A Vulnerability-Centric Requirements Engineering Framework: Analyzing Security Attacks, Countermeasures, and Requirements Based on Vulnerabilities

Golnaz Elahi
University of Toronto
gelahi@cs.toronto.edu

Eric Yu
University of Toronto
yu@ischool.utoronto.ca

Nicola Zannone
University of Toronto
zannone@cs.toronto.edu

Abstract

Many security breaches occur because of exploitation of vulnerabilities within the system. Vulnerabilities are weaknesses in the requirements, design, and implementation, which attackers exploit to compromise the system. This paper proposes a methodological framework for security requirements elicitation and analysis centered on the concept of vulnerability. The framework offers modeling and analysis facilities to assist system designers in analyzing vulnerabilities and their effects on the system; identifying potential attackers and analyzing their behavior for compromising the system; and identifying and analyzing the countermeasures to protect the system. The framework proposes a qualitative goal model evaluation analysis for assessing the risks of vulnerabilities exploitation and analyzing the impact of countermeasures on such risks.

1 Introduction

Developing secure software systems is challenging because errors and misspecifications in requirements, design, and implementation can bring vulnerabilities to the system. Attackers most often exploit vulnerabilities to compromise the system. In security engineering, a vulnerability is an error or weakness of a system or its environment that in conjunction with an internal or external threat can lead to a security failure [1]. For example, vulnerabilities may result from input validation errors, memory safety violations, weak passwords, viruses, or other malware.

In recent years, software companies and government agencies have become particularly aware of security risks that vulnerabilities impose on the system security and have started analyzing and reporting detected vulnerabilities of products and services. For instance, the IBM Internet Security Systems X-Force [14] has detected and analyzed 6,437 new vulnerabilities in 2007, of which 1.9% are critical and 37% are high risk. 20% of the 5-top critical vulnerabilities were found to be unpatched. Of all the vulnerabilities disclosed in 2007, only 50 percent can be corrected through vendor patches, and 90 percent of vulnerabilities could be remotely exploited. These statistics show the critical urgency of the vulnerabilities affecting software services and products. Various web portals and on-line databases of vulnerabilities are also made available to security administrators. For example, the National Vulnerability Database[1], SANS top-20 annual security risks[2], and Common Weakness Enumeration (CWE)[3] provide updated lists of vulnerabilities and weaknesses. The Common Vulnerability Scoring System (CVSS)[4] also provides a method for evaluating the criticality of vulnerabilities.

[1] http://nvd.nist.gov/
[2] http://www.sans.org/top20/
[3] http://cwe.mitre.org/
[4] http://www.first.org/cvss/

In requirements engineering, agent- and goal-oriented approaches have been found useful for understanding security issues that arise from interactions of multiple actors with malicious or non-malicious intentions, e.g., [18, 19, 10]. Knowledge about vulnerabilities, however, have not been taken advantage of in these security requirements engineering framework. In addition, goal models provide suitable basis for attaching vulnerabilities to the required actions and assets to achieve the goals and propagating vulnerabilities among the actors and system participants. By identifying vulnerabilities or classes of vulnerabilities and associating them with the activities and assets that bring them to the system, the analysts understand how weaknesses are brought to system and how flaws in one part of the system are spread out to other parts. Information about potential attacks that exploit vulnerabilities can be linked to requirements to analyze the effects of the exploited vulnerabilities on activities or goals of stakeholders. Analysts also need to decide about alternative countermeasures by analyzing their impacts on attacks.

Analyzing the effects of vulnerabilities on the system allows one to assess the risks of attacks, analyze the efficacy of countermeasures, and decide on patching or disregarding the vulnerabilities by taking advantage of goal model evaluation techniques [11, 6, 13]. By adapting goal model evaluation, the analysts can verify whether top goals of stakeholders are satisfied with the risk of vulnerabilities and attacks and assess the efficacy of security countermeasures against the risks. The evaluation does not only specify if the goals are satisfied; but also makes it possible to understand why and how the goals are satisfied (or denied) by tracing back the result of evaluation to vulnerabilities, attacks and countermeasures. In addition, the resulting security goal models and goal model evaluation can provide a basis for trade-off analysis among security and other quality requirements [9]. New vulnerabilities are continuously being uncovered. By linking requirements, vulnerabilities, and countermeasures to each other in a modeling framework, one can update the models with newly detected vulnerabilities in order to analyze the risks imposed by the new vulnerabilities.

Current state of the art raises the need for a systematic way to link the empirical security knowledge such as information about vulnerabilities, attacks, and proper countermeasure to stakeholders' goals and security requirements. Secure software engineering frameworks rarely use vulnerabilities to elicit security requirements. Existing frameworks focus on various aspects for eliciting security requirements such as design of secure components [16], security issues in social dependencies among actors [19] and their trust relationships [10], attacker behavior [31, 28] and attacker goals [32], and events that can cause failure in the system [2].

In particular, Liu et al. [19] propose a vulnerability analysis approach for eliciting security requirements. However, vulnerabilities in this framework is different from the one defined in the security engineering community (i.e. weaknesses in the system). Liu et al. refer to vulnerabilities as the weak dependencies that may jeopardize the goals of depender actors. Only some security software engineering approaches consider analyzing vulnerabilities, as weaknesses of the systems, during the elicitation of security requirements. For instance, in [20], vulnerabilities are modeled as beliefs inside the boundary of attackers and may contribute positively to the attacks. However, the resulting model does not specify how the vulnerability is brought to the system, what actions or assets cause the vulnerabilities, and which actors are vulnerable in the system. In addition, the impact of countermeasures on the vulnerabilities and attacks are not captured. The CORAS framework [8] also provides a way for expressing how a vulnerability leads to another vulnerability (or to a threat) and how a vulnerability (or combination of vulnerabilities) lead to a threat. However, similar to [20], CORAS does not investigate which design choices, requirements, or processes has brought the vulnerabilities to the system.

In our previous work [9], we have introduced the concepts of vulnerability into a security conceptual modeling method. Vulnerabilities are treated as weaknesses in the structure of goals and activities of intentional agents. Analyzing vulnerabilities together with system requirements and linking attacks to vulnerabilities allow one to analyze how attacks can compromise the system by exploiting vulnerabilities and identify the security mechanisms needed to protect the system. This paper extends and refines our previous work by proposing an agent-

and goal-oriented framework for eliciting and analyzing security requirements by linking empirical knowledge of vulnerabilities to requirements models. In particular, this work provides a framework that enables understanding how vulnerabilities can be exploited to compromise the security of the system. The proposed vulnerability-centric security requirements framework is the result of surveying current critical vulnerabilities in security engineering discipline to understand how vulnerabilities are brought to the system, exploited by the attacks, and handled by the countermeasures.

The structure of the paper is organized as follows. Section 2 introduces the security concepts used in the paper with a particular focus on the concept of vulnerability and related notions. Section 3 introduces the meta-model of the framework, in which security concepts are incorporated into an agent- and goal-oriented modeling framework. Section 4 describes the modeling process, and Section 5 proposes a method for analyzing security requirements based on the goal model evaluation techniques. The modeling and analysis methods described are illustrated by case examples. Section 6 overviews the current state of the art in threat analysis and security requirements engineering. Finally, Section 7 draws a conclusion and discusses future work.

2 Relevant Concepts

This section investigates the conceptual foundation for the security requirements engineering framework proposed in this paper. We identify the necessary and sufficient security conceptual modeling constructs to be used in the meta-model of our framework in Section 3.

In the security engineering, an *asset* is *"anything that has value to the organization"* [15]. Assets can be people, information, software, and hardware [8]. They can be the target of attackers, and consequently, they need to be protected. A *vulnerability* is a weakness in the system which allows an attacker to compromise its correct behavior [1, 27, 26]. Vulnerabilities are brought to the system by adopting a product or executing a service. Identifying and analyzing the backdoors that attacks can exploit to compromise the system helps adopting effective countermeasure that prevent the security failures through the vulnerabilities.

Security professionals consider *threats* as potential ways an attacker can attack a system [29]. Attacks and threats are usually used interchangeably. An *attack* is a set of intentional unwarranted actions which attempts to break the security of a system or a component of a system [29]. Though the general idea of attack is clear, there is no consensus on a precise definition. For instance, Schneider [27] points out that an attack can occur only in presence of a vulnerability. Conversely, Schneier [29] broadens this vision, considering also attacks that can be performed without exploiting vulnerabilities. By analyzing the possible ways in which a system can be attacked, analysts can assess the risk and cost of attacks and understand their impact on system security. Such knowledge helps them in the identification of appropriate countermeasures to protect the system as well.

Attackers can have malicious *goals* and perform *actions* to achieve them. Examples of malicious goals are *disrupt or halt services*, *access confidential information*, and *improperly modify the system* [3]. Schneier [28] argues that understanding who the attackers are and their motivations, goals, and targets, aids designers in adopting proper countermeasures to deal with the real threats. Analyzing the source of attacks helps to better predict the actions taken by the attacker. Schneier illustrates this through an example: When the attackers are terrorists, we have to worry about attackers who are willing to die to achieve their goal. But if we are worried about bored graduate students studying the security of our system, we usually do not have to worry about illegal attacks such as bribery and blackmail [28].

A *countermeasure* is a protection mechanism employed to secure the system [29]. Countermeasures can be actions, processes, devices, solutions, or systems intended to prevent a threat from compromising the system. For instance, they are used to patch vulnerabilities or prevent their exploitation. Countermeasure are selected according to the attacks and vulnerabilities.

Besides the concepts described, there are other concepts relevant to security requirements. For instance, Giorgini et al. [10] integrate concepts from Trust Management, such as permission, trust and delegation, into a Requirements Engineering framework to address authorization issues in the early phases of software development process. Risk analysis frameworks (e.g., [2]) employ the concepts of event to model uncertain circumstances that affect the correct behavior of the system. However, events do not indicate involvement of intentionality; therefore, the event concepts is appropriate to assess risks and safety requirements in critical systems, and it does not allow the analysis of (malicious) intentional behavior. We do not incorporate the concept of risk into the meta-model, since risk is a value that represent the potential consequences of an attack.

However, security is not only limited to the identification of protection mechanisms to address vulnerabilities. Security originates from human concerns and intents [19]; the social issues of organizations where different actors can collaborate or compete to achieve their goals should be considered as part of security requirements analysis [10, 19]. In addition, security is a subjective and personal feeling [30]; therefore, security requirements analysis and security-related decision makings require analyzing personal and organizational goals of the stakeholders participating to the system. For this purpose, we take advantage of agent- and goal-oriented concepts such as intentional actor, goal, and social dependency. There are evidences in the security requirements engineering literature (e.g., [5, 7, 10, 19, 32, 35]) that these concepts provide the means for analysis of organizational and social contexts in which the system-to-be operates. In the next section, we show how security concepts can be integrated in the meta-model underlying the i* agent- and goal-oriented framework.

3 An extended i* Meta-Model

Security is both a system and a social and organizational problem. The ability of the i* framework [35] to model agents, goals, and their dependencies makes it suitable for understanding security issues that arise among multiple malicious or non-malicious agents with competing goals. In addition to modeling actors, i* offers a way to model actors' dependencies, goals, assets, and actions, refinements of goals into the actions and assets, and decomposition of actions. Thus, the i* framework provides the basic setting suitable to represent vulnerabilities brought by actions and assets and propagate them to the elements through the decomposition and dependency links. Moreover, i* enables modeling contribution of goals, actions, and assets on other goals. Such relations can be used to capture the effects of vulnerabilities on the satisfaction of system and stakeholders' goals.

In this section, we present the meta-model for the security requirements engineering framework, which extends the i* meta-model with security concepts. The meta-model includes both the i* Strategic Dependency (SD) diagram, which captures the actors and their dependencies and the i* Strategic Rationale (SR) diagram, which expresses the internal goals and the behavior of actors to achieve their goals. The meta-model also captures the concepts of vulnerability, attack, security countermeasure, and their corresponding relationships with i* constructs.

