

A Soft Constraint-based Approach to the Cascade Vulnerability Problem

Stefano Bistarelli

Istituto di Informatica e Telematica, CNR, Pisa, Italy

`stefano.bistarelli@iit.cnr.it`

Dipartimento di Scienze

Università degli Studi “G. D’annunzio”, Pescara, Italy

`bista@sci.unich.it`

Simon N. Foley

Department of Computer Science

University College Cork, Ireland

`s.foley@cs.ucc.ie`

Barry O’Sullivan

Cork Constraint Computation Centre

Department of Computer Science

University College Cork, Ireland

`b.osullivan@cs.ucc.ie`

Abstract

The security of a network configuration is based not just on the security of its individual components and their direct interconnections, but also on the potential for systems to interoperate indirectly across network routes. Such interoperation has been shown to provide the potential for cascading paths that violate security, in a circuitous manner, across a network. In this paper we show how constraint satisfaction provides a natural approach to expressing the necessary constraints to ensure multilevel security across a network configuration. In particular, soft constraints are used to detect and eliminate the cascading network paths that compromise security. Taking this approach results in practical advancements over existing solutions to this problem. In particular, constraint satisfaction highlights the set of all cascading paths, which we can eliminate in polynomial time by breaking a minimal number of system links to ensure security.

1 Introduction

The composition of secure systems is not necessarily secure. A user may be able to gain unauthorised access to an object by taking a circuitous access route across individually secure but interoperating systems [16]. Determining security is based not just on the individual system authorisation mechanisms but also on how the systems are configured to interoperate. For example, if Alice is permitted to have access to Bob’s files on the Administration system, and Clare is permitted access Alice’s files on the Sales system, then is it safe to support file sharing between these systems? The extent of system interoperation must be limited if the administration security policy states that Clare is not permitted access to Bob’s (administration) files.

The *cascade vulnerability problem* [29] is concerned with secure interoperation, and considers the *assurance risk* of composing multilevel secure systems that are evaluated to different levels of assurance according to the criteria specified in [29]. The transitivity of the multilevel security policy upheld across all secure systems ensures that their multilevel composition is secure; however, interoperability and data sharing between systems may increase the risk of compromise beyond that accepted by the assurance level. For example, it may be an acceptable risk to store only secret and top-secret data on a medium assurance system, and only classified and secret data on another medium assurance system; classified and top-secret data may be stored simultaneously only on ‘high’ assurance systems. However, if these medium assurance systems interoperate at classification secret, then the acceptable risk of compromise is no longer adequate as there is an unacceptable cascading risk from top-secret across the network to classified.

Existing research has considered schemes for *detecting* these security vulnerabilities and for *eliminating* them by reconfiguring system interoperation. While the detection of the cascade vulnerability [14, 15, 16, 18, 23] can be easily achieved, their optimal elimination is NP-complete [16, 17, 18].

We present an approach to using constraints [13, 21] for reasoning about secure interoperation. Constraint solving is an emerging software technology for modelling and solving large-scale optimisation problems [30]. Constraints have been successfully applied to a number of problems in computer security [2, 3, 6, 7, 24]. However, the cascade vulnerability problem, and secure interoperation in general, have not been studied.

The approach that we present in this paper represents a paradigm shift in the modelling, detection and elimination of the cascade vulnerability problem. We present a constraint model that provides a natural description of an arbitrary multilevel secure network. Any solution to the model represents a cascading path through the network, providing significantly more information on its vulnerabilities than the existing approaches, and providing a basis for eliminating the cascade vulnerability problem. Previous approaches [14, 18] detect a single cascading path in polynomial time, but correcting the cascade in an optimal way is NP-complete. Using a constraint model, we can rely on a significant body of successful techniques from the field of constraint processing for finding

the set of cascading paths which, once found, can be eliminated in polynomial time. These results are applicable to secure interoperation in general.

This paper combines and extends work originally published in [8, 9]. The paper is organised as follows. Section 2 provides background on the cascade vulnerability problem and on soft constraints. Section 3 describes a soft-constraint based model for describing multilevel secure networks. Solutions to this constraint model represent the possible information flow paths through the network and Section 4 characterises the cascade vulnerability problem in terms of these solutions. An example is considered in Section 5 and Section 6 proposes a polynomial-time scheme for eliminating cascading paths from the network.

2 Background

2.1 The Cascade Vulnerability Problem

Figure 1 gives an example of a multilevel security (MLS) network configuration with a cascade vulnerability problem [14].




Figure 1: Network configuration with a cascade vulnerability problem.

The network is comprised of multilevel secure systems *Sys.E*, *Sys.F*, *Sys.G* and *Sys.H* storing classified (C), secret (S) and top-secret (T) information as depicted in Figure 1. Each system is accredited according to levels of assurance $C2 < B1 < B2 < B3$ from [28, 29]. For example, *Sys.F* is used to simultaneously store classified, secret and top-secret information and, therefore, (according to [28, 29]) must be evaluated at level B3 or higher, reflecting the high level of confidence that must be placed in the secure operation of the system. This is to counter the risk of an attacker compromising the system and copying top-secret information to classified. *Sys.H*, on the other hand, has been evaluated at the lowest level of assurance C2 and, therefore, may be used only to store single level data.

However, the security-level interoperability defined by the system connections in Figure 1 results in a cascade vulnerability across the network. There

is a risk that an attacker who has the ability to compromise security on B2 or lower assured systems can copy T to S on Sys.E, to S on Sys.H to S to C on Sys.G. This is contrary to the requirement that the level of assurance that T cannot be copied to C should be B3 or higher. This requirement is met by the individual systems but not as a result of their interoperation. A generalised form of the cascade vulnerability problem is defined as follows.

2.1.1 MLS

A multilevel secure system enforces a lattice-based security policy \mathcal{L} of security levels that has ordering relation \leq . Given $x, y : \mathcal{L}$ then $x \leq y$ means that information may flow from level x to level y , for example, $C \leq S \leq T$.

2.1.2 Assurance Levels

Security criteria define a lattice, \mathcal{A} , of assurance levels with ordering \leq . Given $x, y : \mathcal{A}$, then $x \leq y$ means that a system evaluated at y is no less secure than a system evaluated at x , or alternatively, that an attacker that can compromise a system evaluated at y can compromise a system evaluated at x . Let \mathcal{S} define the set of all possible systems. We define $accred : \mathcal{S} \rightarrow \mathcal{A}$ where $accred(s)$ gives the assurance level of system $s : \mathcal{S}$, and is taken to represent the minimum effort required by an attacker to compromise system s .

