A Novel Radial Visualization of Intrusion Detection Alerts

Yang Shi
Central South University,
Tongji University

Ying Zhao
Central South University

Fangfang Zhou
Central South University

Ronghua Shi
Central South University

Yaoxue Zhang
Central South University

Guojun Wang
Guangzhou University

Intrusion detection systems (IDSs) generally produce an overwhelming amount of alerts, which are commonly plagued by issues of false positives. It is cumbersome for network administrators to manually traverse text-based alert logs in order to detect threats. In this work, we present a novel radial visualization of IDSs alerts, IDSPlanet, which helps administrators identify false positives, analyze attack patterns, and understand evolving network situations. Using a planet's geology as a metaphor for the design, IDSPlanet is composed of chrono rings, alert continents, and an interactive core. Accordingly, these components encode the temporal features of alert types, patterns of behavior in affected hosts, and correlations amongst alert types, attackers, and targets, respectively. The visualization provides an informative picture of networks' status. IDSPlanet offers different interactions and monitoring modes, which allow users to investigate in detail as well as to explore overall pattern. Two case studies and two interviews were conducted to demonstrate the usability and effectiveness of our visualization design.

Intrusion detection systems (IDSs) assist network administrators in monitoring network communication and computer system integrity. IDSs are generally deployed along crucial nodes within the network system. When a predefined attack signature is matched, IDSs produce an alert to indicate the detection of a potentially intrusive behavior, such as malicious activities or policy violations. However, methods to detect attack signatures are lacking in accuracy, precision, and sensitivity. IDSs may mislabel normal activities as malicious, causing false positives, or they may fail

Published by the IEEE Computer Society
0272-1716/19/$33.00 ©2019 IEEE

to identify malicious traffic, resulting in false negatives. In addition, it is time-consuming and tedious for administrators to investigate a plethora of textual alerts generated on a daily basis.

As an emerging inter-disciplinary field, network security visualization leverages human perceptive and cognitive abilities to solve challenging problems. Researchers have validated its significant contributions in analyzing network security data and facilitating decision-making.[1] The visualization of IDSs alert data has been an important subject of inquiry for domain researchers. Many visualization tools[2,3] have been developed to help reduce false positives, identify correlation among alerts, and enhance situational awareness.

Radial visualization[4] is one of the most popular methods among various IDS visualization techniques. It is powerful when presenting correlations among *where*, *when*, and *what* attributes in network security activities. Traditionally, many research efforts[5–7] use line binding to visualize correlations, which might easily result in visual clutters and occlusions in layout design especially when data volume increases. Container binding mechanism, on the other hand, dynamically aggregates and layouts hosts who possess similar features, thus provides more concise and structured visual patterns.

The ability to inspect IDSs records of specific nodes in details also constitutes an important feature to IDSs visualization. For example, monitoring and analyzing IDSs records related to the crucial nodes within the network system have a significant impact on understanding overall situational awareness and making reasonable decisions. However, many IDSs radial visualization tools focus on providing an overview of correlations among alerts and ignore the requirement of converting from an overview to a detailed view for network monitoring.

In this work, we present IDSPlanet, a novel radial visualization tool that facilitates the analysis of IDSs alert logs and comprehension of network security. IDSPlanet extends the field of radial visualization by introducing the container binding mechanism drawn from a planet's geology. It contains three major visual components: *chrono rings* corresponding to a planet's planetary rings show temporal patterns of IDS activities; *alert continents* represent a planet's crust, providing spatial information about the nodes experimenting specific attacks, as well as the heterogeneity and the cardinality of such attacks; *interactive core* is a central core zone, which visualizes relationships among different continents and allows a drag and drop interaction for inspections of specific nodes. To evaluate the effectiveness of IDSPlanet, two use-cases and interviews with two professionals were conducted. Our work was found to bestow several beneficial effects. First, IDSPlanet addresses the correlations among alerts by using the container binding mechanism. The functionality of *IP identification* is used to compare behaviors of an IP in different containers. Second, IDSPlanet reinforces the capability to analyze visual patterns of alert types in both temporal and spatial dimensions. *Alert type sorting* function provides different layout alternatives to facilitate underlying pattern recognition. Third, by integrating global and detailed monitoring perspectives, IDSPlanet enables multiresolution observations of IP nodes.

RELATED WORK

After ten years of development, network security visualization has reached great achievements.[1] It focuses on visualizing diverse network security data (*e.g.*, network traffic monitoring data, network events monitoring data, and malware samples data) and deploying typical visual methods (*e.g.*, node-link graphs, heatmaps, treemaps, and parallel coordinates). A great number of interactive visualization tools have been introduced to detect anomalies, discover attack patterns, and assess security status.

Radial visualization is popular among various visualization methods (*e.g.*, pie chart, star coordinates visualization, sun-burst). Its compact layout is aesthetically pleasing while being easy to read, understand, and interact with.[4] In network security visualization, radial visualization can be applied to many domains. For example, Keim *et al.*[8] use sunburst to hierarchically display network traffic flow and firewall configuration. Zhou *et al.*[9] use radviz to conduct time-series analysis multidimensional network data.

Radial visualization has been adapted by IDSs visualization[10,11] to help network administrators understand and analyze alert data. IDSs alert data is a type of network security event-based data. When analyzing textual raw alert data generated by IDSs, the overwhelming number of alerts as well as false positives could easily exhaust network administrators.[10] Radial visualization provides an efficient method to analyze the correlation between time, location, and type of IDSs alert. VisAlert[5] introduces a typical radial IDSs alert visualization, which maps the *where* of host configuration, *when* of time period, and *what* of alert type onto a visualization as a circular map. The map is surrounded by bands representing time filled with color-coded alert bars. However, issues of occlusion and visual clutter arise when lines are used to display *where-what* connections. AlertWheel[6] reduces visual clutter by designating paths along concentric circles when associating hosts with alert types. Although presenting an organized layout, the cobweb-like design makes it difficult for users to identifying correlations within the *where-what* domain.

