Sessions and Separability in Security Protocols^{*}

Marco Carbone¹ and Joshua D. Guttman²

¹ IT University of Copenhagen
carbonem@itu.dk

² Worcester Polytechnic Institute
guttman@wpi.edu

Abstract. Despite much work on sessions and session types in non-adversarial contexts, session-like behavior given an active adversary has not received an adequate definition and proof methods. We provide a syntactic property that guarantees that a protocol has session-respecting executions. Any uncompromised subset of the participants are still guaranteed that their interaction will respect sessions. A protocol transformation turns any protocol into a session-respecting protocol.

We do this via a general theory of separability. Our main theorem applies to different separability requirements, and characterizes when we can separate protocol executions sufficiently to meet a particular requirement. This theorem also gives direct proofs of some old and new protocol composition results. Thus, our theory of separability appears to cover protocol composition and session-like behavior within a uniform framework, and gives a general pattern for reasoning about independence.

Keywords: Sessions, Security Protocols, Strand Spaces.

1 Introduction

A transaction or protocol respects sessions if the local runs of the individual participants always match up globally in a compatible way. When one participant receives any message in a session σ , it should have been sent by another participant acting *within the same session* σ . Session-respecting behavior is often studied using *session types* [21,22]. However, most work in this tradition studies sessions within a benign execution environment.

We adapt those ideas to environments containing active adversaries, who may control the medium of communication [28,15,4]. We define session-respecting behavior in an adversarial environment, offering syntactic conditions that ensure a protocol's behavior respects sessions. We exhibit a transformation that, given any protocol, yields one with session-respecting behavior.

Our central idea is *separability*. In an execution, an adversary may receive a message from one session and deliver it, or its fragments, into another session. In this case, we would like to *separate* the sessions by removing the connection that

---

^{*} Carbone thanks the Danish Agency for Science, Technology and Innovation, and Guttman thanks the US National Science Foundation under grant CNS-0952287. Extended version at http://www.cs.wpi.edu/~guttman/pubs/CG13_long.pdf.

the adversary has created. Separability means we can do this, possibly applying a renaming to one of them so that they involve different parameters. Although the adversary can create connections between different sessions, these connections are inessential. They can be removed modulo renaming. Hence, anything the adversary can achieve in a session, he can also do without relying on any other session. No successful attack requires unwitting support from participants engaged in a different session.

Session separability clarifies the real world effects of a protocol. Suppose a protocol allows a customer to buy merchandise through a broker, who receives a commission from a manufacturer. Can the broker manipulate the protocol so one interaction with the customer allows two interactions with the manufacturer? Can the broker receive his commission for the same transaction twice?

Suppose a compliant customer and manufacturer interact with a dishonest broker in a separable protocol. If messages from the single customer run reach local runs M_1, M_2 for the manufacturer, they will belong to the same session. Alternatively one run M belongs to no session, i.e. it occurs with no involvement of a customer. These conditions are easy for a protocol designer to analyze. To protect against the first, the manufacturer should contribute a fresh random value (“nonce”) to help define the session. Then distinct manufacturer runs always belong to different sessions. To protect against the second, some authentication is needed between manufacturer and customer, a familiar and well-understood problem. Thus, separability reduces the no-double-commission property to simple characteristics of the protocol.

Strand spaces offer a partially ordered model; protocol executions (“bundles”) are annotated directed acyclic graphs [30,20,18]. The edges represent causal relations. We interpret separation properties in terms of the absence of causal paths in these graphs, or the ability to find a related graph without them.

Contributions. Our main result, the *Separability Theorem* (Thm. 16), tells how to take an execution of a protocol and modify it into another execution that satisfies separability. It applies to a range of different *separability specifications*. Each separability specification is a partial order that says which kinds of events are allowed to causally affect which others.

Our result about sessions, Thm. 17, says that the syntactic conditions in Def. 8—mainly concerning “session nonces” that serve to define sessions—entail the premise of Thm. 16. Thus, executions can be made to satisfy a separability specification defined in terms of these nonces. We also provide a transformation that strengthens any protocol to one that satisfies these conditions (Thm. 10).

Thm. 16 also applies to other separability specifications. We apply it to protocols with “disjoint encryption” in three slightly different senses (Thms. 19–21), thereby yielding variants, sometimes sharpenings, of a number of results on preserving security goals under protocol composition [19,1,9,8,11]. Thus, the Separability Theorem formalizes a pattern of reasoning with wide applicability in protocol design and analysis. It unites session-oriented reasoning with protocol composition. All proofs are relegated to the extended version http://www.cs.wpi.edu/~guttman/pubs/CG13_long.pdf.

Related Work. Various flavors of sessions and separability have already played roles within protocol analysis and design. Among approaches based on computational methods, a session notion is often used to define the local runs that authentication should connect, as with Bellare and Rogaway [4]; in some models the sessions are defined by a bitstring that may be chosen by the adversary or built out of random contributions from the participants (e.g. Canetti-Krawczyk [7]). The Universal Composability model also assumes a random value that contributes to each cryptographically prepared unit and acts as a session identifier [6]; for a recent and more flexible alternative, see Küsters and Tuengerthal [24].

If different sessions of a protocol can never affect one another, then this simplifies analysis. The designer can explore the outcomes possible with a single instance of each role in the protocol. Indeed, Lowe’s original proof that his changes to Needham-Schroeder were correct used a separability argument. He proved that any run could involve at most two non-separable instances of either role, and then model-checked the possible two-instance runs [25]. Lowe and Al-laa Kamil [23] use separability to establish properties of TLS, such as that the adversary cannot divert application data from one TLS session to another. Their path-based methods within the strand space framework motivated some of the techniques used below in §4.

Cortier et al. [10] propose a protocol transformation, which they prove correct using session separability. Given a protocol satisfying any security property in an environment with a passive adversary, their transformation returns a new protocol that satisfies that property despite an active adversary. Their transformation adds freshly-generated nonces to the original protocol; this suggested our treatment of nonces in Thm. 10. Their transformation then inserts all of these nonces in with each message of the original protocol, which is signed and then encrypted. Our transformation does not add any additional cryptographic operations, but simply inserts the nonces into any pre-existing cryptographic units. This simpler treatment suffices because we are here exclusively concerned with separability, rather than any particular security goals. Another contrast concerns the adversary model. Their result concerns only sessions in which every participant is compliant, whereas our separability holds for the compliant participants, even in sessions with non-compliant participants.

Arapinis, Delaune, and Kremer [2] also offer a separability argument, leading to a protocol transformation which guarantees secrecy. The transformation adds nonces to each encrypted or signed term, although, together with nonces, it also adds principal names. Only the former is needed for separability; the latter helps with secrecy. Their transformation appears to generalize Lowe’s fix to Needham-Schroeder. A subsequent paper with Ryan [12] investigates tagged password-based protocols, where the “tags” are tuples of session parameters hashed in with the key. They show that their composition is resistant to guessing attacks. Their proofs appear to establish particular instances of our Separability Theorem. We conjecture that our methods reconstruct their results, although proving this would require reformulating behavioral equivalences (in addition to trace properties) within our framework.