2.1.3 Acceptable Risk

Security evaluation criteria also define an acceptable risk function $risk : \mathcal{L} \times \mathcal{L} \rightarrow \mathcal{A}$, such that given $l, l' : \mathcal{L}$ then $risk(l, l')$ defines the minimum acceptable risk of compromise from l to l' ; it represents the minimum acceptable effort required to ‘compromise security’ and copy/downgrade information from level l to level l' . Without any loss of generality we assume that there is no security enforcement at the lowest assurance level $\mathbf{0}$, and thus, if $l \leq l'$ then $risk(l, l') = \mathbf{0}$. For example, the function $risk$ encodes the assurance matrix for the ‘B’ levels (from [28, 29]) 1, 2 and 3, with: 0 representing no security enforcement, as $risk(C, S) = risk(C, T) = risk(S, T) = 0$, $risk(S, C) = 1$, $risk(T, S) = 2$, and $risk(T, C) = 3$, and so forth.

2.1.4 Evaluated Systems

Individual systems must be assured higher than the minimum acceptable risk to compromise the data they store. If a system s can hold information at levels l and l' then $risk(l, l') \leq accred(s)$.

2.1.5 Network Model

A system node in our network model is a pair, l_s , and represents the fact that system s can hold information at level l . Thus, a system is a collection of nodes that represent the data it holds. For example, in Figure 1, Sys.E can store

secret and top-secret information and is represented by nodes S_E and T_E . A network of systems is a weighted graph of these nodes according to how they are connected. An w -weighted arc from l_s to $l_{s'}$ means that it requires at minimum w effort to directly copy information at level l held on system s to level l' on system s' .

2.1.6 Cascading Risks

Arcs are used to represent direct flows within a system and interoperation links between systems. A flow $l \leq l'$ that is permitted on system s is represented as a (assurance) $\mathbf{0}$ -weighted arc from l_s to l'_s ; if a flow is not permitted between levels l and l' that are held on system s then it is represented as an arc weighted as $accred(s)$ from l_s to l'_s .

A link from system s to s' that connects l -level information is represented as a $\mathbf{0}$ -weighted arc from l_s to $l_{s'}$ —all other pairs l_s to $l_{s'}$ not related in this way are either represented as having no arc, or an arc with the maximum assurance value $\mathbf{1}$.

Given pairs l_s and $l'_{s'}$ we then define $effort(l_s, l'_{s'})$ as the minimum effort required to compromise the network and copy and/or downgrade level l information held on system s to level l' information on system s' . As an example, in Figure 1, the effort to ‘copy’ top-secret information on system Sys. E to classified on system Sys. G is $effort(T_E, C_G) = B2$ via the path through Sys. H .

2.1.7 Cascade Freedom

We require that for any pair of systems s, s' and levels l, l' , then we have $risk(l, l') \leq effort(l_s, l'_{s'})$. Given a path in the network from l_s to $l'_{s'}$, then its cascade weight is the maximum weight that directly connects any two nodes on the path. This reflects the minimum effort that will be required by an attacker to copy information from l_s to $l'_{s'}$ by using this path. $effort(l_s, l'_{s'})$ is the minimum of the cascade weights for all paths that connect l_s to $l'_{s'}$.

2.2 Soft Constraints

Several formalisations of the concept of *soft constraints* are currently available. In the following, we refer to the one based on c -semirings [4, 11], which can be shown to generalise and express many of the others [5].

A soft constraint may be seen as a constraint where each instantiation of its variables has an associated value from a partially ordered set that can be interpreted as a set of preference values. Combining constraints will then have to take into account such additional values, and thus the formalism has also to provide suitable operations for combination (\times) and comparison ($+$) of tuples of values and constraints. This is why this formalisation is based on the concept of c -semiring, which is just a set plus two operations.

2.2.1 Semirings

A semiring is a tuple $\langle \mathcal{A}, +, \times, \mathbf{0}, \mathbf{1} \rangle$ such that: 1. \mathcal{A} is a set and $\mathbf{0}, \mathbf{1} \in \mathcal{A}$; 2. $+$ is commutative, associative and $\mathbf{0}$ is its unit element; 3. \times is associative, distributes over $+$, $\mathbf{1}$ is its unit element and $\mathbf{0}$ is its absorbing element. A c-semiring is a semiring $\langle \mathcal{A}, +, \times, \mathbf{0}, \mathbf{1} \rangle$ such that: $+$ is idempotent, $\mathbf{1}$ is its absorbing element and \times is commutative. Let us consider the relation \leq_S over \mathcal{A} such that $a \leq_S b$ iff $a + b = b$. Then it is possible to prove that (see [11]): 1. \leq_S is a partial order; 2. $+$ and \times are monotone on \leq_S ; 3. $\mathbf{0}$ is its minimum and $\mathbf{1}$ its maximum; 4. $\langle \mathcal{A}, \leq_S \rangle$ is a lattice and, for all $a, b \in \mathcal{A}$, $a + b = \text{lub}(a, b)$ (where *lub* is the *least upper bound*). Moreover, if \times is idempotent, then: $+$ distributes over \times ; $\langle \mathcal{A}, \leq_S \rangle$ is a distributive lattice and \times its *glb* (*greatest lower bound*). Informally, the relation \leq_S gives us a way to compare semiring values and constraints. In fact, when we have $a \leq_S b$, we will say that *b is better than a*. In the following, when the semiring will be clear from the context, $a \leq_S b$ will be often indicated by $a \leq b$.

2.2.2 Constraint Problems

Given a semiring $S = \langle \mathcal{A}, +, \times, \mathbf{0}, \mathbf{1} \rangle$ and an ordered set of variables V over a finite domain D , a *constraint* is a function which, given an assignment $\eta : V \rightarrow D$ of the variables, returns a value of the semiring. By using this notation we define $\mathcal{C} = \eta \rightarrow \mathcal{A}$ as the set of all possible constraints that can be built starting from S , D and V .

Note that in this *functional* formulation, each constraint is a function. Such a function involves all the variables in V , but it depends on the assignment of only a finite subset of them. So, for instance, a binary constraint $c_{x,y}$ over variables x and y , is a function $c_{x,y} : V \rightarrow D \rightarrow \mathcal{A}$, but it depends only on the assignment of variables $\{x, y\} \subseteq V$. We call this subset the *support* of the constraint. More formally, consider a constraint $c \in \mathcal{C}$. We define its support as $\text{supp}(c) = \{v \in V \mid \exists \eta, d_1, d_2. c\eta[v := d_1] \neq c\eta[v := d_2]\}$, where

$$\eta[v := d]v' = \begin{cases} d & \text{if } v = v', \\ \eta v' & \text{otherwise.} \end{cases}$$

Note that $c\eta[v := d_1]$ means $c\eta'$ where η' is η modified with the assignment $v := d_1$ (that is the operator $[]$ has precedence over application). Note also that $c\eta$ is the application of a constraint function $c : V \rightarrow D \rightarrow \mathcal{A}$ to a function $\eta : V \rightarrow D$; what we obtain, is a semiring value $c\eta = a$.

