Fortification Through Topological Dominance: Using Hop Distance and Randomized Topology Strategies to Enhance Network Security

Paul Hyden and **Ira S. Moskowitz**
Information Management and Decision Architectures Branch, Code 5580
Naval Research Laboratory, Washington, DC 20375

Stephen Russell
Battlefield Information Processing Branch
US Army Research Laboratory, Adelphi, MD 20783-1197

Abstract

The potential to utilize information that identifies threatening and threatened nodes in a network in order to separate them with greater hop distances by using ring networks was demonstrated by the authors in a previous work. Here these ideas are expanded by generalizing the structures and tactics used to separate threatening and threatened nodes to make them more broadly applicable by using generalized network rings. In addition, the notion of randomization of the network topology is introduced as a dynamic means to increase the technical debt to the attacker, offer real time risk management of the network, and further expand the application of separation techniques in real world networks.

Introduction

As a means of protecting network nodes, we consider tactics to structure a network and manage its underlying topology utilizing capabilities of modern network technologies. These technologies produce a network landscape that provide management efficiencies but generate more challenges to security (Baldini et al. 2012), ith topology being a key lever to improving network security and managing risk (Karsai et al. 2011; Reggiani 2013). Specifically, we build upon and expand on the fundamental observations about network topology and the intersection with the security of the network that were discussed in (Hyden, Moskowitz, and Russell 2016). Here we review many of the key assumptions from that previous work. First, we observe that tools, techniques, and procedures are increasingly available to associate a scalar value of risk with nodes on the network. Increasingly, more and more tools and algorithms exist to summarize and characterize the content of both the people associated with the nodes of the graph and well as the communication flowing on the edges of the graph. Two examples in the mobile networking space are given in (Truong et al. 2014) and (Wei et al. 2014). In (Obert et al. 2014), techniques to determine the probability of attack between node pairs using machine learning are discussed. Second, we note that increasing the hop distance between two nodes can be expected to reduce the risk associated from that node pair. Numerous tools and approaches exist to blunt or eliminate an attack if there is a greater hop distance introduced between an attacking node and a defending node, such as the use of an in-cloud scrubber service, as discussed in(Naresh Kumar et al. 2012) or middleboxes such as "virtual private network gateways, wide area network optimizers, intrusion detection and prevention systems, firewalls, and proxies" as discussed in (Qazi et al. 2013). While conceptually and in some cases practically effective, there are significant challenges to this approach. Those challenges include the cost of the resources necessary to support volume packet inspection and the mere fact that the DoS attack is just being redirected to another node, albeit one that is better prepared to handle the load (Singh, Manickam, and Rehman 2014). While techniques to address the scalability challenge have been proposed (Lua, Wah, and Ng 2014; Geva, Herzberg, and Gev 2014), few solutions seek to exploit topological characteristics. Further, volume-based attacks are only one type of threat. Chen et. al (Cheng et al. 2012) point out several varieties of network threats, suggesting that ensemble detection solutions is necessary, adding further justification for using every advantage.

However, in this paper, we expand upon some of the tools used to apply these observations to making the network more secure beyond prime ring networks. Specifically, we introduce **generalized network rings** as a means to separate threatened nodes from threatening nodes. We also expand our list of key observations about network topology by noting the power that **randomization** can play to generate greater topological complexity that can further obscure the ability of an adversary to attack a network.

Generalized Network Rings

In (Hyden, Moskowitz, and Russell 2016), the notion of a 1-hop graph was introduced, which describes which nodes are accessible to each other in exactly 1-hop when the nodes are connected in a 1-dimensional ring. Table 1,column 3 enumerates all twelve valid possibilities for this graph when $n = 5$. In the table, the node at the left is node 1 and the other nodes are numbered in counterclockwise order from 2 through 5. The name of the graph describes the order of the connections in the graph, so that graph $\{4, 2, 1, 3, 5\}$ de-

Copyright © 2016, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Table 1: 1-hop Graphs for all Circular Rings on 5 Nodes, from (Hyden, Moskowitz, and Russell 2016)