3.1 The i* Meta-Model

The framework proposed in this paper extends the i* framework with the concepts of vulnerability and attack. Figure 1 depicts the i* meta-model together with the security concepts. The proposed framework employs agents and roles as two types of actors. An actor is an active entity that has strategic goals and intentionality within the system or the organizational setting, carries out activities, and produces entities to achieve goals by exercising its knowhow [35]. Actors can be roles or agents. A *role* captures an abstract characterization of the behavior of a social actor within some specialized context or domain of endeavor. Its characteristics are easily transferable

to other social actors. An *agent* is an actor with concrete and physical manifestations and can play some role.

Intentional elements defined by the i* framework are goals, softgoals, tasks, and resources. A *goal* represents the intentional desire of an actor, without specification of how the goal is satisfied. Goals are also called hard goals in contrast to *softgoals* which do not have clear criteria for deciding whether they are satisfied or not. A *task* is a sequence of actions which the actor needs to perform (or depends on other actors to perform it) to achieve a goal. Tasks capure system, stakeholders, or attackers' actions and behavior. A *resource* is a physical or an informational entity and is used to represent assets.

The relations between actors are captured by the notion of *dependency*. Actors can depend on each other to achieve a goal, perform a task, or furnish a resource. For example, in a goal dependency, an actor (the *depender*) depends on another actor (the *dependee*) to satisfy the goal (the *dependum*). In addition to the dependum, two other intentional elements are involved in a dependency. One element represents *why* a depender needs the dependum, and the other element specifies *how* the dependee satisfies the dependum.

The meta-model in Figure 1 also describes the relationships between intentional elements inside the *boundary* of actors. Actors have (soft)goals and rely on other (soft)goals, tasks, and resources to achieve them. Softgoals can be decomposed into more softgoals using *AND/OR decomposition* relations. *Means-end* links are relations between goals and tasks, and indicate that a goal (the *end*) can be achieved by performing alternative tasks (the *means*). Tasks can be decomposed into any other intentional elements through *task decomposition* links. By decomposing a task into (soft)goals, resources, and other tasks, one can express that those (soft)goals need to be satisfied, sub-tasks need to be performed, and resources need to be available to have the root task performed.

Softgoals and other intentional elements can contribute either positively or negatively to the other softgoals. This is represented through *contribution* links. The contribution relation is characterized by attribute *type* which accepts Help (+), Make (++), Hurt (−), Break (−−) and Unknown (?) values. By linking an intentional element to a softgoal by a Make and Break contribution, one can express that satisfaction of the intentional element is enough to fully satisfy or fully deny a softgoal, while Help and Hurt contributions indicate that the intentional element has positive or negative impact, but the impact is not enough to fully satisfy or deny the softgoal. This qualitative approach for modeling contribution to softgoals reflects the fact that softgoals do not have clear-cut satisfaction criteria.

3.2 Attack and Security Countermeasure Extensions to the Meta-Models

The concepts of vulnerability, attacks, effects of vulnerabilities, and impact countermeasures are added to the i* meta-model. In Figure 1, extended elements to the i* meta-model are highlighed with a different color. Adopting a task or employing a resource can bring *vulnerabilities* to the system. For the sake of simplicity, we call such an intentional element a *vulnerable element*. The concept of vulnerability is not limited to specific reported vulnerabilities or to general classes of vulnerabilities. For example, one can model the famous worm called *2000 ILOVEYOU*[5] or general class of *argument injection or modification*. While resources such as National Vulnerability Database provide a list of specific vulnerabilities in various versions of software products developed by various vendors, Common Weakness Enumeration focuses on gathering classes of weaknesses.

Exploitation of vulnerabilities can have *effects* on the same vulnerable element that has brought the vulnerabilities or on other tasks, goals, and resources. The effect is characterized by attribute *type* which specifies how the vulnerability affects a goal, a task, or a resource. The effect

[5]http://www.cert.org/advisories/CA-2000-04.html

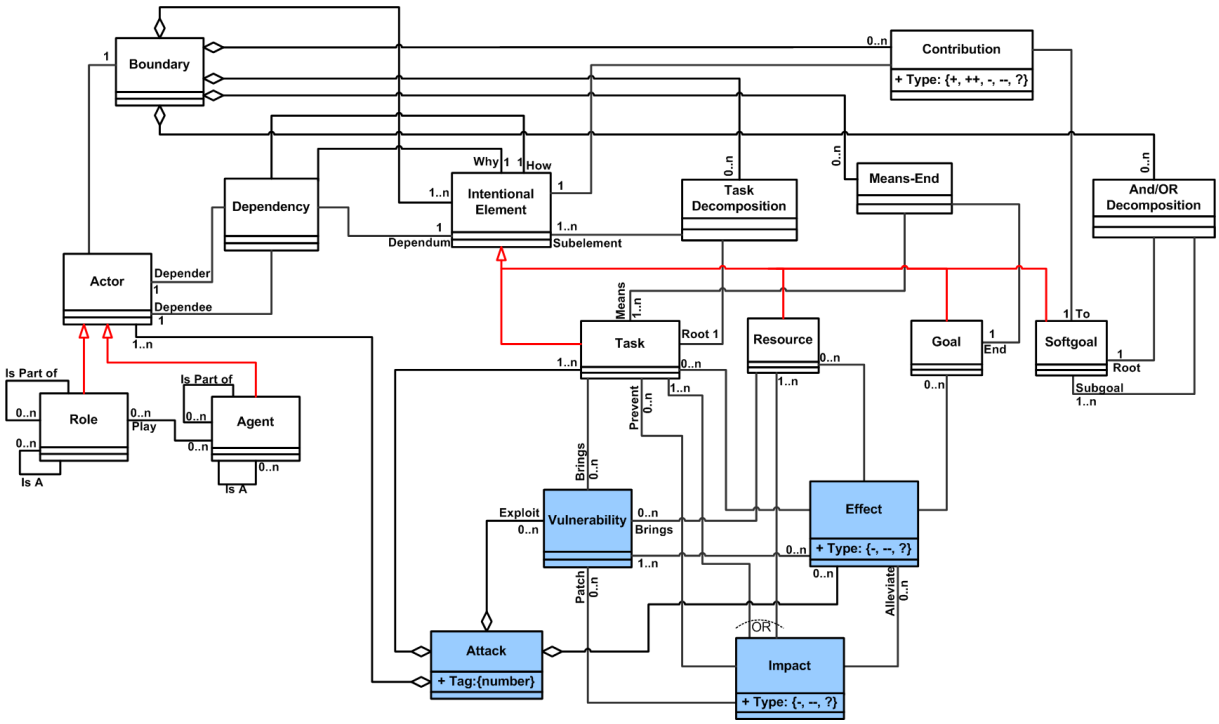

Figure 1: The i* meta-model extended with the concept of vulnerability and attack

types are Hurt (−), Break (−−), and Unknown (?). The effect of vulnerabilities on softgoals is not considered in the meta-model, since softgoals are not directly-measurable goals. The effect of vulnerabilities would be propagated to softgoals from the affected elements that contribute to the softgoals.

The definition of attack and security countermeasure is fundamentally a matter of perspective: a task or a goal counted as malicious can be perceived non-malicious from a different viewpoint. Sequences of actions for mounting an attack are basically similar to sequences of actions performed by legitimate actor. Therefore, the line to differentiate malicious actions from non-malicious ones is arbitrary, and distinguishing malicious goals from non-malicious goals depends on the perspective adopted by the system designer.

Malicious elements have the same semantics of ordinary intentional elements: they can be similar or identical to non-malicious elements. For example, the desire to have a high profit is not a malicious goal, but an actor can achieve such a goal either by working legally and honestly or cheating and frauds. On the other hand, a task can be interpreted as malicious in a condition while it is counted as non-malicious in a different context. For example, one can install a camera for spying into other people privacy, whereas a surveillance camera can be used for security purposes. In this example, the goal for performing tasks indicates if the task is malicious or not.

Other factors can make distinguishing malicious and non-malicious behavior challenging. For instance, the concept of competition, which in general is not malicious, illegitimate, or unauthorized, is similar to the notion of attack because competitions have similar impacts that attacks have. From one organization's perspective, a competitor can be seen as an attacker, and from the competitor viewpoint, the organization is a competitor.

Since malice is a matter of perspective, distinguishing the malicious and non-malicious behavior does not affect the analysis. Therefore, The meta-model presented in Figure 1 is a neutral meta-model that does not distinguish malicious and non-malicious elements. A fragment of this

meta-model is shown in Figure 2. An *attack* is a series of tasks than an actor performs to *exploit* a number of vulnerabilities and has negative effects on other intentional elements. This definition of attacks is based on the definition proposed by Schneider [27] in which vulnerabilities are a key aspect of any attack. This choice is due to the fact that we are mainly interested in analyzing the effects of vulnerabilities on the system. Attacks that are performed without exploiting vulnerabilities can be modeled by introducing a new class of attacks in which their target is a task or a resource instead of a set of vulnerabilities.

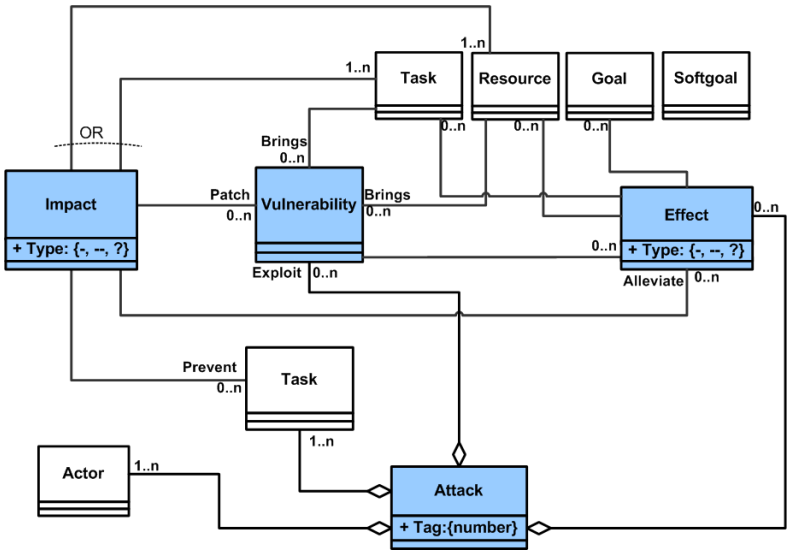

Figure 2: A fragment of the meta-model which focuses on extensions to the i* to express attacks and vulnerabilities. The meta-model does not distinguish malicious and non-malicious behavior.[7]

Resources and tasks can have *impacts* on attacks. Such tasks and resources can be interpreted as security countermeasures; however we do not distinguish them from non-security mechanisms in the meta-model as this distinction does not affect requirements analysis. These impact have an attribute, *type*, which accepts Hurt (−), Break (−−), and Unknown (?) values. Ssecurity countermeasures can be used to *patch* vulnerabilities, *alleviate* the effect of vulnerabilities, or *prevent* the malicious tasks that exploit vulnerabilities or system functionalities. By patching the vulnerability, the countermeasure fixes the weakness in the system. Example of such countermeasure is new updates the software vendors provide for the released products. A countermeasure that alleviates the vulnerability effects does not address the source of the problem, but it intends to reduce the effects of the vulnerability exploitation. For example, a security solution that brings up a backup system in case of working system denial of service alleviates the impact of security failures. Countermeasures can prevent the actions that the attacker performs, which consequently prevents exploitation of the vulnerability by the actions. For example, an authentication solution prevents unauthorized access to assets. Countermeasure may prevent performing system's vulnerable tasks or prevent using vulnerable resources, which results in removing the vulnerability that has been brought to the system by the vulnerable elements. For example, one can disable JavaScript option in the browser to prevent exploitation of malware run by the browser.