A *soft constraint satisfaction problem* (SCSP) is a pair $\langle C, \text{con} \rangle$ where $\text{con} \subseteq V$ and C is a set of constraints: con is the set of variables of interest for the constraint set C , which however may concern also variables not in con . Note that a classical CSP is a SCSP where the chosen c-semiring is: $S_{CSP} = \langle \{false, true\}, \vee, \wedge, false, true \rangle$. Fuzzy CSPs [26] (FCSP) can instead be modelled in the SCSP framework by choosing the c-semiring $S_{FCSP} = \langle [0, 1], \max, \min, 0, 1 \rangle$. Many other “soft” CSPs (Proba-

bilistic, weighted, ...) can be modelled by using a suitable semiring structure ($S_{prob} = \langle [0, 1], max, \times, 0, 1 \rangle$, $S_{weight} = \langle \mathcal{R}, min, +, +\infty, 0 \rangle$, ...).

Figure 2 shows the graph representation of a fuzzy CSP. Variables and constraints are represented respectively by nodes and by undirected (unary for c_1 and c_3 and binary for c_2) arcs, and semiring values are written to the right of the corresponding tuples. The variables of interest (that is the set con) are represented with a double circle. Here we assume that the domain D of the variables contains only elements a and b and c .




Figure 2: A fuzzy CSP.

2.2.3 Combining and projecting soft constraints

Given the set \mathcal{C} , the combination function $\otimes : \mathcal{C} \times \mathcal{C} \rightarrow \mathcal{C}$ is defined as $(c_1 \otimes c_2)\eta = c_1\eta \times_S c_2\eta$. In words, combining two constraints means building a new constraint whose support involves all the variables of the original ones, and which associates with each tuple of domain values for such variables a semiring element which is obtained by multiplying the elements associated by the original constraints to the appropriate sub-tuples. It is easy to verify that $supp(c_1 \otimes c_2) \subseteq supp(c_1) \cup supp(c_2)$.

Given a constraint $c \in \mathcal{C}$ and a variable $v \in V$, the *projection* of c over $V - \{v\}$, written $c \Downarrow_{(V - \{v\})}$ is the constraint c' s.t. $c'\eta = \sum_{d \in D} c\eta[v := d]$. Informally, projecting means eliminating some variables from the support. This is done by associating with each tuple over the remaining variables a semiring element which is the sum of the elements associated by the original constraint to all the extensions of this tuple over the eliminated variables. In short, combination is performed via the multiplicative operation of the semiring, and projection via the additive one.

2.2.4 Solutions

A *solution* of an SCSP $P = \langle C, con \rangle$ is the constraint $Sol(P) = (\bigotimes C) \Downarrow_{con}$. That is, we combine all constraints, and then project over the variables in con .

In this way we get the constraint with support (not greater than) *con* which is “induced” by the entire SCSP. Note that when all the variables are of interest we do not need to perform any projection.

For example, the solution of the fuzzy CSP of Figure 2 associates a semiring element to every domain value of variable x . Such an element is obtained by first combining all the constraints together. For instance, for the tuple $\langle a, a \rangle$ (that is, $x = y = a$), we have to compute the minimum between 0.9 (which is the value assigned to $x = a$ in constraint c_1), 0.8 (which is the value assigned to $\langle x = a, y = a \rangle$ in c_2) and 0.9 (which is the value for $y = a$ in c_3). Hence, the resulting value for this tuple is 0.8. We can do the same work for tuple $\langle a, b \rangle \rightarrow 0.2$, $\langle a, c \rangle \rightarrow 0.2$, $\langle b, a \rangle \rightarrow 0$, $\langle b, b \rangle \rightarrow 0$, $\langle b, c \rangle \rightarrow 0.1$, $\langle c, a \rangle \rightarrow 0.8$, $\langle c, b \rangle \rightarrow 0.2$ and $\langle c, c \rangle \rightarrow 0.2$. The obtained tuples are then projected over variable x , obtaining the solution $\langle a \rangle \rightarrow 0.8$, $\langle b \rangle \rightarrow 0.1$ and $\langle c \rangle \rightarrow 0.8$.

3 Modelling MLS Networks

Consider a network $N = \{E, F, G, H, \dots\}$ of a finite arbitrary number n of systems. This network is represented in our constraint model in terms of all possible paths (of length n and less) that connect the systems. The paths are modelled using $2 \times n$ path variables, where each path variable P_i^s and P_i^d can be instantiated to be one system of the network. A path through the network is represented by a specific instantiation of the variables $[P_1^s, P_1^d, P_2^s, P_2^d, \dots, P_n^s, P_n^d]$. In particular, the instantiation of the pair of nodes P_i^s and P_i^d , for $i := 1 \dots n$, represents the flow from the source P_i^s to the destination P_i^d within the system at the i -th position of the path. Similarly, instantiation of P_i^d and P_{i+1}^s , for $i := 1 \dots n - 1$ represents the flow among the the i -th and the $i + 1$ -th system in the specific instantiated path.

Consider for instance the network $N = \{E, F\}$ represented in Figure 3 involving two systems, Sys.*E* and Sys.*F*, with Sys.*E* handling information at level Top-Secret (T) and Secret (S), and Sys.*F* handling information at level Secret (S) and Confidential (C). We can capture this instance by using 4 path variables: $[P_1^s, P_1^d, P_2^s, P_2^d]$.




Figure 3: A simple network.

3.1 Path Variable Domains

The domain of each path variable defines the set of possible security levels available on each system. In particular, each *source* variable P_i^s contains domain elements marked with s , and each *destination* variable P_i^d contains domain elements marked with d .

The network in Figure 3 has in our model 4 variables $[P_1^s, P_1^d, P_2^s, P_2^d]$ with $\text{dom}(P_i^s) = \{\mathsf{T}_E^s, \mathsf{S}_E^s, \mathsf{S}_F^s, \mathsf{C}_F^s\}$, with $i := 1, 2$, and $\text{dom}(P_i^d) = \{\mathsf{T}_E^d, \mathsf{S}_E^d, \mathsf{S}_F^d, \mathsf{C}_F^d\}$, for $i := 1, 2$. The strategy in our network model is that a network configuration will be represented as a series of constraints between the source and destination variables, representing all possible paths across the network.