Topology Index i	Name	1-hop graph	2-hop graph	$h_2(i)$
1	{4,2,1,3,5}			12
2	{5,2,1,3,4}			11
3	{3,2,1,4,5}			10
4	{5,2,1,4,3}			9
5	{3,2,1,5,4}			8
6	{4,2,1,5,3}			7
7	{2,3,1,4,5}	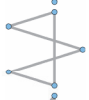		6
8	{5,3,1,4,2}			5
9	{2,3,1,5,4}			4
10	{4,3,1,5,2}			3
11	{2,4,1,5,3}			2
12	{3,4,1,5,2}			1

scribes the graph where node 4 connects to node 2, node 2 connects to node 1, node 1 connects to node 3, node 3 connects to node 5, and node 5 connects to node 4. The value $h_2(i)$ describes the isomorphism between the 1-hop and 2-hop graphs, namely the 2-hop graph corresponding to the 1-hop graph with topology index i is given by $h_2(i)$. For further discussion, please see (Hyden, Moskowitz, and Russell 2016).

Here we consider a generalized form of these types of networks. In this generalization, each node is a member of k different n ring networks simultaneously, rather than exactly one. For this generalized form of the ring network, each node is directly connected to $2k$ other nodes, and can include n^k unique nodes. Each node in the ring network can be interpreted as being labeled with a number in base n, where each position in the base n representation describes which ring the node belongs to in each 'dimension' of the ring.

Example of concepts with $n = 5$ and $k = 2$

In figure 1, we show an example with $n = 5$ and $k = 2$ for a total of 25 nodes. Here we form the graph with 5 layers of $n = 5$ and $k = 1$ generalized network rings. In this 2 dimensional picture, it is convenient to represent these layers with rotated concentric circles to better represent the connections between layers. Note that node 1 is part of the ring of nodes formed by nodes 1,2,3,4, and 5 as well as the ring of nodes formed by nodes 1, 6, 11, 16, and 21. If we instead labeled the nodes by starting with 0 and labeling them with their value in base 5 minus 1, the nodes would be numbered $00, 01, 02, \ldots, 40, 41, 42, 43, 44$. we would see that node 00 is in the ring of nodes formed by nodes 00, 01, 02, 03, and 04 as well as the ring of nodes formed by nodes 00, 10, 20, 30, and 40.

We can think of the number of the node as an address where each position in the address describes which ring the node is a member of in each 'dimension' of the generalized ring. It is also important to note which nodes are separated by the maximum number of hops in the graph. In this example, the maximum hop distance is $k\lfloor n/2 \rfloor = 2\lfloor 5/2 \rfloor = 4$. For node 00, there are 4 such nodes that are at a maximum hop length of 4: nodes 22, 33, 23, and 32.

Promotes hop distance as a means of security between node threat pairs

This new structure gives further control to provide different levels of hop distance spacing. We assume that security between a pair of nodes is provided by increasing the hop distance between two nodes. See (Hyden, Moskowitz, and Russell 2016) for justification. The hop distance is the minimal number of hops required to transmit a packet between a pair of nodes, as a function of the network topology. In the case of $k = 1$ (the one dimensional case), two nodes are separated by a maximum of $\lfloor n/2 \rfloor$ hops. Next, consider Figure 1 which displays the generalized network ring with $n = 5$ and $k = 2$. Now, imagine that every location in the generalized network has an address with k digits, as in the second numbering scheme of Figure 1. To move between two nodes, we have to change the address of the origin node into the address of the destination node.

Suppose we consider the origin node 01 for the $n = 5$ and $k = 2$ case. Let's consider the address of the destination node that would have the worst case in terms of hop distance. In terms of the first address value, the two worst case addresses for the destination are 2 and 4. Hence, from the view of the first address value, the worst case destination nodes are in the set $\{20, 40, 21, 41, 22, 42, 23, 43, 24, 44, 25, 45\}$. In terms of the second address value, the two worst case addresses for the destination are 3 and 0. Hence, from the view of the second address value, the worst case destination node is in the set $\{03, 00, 13, 10, 23, 20, 33, 30, 43, 40, 53, 50\}$. Taking the intersection of these two sets, we see that nodes $20, 40, 43$, and 23 are the four destination nodes that are maximally hop separated. In the worst case, we are at most $\lfloor n/2 \rfloor$ hops from matching any of the k digits. However, in

