Layered Analysis of Security Ceremonies

Giampaolo Bella, Lizzie Coles-Kemp

Layered Analysis of Security Ceremonies

Giampaolo Bella[1,2] and Lizzie Coles-Kemp[3,4]

[1] Dipartimento di Matematica e Informatica, Università di Catania, Italy
[2] Software Technology Research Laboratory, De Montfort University, UK
giamp@dmi.unict.it

[3] Information Security Group, Royal Holloway University of London, UK
[4] School of Computer and Security Science, Edith Cowan University, Australia
lizzie.coles-kemp@rhul.ac.uk

Abstract. A security ceremony expands a security protocol with everything that is considered out of band for it. Notably, it incorporates the user, who, according to their belief systems and cultural values, may be variously targeted by social engineering attacks. This makes ceremonies complex and varied, hence the need for their formal analysis aimed at their rigorous understanding.
Formal analysis in turn requires clarifying the ceremony structure to build a ceremony model. The model defined here spans over a number of socio-technical layers, ranging from a computer network to society. It inspires a layered analysis of security ceremonies, that is layer by layer. This paper focuses on the human-computer interaction layer, which features a socio-technical protocol between a user persona and a computer interface. Future work will be to traverse all layers by formal analysis.

1 Introduction

The awareness that computer security is much more than a technical issue consolidated through the last decade. Today, it can be claimed that accomplishing the security goal in practice requires heterogeneous and combined efforts at least from computer scientists, social scientists, experimental psychologists, cognitive scientists, and web and HCI designers. Most recent flaws are not due to technical deficiencies but to social engineering techniques to fool the human user into making mistakes [1]. This growing awareness inspired the concept of *ceremony* a few years back, introduced by Walker and elaborated further by Ellison [2], which expands a security protocol to include whatever was left out-of band. In particular, a ceremony recognises the role played by the human, who is "likely to do incomplete comparisons of values, for example" [2]. Capturing this incomplete and partial behaviour is essential for ceremony analysis that reflects real-world interactions and behaviours. Socio-technical modelling in security [3] typically models actions on information objects. Whilst this is valuable, it does not account for the different patterns of practice that are influenced by social and personal factors. In recognising patterns of practice, we capture a range of incomplete comparisons of value that better reflect real-world behaviours.

Motivation. Many insights derive from putting a protocol in a context [4,5], which may involve a number of technical and social elements. A simple but representative note is that protocols involving public-key cryptography cannot be used safely in contexts where certification authorities cannot be access reliably. Also, the protocol users cannot be assumed to always act as anticipated by the protocol designers. These notes demonstrate the context sensitive nature of technology. They also indicate that it is more than context at work: the circumstances of technology use, the cultural values, the belief systems and the demographic factors, simplified below as the user's *persona*, play a substantial role in shaping user responses.

Cooper acknowledged human personas in the design process [6,7] but, to our knowledge, a framework that incorporates personas in the *formal analysis* of security ceremonies is not yet available — though strongly desirable. We shall see that our personas are more fine-grained than Cooper's as each user can express a variety of them in front of the same technology at different times. We recall that formal analysis means specification and verification whose syntax and semantics are mathematically founded. Traditional formal analysis of security protocols evaluates them against given security goals. This paper argues that, due to the complex nature of persona interaction with information protocols, it is valuable to graphically map out the interactions in order to visualise them prior to analysing them formally.

The complexity of ceremonies is acknowledged by many. Radke et al. observe that "a different context, even for the same set of protocols, is a different ceremony" [8]. We shall see that such a context may in turn entail various personas. For example, while approaching online banking, a user may express a very cautious persona at first, and a more relaxed one after reassurance from their friends. Karlof et al. appear to deduce that security can be established by constraining the user interaction with a number of *forcing functions* [9] such as the impossibility to continue unless a box is ticked. However, this may not always be viable, for example because there are many perspectives on privacy, and individuals construct notions of privacy in individual ways [10]. Therefore, a population of Internet service users will engage with a service provider's privacy policy in many ways, as privacy fieldstudies show [7,11]. In consequence, the challenge to develop methodologies for the analysis of general ceremonies where humans interact in non-deterministic ways is significant.

