Nightlights: Entropy-based Metrics for Classifying Darkspace Traffic Patterns

Tanja Zseby[1], Nevil Brownlee[2,3], Alistair King[3], and K.C. Claffy[3]

[1] Vienna University of Technology, 1240 Vienna, Austria
[2] University of Auckland, Auckland 1010, New Zealand
[3] CAIDA, UC San Diego, CA 92093, USA

An IP darkspace is a globally routed IP address space with no active hosts. All traffic destined to darkspace addresses is unsolicited and often originates from network scanning or attacks. A sudden increases of different types of darkspace traffic can serve as indicator of new vulnerabilities, misconfigurations or large scale attacks. In our analysis we take advantage of the fact that darkspace traffic typically originates from processes that use randomly chosen addresses or ports (e.g. scanning) or target a specific address or port (e.g. DDoS, worm spreading). These behaviors induce a concentration or dispersion in feature distributions of the resulting traffic aggregate and can be distinguished using entropy as a compact representation. Its lightweight, unambiguous, and privacy-compatible character makes entropy a suitable metric that can facilitate early warning capabilities, operational information exchange among network operators, and comparison of analysis results among a network of distributed IP darkspaces.

Using traffic from five months from a large /8 darkspace monitor, we investigate the use of an entropy vector for IP darkspace traffic classification. As reference we perform an in-depth analysis with the tool iatmon [2] to classify the traffic into 15 different traffic types. We then compare our entropy results to the detailed iatmon analysis. We use the approach and the formula presented in [3] to calculate an estimate for Shannon entropy from IP address and port number distributions: $H(X) = -\sum_{i=1}^{N} \frac{n_i}{S} \cdot log_2\left(\frac{n_i}{S}\right)$, where $i...N$ are the different bins in the frequency distribution (IP addresses or ports). n_i denotes the number of packets that belong to bin i (e.g. all packets with port number 445). X denotes the distribution of a feature (sIP, dIP, $sPort$ or $dPort$), formed by the frequencies $n_1, ...n_N$ of all bins N. S denotes the total number of observations (packets received) in the time interval. In the /8 darkspace we get $N = 2^{24}$ possible destination addresses and therefore $H(dIP)_{max} = 24$.

For each time interval t we compute an entropy vector that contains the four entropy values: $\boldsymbol{H}_t = [H_t(sIP), H_t(dIP), H_t(sPort), H_t(dPort)]$. We expect different changes in the entropy vector ($+\Delta h$ increase, $-\Delta h$ decrease), which provide a unique signature for different darkspace events.

A **multi-source horizonal scan** disperses source IPs and source ports, but concentrates the destination port distribution. $H(dIP)$ dispersion is already close to the maximum (24 bits) in darkspace data, so we expect only small effects on $H(dIP)$ (denoted by $(+\Delta h)$): $\Delta\boldsymbol{H}_t = [+\Delta h, (+\Delta h), +\Delta h, -\Delta h]$. **Backscatter** traffic occurs if victims of a DoS attack are attacked with spoofed source addresses and reply to those spoofed addresses. For backscatter we expect a

concentration of the source IP distribution, because a lot of traffic is sent from relatively few (victim) sources. We expect a source port concentration toward the port that was used as destination port to attack the victim machine, whereas destination ports disperses if the attacker used random source ports. Again we expect only a small effect on $H(dIP)$: $\Delta \boldsymbol{H}_t = [-\Delta h, (+\Delta h), -\Delta h, +\Delta h]$. For a **distributed probe** we expect a source address and source port dispersion, caused by the use of bots or spoofed addresses, and a concentration of destination address and port toward the target: $\Delta \boldsymbol{H}_t = [+\Delta h, -\Delta h, +\Delta h, -\Delta h]$.

We analyse darkspace traffic from 5 month: Nov 2008 (Conficker outbreak), Jan/Feb 2011 and Jan/Feb 2012. We first classify the traffic into 15 traffic classes using an in-depth analysis with iatmon [2]. The output serves as a baseline against which to evaluate our entropy-based inferences. Then we calculate one entropy vector for each hour interval, using the tool Corsaro[1] and the statistical package R[2]. We then compare the detailed iatmon results with the more lightweight entropy analysis to see if new events follow the expected entropy patterns and thus can be classified based on entropy.

Multi-Source Scans: The detailed iatmon analysis of Nov 2008 data reveals an increase of *TCP horizontal scan* packets, caused by the Conficker outbreak, where hosts began to scan port 445 trying to spread the worm [1]. The outbreak of the new worm is clearly visible in entropy vectors, following the expected entropy pattern for a multi-source scan.

