Security Modeling

Estimating Attack Resistance
Quantitative Risk Assessment

Aleksandr Lenin

Motivation

- Are we secure? (now, at this very moment)
- Can the considered infrastructure resist targeted attacks?
- Will it break? If so - under which conditions, when, how?
- There exist no scientifically justified metrics of strength against attacks
- Such metrics exist in other areas - i.e. in civil engineering

We need a simple and reliable method similar to the one used in civil engineering

Are we secure?

We need to determine conditions under which attacking the considered infrastructure is beneficial.

If we could easily determine if attacking is beneficial or not, we would get a simple tool to assess whether the organization is secure against rational attacks, or insecure.

Security Modeling

There are two ways to study a phenomena:

- experimenting with an observable, controllable, measurable object, collecting empirical evidence
- modeling and/or simulations

When conducting experiments is infeasible or economically impractical, modeling is used.

Security is not observable, and we have no evidence that it is even measurable – we do security modeling instead.

Can we come up with a model which could assess if the considered infrastructure is secure?

Security Modeling

- model of the world
- definition of security in the model
- computational methods that verify security
- falsification of model and computational methods
  - Falsifiable given observable results, when a security incident happens
  - Best effort otherwise: assume the model is correct, unless it is falsified

Model of the world

- Threat model
  - threats as events happen with certain probability
  - some of these events result in damage
- Attacker model
  - attacks – treats happening in result of conscious decision-making process
  - how to obtain likelihood for an attack?
  - need to model decision-making process of humans
  - need an attacker behavioral model – i.e. Rational Choice Theory
- Definition of security
  - threats cannot happen
  - the expected damage is minimal
  - the sum of the expected damage and expenses is minimized

The goal of security is not to protect everything from everything else!

Computational methods

- Validate security w.r.t the definition of security in the model
- Yes–no result
- It is safe to under-estimate security
- It is very dangerous to over-estimate security
- Need to avoid false-positive results
- Reliability of analysis results

Rational Choice Theory

Rational Choice Theory is based on a set of assumptions:

- Agents can make preferences over the set of possible alternatives or actions
- The decision-making process is driven by a particular goal
- Agent preferences are consistent across time
- Agent preferences are self-interested – agents undertake actions that maximize personal advantage

Rational Choice Theory

- Irrational behavior (impulsive, stochastic, inconsistent across time)
- Irrational behavior is driven by emotions, beliefs, ideas, ...
- Rational behavior (deterministic)

Rational Choice Theory

- Rational choice requires
  - a goal specification
  - a set of alternatives
- Rational attackers have two alternatives
  - to attack
  - not to attack

An attacker chooses to attack if it is beneficial for him. Therefore, the probability of an occurrence of an attack is:

1 if it is beneficial to attack

0 if it is not beneficial to attack

This decision resembles cost-benefit analysis in project planning.

Economical Feasibility of Attacking

- Preparing and launching an attack requires resource investment
- Successful attacks being some revenue
- Profit is fixed and cannot be influenced by an attacker
- An attacker may minimize expenses (and thus maximize profit) by chosing economically cheapest ways to achieve the goal
- Consistent choice of this self-determined ”best” action
- Utility – the metric to decide the feasibility of attacking
- Utility is the difference between the profit and expenses

Adversarial Model

- Rational profit-oriented malicious actors
- ...driven by monetary profit
- ...launching targeted attacks against the considered organization
- The profit is known to them prior to attacking
- Typically, the reconnaissance phase precedes the infiltration phase of attack (for targeted attacks)
- An attacker collects information about the target organization, possible ways to attack it
- When all the relevant knowledge has been collected, an attacker needs to decide:
  - Is it worth attacking, or
  - an attacker would be better off not even trying to attack the considered organization

Threat Model

An attacker needs to take into account:

- expenses of attack alternatives
- probabilities of success
- profit
- risk appetite / risk aversion

to decide if it is beneficial (economically justified) to attack or not.

Attack Tree Threat Modeling

Let us consider a threat of the loss of market share due to an intellectual property theft.

Way too abstract formulation

Impossible to make informed decision whether it is beneficial to attack or not

The only known variable is the profit

In order to estimate the expenses and the success probability, a structured description of the attack is required.

