Attack-Defense Tree Based Security Analysis and Optimal Defense Synthesis for System Design

Abstract

Attack-Defense Trees (ADTrees) are widely used in the security analysis of software systems. In this paper, we introduce a novel approach to analyze system architecture models via ADTrees and to synthesize an optimal cost defense solution using MaxSMT. We generate an ADTree from the system architecture model with its possible attacks, possible and implemented defenses. We analyze these ADTrees to see if they satisfy their cyber-requirements. We then translate the ADTree into a set of logical formulas, that encapsulate both the logical structure of the tree, and the constraints on the cost of implementing the corresponding defenses, such that a minimization query to the MaxSMT solver returns a set of defenses that mitigate all possible attacks with minimal cost.

1 Introduction

Software security has attracted worldwide attention as society has grown increasingly dependent on computer-based systems. To address the security concerns of systems, many risk analysis techniques have been introduced over the years in order to identify potential system failure and mitigate risks before the system is fielded. *Attack trees* [10] are a prominent methodology to visually depict the security vulnerabilities of a system. They have been used in the analysis of threats against systems in the fields of defense and aerospace. Attack trees capture attacks in a tree structure, where the root node represents the attacker's goal and child nodes refine the goal with details involved in achieving the goal. *Attack-defense trees (ADTrees)* [8] extend attack trees with the notion of defenses against attacks, with the objective of reducing the consequences of attacks. In an ADTree, defense and attack nodes are distinguished node types, and in addition to refinements of nodes via children of the same type of node, child nodes can be *counter-measures* of the opposite kind of parent node. Such trees are able to capture both the attacker and the defenses of a system in an adversarial model, and as such, can be used to analyze the sufficiency of attack mitigation techniques of the system.

Implementing different defenses requires various amount of effort, time, and money. A challenging problem is to select a set of defenses that is able to mitigate all threats while incurring minimal cost of implementation. In this work, we use MaxSMT solvers to synthesize a set of optimal-cost defenses that mitigate all possible attacks on a system. A prototype was implemented in a tool called VERDICT — in conjunction with the model-based architectural analysis (MBAA) component, model-based architectural synthesis (MBAS) calculates a set of defenses for all known attacks at minimal cost. We make the following contributions in this paper.

- We describe an algorithm that converts an AADL system architecture model to an attack-defense tree.
- We describe an evaluation of these ADTrees in terms of a set of cyber-requirements.
- We encode the ADTree along with the costs of implementing defenses as a MaxSMT problem so that the solver can find a least-cost defense solution that satisfies all requirements.
- We present the analysis and synthesis features in the VERDICT toolchain which provides an implementation of the above functionalities.

Section 2 presents our specifications of attacks, defenses, and attack-defense trees. Section 3 presents a translation of AADL models to ADTrees and our method of analyzing whether an ADTree satisfies its requirements. Section 4 describes the interaction with the SMT solver in determining a minimum-cost set of defenses for the system. Section 5 is a presentation of our implementation in the VERDICT toolchain along with an evaluation on a model of a delivery drone system. Section 6 discusses some related work and Section 7 discusses directions our work can move in the future along with a summary.

Severity Level	Acceptable Level of Risk	DAL	Score	Objective (W/Independence)
Catastrophic	1×10^{-9}	A	9	66 (25)
Hazardous	1×10^{-7}	B	7	65 (14)
Major	1×10^{-5}	C	5	57 (2)
Minor	1×10^{-3}	D	3	28 (2)
No Effect	1	E	0	0 (0)

Table 1: Mapping between the severity of consequence, acceptable level of risk, design assurance level (DAL), DAL score (Score) and development objectives (with independence).

Appendix A argues the soundness of our interaction with the SMT solver.

2 Preliminaries

In this section, we formalize our problem and solution space. We consider attacks from the MITRE Common Attack Pattern Enumeration and Classification (CAPEC) library that targets embedded systems and defenses (controls) from the National Institute of Standards and Technology's (NIST's) 800-53 security standard. A toy drone example is used through the paper to illustrate various features.

2.1 Attack and Defense Specification

This work is based on two standards drafted by the Radio Technical Commission for Aeronautics (RTCA) — DO-326A, the Airworthiness Security Process Specification and DO-356A, the Airworthiness Security Methods and Considerations — both providing guidance against threats to aircraft systems. The standards specify the acceptable level of risk corresponding to the level of severity of successful attacks. The severity of successful attacks is categorized into 5 levels based on their effects on the aircraft, crew, and passengers: Catastrophic, Hazardous, Major, Minor, and No Effect. These levels, along with the corresponding levels of risk acceptable in the system, are presented in the first two columns of Table 1. We represent by L the set of severity levels, and by ρ the set of acceptable risk levels. The top-level event of an ADTree represents the attacker's goal, which is measured in terms of confidentiality (`C`), integrity (`I`) and availability (`A`) of the outports of the system. The attacker's goal is to sabotage the system by compromising its components. Attacks on components are ultimately propagated to the outports of the system through internal connections. The system fails if the CIAs of its outports are compromised. To mitigate attacks, the system designer has to implement defenses in components with various degrees of rigor, which can prevent failure of the system. The previously mentioned standards map the rigor of defense implementation, called *design assurance levels (DALs)*, to a security consideration score or DAL score — columns 3 and 4 of Table 1. DAL A is the highest rigor defense and E is the lowest. DALs originated in DO-178 and were reused in DO-254. These standards were developed to ensure that software and complex hardware were developed with enough rigor and could be proved to be absent of bugs with potentially severe consequences. To bring the system within an acceptable risk level of attacks, and to prevent the system from failing with the associated severity level, its developers need to implement the component to the respective DAL score and meet appropriate development objectives from the fifth column. For example, if the failure of software can have a Catastrophic consequence, one must show compliance to 66 objectives as part of the software development process, 25 of which need to be performed by independent developers. The implementation of defenses incurs efforts and cost, which increase with the DAL score, and the goal of this work is to synthesize a set of defenses that mitigate all attacks at an optimal (minimal) cost.

We will use *CIA* to denote the set consisting of the properties confidentiality (`C`), integrity (`I`) and availability (`A`) ($CIA = \{\texttt{C}, \texttt{I}, \texttt{A}\}$), and δ, the set of possible DAL scores ($\delta = \{0, 3, 5, 7, 9\}$).

2.2 Attack-Defense Trees

ADTrees are rooted, labeled, finite trees that represent scenarios of security attacks against a system, and the countermeasures taken against these attacks. The nodes of an ADTree are either *attack nodes* — represented as red circles — or *defense nodes* — represented as green rectangles, and the nodes are labeled either with attack or defense goals, or with logical gates that connect these goals. A node's children represent either *refinements* (represented by straight line edges) of the same node type or *countermeasures* (dotted line edges) of the opposite node type. Refinements can either be conjunctive, (denoted in diagrams with an arc below the parent node) in which case, all the refinements' goals must be achieved for the parent's goal to be achieved; or disjunctive, where at least one of the refined goals must be achieved for the parent's goal to be achieved. Non-refined nodes (leaves) are called *basic actions*. The root of an ADTree represents the attacker's goal, which can be refined down to a logical formula over the basic actions (leaves) by expanding on the refinements performed by nodes (conjunctions and disjunctions).

Example 2.1. Consider the following model that abstracts a simple drone. The model is represented using an architecture diagram in Figure 1. A remote control allows the user to direct the drone. The drone consists of a controller that implements its logic, and a propeller that helps the drone move. As a security measure, the drone consists of a backup controller which

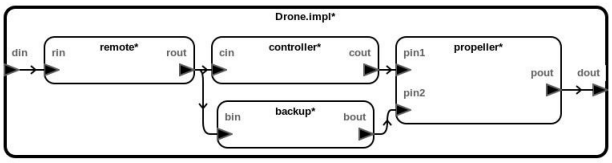

Figure 1: The component diagram for the toy drone model

implements a much simpler logic than the main controller. The user of the remote may invoke the single functionality of the backup which brings the drone back to the location of the remote.

A wireless connection connects the remote to the controller and to the backup, both of which have a wired connection to the propeller. Figure 2a shows the ADTree that models the attacks and defenses of the drone system, supposing that we care about the integrity of the drone's propeller, that is, we want the propeller to move as instructed by the remote, and return back safely to the owner if that isn't feasible. Consider the following attacks:

1. Physical Theft Attack (CAPEC–507) on the remote.

2. A combination of a Software Integrity Attack (CAPEC–184) on the controller, and an Identity Spoofing Attack (CAPEC–151) on the backup controller.

Either Physical Access Control or System Access Control of the remote can defend against Physical Theft, but CAPEC–390 (Bypassing Physical Security) is a *dependent attack* that becomes applicable once Physical Access Control is implemented, and can only be defended against by implementing System Access Control. Three defenses — Remote Attestation, Memory Protection, and Secure Boot — are necessary for the controller to mitigate CAPEC-184 and Heterogeneity alone implemented on the backup controller can protect it against the identity spoofing attack.

