InetVis, a Visual Tool for Network Telescope Traffic Analysis

Jean-Pierre van Riel* Barry Irwin†

Department of Computer Science
Rhodes University
Grahamstown, South Africa, 6140

Abstract

This article illustrates the merits of visual analysis as it presents preliminary findings using InetVis – an animated 3-D scatter plot visualization of network events. The concepts and features of InetVis are evaluated with reference to related work in the field. Tested against a network scanning tool, anticipated visual signs of port scanning and network mapping serve as a proof of concept. This research also unveils substantial amounts of suspicious activity present in Internet traffic during August 2005, as captured by a class C network telescope. InetVis is found to have promising scalability whilst offering salient depictions of intrusive network activity.

Categories and Subject Descriptors: C.2.0 [Computer-Communication Networks]: General—*Security and protection*; I.3.8 [Computer Graphics]: Applications; H.5.2 [Information Interfaces and Presentation]: User Interfaces—*Graphical User Interface (GUI)*.

General Terms: Security.

Keywords: InetVis, network visualization, network monitoring, network security, intrusion detection, visual intrusion signatures, Internet traffic analysis, 3-D (three-dimensional) scatter plot.

1 Introduction

The Internet is a hostile network environment. According to Symantec, “organizations received 13.6 attacks per day”, on average, between July 2004 and December 2004 [Turner et al. 2005: page 11]. During this period, the discovery of 7,360 new and variant forms of malicious software targeted for Microsoft Windows platforms was a record high; and the trend suggests a super-linear growth rate in the advent of viruses, worms and malicious scripts [Turner et al.: 10]. Presumably, these forms of malicious code account for the greater proportion of intrusive Internet traffic. Unlike ‘hands on’ human attempts to break into systems, viruses, worms, and other automated threats, do not intelligently discern between targets of high or low value. In addition to worms, bot networks constitute armies of compromised hosts remotely controlled to perform coordinated and distributed attacks. From this perspective, any system with Internet connectivity serves as a valuable target. Therefore, it is naive to assume that systems containing information of little worth are unlikely targets.

Despite the prevalence of these threats, many Internet users remain ignorant of these covert activities, or ignore them because of technical difficulties in patching and protecting systems against an increasing number of vulnerabilities. Between 1995 and the first quarter of 2005, the CERT® Coordination Centre (CERT/CC) received 17,946 vulnerability reports. Over 3,500 vulnerabilities were recorded for each of the last three years, namely 2002, 2003 and 2004 [CERT 2005]. As a consequence, a multitude of vulnerable systems are open to exploitation and serve as launching platforms for further attacks, thereby proliferating malicious activity.

2 Network Monitoring and Intrusion Detection

For computer security professionals and network administrators, discovering and identifying intrusive activity is far from trivial. Firstly, intrusive activity is intended to go unnoticed. Secondly, it is obfuscated by the complexity of network protocols. Thirdly, it becomes further obscured when hidden in large volumes of traffic.

Intrusion detection requires monitoring network traffic for 'fingerprints' such as anomalous traffic patterns or characteristic packets that are indicative of viral activity, worms, network scanning tools, and the like. Probing network scans across address space are indicative of intrusive reconnaissance, while port scans target particular hosts in search of vulnerability. This activity is not always obvious to detect. The information required to detect this scanning activity can be fragmented across numerous packets and intertwined amongst innocuous traffic. Added to this complication, sophisticated probing methods can cause packets to occur in random sequences with erratic and lengthy timing.

2.1 Conventional Methods

For the most part, current methods for network monitoring are textually based. To understand the output of these textual tools requires technical knowledge and tedious inspection of logs and alerts. This becomes unmanageable for large networks. “In a moderate-sized class B network, log files and packet traces may easily approach terabytes of information each day” [Ball et al. 2004: 56].

