Constraint Differentiation: A New Reduction Technique for Constraint-Based Analysis of Security Protocols*

David Basin
Dep. of Computer Science
ETH Zurich
CH-8092 Zurich, Switzerland
basin@inf.ethz.ch

Sebastian Mödersheim
Dep. of Computer Science
ETH Zurich
CH-8092 Zurich, Switzerland
moedersheim@inf.ethz.ch

Luca Viganò
Dep. of Computer Science
ETH Zurich
CH-8092 Zurich, Switzerland
vigano@inf.ethz.ch

ABSTRACT

We introduce constraint differentiation, a new technique for reducing search when model-checking security protocols. Our technique is based on eliminating certain kinds of redundancies that arise in the search space when using symbolic exploration methods, in particular methods that employ constraints to represent and manipulate possible messages from an active intruder. Formally, we prove that constraint differentiation terminates and is correct and complete, in that it preserves the set of reachable states so that all state-based properties holding before reduction (such as the intruder discovering a secret on the network) hold after reduction. Practically, we have integrated this technique into OFMC, a state-of-the-art model-checker, and demonstrated its effectiveness by extensive experimentation. Our results show that constraint differentiation substantially reduces search and considerably improves the performance of OFMC, enabling its application to a wider class of problems.

Categories and Subject Descriptors

C.2.2 [**Computer-Communication networks**]: Network Protocols—*Protocol verification*; D.2.4 [**Software Engineering**]: Software/Program Verification—*Formal methods, Model checking, Validation*

General Terms

Security, Verification

Keywords

Protocol Verification, Constraints, Partial-Order Reduction

*This work was supported by the FET Open Project IST-2001-39252, "AVISPA: Automated Validation of Internet Security Protocols and Applications".

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
CCS'03, October 27–30, 2003, Washington, DC, USA.
Copyright 2003 ACM 1-58113-738-9/03/0010 ...$5.00.

1. INTRODUCTION

Context. A wide variety of model-checking approaches have recently been developed to analyze security protocols, e.g. [5, 8, 11, 13, 20, 21, 23]. The challenge faced when building such a checker is to handle two kinds of state-explosion. The first kind is caused by the standard Dolev-Yao [12] model that defines an infinite set of messages that the intruder can generate.[1] The second kind stems from the large number of possible interleavings resulting from parallel executions of a protocol by the honest agents and the intruder.

A number of constraint-based approaches have been proposed to tackle the first problem, which employ symbolic representation techniques [16, 1, 20, 14, 8, 11, 7, 5]. Although there are differences between these approaches, they all share in common a symbolic representation of the state space, i.e. sets of (ground) states are represented by terms with variables and constraints on the variables. These constraints describe what terms an intruder must generate from a given set of known messages according to the Dolev-Yao model. Moreover, all these approaches use similar reduction rules to manipulate the constraints; in particular, constraint reduction is demand-driven ("lazy") in the sense that one checks only if a solution for the constraints exists, rather than exploring all solutions. Therefore, we will refer to the technique underlying these constraint-based approaches as the *lazy intruder technique*.

There are also structural similarities in the different realizations of the lazy intruder as we can generally distinguish two *layers of search*: the first layer is a search in the space of constraints to find all (symbolic) solutions for a given set of constraints according to the Dolev-Yao model; the second layer builds on the first one to search in the symbolic search space to determine if an *attack state* is reachable. The first kind of search is performed in a backward fashion ("Can the terms to be generated be reduced to terms the intruder knows?"), while the second kind is performed in a forward fashion ("What states are reachable from the initial state?"). While the different approaches implement the first layer in similar ways, there are differences in the formalism employed to represent the second layer, e.g. [5, 8] use *multiset rewriting* to represent the symbolic state space and its

[1] Using restrictions on the form of messages the intruder may generate (e.g. by allowing only type-correct messages), it is possible to obtain a finite model of the intruder. However, the branching factor of the search tree induced by such a restricted Dolev-Yao intruder is typically still enormous.

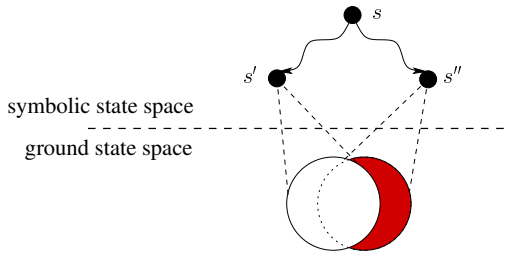

Figure 1: The intuition behind constraint differentiation.

transition relation, [11, 20] use *strand spaces*, and [1, 7] use *process calculi*.

The lazy intruder drastically reduces the search tree generated during protocol analysis, providing an effective solution to the first problem mentioned above, the prolific Dolev-Yao intruder. However, the second problem is still open, namely reducing the state-explosion that results from the exponential number of interleavings induced by parallel protocol executions.

In standard model-checking approaches for concurrent systems, the state-explosion problem is often handled using *partial-order reduction* (*POR*), a technique that reduces the number of interleavings that need to be considered by exploiting independencies between the possible transitions [22]. One might expect that the direct combination of the lazy intruder with partial-order reduction as part of the second layer of the search (i.e. while searching the symbolic state space) would allow us to simultaneously "solve" both state-explosion problems. However, as we will explain below, this combination is not effective: the different transitions of the lazy intruder rarely lead to the same successor state, and therefore there is practically no independence of transitions that can be exploited by POR.

Contribution. We show here how to effectively integrate the lazy intruder and ideas from POR by using independence information from the second layer, i.e. the symbolic transition system, when searching the first layer, i.e. the constraint reduction. We call the resulting technique *constraint differentiation*.

Figure 1 illustrates the intuition behind constraint differentiation. Consider two symbolic states s' and s'' that can be reached from some state s by different interleavings of the same actions, e.g. message exchanges. In practice, although s' and s'' represent different sets of ground states, these sets often substantially overlap. Constraint differentiation exploits these overlaps by restricting the constraints of s'' to those ground states that are not already covered by s', i.e. the shaded part in Figure 1. (Symmetrically, we could restrict s' instead of s''.)

Our work has both theoretical and practical relevance. Theoretically, our contribution is the formalization of constraint differentiation as a search reduction technique integrating the lazy intruder with ideas from POR. We formally prove that constraint differentiation terminates and is correct and complete, in the sense that it preserves the set of reachable states so that all state-based properties holding before reduction (such as the existence of an attack) hold after reduction. It follows that constraint differentiation neither excludes attacks nor introduces new ones.

Constraint differentiation is independent of the technical and conceptual details of the various lazy intruder approaches and underlying protocol models, and can thus be adopted in other approaches. In particular, even though for concreteness we refer to our own variant of the lazy intruder technique [5], it is not difficult to carry over constraint differentiation to other approaches, as the different formalisms for constraint reduction in the first layer of search are all similar. For the second layer of search, we abstract away from particular base formalisms and work only with the requirement that the state space is represented symbolically, in the sense that it can be expressed using terms with variables, and attacks are formulated as reachability problems. To this end, we introduce the notion of a *symbolic transition system*, which provides an interface to the different formalisms.

Our theoretical results have immediate practical applications. As a concrete example, we have integrated the proposed constraint differentiation into our on-the-fly model-checker OFMC [5], a state-of-the-art tool for security protocol analysis based on the lazy intruder. This integration generally reduces the search time by a factor 2 to several orders of magnitude. More importantly, this improvement extends the scope of OFMC so that it not only scales well to the falsification of industrial-strength protocols, but it can also be applied as an effective verification tool since our reductions make it feasible to exhaustively search the (symbolic) state space resulting from several parallel executions of the protocol being analyzed. We have carried out extensive experiments to validate our approach. As a relevant example, we report on our results for the protocol-suite IKE.