As showed by Sindre and Opdahl [17], graphical models become much clearer if the distinction between malicious and non-malicious elements is made explicit and the malicious actions are visually distinguished from the legitimate ones. Sindre and Opdahl show that the use of inverted

[7]The task class has been duplicated to make the meta-model clearer.

elements strongly draws the attention to dependability aspects early on for those who discuss the models. In this regard, an extended meta-model is developed with the assumptions that some actors are attackers and have malicious goals, and other actors employ countermeasures for protecting their goals. Figure 3 presents the extended meta-model, which is derived from the meta-model in Figure 2 by introducing a new type of actor called *attacker* which has malicious intentional elements such as *malicious goals* and *malicious softgoals* inside its boundary. The concept of boundary is added to link the malicious elements to the attacker. An *attack* involves an *attacker*, *malicious tasks* that he performs to exploit a set of *vulnerabilities*, and the *effect* of exploited vulnerabilities on other actors' intentional elements.

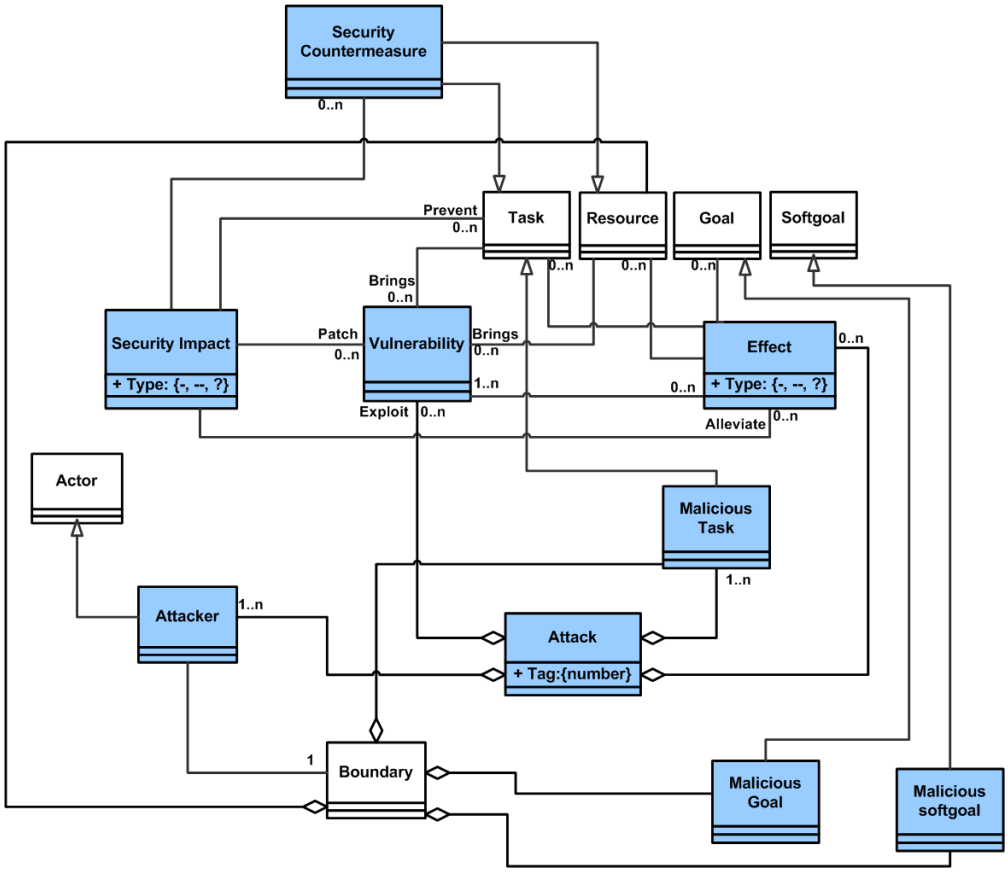

Figure 3: The fragment of the meta-model in which attacks, attackers, malicious behavior, and security countermeasure are explicit elements of the meta-model.

4 Modeling Process

This section presents the security requirements modeling process along with the modeling notation and graphical representation. The resulting models help the analysts to understand the social and organizational dependencies among main stakeholders of the system, their goals, the system architecture, the organization structure [35], and security issues that arise among interaction of actors [9] in the early stages of the development.

Figure 4 summarizes the relationships between the views of the security requirements model and the steps needed to develop those views. The process consists of five steps; each of them re-

sults in a *view* of the security requirements model. Each of these views provides new incremental information:

1. *Requirements view* which captures stakeholders and system actors together with their (soft)goals, the tasks to achieve those goals, required resources, and the dependencies among them.

2. *Vulnerabilities view* which extends the requirements view by adding the vulnerabilities that tasks and resources brings to the system and the impact that their exploitation (or of their combinations) has on the system.

3. *Attacker template view* which captures the behavior of attackers by representing how attackers can exploit vulnerabilities to compromise the system.

4. *Attacker profile view* which captures individual goals, skills, and behavior of a specific class of attackers based on the attacker template view.

5. *Countermeasure view* which captures the security solutions adopted by actors to protect the system and their impacts on attacks and vulnerabilities.

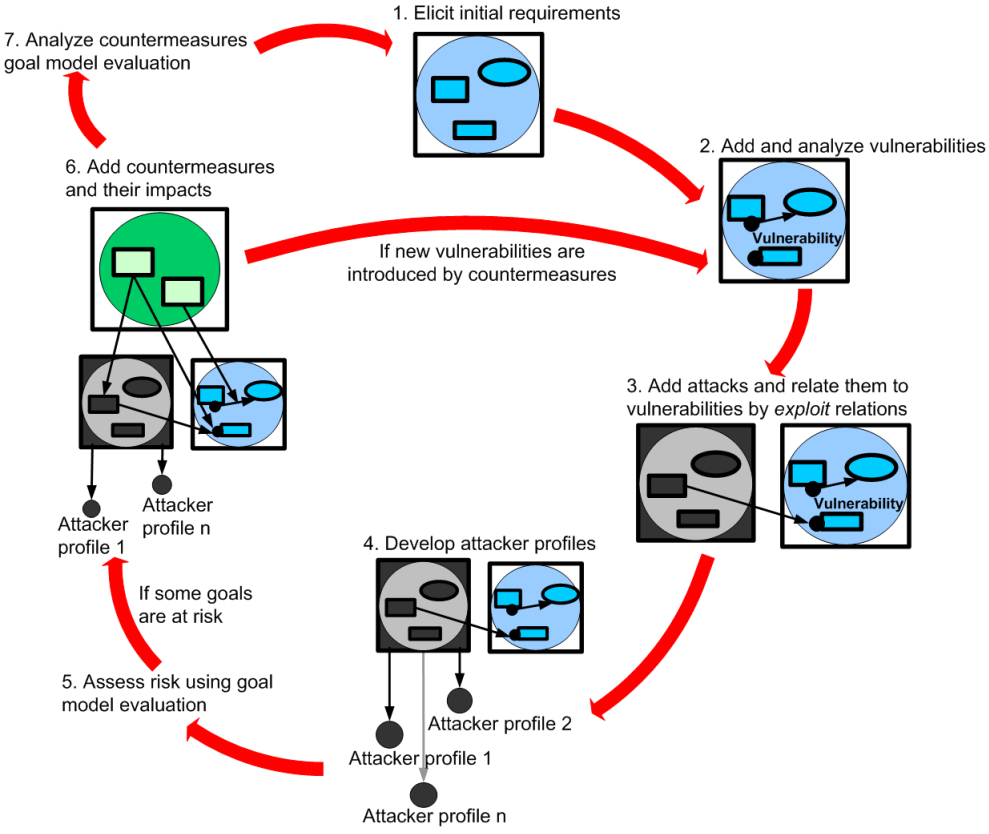

Figure 4: The modeling process

The process for developing security requirements models is incremental: in each step new elements are added to the requirements model to show new aspects. The modeling process starts with the identification of actors, their dependencies, goals, and the tasks and resources necessary to achieve them. Then, the vulnerabilities that tasks and resources bring to the system

are identified and propagated through the goal model. In the third step, possible attacks that can exploit the vulnerabilities are identified and analyzed. Attacker profiles that specify the capabilities and skills of categories of attackers are defined. The model is then evaluated to assess the risk of vulnerabilities exploitation by attackers. If the result of analysis shows that the risks cannot be tolerated by stakeholders, the requirements model is revised by introducing countermeasures and their impacts on vulnerabilities and attacks. Modeling goal, vulnerabilities, attacks, and countermeasures is an iterative process as the adoption of countermeasures may cause the introduction of new vulnerabilities, denial of functionalities or quality goals.

Identification of vulnerabilities, attacks, and countermeasures by the analysts require security knowledge and experience. The proposed framework in this work does not provide guidelines or methods for finding vulnerabilities and attacks and selecting proper countermeasures. The modeling process proposes a way for linking security knowledge such as reports of attacks, list of vulnerabilities, and alternative countermeasures, which are known by security experts.

4.1 Eliciting and Modeling Initial Requirements

Requirements modeling is intended to identify and model stakeholders' needs and system requirements. We take advantage of the i* framework that provides a useful approach for modeling and analyzing stakeholders' and system's goals and system-and-environment alternatives that address achievements of the goals [36, 35]. We do not present the modeling process underlying the i* framework, and details in this regard can be found in [36].

Figure 5 shows the requirements view of a browser which requests the content from a web server to build HTML pages. The *User* depends on a software agent, the *Firefox Browser* to *Browse web sites*. The browser depends on the *User* to *Enter inputs* and depends on *Web server* for *Web page content and JavaScript*. This view describes high level goals and tasks of a browser. For instance, one of the browser's tasks is to *Show the web pages*, and to perform that, the browser needs to *Run the JavaScript with user inputs*. This makes the final customized HTML page, and for this aim, the browser *Request and get pages from the server* and *Get users' input*.

4.2 Modeling and Analyzing Vulnerabilities

Vulnerability modeling intends to understand the weaknesses affecting system requirements. To incorporate specific vulnerabilities or classes of vulnerabilities into the requirements model, we incrementally refine the requirements view by identifying the vulnerable tasks and resources and analyze the effect of vulnerability exploitation. To represent vulnerabilities, the i* modeling notation is enriched with a black circle for the new graphical element. The black circle is chosen to resemble a hole or weakness in the system which leaves a backdoor for attacks. Vulnerabilities are graphically attached to tasks and resources, which implies execution of the task or availability of the resource brings the vulnerability to system. To represent the possible effect(s) of an exploited vulnerability on goals, tasks, and resources a new link is added to the i* relations. The vulnerability effect is visually represented by a dotted line with a label, l, where $l \in \{-, --, ?\}$.

Figure 5 shows the vulnerability view for a browser which requests the content from a web server. One of the browser's tasks is to *Show the web pages*, and to perform that, the browser needs to *Run the JavaScript with user inputs*. The browser *Request and get pages from the server* and *Get users' input*. Each of these tasks bring a vulnerability to the system. By downloading a JavaScript code from the web server, a *Malicious Script* can be downloaded as well. The user inputs can also contain *Malicious input*. As a result, when the browser runs the JavaScript with the user inputs, the browser is exposed to the combination of the *Malicious script* and *Malicious user input* vulnerabilities.