Contribution. In facing such a complexity, the first step of our research was to map the full structure of a security ceremony. The map can be viewed as a model that turns out to feature multiple layers, spanning from a computer network through the human and over to society. Depending on the analyst's focus, certain layers of the model can be conveniently collapsed and assumed to function. Layers can then be analysed in isolation, and future work is to tackle them in combination. Incidentally, "security" is used loosely here to refer to any security property, such as certification, confidentiality and also privacy. A ceremony will be termed accordingly to its main security goal, for example as a certification ceremony, a confidentiality ceremony or a privacy ceremony.

Our model enables a layered view of security ceremonies, and allows for formal analysis of ceremonies from different perspectives. The present paper concentrates on the human-computer interaction layer, III in the model. This layer features a protocol between a user persona and a computer interface, hence a *socio-technical protocol*. The analysis proceeds from the development of a graphical representation of a socio-technical protocol, which in turn demands an implicit encoding of human personas. As a result, our representation can show the various paths of execution of a protocol, each path implicitly encoding a specific persona. This is demonstrated over the socio-technical protocol of an example security ceremony whereby an Internet user registers with a service provider. This ceremony, which is widespread at present as many services require registration, sees the user enter some personal information, hence it is a privacy ceremony as the main security goal is the user's privacy.

It must be remarked that our ceremony model does not intend to prescribe a specific formal method. By contrast, it aims at providing a common canvas on which to use the methods that seem most appropriate. The method we choose for the human-computer interaction layer rests on mathematical induction for the sake of specification, and in inspired by security protocol analysis [12]. Properties of interest can then be assessed through the corresponding induction principle. Known strengths are the support offered by a tool running on a computer, a theorem prover, and the capacity to reason about systems of unbounded size.

Paper summary. A general model of security ceremony is defined (§2), and an example ceremony is outlined (§3). A graphical representation of a ceremony layer is proposed and demonstrated on an example ceremony (§4). Our method of formal analysis of ceremonies is described and demonstrated on that layer of the example ceremony (§5). Conclusions outline on-going and future work (§6).

2 A Security Ceremony Model

As mentioned above, a security ceremony expands a security protocol with the out-of-band, notably with the user [2]. However, a number of layers must be traversed for a security property that the protocol enforces to reach the human. This observation inspires the multiple-layer model of security ceremony pictured in Figure 1, which provides the basis of our formal analysis.

It can be seen that several layers are identified, going beyond the ceremonies between users and systems as described by Ellison. Our model is capable of capturing additional interaction layers between users and technology. This aim is also supported by Whitworth, who advances the importance of the interfaces between humans and machines, and acknowledges the outermost layer whereby society influences by any means, such as word of mouth and publicity, the humans' engagement with technology [13]. Therefore, we are oriented to see the workbench for a socio-technical security analysis as a finer-grained picture.

The layers in Figure 1 feature various abstractions of two example users Alice and Bob, additionally expanded with *Society* — such abstractions will be termed *players*. The (yellow) small boxes indicate the players. From right to left, they

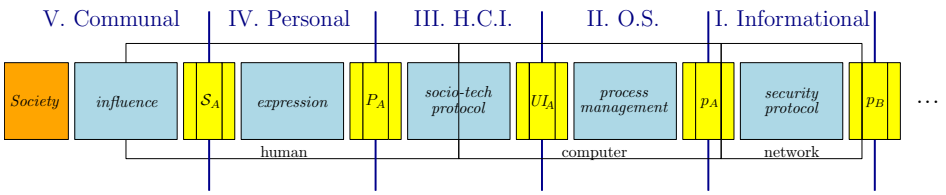

Fig. 1. The security ceremony model underlying our formal analysis

respectively are the computer process p_B running a security protocol with Alice on Bob's behalf and the computer process p_A running the security protocol with Bob on Alice's behalf. Then come the graphical user interface UI_A for her, a generic persona P_A of hers, and finally Alice as a human, that is, her self $\mathcal{S}_A$. The layers can be understood as follows.

Layer I. Informational concerns the security protocol running between computer processes in order to secure Alice and Bob's communications over a potentially insecure network.

Layer II. Operating System manages the inter-process communication between the process that executes the security protocol on behalf of a user and the process that runs the graphical interface presented to that user.

