Secure Overlay Network Design

Li (Erran) Li[1], Mohammad Mahdian[2], and Vahab S. Mirrokni[3]

[1] Bell Laboratories
erranlli@dnrc.bell-labs.com
[2] Microsoft Research
mahdian@microsoft.com
[3] MIT Computer Science and Artificial Intelligence Lab
mirrokni@theory.csail.mit.edu

Abstract. Due to the increasing security threats in the Internet, new overlay network architectures have been proposed to secure privileged services. In these architectures, the application servers are protected by a defense perimeter where only traffic from entities called servelets are allowed to pass. End users must be authorized and can only communicate with entities called access points (APs). APs relay authorized users' requests to servelets, which in turn pass them to the servers. The identity of APs are publicly known while the servelets are typically secret. All communications are done through the public Internet. Thus all the entities involved forms an overlay network. The main component of this distributed system consists of n APs. and m servelets. A design for a network is a bipartite graph with APs on one side, and the servelets on the other side. If an AP is compromised by an attacker, all the servelets that are connected to it are subject to attack. An AP is *blocked*, if all servelets connected to it are subject to attack. We consider two models for the failures: In the *average case model*, we assume that each AP i fails with a given probability p_i. In the *worst case model*, we assume that there is an adversary that knowing the topology of the network, chooses at most k APs to compromise. In both models, our objective is to design the connections between APs and servelets to minimize the (expected/worst-case) number of blocked APs. In this paper, we give a polynomial-time algorithm for this problem in the average-case model when the number of servelets is a constant. We also show that if the probability of failure of each AP is at least 1/2, then in the optimal design each AP is connected to only one servelet (we call such designs *star-shaped*), and give a polynomial-time algorithm to find the best star-shaped design. We observe that this statement is not true if the failure probabilities are small. In the worst-case model, we show that the problem is related to a problem in combinatorial set theory, and use this connection to give bounds on the maximum number of APs that a perfectly failure-resistant design with a given number of servelets can support. Our results provide the *first* rigorous theoretical foundation for practical secure overlay network design.

Keywords: network design, network security, optimization, combinatorics.

1 Introduction

Providing secure and highly available services using the shared Internet infrastructure is very challenging due to security threats in the Internet. Distributed Denial of Service

(DDoS) attacks are a major threat to Internet security. Attacks against high-profile web sites such as Yahoo, CNN, Amazon and E*Trade in early 2000 [7] rendered the services of these web sites unavailable for hours or even days. During the hour long attack against root Domain Name Servers (DNS) in Oct, 2002, only four or five of the 13 servers were able to withstand the attack and remain available to legitimate Internet traffic throughout the strike [13]. Internet service would have started degrading if the attack had been sustained long enough for the information contained in the secondary DNS caches to start expiring—a process that usually takes from a few hours to about two days. A recent attack on June 15, 2004 against Akamai's DNS servers caused several major customers of Akamai's DNS hosting services, including Microsoft Corp., Yahoo Inc., and Google Inc. to suffer brief but severe slowdown [22] in their web performance. The event was marked by being a step beyond "simple bandwidth attacks" on individual web sites to more sophisticated targeting of core upstream Internet routers, DNS servers and bandwidth bottlenecks.

To defend against DDoS attacks, one can trace the attack sources and punish the perpetrators [3, 5, 19, 21, 4, 20, 8, 1, 11]. Due to the large number of compromised hosts (known as Zombies) used in the attack, finding the attack origin can be very difficult. Techniques to prevent DDoS attacks and/or to mitigate the effect of such attacks while they are raging on have been proposed [12, 6, 17, 9, 14, 16, 15]. These mechanisms alone do not prevent DDoS attacks from disrupting Internet services as they are reactive in nature. Recent research efforts [9, 2] have focused on designing overlay network architectures where certain critical elements are hidden from the attackers. The key entities in these architectures are access points (APs), servelets and end application servers. The end application servers are protected by a defense perimeter. Routers at the boundary are installed with filters which only allow traffic from the servelets in. The servelets are hidden from the attackers. Only a subset of access points are allowed to access each servelet. User requests must be authorized by access points and the requests are tunneled to their corresponding servelets via access points. The servelets then communicate with the end application servers. The access points can be geographically well placed to service the end users. The number of access points is assumed to be much larger than the number of secret servelets. All communications go through the public Internet. Thus all the entities involved form an overlay network.

The ability of such distributed systems to service their users is characterized by how many access points can still communicate to the end application servers, should an attack happens. This depends on how the access points are connected to the servelets. Intuitively, if a vulnerable access point connects to all the servelets, once it is compromised, all the servelets will be subject to DDoS attacks. In the worse case, this in turn denies all other access points from accessing the servelets. The network must be designed to resist such attacks. However, how the network should be designed has not been rigorously analyzed. In this paper, we formalize the problem as a combinatorial optimization problem with the objective to maximize the number of surviving access points. We first define our problem settings.