Our previous work improves IDSs radial visualization based on the primary visual layout of VisAlert. For example, ID-SRadar[12] develops temporal representation and introduces an automatic concentric circles configuration method of hosts. NetSecRadar[7] uses Bezier curve binding to help better identify correlations when compared to line binding. However, the improved IDSs radial visualization does not fully resolve the issue of visual clutter, especially when dealing with the big data. We found that the requirement of dynamic aggregation of similar hosts as well as correlation identification between hosts should be addressed. Therefore, our work presented here takes one step further to balance these issues and proposes a novel radial visualization method.

DATA AND DATA PROCESSING

Our database uses IDSs logs generated by Snort, which is one of the most popular open source network IDSs. IDSs alert data is emitted by IDSs when specific predefined attack patterns (signature-based IDSs) or potentially dangerous behaviors (anomaly-based IDSs) have been detected. Each Snort IDSs log entry contains many details. The important parts of a log entry are identified as the text of specific rules violated, the description of the alert, the priority of the alert, date, and IP information. To facilitate correlation identification among *where, when,* and *what* domain, IDSPlanet uses three relevant attributes, including timestamp (when), alert type (what), and host (where) (an IP address is assigned to each host for communication in a computer network. An IP could act as an attacker who initiates attacks or a target who receives attacks). The data processing module of our system reads the data in the specified format from textual logs and writes to the database. The process is divided into three stages, including format recognition, attribute filtering of textual logs, and data formalization. Specifically, attribute filtering focuses on important attributes relevant to timestamp, alert type, and host. We then conduct a series of aggregation functions according to different visualization requirements. For example, visualization of an alert type's temporal variation uses the report frequency of the alert type in a specific time interval [see Figure 1(a)-data]. As shown in Figure 1(b)-data-lower, presentation of an IP node uses the report frequency of the IP related to a certain alert type in a specific time interval. These aggregation results are stored for the purpose of accelerating rendering.

VISUALIZATION DESIGN

Visualization

IDSPlanet is modeled as a celestial object with a composition similar to planets such as Earth and Saturn. The principal components include Chrono Rings, Alert Continents, and Interactive Core, corresponding to a planet's ring systems, crust, and core, respectively. The design is based on geological mappings of *when-where-what* domains in the network security field. *When* refers to the time when the alert happens. *Where* refers to the network node that reports the alert. *What* refers to the type of the alert. Our goal is to visually organize the alert instances and address correlations among the three attributes in a recognizable and informative way. We map the *Where-What* onto a planet's body, while the *When-What* domain is wrapped around it as a planet's ring.

Chrono Rings

Chrono rings present the *when-what* connections and visualize the temporal variations of IDSs alerts. Rings are composed of numerous ring particles, the size and visibility of a ring particle is based on the frequency of alerts generated at a certain interval. Alert incidents of the same type in a certain span are

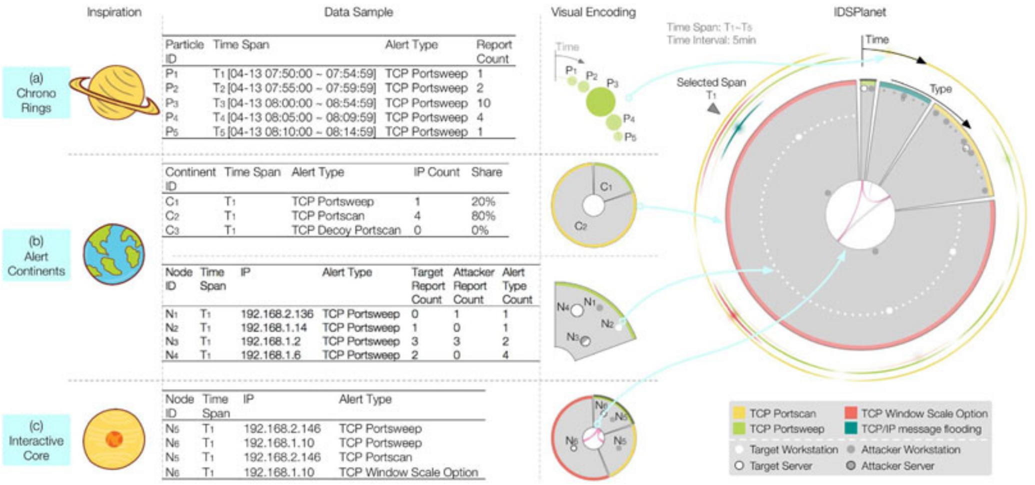

Particle ID	Time Span	Alert Type	Report Count
P_1	T_1 [04-13 07:50:00 ~ 07:54:59]	TCP Portsweep	1
P_2	T_2 [04-13 07:55:00 ~ 07:59:59]	TCP Portsweep	2
P_3	T_3 [04-13 08:00:00 ~ 08:54:59]	TCP Portsweep	10
P_4	T_4 [04-13 08:05:00 ~ 08:09:59]	TCP Portsweep	4
P_5	T_5 [04-13 08:10:00 ~ 08:14:59]	TCP Portsweep	1

Continent ID	Time Span	Alert Type	IP Count	Share
C_1	T_1	TCP Portsweep	1	20%
C_2	T_1	TCP Portscan	4	80%
C_3	T_1	TCP Decoy Portscan	0	0%

Node ID	Time Span	IP	Alert Type	Target Report Count	Attacker Report Count	Alert Type Count
N_1	T_1	192.168.2.136	TCP Portsweep	0	1	1
N_2	T_1	192.168.1.14	TCP Portsweep	1	0	1
N_3	T_1	192.168.1.2	TCP Portsweep	3	3	2
N_4	T_1	192.168.1.6	TCP Portsweep	2	0	4

Node ID	Time Span	IP	Alert Type
N_5	T_1	192.168.2.146	TCP Portsweep
N_6	T_1	192.168.1.10	TCP Portsweep
N_5	T_1	192.168.2.146	TCP Portscan
N_6	T_1	192.168.1.10	TCP Window Scale Option