Layer III. Human-Computer Interaction indicates the socio-technical protocol whereby a user interacts with a graphical interface, typically by filling forms in. This is clearly a technical protocol because of the interaction with technology, but is deeply intertwined with the social protocols regulating the individuals' expressions of social capabilities such as trust, recommendations and advice [14]. Precisely, the user is not involved directly but through one of their personas, expressed through the outermore layer.

Layer IV. Personal pertains to the user expression of a persona to engage with specific technology. According to Cooper [6], a persona is a realistic description of a user of the design, with their specific attitudes and goals. We take a yet finer-grained definition, upon the basis that a user may express various personas, as we shall see (§4). For example, arguably when accessing on-line bank services to pay for bills, users express different personas than when accessing their Facebook accounts to catch up with friends. Or they can be more unwilling to download new applications when in the middle of bank transactions than when attempting to share contents with friends.

Layer V. Communal reflects the reciprocal influence of society over individuals. For example, a national campaign could influence users towards being more careful in opening attachments from unknown sources.

Adjacent layers share a player, who plays in both layers. Each layer features what is termed an *interaction* between a pair of players, such as the socio-technical protocol or the security protocol. The (blue) rectangles indicate the interactions. Players only interact within a layer. Interacting players may not belong to the same user, as is the case of layer I. An early version of this model,

which is also inspired by Whitworth's socio-technical research [13], was already published [15]. The version presented in this paper fixes players and interaction of layer III thanks to the work presented below. Most importantly, it also points out a new interpretation of the model, where layers can be conveniently collapsed as a concertina.

3 Example Ceremony Outline

Privacy ceremonies are particular security ceremonies whose main goal is privacy. Typically, they are studied through fieldwork that relies on surveys as a traditional means of measuring privacy attitudes.

An example of privacy ceremony sees an Internet user register with a service provider. Two surveys indicated a range of privacy practices and attitudes when humans engage with on-line services [16,17]. The surveys confirmed the paradox traditionally found in privacy literature [18,19,20] that users want autonomy over on-line privacy but are prepared to trade their privacy in return for some reward. The VOME project [21] took a multi-method research approach to observe and denote the personas of the engaging users [22,23]. In noting that service users have varying senses about how their private information will be treated by the service provider, the fieldwork identified two main personas. One is goal-driven and accepts the provider's conditions in order to obtain its service. Another common one queries the provider and takes various actions to either obtain more information about the provider's privacy policy or to disengage from the registration process. Each of these personas need to be modelled in such a way that their particular security implications can be analysed.

For example, the current registration page with Amazon implements a very basic privacy ceremony: after the user information is entered, an account can be created with just one click. It is clear that simplicity is boosted with the aim of enhancing usability and therefore engagement. However, arguably those users expressing the second persona outlined above, who seek explanations and reassurances about their privacy, feel that their privacy is not being adequately treated by the present ceremony, and hence are unlikely to engage. The following sections describe our formal approach to capturing these complex interactions.

4 Representing Layer III of Security Ceremonies

Layer III, termed human-computer interaction, sees a socio-technical protocol between a persona and an interface. We develop a graphical representation of this layer taking advantage of an encoding of personas. While this is useful for the formal analysis that will follow (§5), it is valuable on its own as it may be yet more widely understandable than a formal specification. The graphical representation provides a means of extending the informal descriptions of privacy ceremonies. We give an example below based on multi-persona registration using various specified notions, highlighted in italics. Working with security senses enables the

map of persona interaction to be more expressive, enabling the formal analysis to be more nuanced in articulating different persona interactions.

With security ceremonies in general, personas may have various security *senses*. A security sense is *a feeling that contributes to an emotional or attitudinal position at a given time on the security goals that the ceremony aims at achieving*. Therefore, we shall be allowed to talk about confidentiality senses, or authentication senses or privacy senses, depending on the type of ceremony being analysed. Four basic meaningful senses, which could of course be refined and specialised, can be described:

caution is the main sense that personas have at the beginning of their interaction with an interface. This sense is often engendered by the public discussion of security concerns.

puzzlement exemplifies the perception that something may be going wrong or out of the user control. The fieldwork indicates that it may also characterise personas expressed (through layer IV) by experienced users.

confidence is the positive sense that the interaction with the (interface of the) service is going satisfactorily. A common requirement is that confidence should always accompany the completion of a session.

unconfidence is the negative sense that the interaction with the service is not going satisfactorily. A common requirement is that unconfidence should never arise, especially out of session completion.