Definition 1. *A design for a network with* n APs *and* m servelets *is a bipartite graph with APs on one side, and the servelets on the other side. If an AP* fails *(or is* compromised*), it attacks all the servelets that are connected to it and we say that these*

servelets are attacked. *If all servelets connected to an AP are attacked, we say that the AP is* blocked. *By definition, we say any compromised AP is blocked.*

We are interested in designing secure networks in which the number of blocked APs is minimized. We consider two models of failures:

- In the *average case model*, we assume that each AP i fails with a given probability p_i. Our objective is to design the connections between APs and servelets to minimize the expected number of blocked APs[4].
- In the *worst case model*, we assume that there is an adversary that knowing the topology of the network, chooses at most a given number k of APs to compromise. Our objective is to design the connections between APs and servelets to minimize the worse-case number of blocked APs.

This paper presents the *first* theoretical study of secure overlay network design. Our results provide guidelines for practical design of such networks.

The rest of this paper is organized as follows. In Section 2, we study the problem in the average case model. We first prove a lemma on the structure of the optimal design. This lemma restricts the number of possible solutions and gives a polynomial-time algorithm for the problem where the number of servelets is constant. It also implies a polynomial-time algorithm for the case that each AP can be connected to at most one servelet. We prove that if all failure probabilities are large enough (namely, greater than $\frac{1}{2}$), then the optimal design is of this form, and therefore can be found in polynomial time. At the end of Section 2, we give an example that if failure probabilities are not small, then the optimal design is not necessarily star shaped, and in fact, the best star-shaped design can be worse than the optimal design by an arbitrary factor. Finally, in Appendix A, we show hardness results for computing the expected number of blocked APs for a given network. In Section 3, we study the worst case model. We establish a connection between the secure network design problem and a problem in combinatorial set theory, and use this to give the optimal design for one failed AP. For constant number of failed APs, we use the probabilistic method to bound the maximum number of APs that we can support using a fixed number of servelets without blocking any other APs. We conclude in Section 4 with several open questions.

2 The Average-Case Model

In this section, we study the average case model. We give polynomial-time algorithms for this problem in two cases: when the number of servelets is a constant, and when the probability of failure of each AP is at least 1/2. We also demonstrate the difficulty of the problem in Appendix A by showing that even when a design is given, computing the probability that a given AP will be blocked or the expected number of APs that will be blocked is $\#P$-complete.

Our algorithms are based on the following lemma about the structure of the optimal design.

[4] For a detailed justification of this model, please see [2].

Lemma 1. *Assume that APs are ordered in decreasing order of their failure probabilities, i.e., $p_1 \geq p_2 \geq \ldots \geq p_n$. For an AP i, let S_i be the set of servelets connected to i. There exists an optimal design in which for all $i < j < k$, if $S_i = S_k$, then $S_j = S_i$.*

Proof. Assume that there is no optimal design with the desired property. Let $S_1, \ldots, S_n$ be an optimal solution in which for some $i < j < k$, $S_i = S_k$ but $S_j \neq S_i$. Note that $S_i = S_k$ implies that if either i or k fails, then both i and k are blocked. In particular, the expected number of blocked APs given that i fails is equal to the expected number of blocked APs given that k fails and is equal to the expected number of blocked APs given that i and k fail. Let B_{11} be the expected number of blocked APs given that j fails and at least one of i and k fail. Let B_{10} be the expected number of blocked APs given that at least one of i and k fail and j does not fail. Similarly, let B_{01} be the expected number of blocked APs given that j fails but neither i nor k fails and B_{00} be the expected number of blocked APs given that none of i and k and j fails. From this definitions, it is straightforward to see that $B_{11} \geq B_{01}$. The expected number $\mathcal{P}^*$ of blocked APs in an optimal design can be expressed as follows.

$$\begin{aligned}\mathcal{P}^* &= \mathbf{E}[\#\text{blocked APs}] \\ &= p_j(p_i + p_k - p_i p_k)B_{11} + (1-p_j)(p_i + p_k - p_i p_k)B_{10} \\ &\quad + (1-p_i)p_j(1-p_k)B_{01} + (1-p_i)(1-p_j)(1-p_k)B_{00}\end{aligned}$$

Now we prove that the set of servelets of j can be exchanged with the set of servelets of either i or k without increasing the expected number of blocked APs. For contradiction, assume that both these exchanges increase the expected number of blocked APs. The expected number of blocked APs after exchanging i and j can be written as

$$\begin{aligned}\mathcal{P}_1 &= \mathbf{E}[\#\text{blocked APs}] \\ &= p_i(p_j + p_k - p_j p_k)B_{11} + (1-p_i)(p_j + p_k - p_j p_k)B_{10} \\ &\quad + (1-p_j)p_i(1-p_k)B_{01} + (1-p_j)(1-p_i)(1-p_k)B_{00}\end{aligned}$$

Similarly, the expected number of blocked APs after exchanging j and k is

$$\begin{aligned}\mathcal{P}_2 &= \mathbf{E}[\#\text{blocked APs}] \\ &= p_k(p_i + p_j - p_i p_j)B_{11} + (1-p_k)(p_i + p_j - p_i p_j)B_{10} \\ &\quad + (1-p_i)p_k(1-p_j)B_{01} + (1-p_i)(1-p_k)(1-p_j)B_{00}\end{aligned}$$

By our assumption, we have $\mathcal{P}^* < \mathcal{P}_1$ and $\mathcal{P}^* < \mathcal{P}_2$. Therefore,

$$p_j p_k B_{11} + p_i B_{10} + p_j(1-p_k)B_{01} < p_i p_k B_{11} + p_j B_{10} + p_i(1-p_k)B_{01}$$