Backscatter is captured effectively by the entropy vector in our experiments. Figure 1 shows the results from Feb 2012 as an example. It shows the entropies (1st and 2nd graph) and the amount of backscatter packets according to iatmon's classification (3rd graph). As expected, if backscatter increases we observe an increase in $H(dPort)$, a decrease in $H(sPort)$ and $H(sIP)$ and no significant changes for $H(dIP)$. Table 1 lists the correlation coefficients between entropy and backscatter traffic. The observations also conform to the expected behavior. While the increase in backscatter traffic does not always affect the overall packet count (last row in table 1), it always shows significant changes in entropy.

Large **probing** events are also visible in entropy. The iatmon analysis for Jan 2011 shows a large distributed probe originating from many sources (spoofed and/or bots) directed to a specific IP address and port. The new probe traffic is clearly visible in the entropy statistics (figure 2); the increase in new sources

[1] http://www.caida.org/tools/measurement/corsaro/
[2] http://www.r-project.org/

corr. coeff.	**Jan11**	**Feb11**	**Jan12**	**Feb12**	*corr. coeff.*	**Jan11**	**Feb11**	**Jan12**	**Feb12**
bs,H(sIP)	-0.48	-0.38	-0.37	-0.75	*bs,H(dIP)*	0.14	0.29	0.36	0.44
bs,H(sPort)	-0.60	-0.52	-0.58	-0.83	*bs,H(dPort)*	0.69	0.62	0.69	0.91
bs,pktcount	0.28	0.39	0.48	0.76					

Table 1. *Correlation coefficients for time series of amount of backscatter (bs) traffic (as seen by iatmon) and entropy (rows 1,2) and bs traffic and packet count (row 3)*

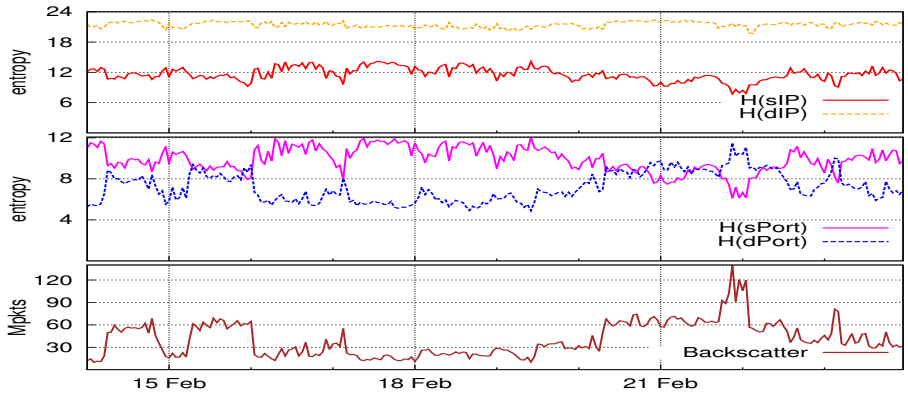

Fig. 1. *Entropy correlation with backscatter traffic (February 2012), showing IP address entropies $H(sIP)$, $H(dIP)$ (1st graph), port entropies $H(sPort)$, $H(dPort)$ (2nd graph), amount of backscatter traffic according to iatmon analysis (3rd graph).*

drives up $H(sIP)$. High concentration of traffic to one address and one port causes $H(dIP)$ and $H(dPort)$ to drop significantly.

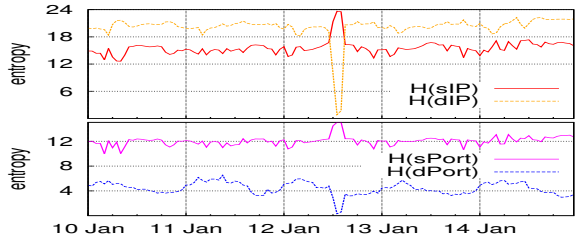

Fig. 2. *Entropy during TCP Probe*

Our results show that entropy-based metrics can reveal noteworthy events in IP darkspace. We plan to further investigate the use of entropy to also detect smaller changes or nested events, and evaluate the utility of this method for early warning and privacy-respecting information sharing among darkspace operators.

References

1. Emile Aben. Conficker/Conflicker/Downadup as seen from the UCSD Network Telescope. Technical report, CAIDA, February 2009.
2. Nevil Brownlee. One-way Traffic Monitoring with iatmon. In *13th Passive and Active Measurement Conference (PAM 2012)*, 2012.
3. Anukool Lakhina, Mark Crovella, and Christophe Diot. Mining Anomalies Using Traffic Feature Distributions. *SIGCOMM Comput. Commun. Rev.*, 35(4):217–228, August 2005.