

On Understanding of Transient Interdomain Routing Failures

Feng Wang[†], Lixin Gao[†], Jia Wang[‡] and Jian Qiu[†]

[†]Department of ECE, University of Massachusetts, Amherst, MA 01002

[‡]AT&T Labs-Research, Florham Park, NJ 07932

[†]{fewang, lgao, jqiu}@ecs.umass.edu, [‡]jiawang@research.att.com

Abstract

The convergence time of the interdomain routing protocol, BGP, can last as long as 30 minutes [14, 15]. Yet, routing behavior during BGP route convergence is poorly understood. BGP can experience transient loss of reachability during route convergence. We refer to this transient loss of reachability during route convergence as transient routing failure. Transient routing failures can lead to end-to-end forwarding failures. Furthermore, the prolonged routing failures can make deploying applications such as voice-over-IP and interactive games infeasible. In this paper, we study the extent to which transient interdomain routing failures occur in the Internet and the duration that these failures can last through both analysis and measurement. We first present a formal model that captures the transient behavior of the interdomain routing protocol. We derive sufficient conditions for and an upper bound for the duration of transient routing failures. Furthermore, we demonstrate the occurrence and duration of transient routing failures in the Internet through measurement. We find that majority of transient failures occur under the commonly applied routing policy setting, and popular and unpopular prefixes can experience transient failures.

1 Introduction

Routing protocols as the “control plane” of the Internet play a crucial role in the end-to-end performance of the Internet. The Internet is divided into thousands of Autonomous Systems (ASes). Routing information is exchanged using the interdomain routing protocol, Border Gateway Protocol (BGP), and routing within an AS is performed using an intradomain routing protocol such as ISIS or OSPF. Studies have shown that intradomain routing can achieve convergence time of a few hundred milliseconds [20]. In contrast, the convergence time of BGP can last as long as 30 minutes [14, 15]. Furthermore, BGP routing instability is pervasive and can occur frequently

[2, 5, 16, 17]. Yet, routing behavior during BGP route convergence is poorly understood.

BGP can experience transient loss of reachability during route convergence. For example, if router *A* uses router *B* to reach a destination, router *A* does not announce its best route to router *B*. This can limit the route visibility of router *B*. Limited route visibility makes it possible for a router to experience transient loss of reachability during the path exploration of the route convergence process. Using the same example as above, suppose that a link failure makes router *B*'s best path infeasible and makes router *B* reach the destination only via router *A*. Router *B* has to withdraw its route from router *A* before router *A* can announce a route to router *B*. That is, router *B* can temporarily lose reachability to the destination during the route convergence process. We refer to this transient loss of reachability during route convergence as *transient routing failure*. Similarly, routing policies and iBGP configurations can limit the route visibility and therefore can lead to transient routing failures.

Transient routing failures can lead to end-to-end forwarding failures. Furthermore, the prolonged end-to-end forwarding failures can make deploying applications such as voice over IP and interactive games infeasible. Therefore, it is important to understand when transient routing failures can occur and how long these transient routing failures can last. However, analysis and measurement studies of transient routing failures can be challenging. First, existing abstract models for BGP focus on route convergence properties or traffic engineering within an AS [4, 8, 13]. The occurrence and duration of transient failures depend on the timing of propagation of route updates, which can be correlated with timing of various events in the network (e.g., link failures, or network configuration changes). Further, timing of routing updates for one prefix is correlated with timing of routing updates for other prefixes since route update rate limiting timers are typically set for each BGP peering session instead of for each prefix. Second, the measurement of transient routing behavior in the Internet requires to differentiate transient routing failures from failures caused by many network changes (e.g., a network or prefix is tem-

porarily unavailable).

Our major contributions are summarized as follows. (i) Contrast to existing models of BGP [6–10, 13], we present an abstract model to capture transient behaviors of BGP, which allows us to scrutinize the detailed interactions between BGP routers. (ii) With the aid of the model, we identify the sufficient conditions for the occurrence of transient routing failures and derive the upper bound of the duration of transient failures. (iii) We derive the upper bound of convergence delay for failover routing changes, which is shown to be much longer than the previous result [16]. (iv) We show that, in a typical BGP system that deploys routing policies conforming to commercial agreements between ASes and applying hierarchical iBGP configurations, any router can experience transient failures. (v) We demonstrate the extent that transient routing failures occur in the Internet by examining a large collection of routing data and configuration files of hundreds of routers. We show that that transient routing failures occur often and can last for a significant period of time. More importantly, we find that transient failures can have a large impact on data traffic of both popular and unpopular prefixes.

The rest of this paper is organized as follows. Section 2 describes our model for investigating transient behavior of a BGP system. Section 3 presents sufficient conditions for transient failures. We study duration of transient failures and convergence delay, and analyze transient failures in a typical BGP system in Sections 4 and 5. Section 6 reports the measurement of transient failures in the Internet. We list related works in Section 7. We conclude the paper with a discussion of future work in Section 8.

2 Abstract Model for BGP Transient Behavior

In this section, we present an abstract model for investigating transient behavior of a BGP system. Our model extends other existing frameworks [11, 12, 21], which capture the long-term stability of BGP, to characterize the detailed timing and interactions in a dynamic BGP system.