Thus,

$$(p_i - p_j)(p_k B_{11} - B_{10} - (1-p_k)B_{01}) > 0$$

Since $p_i \geq p_j$, this implies

$$p_k B_{11} - B_{10} - (1-p_k)B_{01} > 0 \tag{1}$$

Similarly, $\mathcal{P}^* < \mathcal{P}_2$ implies

$$p_i B_{11} - B_{10} - (1 - p_i)B_{01} < 0 \tag{2}$$

By subtracting (1) from (2), we get $(p_i - p_k)B_{11} < (p_i - p_k)B_{01}$, and hence $B_{11} < B_{01}$. However, this is impossible by the definition of B_{11} and B_{01}. □

Using Lemma 1, we can prove the following result.

Theorem 1. *There is a polynomial-time algorithm that constructs the optimal design in the average case model when the number of servelets is at most a constant.*

Proof Sketch. Assume that APs are ordered in the decreasing order of their failure probabilities, i.e., $p_1 \geq p_2 \geq \ldots \geq p_n$. Let S_i denote the set of servelets connected to the AP i. From Lemma 1, we know that there are indices $1 = \alpha_0 < \alpha_1 < \alpha_2 < \cdots < \alpha_s = n+1$ such that for each $j \in [\alpha_i, \alpha_{i+1})$, $S_j = S_{\alpha_i}$, and the sets $S_{\alpha_0}, S_{\alpha_1}, \ldots, S_{\alpha_{s-1}}$ are pairwise distinct. Since the total number of distinct sets of servelets is 2^m, there are at most $\binom{n+2^m}{2^m}(2^m)!$ ways to pick the indices $\alpha_0, \ldots, \alpha_s$ and the corresponding S_i's. This number is bounded by a polynomial in n if m is a constant. Therefore, the algorithm can check all such configurations. Computing the expected number of blocked APs for each configuration can also be done in polynomial time when m is a constant. □

If we can connect each AP to at most one servelet, the resulting graph is a union of stars. We say that the design is *star-shaped* in this case. The following theorem proves that the optimal star-shaped design can be found in polynomial time.

Theorem 2. *The optimal star-shaped design can be computed in polynomial time.*

Proof. Let the failure probabilities of the APs be $p_1 \leq p_2 \leq \ldots \leq p_n$. It is easy to see that the proof of Lemma 1 holds even if the design is restricted to a star-shaped design. This shows that in the optimal star-shaped design we should partition the APs $1, \ldots, n$ into at most $m + 1$ consecutive parts each of which is connected to no servelet or to one of the servelets. This can be done by dynamic programming in polynomial time. We observe that the subset of APs that are connected to none of the servelets should be among the APs with larger failure probability. Let $A[k, t]$ be the minimum (over the choice of the star-shaped design) of the expected number of blocked APs when the set of APs consists of $1, 2, \ldots, k$ and there exists t servelets. Let $B(a, b)$ be the expected number of blocked APs among the APs $a, a + 1 \ldots, b$, if they are all connected to the same servelet (and no other AP is connected to this servelet). Note that $B(a, b)$ can be easily computed in polynomial time for each a and b. It is not hard to see that $A[k, t] = \min\{\min_{1 \leq l \leq k}\{A[l, t-1] + B(l+1, k)\}, \min_{1 \leq l \leq k}\{A[l, t] + k - l\}\}$ and $A[k, 0] = k$. Using this recurrence, the values of $A[k, t]$ can be computed in polynomial time. The value of the best star-shaped design is given by $A[n, m]$. □

It might appear that star-shaped designs are weaker than general designs. The following theorem shows that if all failure probabilities are at least $\frac{1}{2}$, there is an optimal design that is star-shaped.

Theorem 3. *If all failure probabilities are at least $\frac{1}{2}$ then there is a star-shaped optimal design and therefore an optimal design can be found in polynomial time.*

Proof. We start from an optimal design, $\mathcal{D}$, and prove that we can change this design to a star-shaped design without increasing the expected number of blocked APs.

First we prove that we can get rid of all the cycles in the optimal design $\mathcal{D}$. If there is a cycle in $\mathcal{D}$, then there is a chordless cycle C in $\mathcal{D}$ as well. The length of cycle C is even and is at least 4. We consider two cases:

