

Collaborative Attack vs. Collaborative Defense (Extended Abstract)

Shouhuai Xu

Department of Computer Science, University of Texas at San Antonio
`shxu@cs.utsa.edu`

Abstract. We have witnessed many attacks in the cyberspace. However, most attacks are launched by individual attackers even though an attack may involve many compromised computers. In this paper, we envision what we believe to be the next generation cyber attacks — collaborative attacks. Collaborative attacks can be launched by multiple attackers (i.e., human attackers or criminal organizations), each of which may have some specialized expertise. This is possible because cyber attacks can become very sophisticated and specialization of attack expertise naturally becomes relevant. To counter collaborative attacks, we might need collaborative defense because each “chain” in a collaborative attack may be only adequately dealt with by a different defender. In order to understand collaborative attack and collaborative defense, we present a high-level abstracted framework for evaluating the effectiveness of collaborative defense against collaborative attacks. As a first step towards realizing and instantiating the framework, we explore a characterization of collaborative attacks and collaborative defense from the relevant perspectives.

Keywords: cyber security, cyber attack, collaborative attack, collaborative defense.

1 Introduction

Both academic cyber (or Internet) security research and commercial cyber security defense have mainly focused on understanding and combating *individual* cyber attacks. For example, we all should install, and frequently update, computer virus scanners or spyware detectors so as to protect our computers as well as important information stored on them (e.g., personal data files, private communications, credit card numbers, social security numbers, and bank account passwords). Individual attacks have caused severe damages and, for example, there have reportedly many stolen digital identities such as credit card numbers.

In this paper we envision a new class of cyber attacks, called *collaborative attacks*, which might represent the next generation cyber attacks. Collaborative attacks are characterized by the prevalence of coordination before and during attacks. Distributed Denial-of-Service (DDoS) attacks can be seen as a simple

example of collaborative attacks in that they involve a large number of compromised computers, which however are often controlled by a *single* attacker (e.g., a botnet master). Collaborative attacks in general would involve multiple human attackers or criminal organizations that have respective adversarial expertise but may not fully trust each other.¹ Intuitively, collaborative attacks are more powerful than the sum of the underlying individual attacks that can be launched by the individual attackers independently. As an analogy, we may think collaborative attacks as “chemical reactions” and individual cyber attacks as “chemical elements.” Just like new (or fatal) materials can result from carefully designed chemical reactions, severe damage can be caused by well-coordinated less-powerful cyber attacks. In other words, collaborative attacks can exhibit the “ $1 + 1 > 2$ ” phenomenon.

1.1 Our Contributions

In this paper, we make the following contributions:

- We envision the need to consider collaborative attacks. We also envision the need to defend collaborative attacks with collaborative defense because collaborative defense can certainly exhibit the “ $1+1>2$ ” phenomenon as well. This is not only because different defenders may have different expertise, but also because sharing information between the defenders would play a crucial role in successfully and effectively defending collaborative attacks.
- We present a conceptual framework for understanding, characterizing, and evaluating the effectiveness of collaborative defense against collaborative attacks. Within the framework we can ask interesting questions such as: How may we de-incentivize the attackers from launching collaborative attacks, and how can we incentivize the defenders to collaborate on defense?
- Towards a full-fledged realization of the afore-mentioned framework, we explore the attributes of collaborative attacks and the attributes of collaborative defense. It is interesting to note that essentially the attributes of collaborative attacks mirror the attributes of collaborative defense.

The research problems introduced in the present paper may be more important than the content itself. Thus, we certainly hope that the paper will inspire more research activities towards solving them.

¹ In traditional cryptographic and security models, we often assume that all the compromised participants are controlled by a single attacker (for free). This assumption is legitimate when we talk about specific goals (e.g., when we discuss security of a digital signature scheme, we allow the attacker to compromise any private keys other than the one in question) within a relatively small-scale system (e.g., when we discuss security in a mobile ad hoc network). When we talk about attacks in the cyberspace in general, this assumption does not always hold. Indeed, there are likely many cyber criminal organizations that might have different expertise and might be (e.g., economically) motivated to launch collaborative attacks without having any centralized authority whatsoever. As such, collaborative attacks may be seen as a sort of *emergent property*.