In Figure 2b, we label the nodes, attacks, and defenses, and also use the notation from our ADTree definition (Definition 2.1). We also give label *R* for the remote sub-system, *C* for the controller, and *B* for the backup controller. All defenses are implemented to DAL-score 5. □

We use an inductive definition of ADTree from [9].

Definition 2.1. An *ADTree T* is generated by the following

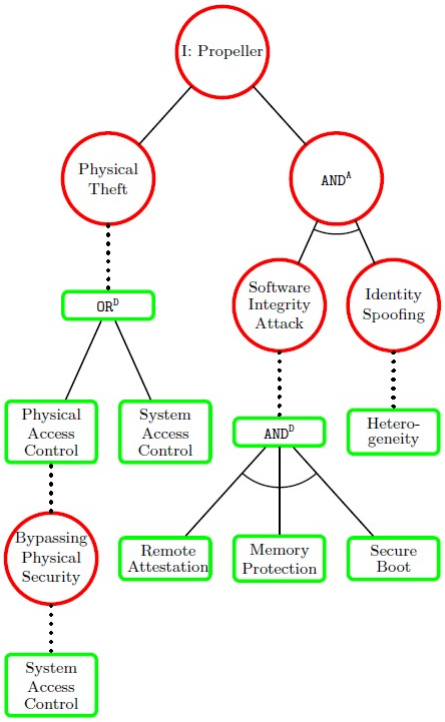

(a) ADTree of drone system

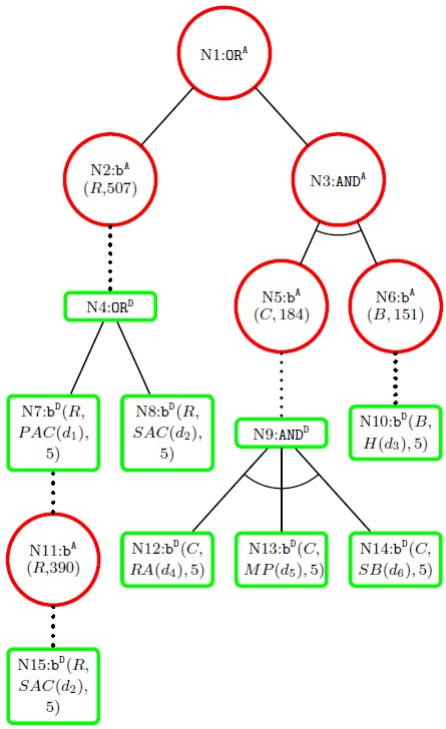

(b) ADTree of drone system - labeled

Figure 2: Example drone system

grammar.

$$
\begin{aligned}
&T \to T^{\mathtt{A}} \mid T^{\mathtt{D}} \\
&T^{\mathtt{A}} \to \mathtt{b}^{\mathtt{A}}(s,a) \mid \mathtt{OR}^{\mathtt{A}}(T^{\mathtt{A}},\ldots T^{\mathtt{A}}) \mid \mathtt{AND}^{\mathtt{A}}(T^{\mathtt{A}},\ldots T^{\mathtt{A}}) \\
&\quad \mid \mathtt{C}^{\mathtt{A}}(T^{\mathtt{A}},T^{\mathtt{D}}) \\
&T^{\mathtt{D}} \to \mathtt{b}^{\mathtt{D}}(s,d,\delta) \mid \mathtt{OR}^{\mathtt{D}}(T^{\mathtt{D}},\ldots T^{\mathtt{D}}) \mid \mathtt{AND}^{\mathtt{D}}(T^{\mathtt{D}},\ldots T^{\mathtt{D}}) \\
&\quad \mid \mathtt{C}^{\mathtt{D}}(T^{\mathtt{D}},T^{\mathtt{A}})
\end{aligned}
$$

Superscripts `A` and `D` represent attack and defense entities respectively. T represents terms or trees, `OR` encapsulates disjunctive refinements of a node, `AND` represents conjunctive refinements of a node, and `C` encapsulates the action of a node and its countermeasure. `b` represents basic actions — for attack nodes, they are parameterized by a component and an attack, and for defense nodes, they are parameterized by a component, a defense, and implemented DAL-score. We define attack trees and defense trees as follows. An attack tree, denoted $T^{\mathtt{A}}$, is an ADTree with root of type `A` and a defense tree, denoted $T^{\mathtt{D}}$, is a tree with root of type `D`. We define a function `root` that returns the root node of an ADTree. □

Defenses are implemented to a particular DAL-score, and the defense nodes (that are basic actions) are parameterized by this DAL-score, along with the component that they defend and the defense itself. DAL-scores can only take values from column 4 of Table 1.

Although our definition allows for any kind of ADTree, in practice, we consider the root node of our ADTrees to always represents an attack. This suits our goal of using ADTrees to analyze the attacker's actions.

An interesting feature of our ADTrees is that we allow repetition of defense and attack nodes. That is, multiple $\mathtt{b}^{\mathtt{A}}$ and $\mathtt{b}^{\mathtt{D}}$ nodes in our trees can have the same label ((s,a) or (s,d,δ)). The only restriction we place is that when two $\mathtt{b}^{\mathtt{A}}$ or $\mathtt{b}^{\mathtt{D}}$ nodes have the same label, their child-node structure must be identical.

3 ADTree Analysis

In this section, we describe how our tool uses ADTrees to analyze a system architecture modeled using AADL (Architecture Analysis and Design Language) [5], which provides a framework and language for early analyses of a system's architecture with respect to performance-critical properties. Our tool builds an ADTree from an AADL model, a specification of possible attacks, possible defenses, implemented defenses, and cyber-requirements to satisfy. This tree is evaluated in terms of the likelihood of success of an arbitrary attacker, given a set of defenses.

3.1 Defense Models

Within VERDICT, the analysis of the AADL model receives information primarily from the Security Threat Evaluation and Mitigation (STEM) component. STEM identifies possible attacks, possible defenses and defenses implemented in the components of the system. STEM provides this data in the form of a defense model $\mathbb{M}$ with two types of relations: an implemented defense model $\mathbb{M}_{\mathbb{I}}$ and an applicable defense model $\mathbb{M}_{\mathbb{A}}$.

A *defense model* $\mathbb{M}$ is a relation containing tuples that relate components of a system to defense–DAL-score pairs, and attack–CIA pairs. Each tuple signifies possibly the applicability of an attack to a component, and either the applicability or implementation of a set of defenses to the same component to respective DAL-scores. We distinguish 3 types of defense models, and define 2 of them as follows. The third, the synthesized defense model, is defined later.

1. An *implemented defense model* $\mathbb{M}_{\mathbb{I}}$ represents defenses currently implemented in the system.

$$
\begin{aligned}
(s,a,\gamma,\Delta) \in \mathbb{M}_{\mathbb{I}} \text{ iff in component } s,\ &\gamma \text{ attack } a \\
&\text{is applicable, and for each} \\
&(d,\delta) \in \Delta, \text{defense } d \text{ is} \\
&\text{implemented to DAL-score } \delta
\end{aligned}
$$

   where $s \in S, a \in A, \gamma \in CIA, d \in D, \delta \in DAL$.
2. An *applicable defense model* $\mathbb{M}_{\mathbb{A}}$ represents defenses applicable in the system.

$$
\begin{aligned}
(s,a,\gamma,\Delta) \in \mathbb{M}_{\mathbb{A}} \text{ iff in component } s,\ &\gamma \text{ attack } a \\
&\text{is applicable, and for each} \\
&(d,\delta) \in \Delta, \text{defense } d \text{ is} \\
&\text{applicable to DAL-score } \delta
\end{aligned}
$$

In the defense models that it provides, STEM guarantees that basic action nodes with the same labels have the same subtrees, as required by our mechanism in section 2.2.

Example 3.1. Consider the ADTree from Example 2.1 in terms of the following cyber requirement:

$$
\begin{aligned}
q:\ (d_{out} : \mathtt{I}) \text{ with Major } &(1 \times 10^{-5}) \\
&\text{severity level}
\end{aligned}
$$

While the tree from Figure 2 models the applicable defense model, we show two different implementations of defenses in Figure 3.