2.1 Formal Model

A BGP system (G, P) contains topology G and routing policy P . The topology of the BGP system is modeled as a graph $G = (V, E)$, where the node set V consists of all BGP-speaking routers, and the edge set E consists of all BGP peering sessions. We include both iBGP and eBGP sessions in order to capture routing dynamics within an AS. Each BGP speaker belongs to one AS and an AS can have one or more BGP speakers.

In our model, we focus on a single destination prefix d that originates from AS 0. Clearly, prefixes in a BGP sys-

tem can interact. First, a supernet prefix is used when a subnet prefix is withdrawn. However, we assume that d does not have a supernet. In Section 6, we will see that majority of prefixes that experience transient failures do not have a supernet prefix. Second, BGP routing updates are exchanged at a time triggered by timers. This can lead to correlated routing updates among multiple prefixes. However, our model will capture all possible sequences that routing updates are exchanged. By focusing on a single destination prefix, we will not lose generality.

In order to capture transient behavior of the BGP system, we define the state of the BGP system in terms of the routes stored by each BGP speaker. That is, each speaker remembers the routes received from its neighbors, and also its best route. As such, we define the system state as a vector $S = (s_1, s_2, \dots, s_n)$, where s_i denotes the set of routes stored at speaker $i = 1, 2, \dots, n$, and the first route in s_i is the best route of speaker i .

The route-selection process proceeds in a distributed and asynchronous fashion, triggered by advertisements and withdrawals of routes. A BGP speaker applies the BGP selection process to pick the best path to d , after applying import policies to the routing update received along one of its BGP sessions or edges. After its decision process changes its best route and its export routing policy allows such an update, the BGP speaker sends an update along an edge. Note that the exact timing that a BGP speaker sends an update along an edge to its neighbor is determined by *Minimum Route Advertisement Interval* (MRAI) Timer.

Formally, we model the BGP route decision process as follows. Once an routing update is triggered by BGP MRAI timers, the routing update is sent to the corresponding BGP speaker and the speaker applies its route selection process. *Triggering* a routing update will cause the corresponding BGP speaking routers (or nodes) to apply the import policy to received routes, to run the BGP path-selection process, and to apply export policy for generating routing updates to the speaker’s neighbors.

Since routing update rate limiting timers operate independently, BGP route updates are triggered asynchronously. For each state, we have a set of routing updates, U , that contain routing updates that will be triggered and sent in a future time. At any time, a subset $A \subseteq U$ of updates can be triggered. We call $T = (t_1, t_2, \dots, t_k)$, an *trigger set*, whereas t_i indicates the time that a routing update will be sent along an edge $u \rightarrow v$. Since a routing update is always sent along an edge, we also use the directed edge to represent the routing update.

Given a state $S = (s_1, s_2, \dots, s_n)$ and a trigger set $T \subseteq U$, where U is the future trigger set for state S , the next state $S' = (s'_1, s'_2, \dots, s'_n)$ and the next future trigger set U' will be derived as follows.

$$S' = \text{DecisionProcess}(S, T)$$

$$U' = (U - T) \cup \text{New}T$$

where $\text{DecisionProcess}(S, T)$ derives the new state by running the import policy and best path selection process for each router that receives a routing update. Other routers that do not receive routing updates remain the same state. $\text{New}T$ is a set of routing updates that are triggered by the decision process. In other words, $\text{New}T$ contains routing updates generated by routers whose state has changed and such changes are allowed to export to neighbors by the export policies. The union operation performs an union on $(U - T)$ and $\text{New}T$ so that newer updates are kept if there is an overlap between the two sets. That is, if a routing update along an edge is in both $(U - T)$ and $\text{New}T$, the update message in $\text{New}T$ will be in the future trigger set U' .

A BGP system can go through a series of state transition after the occurrence of an event. These events include link failures, BGP session reset, link addition, router crash, router recovery, and routing configuration changes. In order to capture the transient behavior for a BGP system after the occurrence of such events, we introduce a state graph that describes all possible transient states of the BGP system, and transition between these transient states. A *state graph* is a directed graph, where each node represents a state and its corresponding trigger set, while an edge between states S and S' represents the transition from state S to S' given a subset of the trigger set U in state S . For example, Figure 1 shows a BGP system with a link failure event and Figure 2 shows the state graph for the BGP system. In Figure 1, we show both export and import routing policies in the BGP system. It means, only paths that are allowed to export are shown, and local preference ranking is shown by the order of paths. In Figure 1, the text around a node lists all shows paths that are allowed to export to the node. The order of paths represents their ranks so that the first one has the highest rank. Additionally, node 1 does not forward the path (1 2 0) to node 3 according to its export policies. Each state graph has an initial state that represents the states of BGP speakers initially and the trigger set associated with the event. A directed path in a state graph represents a trigger sequence. For example, in Figure 2, $(T_0 T_2 T_8 T_{18} T_{21})$ is a trigger sequence, where T_0 contains the routing update from node 0 to node 2 indicating withdrawal of the route to node 0.

A state S in a state graph is a *transient state* if $S' \neq S$ for an $T \subset U$. A state S in a state graph is a *stable state* if $S' = S$ for any $T \subset U$. Griffin *et al* have shown in [13] that in a stable state, the best paths to the destination formed from all BGP speakers is a directed tree where the direction of each edge is the same as the direction that packets traverse to reach the destination. We refer to this direct tree as *best path tree* of the stable state. In this paper, we focus on studying the transient behavior of a BGP system that can always reach a stable state for the given event. We also




Figure 1. A BGP system with a link failure. The text around a node lists all paths that are allowed to export to the node, which reflects the export policy. The order of paths reflects the import policy applied by each node. Note that the paths at each node represent potential routes to the destination, and not all of them will show in routing table at the same time.

assume that the BGP system is in a stable state before the occurrence of the event. The best paths at the initial routing state form a best path tree. We refer to the best path tree as the *best path tree of the initial state*.