From examples of persona interaction with the ceremony, we can abstract that security senses may arise in combination, namely in security *stances*. A stance is a set of senses, that is an *emotional or attitudinal position at a given time on the security goals that the ceremony aims at achieving*. The number of potential stances on four senses is sixteen, which is the cardinality of the powerset of the set of senses. For our treatment, it is useful to number the potential stances, for example as in Figure 2. The fieldwork confirms that some stances are practically rare, such as the empty stance, St0, or contradictory, such as those stances featuring both confidence and unconfidence, St10, St13 and St15. Other stances are rather implausible, such as those featuring both puzzlement and confidence, St8 and St12 (and obviously St15).

After security senses and stances, the possible *action names*, also termed *cues*, used in the ceremony to analyse must be defined. It must be remarked that whenever user and provider are mentioned in the following, they act respectively through their persona and interface. Our example ceremony rests on these cues:

Begins indicates a persona's wish to initiate a session with a service provider, such as typing in its URL to obtain the main interface.

Registers expresses a persona's successful completion of a session.

Aborts expresses a persona's unsuccessful completion of a transaction.

Compels indicates the interface use of forcing functions whereby a user is forced to do some action in order to proceed with the interaction. For example, the provider's privacy policy must be accepted by a tick to continue.

Queries represents a request of additional clarification or information in general, which both persona and interface may make through specific clicks.

St0 = {}
St1 = {caution}
St2 = {puzzlement}
St3 = {confidence}
St4 = {unconfidence}
St5 = {caution, puzzlement}
St6 = {caution, confidence}
St7 = {caution, unconfidence}
St8 = {puzzlement, confidence}
St9 = {puzzlement, unconfidence}
St10 = {confidence, unconfidence}
St11 = {caution, puzzlement, confidence}
St12 = {puzzlement, confidence, unconfidence}
St13 = {confidence, unconfidence, caution}
St14 = {unconfidence, caution, puzzlement}
St15 = {caution, puzzlement, confidence, unconfidence}

Fig. 2. The potential security stances

Explains indicates the release of additional information that both persona and interface may do, such as the provider's (uncommon though desirable) practice of providing additional information about their privacy policy, which could ultimately be negotiated with the user.

An *action* is an occurrence of a cue between two players, such as "P Begins I" to indicate a persona P who begins interacting with an interface I. Actions and their implications on security senses were explored in the qualitative research [23], highlighting that actions influence the security senses of a persona. In particular, each stance depends on the previous stance and on the action just taken. Therefore, it is convenient to define an *episode* as the pair consisting of an action and the stance it leads to. Upon this basis, a *socio-technical protocol* can be defined as a *list of episodes*, hence its players are identified and also the sequence of stances that the user takes throughout. Therefore, *a persona is the sequence of stances taken during a socio-technical protocol*, which form an expression of the user's while the user interacts. The advantage of this approach is that we are able to move beyond simple action-object socio-technical expression and introduce the notion of patterns of practice influenced by different personal and social factors (represented here as sequences of stances).

The dynamics of layer III of the ceremony to analyse can be now represented as a graph, accounting for a number of personas. Our example ceremony is built on the actual fieldwork [23,22]. It maps typical responses to the request to disclose personal data. Layer III of the resulting ceremony, featuring all stances except the contradictory or implausible ones, can be represented as a graph, precisely a tree, which is in Figure 3. Reading the stances on each path describes a persona, while reading the episodes describes a socio-technical protocol.