When an actor depends on another actor for a vulnerable task or resource, the vulnerability is carried to the depender actor by the vulnerable dependum. Figure 6(a) explains the propa-

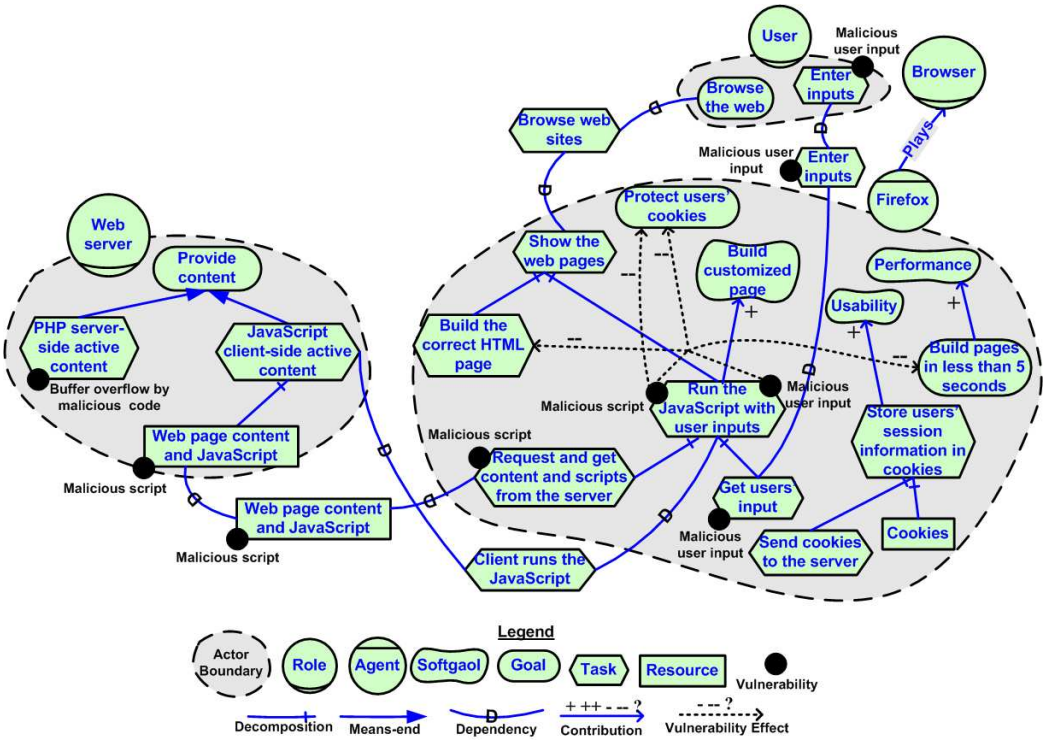

Figure 5: Initial requirements and actors' actions, extended with the vulnerabilities view.

gation of the vulnerability in the reverse direction of dependencies. This figure shows that for dependency relations, the vulnerability V in the dependee's resource $R_{(how)}$ is propagated to the dependum, R_D, and the depender's element, $R_{(why)}$. For example, in Figure 5, the *Malicious Script* vulnerability is brought to the *Firefox* agent because of the dependency link between the browser and *Web server* on *Web page content and JavaScript*. The same argument for resource dependencies and vulnerabilities is valid for task dependencies. For example, when the *Firefox* agent gets the inputs from the user through the dependency link, *Malicious input* is brought to the *Get user input* task of the browser.

Vulnerabilities are also propagated through *decomposition links*. Using decomposition links, analysts refine tasks into more detailed elements with higher resolution information, and the sub-elements describe the up level task in detail. The application of a framework to a number of case studies has shown that it is easier to identify vulnerabilities for concrete sub-elements rather than for high-level abstract ones. Therefore, vulnerabilities are propagated bottom-up from sub-elements to the high level decomposed task.

Figure 6(b) depicts the vulnerability propagation rule through decomposition links. This figure depicts that if a task T_{root} is decomposed into a task T_{child} and a resource R_{child}, respectfully with vulnerabilities V_1 and V_2, the root task would receive both vulnerabilities V_1 and V_2. Vulnerability effects depend on the context of the vulnerable elements. As shown in Figure 6(b), the analyst can either assign the vulnerability effect to the child ($Effect_2$ for V_2) or to the root ($Effect_1$ for V_1) element based on the context. In addition, based on the context, one may determine that the propagated vulnerabilities have a combined effect. Propagating vulnerabilities effects cannot be automatically deducted from the structure of the model and requires human judgement and security experiences.

A concrete example of vulnerability propagation through decomposition links is shown in Figure 5 where the *Run the JavaScript with user inputs* task in decomposed into *Request and*

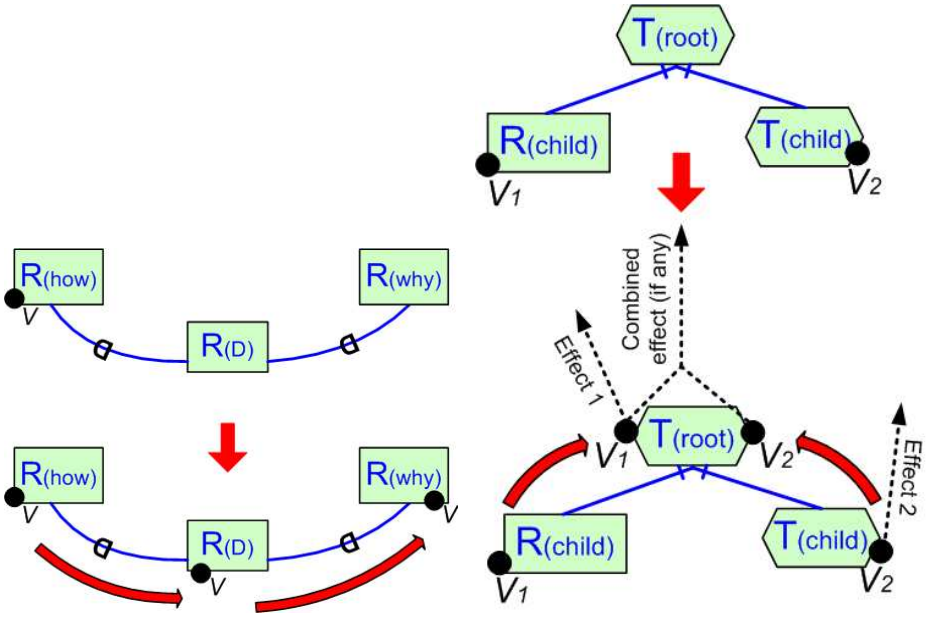

(a) Vulnerability propagation via dependency links

(b) Vulnerability propagation via decomposition links

Figure 6: Vulnerability propagation rules

get content and scripts from the server with *Malicious script* vulnerability and *Get users' input* with *Malicious user input* vulnerability. Accordingly, the root task receives both vulnerabilities. These vulnerabilities or their combination can have various effects on goals and tasks of the actors when running the JavaScript. Figure 7 shows how the vulnerabilities are combined. The effect of exploiting the combination of *Malicious script* in the *malicious user input* vulnerabilities is expressed using the vulnerability effect link with a $--$ (*break*) contribution from the combination of the vulnerabilities to *Protect users' cookies* and *Build the correct HTML page*. In the next section, we describe how the requirements and vulnerabilities views are related to the attacker template and countermeasure views.

4.3 Modeling Attacker Templates

The aims of attacker template modeling is to define a view of the security requirements model that represents the possible ways in which attackers can exploit vulnerabilities to compromise the system and the goals behind these attacks. To build the attack template, designers can take advantages of existing approaches (e.g., attack tree [28] and anti-goals [32]) to develop a tree-like malicious goal model. In addition, catalogs of malicious goals [3] might be useful for driving attacker goals.

As discussed earlier, the modeling notation graphically distinguishes malicious and non-malicious elements using a black shadow in the background of malicious elements as proposed in [19, 9]. The exploitation of a vulnerability by a malicious actor is graphically represented by a link labeled *exploit* from the malicious task to the vulnerability. A vulnerability may have different effects on other goals and mechanisms. Different attacks that exploit a vulnerability may have different effects on other elements. Therefore, to graphically link an attack and the effect of the vulnerability that the attack exploits, the corresponding *vulnerability effect* links for each attack are labeled with the same *tag number* that the exploit link is tagged. In this way, an attack is a quadruple consisting of an *attacker*, the *malicious task* that the attacker performs, a

set of *vulnerabilities*, and their *effect* on the system (see Figure 3).

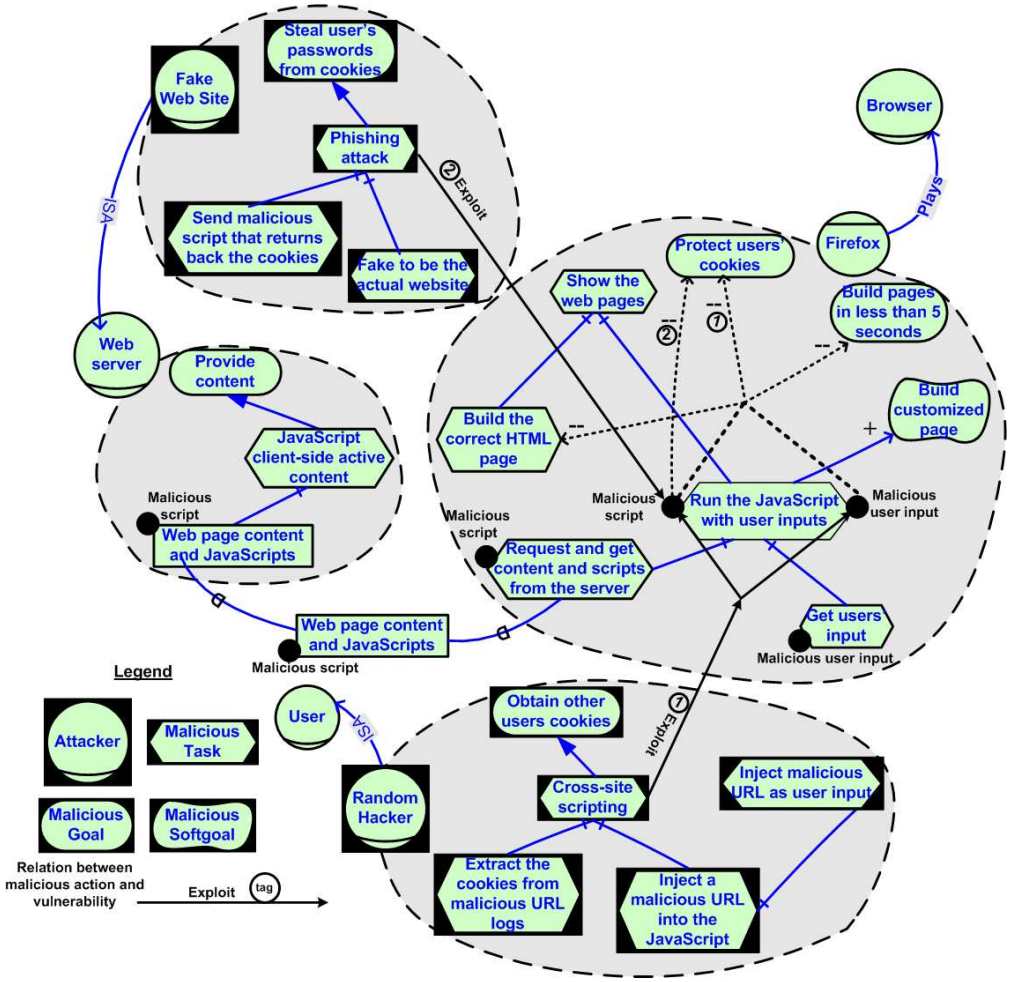

Figure 7: Attacker template view for the browser and web server example.