Figure 1. Design of each visual component of IDSPlanet, including illustration of its inspiration, data sample, and visual encoding. (a) *Chrono rings* correspond to a planet's planetary rings. They show temporal patterns of IDSs activities. (b) *Alert continents* represent a planet's crust. It provides spatial information about the nodes experimenting specific attacks, as well as the heterogeneity and the cardinality of such attacks. (c) *Interactive core* is a central core zone. It presents relationships among different continents and allows a drag and drop interaction for inspections of specific nodes.

gathered as one ring according to their time sequences in clockwise order, starting from the 12 o'clock position. As a result, each ring represents time-varied frequency of a specific type of alert and is assigned a unique color label. For example, as shown in Figure 1(a), each particle represents TCP Portsweep alert incidents occurred at 5-min intervals. In a span of 25 min from T_1 to T_5, five particles are chronologically arranged as a segment of the TCP Portsweep ring and assigned a green label. A particle may expand, brighten, or even detonate if a sufficient spike in the threat level is detected. In other words, the visual intensity of a particle corresponds with its threat level derived from the IDSs alert reports. These visual effects highlight the time frames of high risks and help differentiate between regular patterns and suspicious activities.

Alert Continents

Alert continents help identify correlations within the *where-what* domain and support spatial analysis of alert incidents. The activity and behavior between hosts within the selected time frames are visualized within IDSPlanet's alert continents. The alert continents' layer is analogous to the crust of a planet. IDSPlanet forms several continents based current alert types, with each populated by IP nodes who report at least one incident of the continent's alert. Each alert continent's size is based on the amount of IP nodes it contains. The outermost layer of each alert continent receives a color label according to its alert type. The color is consistent with that of the corresponding chrono ring. For example, as shown in Figure 1(b)-upper, the green TCP Portsweep continent and the yellow TCP Portscan continent form a planet. Based on their populations, the green continent occupies one-fifth of the circle while the yellow continent occupies four-fifths of the circle.

Alert continents focus on the spatial representation of IDSs alerts using the container binding mechanism, in which hosts who report the same alert types are collected. The distance between an IP node and the core of IDSPlanet indicates the frequency of alerts. That is, IP nodes who receive a large amount of attacks are considered more "burdened," causing it to sink toward the core. The size of an IP node shows the degree of variety among its alert types. IP nodes that experiment a wide range of attacks are enlarged. For example, as shown in Figure 1(b)-lower, N_3 who reports the highest total alerts is placed at the nearest position from the core. N_4 who encounters the most alert types is represented as the biggest node. To differentiate between attackers and targets of alerts, IP nodes are rendered with differing amounts of white and gray. The ratio between white and gray slices shows the number of times a node acting as a target or source. For example, a node that is mostly white indicates an IP that acts largely as a source of attacks [see N_1 in Figure 1(b) visual lower], while a node that is mostly gray represents a target of attacks [see N_2 in Figure 1(b)-lower]. The server is addressed by adding a black stroke [see N_3 and N_4 in Figure 1(b)-lower]. By observing group trends and individual activities through various visual encoding, abnormal activity patterns of a subset of hosts could be found.

Interactive Core

By facilitating an in-depth observation of important individuals, the interactive core offers users with the options to review alert incidents from multiple perspectives. Similar to a planetary core, IDSPlanet's interactive core is located at the central of the visualization. The interactive core provides two functionalities, it analyzes correlations among alert types and among IP nodes. First, the interactive core displays correlation among various alert continents. Correlation curves that link the continents are used to indicate the similarity between alert continents. The width of the curve depicts the percentage of identical IPs between two different alert types. For example, as shown in Figure 1(c), two IPs, N_5 and N_6, each reports two types of alerts and both of them report TCP Portsweep alert. When the green TCP Portsweep alert continent is selected, pink correlation arcs show that half of its IPs are identical to those in yellow TCP Portscan continent, and the other half is identical to that in the red TCP Window Scale Option continent. The curves traverse across planet core along the shortest paths with minimal overlapping. The design is based on a minimalist representation that allows for simpler and more direct recognition of correlations between different alert types.

Second, the interactive core supports a detailed view and acts as a monitoring station when suspicious IPs are observed. For example, as shown in Figure 2, when IP 192.168.1.4 is placed in the interactive core at the selected time slice, IDSPlanet switches to individual analysis mode of this IP node. Two types of attacks are found relevant to this IP. IP 192.168.2.143 acts as a source of TCP Portsweep to attack targeted IP 192.168.1.4 while IP 192.168.1.6 is the target of TCP Portscan attack initiated by IP 192.168.1.4. The transmission flows are depicted as flow arcs; pink input flow arc and gray output flow arc. The interactive core supports fine-grained monitoring of suspicious hosts as well as other crucial nodes.

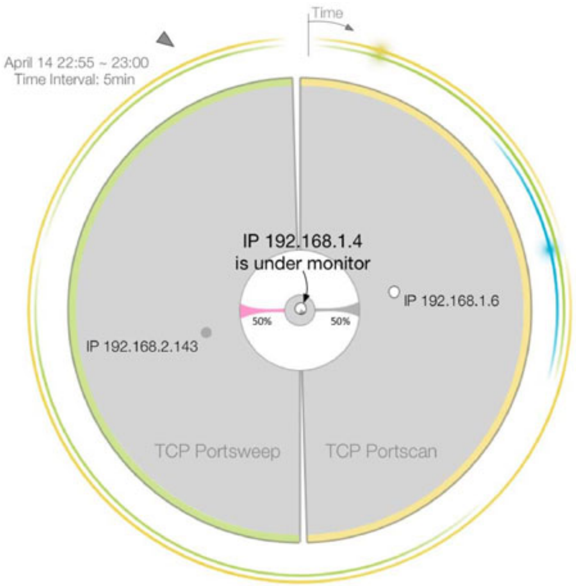

Figure 2. IDSPlanet transitions to a detailed view when IP 192.168.1.14 is placed in the interactive core. Two types of attacks are found relevant to this IP, Green TCP Portsweep and Yellow TCP Portscan. Its transmission flows are depicted as flow arcs. The pink arc represents that the IP acts as the target of the green attack while the gray arc illustrates that the IP initiates the yellow attack.