Case 1: $|C| \geq 6$**.** In this case, let cycle C be $s_1c_1s_2c_2 \dots s_kc_ks_1$, where c_i's are APs and s_i's are servelets. We claim that removing one of the matchings $c_1s_1, c_2s_2, \dots, c_ks_k$ or $c_1s_2, c_2s_3, \dots, c_{k-1}s_k, c_ks_1$ will not increase the expected number of blocked APs. Let $\mathcal{D}_1$ be the design $\mathcal{D}$ after removing the matching $c_1s_1, c_2s_2, \dots, c_ks_k$ and $\mathcal{D}_2$ be the design after removing the matching $c_1s_2, c_2s_3, \dots, c_{k-1}s_k, c_ks_1$. Removing a matching from C will not increase the blocking probability of any AP other than $c_1, c_2, \dots, c_k$. So it is enough to argue that the expected number of blocked APs in $c_1, c_2, \dots, c_k$ decreases as we remove one of these two matchings. Let E_{c_i} for all $1 \leq i \leq k$ be the event that all of servelets that are connected to c_i and are not in the set $\{s_1, s_2, \dots, s_k\}$ are attacked. Let E_{s_i} be the event that at least one of the APs that are connected to servelet s_i fails. The probability of E_{c_i} is denoted by P_{c_i}, and the probability of E_{c_i} and not E_{s_j} is denoted by $P_{c_i\bar{s}_j}$. Similarly, the probability of E_{c_i} and E_{s_j} and not E_{s_l} is denoted by $P_{c_is_j\bar{s}_l}$, etc. Let $\mathcal{P}_{\mathcal{T}}(c_i)$ be the blocking probability of c_i in design $\mathcal{T}$. Then,

$$\begin{aligned}\mathcal{P}_{\mathcal{D}}(c_i) = p_i + (1-p_i)\Big(&P_{c_i} - P_{c_i\bar{s}_i}(1-p_{i-1}) \\ &- P_{c_i\bar{s}_{i+1}}(1-p_{i+1}) + P_{c_i\bar{s}_i\bar{s}_{i+1}}(1-p_{i-1})(1-p_{i+1})\Big).\end{aligned}$$

Furthermore $\mathcal{P}_{\mathcal{D}_1}(c_i) = p_i + (1-p_i)P_{c_is_i}$ and $\mathcal{P}_{\mathcal{D}_2}(c_i) = p_i + (1-p_i)P_{c_is_{i+1}}$.

In order to prove that the expected number of blocked APs is not more in one of the designs $\mathcal{D}_1$ and $\mathcal{D}_2$, it is enough to prove that $\mathcal{P}_{\mathcal{D}}(c_i) \geq \frac{1}{2}(\mathcal{P}_{\mathcal{D}_1}(c_i) + \mathcal{P}_{\mathcal{D}_2}(c_i))$. In order to prove this, it is enough to show the following:

$$\begin{aligned}\mathcal{P} &:= P_{c_i} - P_{c_i\bar{s}_i}(1-p_{i-1}) - P_{c_i\bar{s}_{i+1}}(1-p_{i+1}) + P_{c_i\bar{s}_i\bar{s}_{i+1}}(1-p_{i-1})(1-p_{i+1}) \\ &\geq \frac{1}{2}(P_{c_is_i} + P_{c_is_{i+1}})\end{aligned}$$

Using $P_{c_i} = P_{c_is_i} + P_{c_i\bar{s}_i} = P_{c_is_{i+1}} + P_{c_i\bar{s}_{i+1}}$, we have:

$$\begin{aligned}\mathcal{P} &\geq \frac{1}{2}(P_{c_is_i} + P_{c_i\bar{s}_i} + P_{c_is_{i+1}} + P_{c_i\bar{s}_{i+1}}) - P_{c_i\bar{s}_i}(1-p_{i-1}) - P_{c_i\bar{s}_{i+1}}(1-p_{i+1}) \\ &\geq \frac{1}{2}(P_{c_is_i} + P_{c_is_{i+1}})\end{aligned}$$

where we use the fact that $p_{i-1} \geq \frac{1}{2}$ and $p_{i+1} \geq \frac{1}{2}$.

Case 2: $|C| = 4$**.** Let cycle C be $c_1s_1c_2s_2c_1$. The analysis of this case is very similar to the that of $|C| > 4$. We use the same notation as in the previous case. Again we

prove that removing one the matchings c_1s_1, c_1s_2 or c_1s_2, c_2s_1 will not increase the expected number of blocked APs. Let $\mathcal{D}$, $\mathcal{D}_1$ and $\mathcal{D}_2$ be an optimal design, and this design after removing matchings c_1s_1, c_1s_2 and c_1s_2, c_2s_1, respectively.

$$\begin{aligned}
\mathcal{P}_{\mathcal{D}}(c_i) &= p_i + (1-p_i)(P_{c_i} - (1-p_{i+1})(P_{c_i\bar{s}_1} + P_{c_i\bar{s}_2} - P_{c_i\bar{s}_1\bar{s}_2})) \\
&\geq \frac{1}{2}(p_i + (1-p_i)P_{c_is_1} + p_i + (1-p_i)P_{c_is_2}) \\
&\quad + (1-p_i)(p_{i+1} - \frac{1}{2})(P_{c_i\bar{s}_1} + P_{c_i\bar{s}_2}) \\
&\geq \frac{1}{2}(\mathcal{P}_{\mathcal{D}_1}(c_i) + \mathcal{P}_{\mathcal{D}_2}(c_i)) + (1-p_i)(p_{i+1} - \frac{1}{2})(P_{c_i\bar{s}_1} + P_{c_i\bar{s}_2}) \\
&\geq \frac{1}{2}(\mathcal{P}_{\mathcal{D}_1}(c_i) + \mathcal{P}_{\mathcal{D}_2}(c_i))
\end{aligned}$$

Thus, in at least one of the designs $\mathcal{D}_1$ and $\mathcal{D}_2$, the expected number of blocked APs is less than or equal to the expected number of blocked APs in $\mathcal{D}$.