1.2 Related Work

We are not aware of any prior work on dealing with collaborative attacks, except [5] in which we explore a modeling of coordinated internal and external attacks. Nevertheless, it is worthwhile to note that a somewhat related problem known as *alert correlation* has been extensively studied. However, alert correlation is different from collaborative attacks because it is motivated by the problem that IDSs often overload their human operators with a large number of simple alerts of low-level security-related events, while not being able to provide the often more important succinct and high-level view of multi-stage intrusion incidents [9,8,2]. Various approaches have been proposed for alert correlation, such as temporal correlation [7], spatial correlation for identifying attack sources [4,1], root cause detection [3], prerequisites-consequences correlation [6], and logical correlation [10]. As such, alert correlation mainly deals with attacks launched by a single attacker, rather than dealing with attacks launched by multiple attackers that may have different attack “fingerprints” that facilitate correlation. Nevertheless, alert correlation may be able to help deal with collaborative attacks to some extent as well.

Outline. The rest of the paper is organized as follows. In Section 2 we present a high-level framework for evaluating the effectiveness of collaborative defense against collaborative attack. In Section 3 we elaborate the characterization of collaborative attack, and in Section 4 we elaborate the characterization of collaborative defense. In Section 5 we conclude the paper with open problems for future research.

2 Collaborative Attack vs. Collaborative Defense: An High-Level Evaluation Framework

A networked system S_i may consist of a set of elementary components, which may be specific to the system properties we care. For example, when we only care about which private digital signing keys are compromised in a networked system, we may only take into consideration the software and hardware component instances that could directly or indirectly cause the compromise of the private keys. As a result, a set of individual systems S_1, \dots, S_n can compose a larger system $S = \cup_{1 \leq i \leq n} S_i$, where S_i is the set of component instances in the i th system. Note that the system space may impose a partial order over the power set of the set of the elementary component instances.

For system S , we have a set of assets or targets (e.g., private digital signing keys, passwords or digital identities) whose security is our concern. Let Ω be the power set of the set of the elementary assets. As a consequence of a successful cyber attack, some assets $\omega \in \Omega$ are compromised.

Let $\gamma_i \in \Gamma$ denote the defense resources (or mechanisms) in a networked system S_i , and $\mathcal{D}(\gamma_i)$ the actual defense strategy and tactics used by the defender with defense resources γ_i . Given individual defense γ_i in system S_i , where $i \in I$ for some index set I , the resources for collaborative defense $\cup_{i \in I} \gamma_i$ is well-defined

with respect to the composed system $\cup_{i \in I} S_i$. It should be clear that the defense space, Γ , may impose a partial order over the power set of the set of the elementary defense resources. As such, $\mathcal{D}(\cup_{i \in I} \gamma_i)$ is the collaborative defense given resources $\cup_{i \in I} \gamma_i$. Note that $\mathcal{D}(\cdot)$ is indeed a class of algorithms for computing the defense, although $\mathcal{D}(\cdot)$ may not (always) produce the optimal defense.

Let $\theta_i \in \Theta$ denote the resources used by an individual attack against a networked system S_i with defense resources γ_i and defense algorithm \mathcal{D} , and $\mathcal{A}(\theta_i, \mathcal{D}(\gamma_i))$ the actual attack strategy and tactics used by the attacker with attack resources θ_i . To accommodate the worst-case scenario, we can assume that $\mathcal{A}(\theta_i, \mathcal{D}(\gamma_i))$ is the optimal attack against system S_i , where optimization intuitively means that it will cause the worst (from the defender's perspective) outcome, and will become fully specified later. Given that, $\mathcal{A}(\cdot, \cdot)$ is indeed a class of algorithms computing the optimal attacks based on given attack resources, defense resources and defense algorithm.

Given individual attack resources θ_i against system S_i with defense resources γ_i and defense algorithm \mathcal{D} , where $i \in I$ for some index set I , the resources for collaborative attack $\cup_{i \in I} \theta_i$ is well-defined against the composed system $\cup_{i \in I} S_i$. It should be clear that the attack resource space Θ may impose a partial order over the power set of the set of elementary attack resources. Similarly, based on the combined attack resources $\cup_{i \in I} \theta_i$, the optimal collaborative attack is given by $\mathcal{A}(\cup_{i \in I} \theta_i, \mathcal{D}(\cup_{i \in I} \gamma_i))$.