Interaction

IDSPlanet provides various interactions, among which we address several customized interactions in details. These interactions enable users to obtain comprehension of threat detection and anomaly awareness.

Alert Type Sorting

We provide different layout alternatives for IDSPlanet, through which users could find underlying patterns in the IDSs alerts. Specific measures [i.e., incident time, Pearson product-moment correlation coefficient (PPMCC), and IP quantity] can be used to sort chrono rings and alert continents. When set to the incident time option, the first reported alert type is placed as the innermost chrono ring, followed by successive outer rings of increasing incident time. In alert continents, the alert type of the first reported incident takes its position on the top of the IDSPlanet, at 12 o'clock. Other types are placed based on ascending incident time in clockwise order [see Figure 3(a)]. The PPMCC option sorts chrono rings and alert continents based on the similarity of their temporal fluctuation. The base alert type ring is placed as the innermost chrono ring, and its alert continent is placed at the 12 o'clock position. Other types are placed in descending order according to their temporal similarity coefficient with base alert [see Figure 3(b)]. Figure 3(c) shows the results of the IP quantity sorting which places the alert type with the least incidences as the innermost chrono ring. In alert continents, the smallest continent is placed in the 12 o'clock position, with progressively ascending continent sizes in clockwise order.

Time Refinement

We offer time control with multiple resolutions. Time refinement enables users to focus on analyzing suspicious IDSs alerts and reduce the scope for further investigation. In our case study, we use 60-min

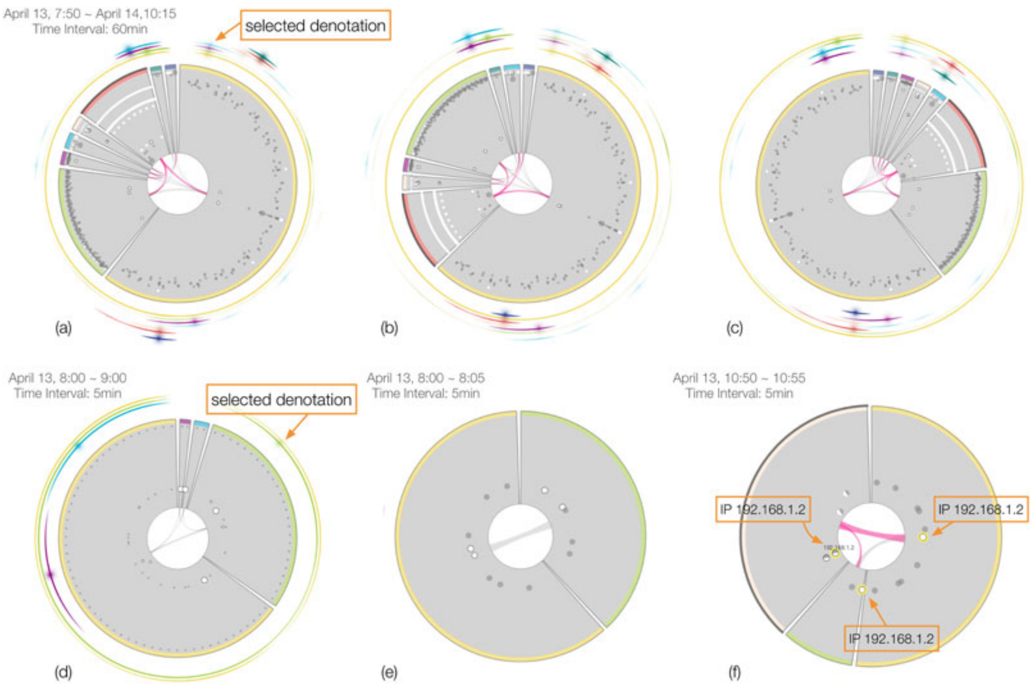

Figure 3. (a)–(c) Global view of IDSPlanet where the alert type sorting method is set to incident time, PPMCC, and IP quantity sorting, respectively. The sorting method decides the order of chrono rings and the order of alert continents (time span: April 13, 2011 7:50 April 15, 2011 10:15. time interval: 60 min). (d) Detailed view of IDSPlanet when the first detonation in (a) occurs. (e) Detailed view of IDSPlanet in 5 min when the second denotation in (d) occurs. (f) IP in the pinkish continent is selected, its instances in green and yellow alert continents are highlighted.

intervals for coarse-grained analysis and 5-min intervals for fine-grained analysis as complements to each other. When the IDSs log data is loaded, users may tend to observe the overall trend over the entire time series. Sixty minutes granularity supports the general observation [see Figure 3(a)]. When a further observation is required, users select time frames of interest. The configuration of IDSPlanet is updated accordingly and the time resolution changes to 5-min granularity [see Figure 3(d)]. Specifically, users could zoom in to a specific time point (i.e., 5 min) on alert continents for detailed analysis [see Figure 3(e)].

IP Identification

IP Identification enables the comparison of the actions that an IP takes part in different alerts, thus giving further insight into its roles and behaviors. Considering that IDSPlanet uses the container binding mechanism for host configuration, a single IP reported multiple alert types necessitates multiple placements on different alert continents. IP identification is designed to help quickly pinpoint and isolate a specific IP address, users can hover their mouse over an IP, causing its other instances in other continents to be highlighted. As shown in Figure 3(f), when IP 192.168.1.2 in the pinkish continent is clicked, we found several instances in other continents. It acts as a target in the green and yellow continents within the selected time frames. In addition, if users hover the mouse over an alert continent, correlation arcs showing the percentage of mutual IPs between these and other types will be highlighted [see pink arcs in Figure 3(f)].

EVALUATION

In this section, we demonstrate the usability and effectiveness of IDSPlanet through two case studies and two interviews with experts (a demo of our work is available at https://youtu.be/zssOc1_LIq4).