2.1.1. *Network Intrusion Detection Systems*

Network intrusion detection systems (NIDS) automate the task of monitoring networks by inspecting traffic for signs of intrusion. To detect suspicious traffic patterns requires correlating information from numerous network packets. Conventional NIDS, such as Snort [Snort], make use of intrusion signatures to detect characteristic traces of traffic known to arise as the result of intrusive activity. Alerts are logged when suspicious activity is found, thereby filtering out traffic considered innocuous and reducing the textual data for review. However, “most successful, algorithmic pattern matching systems are generally limited to recognizing patterns whose general form has been anticipated by the develop-

*g02v2468@campus.ru.ac.za

†b.irwin@ru.ac.za

AFRIGRAPH 2006, Cape Town, South Africa, 25–27 January 2006.
© 2006 ACM 1-59593-288-7/06/0001 $5.00

ers of the algorithms" [Fisk et al. 2003: 1]. A further concern with NIDS is a low rate of detection (false negatives) aggravated by high rate of false positives [Axelsson 2004].

2.2 Visual Methods

Unlike textual alerts and logs, visualizations can help viewers correlate many events at once. By employing visual metaphors that relate network events, the parallel nature of visual cognition allows the viewer to perceive trends and patterns present in network traffic, thereby gleaning a holistic insight. Some researchers proclaim that, "the human mind is capable of very fast visual processing outweighing the data mining capabilities of machines" [Yin et al. 2000: 26]. Furthermore, the pre-attentive nature of visual perception makes the viewer capable of noticing unanticipated patterns. "By incorporating human perception into the data mining process, researchers can detect patterns in data missed by traditional automatic data mining methods" [Teoh et al. 2004: 27].

2.2.1. Scalability Considerations

Visualizations offer a marked advantage over textual methods, but scalability is an apparent concern. Ball et al contend that "there are not currently any techniques that allow the home network or the external network to be very large". They estimate that their visualization, VISAUL, is capable of displaying 2,500 internal hosts versus 10,000 external hosts [Ball et al.: 56, 58]. This claim is suggestive and can be regarded with due scepticism because it does not detail exactly how many network events can be represented.

Scalability issues arise due to 'visual clutter' rendering a visualization display unintelligible as it becomes overwhelmed with too much data. Techniques such as aggregation and implicit filtering can be employed to alleviate visual clutter. However, this approach may unwittingly eliminate unanticipated patterns. Fisk et al state their preference for relying on the human visual system to aggregate the information. Via the 'immersive environment' of their 'Space Shield' visualization, details become apparent as the viewer explores within the visual environment [Fisk et al.: 2]. A second concern is performance. Visualising large amounts of data can make animated visualizations jittery and unresponsive. Scanmap3D version 2.1b is a prime example of this problem, as renders at a rate slower than normal playback speed while displaying only a few thousand events [Clark 2005].

Stephen Lau's 'Spinning Cube of Potential Doom' visualizes TCP connection attempts as an animated 3-D scatter plot of points [Lau 2003, Lau 2004]. By contrast, the majority of network visualizations employ lines to represent connectivity [Ball et al., Clark, Fisk et al., Teoh et al., Yin et al.]. Lines can be considered a suggestive metaphor for representing a connection. Unlike a point, a line conveys the idea of adjoining distinct entities. A point is one dimensional and fails to convey an idea of separation. Given this compromise, points consume less display space and offer better rendering performance. As will be shown in this paper, a 3D scatter plot is a highly scalable approach capable of representing in excess of 800,000 events.

3 InetVis Concept and Features

Although a promising concept, the Spinning Cube of Potential Doom lacked several useful features found in other visualizations, such as variable playback rate, adjustable time window, and filtering. With the intent of building upon Lau's work, InetVis represents network events as coloured points in an animated 3D scatter plot.

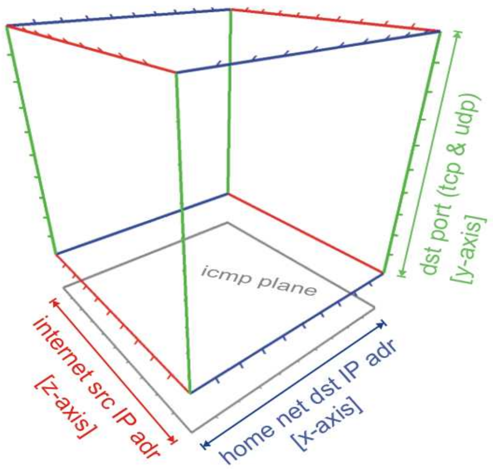

Figure 3-1: InetVis plotting scheme

3.1 Plotting Scheme

The simple plotting scheme for InetVis is illustrated in Figure 3-1. The visualization is intended for viewing events traversing a boundary between an internal (home) network versus the external Internet. InetVis utilizes libpcap (a low-level packet capture library) to extract header information [TCPDump]. The visualization plots TCP and UDP packets within the cube, and plots 'portless' ICMP packets to a flat plane below the cube.