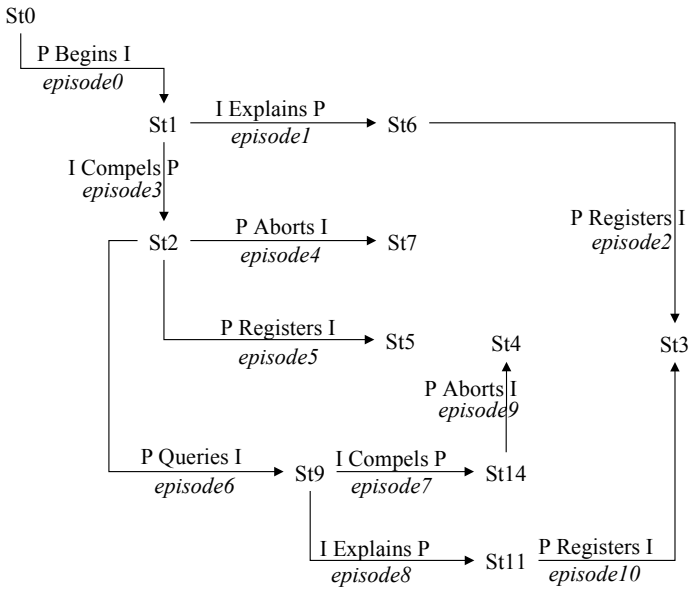

Fig. 3. Graphical representation of layer III of the example ceremony: each path encodes a specific persona interacting with a user interface, i.e. a socio-technical protocol

In particular, path St0, St1, St2, St7 is the persona that aborts with a stance of caution with unconfidence because compelled by a forcing function; path St0, St1, St2, St5 is the persona that gets a stance of caution with puzzlement following registration. The path departing from stance St1, a sense of caution, and leading to registration with stance St3, a sense of confidence, indicates the persona facing an interface where the provider explains their privacy policy adequately and then exits. Another path departing from stance St2, a sense of puzzlement, represents a persona who wants to understand how their privacy is going to be managed, and therefore poses queries. The path branches towards abortion when the provider refuses to provide adequate explanation, or towards registration when the provider accepts to dialogue with the persona on privacy.

Figure 3 communicates several aspects of the ceremony layer. With the mentality of traditional security protocol analysis, it can be understood as a protocol with branches. From the standpoint of socio-technical analysis, it is the complete layer III: each of its paths encodes specific persona and interface, and also the socio-technical protocol between them. Our representation also appears to be flexible. For example, the sequence of actions in a path could be repeated in a parallel path featuring difference stances to account for different personas taking the same actions.

The outcome of formal analysis can then be used to motivate refinements of the underlying social theory as well as modifications to the technical protocols to better support the required stances. Also, it must be remarked that our ceremony

representation does not aim at a descriptive model of human behaviour, as is often found in statistical analysis. Rather, human behaviour is represented as personas, in turn encoded as combinations of actions and stances.

5 Analysing Layer III of Security Ceremonies Formally

This Section explores the type of analysis that may be performed once the Layer III interactions are mapped. The formal analysis of layer III of security ceremonies is challenging because of the variety of socio-technical protocols and personas to account for. In order to include Layer III within formal security analysis, the ceremonies first need to be mapped so that systematic analysis can take place. Our formal analysis at this layer aims at providing a mathematically-solid ground to the analysis of ceremonies. This will help ceremony designers predict the potential effects of the various combinations of stances and actions, and to suggest possible refinements of such combinations in step with a refinement of the underlying social theory. In other words, the formal analysis can be used in many ways, such as to foresee what protocol may lead to a desired stance or, vice versa, to evaluate the stances of a protocol.

A variety of formal approaches can be taken, such as *model checking* [24] and *theorem proving* [25]. Here, we opt for the latter because we aim at deriving formal statements, supported by mathematical proof, about the security stances of the players interacting at layer III. Another reason is to answer queries of the form: if a user registers with a provider without getting any explanation, will he feel confident? The theorem prover we use is Isabelle [25], taking inspiration from vast work conducted to analyse security protocols [12]. The main specification strategy that is adopted is mathematical induction, which allows for the compact specification of systems of unbounded size, such as the natural numbers. The main reasoning strategy is the corresponding induction principle. A brief outline of Isabelle is in Appendix A.

5.1 Formal Specification of the Example Ceremony

This Section takes the graphical representation of a ceremony seen above (§4) and reformulates it using a formal language (whose syntax and semantics are mathematically defined).