Case Study

The data consists of the Snort IDSs logs provided by the IEEE VAST Challenge in 2011 and 2012.[12]

Case 1—Shipping Company IDSs Logs

The first case visualizes the IDSs logs from a corporate network of a major shipping company from April 13, 2011, to April 15, 2011, including 8 alert types and about 20 000 alert records. Figure 4(a) shows a global view of IDSPlanet. The alert types are sorted according to the IP quantity which places the smallest continent at the 12 o’clock position, with progressively ascending continent sizes in clockwise order.

We focus on analyzing visual patterns from the three largest alert continents that contain the most of IPs. Chrono rings show that the yellow TCP Portscan [alert types are henceforth referred to with their colors and abbreviation, as shown in Figure 4(c)] and green portscan2 persists for long durations over regular periods. The red snort decoder occurs only in two time periods, accompanied with two detonations along the ring. This indicates a sudden increase in alert reports, which requires additional investigation. Using this temporal pattern, we decided to focus on the red snort decoder alerts in the following exploration of the spatial features.

For alert continents, we pay attention to the distribution of hosts. We found that the alert types of yellow portscan1 and green portscan2 possess many similarities. For example, they both feature a larger number of attackers (depicted as gray nodes) than targets (depicted as white nodes). In terms of the roles in attacks, several server IPs, displayed as white nodes with black strokes, appear near the core. This indicates that they are the main targets of a distributed attack from a large number of sources. In the red alert continent, a pattern emerges where all the white nodes, which are targets of an attack, are separated into four layers. IP nodes in the same layer keep the same distance to the center, suggesting that they receive an identical amount of attacks. The attacker IPs are those few gray nodes placed near the center. This highly structured attack pattern in a large amount of nodes implies a planned attack and calls for further monitoring.

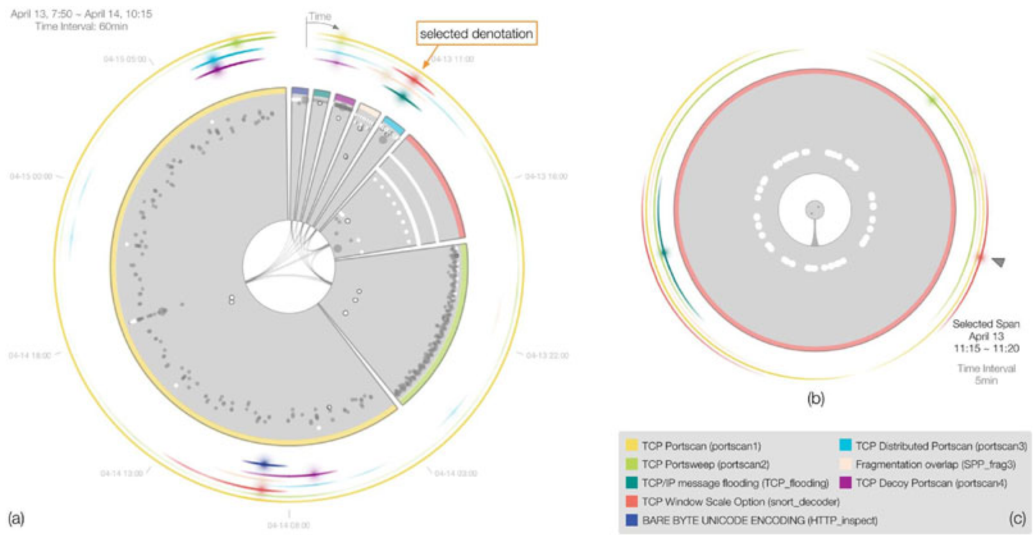

Figure 4. (a) Global view of IDSPlanet where the alert types are sorted according to IP quantity. The alert type with the least incidences is represented by the innermost chrono ring, and the smallest continent is placed in the 12 o'clock position. (b) Detailed view of IDSPlanet when the initial burst of red snort decoder alert on chrono rings occurs. IP 192.168.2.171 and IP 192.168.2.173 are placed in the interactive core. The gray flow arc shows that these two IPs are attackers. The distance between the core and target IP nodes shows that these two IPs attack each of their targets using identical strategies. (c) Legend of color encoding of alert types, including its names and abbreviations.

We transition to the fine-grain view and observe the alerts from 11:15 to 11:20 on April 13, when the initial burst of activities occurs. We found that a huge amount of workstations is attacked by five hosts which share similar behavior patterns. We then place two of the low-variety attackers, IP 192.168.2.171 and IP 192.168.2.173, into the interactive core to investigate their individual behaviors. As shown in Figure 4(b), we found that these two IPs are responsible for only red snort decoder attacks. The pattern of their targets which are all placed equidistantly from the core implies they attack each target using identical strategies. From this observation, we drew the conclusion that the two IPs investigated may be malicious because it attacks each of its many targets a fixed number of times.

In summary, yellow portscan1 and green portscan2 have a higher chance of containing false positives according to its patterns of continuous randomized alerts and a concentrated list of targets. Owing to the inappropriate configuration of attack signatures, IDSs may mark normal communication between workstations and servers as threats or risks. On the other hand, red snort decoder shows symptoms of a malicious targeted attack (*e.g.*, worm) based on the observation of its sudden onset, highly structured alert pattern, and regular attack distribution.

Case 2—Banking IDSs Logs

The second case visualizes three days of IDSs data from the corporate network of a bank from April 5, 2012, to April 7, 2012, consisting of approximately 30 000 alerts with 11 alert types, as shown in Figure 5(a). The alert types are sorted according to PPMCC on temporal features. We choose orange DNS as the base type, other types are placed according to their temporal similarity coefficient with orange DNS.