Figure 7 gives the attacker template view by extending the view in Figure 5 and introducing two possible attackers, *Random hacker* and *Fake Web Site*. *Fake Web Site*, is a *Web server* who intends to *Steal user's passwords from cookies*. The *Fake Web Site* uses *Phishing attack* by exploiting the *Malicious script*. The *Random hacker* is a *User* and inherit users' capabilities. For instance, the *Random hacker* can browse a website and enter inputs. The hacker can use these capabilities for his malicious intents such as *Obtaining other users cookies*. One possible way to obtain cookies of other users is *Cross-site scripting* which consists of *Injecting a malicious URL into the JavaScript* and *Extracting the cookies from malicious URL logs*. To *Inject a malicious URL into the JavaScript*, the hacker *Injects malicious URL as user input* by playing the role of an ordinary user. To specify which malicious task exploits the vulnerability and causes this effect both the exploit and vulnerability effect links are labeled with a tag (number one).

4.4 Identifying and Modeling Countermeasures

By developing requirements, vulnerabilities, and attack template views, analysts have the machinery necessary to evaluate the risks threatening the system. On the basis of risk assessment,

analysts elicit and analyze the security countermeasures needed to protect the system. To model countermeasures, a modeling element is not added to the i* framework, since countermeasures share the same nature with other tasks and resources. Different countermeasures can have a different impact on attacks. A countermeasure can *alleviate* the effect of a vulnerability, *patch* it, or *prevent* malicious tasks or system's functionalities that bring the vulnerabilities. These impacts are modeled through alleviate, patch, and prevent links respectively.

The model in Figure 8 presents the countermeasures for the vulnerabilities and attacks depicted in Figure 7. The countermeasure elements are highlighted using a different color[8]. In Figure 8, the web server employs two security mechanisms: *validate user input* and *Remove HTML tags from use input*. By removing the HTML tags, the malicious code is removed from the user input. This impact is modeled through *prevent* relations between the countermeasures and the malicious task *Inject malicious URL as user input* with "-" label. By validating user input, the *Malicious user input* vulnerability is partially patched.

At the browser side, one can *Disable JavaScript* and use *Anti Phishing tool bar*. Disabling JavaScript prevents performing *Run the JavaScript with user inputs*, hence, the vulnerable task is not performed any more. As a result, the vulnerabilities that are brought by running JavaScript do not exist any more.

4.5 Attacker Profile Definition

Different typologies of attackers may have different capabilities and skills. The idea underlying the attacker profile is to analyze classes of attackers and their behavior against the system. To define capabilities and skills of a class of attackers, the tasks that the attacker can perform, resources that can obtain, and goals that can satisfy are identified and labeled. Intuitively, labels represent the evidences that a goal has been satisfied, a task has been performed, or a resource is available. We refer to Section 5 for details about evaluation labels. Table 1 gives two different attacker profiles for *Random Hacker* and two profiles for *Fake Web Site* introduced in Figure 5. The table indicates which attacker can achieve the tasks by assigning evaluation labels to malicious tasks defined in the attacker template.

Table 1: Attacker profile definition for attacker templates in Figure 5.

Malicious Task	Random Hacker (1)	Random Hacker (2)	Fake Web Site (1)	Fake Web Site (2)
Send malicious script that returns back the cookies	D	D	S	S
Fake to be the actual website	D	D	S	D
Inject a malicious URL into the JavaScript	S	S	D	D
Inject a malicious URL into the JavaScript	S	S	D	D
Extract the cookies from malicious URL logs	S	D	D	D

[8]The highlighted color in the models does not bear any semantical significance and only intends to highlight the countermeasures in the figures.

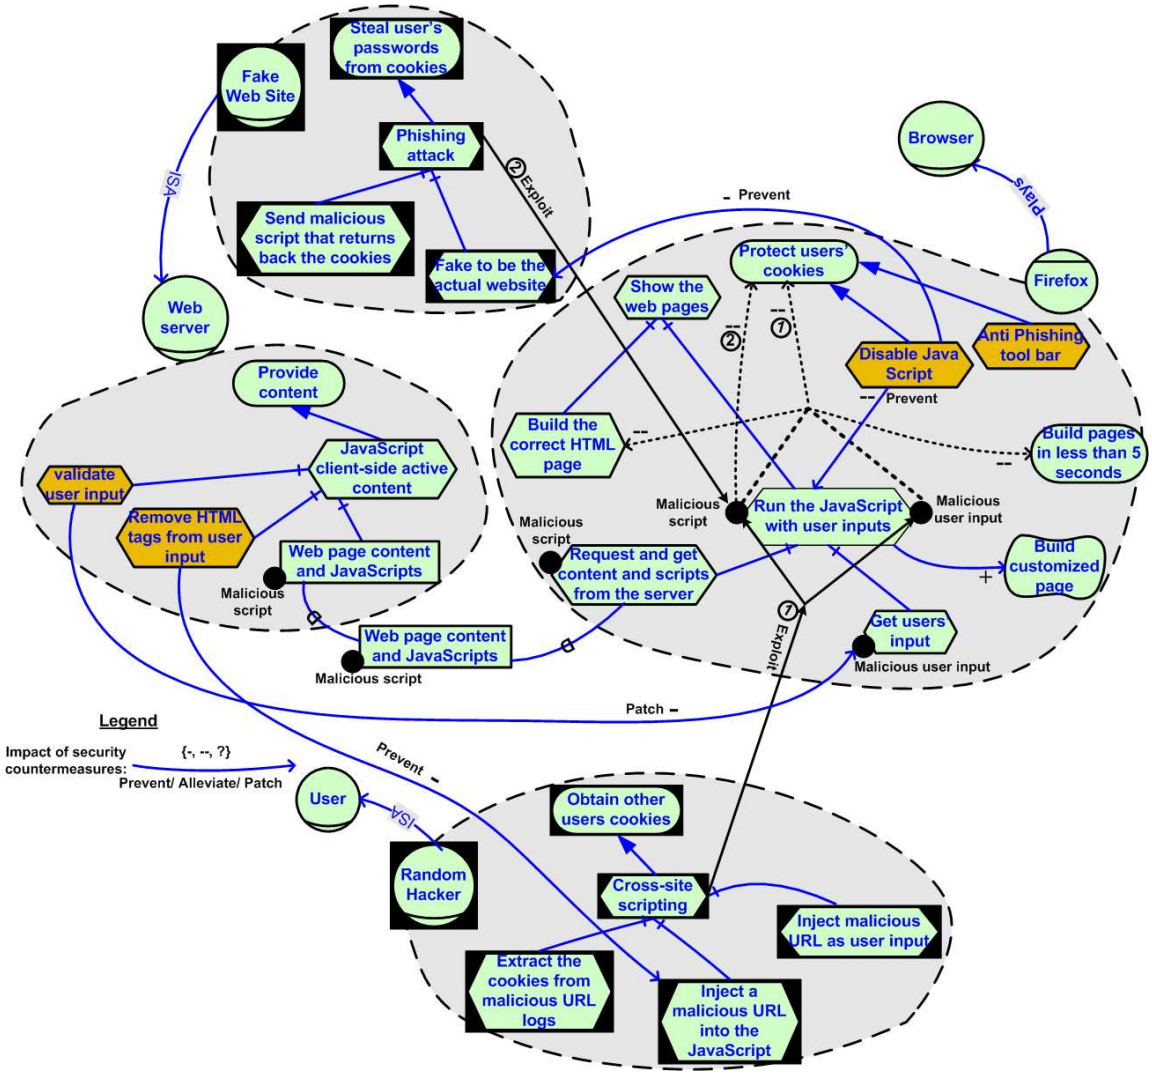

Figure 8: The countermeasure view for the web server and browser example. (The elements with highlighted color are countermeasures)

5 Security Requirements Analysis using Goal Model Evaluation

In addition to the benefits that analysts gain through the modeling process, goal models including vulnerabilities, attacks, and countermeasures provide a basis for security requirements analysis. The purpose of the evaluation is to assess the risks due to attacks for determining the countermeasure necessary to protect the system. While traditional risk analysis methods assess risks by considering quantitative probability and severity of successful attacks [29], we propose analyzing risks by evaluating satisfaction or denial of goals of system and stakeholders. For this purpose, we take advantage of qualitative goal model evaluation techniques. Although a quantitative risk assessment approaches can greatly simplify decision making and provide accurate final results, it can be difficult to apply due to lack of agreed metrics of vulnerabilities and accurate measures, specially in the early stages of development. On the other hand, qualitative evaluation

answers questions with lower resolution information, represented in a qualitative spectrum.

Goal model evaluation is the procedure to verify that actors' top level goals are satisfied by the choices that actors have made and consists of propagating denial or satisfaction evidences through the goal model using a set of rules [6]. Horkoff [13] proposes an i* goal model evaluation method where denial or satisfaction labels are assigned to the leaf nodes and then are propagated through the goal model based on the type of the links between the elements. Evidences for satisfaction or denial are in the scale of values, which ranges from full satisfy (S), partial satisfy (PS), unknown (?), and conflict (C) to partial deny (PD) and full deny (D) which intended order of $S > PS > C >? > PD > D$. In case of conflicts, human judgment is required to resolve the conflicts of contributions during the evaluation process.

In this work, we have enhanced the i* framework with new concepts such as vulnerability, effects of vulnerabilities, malicious elements, countermeasures, and impacts of countermeasures. These new concepts and relations need to be accommodated into the existing evaluation method. To this end, we have adopted and adapted the goal evaluation method in [13]. The syntactical and semantical definition of the new modeling constructs requires revising the evaluation methods accordingly. Employing countermeasures introduces direct contributions of tasks and resources to malicious and non-malicious tasks, vulnerabilities, and effects of vulnerabilities. A security solution may partially prevent a malicious task. For example, increasing the buffer size can only partially addresses buffer overflow unless the buffer size is made infinite. This impact may cause the propagation of partial satisfaction and denial values to hard elements such as tasks and goals. In contrast, the algorithm in [13] assumes that intentional elements contribute only to softgaols. This implies that partial values are only assigned to softgoals. In the current work, we relax this assumption by considering partial values to be assigned also to goals, tasks and resources. Partial values associated with those elements indicate that there is partial evidence of their satisfaction or denial. Labels are also associated with vulnerabilities to represent the evidences that they are exploited or not.

Figure 9 summarizes the security goal model evaluation steps. The evaluation starts with assigning labels to the leaf nodes of both malicious and non-malicious actors. In particular, the labels associated with malicious elements are defined on the basis of the selected attacker profile. Then, the labels are propagated to the upper nodes using propagation rules. The final result of evaluation shows the consequences of attacks and exploitation of vulnerabilities on higher goals. If some stakeholders and system goals are denied because of attacks, analysts need to consider countermeasures and analyze their impact on the system behavior and security. Evaluation labels are propagated through the goal model again to verify the satisfaction or denial of the goals because of the employment of countermeasures. In what follows, we present the propagation rules underlying the security goal evaluation method.

5.1 i* Propagation Rules

This section briefly presents the propagation rules for i* concepts (i.e., dependency, contribution, decomposition, AND/OR, and means-end) based on the work in [13]. In a dependency relation, satisfaction of the depender element relies on the satisfaction of the dependee element. Therefore, dependency links propagate the value of the dependee element to the dependum and then to the depender element. Differently from [13], we do not define rules for dealing with ambiguous scenarios. Analysts should avoid drawing goal models that contains ambiguity such as the ones presented in the left side of Figure 10. Analysts shall revise the models in order to disambiguate the model. The right side of Figure 10 shows some examples of unambiguous scenarios.

Contribution links represent the impact of intentional elements only on softgoals. Table 2 presents the rules for propagating the impact of an intentional element with denial or satisfaction value to the target softgoal through contribution links. However, since we relax the restriction

[9] Softgoal_d, Goal_d, and Task_d are the dependum elements.