A type for players and a datatype of cues are defined:

```
types player = nat
datatype cue = Begins | Queries | Registers | Aborts
             | Compels | Explains
```

Actions are defined as the cartesian product between players, cues and players. An action is thus formalised as a triple, making it possible to denote, for example, actions involving a specific cue among generic players or vice versa:

```
types action = "entity * cue * entity"
```

Another datatype formalises security senses, while a stance can be introduced by an intuitive notation:

```
datatype sense = caution | puzzlement | confidence | unconfidence
types stance = "sense set"
```

We also decide to formalise *episodes* as the cartesian product of actions and stances. An episode then is a pair consisting of an action and a stance:

```
types episode = "action * stance"
```

In order to explain our method, it is useful to introduce two example players, that is to define two constants:

```
consts PP::entity
consts II::entity
```

Here is a specific example episode with a specific stance including two senses:

```
definition ep :: episode where
        "ep == ((PP, Begins, II), {caution, confidence})"
```

It can now be demonstrated how to prove obvious lemmas about the constants:

```
lemma "fst ep = (PP, Begins, II)" by (simp add: ep_def)
lemma "snd ep = {caution, confidence}" by (simp add: ep_def)
lemma "snd ep = {confidence, caution}" by (auto simp add: ep_def)
```

The first lemma states what the first element of example episode `ep` is; its Isabelle proof script consists of a single method (on the same line as the lemma statement) which invokes the simplifier (via `simp`) with an appeal to the definition of the episode, automatically stored as `ep_def`. The second lemma is homologous for the second component. The third lemma is equivalent to the second one, but also confirms that because stances are defined as sets, the security senses they feature can be indicated in any order. It can be observed that this simple form of reasoning based on set theory cannot be handled by the simplifier alone, but necessitates some classical reasoning (via `auto`).

A socio-technical protocol is defined coherently with what was seen above (§4) — each path in the ceremony in Figure 3 is of this type:

```
types protocol = "episode list"
```

Finally, the constant `hci` can be declared:

```
inductive_set hci :: "protocol set"
```

This constant, which formalises the full layer III, is defined by structural induction on the length of a protocol. Its definition is omitted due to space limits.

5.2 Formal Verification of the Example Ceremony

This Section outlines the main guarantees we have proved about our model of layer III of the example ceremony. It can be seen how existential or universal quantifiers are used to strengthen the conclusions. Full proof scripts are omitted.

A first relevant statement is a theorem stating that a persona always aborts with a sense of unconfidence:

```
theorem P_aborts_unconfident:
"⟦ ((P,Aborts,I), sigma) ∈ set stp; stp ∈ hci ⟧
⟹ unconfidence ∈ sigma"
```

More formally, theorem `P_aborts_unconfident` states that, given a generic protocol `stp` of the model, featuring an episode whose action sees a generic agent abort with a generic provider, and whose stance is generic, then that stance can be proved to feature `unconfidence`.

Another theorem states when the sense of unconfidence is in a stance:

```
theorem unconfident_P_when:
"⟦ (alpha, sigma) ∈ set stp; unconfidence ∈ sigma; stp ∈ hci ⟧
⟹ ∃ P I. alpha = (P,Aborts,I) ∨
         alpha = (P,Queries,I) ∨
         alpha = (I,Compels,P)"
```

Precisely, theorem `unconfident_P_when` insists on a generic protocol featuring an episode that is generic except for the fact that its stance includes `unconfidence`. The theorem concludes that, for some specific persona `P` and provider `I`, the action in the precondition could only be of three forms.

A subsidiary lemma is useful to prove subsequent theorems:

```
lemma I_Ca_Co_explains:
"⟦ ((I, gamma, P), sigma) ∈ set stp; caution ∈ sigma; confidence ∈ sigma;
   stp ∈ hci ⟧
⟹ gamma = Explains"
```

Lemmas often have their own significance. This says that whenever a generic episode appears but its stance features `caution` with `confidence`, then the cue that is involved must be `Explains`.

An important theorem states the circumstances for registration to take place:

```
theorem P_registers_confident_when:
"⟦ ((P,Registers,I), {confidence}) ∈ set stp; stp ∈ hci ⟧
⟹ ∃ sigma. ((I,Explains,P), sigma) ∈ set stp  ∧
           caution ∈ sigma ∧ confidence ∈ sigma"
```

More formally, theorem `P_registers_confident_when` insists on a protocol featuring a specific registration episode. It then concludes that a corresponding explanation episode must occur on the same protocol. Further, it specifies the stance of the latter episode, which must include `caution` with `confidence`.