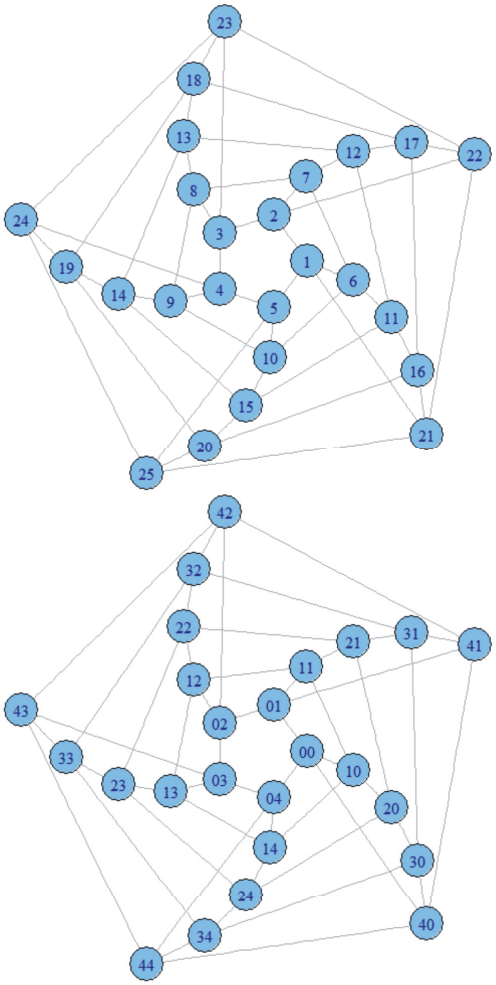

Figure 1: Generalized ring with $n = 5$ and $k = 2$ with two alternative labels

the worst case, this is true for each and every one of the k positions in the address. Hence, in the generalized network ring, two nodes can be separated by a maximum of $k\lfloor n/2 \rfloor$ hops. Figure 2 displays all node pairs that have maximal hop separation for the $n = 5$ and $k = 2$ case.

Generalized network rings for $k = 3$

Although it is more difficult to represent within the restrictions of two dimensions, Figure 3 provides several different perspectives of a 3-D representation of a generalized ring with $n = 5$ and $k = 3$ to give some intuition on the generalized ring for larger values of k. Here the 'layers' visible in the plots are the $n = 5$ and $k = 2$ generalized rings from Figure 1, where each layer is also connected in a ring to its neighboring layers to increase k to 3. These underlying 'layers' are shown with rotation and expanding size to better display some of the connections between layers. This also demonstrates the general approach for constructing the generalized network ring for $k + 1$ when given the generalized network ring for k: form n 'layers' from the generalized network ring for k and connect the nodes in the same position of each layer into another ring by connecting neighboring layers.

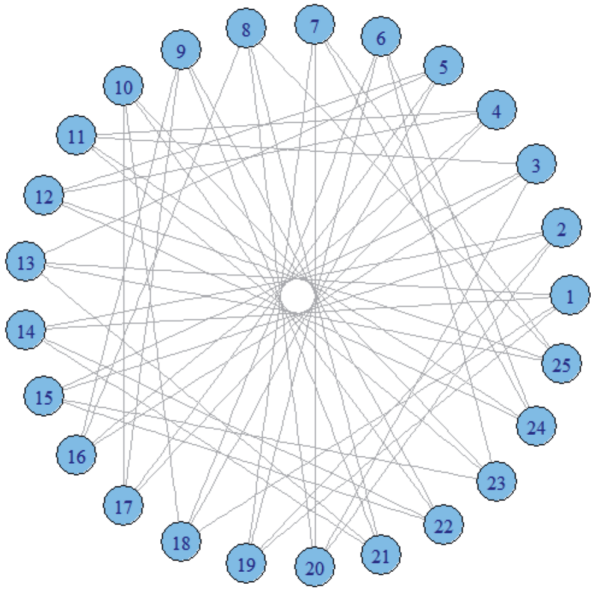

Figure 2: Egdes in this graph indicate node pairs that are at a maximal hop distance of 4 for the generalized ring in figure 1

Generalized network rings allow for more pairs of nodes to be maximally hop-separated