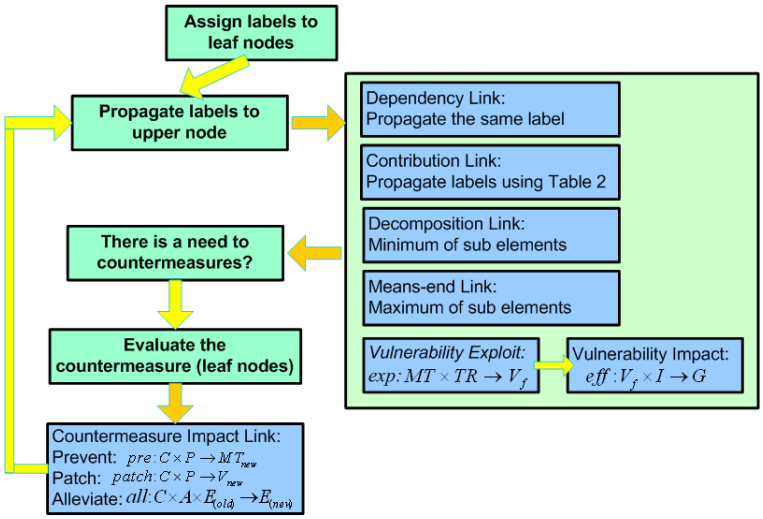

Figure 9: Security goal model evaluation steps

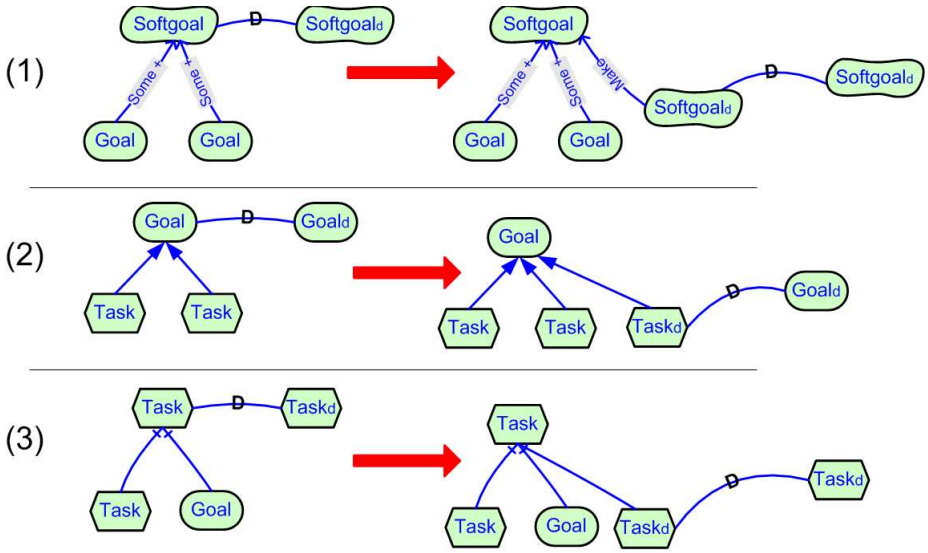

Figure 10: Ambiguous dependency links which analysts need to avoid[9]

of contribution links in the i* syntax, the rules in Table 2 are valid for contribution links to hard elements as well. To propagate the evaluation values from multiple elements to a softgoal, the rules in Table 2 are not enough. Because while one element can make the softgoal satisfied, another element may make the softgoal denied entailing conflicting evidences. In this work, we follow the approach suggested in [13] where human judgment is required to resolve conflicts.

AND/OR links refines a softgoal into one or more softgoals. When a softgoal is refined using AND links, the *minimum* value among the sub softgoals is propagated to the higher softgoal. When the softgoal is refined using OR links, the *maximum* value among the sub softgoals is propagated to the top softgoal. A decomposition link refines a task into one or more intentional sub-elements. To perform the higher task, all sub-elements need to be satisfied. Accordingly, the label that is propagated to the higher task through the decomposition link is the *minimum* value among the values associated with the sub elements. Means-end links specify alternative tasks to

Table 2: Propagation rules for contribution links

Source Label	Contribution Link Type				
Label Name	++	+	−−	−	?
Satisficed (S)	S	PS	D	PD	?
Partially Satisficed (PS)	PS	PS	PD	PD	?
Conflict (C)	C	C	C	C	?
Unknown (?)	?	?	?	?	?
Partially Denied (PD)	PD	PD	PS	PS	?
Denied (D)	D	PD	PS	PS	?

satisfy a goal. Since those tasks are alternative ways to achieve the higher level goal, means-end links work as OR relations: the *maximum* value among alternative tasks is propagated to the goal.

5.2 Vulnerability Exploit and Effect Propagation Rules

Tasks and resources may bring vulnerabilities to the system, and attackers exploit them to compromise the system. In order to propagate the satisfaction or denial labels of malicious tasks to other elements, vulnerability is treated as a *filter*: when the filter is open, a backdoor to the system is open for attackers. To determine if the filter is open or not, both the vulnerable element and the malicious task that exploits the vulnerability need to be analyzed. If the vulnerable task is not executed, the vulnerability cannot be exploited. Similarly, if the resource that brings the vulnerability to the system is not available, then the vulnerability does not exist within the system. At the same time, a vulnerability is exploited if the attack succeeds, indicating that the malicious task has been performed. To determine the exploitation condition of a vulnerability (i.e., if it has been exploited or not), we introduce the function

$$exp: \ MT \times TR \longrightarrow V_f$$

where MT is the evaluation value associated with the malicious task, TR is the value associated with the task or resource that has brought the vulnerability, and V_f represents the exploitation condition of the vulnerability. The evaluation value for V_f is calculated as $V_f = min(MT, TR)$. The satisfaction label (S) for the V_f means that the vulnerability is fully exploited, i.e. the filter is completely open. The denial label (D) indicates that the vulnerability is not exploited and the backdoor to the system is completely closed. Partial labels for V_f indicate that there exist partial evidences about the condition of the filter.

The effect of the vulnerability on other intentional elements is computed on the basis of the exploitation condition of the vulnerability and severity of it effects on the system. For instance, if a vulnerability has not been exploited, its negative effect would not be propagated to other elements. For this purpose, we employ the function

$$eff: \ V_f \times I \longrightarrow E$$

where V_f is the exploitation condition of the vulnerability, I represents severity of the vulnerability effect, and E represents the evidences about the satisfaction or denial of the intentional elements that is affected by the vulnerability. This function shares the same intuition of the propagation rules described in Table 2. In the case where an attack exploits the combination of two or more vulnerabilities, human judgment is required to evaluate the function eff for several combined vulnerability effects.

5.3 Countermeasure Impacts Propagation Rules

Countermeasures can have three different security impacts: they can be used to prevent execution of a task, achievement of a goal, or the availability of a resource; patching vulnerabilities; or alleviating their effects. The propagation of the impact of a countermeasure through a *prevent* link depends on the successful employment of the countermeasure as well as on its efficacy. To evaluate the final impact of the countermeasure on the target element, we introduce the function

$$pre: \ C \times P \longrightarrow E_{(new)}$$

where C is the evaluation label associated with countermeasure, P is the type of the prevent relation, and $E_{(new)}$ is the new evaluation value of the target intentional element affected by the *prevent*. A prevent relation shares the same nature with contribution relations, and accordingly, the function *pre* uses the propagation rules defined in Table 2.

When a countermeasure patches a vulnerability, the vulnerability exploitation condition is modified. The objective of patching a vulnerability is to make the filter closer and consequently to reduce the impact of the attack that exploits the vulnerability:

$$patch: \ C \times P \longrightarrow V_{(new)}$$

where C is the countermeasure evaluation label, P is the contribution type of the patch relation, and $V_{(new)}$ is the new value to be associated to the vulnerability. Similarly to prevent relations, a patch relation shares the same intuition with a contribution link, and the corresponding propagation rules are similar to the ones defined in Table 2. However, we assume that countermeasures only reduce the risks and do not magnify vulnerabilities or attacks. Therefore, propagation rules apply in cases where the countermeasure is partially or fully satisfied, and the impact is not propagated if the countermeasure is partially or fully denied. Once the impact of a countermeasure is propagated through the patch link, the new exploitation condition of the patched vulnerability needs to be propagated to other instances of the vulnerability through decomposition and dependency links.

Finally, a countermeasure may alleviate the effect of an exploited vulnerability. In this case, the contribution value of the alleviate link is combined with the countermeasure evaluation value and the current effect of the vulnerability. This can be represented by the function

$$all: \ C \times A \times E_{(old)} \longrightarrow E_{(new)}$$

where C is the countermeasure evaluation label, A is the contribution type of the alleviation relation, and $E_{(old)}$ and $E_{(new)}$ are the contribution of the vulnerability effect before and after applying the countermeasure respectively. Table 3 defines the rules used by function *all* to compute the $E_{(new)}$ value.

Table 3: Propagation rules for alleviate links

Countermeasure label	Alleviate Contribution	Old Vulnerability Effect	New Vulnerability Effect
S	-	-	-
PS	- -	- -	-
S	- -	- -	? (no impact)
PS	- -	- -	-
S	-	- -	-
PS	-	- -	-
S	- -	-	? (no impact)
PS	- -	-	-

5.4 Evaluation Example

The first step of the evaluation is assignment of evaluation labels to the leaf nodes of malicious and non-malicious actors. Attacker profiles are used for the assignment of evaluation labels to the leaf nodes of malicious actors. Figure 11 shows the result of label propagation on a fragment of Figure 7. The steps of propagation are depicted by the tag numbers assigned to the evaluation labels of each element. After the initial label assignment, labels are propagated to the upper nodes. For example, the *Malicious script* vulnerability attached to *Run the JavaScript with user inputs* task is fully exploited because the vulnerable and malicious tasks that exploit it are fully satisfied. The exploitation condition of the vulnerability together with its effect makes *Protect users' cookies*, fully denied.

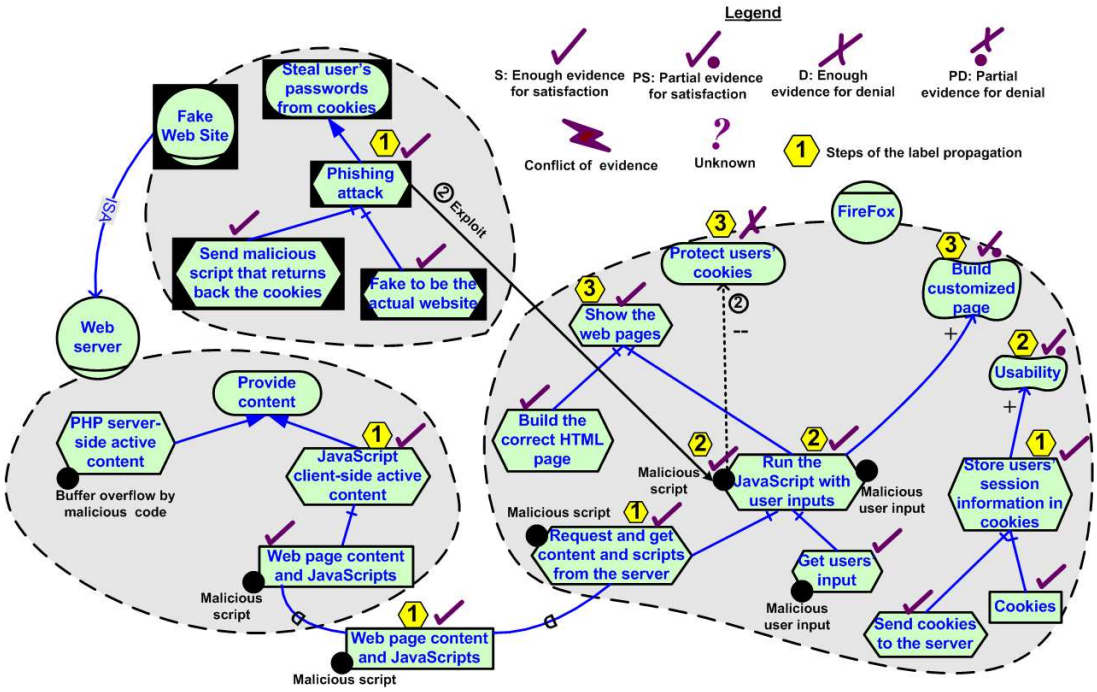

Figure 11: Propagation of the evaluation labels through the attacker template view

In Figure 12, two alternative countermeasures, *Disable JavaScript* and *Anti Phishing tool bar*, are added to the system to analyze their impacts on the system security. Assuming that the user *Disables JavaScript* option, the evaluation process continues by propagating the impact of this countermeasure to the *Run the JavaScript with user inputs* task. As a result, the task that brings the vulnerability is fully denied, and the impact of the vulnerability is not propagated to *Protect users' cookies* goal. On the other hand, the *Build customized page* softgoal is partially denied.