Through the analysis of correlation arcs in the interactive core, the pattern of links in Figure 5(a) shows a division of correlated alert types into two groups. One group contains yellow IPC$, green NTMLSSP, and orange DNS; the other group contains blue IRC and seven other types. The first group of attacks generates persistent attacks for the entire logged duration. All three alert types in this group indicate the same server, IP 172.23.0.10, as the target of attacks. It is highlighted with a yellow circle in Figure 5(b). By dragging it to the interactive core, as shown in Figure 5(c), we found that all

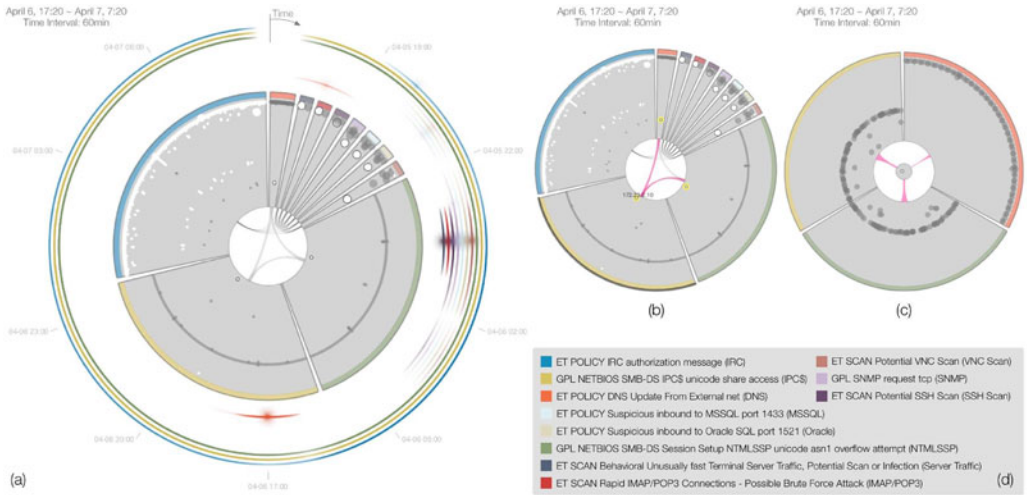

Figure 5. (a) Global view of IDSPlanet where the alert types are sorted according to PPMCC on temporal features. The orange DNS continent is placed at the 12 o'clock position. Other types are placed in descending order according to their temporal similarity coefficient with orange DNS. (b) When IP 172.23.1.10 is selected, its instances in other continents are highlighted. It suggests that the server IP receives attacks of three alert types. When the yellow IPC$ continent is selected, its correlation with other continents is shown as pink arcs in the interactive core. (c) Detailed view of IP 172.23.0.10. It acts as a target and receives three types of attacks. (d) Legend of color encoding of alert types, including its names and abbreviations.

attacking nodes have roughly the same medium size, which implies that they all use multiple forms of attack on the server. The input flows originated from yellow IPC$ and green NTMLSSP continent are thicker than in the orange DNS continent, meaning that these two alert types are detected more often. The pattern suggests either a distributed attack against a single point in the network or faulty false positive alerts caused by the inappropriate configuration of attack signatures.

For the second set, attacks occur within a short burst at midnight and displays a concerted high-intensity pattern. We use 60-min granularity from April 5, 20:00 to April 6, 4:00 during which a large number of attacks are reported, as shown in Figure 6(a). On the chrono ring, we observed that the blue IRC detonation occurs first, followed by other alert types detonating subsequently. It suggests that the blue IRC may have triggered the detonations of the other alerts. We found a notable pattern where the node with IP 172.23.231.69 is attacked as part of a large group during the blue IRC period. Surprisingly, it then takes control of the host and initiates a series of attacks targeting another server using seven other alert types. We, therefore, hypothesized that IP 172.23.231.69 is compromised during blue IRC attack and may be then used in an infected botnet attack against a single high-value target server.

By dragging the infected IP into the interactive core and playing the animation, two significant time frames drew our attention. We found that in the first time frame from 21:00 to 21:20 on April 5, as shown in Figure 6(b), IP 172.23.231.69 is attacked by IP 10.32.5.51, IP 10.32.5.52, and IP 10.32.5.54.

The flow arc between blue IRC continent and IP 172.23.231.69 in the core is pink, indicating a unidirectional flow. This meant that IP 172.23.231.69 only acts as a target during attacks. In the next time frame from 0:00 to 0:25 on April 6, as shown in Figure 6(c), it becomes an attacker that targeted IP 172.32.0.1. Based on the fact that botnet might communicate with external machines through the use of IRC, we concluded that the pattern in the second set of alerts reflects a situation where a botnet detects vulnerabilities within a system and initiates a successful infection. The botnet uses a blue IRC attack before using it as to attack another server on the same network. Examples of this harmful communication behavior may include hijacking websites, phishing for sensitive data, and attempting to steal sensitive information.

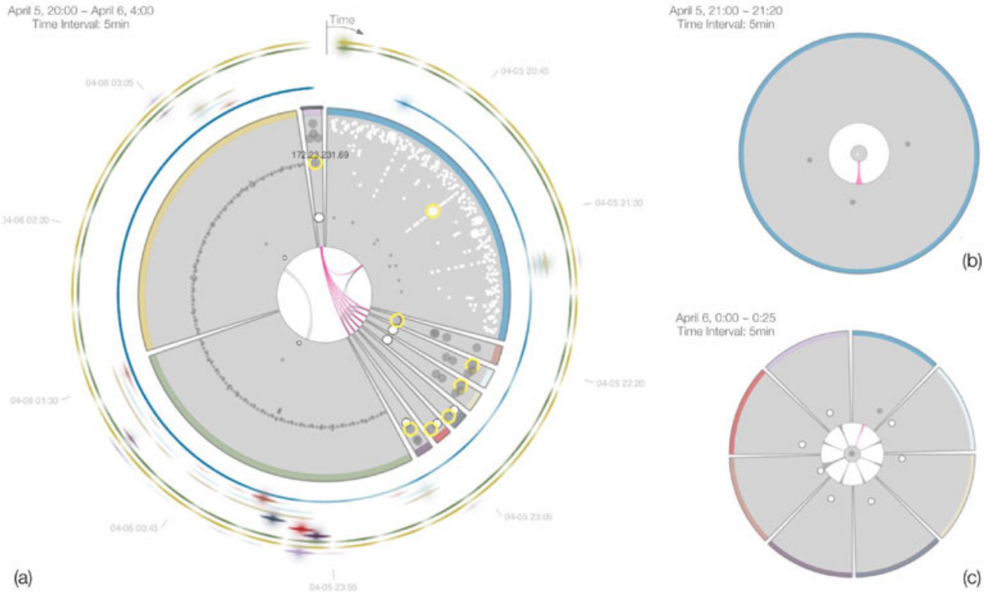

Figure 6. (a) Global view of IDSPlanet focuses on the time period when the high intensity burst occurs. The blue IRC ring shows that the IRC attack persists for a long duration and becomes serious. (b) From April 5, 21:00 to 21:20, IP 172.23.231.69 receives blue IRC attacks. (c) In the next time frame from April 6, 0:00 to 0:25, IP 172.23.231.69 becomes an attacker who initiates IRC and other types of attacks.