Organization. We proceed as follows. In §2 we formalize the lazy intruder and symbolic transition systems. In §3 we formalize the integration of constraint differentiation into this symbolic protocol model. We present experimental results in §4 and draw conclusions and discuss future work in §5. Due to lack of space, examples and proofs have been shortened or omitted; details can be found in [6].

2. CONSTRAINT-BASED PROTOCOL MODELS

We introduce the basic constraint-based model underlying our approach in a bottom-up fashion, starting with the first layer of the search, the constraints and their reduction. Then, for the second layer, we introduce symbolic transition systems by abstracting away the details of the formalism used to represent the symbolic state space.

We will proceed as abstractly as possible, without limiting the design choices of the approach, and will only commit to particular choices (based on [5]) to make the exposition concrete.

2.1 Lazy Intruder and Constraint Reduction

As is standard, messages in our model are elements of an algebra $T_\Sigma(\mathcal{V})$, the set of *terms* built from a signature Σ and a set of variables $\mathcal{V}$. The signature Σ contains the operators for building messages, and for brevity we restrict our presentation here to *pairing*, denoted by $\langle m_1, m_2 \rangle$, and *symmetric encryption*, denoted by $\{\!|m_1|\!\}_{m_2}$; note that arbitrary messages can be used as keys, allowing us to analyze proto-

cols with non-atomic keys. It is straightforward to extend the methods presented in this paper to other operators like asymmetric encryption, hashing, and key-tables. Note that we do not associate any type information with messages in order to allow the detection of type-flaw attacks. Also, like most other approaches, we employ the *free algebra assumption* and assume that syntactically different terms represent different messages.

The notions of *atomic* and *composed message* are defined in the usual way and so are the notions related to substitution and unification, such as *ground term*, *most general unifier (mgu)*, and *matching*; see, e.g., [3]. We denote the application of a substitution σ to a term t by writing $t\sigma$ and denote the composition of substitutions σ_1 and σ_2 by $\sigma_1\sigma_2$, i.e. $t(\sigma_1\sigma_2) = (t\sigma_1)\sigma_2$. The *identity substitution id* is the substitution with $domain(id) = \emptyset$. We say that two substitutions σ_1 and σ_2 are *compatible*, written $\sigma_1 \approx \sigma_2$, if $v\sigma_1 = v\sigma_2$ for every $v \in domain(\sigma_1) \cap domain(\sigma_2)$. For two sets of ground substitutions Σ_1 and Σ_2, we define their *intersection modulo the different domains* as $\Sigma_1 \sqcap \Sigma_2 = \{\sigma_1\sigma_2 \mid \sigma_1 \in \Sigma_1 \wedge \sigma_2 \in \Sigma_2 \wedge \sigma_1 \approx \sigma_2\}$. Since the composition of compatible ground substitutions is associative and commutative, so is the $\sqcap$ operator.

We follow Dolev and Yao [12] and consider the standard asynchronous intruder model where the intruder controls the network but cannot break cryptography. In particular, the intruder can intercept messages and analyze them if he possesses the corresponding keys for decryption, and he can generate messages from his knowledge and send them under any agent name.

For a set M of messages, we define $\mathcal{DY}(M)$ (for Dolev-Yao) to be the smallest set closed under the *generation (G)* and *analysis (A) rules* in Figure 2. The generation rules express that the intruder can compose messages from known messages using pairing and symmetric encryption, while the analysis rules describe how he can decompose messages.

The Dolev-Yao intruder leads to an enormous branching of the search tree when one naïvely enumerates all (meaningful) messages that the intruder can send. The lazy intruder technique significantly reduces the search tree without excluding any potential attacks, by exploiting the fact that the actual value of certain parts of a message is often irrelevant for the receiver. So, whenever the receiver will not further analyze the value of a particular message part, we can postpone during the search the decision about which value the intruder actually chooses for this part by replacing it with a variable and recording a *constraint* on which knowledge the intruder uses to generate the message. We express this information using constraints of the form $from(T, IK)$, meaning that T is a set of terms generated by the intruder from the set of his known messages IK (for "intruder knowledge").

DEFINITION 1. *A* from constraint *(or simply* constraint*) has the form* $from(T, IK)$*, where T and IK are sets of terms. Its semantics is the set of satisfying ground substitutions σ for the variables in the constraint, i.e.* $[\![from(T, IK)]\!] = \{\sigma \mid ground(\sigma) \wedge ground(T\sigma \cup IK\sigma) \wedge (T\sigma \subseteq \mathcal{DY}(IK\sigma))\}$.

A constraint c is satisfiable *if* $[\![c]\!] \neq \emptyset$. *A* (from) constraint set *is a finite set of* (from) *constraints and its semantics is the intersection of the semantics of its elements modulo the different domains, i.e., overloading the $[\![\cdot]\!]$ notation,* $[\![\{c_1, \ldots, c_n\}]\!] = \sqcap_{i=1}^{n} [\![c_i]\!]$. *A constraint* $from(T, IK)$ *is* simple*, written* $simple(from(T, IK))$*, if $T \subseteq \mathcal{V}$. A constraint set is simple if all its constraints are simple. A constraint set C is* well-formed *if one can index its constraints,* $C = \{from(T_1, IK_1), \ldots, from(T_n, IK_n)\}$*, so that the following conditions hold:* (1) $IK_i \subseteq IK_j$ *for* $i \leq j$*, and* (2) $vars(IK_i) \subseteq \cup_{j=1}^{i-1} vars(T_j)$.

Intuitively, (1) requires that the intruder knowledge increases monotonically and (2) requires that every variable that appears in intruder-known terms is part of a message that the intruder created earlier, i.e. variables only "originate" from the intruder. (As we explain in [6], these properties are crucial for the completeness of the reduction functions *Red* and *D-Red* that we introduce below.) All constraint sets that we consider in the remainder of this paper are well-formed, unless stated otherwise.

The core of the lazy intruder technique is to reduce a given constraint set into an equivalent one that is either unsatisfiable or simple (and thus satisfiable, as it is straightforward to show that every simple constraint set is satisfiable since the intruder can always generate some message, e.g. his own name). In Figure 3 we give the generation and analysis rules of [5], which describe how constraint sets can be reduced (again, we consider only pairing and symmetric encryption and decryption for brevity). Note that here and elsewhere, we simplify notation for singletons, writing, e.g., $m_2 \cup IK$ for $\{m_2\} \cup IK$.

A generation or analysis rule r has the form $\dfrac{C', \sigma'}{C, \sigma}\ r$, with C and C' constraint sets and σ and σ' substitutions, and it expresses that (C', σ') can be *derived* from (C, σ), which we denote by $(C, \sigma) \vdash_r (C', \sigma')$. That is, the constraint reduction rules are applied backwards. We extend the semantics function to such pairs, $[\![(C, \sigma)]\!] = \{\sigma\sigma' \mid \sigma' \in [\![C]\!]\}$, and we extend it pointwise to sets of such pairs.

The generation rules G^L_{pair} and G^L_{scrypt} express that the constraint stating that the intruder can generate a message composed from submessages m_1 and m_2 using pairing or symmetric encryption can be replaced by the constraint stating that he can generate both m_1 and m_2. The rule G^L_{unif} expresses that the intruder can use a message m_2 from his knowledge if this message can be unified with the message m_1 he has to generate (note that both the terms to be generated and the terms in the intruder knowledge may contain variables). The reason that the intruder is "lazy" stems from the restriction that the G^L_{unif} rule cannot be applied when the term to be generated is a variable: the intruder's choice for this variable does not matter at this stage of the search and hence we postpone this choice.

