The State of the Art in BGP Visualization Tools: A Mapping of Visualization Techniques to Cyberattack Types

Justin Raynor, Tarik Crnovrsanin, Sara Di Bartolomeo, Laura South, David Saffo, Cody Dunne

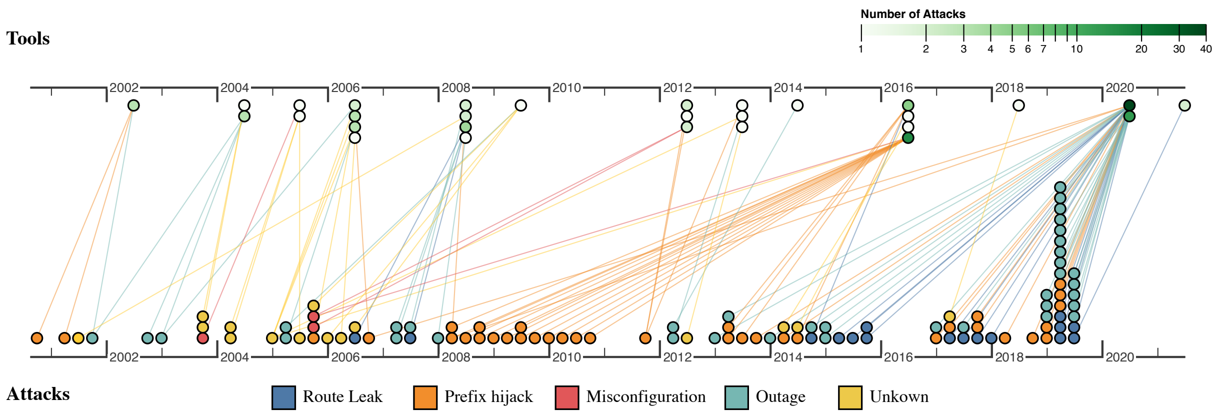

Fig. 1. Timeline of attacks on the BGP system (bottom) and of the tools that have been proposed to visualize them. The bottom nodes and the corresponding edges are colored categorically by the type of attack. Each tool is connected to the specific attacks that were used to demonstrate or validate the tool design, and are colored to show the number of connected attacks.

Abstract—Internet routing is largely dependent on Border Gateway Protocol (BGP). However, BGP does not have any inherent authentication or integrity mechanisms that help make it secure. Effective security is challenging or infeasible to implement due to high costs, policy employment in these distributed systems, and unique routing behavior. Visualization tools provide an attractive alternative in lieu of traditional security approaches. Several BGP security visualization tools have been developed as a stop-gap in the face of ever-present BGP attacks. Even though the target users, tasks, and domain remain largely consistent across such tools, many diverse visualization designs have been proposed. The purpose of this study is to provide an initial formalization of methods and visualization techniques for BGP cybersecurity analysis. Using PRISMA guidelines, we provide a systematic review and survey of 29 BGP visualization tools with their tasks, implementation techniques, and attacks and anomalies that they were intended for. We focused on BGP visualization tools as the main inclusion criteria to best capture the visualization techniques used in this domain while excluding solely algorithmic solutions and other detection tools that do not involve user interaction or interpretation. We take the unique approach of connecting (1) the actual BGP attacks and anomalies used to validate existing tools with (2) the techniques employed to detect them. In this way, we contribute an analysis of which techniques can be used for each attack type. Furthermore, we can see the evolution of visualization solutions in this domain as new attack types are discovered. This systematic review provides the groundwork for future designers and researchers building visualization tools for providing BGP cybersecurity, including an understanding of the state-of-the-art in this space and an analysis of what techniques are appropriate for each attack type. Our novel security visualization survey methodology—connecting visualization techniques with appropriate attack types—may also assist future researchers conducting systematic reviews of security visualizations. All supplemental materials are available at https://osf.io/tupz6/.

1 Introduction

Border Gateway Protocol (BGP) was not built with security in mind. And yet, as the most widely used inter-domain routing protocol on the internet, it is a constant target for myriad threat actors. There has been a proliferation of attacks and anomalies on the BGP system [1], and the number of BGP attacks and anomalies increases each year [46]. With society's growing reliance on internet and networked systems for vital goods and services including online medications, healthcare information, banking transactions, major utilities and more, coupled with the ability of BGP anomalies to affect large swaths of internet resources and millions of users, the need to quickly respond to BGP anomalies is paramount. Even though BGP is governed by a simple set of protocol rules [59], in operational practice BGP routing policy is complex and decentralized—often making any local security adoption challenging to assess for effectiveness. Other challenges exist that add to the inability to implement traditional security approaches. For example, the system produces voluminous data in the form of reachability and update messages where just a few routing updates can generate thousands of BGP messages. Additionally, incidents drastically vary in duration from minutes to several hours making the prospects of properly capturing data from an attack challenging. While there are possible automated solutions for adding security to the BGP system—such as Resource Public Key Infrastructure (RPKI) and Secure BGP (S-BGP) [26,35,38,51]—the deep entrenchment and foundational role of BGP in the internet [64] means that the protocol does not often change and infrastructure and software upgrades are costly. Because of these challenges, visualization tools present a feasible alternative to manage and understand operations, security incidents, and anomalies.

Researchers have developed several BGP visualization tools to aid

- *The authors are all with Northeastern University.
  E-mail: [raynor.j | t.crnovrsanin | dibartolomeo.s | south.l | saffo.d | c.dunne]@northeastern.edu*

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxxx/TVCG.201x.xxxxxxx

expert users in identifying, analyzing and understanding BGP attacks and anomalies. For example, well-known and publicly available tools such as BGPlay[1] [14], and more recently Upstream Visibility [9] and ProBGP [77], have been designed for use by network operators, researchers, and law enforcement. However, there has been less BGP visualization tool development than one might expect given how fundamental BGP is to the internet, as well as its historically-unresolved security issues. A cursory survey of existing tools shows that—although the domain characterization, tasks, and approaches are largely similar—the visualization techniques employed are quite diverse. This diversity of options can cause confusion about which tool to use in which scenario, and, ultimately, lower adoption rates of security approaches [62]. As Butler *et al.* argue, "no solution has yet struck an adequate balance between comprehensive security and deployment cost" [7]. Another reason for lower adoption rates may be because many of the existing BGP visualization tools have not been validated using human-subjects studies and, instead, are "failing to address the focal points of user experience" [65]. Yet, with several recent and high profile BGP incidents [18,36,46], a growing reliance on internet services, and issues that prevent traditional security approaches, creating effective visualizations in this space remains an open challenge.

The purpose of this paper is to (1) serve as a call to action for more visualization work to support BGP security as well as (2) provide an in-depth survey and analysis of existing BGP security visualizations to guide future researchers and practitioners.

In this paper, we use PRISMA guidelines[2] to offer a systematic review of BGP visualization tools. Our review details the state-of-the-art in BGP security visualization and provides the background necessary for designers to build effective and useful visualization tools for current and future threats. We approach this review from three perspectives. The first is to explore how BGP visualization tools have evolved over time. For each tool, we draw connections with the attacks or anomalies that were used to demonstrate and validate the tool's functionality. We draw conclusions about whether and how BGP visualization tools evolve in response to a changing threat model. Second, we categorize the BGP visualization tools by the visualization techniques they employ and the human tasks they are designed to support. Our third approach is to categorize the attacks themselves by type. We offer a novel mapping of visualization techniques to attack types, so that we may identify missing capabilities and opportunities for combining capabilities from existing approaches to better address novel BGP attacks. This study serves not only as a mechanism for understanding the BGP visualization tool space from a historic and evolutionary perspective—including exposing potential capability gaps and future threats—but as a broader survey methodology that can be applied for other security-related tools as well. This study contributes:

1. A comprehensive and systematic survey of the state-of-the-art in BGP visualization, which will assist researchers and practitioners designing the next generation of tools for BGP security. This survey provides:

   (a) A timeline view of BGP visualization tool developments and historic BGP attacks. Understanding the evolution of this domain and the changing threat model can highlight whether and how tools have changed in response.

   (b) A categorization of BGP visualization tools by the component visualization techniques they employ.

   (c) A categorization of BGP attacks and anomalies by type.

   (d) A mapping of the visualization techniques used to the BGP attacks they were meant to address. By analyzing these connections, we can identify which techniques may be most appropriate for known or novel attacks, where there are gaps in current approaches, and where approaches can be combined to help guide future research and development.

2. A novel survey methodology for security related visualization tools, exemplified by our case here with BGP visualization, that directly maps visualization techniques to attack types in order to provide a richer context for future development and a formalization of visualization techniques for this domain.

2 Related Surveys and Motivation

We identified three types of surveys and studies that describe this domain and the state-of-the-art. These are: surveys that focus on overall BGP security, surveys that look at BGP anomaly detection tools, and surveys that more specifically look at BGP visualization tools and techniques. Our study is distinct in that we additionally (1) identify the BGP attack types existing tools were designed to address and (2) how visualization techniques employed in these tools map to those attack types. We provide examples of each of these survey types in order of their specificity and how they relate to our own study. Some of the examples in each survey did not meet our inclusion criteria defined in Section 4. We cite and discuss these examples more thoroughly in Section 4.1. Prior studies motivating the three survey types are summarized below:

Surveys of overall BGP security issues: Butler *et al.* provided an overview survey of current BGP security issues and proposed security solutions such as cryptographic, session security, and filtering techniques [7]. However, they argue that most proposed security solutions are infeasible to deploy because none of them strike a balance between overall security and cost. Gilad *et al.* go further to explain the challenges in depth, as well as realistic goals for implementing security mechanisms [26]. Our survey takes a different approach by instead focusing on current security methods and visualization tools that enable operators to make security decisions—recognizing that the current visualization tool landscape in this domain can benefit from a formalization of techniques.

Surveys of BGP anomaly detection tools: Al-Musawi *et al.* surveyed 20 BGP anomaly detection tools, categorizing them by approach, BGP features used to identify the anomaly, and effectiveness [1]. Additionally, they proposed a taxonomy of BGP anomalies separated by whether the anomalies were direct or indirect, intended or unintended, or a result of link failure. Al-Musawi *et al.* connected these anomaly features with these overall classifications. However, the tools they surveyed are not visualization-centric. Instead, they focus more on algorithmic techniques and pattern matching. In our survey, we recognize that solely algorithmic solutions are not enough to provide operators with actionable intelligence and that visualization can aid in both anomaly detection and decision making.

