**QUEEN'S
UNIVERSITY
BELFAST**

Dynamic Opcode Analysis of Ransomware

Carlin, D., O'Kane, P., & Sezer, S. (2018). Dynamic Opcode Analysis of Ransomware. In *Proceedings of International Conference on Cyber Security and Protection of Digital Services (Cyber Security 2018)* Institute of Electrical and Electronics Engineers Inc.. <https://doi.org/10.1109/CyberSecPODS.2018.8560667>

Published in:

Proceedings of International Conference on Cyber Security and Protection of Digital Services (Cyber Security 2018)

Document Version:

Peer reviewed version

Queen's University Belfast - Research Portal:

[Link to publication record in Queen's University Belfast Research Portal](#)

Publisher rights

© 2018 IEEE.

This work is made available online in accordance with the publisher's policies. Please refer to any applicable terms of use of the publisher.

General rights

Copyright for the publications made accessible via the Queen's University Belfast Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The Research Portal is Queen's institutional repository that provides access to Queen's research output. Every effort has been made to ensure that content in the Research Portal does not infringe any person's rights, or applicable UK laws. If you discover content in the Research Portal that you believe breaches copyright or violates any law, please contact openaccess@qub.ac.uk.

Dynamic Opcode Analysis of Ransomware

Domhnall Carlin*, Philip O’Kane†, Sakir Sezer‡

Centre for Secure Information Technologies, Queen’s University, Belfast, Northern Ireland

Email: *dcarlin05@qub.ac.uk, †p.okane@qub.ac.uk, ‡s.sezer@qub.ac.uk

Abstract—The explosion of ransomware in recent years has served as a costly reminder that the malware threatscape has moved from that of socially-inept hobbyists to career criminals.

This paper investigates the efficacy of dynamic opcode analysis in distinguishing cryptographic ransomware from benignware, and presents several novel contributions. Firstly, a new dataset of cryptoransomware dynamic run-traces, the largest of its kind in the literature. We release this to the wider research community to foster further research in the field.

Our second novel contribution demonstrates that a short run-length of 32k opcodes can provide highly accurate detection of ransomware (99.56%) compared to benign software. Third, our model offers a distinct advantage over other models in the literature, in that it can detect a form of benign encryption (i.e. file zipping) with 100% accuracy against not only ransomware, but also the non-encrypting benignware in our dataset. The research presented here demonstrates that dynamic opcode tracing is capable of detecting ransomware in comparable times to static analysis, without being thwarted by obfuscation tactics.

I. INTRODUCTION

Beginning with Brain, the first PC virus, in 1986, the modern malware pandemic has escalated at an alarming pace. This deliberately disruptive software was formerly the preserve of hobbyists seeking infamy and bragging rights, but has moved to enable a worldwide criminal ecosystem, providing a rich platform for the extortion of individuals, businesses, the public sector and even sovereign states. McAfee [1] state that 1.5 million new ransomware samples were found in Q3 of 2017, a rise of 36% on the previous quarter. This drove the total number of ransomware samples held by McAfee to breach 12 million.

A. The rise of ransomware

Ransomware is deliberately malicious software that denies or limits access to a machine or data, often with just enough access to facilitate payment of the ransom demand. The perfect storm of anonymous payment channels and internet access, available encryption, and increased attack surface has given criminals the opportunity to cash in on the recent data explosion. The incessant demand for connected devices and the subsequent exponential rise in data volumes has fast outpaced security considerations, leaving low-hanging targets for the exploitative[2].

Early versions of ransomware were largely unsuccessful, as they faltered on one or more of the ‘perfect storm’ factors mentioned previously. The AIDS PC virus from 1989 was distributed via floppy disc, but used easily-breakable encryption and easily-traceable payment methods (a post office box). However, even as recently as 2010, ransomware was viewed by some authors as a non-event:

‘The ransomware phenomenon is a reality that has to be monitored but in some ways it is not a mature and complex enough activity that deserves such communication around it. Ransomwares[sic] as a mass extortion mean is certainly doomed to failure...may means[sic] that criminals have evolved to something else and other sources of income.’[3, P.90]

Latter-day ransomware has evolved dramatically, increasing in aggression, sophistication and potential with each strain and iteration.