The applicable defense model $\mathbb{M}_{\mathbb{A}}$ (Figure 2) consists of

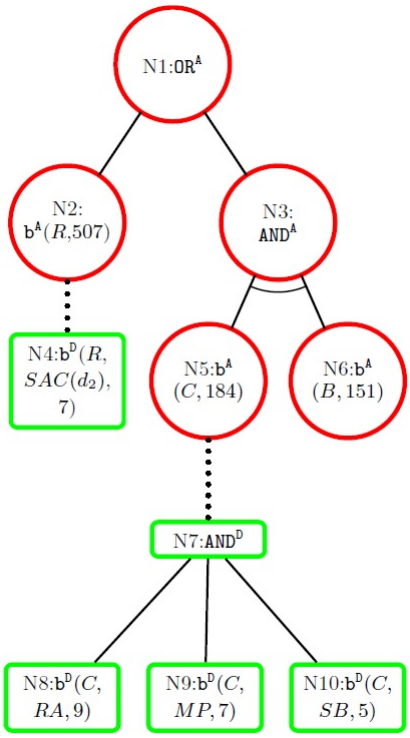

(a) ADTree representing $\mathbb{M}_{\mathbb{I}}$

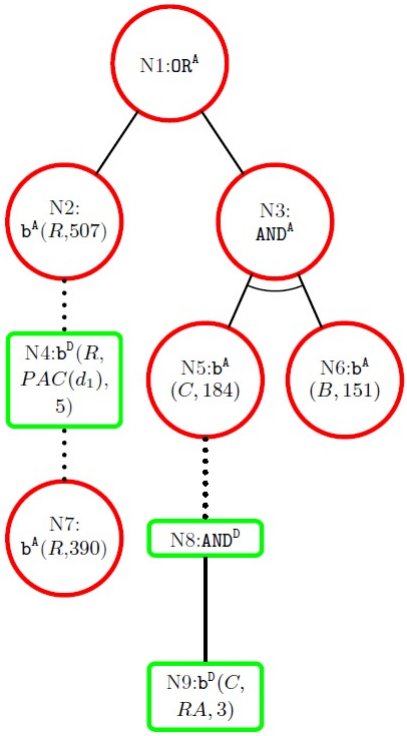

(b) ADTree representing $\mathbb{M}'_{\mathbb{I}}$

Figure 3: ADTree representing defense implementations of the drone system

the following tuples

$$
\begin{aligned}
&(\mathit{Remote},\ \mathit{CAPEC}\text{–}507,\ I,\\
&\{(d_1,5)\}) && \in \mathbb{M}_{\mathbb{A}}\\
&(\mathit{Remote},\ \mathit{CAPEC}\text{–}507,\ I,\\
&\{(d_2,5)\}) && \in \mathbb{M}_{\mathbb{A}}\\
&(\mathit{Backup},\ \mathit{CAPEC}\text{–}151,\ I,\\
&\{(d_3,5)\}) && \in \mathbb{M}_{\mathbb{A}}\\
&(\mathit{Controller},\ \mathit{CAPEC}\text{–}184,\ I,\\
&\{(d_4,5),(d_5,5),(d_6,5)\}) && \in \mathbb{M}_{\mathbb{A}}
\end{aligned}
$$

Figure 3a shows the ADTree from implementation $\mathbb{M}_{\mathbb{I}}$.

$$
\begin{aligned}
&(\mathit{Remote},\ \mathit{CAPEC}\text{–}507,\ I,\\
&\{(d_2,7)\}) && \in \mathbb{M}_{\mathbb{I}}\\
&(\mathit{Controller},\ \mathit{CAPEC}\text{–}184,\ I,\\
&\{(d_4,9),(d_5,7),(d_6,5)\}) && \in \mathbb{M}_{\mathbb{I}}
\end{aligned}
$$

3.2 ADTree Construction

The ADTree construction algorithm `ADTree` operates on the following parameters:

1. *mod*, the AADL model of a system, that specifies ports P, connections C, components S, and cyber-relations R between ports of the system, where cyber relations are internal to a component and specify how the CIA vulnerabilities of inports propagate to the CIA vulnerabilities of outports.
2. Q, the set of cyber requirements, where each cyber-requirement q is a logical formula over (p,γ) atoms, $p \in P, \gamma \in \mathit{CIA}$, with a corresponding level of severity $l \in L$.
3. $\mathbb{M}$, a defense model.

and returns the following output:

- ADTree T corresponding to the requirements in Q on the AADL model *mod* considering attacks and defenses in $\mathbb{M}$

The mutually recursive functions `ADTree_port` and `ADTree_const` from Figure 4 assist in constructing an ADTree. The notation `loc` $\mapsto$ `node` indicates that the node `node` is created at the location `loc` in the tree. An ADtree for a single (p,γ) atom from requirement q is constructed through the call `ADTree_port`$(p,\gamma,\mathbb{M},$ `root`$(\mathit{new\_tree}))$. The function recursively scans through the ports within the architecture of the system, modeled in *mod*, while constructing the ADTree. It calls `ADTree_const` when it encounters a

constituent which is either a component, a connection, or a cyber-relation. `ADTree_const` calls `DTree` from Figure 5 which constructs a defense tree from a subset of the defense model under consideration. Once an ADTree is constructed, `CrushTree`, also from Figure 5, removes redundant nodes from the tree.

`ADTree`(mod, Q, $\mathbb{M}$) constructs an ADTree as follows:

1. For each requirement, q, for each $(p, \gamma) \in q$, call `ADTree_port`(p, γ, $\mathbb{M}$, `root`(*new_tree*)).
2. Combine each tree from the previous step using logical attack nodes $\texttt{AND}^{\texttt{A}}$ and $\texttt{OR}^{\texttt{A}}$ to create the ADTree corresponding to q.
3. Combined the trees from each requirement using an $\texttt{OR}^{\texttt{A}}$ node, and set the level of severity of the root of the tree as the maximum of each of the trees combined from the individual requirements.
4. Call `CrushTree` on the resultant tree.

Notice that our ADTree construction algorithm constructs ADTrees without $\texttt{C}^{\texttt{D}}$ nodes. This is a restriction of our implementation. However, we consider $\texttt{C}^{\texttt{D}}$ nodes in the rest of the paper, for completeness of our approach in evaluating and encoding ADTrees as MaxSMT queries for synthesizing optimal cost solutions.

Example 3.2. Figure 1 shows the architecture diagram for the drone model, which specifies most of the AADL model of the drone system. They are completely specified as follows.

- Ports P: $\texttt{d}_{\texttt{in}}$ and $\texttt{d}_{\texttt{out}}$, $\texttt{r}_{\texttt{in}}$ and $\texttt{r}_{\texttt{out}}$, $\texttt{c}_{\texttt{in}}$ and $\texttt{c}_{\texttt{out}}$, $\texttt{b}_{\texttt{in}}$ and $\texttt{b}_{\texttt{out}}$,and $\texttt{p}_{\texttt{in1}}$, $\texttt{p}_{\texttt{in2}}$ and $\texttt{p}_{\texttt{out}}$, — the inports and outports of the drone system, the remote sub-system, the controller, the backup controller, and the propeller, respectively.
- Connections C: `c1` between `din` and `rin`, `c2` between `rout` and `cin`, `c3` between `rout` and `bin`, `c4` between `cout` and `pin1`, `c5` between `bout` and `pin2`, and `c6` between `pout` and `dout`.
- Components S: `remote`, `controller`, `backup` and `propeller`.
- Cyber-relations R that connect the Integrity of the outport of each sub-system to the Integrity of all in-ports: $\texttt{p}_{\texttt{in1}} : I \wedge \texttt{p}_{\texttt{in2}} : I \rightarrow \texttt{p}_{\texttt{out}} : I$, $\texttt{r}_{\texttt{in}} : I \rightarrow \texttt{r}_{\texttt{out}} : I$, $\texttt{c}_{\texttt{in}} : I \rightarrow \texttt{c}_{\texttt{out}} : I$.

Using the following as inputs,

1. the specification of the model from above,
2. the cyber-requirements Q consisting of the single requirement $q : (\texttt{d}_{\texttt{out}} : \texttt{I})$ with Major severity level
3. either the applicable defense model from Example 2.1 or one of the implemented defense models from Example 3.1

Input: $p \in P$, $\gamma \in CIA$, $\mathbb{M}$, Tree Location `loc`
Output: ADTree T
Algorithm **`ADTree_port`**(p, γ, $\mathbb{M}$, `loc`)
  `loc` $\mapsto \texttt{OR}^{\texttt{A}}$ (`node*`)
  for all incoming constituents *const* to p **do**
    Add child `ADTree_const`(*const*, γ, $\mathbb{M}$, `root`(*new_tree*)) to `node*`
  end for

Input: $const \in S \cup C \cup R$, $\gamma \in CIA$, $\mathbb{M}$, Tree Location `loc`
Output: ADTree T
Algorithm **`ADTree_const`**(*const*, γ, $\mathbb{M}$, `loc`)
  if *const* is a component or a connection with inport p_{in} and outport p_{out} **then**
    `loc` $\mapsto \texttt{OR}^{\texttt{A}}$ (`node*`)
    for all $a \in A \mid (const, a, \_, \_) \in \mathbb{M}$ **do**
      for all $(const, a, \gamma,) \in \mathbb{M}$ **do**
        $(const, a, \gamma, \_) \in D_a$
        `D` $\leftarrow$ `CrushTree` (`DTree`(D_a))
        if `D` is empty **then**
          Add $\texttt{b}^{\texttt{A}}(const, a)$ as child of `node*`
        else
          Add $\texttt{C}^{\texttt{A}}(\texttt{b}^{\texttt{A}}(const, a), \texttt{D})$ as child of `node*`
        end if
      end for
      Add an additional child to `node*` and set it as `loc`
      `ADTree_port`(p_{in}, γ, $\mathbb{M}$, `loc`)
    end for
  else if *const* is a relation $F \rightarrow p_{out} : \gamma$ **then**
    `loc` $\mapsto \texttt{OR}^{\texttt{A}}$ (`node*`)
    for all atom $p_{in} : \gamma \in F$ **do**
      `ADTree_port`(p_{in}, γ, $\mathbb{M}$, `root`(*new_tree*))
    end for
    Connect all the new ADTrees using attack nodes that correspond to the logical connectives in F
  end if