Surveys of BGP visualization tools and techniques: Shiravi *et al.* provided a use-case-centered survey on visualization systems for network security, with a subset of the study focusing on BGP and routing behavior [65]. They cited six BGP tools, five of which are included in our systematic review: [14,41,71,73,74,80]. However, the visualization technique classification only provided high-level labels including node link graphs, histograms, and color maps, and consideration of lower level visualization features is needed to differentiate.

Tamassia *et al.* offered another survey of visualization techniques focusing on the application of graph drawing methods to cybersecurity areas [69]. Again, BGP-relevant work was a subset of the overall survey and only two examples were provided with high-level labels (force-directed and circular graph drawing methods). These examples are included in our review: [50,73].

Ulmer *et al.* provided a more comprehensive and specific survey of BGP visualization tools [75]. They examined 12 tools, nine of which are included in our study: [10,11,14,22,24,43,55,56,64]. Additionally, they provide an overall domain characterization including challenges, data sources and users. Ulmer *et al.* graded the 12 examples based on their accessibility, scalability, ability to discover BGP incidents and other attributes—we, on the other hand, focus on the visualization techniques themselves. Going further back to 2012, Biersack *et al.* presented a survey of nine BGP visual analytics tools, seven of which appear in our survey [14,41,50,64,71,73,80], that not only considers the visualization techniques employed, but also lower-level visualization

[1] RIPEStat BGPlay https://stat.ripe.net/widget/bgplay

[2] PRISMA: Transparent Reporting of Systematic Reviews and Meta-Analyses http://www.prisma-statement.org/

features, interaction methods, and use cases [5]. Additionally, they distinguish between BGP visualizations that are high-level, or overview, lower-level, or local view, and multi-view. However, only nine tools were presented (compared to the 29 we reviewed) and several novel methods have been proposed since.

Finally, Youn *et al.* presented a study on BGP visualization tools focused on cyber situational awareness [82]. Their survey canvassed 10 tools, seven of which appear in our study: [5, 9, 64, 67, 68, 73, 76], and give more specifics on the visualization techniques employed, core functions, level of detail and use cases. However, they do not consider the specific attacks used to validate the tools or attempt to map the visualization techniques to the attack types.

Our study complements these previous surveys by providing a deeper examination of the connection between visualization techniques and attack types, as well as a formalization of the visualization techniques in this domain. We consider all aspects of BGP visualization tools, including their tasks, techniques, and, more particularly, the attacks that they were designed to detect and analyze. We produce insights into how the visualization tools in this domain are evolving with the attacks, and show the application of specific combinations of visual techniques used for specific attack types.

3 BGP Background and Terminology

In this section, we introduce BGP concepts and terminology in order to help familiarize visualization researchers with this domain. We start with a few computer networking basics.

An Internet Protocol (IP) address is a unique designation that defines an end system or router on the internet or a local network [39]. Currently, Internet Protocol version 4 (IPv4), first deployed in 1983, is still used to route the majority of internet traffic today even though its successor, Internet Protocol version 6 (IPv6), was developed in 1998 to increase the possible address space [29]. IPv4 is a 32-bit address divided into four 8-bit octets, while IPv6 uses a 128-bit address divided into sixteen 8-bit octets. An aggregation of these addresses is denoted by a network prefix, which shows the number of left-most bits that are the same in the group of addresses. For example, 155.33.0.0/16 represents the IP addresses 155.33.0.0–155.33.255.255. A collection of these IP prefixes under the control of one entity is called an Autonomous System (AS), which is denoted by a unique AS Number (ASN). For example, one of Northeastern University's networks is AS156 with the IP prefix 155.33.0.0/16, which is composed of 65,536 IP addresses[3].

The Internet Assigned Numbers Authority (IANA)[4] is responsible for the global coordination system of IP addresses and ASNs. IP address space is generally allocated in a hierarchical manner beginning with a Regional Internet Registry (RIR), then subdivided to an Internet Service Provider (ISP), then to organizations and users [31]. According to Huston's BGP 2020 report, there were 860,000 prefixes (up by 6% from the previous year) and 66,800 ASes (up by 5%) [34].

Routers internal to an AS use internal BGP to communicate routing and reachability information with other routers within the AS. External BGP is used by border routers, or, routers that sit along the edge of an AS, to communicate routing and reachability information between ASes. This information comes in the form of four types of BGP messages: Open, Update, Notification, and Keep Alive. The Update and Notification messages—which advertise new routes, update existing routes, or withdraw routes—are the prime targets for BGP attacks and more often the cause of misconfigurations [45]. When discussing security mechanisms and where they are applied, the BGP messages and the data contained in them represent the *data plane*. The way the messages are routed, and the decisions, policies, and rules used in routing is the *control plane*. Since there are no authentication or integrity mechanisms for these messages, they are relatively easy to forge, which can lead to "blackholing", loss of data, issues with resource reachability, and confidentiality [27]. An attacker can intercept data by routing it through an AS of choice or simply drop the data altogether, essentially making online resources unreachable. However, one of the visualization challenges in identifying BGP attacks and anomalies is that both intentional and unintentional incidents can propagate in the same way. Discerning the difference between an attack versus an innocent misconfiguration is difficult not only from a visualization perspective, but from a fundamental security perspective as well.

Additionally, it is important to note that each AS has its own routing policy defined by the AS owner or operator. Accounting for decentrally-executed routing policies, which affect how attacks propagate, coupled with the inherent challenges in representing dynamic networks is another major challenge to the visualization design space. Routing policies favor and prioritize certain routes over others depending on several weighted factors including, for example, shortest or most reliable paths, cost, local preference, geopolitical considerations, and peering relationships, and each AS may either propagate or drop the message depending on its own internal policy. Routers can receive multiple update messages for the same IP prefix and choose which route to use. It is not always apparent what the impact of a false routing message will be or how it will propagate through the network of ASes, which makes modeling or visualizing future BGP behavior challenging.

4 Scope and Methodology

The scope of this paper focuses on BGP security—not in the form of prevention, but rather on incident and anomaly identification and analysis. Even this subset of the domain encompasses a wide variety of techniques, algorithms, and methods for detecting BGP issues. We limit our systematic review to focus on BGP tools that use visualization techniques in order to aid users in identifying, analyzing, and characterizing BGP attacks and anomalies. In this section, we discuss eligibility and inclusion criteria, how we conducted our search, and how we processed and collected the data.

With relatively few papers encompassed in this review (29), one obvious question from the perspective of justifying this type of study is: Why focus on BGP visualization tools in the first place? One key reason for this choice is that this domain from the perspective of management and security is ripe for visualization solutions and techniques. A formalization of these techniques is needed for designers and researchers, not only to survey what has been implemented in the past, but, more importantly, to understand which visualization techniques are commonly used for different types of attacks and anomalies. Additionally, one benefit of having fewer included papers is that we can compare several different attributes in order to define, formalize, and map visualization techniques to real BGP attacks—as well as identify areas that require more research and focus. We believe that our literature search has surfaced a sufficiently-representative sample of BGP visualization tools to accomplish these goals.

4.1 Data Collection Process

In collecting publications to include in this review, we recognized that our domain was a cross section of both visualization and cybersecurity and, thus, papers came from both communities of research. We applied a top-down approach starting by identifying well-known tools using high-level searches in IEEE Xplore, the ACM Digital Library, and Google Scholar. We further refined our search to focus more on current visualization and security conferences, venues, and journals, including IEEE Transactions on Visualization and Computer Graphics (TVCG), IEEE Visualization (VIS), the Eurographics Conference on Visualization (EuroVis), and the IEEE Symposium on Visualization for Cyber Security (VizSec). From each relevant paper, we followed their citations to identify additional publications to include as well as the later incoming citations, which we discovered using Google Scholar. We used keywords to generate the search results including, for example, "BGP", "BGP visualization", "BGP tool", and "BGP visualization techniques". Additionally, we received feedback, guidance, and recommendations from members of the Center for Applied Internet Data Analysis's (CAIDA) BGP Visualization Working Group.

In total, 1,336 records were identified. Of these records, only 58 were identified as security or attack detection related tools that use visualization in some way. Of the 58 records, we identified 29 publications