A final theorem specifies the stance of a registration episode appearing in a protocol where no corresponding explanation episode occurs:

```
theorem P_registers_without_confidence:
"⟦ ((P,Registers,I), sigma) ∈ set stp;
   ∀ sigma'. ((I,Explains,P), sigma') ∉ set stp; stp ∈ hci ⟧ ⟹
   confidence ∉ sigma"
```

Precisely, theorem `P_registers_without_confidence` assumes a protocol with a registration episode featuring a stance `sigma`. Also, it assumes that no explanation episode, for no possible stance `sigma'`, ever occurs on that protocol. Upon these preconditions, it concludes that `sigma` cannot feature the `confidence` sense.

6 Conclusions

The contribution of this paper is threefold: the definition of a detailed model for security ceremonies that relates the technical and the social layers; a graphical representation of the complicated layer that features user personas who interact with human-computer interfaces; the formal analysis of this layer.

There are several benefits of using formal methods at the persona layer. By creating a formal model linking human practices with mental and emotional models that users have about security, it becomes easier to conceptualise of on-line services that respond to particular personas and how this might take place. For example, the modelling of when explanations are required or not required is a valuable tool for the design of on-line services because it provides indicators as to when service providers need and need not intervene with privacy information in the on-line process. In consequence, certain aborts that were often invisible to analysts from the technological, rather than human, perspective may become visible. Aborts are an important aspect of any ceremony because they provide valuable input as to engagement with a particular security or privacy technology.

Formal methods also provide insights that may help detect potential weaknesses in the social and practice theories relating to security and privacy practices. For example, potential ceremony paths that have not surfaced in the field research but that are possible in the formal model can be clearly identified. This can draw field research into considering why certain ceremony paths have not emerged in the grounded research and theorise about why this is so, thus maturing the social and practice theories by identifying the all-important outliers. The formal reasoning also provides a valuable critique to the technology design, pinpointing where the technology design is insufficient for particular personas or where superfluous functionality is provided. Finally, formal analysis is powerful because it is enmeshed in practice-led research. This means that as the fieldwork findings refine, so too does the formal model. Therefore, both model and fieldwork interoperate to refine both the ceremony descriptions and their analysis. The next challenge ahead is to tackle the other layers of our full ceremony model.

Acknowledgements This work was supported by the Technology Strategy Board; the Engineering and Physical Sciences Research Council and the Economic and Social Research Council [grant number EP/G00255/X]. We are grateful to Gabriele Lenzini and Paul Curzon for useful discussions, and to Elahe Kani-Zabihi and Yee-Lin Lai for working in the team that conducted the initial field research used as the basis for this work.

References

1. URL: Microsoft warning over browser security flaw. http://www.bbc.co.uk/news/technology-12325139 (2011)
2. Ellison, C.: Ceremony design and analysis. Technical report, Cryptology ePrint Archive, Report 2007/739 (2007)
3. Pieters, W.: Representing Humans in System Security Models: An Actor-Network Approach. Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications **2** (2011) 75–92
4. Martina, J.E., de Souza, T.C.S., Custódio, R.F.: Ceremonies design for PKIs hardware security modules. In: Proc. of the 9th Brazilian Symposium on Information and Computer System Security (NDSS'09), SBC Press (2009) 115–128
5. Karlof, C., Tygar, J.D., Wagner, D.: Conditioned-safe ceremonies and a user study of an application to web authentication. In: Proc. of the 16th Network and Distributed System Security Symposium (NDSS'09). (2009)
6. Cooper, A.: The Inmates Are Running the Asylum. SAMS Publishing (2004)
7. URL: Primelife project. https://www.primelife.eu (2008)
8. Radke, K., Boyd, C., Nieto, J.G., Brereton, M.: Ceremony analysis: Strengths and weaknesses. In: Proc. of the 26th IFIP International Information Security Conference (IFIP SEC'11). LNCS, Springer (2011)
9. URL: Forcing functions. http://www.interaction-design.org/encyclopedia/forcing_functions.html (2011)
10. Solove, D.J.: Understanding Privacy. Harward University Press (2008)
11. Kumaraguru, P., Cranor, L.F.: Privacy indexes: A survey of westins studies. Technical report, SCS Technical Report Collection (2005)
12. Bella, G.: Formal Correctness of Security Protocols. Information Security and Cryptography. Springer (2007)
13. Whitworth, B.: The Social Requirements of Technical Systems. In: Socio-technical Design and Social Networking Systems. IGI Global (2009) 3–22
14. Jr., J.M.R.: Social Protocols. http://www.w3.org/Talks/980922-MIT6805/SocialProtocols.html (1998)
15. Bella, G., Coles-Kemp, L.: Seeing the Full Picture: the Case for Extending Security Ceremony Analysis. In: Proceedings of the 9th Australian Information Security Management Conference, Secau Security Research Centre Press (2011)
16. URL: Privacy on the internet: Attitudes and behaviours. http://www.vome.org.uk/wp-content/uploads/2010/03/VOME-exploratorium-survey-summary-results.pdf (2010)
17. Hubbard, T., Ampofo, L.: What's out there? an evaluation of online identity management. Technical report (2010)
18. Buchanan, T., Reips, U.D., Paine, C., Joinson, A.N.: Development of measures of on-line privacy concern and protection for use on the internet. Journal of the American Society for Information Science and Technology **58** (2007) 157–165
19. Norberg, P.A., Horne, D.R., Horne, D.A.: The privacy paradox: Personal information disclosure intentions versus behaviors. Journal of Consumer Affairs **41** (2007) 100–126
20. Paine, C., Reips, U.D., Stieger, S., Joinson, A., Buchanan, T.: Internet users' perceptions of 'privacy concerns' and 'privacy actions'. International Journal of Human-Computer Studies **65** (2007) 526–536
21. Coles-Kemp, L., Kani-Zabihi, E.: On-line privacy and consent: A dialogue, not a monologue. In: Proc. of the New Security Paradigms Workshop (NSPW'10), ACM Press (2010)