The analysis of the intruder knowledge is complex as messages may contain variables and hence the applicability of an analysis step may depend on the substitution for the variables. A straightforward way to express analysis in the lazy intruder setting is given by the rule A^L_{scrypt}: for a message $\{\!|m_1|\!\}_{m_2}$ that the intruder attempts to decrypt, we add the content m_1 to the intruder knowledge of the respective constraint (as if the check succeeded) and add a constraint expressing that the symmetric key m_2 needed for decryption must be generated from the same knowledge. This constraint renders the constraint set unsatisfiable when m_2 cannot be generated using the corresponding intruder knowledge.[2]

Note that, in the A^L_{scrypt} rule, we also make the restriction

[2]Note that the side conditions prevent the repeated application of the analysis rules to the same term.

$$\frac{m \in M}{m \in \mathcal{DY}(M)}\ G_{\text{axiom}}\,, \quad \frac{m_1 \in \mathcal{DY}(M) \quad m_2 \in \mathcal{DY}(M)}{\langle m_1,m_2\rangle \in \mathcal{DY}(M)}\ G_{\text{pair}}\,, \quad \frac{m_1 \in \mathcal{DY}(M) \quad m_2 \in \mathcal{DY}(M)}{\{|m_1|\}_{m_2} \in \mathcal{DY}(M)}\ G_{\text{scrypt}}\,,$$

$$\frac{\langle m_1,m_2\rangle \in \mathcal{DY}(M)}{m_i \in \mathcal{DY}(M)}\ A_{\text{pair}_i}\,, \quad \frac{\{|m_1|\}_{m_2} \in \mathcal{DY}(M) \quad m_2 \in \mathcal{DY}(M)}{m_1 \in \mathcal{DY}(M)}\ A_{\text{scrypt}}\,.$$

Figure 2: Dolev-Yao intruder rules

$$\frac{(\mathit{from}(M, m_2 \cup IK) \cup C)\tau, \sigma\tau}{\mathit{from}(m_1 \cup M, m_2 \cup IK) \cup C, \sigma}\ G^L_{\text{unif}}\ (\tau = mgu(m_1, m_2),\ m_1 \notin \mathcal{V})\,, \quad \frac{\mathit{from}(m_1 \cup m_2 \cup M, IK) \cup C, \sigma}{\mathit{from}(\langle m_1,m_2\rangle \cup M, IK) \cup C, \sigma}\ G^L_{\text{pair}}\,,$$

$$\frac{\mathit{from}(m_1 \cup m_2 \cup M, IK) \cup C, \sigma}{\mathit{from}(\{|m_1|\}_{m_2} \cup M, IK) \cup C, \sigma}\ G^L_{\text{scrypt}}\,, \quad \frac{\mathit{from}(M, m_1 \cup m_2 \cup \langle m_1,m_2\rangle \cup IK) \cup C, \sigma}{\mathit{from}(M, \langle m_1,m_2\rangle \cup IK) \cup C, \sigma}\ A^L_{\text{pair}}\ (\{m_1, m_2\} \setminus IK \neq \emptyset),$$

$$\frac{\mathit{from}(m_2, IK) \cup \mathit{from}(M, m_1 \cup \{|m_1|\}_{m_2} \cup IK) \cup C, \sigma}{\mathit{from}(M, \{|m_1|\}_{m_2} \cup IK) \cup C, \sigma}\ A^L_{\text{scrypt}}\ (m_1 \notin IK).$$

Figure 3: Lazy intruder: constraint reduction rules

that the message $\{|m_1|\}_{m_2}$ analyzed may not be used in the generation of the possibly non-atomic key m_2. This is in contrast to similar approaches that, like ours, can handle non-atomic symmetric keys, e.g. [20, 8]. The fact that our restriction does not exclude any solutions is shown as part of the proof of Theorem 1, on the correctness and completeness of constraint reduction, which is given in [6].

Definition 2. *Let $\vdash$ denote the reflexive and transitive closure of the union of the derivation relations $\vdash_r$ for every rule r of Figure 3. The set of pairs of simple constraint sets and substitutions that can be derived from (C, id) is $Red(C) = \{(C', \sigma) \mid ((C, id) \vdash (C', \sigma)) \wedge simple(C')\}$. We call Red the* constraint reduction function.

Theorem 1. *If C is a well-formed constraint set, then $Red(C)$ is finite and $\vdash$ is well-founded, and $[\![C]\!] = [\![Red(C)]\!]$.*

By Theorem 1 we have that *Red* is correct, complete, and recursively computable (since $\vdash$ is finitely branching).

2.2 Symbolic Transition Systems

The various implementations of the lazy intruder are based on different formalisms to represent the second layer of the search, namely the symbolic state space and its transition relation. For example, [5, 8] use multiset rewriting, [11, 20] use strand spaces, and [1, 7] use process calculi. Constraint differentiation is not specialized to any of these approaches. It only requires a state space where states are represented symbolically using terms with variables, together with a state transition function and a goal predicate. Therefore we define the concept of a symbolic transition system as an abstract interface to the different symbolic formalisms.

The idea behind the symbolic approaches is to use symbolic states, i.e. terms with variables, to represent sets of ground states, i.e. sets of terms without variables. The permissible substitutions for variables are described by some kind of constraints, in our case *from* constraints. The semantics of a particular constraint set (i.e. the set of substitutions allowed by the constraints in the set) can then be extended to a symbolic state s (i.e. the set of ground terms represented by s). One can then define a state transition function as usual, but it must agree with the semantics of each symbolic state in the sense that equivalent symbolic states have equivalent successors and are equivalent with respect to a predicate, which in our case describes attacks. Formally:

Definition 3. *A* symbolic transition system *over a countable set Σ of* constant *and* function *symbols and a countable set $\mathcal{V}$ of variables is a 5-tuple $(\mathcal{G}, \mathcal{S}, \mathcal{I}, \mathcal{T}, \mathcal{P})$, where $\mathcal{G} = T_\Sigma$ is the* set of ground states*; $\mathcal{S} = T_\Sigma(\mathcal{V}) \times \mathcal{CS}$ is the* set of symbolic states*, where $\mathcal{CS}$ denotes the* set of all well-formed *from* constraint sets*; $\mathcal{I} \in \mathcal{S}$ is the* initial symbolic state*; $\mathcal{T} : \mathcal{S} \to 2^{\mathcal{S}}$ is a* transition function *on symbolic states; and $\mathcal{P}$ is a predicate on symbolic states. We also refer to $\mathcal{G}$ and $\mathcal{S}$ as the* spaces *of ground and symbolic states, respectively, and call $\mathcal{P}$ the* attack predicate *as it will be used to specify symbolic states representing attacks.*

The semantics of a symbolic state $s = (t, C)$ is defined in terms of the semantics of the constraints: $[\![(t, C)]\!] = \{t' \mid \exists\sigma.\ \sigma \in [\![C]\!] \wedge t' = t\sigma\}$. We straightforwardly extend $\mathcal{T}$, $\mathcal{P}$, and $[\![\cdot]\!]$ pointwise to sets of symbolic states. The symbolic transition system must agree with the semantics of the symbolic states in the following sense: for two sets of symbolic states S_1 and S_2 with $[\![S_1]\!] = [\![S_2]\!]$, we require that $[\![\mathcal{T}(S_1)]\!] = [\![\mathcal{T}(S_2)]\!]$ and $\mathcal{P}(S_1) \iff \mathcal{P}(S_2)$.

The set of reachable symbolic states *is the smallest set that contains the initial symbolic state and is closed under the transition function. A symbolic transition system is* secure *iff no reachable symbolic state satisfies the attack predicate.*

The constraint reduction is extended to symbolic states by $Red((t, C)) = \{(t', C') \mid \exists\sigma.\ (C', \sigma) \in Red(C) \wedge t' = t\sigma\}$, so that we straightforwardly have that $[\![Red(s)]\!] = [\![s]\!]$ for every symbolic state s.