Figure 4: Mutually recursive functions `ADTree_port` and `ADTree_const` are used to construct an ADTree

Input: Set of tuples `D`
Output: Defense tree $T^{\mathtt{D}}$
Algorithm **`DTree`**(D)
  Create an $\mathtt{OR}^{\mathtt{D}}$ (`node*`)
  for all $(const, a, \gamma, \Delta) \in D$ **do**
    Create an $\mathtt{AND}^{\mathtt{D}}$ node (`node#`)
    for all $(d, \delta) \in \Delta$ **do**
      Add $\mathtt{b}^{\mathtt{D}}(const, d, \delta)$ as a child of `node#`
    end for
    Add `node#` as child of `node*`
  end for
  Return the tree under `node*`

Input: ADTree T
Output: ADtree T'
Algorithm **`CrushTree`**(T)
  Scan T top-down, and

  1. If an $\mathtt{OR}^{\mathtt{A}}/\mathtt{OR}^{\mathtt{D}}$ node has a single child, replace the node by its child
  2. If an $\mathtt{OR}^{\mathtt{A}}/\mathtt{OR}^{\mathtt{D}}$ has no child, remove the $\mathtt{OR}^{\mathtt{A}}/\mathtt{OR}^{\mathtt{D}}$ node

Figure 5: `DTree` is used to construct a defense tree from defenses, and `CrushTree` crushes an ADTree by removing unnecessary nodes

`ADTree` constructs either the ADTree from Figure 2b, Figure 3a, or Figure 3b, depending on the input defense model.□

3.3 ADTree Evaluation

An ADTree represents the goal of the attacker, and an evaluation of the tree specifies the likelihood of success of the attacker in achieving this goal. We call the evaluation of the tree its *measure*, and calculate it using the recursive function M.

$$
\begin{aligned}
&M(T) \coloneqq \text{match } T \text{ with}\\
&\quad \mid \mathtt{b}^{\mathtt{A}}(s,a) \to 1\\
&\quad \mid \mathtt{b}^{\mathtt{D}}(s,d,\delta) \to 1e^{-\delta}\\
&\quad \mid \mathtt{OR}^{\mathtt{A}}(T_1,\ldots,T_n) \to \mathtt{max}(M(T_1),\ldots,M(T_n))\\
&\quad \mid \mathtt{OR}^{\mathtt{D}}(T_1,\ldots,T_n) \to \mathtt{min}(M(T_1),\ldots,M(T_n))\\
&\quad \mid \mathtt{AND}^{\mathtt{A}}(T_1,\ldots,T_n) \to \mathtt{min}(M(T_1),\ldots,M(T_n))\\
&\quad \mid \mathtt{AND}^{\mathtt{D}}(T_1,\ldots,T_n) \to \mathtt{max}(M(T_1),\ldots,M(T_n))\\
&\quad \mid \mathtt{C}^{\mathtt{A}}(\mathtt{b}^{\mathtt{A}}(s,a),T^{\mathtt{D}}) \to \mathtt{min}(M(\mathtt{b}^{\mathtt{A}}(s,a)),M(T^{\mathtt{D}}))\\
&\quad \mid \mathtt{C}^{\mathtt{D}}(\mathtt{b}^{\mathtt{D}}(s,d,\delta),T^{\mathtt{A}}) \to \mathtt{max}(M(\mathtt{b}^{\mathtt{D}}(s,d,\delta)),M(T^{\mathtt{A}}))
\end{aligned}
$$

Basic attack nodes are always assigned a value of 1 for likelihood of a successful attack. Assigning a number to the level of attack is quite difficult and would hold true for only a short period of time, and not for the lifetime of a system. According to [7], the issue with deciding the likelihood of various attacks is that "the risk values may be different for different researchers according to the information available and level of analysis. Hence, more emphasis should be put on countermeasures for threats which receive high priority." Thus, we chose to assume a worst-case likelihood for attacks (from the point of view of defendeing the system) and give more fine-grained scores for defenses.

3.3.1 Satisfaction

A cyber-requirement q specifies the severity level l of a *CIA* of an outport of the system.

A defense model $\mathbb{M}$ corresponding to AADL model *mod* *satisfies* q,

$$\mathbb{M} \vdash q,$$

if $M(T_q) \leq \rho$ where $T_q = \mathtt{ADTree}(mod, q, \mathbb{M})$ and ρ is the acceptable level of risk corresponding to l from Table 1. In this case, we also say that T_q satisfies q, or $T_q \vdash q$. Intuitively, implementing the defenses from $\mathbb{M}$ in *mod* results in an ADTree whose attacks are mitigated. Satisfaction of a requirement by a model (resp. ADTree) is naturally extended to a set of requirements.

$$\mathbb{M} \vdash Q \quad \text{if } \forall q \in Q,\ \mathbb{M} \vdash q$$

A tree constructed from $\mathbb{M}_{\mathbb{A}}$ satisfies its requirements, by definition, while one constructed from $\mathbb{M}_{\mathbb{I}}$ may or may not.

Example 3.3. The following are the evaluations of the ADTrees from our applicable and implemented defense models.

$$
\begin{aligned}
\mathtt{ADTree}(mod, q, \mathbb{M}_{\mathbb{A}}) &= 1 \times 10^{-5}\\
\mathtt{ADTree}(mod, q, \mathbb{M}_{\mathbb{I}}) &= 1 \times 10^{-7}\\
\mathtt{ADTree}(mod, q, \mathbb{M}'_{\mathbb{I}}) &= 1
\end{aligned}
$$

Thus, $\mathbb{M}_{\mathbb{A}}$ and $\mathbb{M}_{\mathbb{I}}$ satisfy q while $\mathbb{M}'_{\mathbb{I}}$ doesn't. The following

presents the evaluation of some of the nodes from $\mathbb{M}_{\mathbb{A}}$.

$$\begin{aligned}
M(N15) &= M(\mathtt{b}^{\mathtt{D}}(R,d_2,5)) = 1\times10^{-5}\\
M(N11) &= M(\mathtt{C}^{\mathtt{A}}(\mathtt{b}^{\mathtt{A}}(R,390),N15))\\
&= \mathtt{min}(1,1\times10^{-5}) = 1\times10^{-5}\\
M(N7) &= M(\mathtt{C}^{\mathtt{D}}(\mathtt{b}^{\mathtt{D}}(R,d_1,5),N11)) =\\
&\mathtt{max}(1\times10^{-5},N11) = 1\times10^{-5}\\
M(N4) &= M(\mathtt{OR}^{\mathtt{D}}(N7,N8))\\
&= \mathtt{min}(1\times10^{-5},1\times10^{-5}) = 1\times10^{-5}\\
M(N3) &= M(\mathtt{AND}^{\mathtt{A}}(N5,N6))\\
&= \mathtt{min}(1\times10^{-5},1\times10^{-5}) = 1\times10^{-5}\\
M(N1) &= M(\mathtt{OR}^{\mathtt{A}}(N2,N3))\\
&= \mathtt{max}(1\times10^{-5},1\times10^{-5}) = 1\times10^{-5}
\end{aligned}$$

An evaluation using the applicable defense model is always within the level of severity corresponding to the requirement. The evaluation of $\mathbb{M}'_{\mathbb{I}}$ shows that the implementation does not succeed in stopping the attacker because the bypassing physical security attack isn't defended at all, and neither of CAPEC–184 and CAPEC–51 are defended sufficiently. $\mathbb{M}_{\mathbb{I}}$, on the other hand, is able to satisfy the requirement. □

4 ADTree Synthesis

While the goal of analysis is to construct an ADTree from an AADL model and determine whether the cyber-requirements are satisfied (alternatively, whether the attacks corresponding to the ADTree are mitigated), synthesis constructs an optimal set of defenses based on a cost model for these defenses and (possibly) on the currently implemented defenses.

We define the concepts of synthesized defense models, and cost models.

