Defensive Cyber Deception

SaTC 2019 PI meeting breakout group report
Co-leads: Aron Laszka, Deniz Gurkan, Rakesh Verma (University of Houston)

1. Problem/Domain Summary

Our breakout group discussed research challenges and approaches for cyber deception.

Most security measures focus on preventing intrusions (e.g., access control) or on reacting to them (e.g., intrusion detection and mitigation). Such measures advantage the adversary by letting it surveil its targets and deliberate its actions without facing any opposition before attacking. To diminish the adversary's advantage, we need to take a more proactive approach that engages the adversary early and thwarts its efforts even before the attack.

Cyber deception is an emerging proactive defense approach that aims to thwart the adversary's effort by providing it with false information. Typical approaches for implementing cyber deception include masquerading or hiding real networks, systems, and services and deploying fake ones (e.g., honeypots and honeynets). Defenders may use cyber deception to surveil adversaries in action, detect intrusions (when an intruder probes a fake asset), delay the intruders' lateral movement, etc.

A wide range of techniques and tools are available for implementing deception in practice. Recently, significant effort has been spent on developing models for planning and evaluating deception formally. However, there still remains a significant gap between these models and practical deployments.

2. Key Research Challenges

- **Difference between security by obscurity and deception**
    - Cyber deception may appear similar to *security by obscurity* since both try to diminish the correct information that is available to the adversary. A key research challenge is to distinguish between deception and *security by obscurity*, and to establish formal models and strong security guarantees for deception techniques.
    - Our group noted that in practice, people do use *security by obscurity* and that it does seem to provide some security benefits in some cases (e.g., SSH service on a standard port is likely to be probed frequently, while on a non-standard port it is almost never probed). However, we do not know how much benefit *security by obscurity* provides, and we have no guarantees about these benefits. Our group also discussed that practical cryptography could also be viewed as *security by obscurity* since there is a key to infer, as in RSA, but we hope that attackers do not have sufficient (computational) power to infer this key.

    - Our group also noted that deception terminology is not always used consistently in literature (see, e.g., relation between MTD and deception below).
- **Quantifying the benefit of deception:**
    - Before applying any security measure, practitioners need to be able to perform a cost/benefit analysis. In the case of cyber deception, both quantifying costs and quantifying benefits are challenging (see next point as well); however, the latter appears to be a more challenging research problem.
    - Our group discussed examples of how deception can benefit defenders. Examples include detecting intruders when they probe or exploit fake assets (e.g., honeypots) on an internal network and delaying adversaries by wasting their effort.
- **Quantifying the cost of deception:**
    - Cyber deception can have usability and performance costs. For example, masquerading or hiding real assets can hurt real users since they need to access the real information. On the other hand, deploying fake assets (or generating fake traffic from virtual fake assets) requires resources or may negatively impact performance. To apply cyber deception, practitioners need to be able to understand and quantify these costs. However, capturing usability costs can be very challenging and requires further research.
- **Techniques for implementing deception:**
    - Our group also discussed implementations of cyber deception. We noted that deception does not necessarily require defenders to actually deploy fake assets or to actually change real ones since fake information may be fed to the adversary by other means, e.g., by generating fake traffic from sources that do not exist. In the networking domain, deception can be implemented using SDN, which can be used to create entire fake networks (e.g., adversary may be virtually "moved" to a fake network from a real one). We also noted that deception should be implemented on the least addressed threat vectors and on the most valuable assets.
- **Lack of data on deception in practice:**
    - Currently, the research community lacks access to large-scale datasets on how deception is used in practice, what costs are incurred, and what the benefits (or perceived benefits) are. To some extent, this is due to the fact that deception is often used in practice as *security by obscurity*. Collecting data regarding the usage and impact of cyber deception in practice is a significant research challenge.
- **Moving Target Defense (MTD):**
    - Moving target defense is another proactive defense approach, which continuously changes the attack surface of a system (without widening it), thereby increasing the cost and complexity of attacks (e.g., address space layout randomization).
    - One of the key challenges that our group discussed was the lack of turn-key platforms and frameworks for MTD in certain domains. Some practitioners might not be able to apply a defense approach unless a turn-key solution is available