6 Related Work

Security requirements intend to protect the system against threats and prevent the exploitation of vulnerabilities by attackers. Security Requirements Engineering should thus provide techniques for modeling and analyzing attacks, attackers and vulnerabilities, and eliciting countermeasures in order to protect the system. In this section, we overview the current state of the art in threat analysis and security requirements engineering.

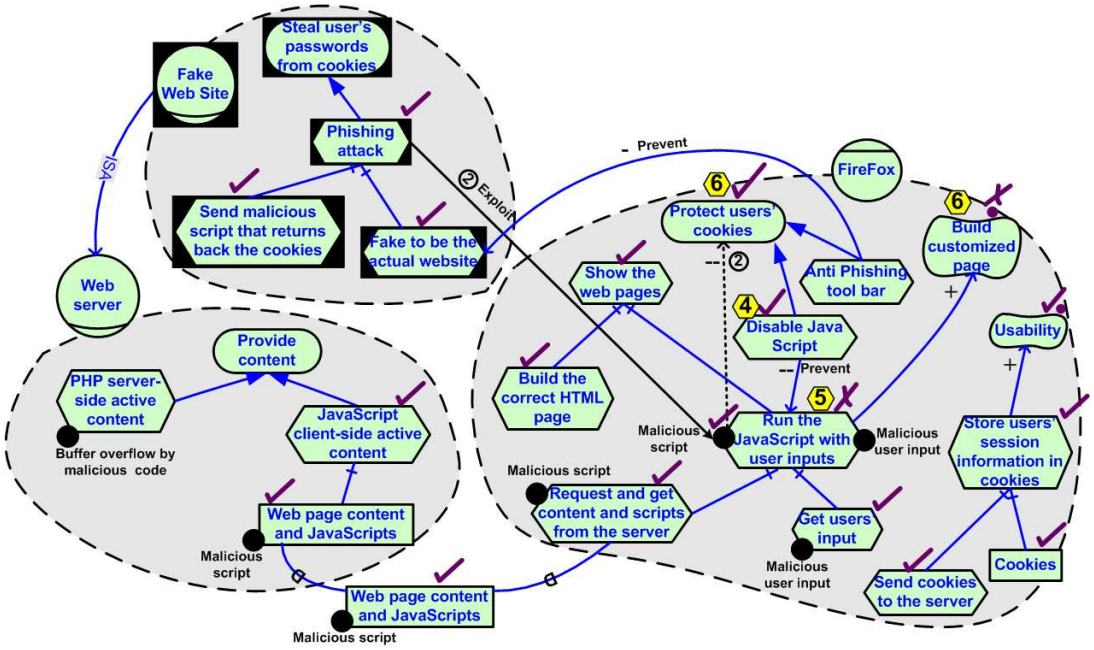

Figure 12: Propagation of the evaluation labels through the countermeasure view

6.1 Threat Analysis

In the security engineering, various modeling techniques have been proposed to analyze the system from the perspective of attackers [22, 25, 28, 34]. Schneier [28] proposes attack tree as a formal and methodical way for analyzing attacks. The root node of an attack tree is the goal of the attacker that is refined using AND/OR relations to understand the possible alternatives used by the attacker to achieve his goal. Schneier argues that designers can take advantages from the understanding of all possible alternatives in which a system can be attacked, who the attackers are, their abilities and motivations in order to select appropriate countermeasures against those attacks. Attack trees can be also annotated with properties of attackers (e.g., skill, access, risk aversion, etc.) and labels representing the cost or probability of achieving a goal. Such properties allow designers to analyze the behavior of classes of attackers by focusing on a certain parts of the attack tree. Attack trees, however, are not linked to other development artifacts, such as design, architecture, and requirements specifications.

Fault Trees Analysis (FTA) [34] is one of the most commonly-used techniques in reliability engineering. The main goal is to assess the likelihood of system failures based on the likelihood of external events. Fault trees visually model logical relationships among infrastructure failures, human errors, and external events that could lead to the system failure. Although FTA enables modeling faults and tracing them to events or errors, it does not provide means to express vulnerabilities of the system and link attacks to them. Moreover, FTA does not support the analysis of the impact of countermeasures on the system.

McDermott [22] proposes to model attack nets as Petri Nets where places represent states or modes of the security-relevant entities within the system, and transitions represent input events, commands, or data that cause one or more security relevant entities to change state. Attack steps are represented by places, transitions are used for the explicit modeling of attacker actions, and tokens are used to indicate the progress of the attack. However, Petri Net models do not support the modeling and analysis of vulnerabilities and countermeasures. The attack model is not linked to other development artifacts, which is an obstacle to elicit security requirements

and design the architecture. Although it is possible to express attackers, their goals, skills, and capabilities cannot be expressed.

Phillips et al. [25] introduced attack graphs to analyze vulnerabilities in computer networks. Attack graphs provide a method for modeling attacks and relating them to the machines in a network and to the attackers. The proposal is based on attack templates, attack profiles, and network configurations. Attack templates describe generic steps in known attacks and conditions which must be hold. The underlying idea is to match the network configuration, attacker profile, and attack templates to generate the attack graph. An attack graphs is an attack template instantiated with particular attackers/users and machines. Thereby, one can analyze an attack graph by identifying the attack paths that are most likely to succeed. Although attack graphs are able to model the steps of an attack, post and pre conditions, required configurations, and capabilities, they do not express the impact of the attacks on system functionalities.

The CORAS project [8] proposes a modeling framework for model-based risk assessment in the form of a UML profile. The profile defines UML stereotypes and rules to express assets, risks targeting the assets, vulnerabilities, accidental and deliberate threats, and the security solutions. In addition to the UML profile, CORAS defines a methodology based on the Unified Process for risk assessment. The analysis method consists of analyzing the target context by developing asset and threat models. The potential attackers, who impose the risks, and vulnerabilities that are exploited by attackers are identified. Risks are prioritized with respect to their severity and likelihood, and the treatments for those risks that are not acceptable are identified. CORAS provides a way for expressing how a vulnerability leads to another vulnerability and how a vulnerability or combination of vulnerabilities lead to a threat. CORAS also provides facilities to relate treatments to threats and vulnerabilities. However, it does not investigate which design choices, requirements, or processes has brought the vulnerabilities to the system.

6.2 Security Requirements Engineering

In recent years, the necessity of considering security from the early phases of the software development process has been recognized. To address this need, traditional requirements engineering framework has been adopted and adapted to support the modeling and analysis of security requirements. Van Lamswerde extended KAOS [7], a goal-oriented security requirements engineering methodology, by introducing the notions of obstacle to capture exceptional behaviors [33] and anti-goal to model intentional obstacles set up by attackers to threaten security goals [32]. Anti-goals are defined as the negation of application-specific instances of generic security goals such as confidentiality, availability, and privacy. Basically, anti-goals represent the goals of attackers. Anti-goals are then refined to form a threat tree on the basis of attackers' goals and capabilities as well as software vulnerabilities. The leaf nodes are either software vulnerabilities or anti-requirements (i.e., anti-goals that are realizable by some attacker). Security requirements are defined as the countermeasures to software vulnerabilities or anti-requirements. The framework does not consider assets as a main concept for eliciting and elaborating security requirements. In addition, vulnerabilities are identified as part of the anti model, while vulnerabilities exist independent of the threats.

A number of approaches based on i*/Tropos [35, 5] have been proposed to address different security aspects. Liu et al. [19] propose to explicitly model the relationships among strategic actors in order to elicit, identify, and analyze security requirements. All actors are assumed potential attackers, which inherit capabilities, intentions, and social relationships of the corresponding legitimate actor. The framework attempts to identify the vulnerable points in the dependency network when an actor behaves maliciously and to understand the measures necessary to protect the system. Attacker identification however is limited to analyzing what roles in the system can impose threat on the dependencies and ignores external attackers. Moreover, the approach does not explicitly describe how countermeasures need to be incorporated into the model and what are their impacts on attacks and other goals. A different perspective has been adopted by

Giorgini et al. [10] who define Secure Tropos for modeling and analyzing authorization, trust and privacy concerns. Secure Tropos extends Tropos with concepts specific to security, namely ownership, permission, delegation, and trust. Secure Tropos, however, addresses security issues within the organization setting rather than dealing with malicious actors.

Some proposals focus on integrating risk analysis into the requirements engineering mainstream. Asnar et al. [2] propose Goal-Risk (GR) Tropos, which extends Tropos with three basic layers: strategy, event, and treatment. The strategic layer analyzes strategic interests of the stakeholders; the event layer analyzes uncertain events along their impacts to the strategy layer; and treatment layer analyzes treatments to be adopted in order to mitigate risks. However, the framework mainly concerns the development of safety critical system and does not consider the intentionality of attackers. Mayer et al. [23] analyze the impacts of risks on business assets and elicit the security requirements for mitigating the risks. In this work, risks are related to the threats and vulnerabilities in the architecture. However, threats are not assigned to an actor and vulnerabilities are not attached to the assets or actions that bring the vulnerability to the system; also threats and vulnerabilities are not related to each other and the impact of countermeasures on the vulnerabilities is not analyzed. Matulevicius et al. [20] improves the Secure Tropos [24][10] modeling language for risk management purposes, where risk is defined as the combination of a threat with vulnerabilities leading to negative impacts on assets. Vulnerabilities are treated as beliefs inside the boundary of attackers which may contribute positively to the successful of an attack. However, similar to the CORAS framework [8], the resulting model does not specify how the vulnerability is brought to the system, by what actions, and by what actors. In addition, the enhanced Secure Tropos models do not capture the impact of countermeasures on the vulnerabilities and attacks.

Haley et al [12] propose a security requirements framework based on constructing the system context, representing security requirements, and developing satisfaction arguments for those requirements. The framework extends problem frames and intends to determine adequate security requirements for the system by considering threats as crosscutting concerns. Functional requirements describe how assets (i.e., objects to be protected) are used within the system, and threats describe how attackers can compromise the security of assets. Security requirements are thus defined as constraints on functional requirements. Once security requirements are elicited, satisfaction arguments are used to verify that security requirements are satisfied by the system as described by the context. This proposal, however, mainly focuses on system requirements and does not provide methodological support for the analysis of the organizational context where the system will operate.

In the UML community, Sindre and Opdahl [31] propose analyzing security requirements by defining misuse cases, inverted UML use cases, which describe functions that the system should not allow. They are depicted as black ovals to distinguish them from traditional use cases. Misuse cases can be linked to use cases to indicate that the use case is exploited by the misuse case, and use cases to misuse cases to indicate that the use case is a countermeasure against the misuse case. This new construct makes it possible to represent actions that the system should prevent together with those actions which it should support. A similar proposal is defined by McDermott and Fox [21], who introduce abuse cases to specify the interactions that their results are harm to system. Differently form misuse cases, abuse cases are distinguished from use cases by representing them in separated models. This does not allow one to analyze the impact of an abuse case on use cases. The security requirements elicitation process underling abuse and misuse cases does not consider why and how security goals are defined without analyzing what may threaten the assets.