The other lazy intruder approaches can be easily recast as symbolic transition systems. For example, in the concrete case of the model underlying the tool OFMC [5], the ground states in $\mathcal{G}$ are multisets of facts which express the local state of agents, the intruder knowledge, and the messages sent on the network that are not yet received. The symbolic states in $\mathcal{S}$ are also multisets of facts, but message terms may contain variables; hence, a symbolic state represents the set of ground states that are obtained by a ground substitution for the variables. The transition function on symbolic states $\mathcal{T}$ is induced by multiset rewrite

rules that describe the behavior of the honest agents, and add constraints to the constraint set when the intruder has to generate a new message (note that by the construction of this transition function it is ensured that in all reachable symbolic states the constraint set is well-formed). Finally, the attack predicate $\mathcal{P}$ is used to express state-based properties of symbolic states. In OFMC, we use $\mathcal{P}$ to formulate standard authentication and secrecy goals, but in the abstract symbolic transition system above we commit neither to particular kinds of attacks nor to a particular formalism to specify attacks (any of the other formalisms previously listed, e.g. strands, could be used), as long as they are formalized as reachability problems.

A symbolic transition system gives rise to a search tree where the root node is the initial state and the children of a node are all states that can be reached with one transition. For every symbolic state, applying the function Red yields a set of equivalent symbolic states with simple constraint sets. This can be exploited for reduction since if the constraint set C of a symbolic state is unsatisfiable, then $Red(C) = \emptyset$ by Theorem 1. In this case, we can safely prune the subtree of the search tree node containing the unsatisfiable constraint. In the next section, we will see that the integration of constraint differentiation into the symbolic transition system is based on a similar form of pruning of the search tree.

Before considering this integration in detail, let us first observe that the lazy intruder can be straightforwardly extended with a technique that we call *step-compression*, which leads to a significant reduction of the search tree without excluding any attacks.[3] Step-compression is based on the idea that since the intruder completely controls the communication network, we can safely assume that every message from an honest agent is automatically intercepted by the intruder (who can always play it back into the network) and that every message that an honest agent receives comes from the intruder. This allows us to restrict the search to transitions where two steps are merged (or "compressed") into one: first, the intruder sends a message to an honest agent and second, the intruder intercepts the agent's reply.

When step-compression is used, the symbolic transition system has the following property: for every transition from a symbolic state $s_1 = (t_1, C_1)$ to a symbolic state $s_2 = (t_2, C_2)$, the constraint sets will have the form $C_2 = C_1 \cup from(m_1, IK)$ for some message m_1, representing the message the intruder sends to an honest agent, and a set of messages IK, representing the knowledge the intruder can use to generate m_1. Also, the intruder knowledge in s_2 is augmented by the agent's reply. We will make use of this property in the constraint differentiation technique.

3. Constraint Differentiation

The lazy intruder technique allows us to significantly reduce the search tree generated by the prolific Dolev-Yao intruder without excluding any attacks (cf. [5]). In particular, our experiments have shown that the search tree induced by the lazy intruder often has roughly the same size as the tree that would be searched when considering a *passive* intruder, i.e. one that listens to the communication on the network but does not manipulate or generate any messages. This is the maximal reduction possible since the search tree of reachable symbolic states must cover all "legal" executions of the protocol between honest agents.

However, even when considering only a small number of sessions that can be executed in parallel, searching the tree that contains all the interleavings of these sessions can be infeasible; we are faced with the standard state-explosion problem of model-checking. In model-checking approaches for concurrent systems, this problem is often handled using *partial-order reduction* (POR), a technique that reduces the number of interleavings that need to be considered by exploiting independencies between the possible transitions [22].

One might expect that the direct combination of the lazy intruder with partial-order reduction would allow us to simultaneously "solve" both state-explosion problems, that of the prolific Dolev-Yao intruder and that of the large number of possible interleavings. However, their direct combination is not effective: the different transitions of the lazy intruder rarely lead to the same successor state and therefore there is practically no independence of transitions that can be exploited by POR. We present a way around this problem by directly using independence information in the constraint reduction; the result is a POR-inspired reduction technique that we call *constraint differentiation*.

To see why the direct combination of partial-order reduction with the lazy intruder is not effective, consider the successor function of the search tree that results from the symbolic transition system (with step-compression) given in the previous section. A direct application of POR would require identifying situations of the form depicted in Figure 4. There are two sequences of transitions. In the left one, the intruder i first sends a message m_1 to an agent a, receiving the answer m_2, and afterwards he sends a message m_3 to an agent b, receiving the answer m_4. In the right sequence, the intruder first talks to b and then to a. The transitions result in the states $s_2 = (t_2, C_2)$ and $s_4 = (t_4, C_4)$, for terms (with variables) t_2 and t_4 and constraint sets C_2 and C_4. We consider the case $t_2 = t_4$, which holds when the transitions are independent in the sense that on ground states the respective order of operations would lead to the same successor states. In this case, for every substitution σ, the represented ground states $t_2\sigma$ and $t_4\sigma$ are the same; however, the constraints, determining the set of permissible substitutions, are different due to the fact that the intruder generated the respective messages at different states of his knowledge. Hence, the direct combination of partial-order reduction with the lazy intruder is not effective.[4]

The observation that leads to a reduction is that in this situation there is an overlapping of the set of ground states, which is represented by the two symbolic states s_2 and s_4, as shown by the shaded part in Figure 4: all those ground states in the semantics of the symbolic states that do not exploit the new knowledge m_2 or m_4 are covered by the other symbolic state. The idea is that we can use independence of transitions by exploiting precisely this overlap. If, for example, we "prefer" the left sequence, then for the state s_4 reached by the other sequence we will only be interested

[3] Note that step-compression is applied not only in all symbolic approaches we know of [1, 5, 7, 11, 14, 16, 20], but also in some non-symbolic approaches, e.g. [2].

[4] One may wonder whether without step-compression, i.e. without this property of the constraints, POR could be more effective. One can show that directly applying POR to a symbolic transition system without step-compression can only result in reductions that are also achieved using step-compression.

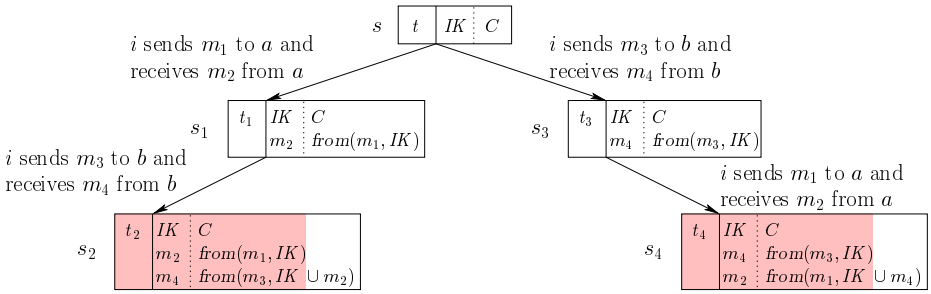

Figure 4: An illustration of constraint differentiation for $t_2 = t_4$ (for each symbolic state s we display t and C as well as the corresponding intruder knowledge IK).

in solutions that are not subsumed by s_2 already, i.e. those where the intruder actually uses the message m_4 that he learned in the first transition to generate the message m_1 in the second transition.

In this way, we can propagate information about independencies obtained on the second search layer, the symbolic transition system, to the first layer, the constraint reduction. We exploit the fact that we only need to consider solutions for a given constraint set that are obtained by using new knowledge. In the example, we could express the fact that the message m_4 needs to be used when creating m_1 by using constraints of the form $D\text{-}from(m_1, IK, m_4)$, which intuitively has the same meaning as the constraint $from(m_1, IK \cup m_4)$, except that we exclude all solutions of $from(m_1, IK)$.