In general, when the network contains $n > 2$ systems, we also need to be able to deal with shorter paths of length $k < n$. To do this, we need to extend the domain of each path variable, $P_i^?$ (where ? stands alternatively for source s and destination d), for any $i > 2$, with some artificial elements. More precisely, we extended the domain $\text{dom}(P_i^?)' = \text{dom}(P_i^?) \cup \{*_1^?, *_2^?, \dots, *_i^?, *_i^?_{-2}\}$. These $*$ elements are added to deal with paths shorter than n . This is necessary because solving an SCSP requires finding an assignment for each variable in the SCSP and we may want to represent paths shorter than the number of nodes in the network.

3.2 Modelling each System

The flows that are possible within a particular system occurring in position i of a network path are modelled as a constraint between the source and destination variables P_i^s and P_i^d . These flows reflect the accesses that are permitted by the system's MLS security mechanism. For example, secret information is permitted to flow to top-secret in Sys. E and, thus, there is a constraint $[P_1^s := \mathsf{S}_E; P_1^d := \mathsf{T}_E s]$, which evaluates to 0, meaning unrestricted information flow. The flows that are possible within a system can be categorised in three ways.

- *Flow_{permitted}* represents the information flows that are permitted by the policy in each node. For example, S may flow to T in Sys. E .
- *Flow_{risk}* represents the information flows that are not permitted by the policy, but for which there is a risk of flow if the system became compromised. For example, the risk of a flow from T to S in system Sys. E is B2 (assurance level 2), corresponding to the level of assurance at which Sys. E has been evaluated.
- *Flow_{invalid}*: represents all the remaining flows that are not valid (that is, are impossible for the given system). For example, a flow from T to C is not possible on Sys. E since the system has not been configured to store information labelled as classified.

Consider an arbitrary system S_i that can occur on position i of a path through the network. Between each pair of variables P_i^s and P_i^d for each system

S_i , we define a soft constraint, $c_{(P_i^s, P_i^d)}$, that gives a weight to each possible (permitted or risk) flow within that system. Various semirings could be used to represent the network and the associated policy. We use the following semiring in this paper, although our results are general and are not limited to this particular one:

$$S_{cascade} = \langle \mathbb{N}, \min, \max, +\infty, 0 \rangle.$$

Given this semiring, the constraint $c_{(P_i^s, P_i^d)}$ representing the flow inside a system S that occurs on position i of the path is defined as follows:

$$c_{(P_i^s, P_i^d)}(s, d) = \begin{cases} \text{accred}(S_i) & (s, d) \in Flow_{risk} \\ & \text{(risk flows);} \\ 0 & (s, d) \in Flow_{permitted} \\ & \text{(permitted flows);} \\ +\infty & \text{otherwise} \\ & \text{(invalid flows).} \end{cases}$$

Since the domain of the variables $P_i^?$ (where ? stands for s and d) has been extended with the elements $\{*_1^?, *_2^?, \dots, *_i^?\}$, we have also to consider these ‘‘artificial’’ values. In particular, we extend the definition of each constraint $c_{(P_i^s, P_i^d)}$ as follows:

$$c_{(P_i^s, P_i^d)}(s, d) = \begin{cases} 0 & (s, d) \in \{(*_1^s, *_1^d), \dots, (*_{i-2}^s, *_i^d)\} \\ & \text{(Artificial permitted flows);} \\ +\infty & \text{otherwise} \\ & \text{(Artificial invalid flows).} \end{cases}$$



Figure 4: Flow constraints within Sys.E and Sys.F.

The constraints between P_1^s and P_1^d and between P_2^s and P_2^d in Figure 4 depicts some of the system flow constraints that model the network configuration in Figure 3. The use of the ‘*’ elements and the representation of connections between systems within the network is considered in the next section.




Figure 5: Flow constraints on the connection from Sys. E to Sys. F .

3.3 Modelling the Network

Security level interconnections between systems result in two classes of flows between systems in a network.

- $Network_{permitted}^k$ represents information flows permitted by the connection policy between each system and represents direct synchronisation flows between systems. For example, the secret-level connection from Sys. E to Sys. F in Figure 3 corresponds to a network permitted flow from S_E to S_F .
- $Network_{invalid}^k$ represents the absence of direct connection between the systems in the network configuration.

For each adjacent pair of systems at positions i and $i + 1$ along a network path, we define a soft constraint, $c_{(P_i^d, P_{i+1}^s)}$, that defines the possible synchronisations between the systems. Note that these constraints are defined between the destination variable of the first system, P_i^d , and the source variable of the second system, P_{i+1}^s . The constraint $c_{(P_i^d, P_{i+1}^s)}$ representing the synchronisation flows between system at positions i and $i + 1$ is defined as follows:

$$c_{(P_i^d, P_{i+1}^s)}(d, s) = \begin{cases} 0 & (d, s) \in Network_{permitted}^k \\ & \text{(Policy permitted synchronisation);} \\ +\infty & \text{otherwise} \\ & \text{(invalid synchronisation).} \end{cases}$$

For example, the network configuration $Network_{permitted}^k$ for Figure 3 is defined as follows.

$$Network_{permitted}^k = \{(S_E^d, S_F^s), (S_F^d, S_E^s)\}$$

Note that the connection model does not consider assurance risks for connections; this can be achieved, if desired, by explicitly modelling the connections by their components (for example, a link encryption device) and corresponding assurance levels.

When connecting systems at P_i^d and P_{i+1}^s it is also necessary to consider the constraints imposed by the artificial elements $*_i^?$. The definition of each constraint $c_{(P_i^d, P_{i+1}^s)}$ is extended as follows:

$$c_{(P_i^d, P_{i+1}^s)}(d, s) = \begin{cases} 0 & (d, s) \in \{(*_1^d, *_2^s), \dots, (*_{i-3}^d, *_{i-2}^s)\} \\ & \cup \{(\#, *_1^s) \text{ for all } \# \in D(P_i^d)\} \\ & \text{(Artificial permitted synchronisation);} \\ +\infty & \text{otherwise} \\ & \text{(Artificial invalid synchronisation).} \end{cases}$$

The extension of this constraint is slightly different to the previous system-level constraints. In particular, it enables us to model the connection between the last real domain element in the path and the first $*_1^s$ -element.

Figure 5 depicts some of the system and network flow constraints that model the network configuration in Figure 3. This constraint network represents a number of paths of length 4 (corresponding to $2 * n$, where n is the number of systems in the configuration). The path represented by the constraint solution $\eta = [P_1^s := T_E^s; P_1^d := S_E^d; P_2^s := S_F^s; P_2^d := C_F^d]$ has assurance/risk values: 2 (top-secret to secret compromise on Sys.E); 0 (secret level system connection); and 1 (secret to classified compromise on Sys.F), on the relating arcs, respectively.