- **Synthesized Defense Model**
  A synthesized defense model $\mathbb{M}_{\mathbb{S}}$ is the set of optimal defenses to implement, output by synthesis.

  if $(s,a,\gamma,\Delta)\in\mathbb{M}_{\mathbb{S}}$, then for each $(d,\delta)\in\Delta$, the implementation of defense d to DAL-score δ in s is part of the optimal solution to mitigate γ attack a

- **Cost Model**
  The *cost model* $\mathbb{C}$ associates a cost with each component–defense–DAL-score triple from the tuples in a defense model. Given defense d, sub-component s, and DAL-score δ, the cost of implementing d in s to δ is the non-negative real number represented by $\mathbb{C}(s,d,\delta)$.

$$\mathbb{C} : S\times D\times DAL\to\mathbb{R}_{\geq0}$$

We define the cost model of a defense model $\mathbb{M}$ as follows.

$$\mathbb{C}(\mathbb{M}) = \forall(s,a,\gamma,\Delta)\in\mathbb{M},\ \sum_{(d,\delta)\in\Delta}\mathbb{C}(s,d,\delta)$$

The only restriction we place on cost models is that costs must be monotonically increasing with respect to the DAL-scores, that is, $\delta_i>\delta_j\to\mathbb{C}(s,d,\delta_i)\geq\mathbb{C}(s,d,\delta_j)$, for any $s\in S$, $d\in D$, and $\delta_i,\delta_j\in DAL$. This reflects the expectation that higher DALs are more expensive to implement (or at least, not cheaper).

The synthesis problem seeks an optimal solution with respect to $\mathbb{C}$. The cost may represent financial cost, time required for implementation, or perhaps some compound or abstract definition of cost. For simplicity, one might consider a cost model that assigns the DAL-score as the cost of a component–defense–DAL-score triple, $\mathbb{C}(s,d,\delta)=\delta$ (for arbitrary s, d, and δ).

Example 4.1. Recollect the requirement q for our example drone system:

$$q:\ (d_{out}:\mathtt{I}) \text{ with Major } (1\times10^{-5}) \text{ severity level}$$

As we informally stated in Example 3.3, one implementation of the defenses doesn't satisfy q, another does, and the applicable defenses also satisfy q, by definition.

$$\begin{aligned}
\mathbb{M}_{\mathbb{A}} &\vdash q\\
\mathbb{M}_{\mathbb{I}} &\vdash q\\
\mathbb{M}'_{\mathbb{I}} &\nvdash q
\end{aligned}$$

Now, we define a cost model $\mathbb{C}$ for the drone system.

$$\mathbb{C}(s,d,\delta) = \begin{cases}
2\delta, & \text{for } (Remote,d_1,\delta)\\
2\delta, & \text{for } (Remote,d_2,\delta)\\
4\delta, & \text{for } (Backup,d_3,\delta)\\
2\delta, & \text{for } (Controller,d_4,\delta)\\
2\delta, & \text{for } (Controller,d_5,\delta)\\
3\delta, & \text{for } (Controller,d_6,\delta)\\
\delta, & \text{otherwise}
\end{cases}$$

□

4.1 Synthesis Problem Statement

The goal of synthesis is to construct a set of defenses to mitigate all attacks with the least cost. We distinguish 3 cases to synthesize solutions for.

1. Ignore implemented defenses. In this case, the job of synthesis is to synthesize a defense model $\mathbb{M}_\mathbb{S}$ from scratch such that $\mathbb{M}_\mathbb{S} \vdash Q$ and $\mathbb{C}(\mathbb{M}_\mathbb{S})$ is minimal. This case finds a globally minimal solution, in the sense that every other solution which mitigates the attack-defense tree must have a cost greater than or equal to the cost of $\mathbb{C}(\mathbb{M}_\mathbb{S})$. It resembles the early design phase of a system, when defenses have not yet been implemented.
2. Use implemented defenses. There are two possible cases to consider here.
   (a) $\mathbb{M}_\mathbb{I} \vdash Q$. In other words, all possible attacks are mitigated and the requirements in Q are satisfied by the implemented defenses in $\mathbb{M}_\mathbb{I}$. In this case, synthesis tries to optimize the implemented defenses. $\mathbb{M}_\mathbb{S}$ is an optimization of $\mathbb{M}_\mathbb{I}$ using any combination of:
      i. eliminating unnecessary defenses — removing tuples from $\mathbb{M}_\mathbb{I}$
      ii. downgrading current defenses — replacing $(s,a,\gamma,\{(d,\delta_i),\Delta_R\})$ in $\mathbb{M}_\mathbb{I}$ with $(s,a,\gamma,\{(d,\delta_j),\Delta_R\})$ such that $\delta_j < \delta_i$

   This case resembles a situation where successful defenses have already been implemented, but can be downgraded or removed to save costs. Here, we restrict addition of new defenses to save costs.
   (b) $\mathbb{M}_\mathbb{I} \nvdash Q$. In other words, the requirements in Q are not satisfied by the implemented defenses in $\mathbb{M}_\mathbb{I}$. In this case, synthesis corrects the implemented defenses with the least amount of change possible. $\mathbb{M}_\mathbb{S}$ is a modification of $\mathbb{M}_\mathbb{I}$ using some combination of:
      i. implementing new defenses — adding triples to $\mathbb{M}_\mathbb{I}$
      ii. upgrading current defenses — replacing $(s,a,\gamma,\{(d,\delta_i),\Delta_R\})$ in $\mathbb{M}_\mathbb{I}$ with $(s,a,\gamma,\{(d,\delta_j),\Delta_R\})$ such that $\delta_j > \delta_i$

   A real-life application of this situation is one where defenses have been implemented, unsuccessfully, and need to be improved to mitigate attacks, at minimal additional cost. The already implemented defenses are considered sunk costs that cannot be recovered and thus are not downgraded or removed.

4.2 MaxSMT Encoding for Synthesis

The problem of optimizing the defense costs is stated as a MaxSMT problem and sent to Z3's MaxSMT solver [3]. The input to the MaxSMT solver is an SMT-LIB [2] script (with some extensions for the optimization commands) that include (i) declarations of variables, (ii) assertions of formulas, and, (iii) an expression over the variables to optimize, given the constraints asserted. Our MaxSMT encoding depends on the case we are encoding from Subsection 4.1.

In all 3 cases, we do the following. For each component–defense pair (s,d) ($(s,\_,\_,\{(d,\_),\_\}) \in \mathbb{M}_\mathbb{A}$), we declare a variable $v_{s,d}$ which stands for the real number representing the synthesized cost of implementing defense d in component s to a particular DAL δ (for each $(s,\_,\_,\{(d,\delta),\_\})$ that we care about, we add a constraint on $v_{s,d}$, as we will show). Since we allow repeated labels, notice that during the creation of these $v_{s,d}$ variables, multiple nodes in the tree might necessitate the creation of the same variable. Some mechanism, such as a hash table, would have to check that variable declarations aren't repeated in the SMT script. Constraints, however, can be repeated.

For each variable $v_{s,d}$, we assert that the cost is non-negative. Then, we encode the $\texttt{ADTree}(mod, Q, \mathbb{M}_\mathbb{A})$ as an assertion, where *mod* is the AADL model of the system, Q is the set of requirements, and $\mathbb{M}_\mathbb{A}$ is the applicable defense model. Since $\mathbb{M}_\mathbb{A}$ satisfies Q, this assertion sets a baseline on the synthesized model. The following function F converts an ADtree to a quantifier-free first-order formula, which is asserted.

$$
\begin{aligned}
F(T) \coloneqq\ & \text{match } T \text{ with} \\
& \mid \texttt{OR}^\texttt{A}(T_1,\ldots,T_n) \to F(T_1) \wedge \ldots \wedge F(T_n) \\
& \mid \texttt{OR}^\texttt{D}(T_1,\ldots,T_n) \to F(T_1) \vee \ldots \vee F(T_n) \\
& \mid \texttt{AND}^\texttt{A}(T_1,\ldots,T_n) \to F(T_1) \vee \ldots \vee F(T_n) \\
& \mid \texttt{AND}^\texttt{D}(T_1,\ldots,T_n) \to F(T_1) \wedge \ldots \wedge F(T_n) \\
& \mid \texttt{C}^\texttt{A}(\texttt{b}^\texttt{A}(s,a),T^\texttt{D}) \to F(\texttt{b}^\texttt{A}(s,a)) \vee F(T^\texttt{D}) \\
& \mid \texttt{C}^\texttt{D}(\texttt{b}^\texttt{D}(s,d,\delta),T^\texttt{A}) \to F(\texttt{b}^\texttt{D}(s,d,\delta)) \wedge F(T^\texttt{A}) \\
& \mid \texttt{b}^\texttt{A}(s,a) \to \bot \\
& \mid \texttt{b}^\texttt{D}(s,d,\delta) \to v_{s,d} \geq \mathbb{C}(s,d,\delta)
\end{aligned}
$$

$\texttt{OR}^\texttt{A}$ nodes are translated to conjunctions and $\texttt{AND}^\texttt{A}$ nodes to disjunctions because the ADTree is concerned with the success of the attacker while the MaxSMT encoding is concerned with the success of defending any possible attack. If an attacker needs a conjunction ($\texttt{AND}^\texttt{A}$) of attacks to succeed, it suffices from the defender's point of view to stop at least one of the attacks successfully, and hence the disjunction in the MaxSMT encoding. The reasoning for using conjunctions for $\texttt{OR}^\texttt{A}$ nodes is similar. Finally, we need to minimize the cost, which is done by using the `minimize` command in the SMT-LIB script:

$$\texttt{minimize} \sum_{s \in S, d \in D} v_{s,d}$$

The variables declarations, assertions and the optimization command are common in all cases. Additions to the assertions are unique to each case of the problem statement and we consider each of the 3 cases (all assertions must be added before the optimization command in the script).