After getting rid of all cycles, $\mathcal{D}$ is a tree. Next, we show that it is possible to change this tree to a star-shaped design without increasing the expected number of blocked APs. Again, we consider two cases:

Case 1: There is a leaf s in tree $\mathcal{D}$ that is a servelet.
In this case, let c be the AP connected to servelet s. Removing all edges of c to servelets other than s will decrease the expected number of blocked APs among APs other than c. Furthermore, the blocking probability of c will not increase, since c has a private servelet s.

Case 2: All leaves of $\mathcal{D}$ are APs.
Consider a connected component of $\mathcal{D}$ which is not a star. Now consider a leaf AP c in this component. AP c is connected to servelet s. Servelet s must have a neighboring AP c' which is connected to at least one other servelet s', for otherwise the component would be a star. We claim that removing the edge $c's'$ decreases the expected number of blocked APs. Let $\mathcal{D}'$ be the tree after removing $c's'$.

The blocking probability of all APs except c' decrease in $\mathcal{D}'$. In the following, we prove that removing $c's'$ also decreases the sum of blocking probabilities of the APs c and c'. Let $P_{c'}$ be the probability that all servelets connected to c', except possibly s, are attacked. Let P_s be the probability that one AP other than c' and c in the neighborhood of s fails. As before, let $\mathcal{P}_{\mathcal{D}}(c)$ be the blocking probability of AP c in the design $\mathcal{D}$. Using the fact that $\mathcal{D}$ is a tree, we have

$$\begin{aligned}
\mathcal{P}_{\mathcal{D}}(c) &= p_c + (1-p_c)(p_{c'} + P_s - p_{c'}P_s) \\
\mathcal{P}_{\mathcal{D}}(c') &= p_{c'} + (1-p_{c'})P_{c'}(p_c + P_s - p_cP_s) \\
\mathcal{P}_{\mathcal{D}'}(c) &= p_c + (1-p_c)P_s \\
\mathcal{P}_{\mathcal{D}'}(c') &= p_{c'} + (1-p_{c'})P_{c'}.
\end{aligned}$$

Therefore,

$$\mathcal{P}_{\mathcal{D}}(c) + \mathcal{P}_{\mathcal{D}}(c') = p_c + (1-p_c)(p_{c'} + P_s - p_{c'}P_s) + p_{c'}$$

$$
\begin{aligned}
&+ (1-p_{c'})P_{c'}(p_c + P_s - p_c P_s) \\
&= \mathcal{P}_{\mathcal{D}'}(c) + \mathcal{P}_{\mathcal{D}'}(c') + (p_{c'} - (1-p_{c'})P_{c'})(1-p_c)(1-P_s) \\
&\geq \mathcal{P}_{\mathcal{D}'}(c) + \mathcal{P}_{\mathcal{D}'}(c'),
\end{aligned}
$$

where in the last inequality we use the fact that $p_{c'} \geq \frac{1}{2}$ and $P_{c'} \leq 1$, and hence $p_{c'} - (1-p_{c'})P_{c'} \geq 0$. This completes the proof of this case.

Using the above operations, we can change the tree-shaped design $\mathcal{D}$ to a star-shaped design without increasing the expected number of blocked APs. Hence, we can change any optimal design to an optimal tree-shaped design and then to an optimal star-shaped design. □

Another case for which we can show that there is an optimal star-shaped design is when the number of servelets is two.

Theorem 4. *If the number of servelets is two, then there is an optimal design that is star-shaped.*

Proof. For simplicity, we prove the theorem assuming all APs have the same failure probability p. The proof in the general case is similar. Let $q = 1-p$. Let A_{00}, A_{10}, A_{01}, and A_{11} be the set of APs connected to none of the servelets, to servelet 1, to servelet 2, and to both servelets in an optimal solution. Let $n_{uv} = A_{uv}$ for $0 \leq u, v \leq 1$ and $n = n_{01}+n_{10}+n_{11}$. Let P_1 be the probability that servelet 1 is not attacked. For $i \in A_{10}$, $P_1 = \mathbf{Pr}[i \text{ is blocked}] = 1-q^{n_{10}+n_{11}}$. For $i \in A_{01}$, $P_2 = \mathbf{Pr}[i \text{ is blocked}] = 1-q^{n_{01}+n_{11}}$. For $i \in A_{11}$, $P_3 = \mathbf{Pr}[i \text{ is blocked}] = 1 - q^{n_{01}+n_{11}} - q^{n_{10}+n_{11}} + q^{n_{01}+n_{10}+n_{11}}$. Thus, the expected number of blocked APs is equal to $\mathcal{P}^* = n_{10}P_1 + n_{01}P_2 + n_{11}P_3$. WLOG, assume that $n_{01} \geq n_{10}$. We prove that moving one of the APs from A_{11} to A_{10} decreases the expected number of blocked APs. Before moving this AP from A_{11} to A_{10},

$$
\mathcal{P}^* = n - (n_{10}+n_{11})q^{n_{10}+n_{11}} - (n_{01}+n_{11})q^{n_{01}+n_{11}} + n_{11}q^{n_{01}+n_{10}+n_{11}}
$$