Jürjens proposes UMLsec [16], a UML profile designed to express security relevant information within UML diagrams. UMLsec objectives are to encapsulate knowledge and make it

[10] In security requirements literature, two different frameworks developed by different researchers are called Secure Tropos [24, 10].

Table 4: Comparison of modeling notations. N indicates that the concept or relation is not considered, and Y indicates the relation is considered explicitly in the notation. P means the relation is implicitly considered or its semantics is not well defined.

Security Requirements Modeling Framework	Vulnerability graphical representation	Relation to vulnerable elements	Effects of vulnerabilities	Relation of vulnerabilities and attacks	Countermeasure's impact on vulnerabilities
CORAS Framework [4]		N	Y	P	P
Secure Tropos by Matulevicius et al. [20]		N	Y	P	N
Risk-Based Security Framework by Mayer et al. [23]		N	P	Y	N
Security extention on i* framework by Elahi et al. [9]		Y	N	P	N
Current proposal for vulnerability-centric requirements engineering		Y	Y	Y	Y

available to developers in the form of a widely used design notation, and to provide formal techniques for the verification of security requirements. The profile is described in terms of UML stereotypes, tags, and constraints that can be used in various UML diagrams such as activity diagrams, statecharts, and sequence diagrams. The stereotypes and tags encapsulate the knowledge of recurring security requirements of distributed object-oriented systems, such as secrecy, fair exchange, and secure communication link. Assigning a stereotype and tag to the model and defining potential threats make it possible to analyze the behavior of subsystems in order to verify the security impact of threats on the system. The defined security requirements are high level and general. This increases the reusability of the extensions in various contexts. On the other hand, the default threats may not hold in every context.

6.3 Discussion and Comparision

Few security modeling notations provide explicit constructs for modeling vulnerabilities and analyzing their impacts on security requirements. As mentioned before, CORAS framework models and analyzes vulnerabilities by linking them to the threats and risks. In [20], vulnerabilities are modeled as beliefs of attackers which may contribute possitively to the attacks. In [23], i* is extended to represent vulnerabilities and their relation with threats and other elements of the i* goal model.

The missing point in these approaches is providing modeling constructs to understand why vulnerabilities are within the system and how they are spread out among the actors. To our knowledge, existing modeling notations do not provide means to assign vulnerabilities to the actions that actors perform or assets they use. The harmful effects of vulnerabilities on stakeholders and system requirements are not expressed by existing proposals. Among the modeling notations that provide explicit constructs for the concept of vulnerability, CORAS is the only method that relates the countermeasures to vulnerabilities. However, CORAS and other vulnerability modeling notations do not analyze the impacts of countermeasures on vulnerabilities. Table 4 compares capabilities and limitations of existing modeling notations and the new proposal in this paper for modeling and analyzing vulnerabilities as part a security requirements engineering framework.

7 Conclusions and Future Work

This paper proposes a requirements engineering framework to support the elicitation of security requirements based on the vulnerabilities that requirements and design decisions bring to the system. The framework comprises of a modeling framework that extends i* with the concept of vulnerability and relations that allow modeling and understanding effects of vulnerabilities on security requirements. Security requirements are expressed in the form of countermeasures to be adopted to prevent attacks, patch vulnerabilities, or alleviate their effect. Together with a modeling notation, the framework provides an evaluation method for assessing vulnerabilities risks and countermeasures efficacy.

This work is still in progress to better support system designers in modeling and analysis of security requirements. Goal models and the i* modeling notation do not capture temporal aspects of the systems. Therefore, the resulting models do not provide a temporal sequence to guide reading and understanding the model for the analysts that view the large and complicated models. A major limitation of the proposed approach concerns scalability issues. Goal models contain multiple actors and dependency chains, and each actor includes several intentional elements and complicated relationships. The resulting models, especially extended with security concepts, can be complicated and hard to understand. This requires development of modeling and analysis tools that provide resolution management and handle the model complexity by providing views of the security requirements model. To manage the complexity of the models, one can filter vulnerabilities that are not exploited as well as their effects. Analysts would benefit from views that focus on a specific attack, vulnerability, or countermeasure of importance, which cuts some other elements out of the model.

The proposed framework assumes that analysts have knowledge about vulnerabilities, potential attacks, and proper countermeasure or can obtain such information. In particular, the analysis of vulnerabilities, such as propagating them through the goal model or identifying their impacts require experience with vulnerable software products and services. However, existing vulnerability databases do not provide the required knowledge for linking vulnerabilities to actions and assets of the system actors and to potential attacks and countermeasures. Therefore, software developers without security expertise may need additional support for applying the proposed framework.

To address these issues, we are building catalogs of attacker templates that defines the behavior of attackers and catalogs of countermeasures that describe how security flaws are addressed in current security practices. The attacker catalogs assume limited skills and capabilities for the attackers and analyze the actions that they can perform for compromising the system in detail. The attacker templates are then instantiated using attacker profiles to study the behavior of particular classes of attackers. We are also developing security metrics based on the risk attitude of system designers and stakeholders as well as on specific application domains to assist designers during the decision making process.

Finally, we are currently performing empirical studies for evaluating the expressiveness of the proposed modeling notation and the accuracy of the analysis method. Human subjects are being asked to use the proposed framework for modeling and analyzing a number of case studies. The modelers are interviewed and models are critically analyzed to draw conclusions about the practical usefulness and expressiveness of the approach.

Acknowledgment

Financial support from Natural Science and Engineering Research Council of Canada and Bell University Labs is gratefully acknowledged.

References

[1] R. Anderson. *Security Engineering: a guide to Building dependable Distributed systems*. John Wiley and Sons, 2001.

[2] Y. Asnar, R. Moretti, M. Sebastianis, and N. Zannone. Risk as Dependability Metrics for the Evaluation of Business Solutions: A Model-driven Approach. In *Proc. of DAWAM'08*, pages 1240–1248. IEEE Press, 2008.

[3] A. Avizienis, J.-C. Laprie, B. Randell, and C. E. Landwehr. Basic Concepts and Taxonomy of Dependable and Secure Computing. *TDSC*, 1(1):11–33, 2004.

[4] F. Braber, I. Hogganvik, M. S. Lund, K. Stolen, and F. Vraalsen. Model-based security analysis in seven steps — a guided tour to the coras method. *BT Technology Journal*, 25(1):101–117, 2007.

[5] P. Bresciani, P. Giorgini, F. Giunchiglia, J. Mylopoulos, and A. Perini. TROPOS: An Agent-Oriented Software Development Methodology. *JAAMAS*, 8(3):203–236, 2004.

[6] L. Chung, B. A. Nixon, E. Yu, and J. Mylopoulos, editors. *Non-Functional Requirements in Software Engineering*. Kluwer Academic Publishing, 2000.

[7] A. Dardenne, A. van Lamsweerde, and S. Fickas. Goal-directed Requirements Acquisition. *Sci. Comput. Programming*, 20:3–50, 1993.

[8] F. den Braber, T. Dimitrakos, B. A. Gran, M. S. Lund, K. Stolen, and J. O. Aagedal. The CORAS methodology: model-based risk assessment using UML and UP. In *UML and the unified process*, pages 332–357. IGI Publishing, 2003.

[9] G. Elahi and E. Yu. A goal oriented approach for modeling and analyzing security trade-offs. In *Proc. of ER'07*, LNCS 4801, pages 375–390. Springer, 2007.

[10] P. Giorgini, F. Massacci, J. Mylopoulos, and N. Zannone. Modeling security requirements through ownership, permission and delegation. In *Proc. of RE'05*, pages 167–176. IEEE Computer Society, 2005.

[11] P. Giorgini, J. Mylopoulos, E. Nicchiarelli, and R. Sebastiani. Formal reasoning techniques for goal models. *J. Data Semantics*, 1:1–20, 2003.

[12] C. Haley, R. Laney, J. Moffett, and B. Nuseibeh. Security requirements engineering: A framework for representation and analysis. *TSE*, 34(1):133–153, 2008.

[13] J. Horkoff. Using i* Models for Evaluation. Master's thesis, University of Toronto, 2006.

[14] IBM Global Technology Services. IBM Internet Security Systems X-Force 2007 Trend Statistics, 2008.

[15] ISO/IEC. Management of Information and Communication Technology Security – Part 1: Concepts and Models for Information and Communication Technology Security Management. ISO/IEC 13335, 2004.

[16] J. Jürjens. *Secure Systems Development with UML*. Springer, 2004.

[17] J. Krogstie, A. L. Opdahl, and S. Brinkkemper. Capturing dependability threats in conceptual modelling. *Conceptual Modelling in Information Systems Engineering*, pages 247–260, 2007.

[18] L. Liu, E. Yu, and J. Mylopoulos. Analyzing security requirements as relationships among strategic actors. In *Proc. of SREIS'02*, October 2002.

[19] L. Liu, E. Yu, and J. Mylopoulos. Security and privacy requirements analysis within a social setting. In *Proc. of RE'03*, page 151. IEEE Computer Society, 2003.

[20] R. Matulevicius, N. Mayer, H. Mouratidis, E. Dubois, P. Heymans, and N. Genon. Adapting secure tropos for security risk management in the early phases of information systems development. In *CAiSE*, pages 541–555, 2008.

[21] J. McDermott and C. Fox. Using abuse case models for security requirements analysis. In *Proc of ACSAC'99*, page 55. IEEE Computer Society, 1999.

[22] J. P. McDermott. Attack net penetration testing. In *Proc. of NSPW'00*, pages 15–21. ACM, 2000.

[23] N. Meyer, A. Rifaut, and E. Dubois. Towards a Risk-Based Security Requirements Engineering Framework. REFSQ-Proc. Of Internat. *Workshop on Requirements Engineering for Software Quality*, 2005.

[24] H. Mouratidis, P. Giorgini, G. Manson, and I. Philp. A natural extension of tropos methodology for modelling security. In *Proceedings of the Workshop on Agent-oriented methodologies, at OOPSLA*, 2002.

[25] C. Phillips and L. P. Swiler. A graph-based system for network-vulnerability analysis. In *Proc. of NSPW'98*, pages 71–79. ACM, 1998.

[26] E. R. Kissel. *Glossary of key information security terms.* NIST IR 7298, 2005.

[27] F. B. Schneider, editor. *Trust in Cyberspace.* National Academy Press, 1998.

[28] B. Schneier. Attack trees. *Dr. Dobb's Journal*, 24(12):21–29, 1999.

[29] B. Schneier. *Beyond Fear.* Springer, 2003.

[30] B. Schneier. The psychology of security. *Commun. ACM*, 50(5):128, 2007.

[31] G. Sindre and L. Opdahl. Eliciting security requirements with misuse cases. *Requir. Eng.*, 10(1):34–44, 2005.

[32] A. van Lamsweerde. Elaborating security requirements by construction of intentional anti-models. In *Proc. of ICSE'04*, pages 148–157. IEEE Computer Society, 2004.

[33] A. van Lamsweerde and E. Letier. Handling Obstacles in Goal-Oriented Requirements Engineering. *TSE*, 26(10):978–1005, 2000.

[34] W. E. Vesely, F. F. Goldberg, N. Roberts, and D. F. Haasl. Fault tree handbook. Technical Report NUREG-0492, U.S. Nuclear Regulatory Commission, January 1981.

[35] E. Yu. *Modeling Strategic Relationships for Process Reengineering.* PhD thesis, University of Toronto, 1995.

[36] E. S. K. Yu. Towards modeling and reasoning support for early-phase requirements engineering. In *Proc. of RE'97*, page 226. IEEE Computer Society, 1997.