Ultimately, we want to fully specify the function $f_{\mathcal{D}, \mathcal{A}} : \Theta \times \Gamma \times \{S\} \rightarrow \Omega$ such that

$$f_{\mathcal{D}, \mathcal{A}}(\theta, \gamma, S) \mapsto \omega,$$

where $\mathcal{A}(\theta, \mathcal{D}(\gamma))$ is an attack against a system S with defense $\mathcal{D}(\gamma)$, and $\omega \in \Omega$ is the outcome of launching attack $\mathcal{A}(\theta, \mathcal{D}(\gamma))$ against system S . Note that \mathcal{A} is optimal if for any \mathcal{D} and \mathcal{A}' , and for all θ, γ and S , we have

$$f_{\mathcal{D}, \mathcal{A}'}(\theta, \gamma, S) \subseteq f_{\mathcal{D}, \mathcal{A}}(\theta, \gamma, S).$$

Since f may impose a partial order over Ω , the above definition of optimization may be adjusted as

$$\text{payoff}(f_{\mathcal{D}, \mathcal{A}'}(\theta, \gamma, S)) \leq \text{payoff}(f_{\mathcal{D}, \mathcal{A}}(\theta, \gamma, S))$$

where payoff is an appropriate payoff function.

Unfortunately, it may be very difficult to fully specify the function $f_{\mathcal{D}, \mathcal{A}}$, which means that we may have to approach it through various "approximations." Still, there is a plenty of interesting questions that can be asked.

- Fixing (collaborative or non-collaborative) defense algorithm \mathcal{D} , attack resources $\theta \in \Theta$ and a system S , is $f_{\mathcal{D}, \mathcal{A}}(\theta, \gamma, S)$ a decreasing function of (collaborative or non-collaborative) defense resources γ ? Under what conditions the function $f_{\mathcal{D}, \mathcal{A}}(\theta, \gamma, S)$ is convex, concave, linear, or exhibits a phase transition with respect to γ ? Note that in general $f_{\mathcal{D}, \mathcal{A}}(\theta, \gamma, S)$ may impose a partial order, meaning that we may need to resort to some payoff function as discussed above.

- Fixing (collaborative or non-collaborative) defense resources $\gamma \in \Gamma$, defense algorithm \mathcal{D} and a system S , is $f_{\mathcal{D},\mathcal{A}}(\theta, \gamma, S)$ an increasing function of (collaborative or non-collaborative) attack resources θ ? Under what conditions the function is convex, concave, linear, or exhibits a phase transition with respect to θ ? Note that $f_{\mathcal{D},\mathcal{A}}(\theta, \gamma, S)$ may impose a partial order in general, meaning that we may resort to some payoff function as mentioned before.
- Fixing a system S , when both (collaborative or non-collaborative) attack and (collaborative or non-collaborative) defense resources are adaptively determined, meaning that both θ and γ are functions of time t and denoted by $\theta(t)$ and $\gamma(t)$, what are the dominating factors that determine the outcome $f_{\mathcal{D},\mathcal{A}}(\theta(t), \gamma(t), S)$? In particular, for a given $\theta(t)$, how can we optimally select adaption methods $\gamma(t)$, perhaps with minimal extra effort, so as to minimize $f_{\mathcal{D},\mathcal{A}}(\theta(t), \gamma(t), S)$? Note that $f_{\mathcal{D},\mathcal{A}}(\theta(t), \gamma(t), S)$ may impose a partial order in general, meaning that we may resort to some payoff function as mentioned before.
- For a given (collaborative or non-collaborative) defense resources $\gamma \in \Gamma$ in a system S , what is the minimal attack α in order for a (collaborative or non-collaborative) attacker to achieve an attack goal $\omega \in \Omega$? In other words, we need to identify the minimal attack effort corresponding to collaborative attack θ such that

$$\alpha = \min\{\theta : f_{\mathcal{D},\mathcal{A}}(\theta, \gamma, S) \mapsto \omega\}.$$

Note that since the collaborative attack space may impose a partial order, there may be multiple such α 's. Nevertheless, if we are able to define a payoff function on the attack space Θ , it is possible to reduce the size of the set of collaborative attacks corresponding to the minimal attack efforts.