Mirroring the development of §2, we proceed as follows: we introduce constraint differentiation by defining this new kind of constraint (D-from constraints), the associated reduction rules (D-from rules), and the constraint reduction function (*D-Red*) that describes the constraints that can be derived with these rules. We then show that the *D-Red* function has properties analogous to those of the *Red* function, namely it is correct, complete, and recursively computable. We conclude the section by integrating the new constraint reduction into the second layer of the search as a transformation of the search tree induced by the symbolic transition system.

3.1 Constraint Reduction with Constraint Differentiation

We now extend Definition 3 of a symbolic transition system with the following form of constraints.

Definition 4. *A* D-from constraint *c has the form*

$$D\text{-}from(T, IK, NIK)\,,$$

where T, IK, and NIK are sets of messages. Its semantics is $[\![c]\!] = [\![\lceil c \rceil]\!] \setminus [\![\lfloor c \rfloor]\!]$, where $\lceil D\text{-}from(T, IK, NIK) \rceil = from(T, IK \cup NIK)$ and $\lfloor D\text{-}from(T, IK, NIK) \rfloor = from(T, IK)$ are functions mapping D-from constraints to from constraints. We say that $D\text{-}from(T, IK, NIK)$ is simple *if $T \subseteq \mathcal{V}$ and $T \neq \emptyset$.*

Extending Definition 1, we now write simply constraint *to refer to a from constraint or a D-from constraint. The semantics of a constraint set $C = \{c_1, \ldots, c_n\}$ is $[\![\{c_1, \ldots, c_n\}]\!] = \sqcap_{i=0}^{n} [\![c_i]\!]$. The definitions of $\lceil \cdot \rceil$, $\lfloor \cdot \rfloor$, and simple are extended straightforwardly to constraint sets, and a constraint set C is* well-formed *iff both $\lceil C \rceil$ and $\lfloor C \rfloor$ are well-formed.*

Intuitively, *NIK* represents new messages that are not in *IK* (it is an acronym for *new intruder knowledge*). As we will use them, *NIK* and *IK* will be disjoint, so that the constraint $D\text{-}from(T, IK, NIK)$ states that the set of terms T must be generated by the intruder using the knowledge in the set $IK \cup NIK$, but we are only interested in solutions that employ new information in *NIK*; hence, we exclude all solutions of $from(T, IK)$. Note that *from* constraints are a special case of *D-from* constraints since $from(T, IK) = D\text{-}from(T, \emptyset, IK)$. The function $\lceil \cdot \rceil$ yields a from constraint set by removing the requirement on new knowledge; therefore $[\![C]\!] \subseteq [\![\lceil C \rceil]\!]$. Similarly, $[\![\lfloor C \rfloor]\!] \cup [\![C]\!] = [\![\lceil C \rceil]\!]$, as $\lfloor \cdot \rfloor$ returns the solutions removed from $\lceil C \rceil$. Note that for a simple constraint set C both $\lceil C \rceil$ and $\lfloor C \rfloor$ are simple (and hence satisfiable). Note that if a constraint set C is simple then $\lceil C \rceil$ is satisfiable.[5]

Let us extend the set of lazy intruder reduction rules of Figure 3 with the reduction rules for the new D-from constraints, for short *D-from rules*, which we display in Figure 5. As representative cases, we explain the three most interesting D-from rules: G^{LD}_{unif2}, A^{LD}_{scrypt2}, and A^{LD}_{scrypt3}. The rules for the analysis of pairs are similar to those for the analysis of encryption and the other rules are simply "liftings" of the original rules to D-from constraints.

As shorthand, we use the terms *T-part*, *IK-part*, and *NIK-part* to refer to the respective parts of a D-from constraint $D\text{-}from(T, IK, NIK)$. The G^{LD}_{unif2} rule expresses the case where the term m_1, that should be generated, can be unified with a term m_2 in the *NIK*-part. Unless m_1 can be derived from the *IK*-part, the condition that some message from the *NIK*-part must be used is satisfied, and thus we are not obliged to use messages in *NIK* for generating further terms in the *T*-part. However, if m_1 can be derived from *IK*, then the application of the rule is inefficient (as it removes a possible restriction), but still correct in the sense that the resulting solution is valid.

The A^{LD}_{scrypt2} rule describes the analysis of an encrypted term $\{\!|m_1|\!\}_{m_2}$ of the *IK*-part, where m_1 is not yet in the *IK*-part of the D-from constraint. Suppose the key m_2 can be derived using the intruder knowledge $IK \cup NIK$ but not from *IK* alone, then m_1 can only be derived when using knowledge in *NIK*. Therefore using m_1 itself means that knowledge from *NIK* is (indirectly) used (unless there is another analyzable term from which m_1 can be derived, but

[5]Unlike for from constraints, it does not always hold that every simple D-from constraint set C is satisfiable (although this is usually the case). However, if C is simple then $\lceil C \rceil$ is satisfiable. Hence, not excluding C from the search is not a problem of correctness but only of efficiency.

$$\frac{\textit{D-from}(m_1 \cup m_2 \cup M, IK, NIK) \cup C, \sigma}{\textit{D-from}(\langle m_1,m_2\rangle \cup M, IK, NIK) \cup C, \sigma}\ G^{LD}_{\text{pair}}, \quad \frac{\textit{D-from}(m_1 \cup m_2 \cup M, IK, NIK) \cup C, \sigma}{\textit{D-from}(\{|m_1|\}_{m_2} \cup M, IK, NIK) \cup C, \sigma}\ G^{LD}_{\text{scrypt}},$$

$$\frac{(\textit{D-from}(M, m_2 \cup IK, NIK) \cup C)\tau, \sigma\tau}{\textit{D-from}(m_1 \cup M, m_2 \cup IK, NIK) \cup C, \sigma}\ G^{LD}_{\text{unif1}}\ (\tau = mgu(m_1, m_2);\ m_1 \notin \mathcal{V}),$$

$$\frac{(\textit{from}(M, m_2 \cup IK \cup NIK) \cup C)\tau, \sigma\tau}{\textit{D-from}(m_1 \cup M, IK, m_2 \cup NIK) \cup C, \sigma}\ G^{LD}_{\text{unif2}}\ (\tau = mgu(m_1, m_2);\ m_1 \notin \mathcal{V}),$$

$$\frac{\textit{D-from}(M, m_1 \cup m_2 \cup \langle m_1,m_2\rangle \cup IK, NIK) \cup C, \sigma}{\textit{D-from}(M, \langle m_1,m_2\rangle \cup IK, NIK) \cup C, \sigma}\ A^{LD}_{\text{pair1}}\ (\{m_1, m_2\} \setminus IK \neq \emptyset),$$

$$\frac{\textit{D-from}(M, IK, \langle m_1,m_2\rangle \cup M' \cup NIK) \cup C, \sigma}{\textit{D-from}(M, IK, \langle m_1,m_2\rangle \cup NIK) \cup C, \sigma}\ A^{LD}_{\text{pair2}}\ (M' = \{m_1, m_2\} \setminus IK;\ \{m_1, m_2\} \setminus (IK \cup NIK) \neq \emptyset),$$

$$\frac{\textit{D-from}(M, m_1 \cup \{|m_1|\}_{m_2} \cup IK, NIK) \cup \textit{from}(m_2, IK) \cup C, \sigma}{\textit{D-from}(M, \{|m_1|\}_{m_2} \cup IK, NIK) \cup C, \sigma}\ A^{LD}_{\text{scrypt1}}\ (m_1 \notin IK),$$

$$\frac{\textit{D-from}(M, \{|m_1|\}_{m_2} \cup IK, m_1 \cup NIK) \cup \textit{D-from}(m_2, IK, NIK) \cup C, \sigma}{\textit{D-from}(M, \{|m_1|\}_{m_2} \cup IK, NIK) \cup C, \sigma}\ A^{LD}_{\text{scrypt2}}\ (m_1 \notin IK),$$

$$\frac{\textit{D-from}(M, IK, m_1 \cup \{|m_1|\}_{m_2} \cup NIK) \cup \textit{from}(m_2, IK \cup NIK) \cup C, \sigma}{\textit{D-from}(M, IK, \{|m_1|\}_{m_2} \cup NIK) \cup C, \sigma}\ A^{LD}_{\text{scrypt3}}\ (m_1 \notin (IK \cup NIK)).$$

Figure 5: The D-from rules for generation and analysis

this is, as in the previous case, only a matter of efficiency, not of correctness). Therefore we have the newly learned term m_1 in the *NIK*-part and the constraint that m_2 can be derived only when using *NIK*.