[3] AS156—Northeastern University: https://whois.ipip.net/AS156

[4] Internet Assigned Numbers Authority (IANA): https://www.iana.org

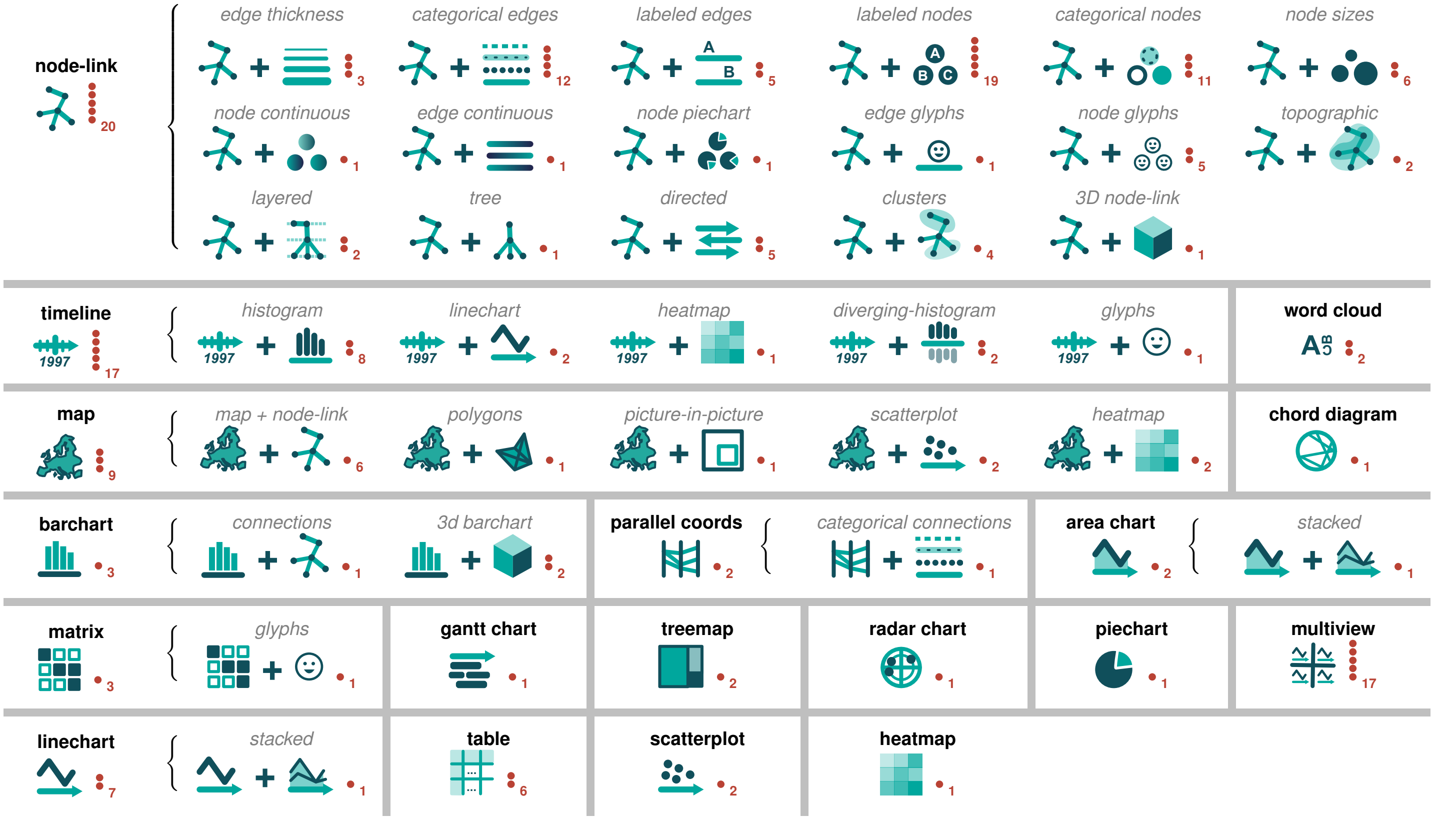

Fig. 2. This table illustrates all the visualization techniques we found used in the papers included in our systematic review. High-level visualization techniques are shown in bold, and curly brackets show sub-categories of high-level techniques. Orange dots show popularity of the technique: one dot • represents ≤ 4 instances, two dots •• represent ≤ 8 instances, three ••• represent ≤ 12, four •••• represent ≤ 16, five ••••• represent ≤ 20.

between 2002 and 2021 that use visualization techniques specifically to allow users to identify, understand, analyze, or characterize BGP attacks and anomalies and eliminated purely algorithmic solutions. In this process, we also identified eight surveys, four additional papers that provide background knowledge and domain characterization, six live projects accessible online, and 13 papers that were deemed out of scope. A full list of venues, surveys, and papers can be found in our resources posted at https://osf.io/tupz6/.

In terms of the papers that were excluded from this survey, we wanted to ensure that we captured the breadth of BGP tools that used visualization techniques while making sure that the papers met our criteria. Several published papers present extraordinary and useful work in the domain of BGP security, but do not use visualization techniques as a mechanism to communicate results or inform decisions. There are lower-level papers focused solely on BGP anomaly detection such as a prefix hijack alert and notification systems [40], the use of deep learning or data mining algorithms for anomaly detection [15, 47], comparing detection methods using decision trees versus naive Bayes methods [16], and detecting anomalies in dynamic networks [52]. We identified papers that were data-centric and more focused on data generation and processing [25], refining and simplifying BGP data [81], or BGP data extraction [6]. Other papers were visualization-specific, but focused on different aspects such as exploration processes when using BGP visualizations [72], visual metaphors [32], and evaluation of specific visualization techniques such as using link weights [42]. Finally, we also explored higher-level papers that were not necessarily domain-specific, such as one that visualized collective anomalies [70] and a system to validate routing policies [44].

There are multiple examples of papers that offer analysis and solutions toward a better and more secure BGP environment. However, we scoped our study to focus solely on visualization tools. Restricting our scope enabled us to provide an analysis of what visualization techniques are used and to connect those techniques to attack types.

4.2 Data Analysis Process

The authors of this study categorized each collected paper using tagging in an iterative process across three sets of attributes (visualization techniques, attack types, visualization tasks). We also consulted with CAIDA's BGP Visualization Working Group. What follows is a description of how we analyzed the BGP tools across these attributes:

1. For visualization technique attributes, we used multiple tags to categorize the papers based on the high-level visualization techniques employed in the tools, lower-level visualization features, and level of user interaction. We mainly relied on the figures of the tools in the papers to determine which visualization techniques were used. We did use live examples when possible. We discuss the top visualization techniques used in Section 6.1.

2. For attack types, we pulled out all BGP attacks and anomalies from the papers that were used to evaluate the performance of the tool or illustrate its use, and used tags to indicate the overall attack type and which papers used which attacks noting when different papers used the same attacks. When the papers were not specific about the details that would allow us to cross-reference the BGP attack or anomaly through a literature search, we noted the overall attack type for which the tool was intended to be used. There were a few papers where neither the attack, anomaly specifics, or attack type were apparent. In those cases, we marked the attack type as unknown or unavailable. We discuss the specific attack types further in Sections 5.1 and 6.2.

3. For visualization tasks, we collected and categorized all of the visualization tasks in the papers by similarity in goals and then categorized them further by overall action and sub-action using Munzner's taxonomy of task abstraction [49]. Visualization tasks are discussed further in Section 5.2.2.

All visualization technique, attack type, and task categorization decisions are available in our supplemental materials at https://osf.io/tupz6/. What follows are descriptions and discussions of the first three visualization products that resulted from this data analysis process, including how they can be used in combination and particular insights:

- Figure 1 is a parallel timeline that temporally depicts the BGP attacks and anomalies from the papers in this survey encoded with discrete colors by attack type, and the BGP visualization tools used to detect them encoded with a continuous color scale to depict the number of attacks used in that particular paper. The connecting edges between the tools and attacks are encoded with the same discrete color of the attack, which allows the reader to better follow the connections. Several insights can be made from this figure. For example:
  - The focus of attack types seems to evolve over time. Earlier tools developed between 2002 and 2009 used BGP attack types that were largely unknown. Tools then focused on prefix hijack attacks almost exclusively until 2016. Then, more recent studies seem to focus on route leaks and outages.
  - The number of BGP attacks used to validate the tools seems to increase over time from very few attacks (one to ten) early on to many (twenty plus) in more recent studies.
  - It is interesting to note that some of the same attacks were used by multiple tools
- Figure 2 shows the top level visualization techniques implemented in the BGP tools, lower level visualization features, such as edge thickness in node-link visualizations, and some combinations of techniques implemented, such as maps combined with node-link graphs. We also provide dot sequences depicting frequency from an overview visual perspective as well as the discrete number of tools out of the 29 studies that use that particular technique. For example, 20 of the 29 tools use node-link graphs, which is the most used technique in this domain, and three of those 20 use edge thickness to encode more particular information.
- Figure 3 shows the distribution of visualization tasks categorized by action and sub-action using Munzner's taxonomy of task abstraction [49]. Further insights can be found in Section 5.2.2.

Tools referenced in Figure 1 and Figure 2 can be cross-referenced with Figure 4, which is discussed further in Section 6, in order to determine the specific combination of visualization techniques used to detect particular attack types. We further provide the references to all of the tools in Figure 1 and Figure 4 and a mapping between top level visualization techniques and attack types in Table 1, which is also discussed further is Section 6. We encourage researchers to derive further insights from these figures in future studies.

5 Literature Review: Domain Characterization

In this section, we present background, domain characterization and attack descriptions of BGP through the lens of visualization design and literature to provide context, background and challenges in this domain.

5.1 BGP Attack and Anomaly Types

Publicly known and documented BGP attacks and anomalies are commonly used to evaluate BGP visualization tools. Sixteen out of 29 of the papers in this survey evaluated their tools in this way while only six of the papers conducted user studies. The prevalence of this approach illustrates that evaluation of visual analytics is challenging [78]. Using known attacks to evaluate BGP tools has benefits from the perspective of understanding an attack and how it is visualized, but also drawbacks with respect to the true ability to effectively identify unknown attacks. As Fischer *et al.* note, "On the one hand, real-world scenarios often have no ground truth, and on the other hand, only experts can identify and validate insights" [23]. Understanding and characterizing the BGP attack surface such as the different dimensions and types of attacks that can occur as well as the vulnerabilities, is critical to visualization designers not only from a design perspective, but also for evaluation purposes. However, there is an apparent issue with balancing the use of already known attacks to evaluate BGP tools and obtaining ground truth about which identified anomalies are actually incidents.

BGP remains a perilously vulnerable domain. Described in *The National Security Strategy to Secure Cyberspace*, BGP vulnerabilities are the "[...] greatest risk of being the target of attacks designed to disrupt or degrade service on a large scale" [30] and several thousand BGP incidents occur each year [60]. As we investigated and classified BGP attacks and anomalies through documented attributes, we categorized and mapped them to visualization techniques. The BGP incidents used to validate BGP visualization tools largely fell into four categories: prefix hijack, route leak, misconfiguration, and physical or logical outages. What follows is an overview of each BGP attack or anomaly. Table 1 shows which tools were designed for which attacks.

Prefix Hijack occurs when an AS incorrectly claims to originate a prefix it was not delegated or owns, causing a multiple origin AS (MOAS) incident where a router has multiple potential paths to reach a resource. By creating an update or announcement message that provides a more attractive route, which agrees more with an AS routing policy than previous updates, attackers can monitor or spy on network traffic, censor access to internet resources, potentially inject malicious software, or even drop network data altogether [75]. Several studies categorize different types of prefix hijacks in different ways. For example, Mitseva *et al.* shows different prefix hijack outcomes through an interception or replay attack, and different means, such as sub prefix hijacks and AS path forgeries [48]. Sermpezis *et al.* propose two different classifications based on the announced AS path and data plane traffic manipulation [63]. Candela *et al.* differentiate between same prefix, or origin AS, and more specific prefixes favored by AS routing policies. Because every paper had different levels of detail, we could only classify incidents with the overall prefix hijack label. Eleven out of the 29 papers in this survey specifically mention prefix hijack as an attack type its tool was designed to support.