Points are spatially located according to three axes:

- Traffic destined to the home (internal) network range is scaled and mapped to the *X-axis* by the destination IP address.
- The entire Internet is plotted along the *Z-axis* according to the source IP address of originating traffic.
- The *Y-axis* relates the destination port in the case of TCP and UDP traffic.

In comparison, the Spinning Cube of Potential Doom visualizes TCP/IP connection attempts by utilizing TCP connection handshake information that is leveraged from Bro IDS logs [Bro]. It renders unsuccessful attempts in a bright rainbow colour map according to detonation port number, and greys successful connection attempts. However, it does not plot UDP and ICMP based traffic, thereby limiting its scope of observation.

Like Lau's Cube, InetVis colours points according to the destination port. As an extension to dimensionality, it can also colour by source port to note scans that originate from one source port, by protocol to distinguish traffic types, or by packet size since probe packets are typically small and often the same size. The point size can be reduced for visualising a large number of events, or increased for prominent viewing of specific events (typically with a filter applied).

3.2 Features

A timer animates the replay of capture files according to playback rate. The maximum replay rate is 86400x (1 day per second) and the minimum is 0.001x (1 millisecond per second). Like the

Space Shield [Fisk et al.], the ability to scale time allows viewers to review traffic in a ‘quick search’ fashion. In searching, the viewer can skip past disinteresting periods of inactivity. When something of interest is noticed, a seek bar can be used to jump back to an appropriate replay position, and the replay speed can be slowed down for meticulous inspection. InetVis is also capable of displaying real-time traffic capture from a live network interface. (Both live monitoring and variable rate replay are noted by Stephen Lau as desirable extensions [Lau 2003].)

The idea of a variable 'time window' is adopted from VisFlowConnect [Yin et al.]. A time window specifies the length of time to continue displaying an event after it has occurred – effectively, a form of time filtering. This is a useful extension because a narrow time window reveals the faster scans whilst a wider window allows slower scans observable.

Similar to the ‘Space Shield' visualization [Fisk et al.], the navigation controls of InetVis allow the viewer to explore within the visual environment by using the mouse to move, rotate and zoom in and out. This facilitates focusing closer on interesting visual phenomena within the display. For further focus on specific phenomena, dynamic filtering can be performed by entering BPF (Berkley Packet Filter) expressions that can filter out unwanted traffic.