The A^{LD}_{scrypt3} rule describes, similarly, the case that the term to analyze is in the *NIK*-part. In this case, the result of the decryption is also added to the *NIK*-part, as the analysis is based on terms in *NIK* and could not be performed without *NIK*. However, to analyze the key m_2 we need not use information in *NIK*, so the generation of m_2 is expressed with a from constraint.

Definition 5. *Let $\vdash^D$ be the extension of the derivation relation $\vdash$ (cf. Definition 2) with the rules in Figure 5. The set of pairs of simple D-from constraint sets and substitutions that can be derived from (C, id) is D-Red$(C) = \{(C', \sigma) \mid ((C, id) \vdash^D (C', \sigma)) \wedge simple(C')\}$.*

3.2 Properties of *D-Red*

In this section, we show that the *D-Red* function has analogous properties to the *Red* function, namely correctness, completeness, and termination. Analogously to Theorem 1, we would like to have the property $[\![C]\!] = [\![\textit{D-Red}(C)]\!]$. However, this does not hold: there may be solutions of $(C', \sigma) \in \textit{D-Red}(C)$ that are not subsumed by C (but by $\lceil C \rceil$). An example is the D-from constraint $C = \textit{D-from}(m, k \cup \{|m|\}_k, m)$. Obviously, the intruder can derive the term m without the new knowledge, hence C is unsatisfiable. However, the rule G^{LD}_{unif2} is applicable, leading to a from constraint with an empty T-part: $C' = \textit{from}(\emptyset, k \cup \{|m|\}_k \cup m)$. As C' is simple, $[\![C']\!] = \{id\}$. Note that although this solution is not contained in $[\![C]\!]$, it is contained in $[\![\lceil C \rceil]\!]$, meaning that *D-Red* returns more solutions than desired, but still correct solutions — it has missed a possible reduction. This is a problem of efficiency, not of correctness and completeness. Hence we relax the above statement and show correctness only with respect to the original approach (the proofs can be found in [6]):

Theorem 2. *If C is a well-formed constraint set, then D-Red(C) is finite and $\vdash^D$ is well-founded, and $[\![\textit{D-Red}(C)]\!] \subseteq [\![\lceil C \rceil]\!]$.*

For the other direction, i.e. completeness, we have that all solutions of C are contained in *D-Red*:

Theorem 3. *$[\![C]\!] \subseteq [\![\textit{D-Red}(C)]\!]$ for a well-formed constraint set C.*

These theorems tell us that *D-Red* is correct, complete, and recursively computable (since $\vdash^D$ is finitely branching). This implies that we can apply *D-Red* to reduce the search space without excluding attacks or introducing new ones. To formalize this reduction of the search space, we integrate constraint differentiation into the search tree induced by the symbolic transition system.

3.3 Integrating Constraint Differentiation with Symbolic Transition Systems

Consider again the tree of Figure 4. This situation (which occurs whenever two independent actions are performed in two different orders, as explained above) characterizes when we apply constraint differentiation: we exploit the fact that the two symbolic states s_2 and s_4 of Figure 4 represent overlapping sets of ground states as shown by the shaded parts in s_2 and s_4.

Figure 6 merges parts of Figures 1 and 4 to illustrate how constraint differentiation works: we pick one, say s_4, of the overlapping states s_2 and s_4 in Figure 4 (where $t_2 = t_4$) and replace the *from* constraint that does not appear in the other constraint set with a *D-from* constraint; this yields the transformed state s'_4. That is, we use differentiation of constraints to restrict the extension of one of the two symbolic states to those ground states that are not covered by the other (as illustrated by the shaded part in the set of ground states). The following theorem shows that s_2 and s'_4 represent the same ground states as s_2 and s_4.

Theorem 4. *Consider two symbolic states $s_2 = (t_2, C_2)$ and $s_4 = (t_2, C_4)$ with constraint sets of the form $C_2 =$*

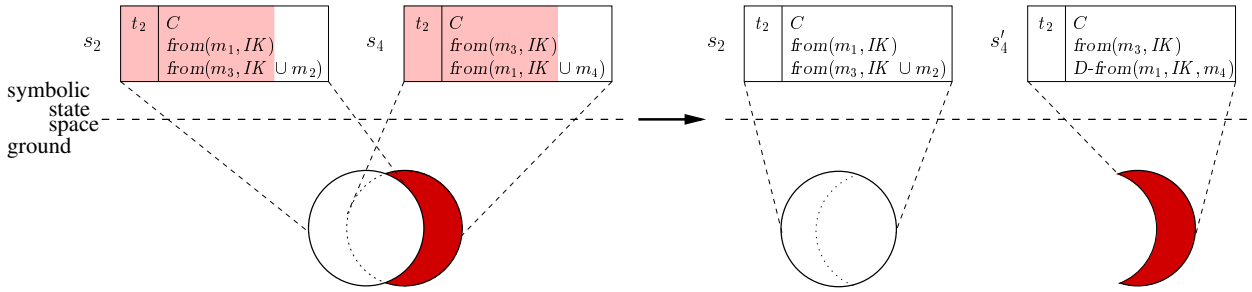

Figure 6: Constraint differentiation at work.

$C \cup from(m_1, IK) \cup from(m_3, IK \cup m_2)$ and $C_4 = C \cup from(m_3, IK) \cup from(m_1, IK \cup m_4)$ for some messages m_1, m_2, m_3, m_4 and some constraint set C. Then $[\![s_2]\!] \cup [\![s_4]\!] = [\![s_2]\!] \cup [\![s'_4]\!]$ for $C'_4 = C \cup from(m_3, IK) \cup D\text{-}from(m_1, IK, m_4)$ and $s'_4 = (t_2, C'_4)$.

PROOF. It suffices to show that $[\![C_2]\!] \cup [\![C_4]\!] = [\![C_2]\!] \cup [\![C'_4]\!]$. Since $\lceil C'_4 \rceil = C_4$ and $[\![\lceil C'_4 \rceil]\!] \supseteq [\![C'_4]\!]$, the direction $\supseteq$ is trivial. To show $[\![C_2]\!] \cup [\![C_4]\!] \subseteq [\![C_2]\!] \cup [\![C'_4]\!]$, we show that for every solution $\sigma \in [\![C_4]\!]$ with $\sigma \notin [\![C_2]\!]$ it holds that $\sigma \in [\![C'_4]\!]$. So assume $\sigma \notin [\![C'_4]\!]$. Then, since $\sigma \in [\![C_4]\!]$, it must hold that $\sigma \in [\![C_4]\!] \setminus [\![C'_4]\!]$, i.e. $\sigma \in [\![\lfloor C_4 \rfloor]\!]$. Hence, $\sigma \in [\![from(m_1, IK)]\!]$ and, since $\sigma \in [\![from(m_3, IK)]\!]$, we have $\sigma \in [\![C_2]\!]$, which contradicts the assumption. □

This theorem allows us to perform a transformation on the search tree that replaces from constraints with more restrictive D-from constraints without changing the set of represented ground states. If under the more restrictive constraint C'_4 the intruder could not use any new message from his knowledge, then even if C_4 is satisfiable, C'_4 is unsatisfiable (so that the shaded part of the set of ground states in Figure 6 is also empty), which we can check using *D-Red*. This is the maximal reduction that can be achieved by constraint differentiation: the node of the state s_4 and its subtree can be completely removed from the search tree as the intruder could not generate anything "interesting", i.e. nothing that he could not have generated before. Note that we can always consider the symmetric situation: if performing the restriction on s_2 rather than s_4 leads to an unsatisfiable constraint set, then we can remove the respective subtree.