Expert Interview

To evaluate our approach further, we performed interviews with two potential users who are network administrators of a large scale university network. Expert A is responsible for the rapid remediation of security problems, whereas expert B is responsible for optimizing the network configurations. Before the interview, we first explained our approach, followed by an interface demonstration using the two user case studies. We then introduced two other IDS visualization tools, IDSRadar,[12] and NetSecRadar.7 The two experts were encouraged to compare these tools and express their opinions based on their user experience.

Both experts stated that IDSPlanet is a useful visual analysis tool for rapid attack detection. Expert A said,

> *The practicability of IDSs is relative low in the current network management. The main reason is that true threats are often obscured by a huge amount of false positives. I think this tool effectively helps analyze and estimate these alerts*

expert B added, *"It seems that IDSPlanet even helps find the inappropriate configuration of alert settings and the corresponding causes."*

When compared with IDSRadar and NetSecRadar, the experts thought that IDSPlanet is capable of displaying massive alerts in a more organized way.

> *The other two tools are more likely to result in visual clutter and thus obscure visual patterns, while IDSPlanet provides more aggregated information and supports understanding status of network events,*

said expert B. In addition, IDSPlanet provides more storytelling visual patterns. Expert A commented,

> *alert continents tell interesting stories through different visual patterns even though they contain numerous alerts and hosts. You could analyze groups of similar continents when set to different alert sorting method. You could find IP layouts of interest and deduce the 'killer' IP and the 'victim' IP. The design of IDSRadar and NetSecRadar obscures the underlying pattern by using lots of line connections.*

They preferred IDSPlanet's various interactions that offer flexible ways to filter data.

> *I found IP identification is very useful. When I click it, I immediately know if this IP is relevant to other alert types. It is especially helpful when you find an IP that acts as the attacker and the target at the same time,*

said expert A. Expert B commented on its multidepth investigation,

> *the two modes in interactive core is truly inspiring. It allows administrators to analyze detailed behaviors of a few specific hosts by switching from the global perspective to the individual perspective. It offers different perspectives to interpret data. One suggestion that I have is to add the functionality of browsing raw data, it would make the tool more practical.*

The experts suggested that the design of IDSPlanet could be improved. When comparing the container binding mechanism with the line binding mechanism, expert A said, *"I found that the linking style used by IDSRadar and NetSecRadar helps find correlations more efficiently when the alert size is small."* Expert B added, *"the container binding mechanism places multiple instances of an IP node in relevant continents; this technique might have negative effects on pattern recognition for small data."* Expert A pointed out that the design of chrono rings is abstract to some extent,

> *its denotation effects highlight the risky moment, but its mild trend can hardly show slight changes. The temporal presentation in IDSRadar is more recognizable and intuitive.*

Expert B thought that the visual encoding of IPs should be combined with more information in network topology,

> *for example, the shape of IP nodes could be used to indicate different network devices, the arrangement of IP nodes could be used to convey the information of subnetwork relationship.*

DISCUSSION

The result of our evaluation demonstrates that IDSPlanet effectively facilitates the detection of false positive alerts as well as latent threats. IDSPlanet provides storytelling visual patterns, which helps users analyze data and understand the situation. In this section, we discuss the advantages and limitations of our approach and suggest potential resolutions to improve it.

Scalability. Although IDSPlanet has the capability of visualizing massive network alerts and monitoring a great number of hosts, it experiences the same scalability issues as other traditional IDSs alert radial visualization. For example, as the number of alert types increases, the size of each continent reduces. When rendering IP nodes within a limited space, scaling up would introduce issues of overlapping and cause difficulty in interaction. For future work, the functionality of manual and automatic filtering should be added to the tool. It helps filter alert types, and thus reduce the data volume to be rendered.

Efficiency. IDSPlanet visualizes results produced by a large number of aggregated computations, the efficiency of data preprocessing would have an impact on the visual analysis. When monitoring large-scale network in real time, optimization of the preprocessing strategy is one of the key focuses.

Usability. One concern raised by experts about IDSPlanet is the absence of other types of network security data, such as traffic and firewall. We suggest to integrate multisource network security data and multiaspect views to IDSPlanet, which would facilitate the sophisticated anomaly detection and comprehensive security status assessment.

Effectiveness. In the two case studies, we performed several tasks in VAST Challenge. The results show that IDSPlanet assists in finding the false alerts and high-risk alerts in case 2011 and the targets of botnet attacks in case 2012, and it conveys valuable information through the presentation of visual patterns. In order to draw stronger conclusions regarding the effectiveness of our system, we plan to quantitatively compare IDSPlanet with IDSRadar and NetSecRadar using a series of predefined tasks in the future work.

CONCLUSION

This paper presents a novel radial visualization, IDSPlanet, which supports network administrators in obtaining insights into large logs of alerts generated by IDSs. First, IDSPlanet optimizes the advantages of radial visualization in being aesthetically concise, highly interactive. It effectively portrays correlations between *What, When, Where* attributes during network security analysis. Second, IDSPlanet replaces line binding with container binding, which not only minimizes visual clutter but also improves identification of relations in the *Where-What* domain. Finally, IDSPlanet combines dynamic host layout and strategic monitor design in radial visualization, providing richer interactions, visual patterns and analysis perspectives.