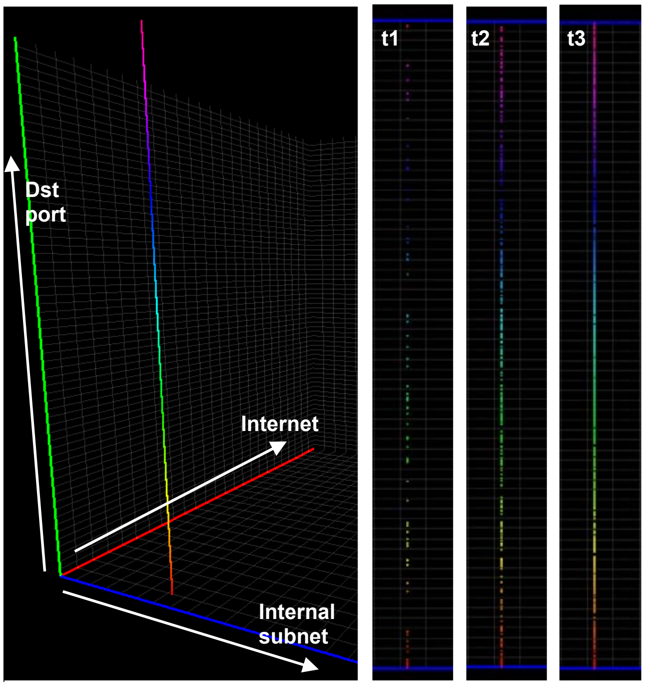

Figure 4-1: Visual signature of a port scan

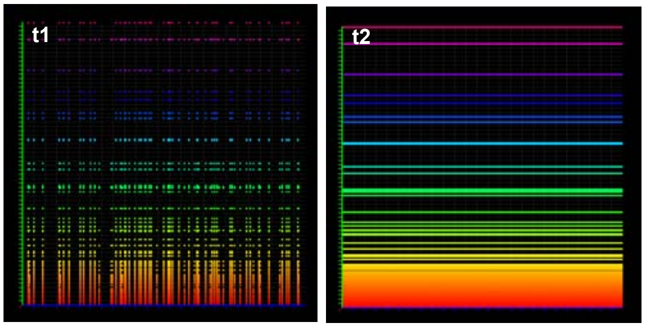

Figure 4-2: Visual signature of a ‘network sweep’

InetVis also offers two projection modes, namely perspective projection and orthographic projection. A perspective projection conveys three dimensionality and depth to reflect a sense of spatial locality. An orthographic projection is useful for obtaining an accurate reflection of geometry and obtaining flat planar views along a particular axis. The reference frame around InetVis includes flexibility to show and hide components such as primary axes, bounding axes, grids and markers. It also has the ability to specify the number of divisions for grids and markers for each axis. These features assists the viewer with an indication on the address and port values.

4 Results and Analysis

4.1 Visual Signatures

The NMap [NMap] scanning tool is used to generate traffic capture data of scanning activity, isolated within a controlled network environment. A port-scan of a single host randomly covers all ports, and Figure 4-1 shows how it progressively forms a solid vertical line from times t1, t2 and t3. Figure 4-2 presents a visual signature for a '’network sweep', viewed with orthographic projection down the z axis. In this type of scan, the network range is randomly covered in segments as it checks for common open service ports. These visual signatures serve as examples of what some probing activity looks like, and appear as anticipated, confirming that InetVis plots and conveys the traffic as intended. Ethereal [Ethereal], a textually based network packet analyser, was used corroborate the proper functioning of the plotting scheme.

4.2 Internet Traffic Analysis with a Network Telescope

A network telescope (also referred to as a darknet) is an 'unused' address range. It does not offer Internet services, nor does it use any. Hence 'darknet', because there ought not to be any observable traffic. Monitoring from this perspective gives a good indication of rogue and misplaced Internet traffic. During August 2005, 867,085 packets were captured whilst monitoring traffic across a class C darknet. As seen in Figure 4-3, a wealth of horizontal lines occurs across the x axis, evidence of network probing attempts that scan across the internal network address range in an attempt to discover vulnerable hosts. Since the network is empty, no port scans are present as scanning techniques first establish the presence of a host before expending time scanning ports.

4.2.1. Bands of Activity

In Figure 4-4a, the orthographic projection down the z axis provides a 2-D perspective of the x-y plane. The vertical bands show that relatively few external Internet ranges contribute heavily to the unwarranted Internet traffic. This is a useful way of identifying hostile IP ranges of the Internet.

4.2.2. Anomalous Diagonals

Network probing can be done in a variety of covert ways that do not form obvious horizontal or vertical lines. Figure 4-4b is an orthographic projection down the z axis and reveals a 45° diagonal that runs across the first half section of the internal network address space. The line is formed by incrementing port numbers in relation to network address. The animated replay indicates that this occurs in a slow random fashion over a number of days. By using a BPF filter expression to view three source networks