22. Coles-Kemp, L., Kani-Zabihi, E.: Practice Makes Perfect: Motivating Confident On-line Privacy Protection Practices. In: Proceedings of the 3rd IEEE International Conference on Social Computing (SocialCom'11), IEEE (2011)
23. Coles-Kemp, L., Kani-Zabihi, E.: Service users' requirements for tools to support effective on-line privacy and consent practices. In: Proc. of the 15th Nordic Conference in Secure IT Systems (NordSec'10). LNCS, ACM Press (2010) 106–120
24. McMillan, K.: Symbolic Model Checking. Kluwer Academic Publisher (1993)
25. Nipkow, T., Paulson, L.C., Wenzel, M.: Isabelle/HOL: A Proof Assistant for Higher-Order Logic. Springer (2002) LNCS Tutorial 2283.
26. URL: Isabelle download page. `http://www.cl.cam.ac.uk/Research/HVG/Isabelle/download.html` (2011)

A A primer on the theorem prover Isabelle

Isabelle is a generic, interactive theorem prover [26]. *Generic* means that it can reason in a variety of formal systems. This paper refers to Isabelle/HOL [25], which supports the formal language *higher-order logic*, a typed formalism that allows quantification over functions, predicates and sets, but has no temporal operators. *Interactive* means that it is not entirely automatic and, rather, requires a good amount of human intervention. But Isabelle also provides much automation. Its *simplifier*, which can be invoked by the proof method `simp`, combines rewriting with arithmetic decision procedures. Its *automatic provers* can solve most simple proof scenarios. For example, the proof method `blast` implements a fast classical reasoner, and `auto` combines that with the simplifier.

Most proofs are conducted interactively. In a typical proof, the user directs Isabelle to perform a certain induction and then to simplify the resulting proof state, which may contains a number of subgoals. Each subgoal can be given to an automatic prover or be reduced to other subgoals by the use of some lemma. Failure to find a proof for a conjecture may simply mean that the user is not skilled enough; otherwise, it may exhibit what in the modelled system contradicts the conjecture and hence help in locating a system bug or an erroneous human intuition. The series of commands used to prove a theorem can be seen as a proof sketch. confidence that the proof is sound comes from inspecting the line of reasoning adopted and the lemmas it requires.

The typical Isabelle proof methods mentioned above, `simp` for the conditional term rewriter and `blast` for the classical reasoner, are combined by the proof method `force`, which only applies to the first subgoal in the proof state, and by `auto`, which applies to the entire state. All these commands allow the installation of additional lemmas to appeal to. A very useful method is `clarify`, performing obvious steps, while `clarsimp` interleaves it with the simplifier. A variant of `clarify` is `safe`, which additionally performs the obvious steps that split up the subgoal.