- Fixing (collaborative or non-collaborative) defense $\gamma \in \Gamma$ in a system S , it is important to characterize under what conditions we have

$$\cup_{i \in I} f_{\mathcal{D},\mathcal{A}}(\theta_i, \gamma, S) \subseteq f_{\mathcal{D},\mathcal{A}}(\cup_{i \in I} \theta_i, \gamma, S)$$

or

$$\text{payoff}(\cup_{i \in I} f_{\mathcal{D},\mathcal{A}}(\theta_i, \gamma, S)) \leq \text{payoff}(f_{\mathcal{D},\mathcal{A}}(\cup_{i \in I} \theta_i, \gamma, S))$$

based on an appropriate payoff function **payoff**. This allows us to disrupt the power of collaborative attacks so as to ensure an “ $1 + 1 \leq 2$ ” effect and to deinceivize the attackers from launching collaborative attacks.

- For a given (collaborative or non-collaborative) attack $\theta \in \Gamma$ against a system S , what is the minimal attack β in order for a (collaborative or non-collaborative) defense to successfully protect asset $\omega \in \Omega$? In other words, we need to identify the minimal defense effort corresponding to collaborative attack θ such that

$$\beta = \min\{\gamma : f_{\mathcal{D},\mathcal{A}}(\theta, \gamma, S) \mapsto \omega\}.$$

Note that since the collaborative defense space may impose a partial order, there may be multiple such β 's. Nevertheless, if we are able to define an

investment function on the defense space Γ , it is possible to reduce the size of the set of collaborative defense corresponding to the minimal investment.

- Fixing (collaborative or non-collaborative) attack $\theta \in \Theta$ against a system S , it is important to characterize under what conditions we have

$$\cup_{i \in I} f_{\mathcal{D}, \mathcal{A}}(\theta, \gamma_i, S) \subseteq f_{\mathcal{D}, \mathcal{A}}(\theta, \cup_{i \in I} \gamma_i, S)$$

or

$$\text{payoff}(\cup_{i \in I} f_{\mathcal{D}, \mathcal{A}}(\theta, \gamma_i, S)) \leq \text{payoff}(f_{\mathcal{D}, \mathcal{A}}(\theta, \cup_{i \in I} \gamma_i, S))$$

based on an appropriate payoff function **payoff**. This gives the defenders incentives to collaborate in defending cyber attacks.

In order to answer the above questions, we need to fully specify the attack space and the defense space. As a first step towards this goal, in what follows we present a characterization of collaborative attack and collaborative defense.

3 A Characterization of Collaborative Attack

We believe that collaborative attack and collaborative defense have much in common, especially they need some kinds of Command & Control (C&C) for coordinating attack and defense, respectively. Given that, we characterize them from the same five perspectives. Specifically, for collaborative attacks we consider the time-aspect of collaborative attack C&C, the space-aspect of collaborative attack C&C, the effect of collaborative attack, the information exchange during collaborative attack, and the privacy aspect of collaborative attack. Figure 1 highlights the perspectives.

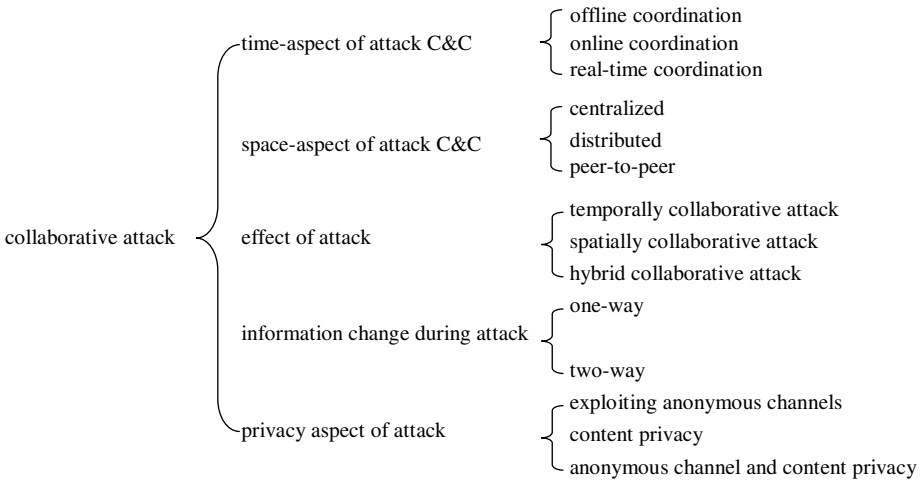


Fig. 1. A characterization of collaborative attack

Attribute 1: Time-aspect of collaborative attack C&C. C&C mechanisms are used for coordinating collaborative attacks. There is a spectrum of coordination methods from a time perspective, ranging from the least sophisticated off-line coordination to the most sophisticated real-time coordination. Various on-line coordination methods reside in between.