In [4] Symantec listed the average ransom demand as \$294 in 2015, rising to \$679 in 2016. This, coupled with the fluctuations in Bitcoin exchange rates, has helped to drive the rise in *ransomware-as-a-service (RaaS)*. Here, customised binaries are available for as little as \$50 and a 10% cut of the proceeds. With such a high potential return, and such ease of access to the offending ransomware, the onslaught may only be beginning. A shift from a shotgun approach to specific big-ticket targets has led to the ransom of the San Francisco metro system, private healthcare practices, the UK National Health Service, hotel key-cards, and specific Big Data stores. With the data explosion and rapid increase in connected devices, the prospect of escalating attacks is becoming all too real.

The remainder of this paper is presented as follows: Section Two provides highlights of the literature to date on dynamic analysis of malware and ransomware, and the objectives of the present work. Section Three briefly describes the methods used to generate the dataset for the present research and the experimental analyses. Section Four presents the results of our experiments and Section Five provides conclusions on the importance of the results.

II. RELATED WORKS

A. Dynamic opcode analysis of malware

Research into overcoming issues inherent in signature detection and obfuscation has focused on observing behaviours of the malware at run-time. Opcodes (operation codes) are the readable machine language instructions that perform CPU operations. Monitoring the run-time opcodes provides a snapshot of the malware’s behaviour, while circumventing obfuscation strategies.

Santos et al [5] inspected n-grams (i.e. combinations of observed opcode sequences) in a static analysis. Accuracy and F-measure were reported as >85% in detecting malware from benign. Runwal, Low and Stamp [6] found graph approaches to detecting metamorphic malware using opcodes to be successful not only between metamorphic malware and benignware, but metamorphic malware and other malware types.

O’Kane et al [7] found that 99.5% of variance in their data was attributed to the top 8 opcodes, providing feature reduction from the original 150 opcodes.

Carlin et al. [8] used run-time opcodes to classify 48,000 malware executables. The effects of run-length (i.e. trace length) and n-grams were also analysed. The authors found that traces of 32k opcode length using $n=1$ showed the highest levels of accuracy (99.05%) using a Random Forest classifier and 10-fold cross-validation. This indicates that dynamic opcode analysis, when applied to a broad range of malware, can speedily and accurately detect the malicious behaviours of unseen malware.

B. Dynamic analysis of ransomware

Dynamic analyses of ransomware have mainly concentrated on taint analysis and network monitoring. Ahmadian, Shahriari and Ghaffarian [9] attempted to detect ransomware prior to the malicious data encryption by detecting domain generation algorithm (DGA) and key exchange activities. Ransomware often employs a DGA to connect to a randomized domain, to avoid hard-coding Command and Control (C&C) domains into the malware. If the DGA generates the valid domain successfully, a connection is made to the C&C, and a key exchange takes place. Using a text analysis approach, the authors sought to check if a domain was probably being generated at random using typical rules for English. This appears to be conceptually flawed, as non-English domains, or even acronyms would potentially be flagged. A second parallel approach monitored traffic on outward network connections, allowing the user the opportunity to deny attempted connections. The authors claim 100% accuracy with 0 false negatives (FN) using their framework. However, the dataset only contained 20 samples of ransomware and the overall accuracy was actually 50% across these 20. While a FN rate of 0 may have been observed, the research did not examine benign files and so false positives could not be reported.

Monika, Zavarsky and Lindskog [10] examined ransomware on both Android and Windows. Using dynamic analysis, the system changes and network communications of each sample were observed. It was found that there was consistency in approach across strains of ransomware, including file writes, registry manipulation and network access.

Cabaj et al [11] investigated six CryptoWall samples dynamically, focussing on the network activity, which allowed the tracking of infected proxies and the infrastructure behind the infection. This work was furthered in [12] with an expanded dataset (350 instances). In this work, multiple domain names resolved to the same IP addresses were found to be responsible for multiple C&C services

Scaife et al [13] employed dynamic analysis on 492 ransomware samples from 14 families. Change-monitoring techniques were used to analyse file type changes, similarity measurement, entropy, file deletion, and file type funnelling. Ransomware was detected with a 100% TP rate, with the median file-loss due to encryption as low as 10 out of 5100 files. This model was then used with simulated user behaviour on 30 Windows applications. The only FP the model

triggered was during benign packing and encryption by the 7Zip application.