Attack Tree Threat Modeling

Attack trees:

- Evolved from reliability analysis (fault trees)
- Fault trees were extensively used by NASA in 1960-s
- Advanced theory and mathematical apparatus behind fault tree analysis
- Were adopted to the security domain by Schneier
- Evolving area of security analysis based on attack trees

Attack Tree

- A hierarchical description of attack steps
- Represents alternatives
- Does not represent ordering of attack steps
- Attack strategy sets the ordering in which attack steps are launched
- Conjunctive and disjunctive refinements
- A monotone Boolean function of its inputs (elementary attack steps)
- Attack suite – a subset of the set of attack steps
- Satisfying attack suite – an attack suite which satisfies the Boolean function

A Bit of History

2005 – Foundations of Attack Trees (Mauw, Oostdijk)

- Formalized attack trees and their semantics
- Formalized propositional Boolean semantics
- Introduced multiset semantics
- Formalized attribute domains and bottom-up propagation rules
- Showed that some of the existing models were inconsistent with their underlying semantics
- Introduced semantical indistinguishability criteria – valid computational procedure must produce the same result for semantically indistinguishable models

A Bit of History

2006 – Rational Choice of Security Measures via Multi-Parameter Attack Trees (Buldas, Laud, Priisalu, Saarepera, Willemson)

- Multiparameter model for attack tree analysis
- Introduced economic reasoning into attack tree analysis
- Game–theoretic treatment of AT analysis
  - Separated attacker model from the threat model
  - Considered different adversarial strategies to achieve the goal (as a game)
  - Considered rational targeted profit-oriented attacks
  - Introduced economical reasoning into attack tree analysis
  - Introduced upper bounds ideology

A Bit of History

Following this ideology, several models appeared:

- Multiparameter model (Buldas *et al.*, 2006)
- Parallel model (Jürgenson, Willemson, 2008)
- Serial model (Jürgenson, Willemson, 2010)
- Fully–adative model (Buldas, Stepanenko, 2011)
- Infinite repetition model (Buldas, Stepanenko, 2011)

A Bit of History

2010 – Foundations of Attack–Defense Trees (Kordy, Mauw, Radomirovic, Schweirzer)

- Every attack node in a tree can be countered by the corresponding defense sub-tree
- Defense trees are actions of the security team
- Showed that Attack–Defense Tree (ADT) and a two–player zero-sum game (in game theory) are equivalent
- Optimal strategies yield equilibrium

A Bit of History

2012 – Failure–Free model (Buldas, Lenin)
2013 – Improved Failure-Free model (Buldas, Lenin)

- Considered fully adaptive adversarial strategies
- In general, finding an optimal strategy in security games belongs to PSPACE.
- Showed that in failure-free games optimal strategies always exist.
- Showed that finding an optimal strategy in the failure–free model is equivalent to solving a Weighted Monotone Satisfiability (WMSAT) problem.
- Showed that WMSAT problem is NP–complete.
- Justified the propagation rules and showed that they do not violate the upper bounds ideology

A Bit of History

2014 – Socio-technical security metrics (Böhme, Van Eeten, Foley, Hadžiosmanović, Lenin, Pape, Pieters)

Being NP–complete is good and bad.

Good in a game–theoretical sense. Other approaches are NP–hard!

Bad in practice – search for the exact result can take an unacceptable amount of time, even considering a moderate–size attack tree.

2015 – Genetic Approximations for the Failure–Free Security Games (Lenin,Willemson,Charnamord)

A Bit of History

2014 – Attacker Profiling (Lenin, Willemson, Permata-Sari)
2014 – Attacker Profiles for Security Risk Analysis (Pieters, Hadžiosmanović, Lenin, Montoya, Willemson)

- Considered different specifications of adversarial profiles in attack tree analysis
- Trivial idea: prune attack-tree branches that do no fit into the attacker profile
- Better idea: employ item–response theory to assist in attacker profiling calculations and update quantitative metrics in attack trees w.r.t. considered attacker profile
- Tried to capture various factors contributing to adversarial time and likelihood by means of item response theory

Attacker Profiling

The more skillful and experienced the attacker is, the more resources are available to the attacker – the more likely he is to succeed in the considered attack.

Similar reasoning may be applied to the skill parameter – the more skillful and experienced the attacker is the less difficult is the attack process for him, the less time it will take to succeed in an attack.

Less skilled attackers given sufficient amount of time may be as efficient in terms of the success likelihood as more skilled attackers who have been given less time for executing the same attack.

Similar logic may be extended to other quantitative annotations as well.

Ongoing Research

2017 – Simple infeasibility certificates for attack trees (Buldas, Lenin, Willemson, Charnamord)

- How can we certify that a WMSAT instance is unsolvable?
- Is there a way to generate efficient certificates of unsolvability? (coNP task)
- Introduced infeasibility certificates for WMSAT problem.
- Ongoing research – positivstellensatz certificates

Ongoing Research

2019 – Attribute Evaluation on Attack Trees with Incomplete Information (Buldas, Gadyatskaja, Lenin, Mauw, Trujillo-Rasua)

- Can we assist analysts in estimating quantitative annotations in attack trees by using the available statistical data from the past?
- How to handle contradictions between the experts' gut feeling and statistical data?
- To what extent past data is valid today? How its quality influences the analysis?
- A new formalism for attack trees – a set of hard (structural) and soft (constraints) predicates.