for it, which may limit the practical impact of MTD research. In the networking domain, SDN might be used (with some caveats); but for some other domains, there are no practical solutions available.
    - Another challenge that MTD faces is its high cost in terms of performance or usability (e.g., moving IP addresses or configurations might result in transient service interruptions).
    - Our group also discussed the distinction between MTD and cyber deception. While some may view MTD as a type of deception, our group agreed that there are important differences, which make MTD and deception separate approaches: MTD aims to deny information to the adversary by changing the system, while deception aims to change the adversary by selectively releasing false and true information. An important implication of this observation is that deception should be tailored to the adversary to maximize its effectiveness.
- **Frequency and port hopping as MTD examples:**
    - Our group first discussed frequency hopping as an example of MTD for wireless communication. To implement frequency hopping, the sender and receiver first need to agree on a hopping sequence (i.e., what frequencies to use and when to switch), which acts similarly to a cryptographic key. Then, the sender transmits over various channels, “hopping” to a new one from time to time according to the sequence. Frequency hopping can be used for secrecy and performance reasons.
    - Next, we discussed port-number hopping. In this case, the port number on which a service is available is changed from time to time (similar to frequency hopping). The service might show shallow data unless someone connects to the right port. Unfortunately, port-number hopping can be difficult to implement for multiple reasons: there might be a gap between the administrators of the boxes and the networks; and the networking and application people might not be open to collaboration.
- **Combining deception and MTD with other defense approaches:**
    - In principle, proper access control would make MTD unnecessary. In this sense, MTD becomes a way of augmenting access control. Similarly, deception can be viewed as augmenting authentication. However, it is not clear what the best way is for combining MTD or deception with other defense approaches, such as access control.
    - Our group also discussed the combination of MTD and deception, e.g., deploying honeypots and then continuously changing the configuration of honeypots and real assets. A key research challenge is providing models and methodology for combining deception and MTD optimally.
- **Quantifying the security benefit of MTD:**
    - The entropy of MTD can be thought of as similar to the strength of a cryptographic key. Can we prove that MTD (with certain entropy) amounts to some well-defined level of protection?
- **Adversarial deception:**

- In addition to defensive deception, our group also discussed deception that is employed by adversaries. Specifically, we discussed how an adversary can masquerade as a real person, e.g., pretending to be a trusted person in a spear-phishing attack.
- To succeed, the adversary must be able to generate e-mails that the victim will not be able to distinguish from e-mails that originate from the impersonated user. Currently, this is possible semi-automatically, as existing results show that users cannot differentiate between real and semi-automatically generated fake e-mails. Results also show that deep-learning based generation has coherence issues (repetitive/redundant). A key research challenge in this direction is measuring the coherence of the attacks.
- Our group also discussed the parallels between e-mail masquerading and text spinning (i.e., text generation) on websites used for search-engine optimization. We also discussed the parallels to chat bots, which currently work well in only limited domains, and to deepfake techniques.

3. Potential Approaches

- **Formal metrics and proofs for cyber deception**
  - For providing formal guarantees about the benefits of deception, our group discussed the idea of using randomized deception. In particular, we discussed the idea of assuming that the adversary knows that deception may be used as well as the deception techniques available (similar to Kerckhoffs' principle for cryptography); however, choosing the particular deployment of deception (i.e., what fake assets are deployed, and how fake and real assets are configured) at random, unknown to the adversary. In this case, the number of possible deployments and configurations is similar to the size of the key space in cryptography. Hence, the level of security provided by deception hinges on the "key size," i.e., logarithm of the number of possible deployments and configurations.
  - We noted that this key size will necessarily be much smaller than typical cryptographic key sizes due to practical constraints. However, even shorter key sizes can be sufficient since the adversary typically cannot perform an "offline" search over this space; instead, it needs to surveil the system to gain information. Relatively small-scale deception also appears to be beneficial from a practical perspective because the "right" message at the "right" time might be able to significantly alter the adversary's decision-making process.
- **Threat modeling:**
  - To tailor deception to adversaries, we need realistic threat models that can capture the adversaries' decision making process, including factors such as bounded rationality or risk aversion. Our group discussed how social and behavioral sciences can help with capturing bounded rationality and risk aversion.

- **Adopting existing techniques to MTD:**
  - To provide turn-key solution, our group discussed how some existing techniques (e.g., randomized load balancing) could be re-purposed for implementing MTD.
- **Measuring usability costs:**
  - Our group also discussed how social and behavioral sciences can help with measuring the usability costs of deception and MTD.

4. Long-Term (> 10 years) Significance

Both in research and in practice, existing security measures tend to focus on the prevention of intrusions or on reacting to them. Clearly, these approaches have their limitations. To manage cyber risks more effectively, we need to take a more proactive approach and augment existing security measures with emerging proactive techniques, such as cyber deception. Deception can play an important role in diminishing the adversary's inherent offensive advantage; however, the research challenges that we outlined above show that there remain fundamental research questions that we must address before deception can be applied methodologically in practice.