2.2 Control Plane and Data Plane Failure States

A router is in *control plane failure state* if it has no route to the destination on control plane. For example, state S_2 in Figure 2 is a control plane failure state for router 2. Whether a router goes through a control plane failure state depends on the trigger set sequence it passes. For example, in Figure 2, router 3 does not go through a control plane failure state for trigger set sequence $(T_0 T_2 T_{10} T_{15})$ while it does for trigger set sequence $(T_0 T_2 T_9 T_{12})$. On the contrary, router 2 will definitely go through a control plane failure state. Given a BGP system and an event, a router experiences *potential control plane transient failure* if there is a path from the initial state to the final stable state such that the path contains a control plane failure state for the router; a router experiences *control plane transient failures for sure* if any path from the initial state to the final stable state contains a control plane failure state for the router.

A *forwarding path* on data plane is the path that packets actually pass from a router to a destination. The forwarding path of a router in a state can be constructed by starting from this router and iteratively appending the next hop router of each router to the path. If a router has no complete forwarding path to the destination or the path contains a loop, the router has a *null path*.

A router is in *data plane failure state* if it has a null path to the destination on data plane. It is clear that if a router goes through a control plane failure state, it is sufficient but not necessary for it to experience a data plane failure state. For example, state S_3 of Figure 2 is a data plane failure state

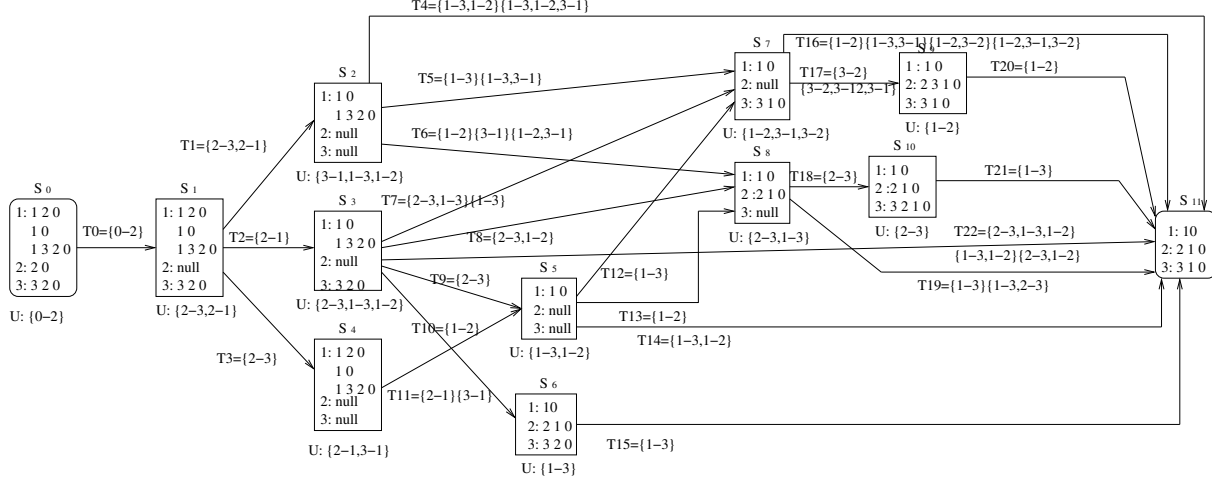


Figure 2. The state graph for the BGP system described in Figure 1.

for router 3 but not a control plane failure state. Similarly, given a BGP system and an event, a router experiences *potential data plane transient failures* if there is a path from the initial state to the final stable state so that the path contains a data plane failure state for the router; a router experiences *data plane transient failures for sure* if any path from the initial state to the final stable state contains a data plane failure state for the router.

3 Sufficient Condition for Transient Failures During Failover

Since a sufficient condition for control plane transient failures is also sufficient for data plane transient failures, we focus on the sufficient condition for the former. Transient failures take place during the failover events that do not disconnected the destination to the network. For simplicity, the events are assumed to be triggered by link failures.

We first introduce *path availability (PA) graph* of a BGP system. A PA graph is a directed graph, consisting of all BGP routers. For two neighboring routers u and v , if the path of one of them is available to the other, there is a directed edge between them. The edge is a solid directed line from u to v if the edge is in the best path tree in the initial state and u installs v 's path. Otherwise, the edge is a directed dashed line from u to v if the best path of v is announced to u and installed at u as *backup*. Note that an edge between two nodes can be bi-directional dashed line if both sides announced their best paths to each other as backup paths. But a solid edge cannot be bi-directional. Each node has one and only one outgoing solid edge.

In a PA graph, a node u is said *reachable* to the destination if there exists a path from u to the destination such that the edges of the path have the same direction as the path and all of them are solid except the one adjacent to the




Figure 3. An example of control plane failures.

destination. A *predecessor/successor* of a node means the predecessor/successor of the node in the best path tree in the initial state.