Mbol et al. [14] employed entropy analysis to detect encryption behaviour of TorrentLocker on JPEG files. With an abnormal behaviour defined as belonging to a process that opens a large number of files, and the structure of the output differing from the input structure. The authors found that the algorithm could differentiate between typical JPEGs and encrypted JPEGs using only the first 128kb with 99.95% accuracy. They postulated that, because sections of JPEGs resemble encrypted files closely, more highly-structured files (e.g. DOC or TXT) could be distinguished even more easily. However, using the abnormal behaviour definition to diagnose file-encryption by ransomware would mostly likely classify typical benign compression or encryption behaviours as malicious. Similarly, malware could bypass the first rule of that definition by starting child processes for each file encryption.

Previous dynamic approaches to ransomware analysis have largely focussed on the effects of the file on the host system, rather than directly observing the behaviour of the file. For example, monitoring the network traffic or file I/O behaviours on disc, rather than instructions the executable is issuing. As the latter technique has proven successful at distinguishing malware from benignware in large datasets and with high accuracy, it would seem logical that ransomware should also yield to such analyses. The research presented here investigates the application of these approaches to ransomware, using a novel dataset with respect to the following research questions:

- 1) Can opcode counts extracted from runtraces offer accurate ransomware detection over a large sample size?
- 2) Does a 32k opcode run-length offer superior accuracy over full-length traces?
- 3) Can the method(s) which successfully detect(s) ransomware behaviour, differentiate these from benign encryption (e.g. file zipping) behaviours?

III. METHODOLOGY

The methodology employed for the work presented here uses the framework created in [8], with slight modifications for use in the context of ransomware. For brevity, this process is briefly restated here.

To acquire samples, the repository at VirusShare.com [15] was selected for its size, modernity and facilities, with 40 million malware samples at the time of writing. For the present paper, the cryptoransom collection was used, containing approximately 36,000 files. When PE files were extracted, 21,378 remained for the experiment. Benign files were harvested from Windows machines, and the SMOTE algorithm [16] was implemented to oversample the minority class, resulting in 3,591 benign files. The benign zipping class was created using 7Zip to pack and encrypt a collection of files. This was oversampled to create a dataset of 1000 benign zipping samples. A clean baseline snapshot of a Windows 10 VirtualBox VM was used to execute each file for nine minutes. This time was used to provide a pragmatic balance between number of samples executed and execution time. The open source debugger OllyDbg v2 was used to trace the dynamic

TABLE I
MACHINE LEARNING METRICS FOR FULL-LENGTH BENIGN AND
RANSOMWARE TRACES

TP	FP	Precision	Recall	F	MCC	ROC	PRC	Class
0.972	0.002	0.989	0.972	0.98	0.977	0.998	0.995	Benign
0.998	0.028	0.995	0.998	0.997	0.977	0.998	1	Ransomware
0.994	0.025	0.994	0.994	0.994	0.977	0.998	0.999	Average

TABLE II
MACHINE LEARNING METRICS FOR 32K RUN-LENGTH BENIGN AND,
RANSOMWARE TRACES

TP	FP	Precision	Recall	F	MCC	ROC	PRC	Class
0.932	0.001	0.982	0.932	0.957	0.955	0.996	0.982	Benign
0.999	0.068	0.996	0.999	0.998	0.955	0.996	1	Ransomware
0.996	0.064	0.996	0.996	0.996	0.955	0.996	0.999	Average

opcodes of each file during runtime, with StrongOD v0.4.8.892 used to mask the presence of the debugger, as per [7]. The host Windows OS was crafted to include the items found in a typical system, including a full recycling bin, browsing history, recent documents, Flash, .Net, Java etc. Further anti-virtualization mitigations were not utilized at the execution stage, as the presence of VBoxGuestAdditions is an easy indicator of the system being under a virtualised environment, but is necessary for automation. One main aim of dynamic investigations of malware is to experience the binary as a user would normally experience it, including virtualized platforms. Lastly, the anti-virtualization attacks by malware may also provide further useful features for detection when operating as low level as opcode instructions. This dataset was then input to a Random Forest classifier, as implemented in WEKA 3.9.

As this dataset is the largest dynamic opcode analysis of ransomware, to the best of our knowledge, we wish to release this dataset to the wider research community in order to further the body of research into mitigating dangerous malware threats.

IV. RESULTS

A. Benignware vs ransomware

To confirm past work that showed a 32k length run-trace provided optimal accuracy with lower computational overhead, both 32k run-traces (32k) and full-length run-traces (FL) were compared. Table I shows the metrics for the full-length traces, and Table II lists the metrics for the 32k set.