For a given configuration, every path through the network must be modelled in this way. Figure 5 represents just one path and its subpaths, starting from Sys.E. Further constraints must be added to the model of the configuration to depict the other paths. For example, paths that start from Sys.F. Figure 5 depicts some of the constraints that represents all paths for the configuration. In more complex network configurations the $*$ elements allow the modelling of alternate paths between systems. For example, in the case of Figure 1, paths (of length four) include $[E; F; G; H]$, $[E; H; G; *]$, and so forth.




Figure 6: Modelling multiple paths.

In addition to ensuring that systems are configured in a valid way, we also need to ensure that no two pairs of path variables represent the same system. This ensures that our model does not capture cyclic paths. Therefore, we need to post an *alldifferent* [25] constraint amongst all the variables in the model. An *alldifferent* constraint ensures that all variables over which it is defined take on different values.

The solutions of the defined SCSP (referred to as the Effort-CSP, \mathcal{E}), that is all the solutions with a weight lower than $+\infty$, represents all of the possible paths through the system. The semiring level associated with each path (solution) gives a measure of the effort required to compromise the network using that specific path.

4 Detecting Cascade Vulnerabilities

To determine whether or not there exists a cascade vulnerability problem, we need to compare the effort required to compromise the network against the risk of compromising the system as a whole. Therefore, we introduce a set of *risk* constraints, $\mathcal{R} = \{r_{(P_1^s, P_i^d)} | i \in \{2, \dots, n\}\}$. The weight of each instance of $r_{(P_1^s, P_i^d)}$ represents the risk associated with the path from P_1^s to P_i^d . The cost of each tuple in these constraints is defined as follows:

$$r_{(S_1^s, S_i^d)}(s, d) = \begin{cases} \mathbf{0} & \text{if } d = *^d_i; \\ risk(s, d) & \text{otherwise.} \end{cases}$$

The set of solutions of the SCSP E (that is the Effort-CSP defined above), each of whose associated semiring level is lower than $+\infty$, represents the set of paths through the network. The semiring level associated with each solution-path of E represents the minimum *effort* required to compromise the network, while the combination of the constraints in R (the Risk-CSP) gives the risk for all the paths. Therefore, a cascading path can be identified as any path η where the risk associated with the path exceeds the effort to compromise it, that is, where the following constraint is satisfied:

$$\otimes \mathcal{R}\eta > \otimes \mathcal{E}\eta$$

Therefore, by adding the above constraint to our constraint model, the existence of a solution to that model indicates that here exists a cascading path. Furthermore, the set of solutions provides the set of cascading paths. This provides us with a basis upon which we can set about removing the cascade vulnerability problem from the network by eliminating all solutions of the model.

5 An Example

In this section we encode the network example described in Figure 1 within the proposed constraints model. Figure 7 depicts the structure of the constraint

relationships in this model. We first present an example of how our model identifies a cascade-free path, and then present an example of detecting a cascading path.




Figure 7: The constraint model structure.

For the purposes of the examples, the risk lattice is assumed to be as follows: $risk(C, S) = risk(C, T) = risk(S, T) = 0$, $risk(S, C) = 1$, $risk(T, S) = 2$, $risk(T, C) = 3$.

Figure 7 presents the structure of the constraint model for an example from [14]. Our model comprises 8 path variables, $P_1^s, P_1^d, P_2^s, P_2^d, P_3^s, P_3^d, P_4^s$, and P_4^d , and 3 risk variables, $r_{(P_1^s, P_2^d)}, r_{(P_1^s, P_3^d)}$ and $r_{(P_1^s, P_4^d)}$. The domain of each path variable, $D(P_i^?)$, is: $\{T_E^?, S_E^?, T_F^?, S_F^?, C_G^?, S_G^?, C_H^?, S_H^?\}$ (where ? stands alternatively for s and d) and $i := 1, \dots, 4$. Note that we also extend each domain using $*_i^?$ values as described above, but do not show this here for conciseness.

5.1 A Cascade-free Path

Consider the following path through the network:

$$\eta = [P_1^s := T_E^s, P_1^d := T_E^d, P_2^s := T_F^s, P_2^d := S_F^d, P_3^s := S_G^s, P_3^d := C_G^d, P_4^s := *_1^s, P_4^d := *_1^d].$$

This scenario is illustrated in Figure 8.

Evaluating the cascade detection constraint we get the following, proving that this path is cascade-free:

$$\otimes \mathcal{R}\eta > \otimes \mathcal{E}\eta \equiv 3 > 3 \equiv False.$$

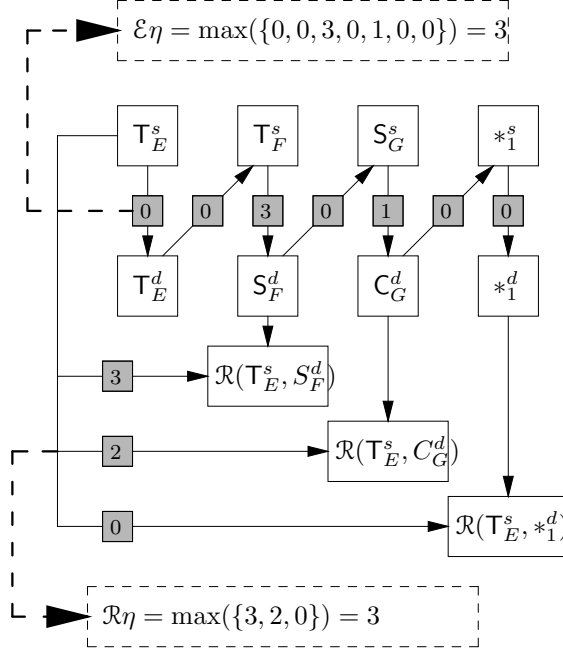


Figure 8: A cascade-free path.

5.2 A Cascading Path

Consider the following path through the network, depicted in Figure 9:

$$\eta = [P_1^s := T_E^s, P_1^d := S_E^d, P_2^s := S_H^s, P_2^d := S_H^d, \\ P_3^s := S_G^s, P_3^d := C_G^d, P_4^s := *1^s, P_4^d := *1^d]$$

Evaluating the cascade detection constraint we get the following:

$$\otimes \mathcal{R}\eta > \otimes \mathcal{E}\eta \equiv 3 > 2 \equiv True$$

Therefore, this path exhibits a cascade vulnerability problem.

6 Eliminating Cascade Vulnerabilities

In order to eliminate the cascade vulnerability problem from an MLS network it is necessary to remove all of the cascading paths that run through it. However, in breaking links (connections) between systems, the services provided by the network are affected. Therefore, when eliminating the cascade vulnerability problem it is preferable to break as few links as possible.