Deniérou et al. in [5] provide a compiler for generating ML code from multiparty protocol specifications. Their main result (Session Integrity Theorem) shows that there is no interference between multiple instances of the same protocol. Such a result could also be modeled in our framework as a variant of protocol independence. As we have mentioned, our approach also seems to capture the essence of several protocol composition results [19,1,9,8,11].

Separability also allows full verification within the bounded-session model of protocol analysis [29,26,3], rendering many classes of problems decidable.

In [17] we offer a logical language that can formalize the security goals that are preserved when omitting separable local runs.

2 Strand Spaces, with a Session Protocol

We first summarize the strand space terminology we will use in this paper. See [30,18] for more detail on strand spaces and our terminology.

We also introduce an example to illustrate separability, the *Trusted Broker Service*, in which a server S acts as a broker to match clients C_1 and C_2 , who are executing different roles. S provides them with a key K to use to initiate an exchange. Clients trust the broker to generate a fresh key; to distribute it only to compliant principals; and to choose an appropriate pairing of clients.

Messages. Let Alg_0 be an algebra equipped with some operators and a set of homomorphisms $\eta: \text{Alg}_0 \rightarrow \text{Alg}_0$. We call the members of Alg_0 **basic values**.

Alg_0 is the disjoint union of infinite sets of *nonces*, *basic keys*, *names*, and *texts*. The operators $\text{sk}(a)$ and $\text{pk}(a)$ map names to signature keys and public encryption keys. K^{-1} maps an asymmetric basic key to its inverse, and a symmetric basic key to itself. Homomorphisms η are maps that respect sorts and the operators $\text{sk}(a)$, $\text{pk}(a)$, and K^{-1} . An infinite set X disjoint from Alg_0 —the **indeterminates**—act like unsorted variables.

The algebra Alg of **messages** is freely generated from $\text{Alg}_0 \cup X$ by two operations: encryption $\{t_0\}_{t_1}$ and tagged concatenation $\text{tag } t_0, t_1$. The tags tag are drawn from some set TAG . For a distinguished tag nil , we write $\text{nil } t_0, t_1$ as t_0, t_1 . In $\{t_0\}_{t_1}$, a non-basic key t_1 is a symmetric key. To reduce cases in proofs, we do not introduce digital signature and hashing as separate operations. We can encode hashes $\text{hash}(t)$ as encrypting t with a public key K_h , where no principal holds the inverse decryption key K_h^{-1} . A digital signature $\llbracket t \rrbracket_K$ is encoded as the concatenation $t, \{\text{hash}(t)\}_K$.

A homomorphism $\alpha = (\eta, \chi): \text{Alg} \rightarrow \text{Alg}$ pairs a homomorphism η on basic values and a function $\chi: X \rightarrow \text{Alg}$; $\alpha(t)$ is defined by the conditions:

$$\begin{aligned} \alpha(a) &= \eta(a), & \text{if } a \in \text{Alg}_0 & & \alpha(\{t_0\}_{t_1}) &= \{\alpha(t_0)\}_{\alpha(t_1)} \\ \alpha(x) &= \chi(x), & \text{if } x \in X & & \alpha(\text{tag } t_0, t_1) &= \text{tag } \alpha(t_0), \alpha(t_1) \end{aligned}$$

We call these homomorphisms **substitutions**, and use them to plug in values for parameters. Indeterminates x are blank slots, to be filled by any $\chi(x) \in \text{Alg}$.

Messages t_1, t_2 have a *common instance* when there exist substitutions α, β that identify them: $\alpha(t_1) = \beta(t_2)$. Alg has the most general unifier property. That is, suppose that for $v, w \in \text{Alg}$, there exist any α, β such that $\alpha(v) = \beta(w)$. Then

there are α_0, β_0 , such that $\alpha_0(v) = \beta_0(w)$, and for all α_1, β_1 , if $\alpha_1(v) = \beta_1(w)$, then α_1 and β_1 are of the forms $\gamma \circ \alpha_0$ and $\gamma \circ \beta_0$.

Strands, Ingredients, and Origination. A **strand** is a sequence of local actions called **nodes**, each of which is either a message *transmission*, written $\bullet \rightarrow$, or else a message *reception*, written $\bullet \leftarrow$. Strands may be written vertically, or horizontally as in Fig. 1. This figure shows the behaviors of an initiating client C_1 and a responding client C_2 with a broker or server S . The protocol, which we call TBS, allows the broker to pair requests from suitable pairs of clients, and distribute a session key to them.

If n is a node, and the message t is transmitted or received, then we write $t = \text{msg}(n)$. Double arrows indicate successive events on the same strand, e.g. $\bullet \Rightarrow \bullet \Rightarrow \bullet$. Each role in Fig. 1, and each column in Figs. 3,5 is a strand.

We write $s \downarrow i$ to mean the i^{th} node along s , starting at $s \downarrow 1$. The **parameters** of s are the basic values and indeterminates in any $\text{msg}(s \downarrow i)$.

The **ingredients** of a message are those subterms that may be reached by descending through concatenations, and through the plaintext but not the encryption keys. The values that **occur** in it descend also through encryptions. We write $\sqsubseteq$ (“is an ingredient of”) and $\ll$ (“occurs in”), resp., for the smallest reflexive, transitive relation such that

$$\begin{aligned} t_1 \sqsubseteq (t_1, t_2) \quad t_2 \sqsubseteq (t_1, t_2) \quad t_1 \sqsubseteq \{t_1\}_{t_2} \\ t_1 \ll (t_1, t_2) \quad t_2 \ll (t_1, t_2) \quad t_1 \ll \{t_1\}_{t_2} \quad t_2 \ll \{t_1\}_{t_2}. \end{aligned}$$

We say that t **originates** on a node n if n is a transmission node, and $t \sqsubseteq \text{msg}(n)$, and for all n_0 , if $n_0 \Rightarrow^+ n$, then $t \not\sqsubseteq \text{msg}(n_0)$. A basic value a is *freshly chosen* if it originates just once. We call it **uniquely originating**. In this case, a was chosen by a participant, without the bad luck of any other principal selecting the same value independently. A key is regarded as uncompromised if it originates *nowhere*. We call a basic value a **non-originating** in a set of nodes B if there is no $n \in B$ such that a originates at n . It may still be used in B even if it does not originate anywhere, since the regular strands may receive and send messages encrypted by K or K^{-1} , thus using K for encryption and decryption, resp.

A message t_0 **lies only inside encryptions in t with keys S** iff, in t 's abstract syntax tree, every path from the root to an occurrence of t_0 traverses an encryption, and if that occurrence is in the plaintext, then the key is in S .

