this graph show the system's queues aggregated together to form the original ASes of the topology before the graph transform, and colored by their inferred level in the AS hierarchy. The yellow AS line is artificially spiked at time $t = 50$, which is denoted by a gray vertical line. At time $t = 51$ the spiked AS drops as it sends its churn out, increasing the neighboring AS's churn. These neighbors then propagate this churn at $t = 52$, even back to the original AS, causing a much larger spike than the one artificially induced.

This leads us to the following conjecture two conjectures, we leave proof as future work.

Conjecture 1. We conjecture that the non-zero churn stabilization is an intrinsic topological property, for a given set of model parameters. The conjecture is based on our initial experiments, where churn seems independent of the initial excitation. For example, a subsequent excitation adds churns momentarily but the system returns to the same non-zero churn level as before. Given previous works, we further conjecture that the first eigenvalue of the adjacency matrix $\lambda_{A,1}$ is sufficient to quantify the effect of network topology on the information propagation [17][2].

Conjecture 2. We also conjecture that oscillations (period and amplitude) are intimately related to topological properties, and are revealed for particular values of the parameters of our model. We conjecture that the periodicity observed is inherent to the system and is created by topological features.

In addition, we intend, as future work, to extend Algorithm 1. We note that BGP update messages come in a variety of forms, including announcement and withdrawals (see [5] for additional update message distinctions). We conjecture that differences in the type of message will affect the propagation of said messages. By including BGP message type information into Algorithm 1, we hope to capture additional behaviors not observed with a single, homogeneous message type.

VI. Conclusions

In this paper, we begin to establish a theoretical basis for non-binary information propagation, using BGP as a case-study. We are motivated by the observations of [8], [5], and [3] and develop a discrete-time, continuous non-linear model of partial information propagation, based in part on the the logistic function (Eq. 3). We extend the power of our model through a novel graph transformation that encodes BGP-like policy and AS-level connectivity information into a single, queue-based graph. Though simulation on the queue-based graph, we examine the tradeoff between network induced churn amplification versus distance-based churn dampening effects. We observe interesting behaviors including die-outs, stabilization and explosive churn growth. With further work, we hope to closely correlate these behaviors with realistic BGP behaviors. We believe this is the first study of partial information propagation and hope that it leads to more theoretical discussions of the nature of information propagation, challenging the binary or categorical view of information dominate in today literature.

Acknowledgment

The authors would like to thank the numerous reviewers who to contributed to this work. This material is based upon work supported by the Army Research Laboratory under Cooperative Agreement No. W911NF-09-2-0053, the National Science Foundation under Grants No. IIS-1017415, CNS-0721736 and CNS-0721889, NETS-0721889 and NECO-0832069 and a Sprint gift. Any opinions, findings, and conclusions or recommendations in this material are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory, the U.S. Government, the National Science Foundation, or other fund- ing parties. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

References

[1] Y. Rekhter and T. Li, "RFC 4271: A border gateway protocol 4 (BGP-4)," [Online]. Available: https://datatracker.ietf.org/doc/rfc4271/, last accessed: Dec. 2010.
[2] D. Chakrabarti, J. Leskovec, C. Faloutsos, S. Madden, C. Guestrin, and M. Faloutsos, "Information Survival Threshold in Sensor and P2P Networks," in *IEEE ICC*, 2007.
[3] B. A. Prakash, N. Valler, D. Andersen, M. Faloutsos, and C. Faloutsos, "BGP-lens: Patterns and Anomalies in Internet Routing Updates," in *ACM SIGKDD*, 2009, pp. 1315–1324.
[4] C. Labovitz, G. R. Malan, and F. Jahanian, "Internet Routing Instability," in *ACM SIGCOMM*, 1997, pp. 115–126.
[5] J. Li, M. Guidero, Z. Wu, E. Purpus, and T. Ehrenkranz, "BGP Routing Dynamics Revisited," *SIGCOMM Comput. Commun. Rev.*, pp. 5–16, 2007.
[6] M. Lad, X. Zhao, B. Zhang, D. Massey, and L. Zhang, "Analysis of BGP Update Burst During Slammer Attack," in *The 5th International Workshop on Distributed Computing*, Dec 2005.
[7] O. Nordstrom and C. Dovrolis, "Beware of BGP attacks," *ACM Computer Comm. Review*, 2004.
[8] T. G. Griffin and G. T. Wilfong, "An Analysis of BGP Convergence Properties," in *SIGCOMM*, 1999, pp. 277–288.
[9] A. Flavel, M. Roughan, N. Bean, and O. Maennel, "Modeling BGP Table Fluctuations," in *International Teletraffic Congress*, ser. Lecture Notes in Computer Science, L. Mason, T. Drwiega, and J. Yan, Eds., vol. 4516. Springer, 2007, pp. 141–153.
[10] N. Feamster, J. Winick, and J. Rexford, "A Model of BGP Routing for Network Engineering," in *ACM Sigmetrics*, 2004.
[11] E. G. Coffman, Z. Ge, V. Misra, and D. Towsley, "Network Resilience: Exploring Cascading Failures withing BGP," in *Allerton Conference on Communications, Computing, and Control*, 2001.
[12] H. W. Hethcote, "The mathematics of infectious diseases," *SIAM Review*, vol. 42, pp. 599–653, 2000.
[13] A. Ganesh, L. Massoulie, and D. Towsley, "The effect of network topology in spread of epidemics," *IEEE INFOCOM*, 2005.
[14] M. Boguna and R. Pastor-Satorras, "Epidemic spreading in correlated complex networks," *Physical Review E*, vol. 66, p. 047104, 2002.
[15] Y. He, M. Faloutsos, S. Krishnamurthy, and M. Chrobak, "Policy-aware topologies for efficient inter-domain routing evaluations," in *IEEE INFOCOM*, 13-18 2008, pp. 2342 –2350.
[16] S. Sternberg, *Dynamical Systems*. Dover, 2010.
[17] Y. Wang, D. Chakrabarti, C. Wang, and C. Faloutsos, "Epidemic spreading in real networks: An eigenvalue viewpoint," *Reliable Distributed Systems, IEEE Symposium on*, vol. 0, p. 25, 2003.