When we apply *D-Red* to a state that results by replacing from constraints with D-from constraints, in the best case the constraint set turns out to be unsatisfiable, so the state (and the respective subtree) can be removed. However, it is also possible that after applying *D-Red* there still remain simple D-from constraints (i.e. with variables in the T-part). This means that it is not yet determined what the intruder will use here, so it is possible that it is some message from *NIK*. Such a D-from constraint is nonetheless useful for the reduction, as it constrains the child nodes by excluding certain solutions: the D-from constraint prevents all later instantiations of the variable in the T-part if these instantiations do not use some message of the *NIK*-part.

4. EXPERIMENTAL RESULTS

To test the practical relevance of constraint differentiation, we have integrated it in our on-the-fly protocol model-checker OFMC, which is based on the lazy infinite-state approach to protocol analysis presented in [4, 5]. In this approach, the search space is represented as a potentially infinite tree that is generated on demand.[6] We have straightforwardly implemented constraint differentiation as a filter that is applied to, and prunes, the search tree. We now describe experiments which show that constraint differentiation significantly improves both the performance and the scope of OFMC.

Even without constraint differentiation, OFMC is a state-of-the-art tool for finding protocol flaws, as shown in [5]: on a 2,4 GHz Pentium-4 PC (with 512 MB RAM, but OFMC is not memory intensive), OFMC takes 5.29 seconds of cumulative CPU time to detect flaws in all of the 35 protocols of the Clark-Jacob-library [9] that are known to be flawed [13]. OFMC also discovered a previously unknown flaw in the Yahalom protocol [5, 6]. Moreover, we have applied OFMC to large-scale protocols including IKE, SET, CHAP, and various other industrial protocols currently being standardized by the Internet Engineering Task Force (IETF).[7] In particular, in [5] we describe our analysis of the H.530 protocol [17], a protocol developed by Siemens and proposed as an Internet standard for multimedia communications. Siemens performed an analysis using the finite-state model-checker Casper/FDR [13, 18], which ran out of memory due to the complexity of the protocol. We have modeled H.530 in its full complexity and have detected a replay attack in 1.62 seconds. The weakness is serious enough that Siemens has changed the protocol.

The integration of constraint differentiation has allowed us to improve the performance and extend the scope of OFMC so that it not only scales well to the falsification of industrial-strength protocols, but it can also be applied as an effective verification tool, exhaustively searching the large space associated with bounded numbers of protocol sessions. We have carried out a large number of experiments to validate our approach.

As a first example, using constraint differentiation the total time for detecting the flaws in the Clark-Jacob library drops from $5.29s$ to $3.36s$, and the detection of the H.530 flaw drops from $1.62s$ to $0.80s$. The improvements achieved by constraint differentiation are more dramatic the more

[6] There is no relation between the lazy intruder and the lazy protocol analysis of [4] except that both are based on demand-driven evaluation.

[7] Several of these protocols employ the Diffie-Hellman key-exchange, which requires relaxing the free algebra assumption. Discussing how we realized this is beyond the scope of this paper.

complex the analyzed protocols and their flaws are, and when performing verification (for a bounded number of sessions). Verifying correct protocols is usually substantially more complex than falsifying flawed protocols, since flaw detection terminates as soon as an error is detected, while for verification the entire search space must be examined. Constraint differentiation has enabled us to perform verification for considerably larger numbers of parallel sessions than was possible with the previous version of the tool.

As a relevant example of an industrial-strength protocol we discuss here our analysis of the protocol-suite IKE [15]. We have used OFMC to analyze the full specification of each of the individual subprotocols of IKE and several combinations of them. By "full" we mean that we did not simplify the structure of the messages, which contain highly complex key-terms. OFMC detected a number of minor weaknesses of IKE, which have also been reported in [19, 25]. To discuss concrete performance results, we consider two subprotocols of IKE, the *Main Mode* and the *Aggressive Mode* of Phase 1 in the pre-shared key variants; the running times of the other subprotocols are similar.

Figure 7 compares the size of plies of the search tree for Aggressive Mode without and with constraint differentiation. We consider six scenarios varying in the number of parallel and consecutive sessions. We consider two ($[a,b],[a,i]$), three (also $[i,a]$) or four (also $[b,i]$) sessions, where $[a,b]$ means that agent a plays the protocol initiator role and agent b the responder role, and i is the intruder.[8] Moreover, sn denotes n consecutive sessions, where s1 (i.e. only one session) is the standard case and s2 means that an honest agent who has finished his part of the protocol session is prepared to engage in a further session with the same partners. Consecutive sessions are valuable for detecting replay attacks as they create a smaller search space than the same sessions in parallel. We display the number of nodes on each ply of the search tree; note that for the "smaller" scenarios (i.e. fewer parallel and consecutive sessions) the depth of the search tree is smaller, hence the empty cells. We also display the total number of nodes in the tree and the CPU time for searching the entire tree (TO denotes *time out* after one day, i.e. 1440 minutes of CPU time).

Figure 7 shows that constraint differentiation is most effective, as measured by the number of nodes that must be searched, when the original search space contains many interleavings of parallel sessions. The savings are most dramatic on the deeper plies of the search tree as the number of interleavings grows exponentially in the original model; since many interleavings are redundant and constraint differentiation can exploit this redundancy, the number of nodes does not necessarily grow exponentially with the depth of the tree. The difference between an exponential growth without constraint differentiation and an often sub-exponential growth with constraint differentiation leads to more dramatic savings the deeper the tree is searched. Note also that without constraint differentiation the number of nodes on each ply typically grows monotonically with the depth of the ply; with constraint differentiation the number grows on the first few plies and then starts to shrink again. This phenomenon is explained by the fact that with constraint differentiation the deeper we are in the tree, the more successor nodes are completely excluded. The intuition behind this is that many transitions possible in the original model do not permit the use of newly learned messages and are thus pruned from the tree by constraint differentiation.

In summary, by employing constraint differentiation, the OFMC tool scales significantly better with the size of the considered scenario. This extends the scope of OFMC and allows its use in verification for bounded scenarios. Figure 8 shows the broader picture, assessing the savings due to constraint differentiation for both main and aggressive mode for different parallel and consecutive sessions. Constraint differentiation reduces the search space significantly in all cases, and in some cases it even enables the analysis of problems that were out of the scope of OFMC and other tools.

[8]Such an explicit declaration of the intruder is only necessary when he shall participate under his real name, modeling a compromised or dishonest agent. He can always, without explicit declaration, send messages under an arbitrary identity.

5. CONCLUSIONS AND FUTURE WORK

Constraint differentiation effectively integrates the lazy intruder with ideas from partial-order reduction. We have proven that this integration does not change the set of represented ground states and hence is correct and complete. We have shown too, empirically, that it leads to dramatic reductions in the size of the state space searched, considerably improving the performance of the model-checker OFMC and making possible new kinds of analysis, i.e. both falsification of industrial-strength protocols and verification of bounded numbers of sessions of such protocols.