After this movement, the expected number of blocked APs is

$$
\mathcal{P} = n - (n_{10}+n_{11})q^{n_{10}+n_{11}} - (n_{01}+n_{11}-1)q^{n_{01}+n_{11}-1} + (n_{11}-1)q^{n_{01}+n_{10}+n_{11}}
$$

Now we have,

$$
\begin{aligned}
\mathcal{P}^* - \mathcal{P} &= q^{n_{01}+n_{11}-1}[(n_{01}+n_{11})(1-q) - 1 + q^{n_{10}+1}] \\
&\geq q^{n_{01}+n_{11}-1}[(n_{01}+n_{11})p - 1 + (1-p)^{n_{10}+1}] \\
&\geq q^{n_{01}+n_{11}-1}(n_{01}+n_{11}-n_{10}-1)p \\
&\geq 0
\end{aligned}
$$

where the last two inequalities are from $(1-p)^{n_{10}+1} - 1 > -p(n_{10}+1)$ and $n_{01}+n_{11} \geq n_{10}+1$. Therefore, we can move all APs from A_{11} to either A_{10} or A_{01} without increasing the expected number of blocked APs. Thus, there is a star-shaped optimal solution. □

The above proof was based on a local operation that removes one of the edges attached to an AP of degree more than one. However, this local operation can increase the expected number of blocked APs when the number of servelets is more than two.

For example, consider a cycle of size six with three APs and three servelets. It is not hard to show that removing any of the edges of this design will increase the expected number of blocked APs. In the following theorem, we show that without an assumption on the failure probabilities or the number of servelets, the optimal design need not be star shaped.

Theorem 5. *There is an instance of the secure network design problem in which the expected number of blocked APs in the optimal design is larger than that of the optimal star-shaped design by an arbitrary factor.*

Proof. Choose a sufficiently large number m, and let $n = \binom{m}{m/2}$ and $p = 1/n^2$. We first analyze the expected number of blocked APs in the best star-shaped design with these parameters. Let n_i denote the number of APs connected to the ith servelet in such a design, and n_0 denote the number of APs not connected at all. The expected number of blocked APs can be expressed as

$$n_0 + \sum_{i=1}^{m} n_i \left(1 - (1-p)^{n_i}\right) \geq \sum_{i=0}^{m} n_i \left(1 - (1-p)^{n_i}\right).$$

There is at least one i, $0 \leq i \leq m$, with $n_i \geq n/(m+1)$. Thus, the above expression is at least

$$\frac{n}{m+1}\left(1 - (1-p)^{n/(m+1)}\right) \geq \frac{n}{m+1}\left(\frac{pn}{m+1} - \frac{p^2n^2}{(m+1)^2}\right) \geq \frac{pn^2}{2(m+1)^2},$$

where the first inequality follows from $(1-p)^s \leq 1 - ps + p^2s^2$.

Now, we propose a different design and analyze the expected number of blocked APs in such a design. For each of the $n = \binom{m}{m/2}$ APs, we pick a distinct subset of $m/2$ servelets, and connect the AP to the servelets in this set. This design guarantees that if only one AP is attacked, then no other AP will be blocked. We use this to bound the expected number of blocked APs. By the union bound, the probability that more than one AP is attacked can be bounded by n^2p^2. In this case, we bound the number of blocked APs by n. Similarly, with probability at most np, exactly one AP is attacked, and in this case only one AP (the one that is attacked) is blocked. Thus, the expected number of blocked APs is at most $n^2p^2 \times n + np \times 1 = 2/n$.

Therefore, the ratio of the expected number of blocked APs in the latter design to the one in the best star-shaped design is at most $4(m+1)^2/n$, which tends to zero as m tends to infinity. □

3 The Worst-Case Model

In this section, we study a model where an adversary selects at most a given number k of APs to compromise, and the objective is to minimize the number of blocked APs in the worst case. We observe that the worst-case model is closely related to the following problem in extremal combinatorics.

Definition 2. *Let $\mathcal{A} = (A_1, A_2, \ldots, A_n)$ be a family of subsets of the universe $U = \{1, 2, \ldots, m\}$. We call the family $\mathcal{A}$* k-union free *if for any $A_{i_0}, \ldots, A_{i_k} \in \mathcal{A}$ such that $i_j \neq i_t$ for $j \neq t$, we have $A_{i_0} \not\subseteq \cup_{1 \leq j \leq k} A_{i_j}$. In particular, a family $\mathcal{A}$ is* 1*-union free if none of the elements of $\mathcal{A}$ is a subset of another. Let $\mathcal{L}_k(m)$ be the maximum number of subsets in a k-union free family of subsets of the universe $\{1, 2, \ldots, m\}$.*

We call a design *perfect* for k failures, if no matter which k APs fail, no other AP is blocked. It is not difficult to see that there exists a perfect design for k failures with m servelets and n APs if and only if $n \leq \mathcal{L}_k(m)$. The following theorem gives lower and upper bounds on the value of $\mathcal{L}_k(m)$. The lower bound in this theorem is proved by Kleitman and Spencer [10] for a more general problem. We include the proof here for the sake of completeness. We also give an upper bound based on Sperner's theorem. Sperner's theorem gives a tight bound on the maximum number of subsets in a 1-union free family of subsets. See also Ruszinkó [18] for an upper bound for a related problem.