Overall the results show that the classifier can distinguish between each class with high levels of accuracy. The F_1 measure was 0.994 (FL) and 0.996 (32k) averaged across the classes, indicating balance between both precision and recall. The shorter run-length did, however, introduce higher levels of FP ratings for ransomware, when compared to the full-length traces. Area-under Receiver Operating Characteristic curve (ROC) indicated excellent classification performance at 99.8% and 99.6% respectively.

B. Benignware vs ransomware vs zipping

To investigate whether the RF model can distinguish ransomware from benign encryption (zipping), a three-class problem was put to the classifier. Again, the 32k run-length was

TABLE III
MACHINE LEARNING METRICS FOR FULL-LENGTH BENIGN,
RANSOMWARE, AND BENIGN ENCRYPTION TRACES

TP	FP	Precision	Recall	F	MCC	ROC	PRC	Class
0.973	0.002	0.989	0.973	0.981	0.978	0.998	0.995	Benign
0.998	0.021	0.995	0.998	0.997	0.982	0.998	1	Ransomware
1	0	1	1	1	1	1	1	Zipper
0.995	0.018	0.995	0.995	0.995	0.982	0.998	0.999	Average

TABLE IV
MACHINE LEARNING METRICS FOR 32K RUN-LENGTH BENIGN,
RANSOMWARE, AND BENIGN ENCRYPTION TRACES

TP	FP	Precision	Recall	F	MCC	ROC	PRC	Class
0.932	0.001	0.98	0.932	0.956	0.954	0.996	0.981	Benign
1	0	1	1	1	1	1	1	Zipper
0.999	0.035	0.996	0.999	0.998	0.975	0.998	1	Ransomware
0.996	0.032	0.996	0.996	0.996	0.975	0.998	0.999	Average

compared to the full-length traces. As with the two class problem above, the classification metrics showed excellent performance in distinguishing the classes, with a TP rate of 99.6%. The highlight of these results is the 100% rating for each score within the Zipper class, i.e. 100% of the time, the classifier could distinguish not only benignware from ransomware, but zipping from both other classes.

C. Accuracies

The accuracies (stated as correctly classified instances out of all classifications) for each run-length (full vs 32k) and all classes are depicted in fig 1. Investigating at the 32k run-length shows higher accuracy levels than using the full-length traces, for both the two-class and three-class analyses. As the zipping class is detected with 100% accuracy, this has improved the overall accuracy in the three-class analysis, with the other classes showing almost identical results (Benign +0.1%) between the two-class and three-class problems. Overall, the 32k run-length benign vs ransomware vs zipping analysis showed the highest level of accuracy.

Fig. 1. Accuracies (correctly classified instances) for all classes and run-lengths

V. CONCLUSIONS

This paper makes several novel contributions to the current body of knowledge. First, we present a dynamic opcode analysis of ransomware on the largest dataset of its kind in the literature, to the best of our knowledge. We wish to release this dataset to the wider research community to help propagate future work in the field.

Second, we confirm findings of our previous work that 32k run-length is most accurate. This is important, particularly in the context of ransomware, as speedy detection and mitigation of the threat is key in preventing data loss. Further, this serves as confirmation that our methodological decision to limit run-time to nine minutes is more than adequate to capture the vast majority of information required to enable a classifier to correctly predict a class.

Third, the three-class problem showed the most accurate results, due to the 100% accuracy of the zipping class. It is widely accepted in machine learning that increasing numbers of classes typically increases misclassification. Despite this, the comparison of ransomware with benign and benign encryption showed the highest accuracy. A key advantage of this model over past work is the ability of our model to detect ransomware, benignware and benign encryption behaviours, with 100% accuracy in the latter case. Future work could refactor this as a two-class problem, incorporating the benign and zipping classes as one. Similarly, the ransomware class could be investigated with respect to other non-cryptographic malware.

This research confirms our past work into the advantages of dynamic opcode analysis in the context of malware detection. Within the realm of ransomware, fast detection is essential, and our model demonstrates that short run-times are capable of offering high detection rates, while mitigating obfuscation techniques.