Protocols. A **protocol** Π is a finite set of strands, called the **roles** of the protocol. A **regular** strand for Π is any instance of one of the roles of Π , i.e. the result $\alpha(\rho)$ of some substitution α on the parameters of a role $\rho \in \Pi$. Fig. 1 is an example protocol. A protocol may also specify some parameters of a role that are always non-originating or uniquely originating. We will also formalize *adversary* behavior by strands (which use inverse). We stipulate a syntactic constraint:

Assumption 1. *If $\rho \in \Pi$, then the key inverse symbol does not appear in any message $\text{msg}(\rho \downarrow i)$. Moreover, $\text{sk}(A) \not\sqsubseteq \text{msg}(\rho \downarrow i)$. If $\{t\}_K \ll \text{msg}(\rho \downarrow i)$ for $\rho \in \Pi$, then K is either a basic value or an encryption (not a concatenation).*

Fig. 1. Trusted Broker Service Protocol, TBS

Fig. 2. Part I: Adversary roles to generate basic value a ; encrypt and decrypt; concatenate and separate. Part II: A compound adversary activity

The Adversary. **Adversary strands** consist of zero or more reception nodes followed by one transmission node. The adversary obtains the transmitted value as a function of the values received; or creates it, if there are no reception nodes. The adversary can choose basic values, and operate on complex values using the strands shown above in Fig. 2. These are often used in patterns, e.g. as in Fig. 2 Part II, which transport information along paths. Six strands are shown. Two are of length 1, in which the adversary transmits keys, namely his own private decryption key $\text{pk}(C)^{-1}$ and B 's public encryption key $\text{pk}(B)$. Two are a (leftmost) decryption strand and a (rightmost) encryption strand. The second node on a decryption or encryption strand is called the **key node**, since it receives the key used to perform the cryptographic operation.

In the middle are a separation strand that breaks A, N_a into its two parts, followed by a concatenation strand that puts them back together. These strands are unnecessary here. We include them here to illustrate that the adversary can always break a concatenation down to non-concatenated parts, i.e., either basic values or encryptions (see Assumption 3).

Fig. 4. Open bundle separating Fig. 3

allowing some (untrusted) auxiliary process to propose a matching. In Fig. 3, the adversary reuses the nonces N_1, N_2 to start a second server strand. However, we can fix this, separating the second server strand, just by renaming these nonces to new values N'_1, N'_2 . This yields the new bundle in Fig. 4, in which the adversary can supply the message coming from the upper right. Fig. 4 is an open bundle, as shown without adversary activity.

These figures are annotated (possibly open) bundles with various choices of **unique**, **non**. An interesting choice would be $\text{unique} = \{N_1, N_2, K, N'_1, N'_2\}$ for Fig. 4 and $\text{unique} = \{N_1, N_2, K\}$ for Fig. 3. A relevant choice for both $\text{non} = \{\text{sk}(S), \text{pk}(C_1)^{-1}, \text{pk}(C_2)^{-1}, \text{pk}(S)^{-1}\}$.

In studying separability we are interested in bundles equipped with a choice of fresh and uncompromised values. Hence, we will assume that all bundles are annotated with sets of uniquely originating and non-originating values **unique**, **non**. When using “bundle” and $\mathcal{B}$, we will mean “annotated bundle” as defined above.

The core pattern for separating a session is:

- removing dependence on an existing session;
- renaming some freshly chosen items in one or more local runs;
- allowing the adversary to supply incoming messages in these runs.

When a protocol ensures that this pattern will succeed in separating behaviors, it has session behavior.

However, this is not always possible. As an example, consider the protocol TBSMINUS, which is just like TBS, except that the session nonces N_1, N_2 are omitted in all the messages. Here we can have the essentially inseparable bundle Fig. 5. No amount of renaming and pruning edges will produce a bundle in which C_2 and C'_2 do not both depend on the same strand C_1 .

We assume (i) public encryption keys may be freely sent or used by anyone, including the adversary; and (ii) when a value a originates uniquely, and is used on a different regular strand as part of a key for encryption, then it has been received as an ingredient on that strand. When $a \sqsubseteq \text{msg}(m_1)$, this conclusion follows from the definition of unique origination. We also assume (iii) that a basic value is not received from a later transmission when it could be received from an earlier one. If a bundle violates this property, we can fix it by rerouting arrows to start from earlier nodes.

and $N'_2 \mapsto N_2$, and is elsewhere the identity, is a local renaming, which we will write $[N'_1 \mapsto N_1, N'_2 \mapsto N_2]$. It is a renaming (the identity) when restricted to the parameters that appear in S_1 , the left half, since N'_1, N'_2 do not appear on the left. Moreover, it is a renaming when restricted to S_2 , the right half, too, since N_1, N_2 do not appear as parameters on the right. Thus, it is injective on the parameters appearing in S_2 .

Every renaming is a local renaming, but a local renaming α is not a true (“global”) renaming when $\alpha(x) = \alpha(y)$ holds for parameters x, y to nodes in different partition classes S_j, S_k . We often think of the action of a local renaming backward, viewing its source as the result of “renaming apart” values that are equated in its target. If we view $[N'_1 \mapsto N_1, N'_2 \mapsto N_2]$ as if it were acting on Fig. 3 to yield Fig. 4, then it is “renaming apart” different occurrences of N_1, N_2 .

One open bundle lies below another if, after applying a local renaming forward, their regular nodes are the same, as are their uniquely originating and non-originating values, but one precedence order is a suborder of the other:

Definition 5. 1. $\mathcal{C}$ lies below $\mathcal{B}$ via α iff, for some $S_1, \dots, S_i$, α is a local renaming for $\mathcal{C}$ with respect to $S_1, \dots, S_i$, and:

- (a) $\alpha(\text{regnodes}(\mathcal{C})) = \text{regnodes}(\mathcal{B})$;
- (b) For all $n_0, n_1 \in \text{regnodes}(\mathcal{C})$, $n_0 \preceq_{\mathcal{C}} n_1$ implies $\alpha(n_0) \preceq_{\mathcal{B}} \alpha(n_1)$;
- (c) $\alpha^{-1}(\text{unique}(\mathcal{B})) = \text{unique}(\mathcal{C})$; and
- (d) $\alpha^{-1}(\text{non}(\mathcal{B})) = \text{non}(\mathcal{C})$

2. $\mathcal{C}$ lies below $\mathcal{B}$ if it does so via some α .

3. $\mathcal{B}$ and $\mathcal{C}$ are **equivalent** iff each lies below the other via renamings α, β , and $\alpha \circ \beta$ is the identity on the parameters involved.

For instance, Fig. 4 lies below Fig. 3 via $[N'_1 \mapsto N_1, N'_2 \mapsto N_2]$, given the choices of **unique**, **non** mentioned after Defn. 2.

If $\mathcal{C}$ lies below $\mathcal{B}$, then $\mathcal{C}$ differs from $\mathcal{B}$ only in having a sparser ordering, and in not yet having equated some parameters that have been equated in $\mathcal{B}$. We can think of $\mathcal{C}$ as a simplified version of $\mathcal{B}$. It is less informative in that the information that these parameters are equal has not yet been added.