The *other* impact of generalizing the ring network structure is that more pairs of nodes can be simultaneously maximally separated. Suppose we can assign a scalar to the threat imposed from node i to node j for all pairs of nodes $i, j \in N$. One algorithm for structuring the network is to place the highest risk pair of nodes at a maximal hop distance, and then proceed down the list until no more node pairs can be accommodated with a maximal hop distance. In the case of a 1 dimensional ring on n nodes, $\lfloor n/2 \rfloor$ pairs of nodes can be maximally separated at the same time. In the generalized network ring, $k\lfloor n/2 \rfloor$ pairs of nodes can be maximally separated. In fact, we can be more sophisticated than this and employ an algorithm to minimize the total risk of the node pairs selected for maximal distance for any particular selection of edges, where the total risk is the sum of all risks incurred by a given network topology.

The total number of valid networks under these parameters is very large. First, consider the number of ways to divide the 25 nodes into groups of size 5, noting that the labeling of each five rings are not significant. This takes the form of a multinomial coefficient.

$$\begin{aligned}
&\frac{25!}{5!20!}\frac{20!}{5!15!}\frac{15!}{5!10!}\frac{10!}{5!5!}\frac{5!}{5!0!}\frac{1}{5!} \\
&= \frac{25!}{5!5!5!5!5!} \\
&= 5194672859376
\end{aligned}$$

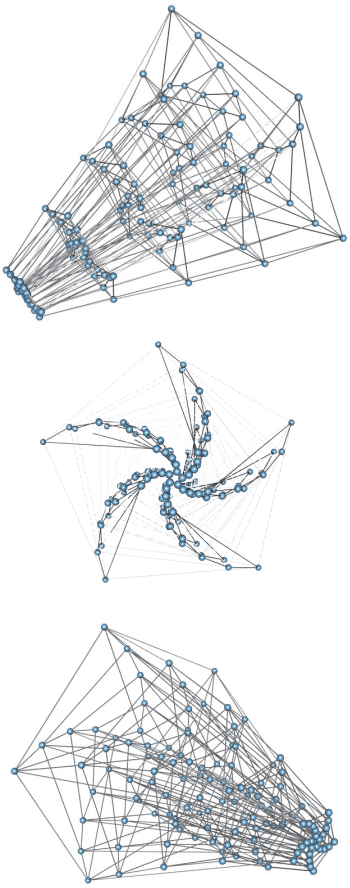

Figure 3: 3D image of a generalized ring with $n = 5$ and $k = 3$ from three different perspectives

Once one set of rings are determined, one member of each rings is selected to form each of the other five rings, again noting that the label of each of these rings is not significant.

$$5^5 4^5 3^5 2^5 1^5 \frac{1}{5!} = (5!)^4$$

Hence, the total number of combinations is

$$\frac{25!}{5!5!5!5!5!}(5!)^4 = \frac{25!}{5!5!} \approx 1.07716736412021 \cdot 10^{21}$$

Now, once we have determined the membership of each node to exactly two rings, the ring can be wired in 12 different ways, as shown in table 1. This is true for each of the 10 rings, which increases the number of combinations by a factor of 12^{10}

Hence, the total number of combinations is

$$\frac{12^{10}25!}{5!5!} \approx 6.66953640144369 \cdot 10^{31}$$

Clearly, with this many combinations, it will be necessary to employ heuristics to estimate the best configuration from a security perspective.

Performance versus Security

It is important to strike a balance between performance and security when introducing these concepts to network topology. We want node pairs that convey a large volume of traffic to be close together, while we want to separate node pairs where one node poses a high threat to another node. Heuristically, the ring structure allows us to co-locate nodes in the same ring when they share a lot of traffic, yet still introduce separation between node pairs that have a higher level of risk. As nodes are members of k different rings, we can conceptually co-locate nodes across k different dimensions, providing extra degrees of freedom.

In the case that we stray from this regular structure and introduce direct links between node pairs, the paper (Moskowitz, Hyden, and Russell 2016) offers a way to characterize the risk of introducing this link.

Techniques for fitting general networks

In general, the number of network nodes will not precisely align with a suitable value of n^k where n is prime and k is an positive integer. However, several strategies exist to bring these values into balance.