- **Off-line coordination:** The attackers command a set of adversarial computers (e.g., bots) to launch a future attack against some predetermined target. During the attack, there are no communications between the attackers and the adversarial computers, nor communications between the adversarial computers themselves. This means that the course of the attack process will not be adjusted according to the situation. Distributed Denial-of-Service (DDoS) attacks are often launched via an off-line coordination method.
- **On-line coordination:** In addition to off-line coordination, during an attack there may be communications between the attackers and the adversarial computers, or communications between the adversarial computer themselves. Moreover, newly compromised computers can become adversarial computers and launch attacks against other computers. On-line coordination gives the attackers extra power because the attackers can more effectively utilize their resources. For example, when the attackers realized that the same effect can be achieved with a subset of the adversarial computers, the attackers can withdraw some of them so as to reduce the chance they are caught (and punished).
- **Real-time coordination:** In this case, both the attackers and the adversarial computers, initially and later compromised alike, are always updated with the current global system state information (e.g., which computers have yet to be compromised). As such, the attackers can arbitrarily orchestrate an ongoing attack through a real-time C&C mechanism. Therefore, the attackers can command any adversarial computer to attack any target computer in a real-time fashion.

In general, off-line coordination is less powerful than on-line coordination, which in turn is less powerful than real-time coordination because off-line coordination does not accommodate any situational awareness in the cyberspace, whereas real-time coordination accommodates the most situational awareness in the cyberspace. In particular, attacks under on-line or real-time coordination can be made stealthy by avoiding any heavy or unnecessary use of resources of (compromised) computers and networks.

Attribute 2: Space-aspect of collaborative attack C&C. This captures the whereabouts of the C&C system. There are three kinds: centralized, distributed, and peer-to-peer.

- **Centralized C&C:** In this case, there is a single attacker that is coordinating the collaborative attacks, which may involve the adversarial computers controlled by multiple attackers (e.g., multiple attackers designate a single commander to exploit their botnets to launch attacks). Traditional IRC-based botnet C&C can be seen as a special example of centralized C&C.

- **Distributed C&C:** In this case, there are multiple attackers for commanding the adversarial computers to launch attacks. The commanding attackers may formulate some topology, which may reflect the relationship between them. For example, there may be a hierarchical structure (e.g., a tree) between the commanding attackers such that the “leaf” attackers actually deliver commands to the adversarial computers.
- **Peer-to-peer C&C:** In this case, there are multiple attackers that play equal roles. They can formulate a logical (i.e., a command is approved by multiple of them) and/or physical (i.e., the network connecting them formulates a graph) peer-to-peer network. Clearly, it is difficult to shut down such C&C networks, which have recently been exploited by some botnets.

In general, centralized C&C is less sophisticated than distributed C&C, which in turn is less sophisticated than peer-to-peer C&C. An important research problem is to identify and exploit the weaknesses of distributed and peer-to-peer C&Cs, if any, to better defend against them.

Attribute 3: Effect of collaborative attacks. The effect of collaborative attacks can be classified as spatially collaborative attacks, temporally collaborative attacks, and hybrid collaborative attacks.

- **Spatially collaborative attacks:** The set of adversarial computers, which are located in different geographic or network places, are coordinated to launch attacks against a target at (roughly) the same time. DDoS attacks often bear this characteristic.
- **Temporally collaborative attacks:** The attack may proceed in a well orchestrated fashion. For example, the first step is to shut down the IDS employed in the target system by one attacker, the second step is to disable the virus scanners installed at the target system by another attacker, and the final step is to launch the real attack against the target by yet another attacker (e.g., stealing confidential data from a data center without being noticed by the defender). Each step may be accomplished through a different set of adversarial computers, which may reside at different geographic or network places.
- **Hybrid collaborative attacks:** These attacks bear the characteristics of the spatially collaborative attacks and temporally collaborative attacks.

In general, spatially collaborative attacks are not compatible with temporally collaborative attacks. However, both of them are not as sophisticated as hybrid collaborative attacks.

Attribute 4: Information exchange during collaborative attacks. During an collaborative attack, information may be exchanged between the commanding attackers, between the commanding attackers and the adversarial computers, and between the adversarial computers. There are two kinds of information exchanges.