Theorem 6. *For every k and m,*

$$(1 - \frac{k^k}{(k+1)^{k+1}})^{-m/(k+1)} \leq \mathcal{L}_k(m) \leq k \left(1 + \binom{m}{\frac{m}{2}}^{\frac{1}{k}}\right) = O(k 2^{m/k} m^{-1/(2k)}).$$

Proof. We start by proving the upper bound. Let $\mathcal{A} = (A_1, A_2, \ldots, A_n)$ be a k-union-free family of subsets of $\{1, 2, \ldots, m\}$. Consider unions of k distinct sets from $\mathcal{A}$. We claim that no two such unions, say $A_{i_1} \cup \cdots \cup A_{i_k}$ and $A_{j_1} \cup \cdots \cup A_{j_k}$, are equal unless $\{i_1, \ldots, i_k\} = \{j_1, \ldots, j_k\}$. The reason for this is that if two such unions are equal and there is an index i_l not contained in $\{j_1, \ldots, j_k\}$, then we have $A_{i_l} \subseteq A_{j_1} \cup \cdots \cup A_{j_k}$, contradicting the assumption that $\mathcal{A}$ is k-union-free. Therefore, the collection of sets that are obtained by taking the union of k distinct sets in $\mathcal{A}$ contains exactly $\binom{n}{k}$ distinct sets. Furthermore, similar reasoning shows that no set in this collection is contained in another. Therefore, by Sperner's theorem, this collection can contain at most $\binom{m}{m/2}$ sets. Thus,

$$\binom{n}{k} \leq \binom{m}{m/2} \Rightarrow n \leq k \left(1 + \binom{m}{\frac{m}{2}}^{\frac{1}{k}}\right) = O(k 2^{m/k} m^{-1/(2k)}),$$

completing the proof of the upper bound.

To prove the lower bound, we use the probabilistic method to construct a k-union-free collection of sets of the required size. Fix $p = \frac{1}{k+1}$, and pick each of the n sets in the collection by picking each element in $\{1, \ldots, m\}$ independently with probability p. Therefore, for a given set of indices $i_0, i_1, \ldots, i_k$, the probability that $A_{i_0} \subseteq A_{i_1} \cup \cdots \cup A_{i_k}$ is precisely $(1 - p(1-p)^k)^m = (1 - \frac{k^k}{(k+1)^{k+1}})^m$. Therefore, by the union bound, the probability that the collection is not k-union-free is less than $n^{k+1}(1 - \frac{k^k}{(k+1)^{k+1}})^m$. Hence, if we pick $n \leq (1 - \frac{k^k}{(k+1)^{k+1}})^{-m/(k+1)}$, there is a nonzero probability that the resulting collection is k-union-free. This completes the proof of the lower bound. □

Note that the above theorem suggests a randomized algorithm for our network design problem: put each edge in the graph with probability $\frac{1}{k+1}$. We can bound the

expected number of blocked APs resulting from this randomized algorithm using similar ideas of the proof of the above theorem. For small values of k, this algorithm works exponentially better than the optimal star-shaped design.

The only case where we know the exact value of $\mathcal{L}_k(m)$ is when $k = 1$. In this case, we can prove the following stronger theorem.

Theorem 7. *If $k = 1$, then there is a design in which the maximum number of APs an adversary can block is at most $\left\lceil n/\binom{m}{\lfloor m/2 \rfloor} \right\rceil$. Conversely, for every design for such a network, there is a strategy for the adversary to block at least $\left\lceil n/\binom{m}{\lfloor m/2 \rfloor} \right\rceil$ APs.*

Proof Sketch. We can obtain a design for $k = 1$ by duplicating each of the $\binom{m}{\lfloor m/2 \rfloor}$ subsets of size $\lfloor m/2 \rfloor$ of the set of servelets $\left\lceil n/\binom{m}{\lfloor m/2 \rfloor} \right\rceil$ times, and associate an AP to each subset. To prove the other direction, we use the fact that the collection of all subsets of a set of size m can be partitioned into $\binom{m}{\lfloor m/2 \rfloor}$ chains. Therefore, in every design there are at least $\left\lceil n/\binom{m}{\lfloor m/2 \rfloor} \right\rceil$ APs that are connected to sets of servelets belonging to the same chain. Hence, if the adversary compromises the AP connected to the subset at the top of this chain, all other APs connected to the subsets in this chain will fail. □

4 Conclusion

In this paper, we presented the first theoretical study of the secure network design problem. We showed that in the average case model, when failure probabilities are large (greater than $\frac{1}{2}$), there is an optimal star-shaped design, and such a design can be computed in polynomial time. On the other hand, there are instances with small failure probabilities where the optimal star-shaped design is arbitrarily worse than the optimal design. The case of small failure probabilities seems to be related to the stronger model where an adversary is allowed to select at most k APs to compromise. We observed that in this model, a random design performs considerably better than the optimal star-shaped design.

We still do not know of any hardness result or a polynomial-time algorithm for the general case of the secure network design problem, although the connection between this problem and the problem of finding a tight bound on the size of the largest k-union-free family of sets (which is a long-standing open problem) suggests that computing the exact optimum is difficult. Even an approximation algorithm for this problem, or tighter bounds for the k-union-free problem, would be interesting. Lovasz Local Lemma gives us a small improvement in the lower bound, but more significant improvement seem to require new techniques. Finally, it would be interesting to prove Theorem 3 with a weaker assumption (e.g., that probabilities are greater than a small constant), or show that such a generalization is not true.