Dummy nodes The simplest approach is to simply define a larger network than needed for the number of nodes in the network, and introduce dummy nodes in the network.

Different values of n In general, the value of n comprising the rings that form the general network do not have to be equal. Hence, for a value of k, we can choose $n_1, n_2, \ldots, n_k$ such that the values of n are distinct so that the total number of nodes in the network is $\prod_{i=1}^{k} n_i$

Subnetworks Another strategy is to simply compose the network into distinct subnetworks, and either keep them disconnected, or connected across specific bridges.

Randomized Periodic Network Connections In the presence of **randomized periodic connections**, we can form subnetworks that, over longer time periods, offer connectivity across the entire network. These are connections that exist with less than probabiliy 1 at any single time step, but exist in the limit with probability 1 over several time steps.

Techniques for Topological Randomization

Another strategy that makes it difficult for an adversary to position for an attack is the introduction of **randomized** network topology. Here, rather than treating the topology as fixed, we introduce randomization to the assignment of a topology. This presents a network topology to a potential attacker that is stochastic. This implies that the attacker must expend more resources to position for an attack from many more nodes than in the deterministic case, increasing the technical debt incurred to mount an attack. Here we are also using a fundamental concept from game theory of a mixed strategy.

To build intuition for the value with regard to security, let us consider an extreme example where the topology of the network is reassigned at every time step. Now, unless the source of the traffic is exactly one hop away from the

destination of the traffic, the likelihood that a node that will be exactly one hop away from its destination in the next time step is the same: $\frac{2k}{n^k-1}$.

Here is the reasoning. Given that the packet is not already in the correct destination, and all the permutations have equal likelihood, then the correct destination is uniform over the other $n^k - 1$ spots in the ring. Since $2k$ of those spots are connected to the current node, the probability is $\frac{2k}{n^k-1}$ that the data packet will be delivered to the correct destination in the next time step. This results in a geometric distribution on the number of hops until reaching the final destination.

Reducing the frequency to changing every k time steps means that traffic will be reachable for hop lengths of k or less in a deterministic way. This effectively makes it much more difficult to successfully transmit traffic on hop lengths greater than k

By using a mix of controls

1. The size and topology of the ring
2. The inclusion or exclusion of nodes in different rings
3. The frequency of randomization

we can arbitrarily adjust the probability that any node threatens another node. As we gain information about attack-target node pairs, we can adjust this probability to match the given risk and the appetite for accepting such risk. Of course, there are many strategies we can employ to randomize the network structure to achieve different effects. The following section describes a sample strategy allowed by combining generalized network topologies with randomization.

Sample Strategy: Indefinitely delaying traffic between a pair of nodes

Suppose we have introduced this randomized network topology scheme to the management of our network. Further, suppose we have been made aware of a dangerous packet that originated from a threat node that is aimed at a specific target node that is currently more than 2 hops away. Within the scheme of the random topology selection, the network managers may limit the topologies selected in the next time periods to topologies where the hop distance between the current location of the dangerous packet and the target node is always at least 2. In this way, we can indefinitely delay threatening traffic from reaching the target node until the nature of the dangerous packet can be determined, the packet can be directed to a scrubber, or the target node can be hardened against the threat. With judicious use of this capability, it is even possible that the attacker will not know that network security diagnosed and countered the threat, so that the attacker will be less likely to deploy various means of deception.

Risk estimation as a function of network topology variation

As a byproduct of varying the network topology, we provide a means to actively assess node pair threats as the topology changes. Some threatening nodes may simply target nodes opportunistically as a function of simply being close by. Other nodes may have specific targets based on the features of the node. In the course of our randomized periodic topology, we can collect this data and manage our future topology selections to adjust for these changing risk assessments.

Improving performance for privileged traffic

Suppose a network packet carried with it information which allowed the node to know the direction that would best move the packet toward its intended destination for the particular randomized network topology. This could be done by providing encrypted information about the future states of the network to trusted nodes. In that case, we could improve the likelihood that longer hop traffic reaches its destination in fewer hops. However, traffic without the foreknowledge of the future network state would be exposed with statistically observable longer path times.