Attack Tree

- Find the minimal cost of executing an attack scenario
- ...by finding the cheapest path through an attack graph
- A complex unconstrained combinatorial optimization task
  - WMSAT task is NP-complete
  - typically belongs to PSPACE
- Can we exploit the structure?
- Efficient propagation rules – attribute domain

Attack Tree Propagation Rules

An attribute domain defines the propagation rules, and is a triplet $(\mathbb{D}, \wedge, \vee)$, where

- $\mathbb{D}$ is the domain, e.g. $\mathbb{R}_{\geqslant 0}$ for cost
- $\wedge, \vee$ are operators applied in conjunctive and disjunctive nodes

The `minCost` attribute domain is a triplet $(\mathbb{R}_{\geqslant 0}, +, \min)$

Intuition:

- Cost is non-negative, therefore the domain is $\mathbb{R}_{\geqslant 0}$
- In conjunctive nodes, the costs are summed up, hence operator $(+)$ for $\wedge$ part of the attribute domain
- Disjunctive nodes represent a choice – hence min operator corresponds to the choice of the cheapest attack

Attack Tree Propagation Rules

It is safe to underestimate costs. Overestimating costs may produce false-positive results – very dangerous.

In order to avoid false-positives, we need to make sure that the propagation rules do not overestimate costs.

What we wish to achieve can be expressed as

$$\mathcal{E}(x_1 \wedge x_2 \wedge \ldots \wedge x_k) \leqslant \mathcal{E}(x_1) + \ldots + \mathcal{E}(x_k) \ ,$$
$$\mathcal{E}(x_1 \vee x_2 \vee \ldots \vee x_k) \leqslant \min\Big\{\mathcal{E}(x_1), \mathcal{E}(x_2), \ldots, \mathcal{E}(x_k)\Big\} \ .$$

What we actually have is

$$\mathcal{E}(x_1 \wedge x_2 \wedge \ldots \wedge x_k) \leqslant \mathcal{E}(x_1) + \ldots + \mathcal{E}(x_k) \ ,$$
$$\mathcal{E}(x_1 \vee x_2 \vee \ldots \vee x_k) = \min\Big\{\mathcal{E}(x_1), \mathcal{E}(x_2), \ldots, \mathcal{E}(x_k)\Big\} \ .$$

These propagation rules do not produce false-positive results, and are therefore reliable.

Attack Strategy

- An attack strategy is a rule, which in every state either suggests the next elementary attack to launch, or to give up trying.
- Every attack step trial may succeed or fail with some probability
- In real life attackers can re-run failed attacks again an arbitrary number of times
- Exponential number of potential sequences of attack steps – a computationally infeasible task
- Are there any simplifications we could potentially make?

Attack Costs Revisited

- Let the attacker be overpowered
- Assume an attacker who has infinite resources
- Assume an attacker who can re-run failed attacks again infinitely until they eventually succeed

This reasoning follows the ideology of avoiding false-positives.

Attack Costs Revisited

Intuition:
By considering overpowered adversaries we are guaranteed not to underestimate them. If the system is shown to be secure against such overpowered adversaries, this implies that it is protected against real-life less powerful attacks.

Indeed, we over-secure our systems and make extra (maybe unnecessary) investments into security, but it is better to be safe than sorry.

Attack Costs Revisited

If $\mathcal{C}$ is the cost of a single attack step execution, then the cost of an infinite series of trials is

$$\mathcal{E}(x) = px + (1-p)x + (1-p)^2x + \ldots + (1-p)^n x = \frac{\mathcal{C}}{p} \ ,$$

where p is the probability of success for a single attack strep trial.

Use value of adjusted costs $\mathcal{E} = \frac{\mathcal{C}}{p}$ as the input value for `minCost` calculations.

Similar to cost–success ratio used in cryptography – intuitive and well understood.

Attack Feasibility Analysis

- Computational methods use bound-oriented approach by calculating the lower bound of adversarial expenses
- Comparing it to the adversarial profit allows to make informed decisions w.r.t feasibility of attacking
- This, in turn, allows to make informed decisions about whether the infrastructure is secure against rational targeted profit-oriented attacks.
- The methodology is easy to use, does not require expert knowledge.
- Simple and reliable approach to security engineering.