ACKNOWLEDGMENTS

This work was supported by the National Natural Science Foundation of China under No. 61872388, 61802283 and 61672538. We thank the anonymous reviewers for their comments and IEEE VAST Challenge for the data. Y. Zhao is the corresponding author of this work.

REFERENCES

1. H. Shiravi *et al.*, “A survey of visualization systems for network security,” *IEEE Trans. Vis. Comput. Graph.*, vol. 18, no. 8, pp. 1313–1329, Aug. 2012.
2. H. Koike and K. Ohno, “Snortview: Visualization system of snort logs,” in *Proc. ACM Workshop Vis. Data Mining Comput. Secur.*, 2004, pp. 143–147.
3. K. Abdullah, C. Lee, G. Conti, J. A. Copeland, and J. Stasko, “Ids rainstorm: Visualizing ids alarms,” in *Proc. IEEE Workshop Vis. Comput. Secur.*, 2005, pp. 1–10.
4. G. M. Draper, Y. Livnat, and R. F. Riesenfeld, “A survey of radial methods for information visualization,” *IEEE Trans. Vis. Comput. Graph.*, vol. 15, no. 5, pp. 759–776, Sep./Oct. 2009.
5. Y. Livnat, J. Agutter, S. Moon, R. F. Erbacher, and S. Foresti, “A visualization paradigm for network intrusion detection,” in *Proc. 6th Annu. IEEE SMC Inf. Assur. Workshop*, 2005, pp. 92–99.
6. M. Dumas, J.-M. Robert, and M. J. McGuffin, “Alertwheel: Radial bipartite graph visualization applied to intrusion detection system alerts,” *IEEE Netw.*, vol. 26, no. 6, pp. 12–18, Nov./Dec. 2012.
7. F. Zhou, R. Shi, Y. Zhao, Y. Huang, and X. Liang, “Netsecradar: A visualization system for network security situational awareness,” in *Proc. Cyberspace Safety Secur.*, 2013, pp. 403–416.
8. D. Keim *et al.*, “Monitoring network traffic with radial traffic analyzer,” in *Proc. IEEE Symp. Vis. Anal. Sci. Technol.*, 2006, pp. 123–128.
9. F. Zhou, W. Huang, Y. Zhao, Y. Shi, X. Liang, and X. Fan, “Entvis: A visual analytic tool for entropy-based network traffic anomaly detection,” *IEEE Comput. Graph. Appl.*, vol. 35, no. 6, pp. 42–50, Nov./Dec. 2015.
10. R. F. Erbacher, K. Christensen, and A. Sundberg, “Designing visualization capabilities for ids challenges,” in *Proc. IEEE Workshop Vis. Comput. Secur.*, 2005, pp. 121–127.
11. E. Bertini, P. Hertzog, and D. Lalanne, “Spiralview: Towards security policies assessment through visual correlation of network resources with evolution of alarms,” in *Proc. IEEE Symp. Vis. Anal. Sci. Technol.*, 2007, pp. 139–146.
12. Y. Zhao, F. Zhou, X. Fan, X. Liang, and Y. Liu, “Idsradar: A real-time visualization framework for ids alerts,” *Sci. China Inf. Sci.*, vol. 56, no. 8, pp. 1–12, 2013.

ABOUT THE AUTHORS

Yang Shi received the M.S. degree in entertainment technology from Carnegie Mellon University, Pittsburgh, PA, USA, and the Ph.D. degree in computer science and technology from Central South University, Changsha, Hunan, China. This presented work was done while she was with Central South University. She is currently an assistant researcher with Intelligent Big Data Visualiation Lab, Tongji University, Shanghai, China. Her research interests include information visualization and human computer interaction. Contact her at yangshi.idvx@tongji.edu.cn.

Ying Zhao is an Associate Professor with the School of Information Science and Engineering, Central South University, Changsha, China. His research interests include visual analytics and information security. He received the Ph.D. degree in computer science and technology from Central South University. Contact him at zhaoying@csu.edu.cn.

Fangfang Zhou is a Professor with the School of Information Science and Engineering, Central South University, Changsha, China. Her research interests include visualization. She received the Ph.D. degree in control science and control engineering from Central South University. Contact her at zff@csu.edu.cn.

Ronghua Shi is currently a Professor with the School of Information Science and Engineering, Central South University, Changsha, China. His research interests include information security, quantum cryptography, and network security. He received the B.S., M.S., and Ph.D. degrees in electrical engineering from Central South University, in 1986, 1989, and 2007, respectively. Contact him at shirh@csu.edu.cn.

Yaoxue Zhang is currently a Professor with the Department of Computer Science, Central South University, Changsha, China, and also a Professor with the Department of Computer Science and Technology, Tsinghua University, Beijing, China. He has authored or coauthored more than 200 technical papers in international journals and conferences, as well as nine monographs and textbooks. He is a Fellow of the Chinese Academy of Engineering and the President of the Central South University. His current research interests include computer networking, operating systems, ubiquitous/pervasive computing, transparent computing, and active services. He received the B.S. degree from Northwest Institute of Telecommunication Engineering, Xi'an, China, and the Ph.D. degree in computer networking from Tohoku University, Sendai, Japan, in 1989. Contact him at zyx@csu.edu.cn.

Guojun Wang is currently the Pearl River Scholarship Distinguished Professor with Guangzhou University, Guangzhou, China. He was a Professor with Central South University, Changsha, China; a Visiting Scholar with Temple University, Philadelphia, PA, USA, and Florida Atlantic University, Boca Raton, FL, USA; a Visiting Researcher with the University of Aizu, Aizuwakamatsu, Japan, and a Research Fellow with Hong Kong Polytechnic University, Hong Kong. His research interests include cloud computing, trusted computing, and information security. He received the B.Sc. degree in geophysics, the M.Sc. degree in computer science, and the Ph.D. degree in computer science from Central South University, Changsha, China. He is a distinguished member of the CCF, and a member of the ACM and IEICE. Contact him at csgjwang@gzhu.edu.cn.