We see room for further improvements of the constraint differentiation technique. One promising idea is the use of the *D-from* constraints for heuristic search by focusing on "interesting" protocol interleavings, i.e. ones where at every step the intruder uses knowledge obtained during the previous step. We would also like to exploit results concerning bounds on the number of agents and parallel sessions that need to be considered for protocol verification, e.g. [24, 10].

6. REFERENCES

[1] R. Amadio and D. Lugiez. On the Reachability Problem in Cryptographic Protocols. In *Proc. Concur'00*, LNCS 1877, pp. 380–394. Springer, 2002.

[2] A. Armando and L. Compagna. Automatic SAT-Compilation of Protocol Insecurity Problems via Reduction to Planning. In *Proc. FORTE 2002*, LNCS 2529, pp. 210–225. Springer, 2002.

[3] F. Baader and T. Nipkow. *Term Rewriting and All That*. Cambridge University Press, 1998.

[4] D. Basin. Lazy Infinite-State Analysis of Security Protocols. In *Proc. CQRE'99*, LNCS 1740, pp. 30–42. Springer, 1999.

[5] D. Basin, S. Mödersheim, and L. Viganò. An On-The-Fly Model-Checker for Security Protocol Analysis. To appear in *Proc. Esorics 2003*, Springer.

[6] D. Basin, S. Mödersheim, and L. Viganò. Constraint Differentiation: A New Reduction Technique for Constraint-Based Analysis of Security Protocols (Extended Version). Technical Report 405, Dep. of Computer Science, ETH Zurich, 2003. `www.inf.ethz.ch/research/publications/`

[7] M. Boreale and M. G. Buscemi. A Framework for the Analysis of Security Protocols. In *Proc. CONCUR 2002*, LNCS 2421, pp. 483–498. Springer, 2002.

IKE Aggressive Mode Pre-Shared Key												
Mode:	without CD						with CD					
Scenario:	$[a,b],[a,i]$		$[a,b],[a,i],[i,a]$		$[a,b],[a,i],[i,a],[b,i]$		$[a,b],[a,i]$		$[a,b],[a,i],[i,a]$		$[a,b],[a,i],[i,a],[b,i]$	
Ply	s1	s2	s1	s2	s1	s2	s1	s2	s1	s2	s1	s2
1	3	3	4	4	5	5	3	3	4	4	5	5
2	7	7	14	14	23	23	5	5	10	10	16	16
3	13	14	43	45	97	100	7	8	19	21	40	43
4	17	27	112	139	368	420	6	12	30	44	86	111
5	15	53	238	422	1228	1727	5	17	35	81	150	261
6	15	101	393	1262	3501	6989	3	18	31	139	218	578
7		191	483	3699	8232	27835		20	22	215	241	1174
8		410	420	10637	15288	108927		23	8	319	203	2290
9		720		29783	21168	417862		22		436	136	4112
10		960		79939	18900	1565354		12		527	48	7025
11		990		201861		5695140		9		602		11062
12		990		467533		TO		5		576		16390
13				929500		TO				428		22544
14				1583582		TO				233		27443
15				2132130		TO				177		31024
16				1801800		TO				53		29595
17						TO						10531
18						TO						10531
19						TO						7857
20						TO						2371
Nodes	71	4467	1708	7242353	68811	TO	30	155	160	3866	1144	197426
Time	0.16s	13.66s	4.64s	40655.50s	3m41s	TO	0.08s	0.49s	0.49s	21.60s	4.17s	26m30s

Figure 7: Comparison of OFMC without and with constraint differentiation (CD) for IKE Aggressive Mode Pre-Shared Key: the nodes for each ply of the search tree and search time.

		IKE Main Mode Pre-Shared Key					IKE Aggressive Mode Pre-Shared Key				
Scenario	sn	Ply	without CD		with CD		Ply	without CD		with CD	
			Time	Nodes	Time	Nodes		Time	Nodes	Time	Nodes
$[a,b]$	1	7	0.11	36	0.08	24	4	0.02	7	0.01	6
	2	14	0.30	91	0.20	61	8	0.04	18	0.02	12
$[a,b],[a,i]$	1	11	45.41	8969	9.18	1902	6	0.15	71	0.08	30
	2	22	TO	TO	11m25	98271	12	13.82	4467	0.48	155
$[a,b],[i,b]$	1	10	10.99	2034	3.93	704	6	0.18	75	0.08	28
	2	20	185m52	1014127	3m45	21155	12	18.51	4897	0.57	159
$[a,b],[a,b]$	1	14	22m43	169531	2m38	19642	8	1.81	689	0.35	123
	2	28	TO	TO	429m09	2020491	16	68m56	906789	6.65	1508

Figure 8: Comparison of search time and tree size without and with constraint differentiation (CD) for IKE Main Mode and Aggressive Mode Pre-Shared Key.

[8] Y. Chevalier and L. Vigneron. Automated Unbounded Verification of Security Protocols. In *Proc. CAV'02*, LNCS 2404, pp. 324–337. Springer, 2002.

[9] J. Clark and J. Jacob. A Survey of Authentication Protocol Literature: Version 1.0, 17. Nov. 1997. URL: `www.cs.york.ac.uk/~jac/papers/drareview.ps.gz`.

[10] H. Comon-Lundh and V. Cortier. Security Properties: Two Agents are Sufficient. In *Proc. ESOP'2003*, LNCS 2618, pp. 99–113. Springer, 2003.

[11] R. Corin and S. Etalle. An Improved Constraint-Based System for the Verification of Security Protocols. In *Proc. SAS 2002*, LNCS 2477, pp. 326–341. Springer, 2002.

[12] D. Dolev and A. Yao. On the Security of Public-Key Protocols. *IEEE Trans. Information Theory*, 2, 1983.

[13] B. Donovan, P. Norris, and G. Lowe. Analyzing a Library of Security Protocols using Casper and FDR. In *Proc. FMSP'99*, 1999.

[14] M. Fiore and M. Abadi. Computing Symbolic Models for Verifying Cryptographic Protocols. In *Proc. CSFW'01*. IEEE Computer Society Press, 2001.

[15] D. Harkins and D. Carrel. RFC 2409: The Internet Key Exchange (IKE). 1998.

[16] A. Huima. Efficient Infinite-State Analysis of Security Protocols. In *Proc. FMSP'99*, 1999.

[17] ITU-T Recommendation H.530: Symmetric Security Procedures for H.510 (Mobility for H.323 Multimedia Systems and Services). 2002.

[18] G. Lowe. Casper: a Compiler for the Analysis of Security Protocols. *J. Computer Security*, 6, 1998.

[19] C. Meadows. Analysis of the Internet Key Exchange Protocol Using the NRL Protocol Analyzer. In *Proc. 1999 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, 1999.

[20] J. K. Millen and V. Shmatikov. Constraint Solving for Bounded-Process Cryptographic Protocol Analysis. In *Proc. CCS'01*, pp. 166–175, ACM Press, 2001.

[21] J. C. Mitchell, M. Mitchell, and U. Stern. Automated Analysis of Cryptographic Protocols Using Murphi. In *Proc. IEEE Symposium on Security and Privacy*, pages 141–153, 1997.

[22] D. Peled. Ten Years of Partial Order Reduction. In *Proc. CAV'98*, LNCS 1427, pp. 17–28. Springer, 1998.

[23] D. Song, S. Berezin, and A. Perrig. Athena: A Novel Approach to Efficient Automatic Security Protocol Analysis. *J. Computer Security*, 9, 2001.

[24] S. D. Stoller. A Bound on Attacks on Authentication Protocols. In *Proc. 2nd IFIP Int. Conf. on Theoretical Computer Science*, pp. 588–600. Kluwer, 2002.

[25] J. Zhou. Further Analysis of the Internet Key Exchange Protocol. *Computer Communications*, 23(17):1606–1612, 2000.