Figure 3(a) gives an example of control plane transient failures. Figure 3(b) shows the corresponding PA graph. Initially, nodes 2 and 3 use node 1 to reach node 0. When the link between nodes 1 and 0 fails, node 2 possibly experience a control plane transient failure if node 1 sends the withdrawal message to node 2 earlier than the arrival of the path (3 0) from node 3. Note that node 3 announces path (3 0) only after it receives the withdrawal of path (1 0) from node 1. Node 1 will temporarily lose all its paths and experiences a control plane transient failure for sure.

With the PA graph, we have the following sufficient conditions for control plane transient failures. The detailed proofs are omitted and can be found in technical report [23].

Theorem 1 A node u in a BGP system will experience potential control plane transient failures when a link l fails if in the corresponding PA graph, (i) l is a solid line; and (ii) if l is removed, any node in node u 's best path is not reachable to the destination.

Theorem 2 A node u in a BGP system experiences a con-

control plane transient failure for sure when a link l fails if (i) for each predecessor of u , u has at most one path to the destination through it; and (ii) u has only one path to the destination not through u 's predecessors but through l .

Theorem 3 A node u in a BGP system experiences a control plane transient failure for sure when a link l fails if (i) after the removal of l , node u is reachable via u 's predecessors only; and (ii) the predecessors of node u prefer the paths through u over any other paths.

Note that Theorem 2 provides a topological sufficient condition for transient failures for sure, while Theorem 3 relaxes the topological constraint but imposes an additional routing policy constraint.

4 Duration of Transient Failures and Convergence Delay

Transient failures' duration and convergence delay can be identified with state graph. A control plane transient failure for node u can be precisely represented with a trigger set sequence or a path in the state graph, in which the last state in the path is the only state in which u installs a *route* to the destination. Similarly, a data plane transient failure for node u can be precisely represented with a trigger set sequence or a path, in which the last state in the path is the only state in which u has a *data forwarding path* to the destination.

In order to capture the timing for a state to transit to the next, we annotate each edge of the state graph with a *weight* that represents the time that it takes to activate the trigger set. The *delay of a path led by a trigger set sequence* is defined as the sum of weights of edges on the path. More precisely, given two neighboring nodes u and v in a BGP system, the MRAI timer of u configured for v is denoted by M_{uv} and the weight for the edge in the direction from u to v is d_{uv} . Obviously, $d_{uv} \leq M_{uv}$. Note that an edge in the BGP system can be either an iBGP or an eBGP session. The value of M varies for different edges. Suppose node v_k reaches node v_1 through path $P_{(v_k, v_1)} = (v_k v_{k-1} \dots v_1)$. The delay of path $P_{(v_k, v_1)}$ is denoted by $d_{v_k v_1} = \sum_{i=k}^2 d_{v_i v_{i-1}}$. Obviously, $d_{v_k v_1} \leq \sum_{i=k}^2 M_{v_i v_{i-1}}$.

Given a BGP system and an event, the *duration of transient control plane failure* is the delay of the path led by the trigger set sequence of the transient failure and the *duration of transient data plane failure* is the delay of the path led by the trigger set sequence of the transient failure on the data plane. The *control plane convergence delay* is the delay of the path led by the trigger set sequence from the initial state to the final state and the *data plane convergence delay* is the delay of the path led by a trigger set sequence from the initial state to the final forwarding state, where the final

forwarding state is the first state on the path whose corresponding data forwarding state is the same as the final state.

Although [15] has shown that routing convergence delay might be long, their theoretical bound for convergence delay is limited to the events that lead to disconnectivity or bring a route back. However, we argue that the most common events in the Internet are those that lead to failover.

Next, we derive upper bounds of transient failure durations for the case specified by Theorem 2. The upper bounds for the other conditions can be derived in the similar way.

We first locate the nodes where node u can obtain a path, i.e., the node that can be activated to provide u failover path information. After link failure occurs, node u sends a withdrawal to its predecessors in the best path tree at the initial state. The predecessors will forward the withdrawal message to their predecessors if they do not have alternate paths, and so on. If any predecessor can reach the destination via a path other than the best path, it has an alternate path in its routing table in the initial state. For each predecessor of u , there is at most one such node. We denote these nodes with a set $B = \{\beta_1, \beta_2, \dots, \beta_m\}$. Once one of them switches to the alternate path after being activated, the new path information will be advertised to its neighbors. One of the new paths will be finally installed at u .

With the above notions, we have the following theorems. The proofs can be found in [23].

Theorem 4 The control plane transient failure duration at node u is bounded by

$$\min_{\beta \in B} (d_{u\beta} + d_{\beta u})$$

Theorem 5 The data plane transient failure duration at node u is bounded by

$$d_{lu} + \min_{\beta \in B} (d_{u\beta} + d_{\beta u}) = \min_{\beta \in B} (d_{l\beta} + d_{\beta u})$$

where l is the failed link.

Theorem 6 The convergence delay on the control plane and the data plane at node u is bounded by

$$d_{lu} + \max_{\beta \in B} (d_{u\beta} + d_{\beta u}) = \max_{\beta \in B} (d_{l\beta} + d_{\beta u})$$

where l is the failed link.