4.2.1 Case 1

Since we ignore implemented defenses, the constraint from $\mathbb{M}_{\mathbb{A}} - (F(\texttt{ADTree}(mod, Q, A, \mathbb{M}_{\mathbb{A}})))$ suffices to restrict the synthesized solution to one that mitigates all attacks. Additionally, the optimization command assures a global optimum.

4.2.2 Case 2(a)

Since the implemented defenses satisfy the requirements, we assert constraints from $\mathbb{M}_{\mathbb{I}}$ — for each $(s, \_, \_, \{(d, \delta), \_\}) \in \mathbb{M}_{\mathbb{I}}$, assert $v_{s,d} \leq \mathbb{C}(s, d, \delta)$. We also restrict implementation of new defenses — for each $(s, a, \gamma\{(d, \delta), \Delta_R\}) \in \mathbb{M}_{\mathbb{A}}$ such that there exists no $(s, \_, \_, \{(d, \_), \_\}) \in \mathbb{M}_{\mathbb{I}}$, assert $v_{s,d} = 0$. Given the lower bounds from $\mathbb{M}_{\mathbb{A}}$, and the upper bounds from $\mathbb{M}_{\mathbb{I}}$, the MaxSMT solver finds the minimal cost solution, without adding any new defenses.

4.2.3 Case 2(b)

Since the implemented defenses do not satisfy the requirements and the cost of implementing them is considered sunk, we assert them as lower bounds — for each $(s, \_, \_, \{(d, \delta), \_\}) \in \mathbb{M}_{\mathbb{I}}$, assert $v_{s,d} \geq \mathbb{C}(s, d, \delta)$. For defenses that don't work, the constraints from $\mathbb{M}_{\mathbb{A}}$ supersede the lower bound specified by the constraints from $\mathbb{M}_{\mathbb{I}}$.

The MaxSMT encoding for each case is summarized in Figure 6.

Case 1:

For each $s \in S, d \in D$, `declare-var` $v_{s,d}$
For each $v_{s,d}$, `assert` $v_{s,d} \geq 0$
`assert` $F(\texttt{ADTree}(mod, Q, A, \mathbb{M}_{\mathbb{A}}))$
`minimize` $\sum^{s \in S, d \in D} v_{s,d}$

Case 2(a):

For each $s \in S, d \in D$, `declare-var` $v_{s,d}$
For each $v_{s,d}$, `assert` $v_{s,d} \geq 0$
`assert` $F(\texttt{ADTree}(mod, Q, A, \mathbb{M}_{\mathbb{A}}))$
For each $(s, d, \delta) \in \mathbb{M}_{\mathbb{I}}$, `assert` $v_{s,d} \leq \mathbb{C}(s, d, \delta)$
For each $(s, d, \delta) \notin \mathbb{M}_{\mathbb{I}}$, `assert` $v_{s,d} = 0$
`minimize` $\sum^{s \in S, d \in D} v_{s,d}$

Case 2(b):

For each $s \in S, d \in D$, `declare-var` $v_{s,d}$
For each $v_{s,d}$, `assert` $v_{s,d} \geq 0$
`assert` $F(\texttt{ADTree}(mod, Q, A, \mathbb{M}_{\mathbb{A}}))$
For each $(s, d, \delta) \in \mathbb{M}_{\mathbb{I}}$, `assert` $v_{s,d} \geq \mathbb{C}(s, d, \delta)$
`minimize` $\sum^{s \in S, d \in D} v_{s,d}$

Figure 6: A summary of the MaxSMT encoding of the defense optimization problem

4.3 SMT Model Evaluation

All our calls to the MaxSMT solver are expected to be satisfiable. A solution where all possible defenses are implemented to the highest possible DAL would trivially satisfy the problem (while being unnecessarily expensive):

$$\forall (s, a, \gamma, d, \delta) \in \mathbb{M}_{\mathbb{A}}, \\ (s, a, \gamma, d, 9) \in \mathbb{M}_{\mathbb{S}}$$

The response from the solver varies in its optimization of defense cost. The variables $v_{s,d}$ in our SMT encoding model the cost of implementing defense d in component s to some DAL-score. Thus, when the SMT solver returns an optimal solution as a model, it returns an optimal cost for each component-defense pair. We need to build $\mathbb{M}_{\mathbb{S}}$ from this for which we need the DAL-score for each component-defense pair. We define the inverse cost function $\mathbb{C}^{-1}$ that given a component-defense-cost triple, returns the DAL-score to implement the defense to in the component. Since a component–defense pair could have the same cost for multiple DAL-scores (the monotonicity requirement does not prevent this), the inverse isn't over an injective function. We break ties by preferring higher DAL-scores, given equal costs.

$$\mathbb{C}^{-1}(c,s,d) = \max\{\delta_i \mid \mathbb{C}(s,d,\delta_i) = c\}$$

Thus, as a minimal cost solution, for each component s and defense d, the SMT solver returns a cost as the real value of $v_{s,d}$. $\mathbb{M}_{\mathbb{S}}$ is then constructed as follows.

$$\forall v_{s,d}, \forall a \in A, \text{ such that } (s,a,\gamma,\{(d,\delta),\Delta_R\}) \in \mathbb{M}_{\mathbb{A}},$$
$$(s,a,\gamma,\{(d,\mathbb{C}^{-1}(v_{s,d},s,d))\}) \in \mathbb{M}_{\mathbb{S}}$$

This is the minimal cost defense model synthesized by the MaxSMT solver.

Example 4.2. We construct synthesized defense models from the satisfiable solution returned by the SMT solver for the toy drone system using the cost model in Example 4.1 as follows.

- **Case 1** Without any additional restrictions, the SMT solver returns values 0, 10, 20, 0, 0 and 0 respectively, for $rd1$, $rd2$, $bd3$, $cd4$, $cd5$ and $cd6$, which are its recommended costs for the defenses. Applying the cost inverse function, we have the following DALs to implement the components to.

$$\mathbb{C}^{-1}(0,R,d_1) = 0$$
$$\mathbb{C}^{-1}(10,R,d_2) = 5$$
$$\mathbb{C}^{-1}(20,B,d_3) = 5$$
$$\mathbb{C}^{-1}(0,C,d_4) = 0$$
$$\mathbb{C}^{-1}(0,C,d_5) = 0$$
$$\mathbb{C}^{-1}(0,C,d_6) = 0$$

  This is an optimal cost solution unrestricted by any implementation constraints. The total cost is 30.
- **Case 2(a)** The SMT solver returns cost 0 for $rd1$ and $bd3$, cost 10 for $rd2$, $cd4$, and $cd5$, and 15 for $cd6$. Applying the cost inverse function, we have that d_1 and d_2 are to be implemented to DAL 0 and 5 in the remote, d_3 to Dal 0 in the backup controller, and d_4, d_5, and d_6 all to DAL 5 in the main controller. Here, since the implemented defenses already satisfy the requirement, new ones aren't added, and instead, synthesis suggests reductions. The global optimal solution would choose d_3 over d_4, d_5 and d_6, but since the latter are already implemented, synthesis only suggests DAL reductions where applicable (to d_4 and d_5). The total cost of the synthesized solution is 45, which is cheaper than the implementation which costs 61.
- **Case 2(b)** The SMT solver returns costs 10, 10, 20, 6, 0 and 0 for $rd1$, $rd2$, $bd3$, $cd4$, $cd5$, and $cd6$ which translate to DALs 5, 5, 5, 3, 0 and 0, respectively. Since the unsatisfactory defenses have already been implemented, their cost is considered a sunk cost (16 here). The SMT solver specifies what defenses need to be added to satisfy the requirements — d_2 and d_3 in this case. The total cost of the synthesized solution is 46.