The set of solutions to the constraint model presented earlier provides all of the cascading paths within the network. Therefore, in order to remove all

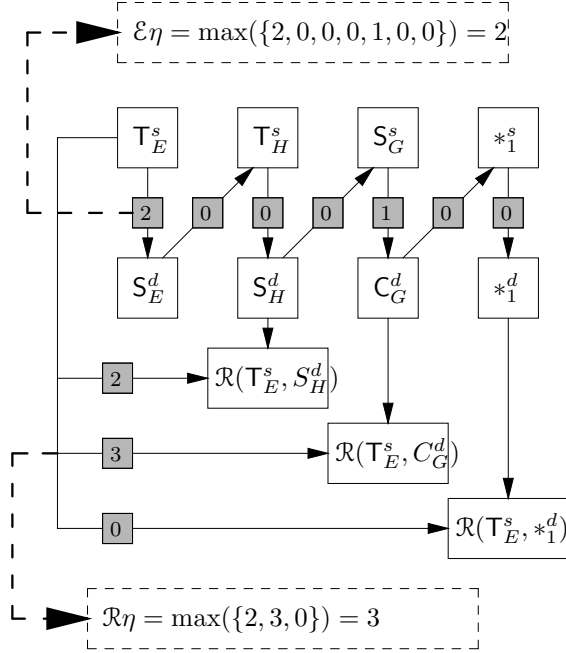


Figure 9: A cascading path.

cascade vulnerabilities from the network, we need to eliminate all solutions to the constraint model.

The problem of finding the *minimum number* of links to break in order to remove all of the cascading paths can be reformulated within our framework as the problem of finding the minimum number of unary constraints (each removing a link) that have to be added to the problem to make it unsatisfiable. This problem has been solved in [1] by solving the corresponding *minimum hitting-set* problem. Given a collection J of subsets of a finite set K , the minimum-hitting set of J is the smallest (cardinality) set $K' \subseteq K$ from J such that K' contains at least one element from each subset in J . Unfortunately, finding the *minimum hitting-set* is NP-complete.

Our approach to solving the problem takes a slightly different approach to [1]. Rather than finding a subset of the links with minimum cardinality, we focus on finding a set of links that is *minimal* in the sense that no link from such a set can be re-introduced without resulting in a cascading path. That is, we are looking for a set of unary constraints to add to our constraint model (with each unary constraint removing a link) such that the resultant CSP is unsatisfiable, and such that no proper subset of this set would give rise to a satisfiable CSP.

Definition 1 (Minimal Set of Removed Links) A set of links that we remove is *minimal* if:

1. the resultant network is cascade free, and
2. no proper subset of this set gives rise to a network without the cascade vulnerability problem. \triangle

Central to our approach to eliminating cascading paths is the notion of a *cascading path generator*.

Definition 2 (Cascading Path Generator) A cascading path generator, η , is a path involving a sequence of assignments to path variables that is not a super-sequence of another cascading path. Note that we ignore the artificial values $*_{s/d}^?$. \triangle

The set of cascading path generators are representative of all cascading paths in the set of solutions to our constraint model. In fact, each cascading path is a cascading path generator itself or one of its extensions. If an extension of a cascading path generator introduces no more critical links, the cascading effect on this path is removed when we remove the problem from the cascading path generator. If other critical links are added, the new links will also appear in another cascading path generator and will be solved when eliminating cascading vulnerability problem that generator. Thus, to focus attention on the causes of the cascade vulnerabilities it is sufficient to limit any elimination strategy to the set of cascading path generators. These can be obtained from the set of solutions in polynomial-time.

Each cascading path generator has a very important property: namely, that it is sufficient to remove one link on the path in order to remove the cascading effect associated with it.

Theorem 1 *Given a cascading path generator, η , removing any of the links on that path removes the cascading effect along that path.*

Proof This follows from the definition of cascading path generator. A cascading path generator represented by η does not, by definition, contain any sub-path that is also a cascading path. Therefore, we can safely remove *any* link in η and be sure the cascade along η (and any other cascading paths that η may generate) will be removed. \blacksquare

Starting from the set of all of the cascading path generators our goal is to find a set of links to be removed (unary constraints to be added) in order to obtain a cascade free network (a CSP with no solutions). Moreover we want to find a *minimal* set of links, as discussed earlier.

Example 1 Consider multilevel secure systems E, F, G, H, I and J managing information at security levels e, f, g, h, i, j and k that are configured according to Figure 10. Each system is evaluated at the lowest level of assurance 1, while the risk function specifies $risk(x, y) = 1$ for any level x and y , except for the cases

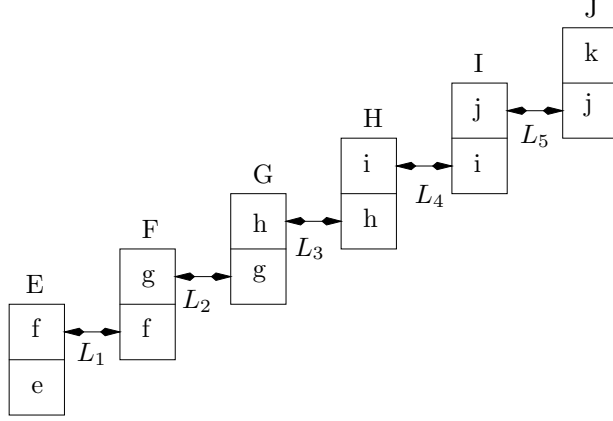


Figure 10: Multiple cascading path generators.

of $risk(e, h) = 2$, $risk(e, i) = 2$, $risk(f, i) = 2$, $risk(g, j) = 2$. This configuration results in the following set of cascading path generators:

$$\begin{aligned}
 P_1 &= \langle L_1, L_2 \rangle \\
 P_2 &= \langle L_2, L_3 \rangle \\
 P_3 &= \langle L_3, L_4 \rangle \\
 P_4 &= \langle L_4, L_5 \rangle.
 \end{aligned}$$

Note that while path $P_5 = \langle L_1, L_2, L_3 \rangle$ is also a cascading path it is not a cascading path generator: P_5 is covered by cascading path generators $\langle L_1, L_2 \rangle$ and $\langle L_2, L_3 \rangle$; removal of any link from each of these paths will also ensure elimination of cascading path P_5 .

The sets $\{L_1, L_3, L_5\}$, $\{L_2, L_3, L_5\}$ and $\{L_2, L_4\}$ are all minimal sets of links that we can remove to eliminate all cascade paths in the system. Notice that the cascading path $\{L_2, L_3, L_4\}$ is not minimal because by removing L_3 we can still obtain a cascade-free configuration. \triangle

Following the approach in [12, 20] used to find minimal conflict sets, we describe an approach to find a minimal set of links to be removed from the network. The following algorithm is used to select an approximation of a minimal set of removed links (the algorithm is known in the literature as *greedy minimal hitting-set* [19, 22, 27]):

1. Maintain a counter for each link involved in the set of minimal cascading paths that need to be removed;
2. Remove the most common link (the link with the highest counter), thus removing all minimal cascading paths involving that link; in case of a tie a random one is selected;

3. Update the link counters built in Step 1 to reflect the effect of reducing the set of minimal cascading paths that we need to consider;
4. Continue removing links and updating the link counters until all cascading paths have been removed.