5 Transient Failures in a Typical BGP System

A typical BGP system means that every router in the system applies *common routing policies*. Routing policies are typically guided by commercial relationships between ASes. In a *provider-to-customer* relationship, the customer

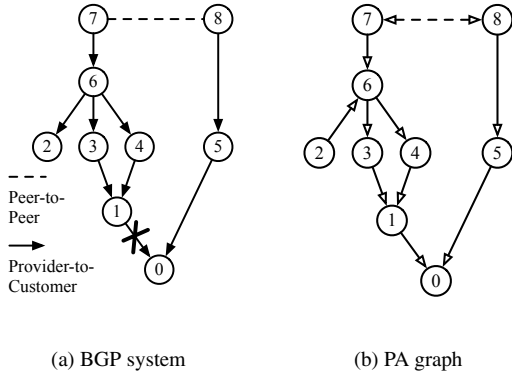


Figure 4. Transient failures take place in a typical BGP system

pays the provider for access to the rest of the Internet. Two ASes with a *peer-to-peer* relationship exchange traffic between their respective customers. In a typical BGP system, the provider-to-customer relationships form a hierarchy of ASes, that is, the directed graph formed by provider-to-customer relationships does not contain a cycle. The export routing policies are typically guided by the *no-valley* routing policy, in which an AS does not export its provider or peer routes to its providers or peers. An AS always announces its customer routes to all its neighbors, and its peer or provider routes to its customers. The import routing policies are guided by the *prefer-customer* routing policy, in which each AS prefers its customer routes over its peer or provider routes.

5.1 Transient Failures at AS Level

Assume that each AS has a single router. We will relax this assumption in Section 5.2 and consider transient failures at all routers inside an AS (i.e., at router level).

Transient failures can be prevalent in a typical BGP system. For instance, Figure 4(a) shows a typical BGP system, in which we assume that each AS contains only one router and the destination is in AS 0. Figure 4(b) illustrates the corresponding PA graph. Suppose that the link between nodes 1 and 0 breaks, then node 6 satisfies the sufficient conditions in Theorem 3, nodes 1, 3, and 4 satisfy the sufficient conditions in Theorem 2, and node 2 satisfies the conditions in Theorem 1. Therefore, nodes 1, 3, 4, and 6 will definitely experience transient failure and node 2 may experience the failure.

Now, we describe necessary conditions for transient failures occurred in tier-1 and tier-2 ASes since our measurement studies in Section 6 focus on tier-1 and tier-2 ASes. Note that a tier-2 AS is an AS that has at least one tier-1 AS as its provider.

Theorem 7 *In a typical BGP system, a tier-1 AS cannot*

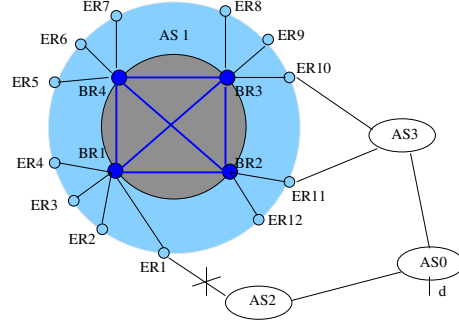


Figure 5. A tier-1 AS with a hierarchical iBGP structure.

experience transient failures.

Theorem 8 *In a typical BGP system, if a tier-2 AS experiences transient failures after a link failure, it must use a customer link to reach the destination before the link failure.*

The proofs of the theorems can be found in [23].

5.2 Transient Failures at Router Level

Our previous analysis shows that tier-2 ASes will experience transient failures, and tier-1 ASes will not experience transient failures under the assumption that those ASes contain only one router. Here, we relax this assumption, i.e., we consider transient failures for an AS containing a set of routers. We focus on transient failures occurring at routers in tier-1 ASes.

Although a tier-1 AS cannot experience transient failures at AS level, routers in a tier-1 AS can. Tier-1 ASes typically have a fully meshed or a hierarchical iBGP structure. Here we focus on describing a hierarchical iBGP structure. Similarly, we can describe scenarios that routers in a fully meshed iBGP structure experience transient routing failures. In a hierarchical iBGP structure, route reflectors are deployed and it consists of a set of *Backbone Routers* (BRs), which are BGP route reflectors, and a set of *Edge Routers* (ERs), which are route reflectors' clients. An edge router could be an *access router* that connects to customer network, or a *peer router* that connects to peer network. For example, Figure 5 shows four backbone routers, which are fully meshed. Each BR connects to a set of edge routers. In terms of route export and import policies, BRs have a peer-to-peer like relationship between each other, i.e., a BR does not transfer the routes between two other BR routers. A BR and its ERs have a provider-to-customer like relationship, i.e., they import and export routes from and to each other without discrimination.

Let's consider the case that all routers use the same egress point to reach the destination. For example, in Fig-

Figure 5, if AS 2 is a customer of AS 1, and AS 3 is a peer of AS 1, according to prefer-customer routing policy, the path via ER1 is assigned higher local preference value than those via ER10 and ER11. As a result, all routers inside AS 1 will use the path via ER1 to reach the destination. Once the link between ER1 and AS 2 fails, all routers inside the AS experience transient failures. According to Theorem 2, all BR1, BR2, BR3, and ER1 will experience transient failure for sure because they have only one path through ER1 to reach the destination before the failure, and all of their predecessors have only one path to reach the destination. All ERs except ER1, ER10, and ER11, and BR4 will experience potential transient failures according to Theorem 1. Therefore, those routers cannot reach the destination along the path through ER11 as soon as the failure occurring. The above example shows if all of routers inside a tier-1 AS use the same egress point to reach a destination, they will experience transient failures if the egress point cannot reach the destination.