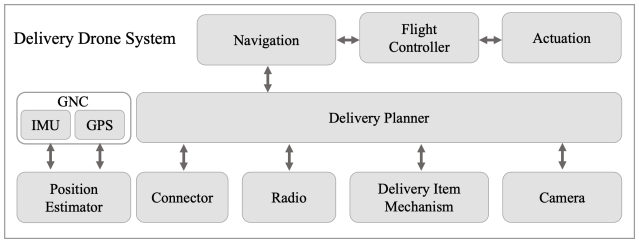

Figure 7: A notional architecture diagram for the delivery drone model

Notice that the same defense can be applicable to a component to defend 2 different attacks. For example, system access control defends against both CAPEC–507 and CAPEC–390. Because our encoding doesn't take into account attacks (and it doesn't need to), once synthesis suggests to implement such a defense, we add all possible occurrences of it to $\mathbb{M}_{\mathbb{S}}$. While this redundance is necessary for soundness of the formalism, it can be ignored during implementation. In fact, it isn't necessary to map synthesized defenses to attacks they mitigate at all, we do it in the formalism just to be able to make synthesized solutions comparable with applicable and implemented solutions. □

5 Evaluation

A prototype of ADTree-based security analysis and synthesis was implemented in the VERDICT toolchain, which is an AADL plugin for the OSATE tool [1]. We perform an evaluation on a high-fidelity architecture model of a delivery drone to demonstrate the capabilities of the tool. The VERDICT tool and the AADL model is publicly available.[1] The delivery drone is part of a delivery mission that delivers packages in neighborhoods. The mission consists of a delivery truck, an operator, multiple drones and packages to deliver. When the truck arrives at a neighborhood, a drone is launched to deliver packages to nearby homes. The drone navigates to the delivery site using GPS. Upon arriving at a site, it captures a picture of the site to ensure it is free of obstacles and is safe to release the package. For a high-valued package, it needs to confirm with the operator in the truck via radio before dropping off the package.

A notional architecture for the delivery drone is shown in Figure 7. The model consists of 11 inter-connected components and is annotated with meta-level properties, defenses properties, cyber relations and cyber requirements. Meta-level properties are properties like component type and pedigree, that come built-in with the AADL properties. Given this system, the STEM component of the VERDICT toolchain iden-

[1]The delivery drone AADL model is available at https://github.com/baoluomeng/USENIX_paper/tree/master/DeliveryDrone

```
CyberReq {
    id = "CyberReq02"
    description = "The drone shall be resilient to maliciously
                   commanded improper delivery of a package"
    condition = delivery_status:I
    severity = Hazardous
};
```

Figure 8: A cyber requirement example

Synthesis results (not using existing solution): total cost: 385.0

Designer Action	Component/Connection	Defense Property	Target DAL	Target Cost
Implement	actuation	systemAccessControl	7	7.0
-do-	camera	systemAccessControl	7	7.0
-do-	deliveryItemMechanism	supplyChainSecurity	7	7.0
-do-	deliveryItemMechanism	systemAccessControl	7	7.0
-do-	deliveryItemMechanism	tamperProtection	7	7.0
-do-	deliveryPlanner	antiJamming	7	7.0
-do-	deliveryPlanner	dosProtection	7	7.0
-do-	deliveryPlanner	failSafe	7	7.0
-do-	deliveryPlanner	heterogeneity	7	7.0

(a) Synthesis solution for Case 1 (partial result is shown)

Synthesis results: Existing solution UNSAT, cost: 455.0 -> 469.0 (increase of 14.0)

Designer Action	Component/Connection	Defense Property	Original DAL	Target DAL	Original Cost	Target Cost	Delta Cost
Implement	deliveryItemMechanism	supplyChainSecurity	0	7	0.0	7.0	7.0
-do-	deliveryItemMechanism	tamperProtection	0	7	0.0	7.0	7.0

(b) Synthesis solution for Case 2

Synthesis results: Existing solution SAT, merit assignment, cost: 700.0 -> 385.0 (reduction of 315.0)

Designer Action	Component/Connection	Defense Property	Original DAL	Target DAL	Original Cost	Target Cost	Delta Cost
Remove	actuation	physicalAccessControl	7	0	7.0	0.0	-7.0
-do-	actuation	supplyChainSecurity	7	0	7.0	0.0	-7.0
-do-	c1	encryptedTransmission	6	0	6.0	0.0	-6.0
-do-	c16	deviceAuthentication	9	0	9.0	0.0	-9.0
-do-	c16	encryptedTransmission	9	0	9.0	0.0	-9.0
-do-	camera	physicalAccessControl	7	0	7.0	0.0	-7.0
-do-	camera	supplyChainSecurity	7	0	7.0	0.0	-7.0
-do-	connector	inputValidation	7	0	7.0	0.0	-7.0
-do-	connector	logging	7	0	7.0	0.0	-7.0

(c) Synthesis solution for Case 3 (partial result is shown)

Figure 9: Synthesis solutions for the delivery drone model

tifies possible CAPEC attacks and NIST 800-53 defenses. These attacks and defenses are fed to the synthesis tool for further processing. The defense property is a numerical DAL-score from δ, and represents the rigor (DAL) of implemented defense in each component of the system, which is used to construct the $\mathbb{M}_{\mathbb{I}}$. Cyber relations and requirements are declared in a language annex for AADL – VERDICT. Two cyber requirements are specified to ensure a successful mission that delivers a package to the intended location. One of them, shown in Figure 8, states that the drone shall be resilient to maliciously commanded improper delivery of a package. The satisfaction of this requirements depends on the integrity of the output "*delivery_status*", which is used as a starting point from which the system architecture is traced, to build the ADTree for analysis. Further, the consequence of successful attack is Hazardous, requiring corresponding defenses to be implemented to DAL-score 7.

To demonstrate the optimal defense synthesis capabilities, we invoke the Synthesis tool on the model for the three cases using the default cost model, one where the cost for each defense-DAL pair is just the DAL score.

- **Case 1** The implemented defenses are ignored, and a *global* optimal solution is returned. Synthesis suggest a list of defenses with minimal costs to be implemented to DAL 7 so that all cyber requirements can be satisfied. A partial solution is shown in Figure 9a, due to space restrictions. The total cost for the implementation is 385 unit.
- **Case 2** Implemented defenses are taken into consideration by Synthesis, and these don't satisfy the requirements.(This corresponds to Case 2(b) from our problem statement). Synthesis suggests implementing two defenses for the "*deliveryItemMechanism*" component: Supply Chain Security and Tamper Protection, both to DAL 7, which would allow for the requirements to be satisfied. These would mitigate CAPEC-439 (Manipulation During Distribution), which can be mitigated using the above mentioned defenses. Synthesis's solution is shown in Figure 9b.
- **Case 3** Once the suggested defenses in case 2 are implemented in the model, they would be considered by Synthesis sufficient to satisfy all cyber requirements. In this case, Synthesis does "merit assignment" which is to suggest downgrades/removals of defenses (Case 2(a) from our problem statement) to save costs. Figure 9c shows the output from Synthesis which suggests removal of multiple defenses.

6 Related Work

In our ADTrees, we use nodes with repeated labels — that is, there can exist multiple nodes in our tree that have the same label. Bossuat et al. [4] extend ADTrees to AD-DAGs to deal with repeated labels. In our work, by guaranteeing that these nodes will have the same child-structure, we are able to maintain the ADTree formalism, and also maintain soundness by handling repetitions during our SMT-encoding.

Fila et al. [6] and Kordy et al. [9] find an optimized set of defenses to mitigate an attack defense tree using integer linear programming. We use an SMT-based optimization approach, and also build our trees from AADL models of the system. Additionally, we are able to incorporate implementations of defenses that may or may not satisfy the requirements specified by the ADTree and suggest solutions based on these variations (cases 2(a) and 2(b) from Section 4.1).

We use the formalism of attack-defense trees introduced by Kordy et al. [8] to specify our ADTrees.

7 Conclusion and Future Work

We propose a security analysis technique for system architecture designs via Attack-Defense Trees, and a novel technique

to synthesize optimal cost defenses for the components of a model. We translate the AADL model of a system into an ADTree, and encode this ADTree along with the cost of implementing its defenses into a MaxSMT query, such that a satisfying model of the SMT query is a minimum-cost defense for the system, that mitigates all applicable attacks.

We utilize advancements in the ADTree literature, and SMT technology, in building our formalism of the process of converting an AADL model to an ADTree and then to an optimization query to a MaxSMT solver. We provide an implementation of our technique as the Synthesis functionality in the VERDICT tool chain.

One potential extension to our formalism and our tool is to allow a single defense to defend attacks over multiple components and connections – *extensibility* of defenses.

Acknowledgement & Disclaimer

Distribution Statement "A" (Approved for Public Release, Distribution Unlimited). This research was developed with funding from the Defense Advanced Research Projects Agency (DARPA). The views, opinions and/or findings expressed are those of the author and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

References

[1] The OSATE Tool. https://osate.org/about-osate.html, 2021.

[2] Clark Barrett, Pascal Fontaine, and Cesare Tinelli. The Satisfiability Modulo Theories Library (SMT-LIB). `www.SMT-LIB.org`, 2016.

[3] Nikolaj Bjørner, Anh-Dung Phan, and Lars Fleckenstein. νz - an optimizing smt solver. In Christel Baier and Cesare Tinelli, editors, *Tools and Algorithms for the Construction and Analysis of Systems*, pages 194–199, Berlin, Heidelberg, 2015. Springer Berlin Heidelberg.

[4] Angèle Bossuat and Barbara Kordy. Evil Twins: Handling Repetitions in Attack–Defense Trees: A Survival Guide. In Peng Liu, Sjouke Mauw, , and Ketil Stolen, editors, *Graphical Models for Security*, volume LNCS, pages 17–37, Santa Barbara, United States, August 2017. Springer.