Lemma 6. “Lies below” is a well-founded partial order to within isomorphism:

- 1. “Lies below” is reflexive and transitive.
- 2. $\mathcal{C}$ and $\mathcal{B}$ each lie below the other iff their regular parts are isomorphic.
- 3. If $\langle \mathcal{B}_i \rangle_i$ is an infinite sequence of bundles such that $i < j$ implies $\mathcal{B}_j$ lies below $\mathcal{B}_i$, then for some i, k , $i < k$ and $\mathcal{B}_i$ lies below $\mathcal{B}_k$.

3 Formalizing Sessions

We now turn to defining when TBS and similar session-oriented protocols are separable. Suppose that Π is a protocol, and $P: \Pi \rightarrow \text{Nonce} \cup \text{Key}$ is a function that chooses a parameter for each role. As an example, if Π is TBS, we would be interested in the function P that assigns N_1 to the first client role; N_2 to the second client role; and K to the server (broker) role.

We say that x is a **session parameter** if $x \in \text{range}(P)$. P associates each role to the session parameter that it chooses. We call $P(\rho)$ ρ 's **proper session parameter**, and we require that $P(\rho)$ originates on ρ .

If x is a session parameter, x is **acquired at step i** if $x \ll \rho \downarrow i$ but $x \not\ll \rho \downarrow j$ for $j < i$. It is **acquired by step k** if it is acquired at step i for some $i \leq k$. A parameter x is **key material at step i** if $x \ll K$ and $\{t\}_K \ll \text{msg}(\rho \downarrow i)$.

As a convention, we will assume that the parameters of each role have been chosen (by a renaming if necessary) so that corresponding session parameters on different roles have the same name. We could of course avoid this convention at the cost of added notation, in the form of a function which would supply the necessary correlations.

No Ambiguity. TBS uses the session parameters unambiguously in each encryption. No encryption in the protocol could be misinterpreted by a receiver so as to interchange the session parameters. For instance, N_1 and N_2 always appear in the same order, and K always appears after them or in key position.

No Ambiguity: If encryptions $\{t\}_K \ll \text{msg}(\rho \downarrow i)$ and $\{t'\}_{K'} \ll \text{msg}(\sigma \downarrow j)$ have a common instance $\alpha(\{t\}_K) = \beta(\{t'\}_{K'})$, then ρ and σ have acquired the same session parameters by steps i and j resp., and $\alpha(x) = \beta(x)$ for each of those session parameters.

We here follow our convention that corresponding session parameters on different roles have been given the same parameters names.

Contribution. Every encrypted unit involves the session parameters. This is akin to the tagging property [12], except that the session parameters do not have to contribute to the key. The last message of TBS is $\{\text{Hello}, C_1, N_1, N_2\}_K$. Two session parameters are in the plaintext, while K is the encryption key. All the session parameters could be concentrated in the key; $\{\text{Hello}, C_1\}_{\text{hash}(N_1, N_2, K)}$ would also work. Alternatively, they could all be concentrated in the plaintext, with some public key used for encryption.

In this protocol, the participants agree on all of the session parameters at the start. They then use them throughout the remainder of the protocol. A protocol can also have some participants agree on their session parameters, while other participants join the session later. These “late arrivals” allow for an attractive flexibility in the session-type literature [13]. Of course, the encrypted units *before* the late arrivals are expected to contain only the session parameters that have already been seen at that point.

Contribution to Encryptions: If $\{t\}_K \ll \text{msg}(\rho \downarrow i)$ and session parameter x is acquired by step i , then $x \ll \{t\}_K$.

The No-Vs Property. The observation that session parameters may be acquired piecemeal is an important insight. It implies that “same session,” which sounds like an equivalence relation, is in fact misleading. A partially defined session with session parameters $x_1, \dots, x_i$ may affect any of its possible extensions with an additional session parameter $x_1, \dots, x_i, x_{i+1}$. However, any one of those

extensions is incompatible with those having a different value x'_{i+1} . Indeed, messages from a step with extended session parameters $x_1, \dots, x_i, x_{i+1}$ should not affect an event with session parameters $x_1, \dots, x_i$. If they did, the latter could also affect a distinct extension $x_1, \dots, x_i, x'_{i+1}$. Thus, transitively, there could be effects from an event with parameters $x_1, \dots, x_i, x_{i+1}$ to one with parameters $x_1, \dots, x_i, x'_{i+1}$. That would be contrary to the session discipline.

For this reason, we regard the “may influence” relation on partially defined sessions as a partial order (on the sessions) or as a pre-order (on the transmission and reception events within the sessions). We will write $n_1 \rightsquigarrow n_2$ when an event n_1 may influence an event n_2 .

We require non-influence to persist, specifically when n_1 selects a fresh value that is a parameter to n_2 . We formulate this as a “no Vs” condition. Whenever we have a V in the may-influence relation, this is not an open V , but a closed triangle-like configuration, for any $n_3 \succeq n_2$:

$$\begin{array}{ccc} n_1 & \rightsquigarrow & n_3 \\ n_2 & \rightsquigarrow & \end{array} \quad \text{implies} \quad \begin{array}{ccc} n_1 & \rightsquigarrow & \\ \vee & & \\ n_2 & \rightsquigarrow & n_3. \end{array} \quad (1)$$

A node n_2 that I cannot influence cannot influence a later node that I can influence, at least when I have uniquely originated a value found in that node. This no Vs property turns out to be crucial to proving the Separability Theorem, whose proof tries to create new bundles by local renamings.

To see what could go wrong, suppose the TBS server received the two parts of its first incoming message on separate nodes: $(c11, C_1, N_1) \Rightarrow (c21, C_2, N_2) \Rightarrow \dots$. Then an adversary could deliver C_2 ’s nonce N_2 as if it were from C_1 , to the first server node n_2 . C_2 ’s first node n_1 should not influence n_2 , since n_1 has the C_2 nonce defined, whereas n_2 does not; n_2 has only the C_1 nonce parameter defined. However, if the adversary re-delivers the same nonce on the server’s second node n_3 , then C_2 ’s first node n_1 can influence this second server node n_3 . Node n_3 has the same value for the only session parameter defined on n_1 . This is precisely the open V situation, where $n_1 \not\rightsquigarrow n_2 \rightsquigarrow n_3$, and $n_1 \rightsquigarrow n_3$.

Acquisition. In order to ensure the no-Vs property syntactically, some properties are needed, constraining how session parameters are acquired. First, some session parameters $\bar{x}$ are received in a principal’s first reception. These may be transported without encryption, such as N_1 and N_2 in TBS. This is why S receives both N_1 and N_2 in a single message in its first node. Second, there are no transmissions after a reception and before transmitting a strand’s proper session nonce. Third, when a session includes late-arriving participants, values freshly chosen after a late arrival in the session will be transmitted under encryptions that cannot be compromised. Various techniques are available for proving this [20,18], but here we will just use a simple sufficient condition, namely that the decryption key is non-originating. These messages will be received by participants that have already joined the session; i.e. their proper nonces have already been chosen, and must also appear in this encryption by the *Contribution* requirement. This is a per-bundle requirement, for a bundle $\mathcal{B}$.

Parameter Acquisition: Session parameters divide into two groups, $\bar{x}$ and $\bar{y}$.