5.3 Transient Failure Duration and Convergence Delay

In Section 4, we derive the upper bound of transition failure duration and convergence delay for the case specified by the sufficient conditions in Theorem 2. But for the case in Theorem 3, it is difficult to precisely identify a set B who provides failover path to node u because node u has a set of paths to the destination through the failure link l in a general BGP system. However, in a typical BGP system, we are able to define a similar node set B_{T1} that consists of routers 1) in the (direct or indirect) providers of u , 2) in tier-1 providers, 3) use their direct customer or peer links to reach the destination. Note that the nodes from whom u obtains a failover path should be the nodes in the (direct or indirect) providers of u that do not experience transient failures if u is not in a tier-1 AS, the delay that u obtains a backup path from its provides will be no worse than the delay that u obtains a backup path from a node in B_{T1} . Therefore, with the notion of B_{T1} , we are able to formulate the upper bounds of transient failure duration and convergence delay similar to Theorems 4~6. For details, please refer to [23].

At last, we compare our result with [14], in which the convergence delay of a fail down event is upper bounded by $n \times MRAI$, where n is the length of the longest path to the destination. Our result shows that convergence time can be longer for failover events since the longest path between failure link l and the β node could overlap the longest path from node β to node u . For instance, suppose the link between nodes 0 and 1 shown in Figure 6 fails, the longest path between node 4 and failure node 1 is 2. However, the convergence delay can be as long as $4 \times MRAI$ because




Figure 6. A simple topology with commonly applied routing policy.

the length of the longest path between nodes 1 and 2 (a β node) is 2, and the length of the longest path between nodes 2 and 4 is also 2.

6 Measuring Transient Failures

In this section, we measure the transient routing failures in the Internet. We study the extent to which transient routing failures occur in the Internet by examining routing data collected from a number of tier-1 and tier-2 ISPs. In addition, we investigate the impact of routing policies on transient failures and the popularity of the prefixes experiencing transient failures. In this section, we show the transient routing failures for a tier-1 ISP and the impact of routing policies on routing failures. The rest of the results can be found in [23].

6.1 Measurement Infrastructure

Our measurement uses both public and proprietary data. In particular, we collect the following two sets of data. The first set of data contains the BGP updates, routing table snapshots, and configuration files from a large tier-1 commercial IP backbone with hundreds of edge routers connecting to customer and peer networks. The routing updates are collected by using a BGP monitor that has iBGP sessions (running over TCP) to some top-level backbone routers and to edge routers connecting to peer networks. A snapshot of BGP routing table and the configuration file from each router are collected on a daily basis. These data were collected over the period of 20 days in July 2004. The second set of data contains BGP updates from Oregon RouteViews collected over the entire month of July 2004. We select 4 tier-2 ASes, which are inferred based on the inferred AS relationships [6].

6.2 Routing Failures in the Internet

We consider the router experiencing routing failure to reach the prefix if we observe a withdrawal message for a

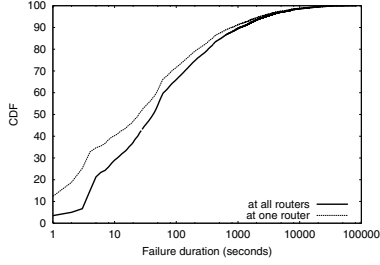


Figure 7. Cumulative distribution of routing failure duration.

prefix at that router. We compute the duration of a routing failure as the time period between the withdrawal message and the first announcement of the prefix after the withdrawal.

We analyze routing updates collected from the tier-1 AS and observe that the monitored routers (52 routers) experience a large number of routing failures. Figure 7 shows the cumulative distributions of routing failure duration at all monitored routers and one backbone router (the x -axis is plotted in log-scale). We observe that most of the routing failures are short-lived. More than 60% of routing failures last less than 100 seconds, and about 50% of failures last less than 30 seconds. This observation holds for both backbone and edge routers. Similar observation also holds on tier-2 ASes. More details can be found in [23].

6.3 Transient Failures in Tier-1 Network

From Figure 7, we find that the Internet has experienced a large number of path failures. Now we analyze transient failures in tier-1 network as following steps.

6.3.1 Identifying Candidate Transient Failures

Among all the observed routing failures, we first identify candidate transient failures. If the paths used before and after a short-lived routing failures (i.e., failures that last less than 30 seconds) are learned from different edge routers, we consider the failure as a *candidate* transient failure. Otherwise, if the paths before and after a routing failure are learned from the same edge router, we consider the failure as *faildown* because there is no alternate path to reach the destination during the failure. Here, we focus on short-lived routing failures. The reason is that an edge router can take as long as $4 \times MRAI$ time to obtain an alternate path from another edge router, and the default MRAI timer among iBGP session is 5 seconds¹. The duration is bounded by

¹Note that not all router vendors implement MRAI timer, such as Juniper routers. On the other hand, MRAI is tunable. Here, we consider the worst case that MRAI timer takes effect at each router.




Figure 8. Percentage of candidate transient failures among all routing failures that last less than 30 seconds.

Theorem 4. If we consider additional MRAI time delay between the edge router and the monitor, at the worst case, the control plane transient failure duration could last about 30 seconds.