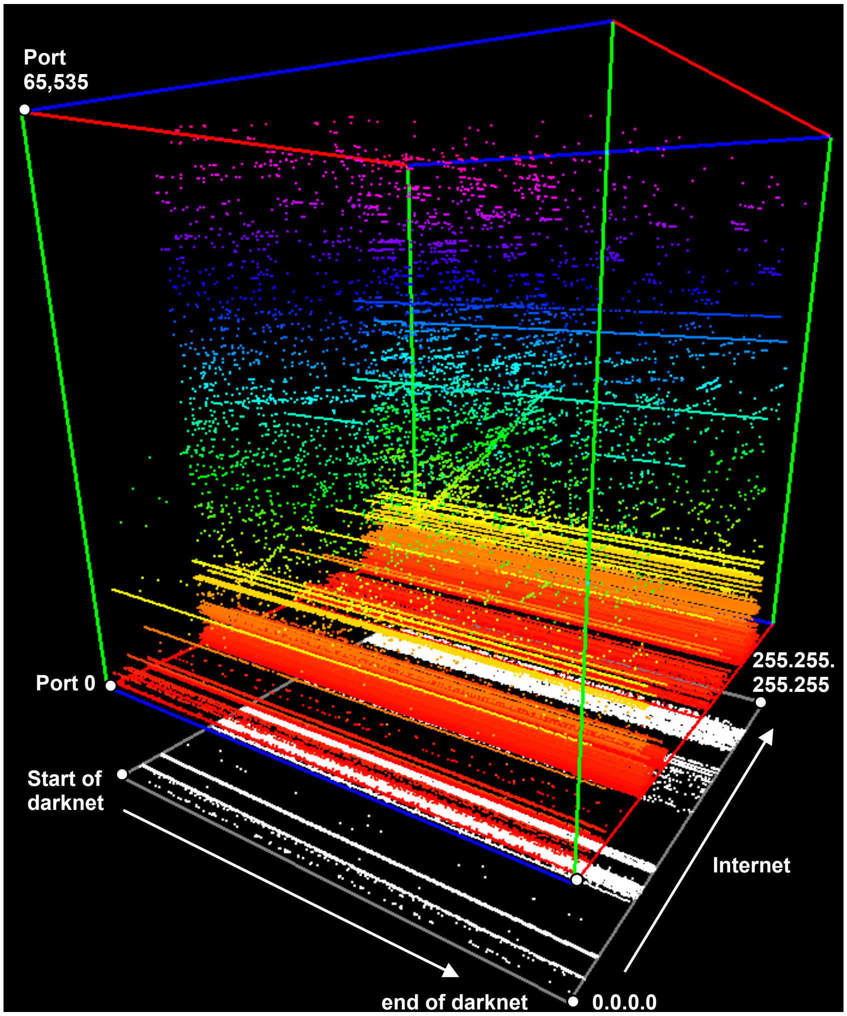

Figure 4-3: 857,085 packets observed during August 2005

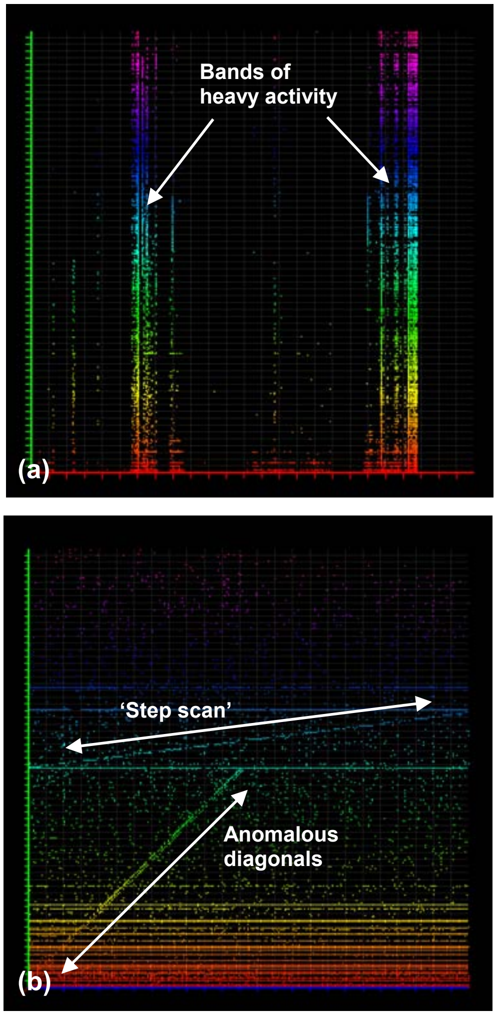

Figure 4-4: Orthographic Projections

(61.0.0.0/8, 220.0.0.0/7 and 218.0.0.0/8) with all other traffic excluded, three distinct diagonal lines clearly emerge. This suggests that the anomalous phenomena are not isolated events, and could possibly be the result of a particular type of worm or scanning tool (or a peculiar network error).