1. If x in $\bar{x}$ is acquired on reception node $\rho \downarrow i$, then i is the earliest reception node on ρ .
2. If x in $\bar{x}$ is acquired on transmission node $\rho \downarrow i$, and $\rho \downarrow k$ is any reception node with $k < i$, then there is no transmission node between them.
3. Let y in $\bar{y}$ be acquired (by reception or transmission) on $\rho \downarrow i$, and let $k \geq i$. There is a set $\text{LAK}(\mathcal{B})$ of *late-arrival protection keys of $\mathcal{B}$* such that:
  - (a) If $\alpha(\rho \downarrow k) \in \text{nodes}(\mathcal{B})$, then $\alpha(y)$ lies only inside encryptions in $\text{msg}(\alpha(\rho \downarrow k))$ with keys K where $\alpha(K^{-1}) \in \text{LAK}(\mathcal{B})$.
  - (b) If $a \in \text{unique}_{\mathcal{B}}$ is any value acquired on $\alpha(\rho \downarrow k)$, $\alpha(y)$ lies only inside encryptions in $\text{msg}(\alpha(\rho \downarrow k))$ with keys K where $\alpha(K^{-1}) \in \text{LAK}(\mathcal{B})$.

Condition 3 ensures that y always appears together with all previously defined session parameters. We focus on bundles in which, for any late arrivals to the session in a bundle $\mathcal{B}$, the strands still active then are all uncompromised, i.e. $\text{LAK}(\mathcal{B}) \subseteq \text{non}_{\mathcal{B}}$. In TBS, all session parameters belong to the first group $\bar{x}$, as all of the roles acquire them from their peers on their first reception. For protocol design, it is desirable that the session key can double as S 's session parameter, traveling in the encrypted messages from the server.

May-Influence Relations. Curiously, the Separability Theorem, Thm. 16 depends only on two properties of a reflexive, transitive *may-influence* relation, namely, the no V s property, and the fact that forward influence on a strand is always permitted. Because of this generality, we sought to specify various degrees of separability, i.e. to specify how sparse a bundle we would like to obtain in the “lies below” ordering. To parametrize our reasoning, we define a *may-influence* relation to be a pre-ordering $n_1 \rightsquigarrow n_2$ on regular nodes with these two properties. It specifies the upper bound on which nodes of Π may influence each other.

Definition 7. Let $\mathcal{B}$ be an (annotated) bundle for protocol Π . Then a preorder $\rightsquigarrow$ is a **may-influence relation** for $\mathcal{B}$ iff for all $n_1, n_2, n_3 \in \text{regnodes}(\mathcal{B})$,

1. if $n_1 \Rightarrow n_2$ then $n_1 \rightsquigarrow n_2$; and
2. “No V s,” Eqn. 1: Suppose (i) $a \in \text{unique}_{\mathcal{B}}$ originates at n_1 and $a \ll \text{msg}(n_2)$ and (ii) $n_2 \rightsquigarrow n_3$ and $n_2 \preceq_{\mathcal{B}} n_3$. If $n_1 \rightsquigarrow n_3$, then $n_1 \rightsquigarrow n_2$.

$\mathcal{B}$ **obeys** $\rightsquigarrow$ iff, for all $m, n \in \text{regnodes}(\mathcal{B})$, $m \preceq_{\mathcal{B}} n$ implies $m \rightsquigarrow n$.

Π **obeys** $\rightsquigarrow$ subject to Φ if, for every Π -bundle $\mathcal{B}$ satisfying Φ , there is a Π -bundle $\mathcal{C}$ satisfying Φ such that $\mathcal{C}$ lies below $\mathcal{B}$ and $\mathcal{C}$ obeys $\rightsquigarrow$.

When $m \rightsquigarrow n$, we say that m is *permitted to influence* n .

When $m \Rightarrow n$, m must be allowed to influence n , since it is impossible to prevent the influence; hence condition 1 on influence functions. Condition 2 prohibits open, V-shaped configurations. One leg of the V starts at a 's origin n_1 , and the other at n_2 , and the legs join at a jointly influenced n_3 . When $a \ll \text{msg}(n_2)$, then n_1 must be permitted to influence n_2 . If a 's origin n_1 cannot influence n_2 , then their causal consequences must remain separated thereafter.

Π obeys $\rightsquigarrow$ if Π -bundles either already obey the ordering constraint, or some bundle lying below is sparse enough to obey it. In weakening the order $\preceq$, we

are allowed to select preimages under local renamings. We use the constraints Φ to record assumptions about freshly chosen nonces and uncompromised keys.

Protocols with Session Parameters. We can now define:

Definition 8. A bundle $\mathcal{B}$ satisfies Φ_s , the **session constraint**, if the late arrival protection keys $\text{LAK}(\mathcal{B}) \subseteq \text{non}_{\mathcal{B}}$ and, for every node $\alpha(\rho \downarrow i) \in \text{nodes}(\mathcal{B})$, where ρ acquires its proper session nonce at step i , $\alpha(P(\rho)) \in \text{unique}_{\mathcal{B}}$.

Π has **session parameters** P for $\mathcal{B}$ if No ambiguity, Contribution to encryptions, and Parameter acquisition hold for Π , P , and $\mathcal{B}$.

The **session may-influence** relation $\sim_s$ holds between Π -nodes n_1 and n_2 , written $n_1 \sim_s n_2$, iff (i) $n_1 = \alpha(\rho \downarrow i)$ and $n_2 = \beta(\sigma \downarrow j)$ where $\rho, \sigma \in \Pi$; (ii) every session parameter x that has been acquired by step i on ρ has been acquired by step j on σ ; and (iii) $\alpha(x) = \beta(x)$ for each session parameter x acquired by step i on ρ .

Essentially, $n_1 \sim_s n_2$ means that the partial function assigning session parameters to values in node n_1 is a subfunction of the partial function assigning session parameters to values in node n_2 . The may-influence relation is fixed by the ordering of definedness on these partial functions.

Lemma 9. If $\mathcal{B}$ is a Π -bundle satisfying Φ_s , and Π has session parameters P in $\mathcal{B}$, then $\sim_s$ is a may-influence relation for $\mathcal{B}$.

A Transformation Yielding Protocols with Session Parameters. A simple transformation can produce protocols with session parameters.

The transformation has two parts. The first part prepends before σ a node that transmits a session parameter, and a node that receives a concatenated tuple containing session nonces from each of the other roles:

$$+N_i \Rightarrow -(N_1, \dots, N_{i-1}, N_{i+1}, \dots, N_k) \Rightarrow \sigma$$

In the second part, we transform all encrypted units $\{t\}_K$ contained in σ , to $\{t, \tilde{N}\}_K$, where $\tilde{N}$ is the sequence of all the session nonces introduced in the first step. Thus, letting $\mathcal{T}_{\tilde{N}}$ be this transformation,

Theorem 10. $\mathcal{T}_{\tilde{N}}(\Pi)$ has session parameters for each $\mathcal{T}_{\tilde{N}}(\Pi)$ -bundle $\mathcal{B}$.