[5] Peter H. Feiler, Bruce Lewis, Steve Vestal, and Ed Colbert. An overview of the sae architecture analysis & design language (aadl) standard: A basis for model-based architecture-driven embedded systems engineering. In Pierre Dissaux, Mamoun Filali-Amine, Pierre Michel, and François Vernadat, editors, *Architecture Description Languages*, pages 3–15, Boston, MA, 2005. Springer US.

[6] Barbara Fila and Wojciech Wideł. Exploiting attack–defense trees to find an optimal set of countermeasures. In *2020 IEEE 33rd Computer Security Foundations Symposium (CSF)*, pages 395–410, 2020.

[7] Ahmad Y Javaid, Weiqing Sun, Vijay K Devabhaktuni, and Mansoor Alam. Cyber security threat analysis and modeling of an unmanned aerial vehicle system. In *2012 IEEE Conference on Technologies for Homeland Security (HST)*, pages 585–590. IEEE, 2012.

[8] Barbara Kordy, Sjouke Mauw, Saša Radomirović, and Patrick Schweitzer. Foundations of attack–defense trees. In Pierpaolo Degano, Sandro Etalle, and Joshua Guttman, editors, *Formal Aspects of Security and Trust*, pages 80–95, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.

[9] Barbara Kordy and Wojciech Widel. How well can i secure my system? In *IFM*, 2017.

[10] Sjouke Mauw and Martijn Oostdijk. Foundations of attack trees. In Dong Ho Won and Seungjoo Kim, editors, *Information Security and Cryptology - ICISC 2005*, pages 186–198, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.

A Translation Soundness

Here, we argue about the soundness of the SMT encoding from Section 4.2. An attack-defense tree models the applicable attacks on the components and connections of a system, and the applicable/implemented defenses that mitigate these attacks. We encode one or more ADTrees as a set of constraints that we send to a MaxSMT solver. The solver may return one of 3 possible results:

1. `sat` - the problem instance is satisfiable. There exists a minimum cost solution which we can construct from the satisfiable model from the solver.
2. `unsat` - the problem instance is unsatisfiable. There exists no solution, given the constraints presented to the solver. The constraints are contradictory. We argue that our encoding never produces such a result, subject to certain assumptions.
3. `unknown` - the solver isn't able to give a conclusive response. When the solver fails, our method fails as well.

The following theorems relates the notion of satisfiability defined in Section 4 to the satisfiability of a formula, and argue about the soundness of our encoding. As a matter of notation, notice that an ADTree satisfies one or more requirements, so an ADTree is said to be satisfying if it mitigates its attacks and unsatisfying otherwise. A formula is satisfiable if there exists a satisfying model for its variables, and unsatisfiable if there doesn't exist any.

Theorem 1. An unsatisfying tree — one that doesn't mitigate the attacks specified by the requirements — is translated to an unsatisfiable set of formulas.

Proof. An ADTree can be made unsatisfying by two possible sources.

1. Undefended attack nodes
2. Attack nodes defended via insufficient defense nodes (DAL-score not high enough)

Case 1 For ADTree T with undefended attack nodes, $F(T)$ is unsatisfiable, and the proof is by induction on F. The most important case is the base case of $\mathtt{b}^{\mathtt{A}}$ nodes. A lone $\mathtt{b}^{\mathtt{A}}$ node (one that isn't encapsulated in a $\mathtt{C}^{\mathtt{A}}$ or $\mathtt{C}^{\mathtt{D}}$ node) is an undefended attack node, which F translates to $\bot$, an unsatisfiable formula. The $\mathtt{b}^{\mathtt{D}}$ node is always satisfying (explained later in the proof) and the other node combinators simply combine translations of child nodes using conjunctions and disjunctions.

Case 2 ADTrees defended via insufficient defense nodes will be converted to satisfiable formulas using F. This is because F models a constraint on the cost of implementing the defense, not its ability to mitigate an attack. A $\mathtt{b}^{\mathtt{D}}$ node is translated to an inequality between a Real-valued variable and its cost - a Real number. Independently, it is satisfiable, and can only be made unsatisfiable when considered along with a contradictory inequality/equality. However, the only kind of inequalities F adds are $\geq$ inequalities between a variable on the LHS and a constant on the RHS. Any two inequalities $v \geq x$ and $v \geq y$ over a variable v and constants x and y are satisfiable, since the satisfying model will contain a value for v which is greater than or equal to the maximum of x and y, and this value will also be greater than or equal to the other constant.

The conversion of an unsatisfying ADTree to a satisfiable query is a source of unsoundness. To prevent this from happening, our encoding only calls F on ADTrees built from applicable defenses — and these are satisfiable by definition.

The non-negativity constraints don't introduce unsoundness — since each of the $v_{s,d}$ variables represent the cost of implementing defense d in component s, we expect these values to be non-negative.

The constraints added by Case 2(a) of our problem statement (4.1) take a currently, satisfying ADTree, and specify upper bounds on the costs of the defenses in the ADTree. If $v_{s,d}$ is currently some value x, $v_{s,d} \leq x$ introduces no unsoundness, since this constraint still allows for $v_{s,d} = x$. Additionally, constraints are added restricting currently unimplemented defenses from being implemented. Since the current implementation is already satisfying in this case, not allowing new defenses doesn't introduce unsatisfiability.

The constraints added by Case 2(b) also don't introduce any unsoundness, since constraints of the form $v_{s,d} \geq x$ don't add unsoundness (same argument as for independent $\mathtt{b}^{\mathtt{D}}$ nodes above).

Therefore, the SMT encoding translates an unsatisfying ADTrees to unsatisfiable formulas. □

Theorem 2. A satisfying ADTree is translated to a satisfiable set of formulas, and the satisfying SMT model translates to a cheapest (satisfying) defense implementation.

Proof. This reduces to proving that our encoding only encodes the necessary and sufficient conditions of the ADTree as a logical constraint. We argue this for F incuctively.

- Independently, an attack node $\mathtt{b}^{\mathtt{A}}$ is unsatisfying. It is translated to $\bot$ which is an unsatisfiable formula.
- Independently, a defense node $\mathtt{b}^{\mathtt{D}}$ is satisfying. It is translated to to an inequality constraint $v_{s,d} \geq x$ for some constant x, and this constraint is satisfiable, because

our translation only introduces constraints of the form $v_{s,d} \geq x$ which can't contradict each other.

- A defense node with a countermeasure attack tree is satisfying if the defense node is satying, and the countermeasure attack tree is also satisfying. $\texttt{C}^{\texttt{D}}$ nodes are therefore encoded as conjunctions of their respective child nodes.
- An attack node with a countermeasure defense tree is translated to a disjunction of the translations of the attack node and the defense tree. The attack node is translated to $\bot$, so the satisfiability of $F(\texttt{C}^{\texttt{D}})$ is reduced to the satisfiability of the translation of its defense tree child.
- Within a defense tree, an $\texttt{AND}^{\texttt{D}}$ indicates that all child defense trees have to be implemented to defend against the attack in question, and $\texttt{OR}^{\texttt{D}}$ indicates that at least one of the child defense trees have to be implemented. As a consequence, $\texttt{AND}^{\texttt{D}}$ nodes are translated to conjunctions, and $\texttt{OR}^{\texttt{D}}$ nodes to disjunctions.
- For $\texttt{AND}^{\texttt{A}}$ and $\texttt{OR}^{\texttt{A}}$ nodes, the translations are reversed. An $\texttt{AND}^{\texttt{A}}$ node indicates that the attacker needs all the child attack trees to succeed for the parent to succeed. From the point of view of the defender (which is how an ADTree is analyzed), it suffices to defend at least one of the child attacks. Thus, an $\texttt{AND}^{\texttt{A}}$ node is translated to a disjunction. For similar reasons, an $\texttt{OR}^{\texttt{A}}$ node is translated to a conjunction.

F is applied to an ADTree constructed from the applicable defense model. The applicable defense model consists of the minimal DAL necessary for a defense to mitigate its respective attack. Thus, the defense nodes in the ADTree consist of the minimal DAL necessary for a particular defense to work. F asserts the cost of implementing this defense–DAL-score pair as a lower bound for the defense in its corresponding component. Therefore, the constraints from F are necessary and sufficient. Given the minimization command to the MaxSMT solver, it will find a least cost model.

Case 2(a) adds upper bounds from currently satisfying implementations, so it still allows for a minimal cost model, with the restriction that new defenses may not be added.

Similarly, Case 2(b) adds lower bounds for already implemented defenses since we consider them a sunk cost, and minimizes given this restriction. □

We have argued that the problem instance is soundly translated to MaxSMT, and that we can trust its result. We finish by arguing that our cost-inverse function soundly extracts the correct defense implementations for a minimal cost solution. This follows from the fact that our cost model is monotonically increasing, and that the cost-inverse function breaks the only possible ties from the costs by preferring higher DAL-scores when the cost of implementing a defense to a higher DAL-score is the same as that of implementing it to a lower DAL score (Section 4.3).