REFERENCES

- [1] "McAfee Labs Threats Report December 2017," McAfee, Tech. Rep., 2017. [Online]. Available: <https://www.mcafee.com/us/resources/reports/rp-quarterly-threats-dec-2017.pdf>
- [2] P. O'Kane, S. Sezer, and D. Carlin, "The evolution of ransomware," *IET Networks*, 2018, In press.
- [3] A. Gazet, "Comparative analysis of various ransomware virii," *Journal in Computer Virology*, vol. 6, no. 1, pp. 77–90, Feb 2010. [Online]. Available: <http://link.springer.com/10.1007/s11416-008-0092-2>
- [4] "An ISTR Special Report: Ransomware and Businesses 2016," Symantec, Tech. Rep., 2016. [Online]. Available: www.symantec.com/content/en/us/.../ISTR2016_Ransomware_and_Businesses.pdf
- [5] I. Santos, F. Brezo, B. Sanz, C. Laorden, and P. Bringas, "Using opcode sequences in single-class learning to detect unknown malware," *IET Information Security*, vol. 5, no. 4, p. 220, 2011. [Online]. Available: <http://digital-library.theiet.org/content/journals/10.1049/iet-ifs.2010.0180>
- [6] N. Runwal, R. M. Low, and M. Stamp, "Opcode graph similarity and metamorphic detection," *Journal in Computer Virology*, vol. 8, no. 1-2, pp. 37–52, May 2012. [Online]. Available: <http://link.springer.com/10.1007/s11416-012-0160-5>
- [7] P. O'Kane, S. Sezer, K. McLaughlin, and E. G. Im, "SVM training phase reduction using dataset feature filtering for malware detection," *IEEE transactions on information forensics and security*, vol. 8, no. 3-4, pp. 500–509, 2013.
- [8] D. Carlin, P. O'Kane, and S. Sezer, *Dynamic Analysis of Malware Using Run-Time Opcodes*, ser. Data Analytics. Springer International Publishing, 2017, ch. chapter 4, pp. 99–125. [Online]. Available: http://link.springer.com/10.1007/978-3-319-59439-2_4
- [9] M. M. Ahmadian, H. R. Shahriari, and S. M. Ghaffarian, "Connection-monitor connection-breaker: A novel approach for prevention and detection of high survivable ransomwares," in *Information Security and Cryptology (ISCISC), 2015 12th International Iranian Society of Cryptology Conference on*. IEEE, Sep 2015, pp. 79–84. [Online]. Available: <http://ieeexplore.ieee.org/document/7387902/>
- [10] Monika, P. Zavorsky, and D. Lindskog, "Experimental analysis of ransomware on windows and android platforms: Evolution and characterization," *Procedia Computer Science*, vol. 94, pp. 465–472, 2016. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S1877050916318221>
- [11] K. Cabaj, P. Gawkowski, K. Grochowski, and D. Osojca, "Network activity analysis of CryptoWall ransomware," *Przegląd Elektrotechniczny*, vol. 91, no. 11, pp. 201–204, 2015.
- [12] K. Cabaj, P. Gawkowski, K. Grochowski, and A. Kosik, "Developing malware evaluation infrastructure," ser. *Annals of Computer Science and Information Systems*, vol. 8. IEEE, Oct 2016, pp. 981–989. [Online]. Available: <https://fedesis.org/proceedings/2016/dr/490.html>
- [13] N. Scaife, H. Carter, P. Traynor, and K. R. B. Butler, "Cryptolock (and drop it): Stopping ransomware attacks on user data." IEEE, Jun 2016, pp. 303–312. [Online]. Available: <http://ieeexplore.ieee.org/document/7536529/>
- [14] F. Mbol, J.-M. Robert, and A. Sadighian, *An Efficient Approach to Detect TorrentLocker Ransomware in Computer Systems*, ser. *Lecture Notes in Computer Science*. Springer International Publishing, 2016, vol. 10052, ch. chapter 32, pp. 532–541. [Online]. Available: http://link.springer.com/10.1007/978-3-319-48965-0_32
- [15] J.-M. Roberts. (2014) Virusshare.com. [Online]. Available: <http://www.virusshare.com>
- [16] N. V. Chawla, N. Japkowicz, and A. Kotcz, "Editorial, special issue on learning from imbalanced data sets," *ACM SIGKDD Explorations Newsletter*, vol. 6, no. 1, p. 1, Jun 2004. [Online]. Available: <http://portal.acm.org/citation.cfm?doid=1007730.1007733>