It is easy to verify that *No Ambiguity*, *Contribution to Encryptions*, and *Parameter Acquisition* are all true, where the late-arriving parameters $\bar{y}$ are vacuous.

4 The Separability Theorem

Penetrator Paths. The ways that adversary strands manipulate messages are tightly constrained by their syntactic forms. We introduce *penetrator paths* to be able to express these relations conveniently.

Definition 11. A **key node** is the middle node on an adversary encryption or decryption strand, which receives the key to be used (Fig. 2).

A **penetrator path** in $\mathcal{B}$ is a sequence $p = \langle n_0, n_1, \dots, n_k \rangle$ with $k > 0$ and each $n_i \in \text{nodes}(\mathcal{B})$, such that:

1. $n_1, \dots, n_{k-1}$ are all penetrator nodes;
2. if n_i is a reception node and $i < k$, then n_{i+1} is a transmission node and $n_i \Rightarrow^+ n_{i+1}$;
3. if n_i is a transmission node, then $n_i \rightarrow n_{i+1}$ in $\mathcal{B}$.

We often focus on the penetrator paths that stretch from a regular node to a regular node, traversing penetrator strands. These represent activities of the adversary that extract useful materials from regular transmissions, and use them to construct messages to satisfy regular receptions.

We write $p(i)$ for the node n_i , and $|p|$ for k , the number of arrows traversed by p , so $p(|p|)$ is the last node on p . Two paths are shown in Fig. 2. In both cases, $p(0)$ is the hollow circle at the upper left, indicating an unshown regular node, and $p(9)$ is the hollow circle at lower right. One path traverses the edge A in the middle, and the other traverses N_a . We generally write $\text{first}(p)$ and $\text{last}(p)$ for $p(0)$ and $p(|p|)$.

Definition 12. The penetrator path p is **direct** if no key node appears in p , except possibly as $\text{last}(p)$.

$\mathcal{B}$ is **normal** if, on every direct penetrator path, each destructive penetrator strand (decryption, separation) appears before any constructive strand (encryption, concatenation).

The penetrator paths in Fig. 2 are direct. We speak of an *extended path* when we wish to emphasize that it may not be direct.

Lemma 13 ([20]). Every bundle $\mathcal{B}$ has a normal bundle $\mathcal{C}$ lying below $\mathcal{B}$ via the identity id . If $\mathcal{C}$ is any normal bundle, and p is a direct penetrator path in $\mathcal{C}$, then there is a pair of nodes $p_j \rightarrow p_{j+1}$ such that, for all $i \leq j \leq k$:

1. $\text{msg}(p(i)) \sqsubseteq \text{msg}(\text{first}(p))$ and $\text{msg}(p(k)) \sqsubseteq \text{msg}(\text{last}(p))$;
2. If $p(i)$ is an adversary node, then $p(i)$ lies on a destructive strand (decryption, separation); and
3. If $p(k+1)$ is an adversary node, then $p(k+1)$ lies on a constructive strand (encryption, concatenation).

This lemma still holds in our current context, which includes compound keys, because it is restricted to *direct* paths p . Since a key node in p must be the last node, and we never continue along its encryption or decryption strand, we never encounter any case different from those already shown in the proof in [20].

By this lemma, when proving that there exists a bundle lying below $\mathcal{B}$ with a particular property, it is sound to silently assume that $\mathcal{B}$ is normal.

The **bridge term** of a direct penetrator path p in a normal $\mathcal{B}$ is the message $\text{msg}(p(j))$ on the edge that follows all destructive penetrator strands and

precedes all constructive penetrator strands. We will write $\text{bt}(p)$ to refer to the bridge term of p . A single communication edge $\text{first}(p) \rightarrow p(1)$, with no adversary strands in between, is a direct path of length 1; $\text{bt}(p) = \text{msg}(\text{first}(p)) = \text{msg}(p(1))$. The two edges leading to n_1 and n_2 in Fig. 3 are examples with the concatenated bridge terms $\text{c11}, N_1, C_1$ and $\text{c21}, N_2, C_2$. The bridge terms for the two direct paths shown in Fig. 2 are A and N_a . The adversary can always break concatenations down in this way:

Assumption 3. *If $p(i) \rightarrow p(i+1)$ is a bridge term in bundle $\mathcal{B}$, then $\text{msg}(p(i))$ is either an encryption or a basic value, but not a concatenation.*

For any bundle $\mathcal{C}$, there is an equivalent $\mathcal{B}$ in which the adversary separates every concatenated value to its basic or encrypted parts, and then subsequently reconcatenates these parts, as in Fig. 2, Part II [20, Prop. 9]. Assumption 3 restricts our attention to these equivalent but more convenient $\mathcal{B}$.

The direct paths form a framework that supports the extended paths:

Lemma 14. *Let $\mathcal{B}$ be a bundle, and p an extended penetrator path in $\mathcal{B}$ that is not direct. Let $p(i)$ be the earliest key node along p .*

1. *The part of p leading to $p(i)$ forms a direct path.*
2. *Let $p(j)$ be any key node along p , lying on an encryption or decryption strand s , $m_1 \Rightarrow p(j) \Rightarrow m_3$. There are direct paths q such that m_1, m_3 lie on q .*
3. *If s is an encryption strand, then $\text{msg}(p(j)) \ll \text{msg}(\text{last}(q))$. If s is a decryption strand, then $\text{msg}(p(j)) \ll \text{msg}(\text{first}(q))$.*

The Separability Theorem. An extended path p is *critical* iff its source $\text{first}(p)$ is not permitted to influence its target $\text{last}(p)$.

We wish to remove the critical paths, since this will reduce a bundle to one that obeys the influence specification. If the adversary uses a path to influence a node, contrary to our $\leadsto$, we want to clip this path. If we can always remove these paths, and replace a Π -bundle containing critical paths with one with no critical paths, then even the adversary gets no advantage from critical paths. No violation of the influence specification is essential. Everything that can happen in Π can happen without violating the influence specification. If this is true in Π , we can assume $\leadsto$ when analyzing Π ; nothing that matters will be left out.

A sufficient condition for this to hold is that Π 's executions be “reparable:”

Definition 15. *A path p is $\leadsto$ -critical in $\mathcal{B}$ iff $\text{first}(p) \not\leadsto \text{last}(p)$.*

$\mathcal{B}$ is $\leadsto$ -reparable iff $\leadsto$ is a may-influence relation for $\mathcal{B}$, and every $\leadsto$ -critical path p has a bridge $p(i) \rightarrow p(i+1)$ where $\text{msg}(p(i)) = a$ is a basic value.

When $\leadsto$ is understood, we omit it and write “critical” or “reparable.” We can assume no bridge term of p is a concatenation by Assumption 3. Thus, when p is reparable, $\text{bt}(p)$ is a basic value. In Fig. 3, the most interesting bridge terms are N_1 and N_2 , which are the uniquely originating values.