- **One-way:** In this case, information may only be sent from one participant to another (e.g., the adversarial computers always report to the respective commanding attackers about their progress), but not in the other direction

(i.e., the commanding attackers may not send direct commands to the adversarial computers). This is possible because sending information to the adversarial computers, which is often a large number, may increase the chance that the commanding attackers or the adversarial computers are detected.

- **Two-way:** In this case, information may be sent from any computer to any other. This allows the sharing of situational awareness, which may be needed in order to launch sophisticated attacks.

In general, one-way information exchange is less powerful and less sophisticated than two-way information exchange.

Attribute 5: Privacy aspect of collaborative attacks. Attackers may abuse some advanced techniques to launch more sophisticated attacks. For example,

- **Exploiting anonymous channels:** In this case, the attackers exploit anonymous channels or their variants, which may or may not be deployed for legitimate uses, to conduct stealthy communications.
- **Enforcing content privacy:** In this case, the attackers exploit cryptographic techniques or their variants to protect the content of their communications (e.g., commands for attacks).
- **Exploiting anonymous channels and enforcing content privacy:** The attackers may exploit anonymous channels and enforce content privacy.

In general, attack exploiting such techniques are often difficult to deal with.

4 A Characterization of Collaborative Defense

In parallel to the characterization of collaborative attacks, we characterize collaborative defense from the same five perspectives, namely the time-aspect of collaborative defense C&C, the space-aspect of collaborative defense C&C, the effect of collaborative defense, the information exchange during collaborative defense, and the privacy aspect of collaborative defense. Figure 2 highlights the perspectives.

Attribute I: Time-aspect of collaborative defense C&C. C&C mechanisms are used for coordinating collaborative defense. There is also a spectrum of coordination methods.

- **Off-line coordination:** The defenders coordinate their defenses regardless of the specific attacks.
- **On-line coordination:** In addition to off-line coordination, there may be communications between the defenders during an attack so as to share information about situational awareness.
- **Real-time coordination:** In this case, the defenders are always updated with the current global system state information. As such, the defenders can orchestrate an ongoing defense through a real-time C&C mechanism.

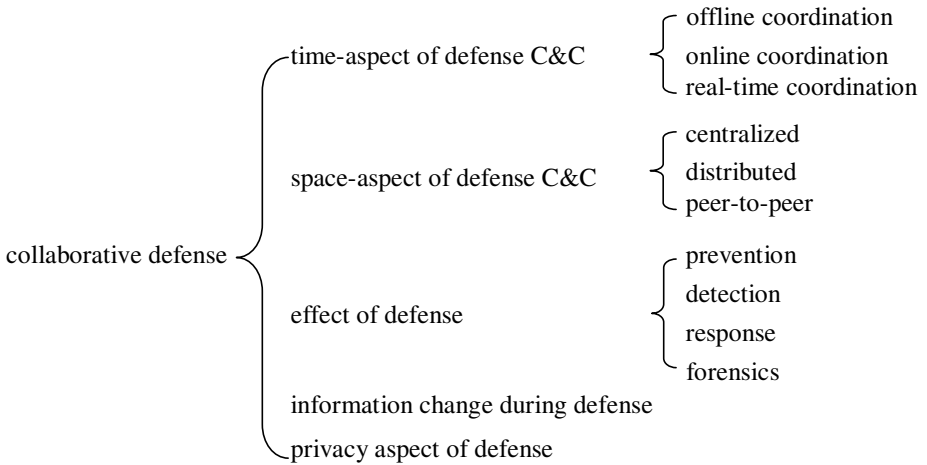


Fig. 2. A characterization of collaborative defense

In general, off-line coordination is less powerful than on-line coordination, which in turn is less powerful than real-time coordination because off-line coordination does not accommodate any situational awareness in the cyberspace, whereas real-time coordination accommodates the most situational awareness in the cyberspace. It is stressed, however, that the situational awareness may be misleading when sophisticated attackers can exploit deception to fool the detection sensors, and thus cause severe problems.

Attribute II: Space-aspect of collaborative defense C&C. This captures where the C&C system is located.

- **Centralized C&C:** In this case, there is a designated defender that is coordinating the collaborative defense, which may involve the resources of multiple defenders.
- **Distributed C&C:** In this case, there are multiple defenders that may formulate some topology, which may reflect the relationship between the defenders. For example, there may be a hierarchical structure (e.g., a tree) between the defenders.
- **Peer-to-peer C&C:** In this case, there are multiple defenders that play equal roles. They can formulate a logical (i.e., a command is approved by multiple of them) and/or physical (i.e., the network connecting them formulates a graph) peer-to-peer network.