4.2.3. Step Scan

Another interesting phenomenon in Figure 4-4b, appears to scan segments of address space on the same port, stepping up the port range between segments. The scan progresses very slowly throughout August and appears incomplete because it continues into September. Like the anomalous diagonals it also originates from two distinct Internet sources, indicating it is not a lone incident as the filtered image as Figure 4-5 clearly shows.

4.2.4. Creepy Crawly Scan

A covert probing scan, dubbed the 'creepy crawly' scan is illustrated in Figure 4-6 at 18 hour intervals with a 36 hour time window (viewed with a filtered orthographic projection down the z axis). A wider time window of 15 days results in a solid line showing that the scan comprehensively covers every IP address. The narrower time window illustrates how the 'creepy crawly' traverses the address range in segments forming a dashed line that persists throughout August.

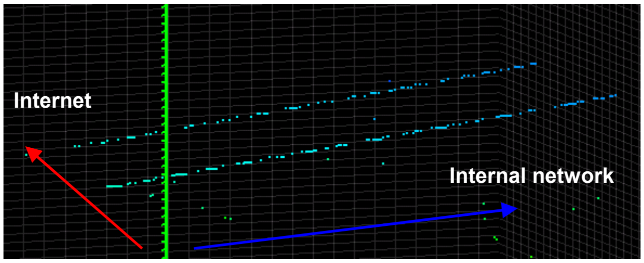

Figure 4-5: Two distinct 'step scans'

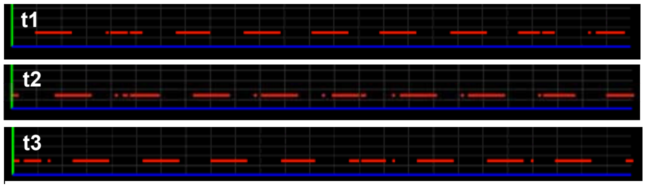

Figure 4-6: 'Creepy crawly' scan progression

4.3 Performance and Scalability

InetVis readily replayed and displayed all the data in a 121.5 MB capture file of 857,085 packets (as shown in Figure 4-3). All of this could rapidly be reviewed in less than a minute at a replay rate of 86400x and a time window of a month, whereby an approximate average of 30,000 new events were displayed every second (3.0GHz Pentium IV, 1GB RAM, NVida GeForce PCX750). The visualization maintained a fixed frame-rate of 25 frames per second until approximately 450,000 events. After this, it was unable to maintain 86400x, but was still manageable rendering approximately 10 frames per second. If a slower replay rates or a narrower time window was selected, no performance impact was noted.

5 Possible Extensions

At present, InetVis functions well, but can be improved. Visualizations, such as VISUAL [Ball et al.], VisFlowConnect [Yin et al.], and the 'space shield' visualization [Fisk et al.], aggregate packet events into connection flows, whereas Scanmap3D and InetVis represent packet events. Connection flows can reduce the amount of graphical elements that need to be rendered, thereby improving rendering performance – especially when complex objects are to be drawn. The performance impact of rendering packet events is not overly severe for InetVis because rendering points is less costly in contrast to the more elaborate visual metaphors. InetVis simply redraws points over each other as packet events occur, effectively conveying traffic flow information.

Extensions to the plotting scheme could also further improve the effectiveness of InetVis. The ability to scale the port range (y-axis) with a logarithmic plot would help spread out the clutter that resides in heavily utilized lower port range (i.e. ports 0-1024). Also, mapping colour by event time-stamp age may visually convey the chronological order of events. Opacity could be used in this regard, where older events appear more transparent.