For each routing failure, we define the *best path before failure* as the path in the last route announcement before the withdrawal. The path has the highest local preference or shortest AS path, and should be stable for certain period of time (e.g., 5 minutes). If we are not able to find the best path before a failure, we consider that failure is related to a previous failure and ignore it in our analysis. According to our definition of transient failure, we consider the *path after failure* as the first path after the failure.

Figure 8 shows the percentage of candidate transient failures among all routing failures that last less than 30 seconds. We find that 55% to 85% of the routing failures observed at each monitored router are candidate transient failures. The result indicates that majority of short-lived routing failures can be transient failures.

6.3.2 Verifying Candidate Transient Failures

We examine if candidate transient failures are due to failover event. For each candidate transient failure, we check if there is an edge router that still has available path to the destination at the time the failure occurs. We consider the time period from 70 seconds before the failure to 70 seconds after it. If there is an edge router that still has an available path to the corresponding destination, we call the candidate transient failure as a *verified transient failure*. Figure 9 shows the percentage of verified transient failures out of all the candidate transient failures. We find that at most of routers, about 60% of candidate transient failures can be verified. In the remaining of the paper, the analysis is conducted on the verified transient failures.

Figure 10 shows the cumulative distribution of the failure durations for the verified transient failures occurring at all routers and at one backbone router. We observe that for both cases, transient failures are short-lived. More than 95% of transient failures last less than 10 seconds, and more than

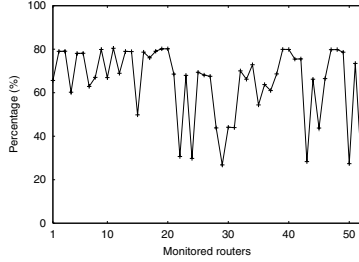


Figure 9. Percentage of candidate transient failures that are verified.

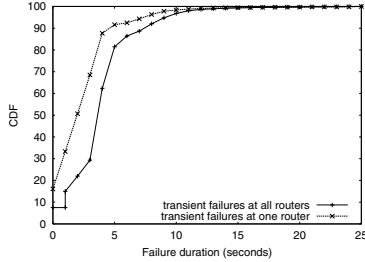


Figure 10. Cumulative distribution of failure duration for verified transient failures.

85% of them last less than 5 seconds.

Note that prefixes can be nested and both the supernet and subnet can exist in the routing table. For example, 10.1.16.0/20 is a subnet of 10.1.0.0/16, and 10.1.0.0/16 is the *supernet* of 10.1.16.0/20. If there is a withdrawal on the route corresponding to the subnet, the destination may still be reachable via the routes corresponding to the supernets. Therefore, the subnet will not experience routing failure. The subnet will experience transient failure only when all the supernets are not available. We examine whether those prefixes that experience transient failures have supernets in the routing table. We find that more than 78% of prefixes that experience transient failures do not have supernet in the routing table. For those prefixes that have supernets (about 22%), we find that all of their supernets will experience at least one routing failures during our measurement. It means that those prefixes are still likely to experience transient failures.

6.4 Impact of Routing Policies on Transient Failures

In Section 5, we study how commonly applied routing policies affect transient failures. However, a tier-1 AS may be configured to prefer a path from its peer over the path from its customer (e.g., backup routing policy). For all verified transient failures, we examine if those transient failures are caused by violation of commonly applied routing policies. We examine the local preference and AS path




Figure 11. Percentage of transient failures that violates *prefer-customer* routing policy.

length of the best route before the failure with the route after the failure, and the type of routers (i.e., peer routers or access routers) from which the paths are learned. We consider a router is configured with *prefer-customer* routing policy if it always assigns higher local preference to the access router than to the peering router. Otherwise, if the router assigns higher local preference to the peering router, we consider it as violation of *prefer-customer* routing policy. Figure 11 shows the percentage of transient failures that violate *prefer-customer* routing policy at each router. We find that only a small fraction of transient failures is due to violation of *prefer-customer* routing policy.

7 Related Work

Previous studies focus on understanding of stability of interdomain routing and measuring end-to-end path performance. Several abstract models [6–10, 13] for routing convergence properties aims to capture the long-term routing stability, and ignore the details of the transient behavior. Griffin et al. show that routing policy conflicts could lead to protocol divergence and characterize sufficient conditions for BGP route convergence [9, 10, 13]. Gao and Rexford [6, 8] exploit AS commercial relationships to ensure the convergence of the BGP system. Our model differs from these existing models in the sense that we aim to capture the transient behavior of BGP, and identify the potential of transient routing failures.

Labovitz *et al.* analyze the convergence delay of BGP and derive theoretical upper and lower bounds for the convergence delay [14–16]. However, their works focus on the convergence delay when a network prefix becomes available or unavailable. Obradovic developed a real-time BGP model to analyze convergence delay under the hierarchical AS relationships [18]. Our work focuses on the convergence delay during the path exploration process of a link failure event.

Correlation between end-to-end path failures and routing instability have been studied through measurement. Paxson identified Internet failures and discovered that routing in-

stability can disrupt end-to-end connectivity [19]. Feamster *et al.* studied the location and duration of end-to-end path failures and correlated end-to-end path failures with BGP routing instability [3]. Their results show that most path failures last less than 15 minutes and most failures that coincide with BGP instability appear in the network core.

Teixeira *et al.* [22] found that routing changes are responsible for the majority of the large traffic variations within a large ISP network. Routing failure within an AS has been studied in [1] by characterizing failures that are correlated with IS-IS routing updates. Our work complements this study by focusing on interdomain routing failures.