References

1. M. Adler. Tradeoffs in probabilistic packet marking for ip traceback. In *Proc. ACM Symposium on Theory of Computing (STOC)*, May 2002.
2. T. Bu, S. Norden, and T. Woo. Trading resiliency for security: Model and algorithms. In *Proc. IEEE International Conference on Network Protocols(ICNP)*, 2004.
3. H. Burch and B. Cheswick. Tracing anonymous packets to their approximate source. In *Proc. USENIX LISA*, pages 319–327, December 2000.
4. D. Dean, M. Franklin, and A. Stubblefield. An algebraic approach to IP traceback. In *Proc. NDSS*, pages 3–12, February 2001.
5. T. Doeppner, P. Klein, and A. Koyfman. Using router stamping to identify the source of IP packets. In *Proc. ACM CCS*, pages 184–189, November 2000.
6. P. Ferguson. *Network Ingress Filtering: Defeating Denial of Service Attacks Which Employ IP Source Address Spoofing.* RFC 2267, January 1998.
7. L. Garber. Denial-of-service attacks rip the Internet. *IEEE Computer*, 33(4):12–17, April 2000.
8. M. T. Goodrich. Efficient packet marking for large-scale IP traceback. In *Proc. ACM CCS*, pages 117–126, November 2002.
9. A. D. Keromytis, V. Misra, and D. Rubenstein. SOS: Secure overlay services. In *Proc. ACM SIGCOMM*, pages 61–72, August 2002.
10. D. Kleitman and J. Spencer. Families of k-independent sets. *Discrete Mathematics*, 6:255–262, 1973.
11. J. Li, M. Sung, J. Xu, and L.E. Li. Large-scale ip traceback in high-speed internet: Practical techniques and theoretical foundation. In *Proc. IEEE Symposium on Security and Privacy*, pages 115–129, 2004.
12. R. Mahajan, S. Bellovin, S. Floyd, J. Ioannidis, V. Paxson, and S. Shenker. Controlling high bandwidth aggregates in the network. *ACM Computer Communication Review*, 32(3):62–73, July 2002.
13. D. McGuire and B. Krebs. Attack on internet called largest ever. http://www.washingtonpost.com/wp-dyn/articles/A828-2002Oct22.html, October 2002.
14. Jelena Mirkovic, Gregory Prier, and Peter Reiher. Attacking DDoS at the source. In *Proc. IEEE ICNP*, pages 312–321, November 2002.
15. Jelena Mirkovic, Max Robinson, Peter Reiher, and Geoff Kuenning. Alliance formation for ddos defense. In *Proc. New Security Paradigms Workshop, ACM SIGSAC*, August 2003.
16. Christos Papadopoulos, Robert Lindell, John Mehringer, Alefiya Hussain, and Ramesh Govidan. COSSACK: coordinated suppression of simultaneous attacks. In *DISCEX III*, pages 22–24, April 2003.
17. K. Park and H. Lee. On the effectiveness of route-based packet filtering for distributed DoS attack prevention in power-law Internets. In *Proc. ACM SIGCOMM*, pages 15–26, August 2001.
18. M. Ruszinkó. On the upper bound of the size of the r-cover-free families. *Journal of Combinatorial Theory, Series A*, 66:302–310, 1994.
19. S. Savage, D. Wetherall, A. Karlin, and T. Anderson. Practical network support for IP traceback. In *Proc. ACM SIGCOMM*, pages 295–306, August 2000.
20. A. Snoeren, C. Partridge, et al. Hash-based IP traceback. In *Proc. ACM SIGCOMM*, pages 3–14, August 2001.
21. D. Song and A. Perrig. Advanced and authenticated marking schemes for IP traceback. In *Proc. IEEE INFOCOM*, pages 878–886, April 2001.

22. J. Vijayan. Akamai attack reveals increased sophistication: Host's DNS servers were DDoS targets, slowing large sites. http://www.computerworld.com/securitytopics/security/story/0,10801,93977p2,00.html, June 2004.

A Expected number of blocked APs: A Hardness result

In terms of hardness, we can show that given a particular design, it is hard to compute the probability that a given AP is blocked, and the expected number of APs that will be blocked.

Theorem 8. *The following two problems are #P-complete:*

- *Given a design and assuming uniform failure probabilities of $p = 1/2$, compute the probability that a given AP i will be blocked.*
- *Given a design and assuming uniform failure probabilities of $p = 1/2$, compute the expected number of APs that will be blocked.*

Proof Sketch. For the first problem, we can give a reduction from the problem of computing the number of solutions of a set-cover instance. The second problem can be reduced to the first by adding a "private servelet" for each AP except one. □

Even though finding the exact expected number of blocked APs is hard, it is not hard to approximate within a factor of $1 + \epsilon$ for any positive constant ϵ by sampling polynomially many times and taking the average. Note that the above theorem does not show any hardness result for finding the optimal network. The complexity of this problem for general failure probabilities is still open.