Route Leak occurs when an AS learns of a new route and propagates the routing announcement beyond its intended scope, violating the policies of the receiver, sender, or one of the ASes along the announced AS path [66]. This type of attack can send data through an alternate path, which enables possible eavesdropping or data manipulation, but most often occurs by accident or misconfiguration. Nine of the 29 papers mention route leaks as an attack type its tool was designed to support with six of the 9 also supporting prefix hijacks.

Misconfiguration occurs relatively frequently and can have spectacular and public consequences (most recently, for example [3]). They are defined as an intended or unintended production or suppression of BGP announcement messages and are reported as the source of close to three in four of all new prefix announcements [45]. Occurring at the router level, they can cause adverse reachability impacts from increased routing loads, disruptions in connectivity, and policy violations. Even though BGP misconfigurations are listed in the literature as the most common, only three of the 29 papers contain this attack type.

Physical or Logical Outages are a common occurrence in network routing and can take several forms, including downed routers and cable cuts. Whether they manifest in a physical way, such as a backhoe or ship cutting a cable, or in a logical or digital manner through bombarding a network resource via a denial of service (DoS) attack [37] or a worm [21], Bellovin and Gansner show that these types of outages can be used to force internet traffic through an attacker's AS of choice to either capture or black hole the data [4]. Additionally, an outage of a commonly used router or network resource can create several thousand BGP update messages as data reroutes around the obstacle. Fifteen of the 29 papers in this survey, which is the most of any other category, mention physical or logical outages as an attack type the tools support.

Unfortunately, nine of the 29 papers in this survey did not mention an attack type specifically, or the attack types were unknown or unavailable. Some papers had lower levels of detail and sub categories of attacks, such as prefix matching or more specific prefix attacks such as in Candela *et al.* [9]. Additionally, these higher-level attacks can sometimes overlap in terms of observable routing behavior. One of the main challenges with respect to BGP attack and anomaly types is the ability to discern normal from malicious routing behavior [67]. In fact, all of the incidents can be accidental as well as malicious [19]. Additionally, based on the attack definitions, it is relatively easy for researchers to incorrectly classify or mislabel incidents in this domain, and one root attack may cause several instances of different types of BGP incidents as it propagates through a larger network. These often confused definitions are another motivation for our study, which formalizes the visualization techniques to identify specific attack types. Better information on BGP attacks and anomalies used to validate or demonstrate tool capabilities is necessary in this domain.

5.2 Users, Tasks and Data

The characterization of the BGP security domain is largely consistent across the papers in this survey, even though the visualization techniques are diverse. In this section, we provide descriptions for the intended users, tasks, data, and current challenges.

5.2.1 Users

In general, BGP visualization tools are intended for expert users or users with strong domain knowledge. Ulmer *et al.*, for example, define the primary target users of their tool as administrators and prefix owners, with regulators, ISPs, and Internet exchange point (IXP) Managers as secondary users [77]. Additionally, Di Donato *et al.* focuses on network operators [22]. However, there is also a need for visualization tools geared toward practitioners who communicate with the general public, such as policy makers and journalists. These types of users need to communicate the effects of not only large impact incidents and outages that affect larger populations, such as the most recent Facebook outage that reportedly affected 6 billion users [2], but also more targeted effects on smaller providers and companies.

5.2.2 Tasks

From an overview perspective, BGP visualization tools are intended to help network operators, researchers and law enforcement to identify and analyze BGP anomalies, routing patterns, and behaviors in order to investigate and understand what actions are necessary to restore network services. However, not all tools in this survey have clearly defined tasks, and we derived some tasks based on the tool description and paper goals. The lack of clearly defined visualization tasks is a key problem for visualization designers in this domain. A more detailed level of fidelity in visualization tasks will better aid researchers in this domain. Still, there is a high degree of commonality between the tasks and requirements for each tool in this survey. Even though it was not always possible to derive tasks for every tool, the collected tasks and categorization is a good representation of the task needs of this domain because of the clear commonalities that exist between tools.

For this survey, we recorded both explicitly defined visualization tasks and derived other tasks from tool requirements. We arrived at 72 unique tasks and grouped the tasks based on common goals, such as the ability to drill down, detect changes, diagnose issues, and explore. Next, we categorized the tasks based on their actions and sub-actions using Munzner's task abstraction taxonomy [49]. A full list of these tasks and categorizations are available at https://osf.io/tupz6/ and Figure 3 is a summary of this work.

Several insights can be inferred from this chart including that most of the tasks fall within the Analysis and Query actions with Consume and Identify as the top sub-actions. The higher amounts of tasks that fall into the Consume, Produce and Identify sub-actions suggests that users have a stake in the BGP anomaly identification and analysis process. This aligns with Fischer *et al.* 's argument that algorithms are not enough in this domain to provide actionable insights about the threat landscape [23]. Additionally, over 20% of the tasks had to do with the overall goal of incident and anomaly detection. This, coupled with the low number of tasks categorized with the Search action, suggests that users desire to be alerted about BGP anomalies as opposed to having to search for those occurrences. In general, users want to be part of the analysis and diagnosis process, but expect a degree of automated detection and actionable intelligence in order to know how to fix any issues. As Papadopoulos *et al.* argue about solely algorithmic approaches, such as data mining and signature based methods, the analyst is left out of the detection procedure and as a result cannot change parameters and dynamically redefine the results [56]. Further, because of the difficulty of discerning suspicious versus normal routing behavior, a user in the loop is needed.

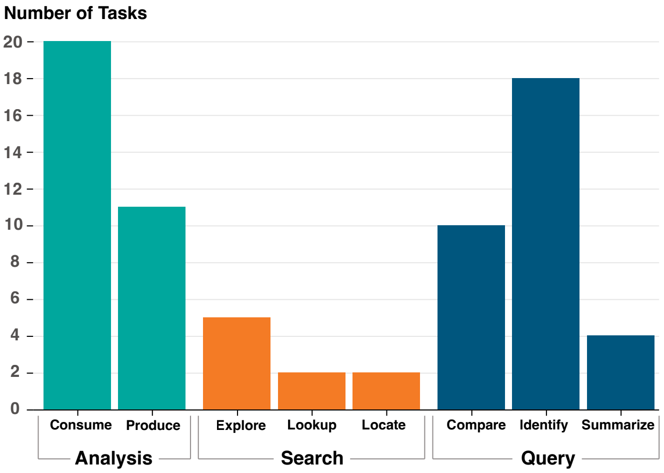

Fig. 3. Distribution of visualization tasks of the tools in this survey categorized by action and sub-action using Munzner's task abstraction taxonomy [49].

5.2.3 Data

In papers reviewed and included in this survey, data sources are few, and they overlap in several canonical BGP incidents used to evaluate multiple visualization tools. Although most papers agree and seem to favor a user in the loop to interpret visual results and take appropriate corrective actions, processing large volumes of data from heterogeneous sources raises serious security concerns [55].

The two most common data sources for BGP data are the University of Oregon's Route Views Project[5] and RIPE Network Coordination Centre's Routing Information Service (RIS)[6] comprising 13 and 15 out of 29 papers respectively. Seven of the 29 papers integrate both data sources into their visualizations [9, 11, 13, 14, 17, 41, 43, 74]. The third most common data source comes from select datasets produced by CAIDA[7] with only a handful of others using BGPmon[8], Spam-Tracer [79], and the Maxmind GEOIP2 Database[9]. Several BGP visualization tools in this survey use data from several sources to integrate multiple perspectives throughout the network and more detailed data for selected network attributes. Authors of only four papers in this survey produced or collected their own BGP data through their own collector or simulated data [20, 68, 71, 80]. One of the major challenges in this domain is how few data sources exist and how relatively difficult it is to produce unique datasets. As Wong *et al.* argue, "Most BGP studies so far have focused on globally visible problems using data from publicly available servers such as RouteViews" [80]. In this way, there is an inherent problem in validating security tools with commonly obtained and available datasets.

Another major challenge with respect to BGP security and data is the location of the data collector, which limits a visualization's perspective and which ASes are visible [43]. The Route Views Project, for example, has 40 collectors in its network, and visualizations that

[5] Route Views Project http://www.routeviews.org/routeviews/

[6] Routing Information Service https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris

[7] CAIDA https://www.caida.org/

[8] BGPMon https://www.bgpmon.net/

[9] Maxmind https://www.maxmind.com/en/geoip2-databases

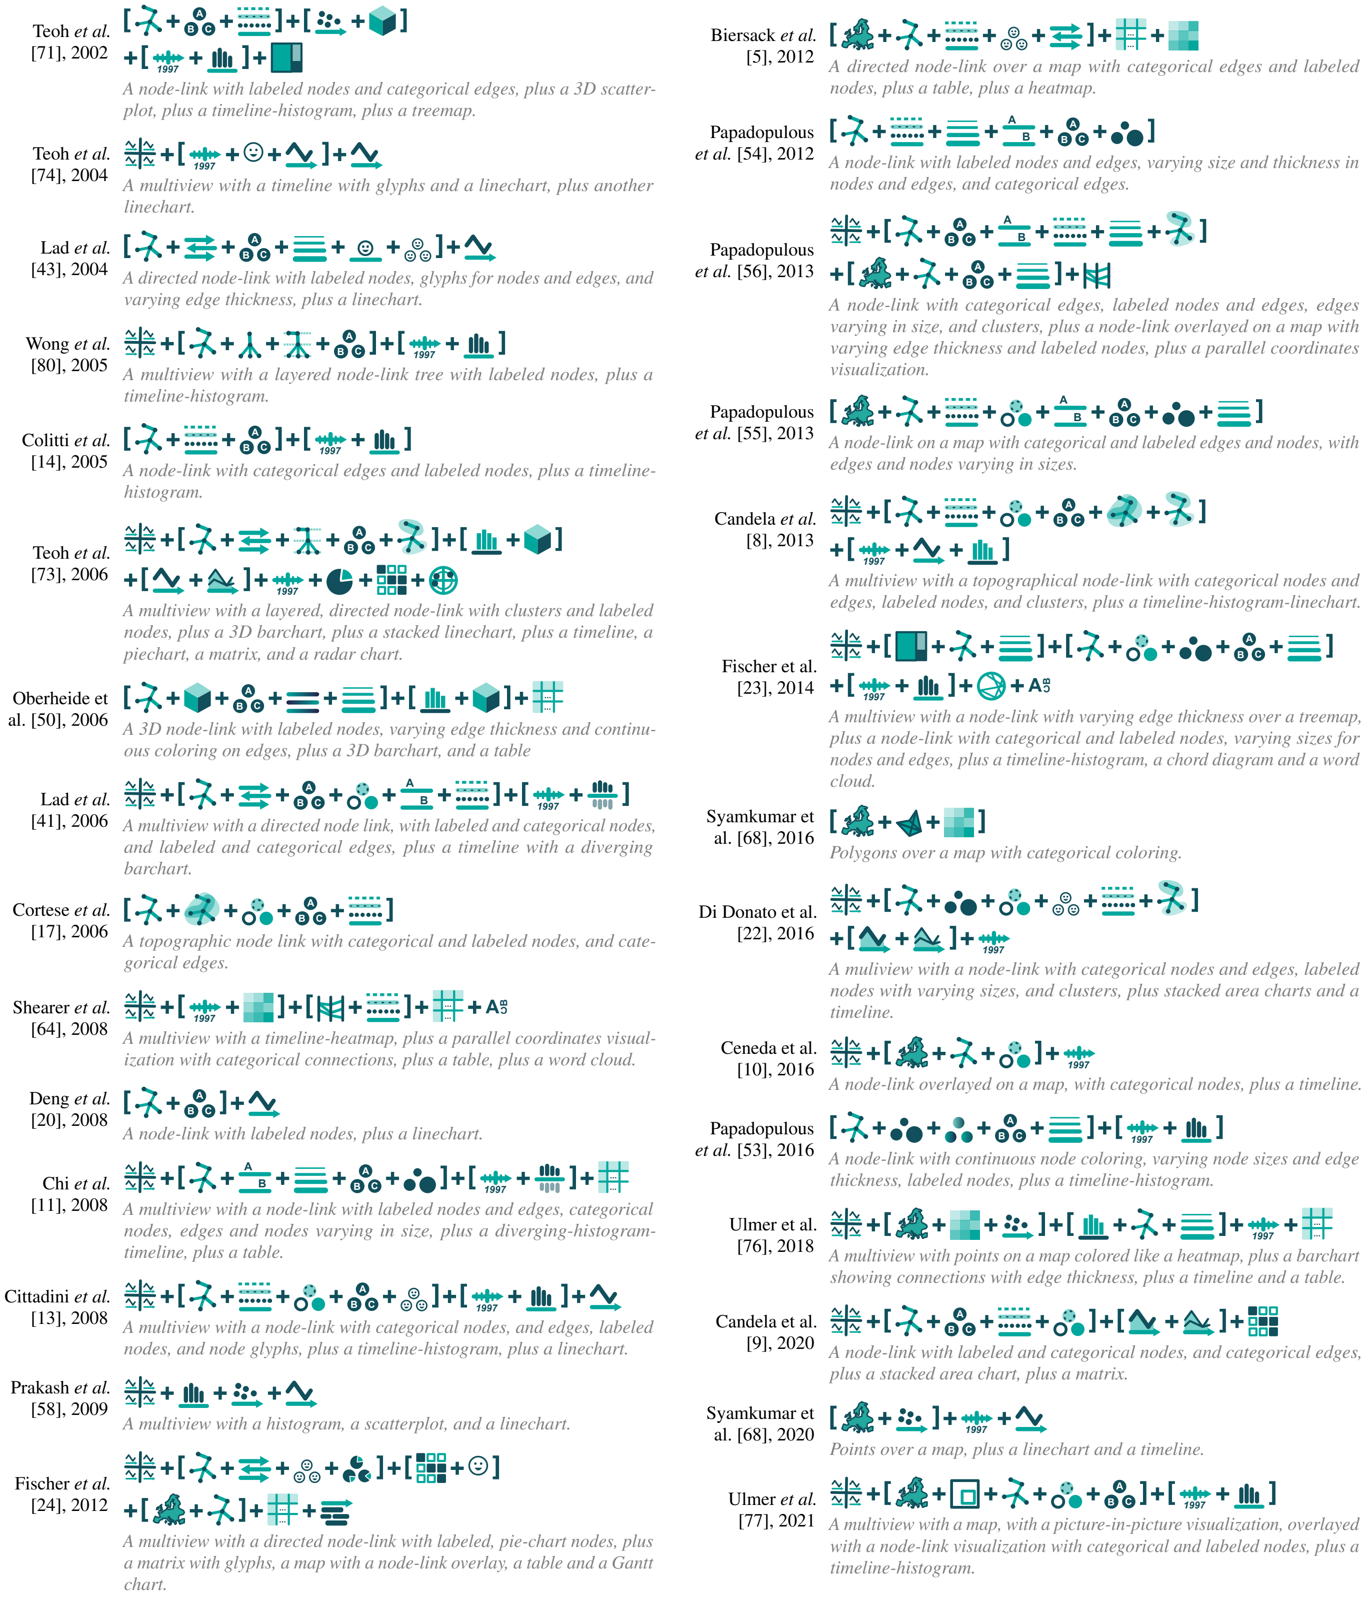

Teoh *et al.* [71], 2002

A node-link with labeled nodes and categorical edges, plus a 3D scatterplot, plus a timeline-histogram, plus a treemap.

Teoh *et al.* [74], 2004

A multiview with a timeline with glyphs and a linechart, plus another linechart.

Lad *et al.* [43], 2004

A directed node-link with labeled nodes, glyphs for nodes and edges, and varying edge thickness, plus a linechart.

Wong *et al.* [80], 2005

A multiview with a layered node-link tree with labeled nodes, plus a timeline-histogram.

Colitti *et al.* [14], 2005

A node-link with categorical edges and labeled nodes, plus a timeline-histogram.

Teoh *et al.* [73], 2006

A multiview with a layered, directed node-link with clusters and labeled nodes, plus a 3D barchart, plus a stacked linechart, plus a timeline, a piechart, a matrix, and a radar chart.

Oberheide et al. [50], 2006

A 3D node-link with labeled nodes, varying edge thickness and continuous coloring on edges, plus a 3D barchart, and a table

Lad *et al.* [41], 2006

A multiview with a directed node link, with labeled and categorical nodes, and labeled and categorical edges, plus a timeline with a diverging barchart.

Cortese *et al.* [17], 2006

A topographic node link with categorical and labeled nodes, and categorical edges.

Shearer *et al.* [64], 2008

A multiview with a timeline-heatmap, plus a parallel coordinates visualization with categorical connections, plus a table, plus a word cloud.

Deng *et al.* [20], 2008

A node-link with labeled nodes, plus a linechart.

Chi *et al.* [11], 2008

A multiview with a node-link with labeled nodes and edges, categorical nodes, edges and nodes varying in size, plus a diverging-histogram-timeline, plus a table.

Cittadini *et al.* [13], 2008

A multiview with a node-link with categorical nodes, and edges, labeled nodes, and node glyphs, plus a timeline-histogram, plus a linechart.

Prakash *et al.* [58], 2009

A multiview with a histogram, a scatterplot, and a linechart.

Fischer *et al.* [24], 2012

A multiview with a directed node-link with labeled, pie-chart nodes, plus a matrix with glyphs, a map with a node-link overlay, a table and a Gantt chart.

Biersack *et al.* [5], 2012

A directed node-link over a map with categorical edges and labeled nodes, plus a table, plus a heatmap.

Papadopulous *et al.* [54], 2012

A node-link with labeled nodes and edges, varying size and thickness in nodes and edges, and categorical edges.

Papadopulous *et al.* [56], 2013

A node-link with categorical edges, labeled nodes and edges, edges varying in size, and clusters, plus a node-link overlayed on a map with varying edge thickness and labeled nodes, plus a parallel coordinates visualization.

Papadopulous *et al.* [55], 2013

A node-link on a map with categorical and labeled edges and nodes, with edges and nodes varying in sizes.

Candela *et al.* [8], 2013

A multiview with a topographical node-link with categorical nodes and edges, labeled nodes, and clusters, plus a timeline-histogram-linechart.

Fischer et al. [23], 2014

A multiview with a node-link with varying edge thickness over a treemap, plus a node-link with categorical and labeled nodes, varying sizes for nodes and edges, plus a timeline-histogram, a chord diagram and a word cloud.

Syamkumar et al. [68], 2016

Polygons over a map with categorical coloring.

Di Donato et al. [22], 2016

A muliview with a node-link with categorical nodes and edges, labeled nodes with varying sizes, and clusters, plus stacked area charts and a timeline.

Ceneda et al. [10], 2016

A node-link overlayed on a map, with categorical nodes, plus a timeline.

Papadopulous *et al.* [53], 2016

A node-link with continuous node coloring, varying node sizes and edge thickness, labeled nodes, plus a timeline-histogram.

Ulmer et al. [76], 2018

A multiview with points on a map colored like a heatmap, plus a barchart showing connections with edge thickness, plus a timeline and a table.

Candela et al. [9], 2020

A node-link with labeled and categorical nodes, and categorical edges, plus a stacked area chart, plus a matrix.

Syamkumar et al. [68], 2020

Points over a map, plus a linechart and a timeline.

Ulmer *et al.* [77], 2021

A multiview with a map, with a picture-in-picture visualization, overlayed with a node-link visualization with categorical and labeled nodes, plus a timeline-histogram.

Fig. 4. The visualizations used in the papers included in this systematic review, described with glyphs. A plus symbol + indicates two features used in combination, while square brackets [ ] enclose a set of visualization techniques and marks used in a visualization. Thus, a paper with multiple visualizations will have multiple enclosures. Descriptions under all the equations can help in understanding the symbols, and all the symbols used are illustrated in Figure 2. An interactive version of this figure can be found at https://osf.io/tupz6/.

use this data source are limited to the AS relationships relative to these collectors. Tools use this data to build visualizations through the announced BGP routes and AS peering relationships. To solve this problem, Sermpezis *et al.* argue that “Adding monitors decreases detection delay and increases visibility of hijacks” [63]. With some visualization tasks requiring both a local network view of close-by or peer ASes, and a more global vantage point, Teoh *et al.* argue that there is a need for a system that integrates multiple viewpoints [73]. In sum, more data sources, more data monitors and collectors, and an integration of both local and global vantage points are needed to improve BGP visualization tool effectiveness.