In general, centralized C&C is less robust than distributed C&C, which in turn is less robust than peer-to-peer C&C.

Attribute III: Effect of collaborative defense. Collaborative defense should apply to the whole lifecycle of networked systems.

- **Collaborative prevention:** The defenders collaboratively prevent attackers from launching successful attacks.
- **Collaborative detection:** Defenders can share information about suspicious activities against their own networked systems so as to detect attacks that may be launched by multiple collaborative attackers.
- **Collaborative response:** During an attack, the defenders can collaboratively deal with attacks by allocating defense resources. For example, one defender's network may have been recruited as a botnet to launch attacks against another defender's network. Shutting down the botnet computers would help eliminate the attacks against the victim.
- **Collaborative forensics:** After the fact that multiple networks have been attacked, the defenders can share information so as to answer questions such as: When did an attack occur? How did it occur? How long did the attack last? What are the consequences (e.g., Which computers were broken? What information was stolen?) What are the possible attackers, supposing we know that different attackers have their fingerprints in, for example, their malware?

Attribute IV: Information exchange during defense. During defense, information may be exchanged between the defenders. In general, the information exchange should be two-way, meaning that any defender can send information to any other defender.

Attribute V: Privacy aspect of defense. There may be privacy issues when the defenders collaborate in defending attacks. In particular, one defender may not be willing to share some information about their assets (e.g., their internal network configurations).

5 Summary and Future Work

We envisioned what we believe to be the next generation cyber attacks, called collaborative attacks. To counter collaborative attacks, we might need collaborative defense. We presented a framework for understanding, characterizing and evaluating the effectiveness of collaborative defense against collaborative attacks. As a first step towards realizing and instantiating the framework, we explored a characterization of collaborative attacks and collaborative defense from the relevant perspectives.

As demonstrated in the paper, there are many challenging and important research problems. Thus, we hope that this paper will inspire active research toward understanding and adequately addressing collaborative attacks.

Acknowledgement. The chemical reaction analogy of collaborative attacks was inspired by a conversation with Bharat Bhargava.

The work was supported in part by AFOSR and NSF. The views and opinions contained in the paper are those of the author and should not be interpreted as, in any sense, the official policies or endorsements of the government or the agencies.

References

1. Allman, M., Blanton, E., Paxson, V., Shenker, S.: Fighting coordinated attackers with cross-organizational information sharing. In: *HOTNETS 2006* (2006)
2. Green, J., Marchette, D., Northcutt, S., Ralph, B.: Analysis techniques for detecting coordinated attacks and probes. In: *Proceedings of the Workshop on Intrusion Detection and Network Monitoring*, pp. 1–9 (1999)
3. Julisch, K.: Clustering intrusion detection alarms to support root cause analysis. *ACM Trans. Inf. Syst. Secur.* 6(4), 443–471 (2003)
4. Katti, S., Krishnamurthy, B., Katabi, D.: Collaborating against common enemies. In: *Proceedings of the 5th ACM SIGCOMM conference on Internet Measurement (IMC 2005)*, p. 34 (2005)
5. Li, X., Xu, S.: A stochastic modeling of coordinated internal and external attacks (manuscript in submission)
6. Ning, P., Cui, Y., Reeves, D.: Constructing attack scenarios through correlation of intrusion alerts. In: *Proceedings of the 9th ACM conference on Computer and communications security (CCS 2002)*, pp. 245–254 (2002)
7. Ourston, D., Matzner, S., Stump, W., Hopkins, B.: Coordinated internet attacks: responding to attack complexity. *Journal of Computer Security* 12(2), 165–190 (2004)
8. Valdes, A., Skinner, K.: Probabilistic alert correlation. In: *Proceedings of the 4th International Symposium on Recent Advances in Intrusion Detection (RAID 2000)*, pp. 54–68 (2001)
9. Valeur, F., Vigna, G., Kruegel, C., Kemmerer, R.: A comprehensive approach to intrusion detection alert correlation. *IEEE Trans. Dependable Secur. Comput.* 1(3), 146–169 (2004)
10. Zhou, J., Heckman, M., Reynolds, B., Carlson, A., Bishop, M.: Modeling network intrusion detection alerts for correlation. *ACM Trans. Inf. Syst. Secur.* 10(1), 4 (2007)