Visual drill downs should also prove useful. Given that the Y-axis plots 65535 port numbers, some of the smaller port scans are unnoticeable, and so the ability to scale and view a specified port range would be useful. Furthermore, selecting and scaling an external Internet range and scaling it to the red axis would allow closer investigation of heavy sources of activity (as is done in VisFlowConnect [Yin et al.]).

6 Conclusion

The 3D scatter-plot of packet events is found to be impressively scalable and shows promise in comparison to the more elaborate visualizations with metaphorically complex representations. Features like variable replay rate, adjustable time window, navigation (for deeper exploration), and dynamic filtering enhance the ability to pick up and detect a varied range of scanning activity. As seen in the images of Section 4, InetVis saliently conveys traffic patterns and reveals suspect network probing – intrusive reconnaissance that often precedes an attack. The images presented from the darknet traffic capture clearly attest that the Internet is a hostile a network environment. In conclusion, this preliminary research illustrates the deft application of visualizing network traffic by making intrusive activity intended to be invisible visible.

References

AXELSON, S. 2004. Combining a Bayesian Classifier with Visualization: Understanding the IDS. In *Proceedings of the 2004 ACM workshop on Visualization and data mining for computer security*, ACM Press, New York. 99-108.

BALL, R., FINK, G.A., NORTH, C. 2004. Home-centric visualization of network traffic for security administration. In *Proceedings of the 2004 ACM workshop on Visualization and data mining for computer security*, ACM Press, New York. 55-64.

BRO. 2005. Bro NIDS. http://www.bro-ids.org (03/06/2005).

CERT/CC – CERT® COORDINATION CENTER. 2005. Carnegie Mellon Software Engineering Institute. Carnegie Mellon University. http://www.cert.org/ (27/05/2005).

CLARK, D. 2005. Scanmap3D-2.1b and Scanmap3D-3.0. Open source Snort visualization tool. http://scanmap3d.sourceforge.net/ (29/05/2005).

ETHEREAL. 2005. Packet capture and textual review tool http://www.ethereal.com/ (07/11/2005).

FISK, M., SMITH, S.A., WEBER, P.M., KOTHAPALLY, S., CAUDELL, T.P. 2003. Immersive Network Monitoring. PAM2003 – Passive and Active Measurement 2003, NLANR/MNA (National Laboratory for Applied Network Research / Measurement and Network Analysis Group). http://public.lanl.gov/mfisk/papers/pam03.pdf (25/05/2005).

LAU, S. 2004. The Spinning Cube of Potential Doom. In *Communications of the ACM* archive, Volume 47, Issue 6, ACM Press, New York. 25-26.

LAU, S. 2003. The Spinning Cube of Potential Doom. Online article, November 10th, 2003. National Energy Research Scientific Computing Center (NERSC) website: http://www.nersc.gov/nusers/security/TheSpinningCube.php (18/09/2005).

NMAP. 2005. NMap network scanning utility. http://www.insecure.org/nmap/ (07/11/2005).

SNORT. 2005. Snort network intrusion prevention and detection system. http://www.snort.org/ (12/11/2005).

TCPDUMP. 2005. TCPDump packet capture utility and LibPCap packet capture library. http://www.tcpdump.org (06/11/2005).

TEOH, S.T., MA, K-L., WU, S.F., JANKUN-KELLY, D.T.J. 2004. Detecting Flaws and Intruders with Visual Data Analysis. In *IEEE Computer Graphics and Applications*, Volume 24, Issue 5.

TURNER, D. (EXEC. ED.), ENTWISE, S. (ED.), ET AL (SYMANTEC). 2005. Symantec Internet Security Threat Report, Trends for July 04 – December 04, Volume VIII, March 2005. http://enterprisesecurity.symantec.com/content.cfm?articleid=1539 (20/04/2005).

YIN, X., YURCIK, W., TREASTER, M., LI, Y., LAKKARAJU, K. 2004. VisFlowConnect: netflow visualizations of link relationships for security situational awareness. In *Proceedings of the 2004 ACM workshop on Visualization and data mining for computer security*, ACM Press, New York. 35-44.