6 Results: Categorization by Visualization Technique and Attack Type

In this section, we present our findings from categorizing the BGP visualization tools and the attack types for which the tools were developed. We present our results in the form of the component visualization techniques and combinations used in each tool shown in Figure 4. We used the glyphs introduced in Figure 2 to denote each visualization technique and pictorial equations to describe the combinations of techniques used in each tool. Finally, we provide a mapping of the higher level techniques implemented in each tool to the actual attack types shown in Table 1. This enables the ability to find which techniques and combinations of techniques are most used to detect certain types of attacks and where the literature is lacking. What follows is a further discussion of these results.

6.1 BGP Tool Categorization by Visualization Technique

Iteratively categorizing BGP tools by both the top level and lower level visualization techniques implemented enabled us to understand which combinations of techniques are used for particular attack types and the prevalence and frequency with which they are used. The top four techniques observed in this survey were node-link, multi-view, timeline and map. In this section we describe and give examples of each.

Node-link diagrams, which depict a network graph and the relationships among the nodes using edges, were used most frequently in tools in this survey (20 out of 29 papers). Additionally, node-link diagrams represent the most number of different encodings used as visual features to enhance the visualizations. For example, 19 of the 20 tools that use node-link graphs have labeled nodes and 12 of the 20 use categorical edges, which suggests a high user preference for level of detail or detail on-demand. In all cases, node-link graphs were combined with some other form of visualization feature or encoding to provide more information, such as node labeling and categorical edges [71], varying thickness and continuous coloring of edges [50], and labeled nodes with varying sizes [23]. It is also interesting to note the higher frequency of pairings of node-link graphs with timelines or temporal representations (in fact, 14 of the 20 node-link tools use both in some combination), which implies that users need an understanding of the temporal changes to diagnose and identify intentional attacks from accidental routing behavior.

Multi-view visualizations, which combine different perspectives and views, such as high- and low-level graphs, and temporal aspects or topographic representations into one dashboard, represent the second highest visualization technique (17 out of 29 tools) in this survey. Multi-view visualizations often combine scalable and informative approaches for long-term analysis [5], and frequently have linked views that show the relationships between data representations. For example, as one of only three tools designed to analyze BGP misconfigurations, Threshold and Merge Prefixes (TAMP) [80], similar to BGPlay [14], allows the user to playback or animate changes seen in a layered node-link tree linked with a timeline histogram. As Teoh *et al.* argue, linking multiple representations can help identify patterns, anomalies, and different aspects of the data [71].

Timelines are tied for the second most used technique in this survey (17 out of 29 papers) and add temporal aspects to visualization tools necessary to understand the evolution of the AS graph [14]. They are often combined with histograms (10 of the 17 examples in this survey) when showing volume differences of update messages over time.

Maps are the fourth most used visualization technique (nine out of 29 papers in this survey) and represent a somewhat controversial topic in this domain. Bigfoot [67] relies on large changes in geographic footprint of AS routes. However, Papadopoulos *et al.* note that “not all the ASes have a publicly known country of origin” [55]. On the other hand, Ulmer *et al.* argue the advantages of geographic visualizations over conventional node-link diagrams [77] and, how the additional location information helps to distinguish suspicious behavior from an unintended misconfiguration [76]. However, there is an alternate opinion that there is not great meaning in geographic distances versus those that can be represented in a node-link graph. Colitti *et al.* found that “geometrical distances between ASes are largely independent from topological distances and there is no natural way to impose the position of the target AS with respect to the other ASes” [14]. In general, map-based visualizations focus on the user task of recognizing and identifying anomalies. By using this technique, visualization designers rely on the hypothesis that changes in geographic features are noticeable to users. However, the overall effectiveness and utility of using maps in BGP security visualization tools is still argued and relatively unknown.

Category Gaps: There are no examples of BGP visualization tools that do not use one of the top four visualization techniques in some way. However, there are large gaps in Table 1 and other visualization techniques not used in this survey that could contribute to this domain. One of the goals of this study is to not only understand which techniques are most effective for different attack types, but also to identify the gaps in the literature and where visualization research can aid in developing tools for future BGP security threats.

6.2 BGP Attack Categorization by Type

During the categorization of attacks into four types, we found that BGP terminology across the papers was notably inconsistent. For example, one common incident that was labeled differently in at least four of the twenty-nine papers was the TTNET Türk Telekom incident occurring on December 24th, 2004, in which AS9121 announced around 100 prefixes causing large swaths of the internet to be unreachable for most users for several hours [57]. Papers in this survey labeled this incident as both a Route Leak as well as a Prefix Hijack, and definitions of the two overlap in some instances [41, 50, 53, 73].

In some cases, we resolved this problem by using the aggregation of attack descriptions to arrive at a more specific categorization. For example, three of the twenty-nine papers [9, 53, 64] used an incident occurring in February 2008 in which Pakistan attempted to block YouTube from users accessing the website within their country [33]. Through the culmination of the descriptions and data provided in these three papers, we can categorize this incident as a Prefix Hijack and that the method used was announcing a more specific prefix. In some cases, there was either a lack of information or documentation in the papers on particular incidents or there was a disagreement in definitions between papers. Multiple studies exist that categorize BGP attacks and anomalies based on several factors including on which plane they exist (the data or control plane) and which facet they manipulate or effect (e.g. routing data, path, prefix origination etc.) [12, 48, 63]. More recently, Hammood *et al.* suggested a more complex categorization strategy that utilizes different feature selection algorithms to find out the most effective BGP features in order to identify types of anomalies [28]. Without more information or data on the actual BGP attack or incident itself, it is not possible to produce a more granular categorization.

Because of this paradox, we identified BGP incidents from the papers and tagged them with an overall category of Prefix Hijack, Route Leak, Misconfiguration, or Physical or Logical Outage. We added a lower-level categorization when possible from either the paper itself, combinations of descriptions from multiple papers, or collaborating open source articles.

Finally, it is interesting to note the focus of attack types, which does not follow the actual reported order of prevalence in the wild discussed in 5.1. In all, Outages represented the most supported attack type (12 of 29), closely followed by Prefix Hijack (11 of 29), and then Route Leak (9 of 29), and Misconfiguration (only 3 of 29).

Visualization	Prefix Hijack	Outage	Route Leak	Misconfig.	Unknown or Unavailable
Multi-View	[9,24,41,56,64,73]	[9,11,13,23,41,58,64,73]	[9,11,41,73,77]	[80]	[8,10,22,58,74,76]
Node-Link	[9,24,41,50,53,56,71,73]	[9,13,41,43,55,71,73]	[9,41,50,53,73]	[53,54,80]	[8,10,14,17,20,22]
Timeline	[41,53,64,71,73]	[11,13,23,41,64,68,71,73]	[11,41,53,68,73,77]	[53,80]	[8,10,14,22,74,76]
Map	[5,24,56,67]	[55,68]	[67,68,77]		[10,76]
Table	[5,24,50,64]	[11,64]	[11,50]		[76]
Line Chart	[73]	[13,43,58,68,73]	[68,73]		[20,58,74]
Matrix View	[9,24,73]	[9,73]	[9,73]		
Area Chart	[9]	[9]	[9]		[22]
Word Cloud	[64]	[23,64]			
Parallel Coordinates	[56,64]	[64]			
Bar Chart	[50,73]	[73]	[50,73]		[76]
Tree Map	[71]	[23,71]			
Chord Diagram		[23]			
Heatmap	[5]				
Gantt Chart	[24]				
Scatter Plot	[71]	[58,71]			[58]
Historograms		[58]			[58]
Radar Chart	[73]	[73]	[73]		
Pie Chart	[73]	[73]	[73]		

Table 1. The papers we included in the systematic review, classified by high-level visualization technique and attack type.

6.3 Discussion: Tool Strengths and Weaknesses

In this section we qualitatively point out what current BGP visualization tools do well in addition to where there is room for improvement:

1. Domain characterization is solid and consistent between tools. There is not a lot of variance in how the users, data, and tasks are represented in the literature, which suggests that these aspects are well understood.

2. The majority of tools observed in this survey did not conduct user studies. This lack of user-centered design causes a lack of confidence in how effective the visualization techniques actually are in accomplishing the defined tasks.

3. Most tools excel at detecting changes (either through changes in a node-link graph or changes in geography and routing). However, the tools lack in their ability to classify the attack or anomaly and determine if the observed issue is intentional. The overall goal is for the visualization to communicate the anomaly to the user such that they can take appropriate corrective actions. However, with both intentional attacks or even configuration changes and unintentional errors propagating in the same way, more work is needed to understand the difference.

4. There is a lack of BGP visualization tools able to model behavior or downstream effects and impacts. Because routing policy is decentralized, it is very difficult to visualize what effect a networking change will have. However, this need within the domain continues to exist.

5. More tools should be more publicly available. This will enable future researchers to build on existing capability.

6. If the use of historic attacks to evaluate tools continues, studies should provide greater detail of the attacks themselves, which will strengthen formalizations of methods and research.

7. Tools in this survey sometimes lacked a solid visualization task abstraction, which would greatly benefit the visualization design. A strong task abstraction and task clarity are paramount to effective visualization design methodology [61].

8. New BGP visualization tools should find and utilize new sources of BGP data and develop the ability to integrate multiple data sources, which will help enrich new solutions.

We intend this study to help motivate future studies, provide the groundwork to accomplish some of these lacking areas, and to emphasize the current state-of-the-art with a formalization of visualization techniques for this domain.

7 Conclusion

In this study, we have presented a rich domain context, mapping of visualization techniques to attack types, and an initial formalization of visualization techniques for the BGP security domain. A number of challenges and future opportunities were identified in Section 6.3. This study shows the need for a better connection or mapping between the visualization techniques used and the actual attacks that were designed to be analyzed by the tools themselves. As fundamental as BGP is to routing data on the Internet, there are few BGP visualization tools in the community. We present this systematic review to encourage new visualization experts, practitioners and designers to contribute to this space. Additionally, we offer a novel methodology of mapping visualization techniques to attack types as a framework for future security related visualization studies. BGP security is ripe for more visualization implementations. Through linking the techniques used to actual attacks, we contribute the first formalization of visual techniques for the BGP security domain.

Acknowledgments

We thank Kimberly “kc" Claffy and Bradley Huffaker at CAIDA for their insights and suggestions. We thank John Alexis Guerra Gómez at Northeastern University for the interactive version of our survey. Thank you to Jane Kokernak at the Khoury College of Computer Sciences for editing and all members of the Data Visualization Lab at Northeastern.