8 Conclusions

In this paper, with the aid of a formal BGP model, we investigate the nature of transient behavior of the interdomain routing protocol. We find that network changes that do not cause prefixes unreachable might still cause transient loss of reachability for these prefixes. Both the analytical and measured results show the existence of such transient routing failures in today's Internet routing system. Our measurement results demonstrate that more than half of the observed failures are transient routing failures. These routing failures can last up to 100 seconds and affect both the unpopular and popular prefixes. Therefore, transient routing failures can have a significant impact on the end-to-end performance in the Internet.

Acknowledgements

We are grateful to Jay Borkenhagen, Jennifer Rexford, Aman Shaikh, Oliver Spatscheck, and Zhi-Li Zhang for their valuable comments and suggestions. The work is partially supported by NSF grants CNS-0325868, ANI-0208116, ANI-0085848, and the Alfred P. Sloan Fellowship.

References

- [1] C. Boutremans, G. Iannaccone, S. Bhattacharyya, C. Chuah, and C. Diot. Characterization of Failures in an IP Backbone. In *Proceedings of ACM SIGCOMM Internet Measurement Workshop*, November 2002.
- [2] D. F. Chang, R. Govindan, and J. Heidemann. The temporal and topological characteristics of BGP path changes. In *Proceedings of IEEE ICNP*, November 2003.
- [3] N. Feamster, D. Andersen, H. Balakrishnan, and M. Kaashoek. Measuring the Effects of Internet Path Faults on Reactive Routing. In *Proceedings of ACM SIGMETRICS*, San Diego, CA, June 2003.
- [4] N. Feamster, J. Winick, and J. Rexford. A Model of BGP Routing for Network Engineering. In *Proceedings of ACM SIGMETRICS*, 2004.
- [5] A. Feldmann, O. Maennel, Z. M. Mao, A. Berger, and B. Maggs. Locating Internet Routing Instabilities. In *Proceedings of ACM SIGCOMM*, 2004.
- [6] L. Gao. On Inferring Autonomous System Relationships in the Internet. *IEEE/ACM Transactions On Networking*, 9(6), December 2001.
- [7] L. Gao, T. Griffin, and J. Rexford. Inherently Safe Backup Routing with BGP. In *Proceedings of IEEE INFOCOM*, 2001.
- [8] L. Gao and J. Rexford. A Stable Internet Routing without Global Coordination. *IEEE/ACM Transactions On Networking*, 9(6):681–692, December 2001.
- [9] T. Griffin, F. B. Shepherd, and G. T. Wilfong. Policy Disputes in Path-Vector Protocols. In *Proceedings of the IEEE ICNP*, Toronto, Canada, November 1999.
- [10] T. Griffin and G. T. Wilfong. A Safe Path Vector Protocol. In *Proceedings of IEEE INFOCOM*, pages 490–499, 2000.
- [11] T. G. Griffin, A. D. Jaggard, and V. Ramachandran. Design Principles of Policy Languages for Path-vector Protocols. In *Proceedings of ACM SIGCOMM*, August 2003.
- [12] T. G. Griffin, F. B. Shepherd, and G. Wilfong. The Stable Paths Problem and Interdomain Routing. *IEEE/ACM Transactions on Networking*, 10(2):232–243, April 2002.
- [13] T. G. Griffin and G. T. Wilfong. An Analysis of BGP Convergence Properties. In *Proceedings of ACM SIGCOMM*, Cambridge, MA, August 1999.
- [14] C. Labovitz and A. Ahuja. The Impact of Internet Policy and Topology on Delayed Routing Convergence. In *Proceedings of IEEE INFOCOM*, Anchorage, Alaska, April 2001.
- [15] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian. Delayed Internet routing convergence. *IEEE/ACM Transactions on Networking*, 9(3):293–306, June 2001.
- [16] C. Labovitz, A. Ahuja, and F. Jahanian. Experimental Study of Internet Stability and Backbone Failures. In *Proceedings of FTCS*, pages 278–285, 1999.
- [17] C. Labovitz, G. R. Malan, and F. Jahanian. Internet Routing Instability. *IEEE/ACM Transactions on Networking*, 6(5):515–528, 1998.
- [18] D. Obradovic. Real-time Model and Convergence Time of BGP. In *Proceedings of IEEE INFOCOM*, 2002.
- [19] V. Paxson. End-to-end routing Behavior in the Internet. *IEEE/ACM Transactions on Network*, 5(5):601–615, 1997.
- [20] A. Shaikh and A. Greenberg. OSPF Monitoring: Architecture, Design and Deployment Experience. In *Proceedings of USENIX Symposium on Networked Systems Design and Implementation*, March 2004.
- [21] J. Sobrinho. Network Routing with Path Vector Protocols: Theory and Applications. In *Proceedings of ACM SIGCOMM*, August 2003.
- [22] R. Teixeira, N. Duffield, J. Rexford, and M. Roughan. Traffic Matrix Reloaded: Impact of Routing Changes. In *Proceedings of PAM'05*, BOSTON, MA, March 2005.
- [23] F. Wang, L. Gao, J. Wang, and J. Qiu. On Understanding of Transient Interdomain Routing Failures. Technical Report TR-05-CSE-01, Department of ECE, University of Massachusetts, Amherst, April 2005.