This is a polynomial-time procedure that gives an approximation of a minimal set of removed links. Starting from this approximation we must now verify that it is a *minimal* set.

This is done using the technique proposed in [12, 20]. The strategy behind this technique is as follows. Given a set of links $\{l_1, l_2, \dots, l_n\}$ generated by the procedure above, then l_n must be part of a minimal set of links, since it was required to eliminate all cascading paths. The process is then repeated, this time starting from l_n , and building the set $\{l_n, l_1, l_2, \dots, l_k\}$ until all cascading paths have again been removed. Since adding a constraint corresponding to l_k lead to an unsatisfiable CSP, l_k must also be part of the minimal set too. At the next step we start by incrementally building the set $\{l_n, l_k, l_1, l_2, \dots, l_m\}$, until the initial set we use for the current iteration eliminates all cascading paths. An example is presented below to illustrate the strategy.

Step	Removed links	Paths remaining	Partial minimal set of links to break
1.	$\{L_3\}$	$\{P_1, P_4\}$	$\{\}$
2.	$\{L_3, L_2\}$	$\{P_4\}$	$\{\}$
3.	$\{L_3, L_2, L_5\}$	$\{\}$	$\{L_5\}$
4.	$\{L_5\}$	$\{P_1, P_2, P_3\}$	$\{L_5\}$
5.	$\{L_5, L_3\}$	$\{P_1\}$	$\{L_5\}$
6.	$\{L_5, L_3, L_2\}$	$\{\}$	$\{L_5, L_2\}$
7.	$\{L_5, L_2\}$	$\{P_3\}$	$\{L_5, L_2\}$
8.	$\{L_5, L_2, L_3\}$	$\{\}$	$\{L_5, L_2, L_3\}$ (minimal)
1.	$\{L_3\}$	$\{P_1, P_4\}$	$\{\}$
2.	$\{L_3, L_2\}$	$\{P_4\}$	$\{\}$
3.	$\{L_3, L_2, L_4\}$	$\{\}$	$\{L_4\}$
4.	$\{L_4\}$	$\{P_1, P_2\}$	$\{L_4\}$
5.	$\{L_4, L_3\}$	$\{P_1\}$	$\{L_4\}$
6.	$\{L_4, L_3, L_2\}$	$\{\}$	$\{L_4, L_2\}$
7.	$\{L_4, L_2\}$	$\{\}$	$\{L_4, L_2\}$ (minimal)

Table 1: Computing minimal sets of links to break all cascading paths.

Example 2 Consider again the system configuration depicted in Figure 10. Using the greedy minimal hitting-set algorithm, defined above, we could¹ find the following approximation of the minimal set of removed links: $r_1 = \{L_3, L_2, L_5\}$.

¹Notice that more than one result can be obtained if some ties are encountered as we consider the link counters.

We now apply the iterative technique [12, 20] to computing a minimal conflict as described above. Link L_5 leads to an inconsistent CSP and, therefore, L_5 must be in the minimal set. An alternative set of links to remove is built, starting from L_5 and following the same order as before. In removing link L_5 we remove only path P_4 and then we remove link L_3 , removing paths P_2 and P_3 . To remove all the paths we have to also remove link L_2 , which removes the path P_1 . L_2 is the last removed link, so it has to be in the minimal set with L_5 .

Starting again, this time from set $\{L_5, L_2\}$, which removes paths P_1, P_2 and P_4 . Cascading path generators remain in the network and thus we continue to select more links for removal. Selecting L_3 results in the removal of P_3 and, since the initial set that we used for the current iteration eliminates all cascading paths, we can say that the set $r_1 = \{L_3, L_2, L_5\}$ is a minimal set of links we can use.

In this case the greedy algorithm gave a minimal set of links. This is not always the case. Suppose the greedy algorithm first removes the link L_3 (that has a counter of 2), then removes link L_2 (that has a counter of 1) and then removes L_4 instead of L_5 (both have counter value 1). We obtain the set $r_2 = \{L_3, L_2, L_4\}$ as the approximation of the minimal set of removed links. Running the algorithm we first remove L_3 (removing paths P_2 and P_3), then L_2 (removing path P_1) and then link L_4 (removing the last path P_4). We now start from the L_4 , but must also remove links L_3 and L_2 to remove all paths. Now, starting with the set $r_2^* = \{L_4, L_2\}$ we remove all the paths, and thus r_2^* is a minimal cascading path.

Table 1 depicts the runs of the algorithm for the two examples. △

The following results regarding this procedure follow from [12, 20].

Theorem 2 (Soundness) *The network we obtain from applying the procedure is cascade free.*

Proof The procedure used to remove links is applied until no cascading path are present. Therefore, the network is cascade free. ■

Theorem 3 (Minimality) *The number of links that are removed from the network is minimal.*

Proof Notice that we can eliminate the cascading effect from a cascading path generator by removing any one of its links. The minimality of this set follows from [12, 20]. ■

7 Conclusion

In this paper we have presented a new approach to detecting and eliminating the cascade vulnerability problem in multilevel secure systems based on soft

constraints. Soft constraints have been successfully applied to other problems in computer security. The Role-Based Access Control policy model described in [3] uses soft-constraints to define authorisation but does not consider the issue of secure/cascading authorisation. [6, 24] considers how soft constraints might be used to specify noninterference-style security properties for systems. In [2] soft constraints are used to represent confidentiality and authentication properties of security protocols. A soft constraint based model is used in [10] to analyze the secure interoperation problem in [16]. These results, and the results in this paper, demonstrate the usefulness of constraints as a general purpose modelling technique for security.

The approach that we present in this paper represents a paradigm shift in the modelling, detection and elimination of the cascade vulnerability problem. In particular, our constraint model provides a natural and declarative description of an arbitrary multilevel secure system. Any solution to the model represents a cascading path, which provides significantly more information regarding the vulnerabilities in the network than the existing approaches. The set of solutions to the proposed constraint model provides a basis for removing the cascade vulnerability problem. Previous approaches [14, 18] detect a single cascading path in polynomial time, but correcting the cascade in an optimal way is NP-complete. As described above, detecting all paths in the constraint model is NP-hard, however elimination of a minimal number of links is polynomial.