Conclusion

Our strategy is to manage the security and risk of the network through four key observations. First, information about the network, via a combination of directly at the traffic and by leveraging metadata about the equipment, users, and risks that compose the network, make it possible to roughly estimate threatening and threatened nodes. Second, separating threatened nodes from threatening nodes with a greater hop distance reduces the threat to the target node. Third, the separation of threatened and threatening nodes can be efficiently managed by utilizing generalized ring structures. Finally, introducing randomization to the network topology increases the technical debt to the attacker and further expands the applicability of separation strategies to real world networks and offers the promise of real time risk management.

References

Baldini, G.; Sturman, T.; Biswas, A.; Leschhorn, R.; Godor, G.; and Street, M. 2012. Security aspects in software defined radio and cognitive radio networks: A survey and a way ahead. *Communications Surveys Tutorials, IEEE* 14(2):355–379.

Cheng, T.-H.; Lin, Y.-D.; Lai, Y.-C.; and Lin, P.-C. 2012. Evasion techniques: Sneaking through your intrusion detection/prevention systems. *Communications Surveys Tutorials, IEEE* 14(4):1011–1020.

Geva, M.; Herzberg, A.; and Gev, Y. 2014. Bandwidth distributed denial of service: Attacks and defenses. *Security Privacy, IEEE* 12(1):54–61.

Hyden, P.; Moskowitz, I. S.; and Russell, S. 2016. Using network topology to supplement high assurance systems. In *IEEE 17th International Symposium on High Assurance Systems Engineering (HASE), 2016*. Orlando, Florida(USA): IEEE.

Karsai, M.; Kivelä, M.; Pan, R. K.; Kaski, K.; Kertész, J.; Barabási, A.-L.; and Saramäki, J. 2011. Small but slow

world: How network topology and burstiness slow down spreading. *Phys. Rev. E* 83:025102.

Lua, R.-P.; Wah, C. H.; and Ng, W. K. 2014. Cornstarch effect: intensifying flow resistance for increasing ddos attacks in autonomous overlays. In *Consumer Communications and Networking Conference (CCNC), 2014 IEEE 11th*, 537–538.

Moskowitz, I. S.; Hyden, P.; and Russell, S. 2016. Network topology and mean infection times. PREPRINT, under review.

Naresh Kumar, M.; Sujatha, P.; Kalva, V.; Nagori, R.; Katukojwala, A.; and Kumar, M. 2012. Mitigating economic denial of sustainability (edos) in cloud computing using in-cloud scrubber service. In *2012 Fourth International Conference on Computational Intelligence and Communication Networks (CICN)*, 535–539.

Obert, J.; Pivkina, I.; Huang, H.; and Cao, H. 2014. Dynamically differentiated multipath security in fixed bandwidth networks. In *Military Communications Conference (MILCOM), 2014 IEEE*, 88–93. IEEE.

Qazi, Z. A.; Tu, C.-C.; Chiang, L.; Miao, R.; Sekar, V.; and Yu, M. 2013. Simple-fying middlebox policy enforcement using sdn. In *Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM*, SIGCOMM '13, 27–38. New York, NY, USA: ACM.

Reggiani, A. 2013. Network resilience for transport security: Some methodological considerations. *Transport Policy* 28(0):63 – 68. Special Issue on Transportation Pricing Policies Special Issue on Transport Security - Policies and Empirical Perspectives.

Singh, P.; Manickam, S.; and Rehman, S. 2014. A survey of mitigation techniques against economic denial of sustainability (edos) attack on cloud computing architecture. In *Reliability, Infocom Technologies and Optimization (ICRITO) (Trends and Future Directions), 2014 3rd International Conference on*, 1–4.

Truong, H. T. T.; Lagerspetz, E.; Nurmi, P.; Oliner, A. J.; Tarkoma, S.; Asokan, N.; and Bhattacharya, S. 2014. The company you keep: Mobile malware infection rates and inexpensive risk indicators. In *Proceedings of the 23rd International Conference on the World Wide Web*, 39–50. International World Wide Web Conferences Steering Committee.

Wei, Z.; Tang, H.; Yu, F.; and Mason, P. 2014. Trust establishment based on bayesian networks for threat mitigation in mobile ad hoc networks. In *2014 IEEE Military Communications Conference (MILCOM)*, 171–177.