REFERENCES

[1] B. Al-Musawi, P. Branch, and G. Armitage. BGP anomaly detection techniques: A survey. *IEEE Communications Surveys Tutorials*, 19(1):377–396, 2017. doi: 10.1109/COMST.2016.2622240

[2] Anonymous. Major Facebook outage "triggered by BGP peering configuration snafu", Oct 2021.

[3] B. Barrett. Why Facebook, Instagram, and WhatsApp all went down today, Oct 2021.

[4] S. M. Bellovin and E. R. Gansner. Using link cuts to attack internet routing. 2003.

[5] E. Biersack, Q. Jacquemart, F. Fischer, J. Fuchs, O. Thonnard, G. Theodoridis, D. Tzovaras, and P.-A. Vervier. Visual analytics for BGP monitoring and prefix hijacking identification. *IEEE Network*, 26(6):33–39, 2012. doi: 10.1109/MNET.2012.6375891

[6] D. Blazakis, M. Karir, and J. Baras. BGP-Inspect - extracting information from raw BGP data. In *2006 IEEE/IFIP Network Operations and Management Symposium NOMS 2006*, pp. 174–185, 2006. doi: 10.1109/NOMS.2006.1687549

[7] K. Butler, T. R. Farley, P. McDaniel, and J. Rexford. A survey of BGP security issues and solutions. *Proceedings of the IEEE*, 98(1):100–122, 2010. doi: 10.1109/JPROC.2009.2034031

[8] M. Candela, M. D. Bartolomeo, G. D. Battista, and C. Squarcella. Dynamic traceroute visualization at multiple abstraction levels. In *International Symposium on Graph Drawing*, pp. 496–507. Springer, 2013.

[9] M. Candela, G. Di Battista, and L. Marzialetti. Multi-view routing visualization for the identification of BGP issues. *Journal of Computer Languages*, 58:100966, 2020. doi: 10.1016/j.cola.2020.100966

[10] D. Ceneda, M. Di Bartolomeo, V. Di Donato, M. Patrignani, M. Pizzonia, and M. Rimondini. RoutingWatch: Visual exploration and analysis of routing events. In *NOMS 2016 - 2016 IEEE/IFIP Network Operations and Management Symposium*, pp. 591–597, 2016. doi: 10.1109/NOMS.2016.7502863

[11] Y.-J. Chi, R. Oliveira, and L. Zhang. Cyclops: The AS-level connectivity observatory. *SIGCOMM Comput. Commun. Rev.*, 38(5):5–16, sep 2008. doi: 10.1145/1452335.1452337

[12] S. Cho, R. Fontugne, K. Cho, A. Dainotti, and P. Gill. BGP hijacking classification. In *2019 Network Traffic Measurement and Analysis Conference (TMA)*, pp. 25–32, 2019. doi: 10.23919/TMA.2019.8784511

[13] L. Cittadini, T. Refice, A. Campisano, G. Di Battista, and C. Sasso. Measuring and visualizing interdomain routing dynamics with BGPATH. In *2008 IEEE Symposium on Computers and Communications*, pp. 780–787, 2008. doi: 10.1109/ISCC.2008.4625641

[14] L. Colitti, G. D. Battista, F. Mariani, M. Patrignani, and M. Pizzonia. Visualizing interdomain routing with BGPlay. In *JOURNAL ON GRAPH ALGORITHMS AND APPLICATIONS*, pp. 2004–2005, 2005.

[15] G. Comarela and M. Crovella. Identifying and analyzing high impact routing events with PathMiner. In *Proceedings of the 2014 Conference on Internet Measurement Conference*, IMC '14, p. 421–434. Association for Computing Machinery, New York, NY, USA, 2014. doi: 10.1145/2663716.2663754

[16] R. Copstein and N. Zincir-Heywood. Temporal representations for detecting BGP blackjack attacks. In *2020 16th International Conference on Network and Service Management (CNSM)*, pp. 1–7, 2020. doi: 10.23919/CNSM50824.2020.9269055

[17] P. F. Cortese, G. Di Battista, A. Moneta, M. Patrignani, and M. Pizzonia. Topographic visualization of prefix propagation in the internet. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):725–732, 2006. doi: 10.1109/TVCG.2006.185

[18] J. Crabtree. Council post: BGP attacks pose a substantial operation risk – are enterprises paying attention?, Jan 2021.

[19] M. Danny. Routing without rumor: Securing the internet's routing system, 2022.

[20] W. Deng, P. Zhu, and X. Lu. ROUSSEAU: A monitoring system for inter-domain routing security. In *6th Annual Communication Networks and Services Research Conference (cnsr 2008)*, pp. 255–262, 2008. doi: 10.1109/CNSR.2008.65

[21] S. Deshpande, M. Thottan, and B. Sikdar. Early detection of BGP instabilities resulting from internet worm attacks. In *IEEE Global Telecommunications Conference, 2004. GLOBECOM '04.*, vol. 4, pp. 2266–2270 Vol.4, 2004. doi: 10.1109/GLOCOM.2004.1378412

[22] V. Di Donato, M. Patrignani, and C. Squarcella. NetFork: Mapping time to space in network visualization. In *Proceedings of the International Working Conference on Advanced Visual Interfaces*, AVI '16, p. 92–99. Association for Computing Machinery, New York, NY, USA, 2016. doi: 10.1145/2909132.2909245

[23] F. Fischer, J. Davey, J. Fuchs, O. Thonnard, J. Kohlhammer, and D. A. Keim. A Visual Analytics Field Experiment to Evaluate Alternative Visualizations for Cyber Security Applications. In M. Pohl and J. Roberts, eds., *EuroVis Workshop on Visual Analytics*. The Eurographics Association, 2014. doi: 10.2312/eurova.20141144

[24] F. Fischer, J. Fuchs, P.-A. Vervier, F. Mansmann, and O. Thonnard. VisTracer: A visual analytics tool to investigate routing anomalies in traceroutes. In *Proceedings of the Ninth International Symposium on Visualization for Cyber Security*, VizSec '12, p. 80–87. Association for Computing Machinery, New York, NY, USA, 2012. doi: 10.1145/2379690.2379701

[25] P. Fonseca, E. S. Mota, R. Bennesby, and A. Passito. BGP dataset generation and feature extraction for anomaly detection. In *2019 IEEE Symposium on Computers and Communications (ISCC)*, pp. 1–6, 2019. doi: 10.1109/ISCC47284.2019.8969619

[26] Y. Gilad, T. Hlavacek, A. Herzberg, M. Schapira, and H. Shulman. Perfect is the enemy of good: Setting realistic goals for BGP security. In *Proceedings of the 17th ACM Workshop on Hot Topics in Networks*, HotNets '18, p. 57–63. Association for Computing Machinery, New York, NY, USA, 2018. doi: 10.1145/3286062.3286071

[27] V. Giotsas, G. Smaragdakis, C. Dietzel, P. Richter, A. Feldmann, and A. Berger. Inferring BGP blackholing activity in the internet. In *Proceedings of the 2017 Internet Measurement Conference*, IMC '17, p. 1–14. Association for Computing Machinery, New York, NY, USA, 2017. doi: 10.1145/3131365.3131379

[28] N. H. Hammood and B. Al-Musawi. Using BGP features towards identifying type of BGP anomaly. In *2021 International Congress of Advanced Technology and Engineering (ICOTEN)*, pp. 1–10, 2021. doi: 10.1109/ICOTEN52080.2021.9493491

[29] M. T. Hossain. A review on ipv4 and ipv6: A comprehensive survey. 01 2022. doi: 10.13140/RG.2.2.18673.61284

[30] W. House and U. S. of America. National strategy to secure cyberspace. 2003.

[31] K. Hubbard, D. J. Postel, M. Kosters, D. Karrenberg, and D. R. Conrad. Internet registry IP allocation guidelines. RFC 2050, Nov. 1996. doi: 10.17487/RFC2050

[32] B. Huffaker, D. Plummer, D. Moore, and K. Claffy. Topology discovery by active probing. In *Proceedings 2002 Symposium on Applications and the Internet (SAINT) Workshops*, pp. 90–96, 2002. doi: 10.1109/SAINTW.2002.994558

[33] P. Hunter. Pakistan YouTube block exposes fundamental internet security weakness: Concern that Pakistani action affected YouTube access elsewhere in world. *Computer Fraud & Security*, 2008(4):10–11, 2008. doi: 10.1016/S1361-3723(08)70065-4

[34] G. Huston. BGP in 2020 – the BGP table, 2021. Accessed: 2022-01-20.

[35] J. Israr, M. Guennoun, and H. T. Mouftah. Credible BGP - extensions to BGP for secure networking. In *Proceedings of the 2009 Fourth International Conference on Systems and Networks Communications*, ICSNC '09, p. 212–216. IEEE Computer Society, USA, 2009. doi: 10.1109/ICSNC.2009.74

[36] S. Janardhan. More details about the October 4 outage, Oct 2021.

[37] M. Jonker, A. Pras, A. Dainotti, and A. Sperotto. A first joint look at DoS attacks and BGP blackholing in the wild. In *Proceedings of the Internet Measurement Conference 2018*, IMC '18, p. 457–463. Association for Computing Machinery, New York, NY, USA, 2018. doi: 10.1145/3278532.3278571

[38] J. Karlin, S. Forrest, and J. Rexford. Pretty good BGP: Improving BGP by cautiously adopting routes. In *Proceedings of the Proceedings of the 2006 IEEE International Conference on Network Protocols*, ICNP '06, p. 290–299. IEEE Computer Society, USA, 2006. doi: 10.1109/ICNP.2006.320179

[39] K. R. Kurose. *Computer Networking: A Top-Down Approach by James*. 2017.

[40] M. Lad, D. Massey, D. Pei, Y. Wu, B. Zhang, and L. Zhang. PHAS: A prefix hijack alert system. In *USENIX Security symposium*, vol. 1, p. 3, 2006.

