
Abstract: Approaches to Anomaly Detection using Host Network-Traffic Traces

John Mark Agosta, Jaideep Chandrashekar,
Frédéric Giroire, Carl Livadas & Jing Xu *

Intel Research
Intel Corporation, Santa Clara, CA 95054

1 Distributed Anomaly Detection

As part of a larger project on distributed anomaly detection using local host traffic, (See Agosta *et al.* [2005], Dash *et al.* [2006]) we review some current analysis and modeling of enterprise traffic flows. The origins of this work grew out of a project on *individual* host anomaly detection for worm spreading behavior. The *Distributed Detection and Inference* (DDI) project improved on this to consider anomalies among a host's neighbors, greatly increasing detection accuracy. To ground the project we've completed a comprehensive traffic data collection effort over several weeks on several hundred individual enterprise hosts. Here we review several threads of the work, and make some observations for better detection methods relevant to the overall project.

Strong dependencies among network flows lead to bursty network traffic, which makes modeling and prediction hard. This is true to the extent that machine idle and off periods are hardly evident in network traffic traces. (Is a quiet period due to the machine being off? Is a burst of activity just due to a periodic background service?) A novel aspect of our trace collection is to record machine state, e.g. user and power events, in addition to conventional network packet signatures. We consequently were able to infer classes of (in-)activity that otherwise taint traffic statistics.

2 Unruly Traffic Distributions

The well-know heavy-tailed property of network traffic is apparent both in self-similarity of network traffic at varying timescales, and Hurst parameter estimates of flow counts in time intervals. We've discovered another manifestation of the property in the threshold settings to achieve a desired detection rate: The threshold setting varies less than proportionally with the time interval with which flow counts are binned. Intuitively this is seen as a consequence of the bursty nature of the traffic. Short duration spikes contain a significant portion of the counts seen in any interval, thus the counts are largely due to the intensity of the spikes, and less sensitive to the low level of activity that fills the duration of the interval.

We are able to quantify this property by fitting mixture distributions to traffic histograms using EM. This reveals that in some ranges of the distribution tail, a power-law fit is not appropriate, and, surprisingly, more conventional distributions fit better. Hence threshold

*Within the larger program we eagerly acknowledge the guidance that Eve Schooler and Nina Taft have offered, and the numerous contributions of Denver Dash, in addition to those of all team members cited as our co-authors in the references.

setting becomes more predictable, both due to the “tighter” distribution in the range of interest and due to the decreased sensitivity to the time-interval size. The possibility that gross traffic can be classified or conditioned on type, service or user to find such better behaved characteristics is a parallel area of investigation, described in the following section.

3 Conditioning Anomalies on Classes of Traffic

Traditional anomaly detectors that we see on end-hosts today are very simplistic: the time series of a traffic/protocol feature (or features) are thresholded upon a pre-assigned value that is constant over users. Instances where the threshold is crossed are flagged. In most enterprise deployments, an alert is then sent (possibly batched) to a central console.

The trouble with existing anomaly detectors is that they take a rather narrow view of what an anomaly means. That is, they do not attempt to look at known dependencies and other sources of information, relying solely on network traffic observations over a short time scale, to “explain away” observed anomalies. As an example, a transient spike in HTTP (web) traffic would look anomalous, but less so when coupled together with the observation that the user is interacting with a web-browser at the time. Many other examples exist, including learning dependencies (or temporal sequences) in traffic and applying the same toward anomaly detection.

Clearly, a wider view of the problem is required, possibly incorporating other sources of data at different levels. We explore the hypothesis that a vast majority of “spikes” in the traffic are easily explainable, by network management traffic, periodic user activity, or other known phenomena. Details are given in Giroire *et al.* [2007]. Notionally, we want a methodology to identify spikes in traffic and to quickly classify them as benign or suspicious. The end-goal is to build more specific detectors that have a very low false positive rate without compromising detector accuracy.

4 Future Directions

Since the distributed detection model does inference at both the host level and aggregate levels, characteristics, such adaptation can be placed where they have best effect. This was explored for an adaptive host detector in Agosta *et al.* [2007]. Similarly there is a tradeoff between decision thresholds at host and aggregate levels where it is shown that a “chatty” host detector, leads to better system accuracy. In other areas of investigation we have begun experimenting with Continuous Time Bayes Network (CTBN) models, to relax the discrete time interval constraint in current models.

References

- John Mark Agosta, Denver H. Dash, , Eve Schooler, and Branislav Kveton. Distributed network attack detection. In *ARCS, Santa Fe Institute*, November 2005.
- John Mark Agosta, Carlos Diuk, Jaideep Chandrashekar, and Carl Livadas. An adaptive anomaly detector for worm detection. In *Proceedings of sysML-07*, 2007.
- Denver Dash, Branislav Kveton, John Mark Agosta, Eve Schooler, Jaideep Chandrashekar, Abraham Bachrach, and Alex Newman. When gossip is good: Distributed probabilistic inference for detection of slow network intrusions. In *AAAI06*, Boston, MA, USA, July 2006. AAAI Press.
- Frédéric Giroire, Jaideep Chandrashekar, Thomas Karagiannis, Dina Papagiannaki, Nina Taft, and Eve Schooler. A case for personalization of end-host anomaly detectors. In submission, July 2007.