Theorem 16 (Separability). *For every $\leadsto$ -reparable Π -bundle $\mathcal{B}$, there is a Π -bundle $\mathcal{C}$ lying below $\mathcal{B}$ such that $\mathcal{C}$ obeys $\leadsto$.*

Separability for Protocols with Session Parameters. We will first apply Thm. 16 to the main case of protocols with session parameters, and $\sim_s$. The key thing is to show that every critical path is reparable. The main reason why this is true is that—unless $\text{first}(p) \sim_s \text{last}(p)$ and $\text{last}(p) \sim_s \text{first}(p)$ —all encryptions at the two ends contain different sets of session parameters. Thus, the bridge terms are basic values.

Theorem 17. *If Π is a protocol with session parameters, then every Π -bundle satisfying Φ_s is $\sim_s$ -reparable. Hence, by Thm. 16, Π obeys $\sim_s$ subject to Φ_s .*

If each strand succeeds in choosing its session parameter freshly, then no two instances of the same role are related by the causal order in a reduced bundle, i.e. one obeying $\sim_s$. This holds because any two instances supply different values for the session parameter, which are thus incompatible in $\sim_s$.

Theorem 18. *Suppose that Π is a protocol with session parameters, and $\mathcal{B}$ obeys $\sim_s$ and satisfies Φ_s . Then $s_1 \downarrow i \not\leq s_2 \downarrow j$ when (i) $s_1 = \alpha(\rho)$ and $s_2 = \beta(\rho)$; and (ii) $P(\rho)$ is acquired on ρ by step $\min(i, j)$.*

5 Protocol Independence

We turn now from our focus on sessions to combining protocols. We organize the results by the choice of may-influence relation.

The Discrete May-Influence Relation. Let Π_1 and Π_2 be protocols, i.e. sets of strands satisfying the assumptions mentioned in 2–3. For simplicity assume that the protocols are disjoint, in the sense that no strand (or initial segment) is an instance of a role of Π_1 and also an instance of a role of Π_2 . Let $\Pi = \Pi_1 \cup \Pi_2$ be the protocol that contains all the roles of Π_1 and Π_2 .

Define $n_1 \sim_1 n_2$ to hold for $n_1, n_2 \in \text{regnodes}(\Pi)$ just in case $n_1, n_2 \in \text{regnodes}(\Pi_1)$ or $n_1, n_2 \in \text{regnodes}(\Pi_2)$. That is, nodes of the two source protocols may not influence each other.

We can use this *may-influence* relation to infer a protocol independence result, à la [1,9]. Define Π_1, Π_2 to have *sharply disjoint encryption* if

1. every key used for encryption on any node of either is a basic value; and
2. if e_1 is any encryption occurring in Π_1 and e_2 is any encryption occurring in Π_2 , then e_1 and e_2 have no common instance.

The two conditions here are essentially syntactic. Condition 2 says that unification fails for the two encryptions. One way to satisfy condition 2 is using tags. If Π_1, Π_2 may have distinct tags τ_1, τ_2 , such that every encryption in Π_i begins with tag τ_i , then condition 2 is certainly satisfied.

Theorem 19. *If Π_1, Π_2 have sharply disjoint encryption, then all $\Pi_1 \cup \Pi_2$ bundles are $\sim_1$ -reparable. Hence, by Thm. 16, $\Pi_1 \cup \Pi_2$ obeys $\sim_1$.*

This is the essential idea behind [1,9]. The clever extension to algebras with convergent subterm rewrite rules in Ciobaca and Cortier’s [8] appears to involve related ideas.

In our formalism, condition 1 is in fact unnecessary:

Theorem 20. *Let Π_1, Π_2 satisfy Condition 2 of sharply disjoint encryption, and let $\mathcal{B}$ be any bundle of $\Pi_1 \cup \Pi_2$. There is a bundle $\mathcal{C}$ lying below $\mathcal{B}$ such that $\mathcal{C}$ is $\sim_1$ -reparable. Hence, by Thm. 16, $\Pi_1 \cup \Pi_2$ obeys $\sim_1$.*

This shows a pitfall in interpreting strand-based results in the applied pi-calculus. In applied pi, letting $w = \text{hash}(k_1, k_2)$, the two protocols P_1 and P_2 :

$$P_1 = \nu k_1 s . \langle k_1 \rangle . \langle \{t1, s\}_w \rangle \quad P_2 = \nu k_2 s . \langle k_2 \rangle . \langle \{t2, s\}_w \rangle$$

compose to yield $\nu k_1 k_2 s . P_1 \mid P_2$. In strands, by contrast, parameters in individual roles are essentially locally bound, since their possible instances are all substitution instances. Thus, there is no sense in which the two roles share the “same” k_1, k_2 . Moreover, ν -binding expresses a notion of local choice that is somewhat different from both our unique origination and non-origination. It appears to be that the adversary never originates the ν -bound value. Thus, this result appears to be strong, but not truly comparable to results such as [9].

Another limitation of our result is that it is proved for a particular message algebra, and an adversary model for that, rather than for a class of algebras. We conjecture that there is a substantial class for which the lemmas of §§2, 4 hold, and that our results will hold throughout that class.

A One-Way Influence Relation. Here we consider an asymmetric relation between the protocols Π_1, Π_2 . Our goal is to ensure that adding the auxiliary protocol cannot undermine the main protocol Π_1 . In many cases, Π_2 consumes cryptographically prepared units such as digital signatures or encrypted tickets (as in Kerberos), for instance, when it resumes sessions created by the main protocol. Thus, the main protocol may influence the auxiliary, but the reverse should not occur [19]. Let $\Pi = \Pi_1 \cup \Pi_2$, and define $n_1 \sim_2 n_2$ to hold for $n_1, n_2 \in \text{regnodes}(\Pi)$ just in case $n_1 \in \text{regnodes}(\Pi_1)$ or $n_2 \in \text{regnodes}(\Pi_2)$.

With a more delicate definition of *disjoint encryption*, and stipulating the condition Φ that no uniquely originating value is contributed by Π_2 , we obtain:

Theorem 21. *Let Π_1, Π_2 satisfy disjoint encryption, and let $\mathcal{B}$ be any bundle of $\Pi_1 \cup \Pi_2$ satisfying Φ . There is a bundle $\mathcal{C}$ lying below $\mathcal{B}$ such that $\mathcal{C}$ is $\sim_2$ -reparable. Hence, by Thm. 16, $\Pi_1 \cup \Pi_2$ obeys $\sim_2$.*

We may also use this second form of protocol independence to explain the “sequential composition” of Datta et al. [11]. Here, the nodes of the auxiliary protocol are placed after nodes of the primary protocol, but on the same strands; the formalization is unchanged. In particular, h maps nodes of the primary protocol to π_1 and nodes of the secondary protocol to π_2 . The Clause 1 in Defn. 7 allows this to work when nodes of the secondary protocol never appear before a node of the primary protocol on any strand.

Vertical Composition. Suppose that a protocol achieves a goal, assuming that it uses channels that provide particular kinds of protection against the adversary, e.g. that the adversary cannot spoof messages on these channels, or cannot snoop on their contents. Does that yield a secure protocol when these channels are replaced by subprotocols that ensure that the assumptions are met? This is the “vertical composition problem” [14,16,27]. Our methods seem highly relevant to this problem, but they require a way to express the channel assumptions as restrictions on the set of relevant bundles. We plan to explore this.