[41] M. Lad, D. Massey, and L. Zhang. Visualizing internet routing changes. *IEEE Transactions on Visualization and Computer Graphics*, 12(6):1450–1460, 2006. doi: 10.1109/TVCG.2006.108

[42] M. Lad, R. Oliveira, D. Massey, and L. Zhang. Inferring the origin of routing changes using link weights. In *2007 IEEE International Confer-*

ence on Network Protocols, pp. 93–102, 2007. doi: 10.1109/ICNP.2007.4375840

[43] M. Lad, L. Zhang, and D. Massey. Link-Rank: a graphical tool for capturing BGP routing dynamics. In *2004 IEEE/IFIP Network Operations and Management Symposium (IEEE Cat. No.04CH37507)*, vol. 1, pp. 627–640 Vol.1, 2004. doi: 10.1109/NOMS.2004.1317749

[44] A. Lutu, M. Bagnulo, and O. Maennel. The BGP visibility scanner. In *2013 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, pp. 115–120, 2013. doi: 10.1109/INFCOMW.2013.6562877

[45] R. Mahajan, D. Wetherall, and T. Anderson. Understanding BGP misconfiguration. *SIGCOMM Comput. Commun. Rev.*, 32(4):3–16, aug 2002. doi: 10.1145/964725.633027

[46] A. Mathurin. A regional look into BGP incidents in 2020, Mar 2021.

[47] K. McGlynn, H. B. Acharya, and M. Kwon. Detecting BGP route anomalies with deep learning. In *IEEE INFOCOM 2019 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, pp. 1039–1040, 2019. doi: 10.1109/INFCOMW.2019.8845138

[48] A. Mitseva, A. Panchenko, and T. Engel. The state of affairs in BGP security: A survey of attacks and defenses. *Computer Communications*, 124:45–60, 2018. doi: 10.1016/j.comcom.2018.04.013

[49] T. Munzner and E. Maguire. *Visualization analysis and design*. A K Peters visualization series. CRC Press, Boca Raton, FL, 2015.

[50] J. Oberheide, M. Karir, and D. Blazakis. VAST: Visualizing autonomous system topology. In *Proceedings of the 3rd International Workshop on Visualization for Computer Security*, VizSEC '06, p. 71–80. Association for Computing Machinery, New York, NY, USA, 2006. doi: 10.1145/1179576.1179592

[51] P. v. Oorschot, T. Wan, and E. Kranakis. On interdomain routing security and pretty secure BGP (PsBGP). *ACM Trans. Inf. Syst. Secur.*, 10(3):11–es, jul 2007. doi: 10.1145/1266977.1266980

[52] J.-c. Pan, D.-m. Han, F.-z. Guo, D.-W. Zhou, N. Cao, J.-r. He, M.-l. Xu, and W. Chen. RCAnalyzer: visual analytics of rare categories in dynamic networks. *Frontiers of Information Technology & Electronic Engineering*, 21(4):491–506, 2020.

[53] S. Papadopoulos, K. Moustakas, A. Drosou, and D. Tzovaras. Border gateway protocol graph: detecting and visualising internet routing anomalies. *IET Information Security*, 10(3):125–133, 2016. doi: 10.1049/iet-ifs.2014.0525

[54] S. Papadopoulos, K. Moustakas, and D. Tzovaras. Hierarchical visualization of BGP routing changes using entropy measures. In *International Symposium on Visual Computing*, pp. 696–705. Springer, 2012.

[55] S. Papadopoulos, K. Moustakas, and D. Tzovaras. BGPViewer: Using graph representations to explore BGP routing changes. In *2013 18th International Conference on Digital Signal Processing (DSP)*, pp. 1–6, 2013. doi: 10.1109/ICDSP.2013.6622756

[56] S. Papadopoulos, G. Theodoridis, and D. Tzovaras. BGPfuse: Using visual feature fusion for the detection and attribution of BGP anomalies. In *Proceedings of the Tenth Workshop on Visualization for Cyber Security*, VizSec '13, p. 57–64. Association for Computing Machinery, New York, NY, USA, 2013. doi: 10.1145/2517957.2517965

[57] A. C. Popescu, B. J. Premore, and T. Underwood. Anatomy of a leak: AS9121. *Renesys Corp., http://www.renesys.com/tech/presentations/pdf/renesys-nanog34. pdf*, 2005.

[58] B. A. Prakash, N. Valler, D. Andersen, M. Faloutsos, and C. Faloutsos. BGP-Lens: Patterns and anomalies in internet routing updates. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '09, p. 1315–1324. Association for Computing Machinery, New York, NY, USA, 2009. doi: 10.1145/1557019.1557160

[59] Y. Rekhter, S. Hares, and T. Li. A Border Gateway Protocol 4 (BGP-4). RFC 4271, Jan. 2006. doi: 10.17487/RFC4271

[60] L. Sani. One year of BGP (in)security, Jul 2020.

[61] M. Sedlmair, M. Meyer, and T. Munzner. Design study methodology: Reflections from the trenches and the stacks. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2431–2440, 2012. doi: 10.1109/TVCG.2012.213

[62] P. Sermpezis, V. Kotronis, A. Dainotti, and X. Dimitropoulos. A survey among network operators on BGP prefix hijacking. *SIGCOMM Comput. Commun. Rev.*, 48(1):64–69, apr 2018. doi: 10.1145/3211852.3211862

[63] P. Sermpezis, V. Kotronis, P. Gigis, X. Dimitropoulos, D. Cicalese, A. King, and A. Dainotti. ARTEMIS: Neutralizing BGP hijacking within a minute. *IEEE/ACM Trans. Netw.*, 26(6):2471–2486, dec 2018. doi: 10.1109/TNET.2018.2869798

[64] J. Shearer, K.-L. Ma, and T. Kohlenberg. BGPeep: An IP-space centered view for internet routing data. In J. R. Goodall, G. Conti, and K.-L. Ma, eds., *Visualization for Computer Security*, pp. 95–110. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.

[65] H. Shiravi, A. Shiravi, and A. A. Ghorbani. A survey of visualization systems for network security. *IEEE Transactions on Visualization and Computer Graphics*, 18(8):1313–1329, 2012. doi: 10.1109/TVCG.2011.144

[66] K. Sriram, D. Montgomery, D. R. McPherson, E. Osterweil, and B. Dickson. Problem definition and classification of BGP route leaks. RFC 7908, June 2016. doi: 10.17487/RFC7908

[67] M. Syamkumar, R. Durairajan, and P. Barford. Bigfoot: A geo-based visualization methodology for detecting BGP threats. In *2016 IEEE Symposium on Visualization for Cyber Security (VizSec)*, pp. 1–8, 2016. doi: 10.1109/VIZSEC.2016.7739583

[68] M. Syamkumar, Y. Gullapalli, W. Tang, P. Barford, and J. Sommers. BigBen: Telemetry processing for internet-wide event monitoring. *arXiv preprint arXiv:2011.10911*, 2020.

[69] R. Tamassia, B. Palazzi, and C. Papamanthou. Graph drawing for security visualization. In *International Symposium on Graph Drawing*, pp. 2–13. Springer, 2008.

[70] J. Tao, L. Shi, Z. Zhuang, C. Huang, R. Yu, P. Su, C. Wang, and Y. Chen. Visual analysis of collective anomalies through high-order correlation graph. In *2018 IEEE Pacific Visualization Symposium (PacificVis)*, pp. 150–159, 2018. doi: 10.1109/PacificVis.2018.00027

[71] S. T. Teoh, K.-L. Ma, S. Wu, and X. Zhao. Case study: Interactive visualization for internet security. In *IEEE Visualization, 2002. VIS 2002.*, pp. 505–508, 2002. doi: 10.1109/VISUAL.2002.1183816

[72] S. T. Teoh, K.-L. Ma, and S. F. Wu. A visual exploration process for the analysis of internet routing data. In *Proceedings of the 14th IEEE Visualization 2003 (VIS'03)*, VIS '03, p. 69. IEEE Computer Society, USA, 2003. doi: 10.1109/VISUAL.2003.1250415

[73] S. T. Teoh, S. Ranjan, A. Nucci, and C.-N. Chuah. BGP Eye: A new visualization tool for real-time detection and analysis of BGP anomalies. In *Proceedings of the 3rd International Workshop on Visualization for Computer Security*, VizSEC '06, p. 81–90. Association for Computing Machinery, New York, NY, USA, 2006. doi: 10.1145/1179576.1179593

[74] S. T. Teoh, K. Zhang, S.-M. Tseng, K.-L. Ma, and S. F. Wu. Combining visual and automated data mining for near-real-time anomaly detection and analysis in BGP. In *Proceedings of the 2004 ACM Workshop on Visualization and Data Mining for Computer Security*, VizSEC/DMSEC '04, p. 35–44. Association for Computing Machinery, New York, NY, USA, 2004. doi: 10.1145/1029208.1029215

[75] A. Ulmer, J. Kohlhammer, and H. Shulman. Towards enhancing the visual analysis of interdomain routing. In *VISIGRAPP (3: IVAPP)*, pp. 209–216, 2017.

[76] A. Ulmer, M. Schufrin, D. Sessler, and J. Kohlhammer. Visual-interactive identification of anomalous IP-block behavior using geo-IP data. In *2018 IEEE Symposium on Visualization for Cyber Security (VizSec)*, pp. 1–8, 2018. doi: 10.1109/VIZSEC.2018.8709182

[77] A. Ulmer, D. Sessler, and J. Kohlhammer. ProBGP: Progressive visual analytics of live BGP updates. *Computer Graphics Forum*, 40(3):37–48, 2021. doi: 10.1111/cgf.14287

[78] J. J. van Wijk. Evaluation: A challenge for visual analytics. *Computer*, 46(7):56–60, 2013. doi: 10.1109/MC.2013.151

[79] P.-A. Vervier and O. Thonnard. SpamTracer: How stealthy are spammers? In *2013 Proceedings IEEE INFOCOM*, pp. 3477–3482, 2013. doi: 10.1109/INFCOM.2013.6567184

[80] T. Wong, V. Jacobson, and C. Alaettinoglu. Internet routing anomaly detection and visualization. In *2005 International Conference on Dependable Systems and Networks (DSN'05)*, pp. 172–181, 2005. doi: 10.1109/DSN.2005.57

[81] H. Yan, R. Oliveira, K. Burnett, D. Matthews, L. Zhang, and D. Massey. BGPmon: A real-time, scalable, extensible monitoring system. In *2009 Cybersecurity Applications Technology Conference for Homeland Security*, pp. 212–223, 2009. doi: 10.1109/CATCH.2009.28

[82] J. Youn, H. Oh, J. Kang, and D. Shin. Research on cyber IPB visualization method based on BGP archive data for cyber situation awareness. *KSII Transactions on Internet and Information Systems (TIIS)*, 15(2):749–766, 2021.