While constraint solving is NP-complete in general, this has not detracted from its uptake as a practical approach to solving many real-world problems [30]. Using a constraint model, we can rely on a significant body of successful techniques for finding the set of cascading paths, which once found, can be eliminated in polynomial time. These results are applicable to secure interoperation in general.

Acknowledgments

The work in this paper builds upon earlier work presented in two conference papers by the authors [8, 9]. We are indebted to the anonymous referees whose comments helped us to improve the presentation considerably.

This work has received partial support from the Italian MIUR project “Constraint Based Verification of Reactive Systems” (COVER) and from Enterprise Ireland under their Basic Research Grant Scheme (Grant Numbers SC/02/289 and SC/2003/007) and their International Collaboration Programme (Grant Number IC/2003/88).

References

- [1] J. Bailey and P.J. Stuckey. Discovery of minimal unsatisfiable subsets of constraints using hitting-set dualization. In *Proc. Practical Aspects of*

- Declarative Languages (PADL 2005)*, volume 3350 of *Lecture Notes in Computer Science*, pages 174–186. Springer, 2005.
- [2] G. Bella and S. Bistarelli. Soft constraint programming to analysing security protocols. *Theory and Practice of Logic Programming (TPLP)*, (Special Issue on Verification and Computational Logic), 4(5):1–28, 2004.
 - [3] V.G. Bharadwaj and J.S. Baras. Towards automated negotiation of access control policies. In *Proc. of IEEE Workshop Policies for Distributed Systems and Networks*, pages 77–80, June 2003.
 - [4] S. Bistarelli. *Semirings for Soft Constraint Solving and Programming*, volume LNCS 2962. Springer, 2004.
 - [5] S. Bistarelli, H. Fargier, U. Montanari, F. Rossi, T. Schiex, and G. Verfaille. Semiring-based CSPs and Valued CSPs: Frameworks, properties, and comparison. *CONSTRAINTS: An international journal. Kluwer*, 4(3):199–240, 1999.
 - [6] S. Bistarelli and S.N. Foley. Analysis of integrity policies using soft constraints. In *Proc. of IEEE Workshop Policies for Distributed Systems and Networks*, pages 77–80, June 2003.
 - [7] S. Bistarelli and S.N. Foley. A constraint based framework for dependability goals: Integrity. In *Proc. of 22nd International Conference on Computer Safety, Reliability and Security (SAFECOMP2003)*, volume 2788 of *Lecture Notes in Computer Science*, pages 130–143. Springer, 2003.
 - [8] S. Bistarelli, S.N. Foley, and B. O’Sullivan. Detecting and eliminating the cascade vulnerability problem from multi-level security networks using soft constraints. In *Proceedings of AAAI/IAAI-2004 (16th Innovative Applications of AI Conference)*, pages 808–813. AAAI Press San Jose, July 2004.
 - [9] S. Bistarelli, S.N. Foley, and B. O’Sullivan. Modelling and detecting the cascade vulnerability problem using soft constraints. In *Proceedings of ACM Symposium on Applied Computing (SAC-2004)*, pages 383–390. ACM Press, March 2004.
 - [10] S. Bistarelli, S.N. Foley, and B. O’Sullivan. Reasoning about secure inter-operation using soft constraints. In *Proceedings of FAST-2004 Workshop on Formal Aspects of Security and Trust*, 2004.
 - [11] S. Bistarelli, U. Montanari, and F. Rossi. Semiring-based Constraint Solving and Optimization. *JACM*, 44(2):201–236, 1997.
 - [12] J.L. de Siqueira N. and J.-F. Puget. Explanation-based generalisation of failures. In *Proc. 8th European Conference on Artificial Intelligence, ECAI 88*, pages 339–344. Pitmann Publishing, 1988.
 - [13] R. Dechter. *Constraint Processing*. Morgan Kaufmann, 2003.

- [14] J.A. Fitch and L.J. Hoffman. A shortest path network security model. *Computers and Security*, 12(2):169–189, 1993.
- [15] S.N. Foley. Conduit cascades and secure synchronization. In *ACM New Security Paradigms Workshop*, 2000.
- [16] L. Gong and X. Qian. The complexity and composability of secure interoperation. In *Proceedings of the Symposium on Security and Privacy*, pages 190–200, Oakland, CA, May 1994. IEEE Computer Society Press.
- [17] S. Gritalis and D. Spinellis. The cascade vulnerability problem: The detection problem and a simulated annealing approach to its correction. *Microprocessors and Microsystems*, 21(10):621–628, 1998.
- [18] J.D. Horton, R. Harland, E. Ashby, R.H. Cooper, W.F. Hyslop, B.G. Nickerson, W.M. Stewart, and O.K. Ward. The cascade vulnerability problem. *Journal of Computer Security*, 2(4):279–290, 1993.
- [19] D.S. Johnson. Approximation algorithms for combinatorial problems. *Journal of Computer and System Sciences*, 9:256–278, 1974.
- [20] U. Junker. QUICKXPLAIN: Conflict detection for arbitrary constraint propagation algorithms. In *IJCAI’01 Workshop on Modelling and Solving problems with constraints*, Seattle, WA, USA, August 2001.
- [21] V. Kumar. Algorithms for constraint-satisfaction problems : A survey. *AI Magazine*, 13(1):32–44, 1992.
- [22] L. Lovasz. On the ratio of optimal integral and fractional covers. *Discrete Mathematics*, 13:383–390, 1975.
- [23] J.K. Millen and M.W. Schwartz. The cascading problem for interconnected networks. In *4th Aerospace Computer Security Applications Conference*, pages 269–273, December 1988.
- [24] A. Di Pierro, C. Hankin, and H. Wiklicky. On approximate non-interference. In *Proc. of Workshop on Issues in the Theory of Security*. IFIP WG1.7, 2002.
- [25] J.-C. Regin. A filtering algorithm for constraints of difference in cps. In *Proceedings AAAI-94*, pages 362–367, 1994.
- [26] T. Schiex. Possibilistic constraint satisfaction problems, or “how to handle soft constraints?”. In *Proc. 8th Conf. of Uncertainty in AI*, pages 269–275, 1992.
- [27] P. Slavík. A tight analysis of the greedy algorithm for set cover. In *Proc. of the 28th ACM Symposium on Theory of Computing (STOC)*, pages 435–441, 1996.

- [28] TCSEC. Computer security requirements – guidance for applying the department of defense trusted computer system evaluation criteria in specific environments. Technical Report CSC-STD-003-85, National Computer Security Center, 1985. Orange Book.
- [29] TNI. Trusted computer system evaluation criteria: Trusted network interpretation. Technical report, National Computer Security Center, 1987. Red Book.
- [30] M. Wallace. Practical applications of constraint programming. *Constraints*, 1(1–2):139–168, 1996.