Conclusion. Two further main areas of future work remain. The more substantial is to adapt this approach to cover a notion of observational equivalence. This appears to involve enriching the adversary model to include a strand that detects the equality of two basic values. We also intend to soften the no V s condition, which is tighter than necessary. For instance, it permits a tuple of messages to be received in a unit, but prohibits these same messages from being received successively, even when there are no intervening transmissions. More careful methods should relax this condition.

Acknowledgments. We are extremely grateful to Véronique Cortier, John Ramsdell, Paul Rowe, and the anonymous referees at POST.

References

1. Andova, S., Cremers, C., Gjøsteen, K., Mauw, S., Mjølsnes, S., Radomirović, S.: Sufficient conditions for composing security protocols. *Information and Computation* (2007)
2. Arapinis, M., Delaune, S., Kremer, S.: From One Session to Many: Dynamic Tags for Security Protocols. In: Cervesato, I., Veith, H., Voronkov, A. (eds.) *LPAR 2008*. LNCS (LNAI), vol. 5330, pp. 128–142. Springer, Heidelberg (2008)
3. Armando, A., Basin, D., Boichut, Y., Chevalier, Y., Compagna, L., Cuellar, J., Hankes Drielsma, P., Heám, P.C., Kouchnarenko, O., Mantovani, J., Mödersheim, S., von Oheimb, D., Rusinowitch, M., Santiago, J., Turuani, M., Viganò, L., Vigneron, L.: The AVISPA Tool for the Automated Validation of Internet Security Protocols and Applications. In: Etessami, K., Rajamani, S.K. (eds.) *CAV 2005*. LNCS, vol. 3576, pp. 281–285. Springer, Heidelberg (2005)
4. Bellare, M., Rogaway, P.: Entity Authentication and Key Distribution. In: Stinson, D.R. (ed.) *CRYPTO 1993*. LNCS, vol. 773, pp. 232–249. Springer, Heidelberg (1994)
5. Bhargavan, K., Corin, R., Deniélou, P.-M., Fournet, C., Leifer, J.J.: Cryptographic protocol synthesis and verification for multiparty sessions. In: *IEEE Computer Security Foundations Symposium* (2009)
6. Canetti, R.: Universally composable security: A new paradigm for cryptographic protocols. Technical Report 2000/067, IACR (October 2001), appeared in *FOCS* (2001)
7. Canetti, R., Krawczyk, H.: Analysis of Key-Exchange Protocols and Their Use for Building Secure Channels. In: Pfitzmann, B. (ed.) *EUROCRYPT 2001*. LNCS, vol. 2045, pp. 453–474. Springer, Heidelberg (2001)
8. Ciobăcă, Ș., Cortier, V.: Protocol composition for arbitrary primitives. In: *CSF*, pp. 322–336. IEEE Computer Society Press (July 2010)

9. Cortier, V., Delaune, S.: Safely composing security protocols. *Formal Methods in System Design* 34(1), 1–36 (2009)
10. Cortier, V., Warinschi, B., Zălinescu, E.: Synthesizing Secure Protocols. In: Biskup, J., López, J. (eds.) *ESORICS 2007*. LNCS, vol. 4734, pp. 406–421. Springer, Heidelberg (2007)
11. Datta, A., Derek, A., Mitchell, J.C., Pavlovic, D.: A derivation system and compositional logic for security protocols. *Journal of Computer Security* 13(3), 423–482 (2005)
12. Delaune, S., Kremer, S., Ryan, M.D.: Composition of password-based protocols. In: *Proceedings of the 21st IEEE Computer Security Foundations Symposium (CSF 2008)*, pp. 239–251. IEEE Computer Society Press (June 2008)
13. Deniéloù, P.-M., Yoshida, N.: Dynamic multirole session types. In: *POPL*, pp. 435–446 (2011)
14. Dilloway, C., Lowe, G.: Specifying secure transport channels. In: *CSF*, pp. 210–223. IEEE (2008)
15. Dolev, D., Yao, A.: On the security of public-key protocols. *IEEE Transactions on Information Theory* 29, 198–208 (1983)
16. Groß, T., Modersheim, S.: Vertical protocol composition. In: *CSF*, pp. 235–250. IEEE (2011)
17. Guttman, J.D.: Security Goals and Protocol Transformations. In: Mödersheim, S., Palamidessi, C. (eds.) *TOSCA 2011*. LNCS, vol. 6993, pp. 130–147. Springer, Heidelberg (2012)
18. Guttman, J.D.: Shapes: Surveying crypto protocol runs. In: Cortier, V., Kremer, S. (eds.) *Formal Models and Techniques for Analyzing Security Protocols*. Cryptology and Information Security Series, IOS Press (2011)
19. Guttman, J.D., Thayer, F.J.: Protocol independence through disjoint encryption. In: *Computer Security Foundations Workshop*. IEEE CS Press (2000)
20. Guttman, J.D., Thayer, F.J.: Authentication tests and the structure of bundles. *Theoretical Computer Science* 283(2), 333–380 (2002)
21. Honda, K., Vasconcelos, V.T., Kubo, M.: Language Primitives and Type Discipline for Structured Communication-Based Programming. In: Hankin, C. (ed.) *ESOP 1998*. LNCS, vol. 1381, pp. 122–138. Springer, Heidelberg (1998)
22. Honda, K., Yoshida, N., Carbone, M.: Multiparty asynchronous session types. In: *Proc. of POPL*, vol. 43(1), pp. 273–284. ACM (2008)
23. Kamil, A., Lowe, G.: Analysing TLS in the strand spaces model. *Journal of Computer Security* 19(5), 975–1025 (2011)
24. Küsters, R., Tuengerthal, M.: Composition theorems without pre-established session identifiers. In: *CCS*, pp. 41–50. ACM (2011)
25. Lowe, G.: Breaking and Fixing the Needham-Schroeder Public-Key Protocol Using FDR. In: Margaria, T., Steffen, B. (eds.) *TACAS 1996*. LNCS, vol. 1055, pp. 147–166. Springer, Heidelberg (1996)
26. Millen, J.K., Shmatikov, V.: Constraint solving for bounded-process cryptographic protocol analysis. In: *CCS*, pp. 166–175. ACM (2001)
27. Mödersheim, S., Viganò, L.: Secure Pseudonymous Channels. In: Backes, M., Ning, P. (eds.) *ESORICS 2009*. LNCS, vol. 5789, pp. 337–354. Springer, Heidelberg (2009)
28. Needham, R., Schroeder, M.: Using encryption for authentication in large networks of computers. *CACM* 21(12) (December 1978)
29. Rusinowitch, M., Turuani, M.: Protocol insecurity with finite number of sessions is NP-complete. In: *Computer Security Foundations Workshop*, pp. 174–187 (2001)
30. Thayer, F.J., Herzog, J.C., Guttman, J.D.: Strand spaces: Proving security protocols correct. *Journal of Computer Security* 7(2/3), 191–230 (1999)