

## 1 Background and Motivation

Every day we exchange sensitive information with other individuals and organizations, often over open communication networks like the internet. Whenever we engage in e-commerce, login to our web-based email service, or sign up for special offers we are required to provide our credit card number, password, or home address to the other party. How can we be sure that sensitive information does not get released to third parties who should not have access to it? The answer depends in large part on the threat model.

The standard solution to securing an open network is to use a cryptographic protocol, such as SSL, Kerberos or WPA, which might allow agents to mutually authenticate each other and to establish new cryptographic keys used to encrypt the subsequent communication. The main goal of such a protocol is to establish a communication channel that prevents an intruder from learning or manipulating any of the data that is sent along it. Unfortunately, cryptographic protocols can be quite complex, and it is not uncommon for flaws or weaknesses to be found in protocols that were thought to be secure.

To make matters worse, new cryptographic protocols are continually being developed. Although newly proposed protocols undergo considerable review by standardization bodies such as the National Institute of Standards and Technology, flawed protocols can still survive the process. The abstract, mathematical analysis of cryptographic protocols has reached a point where it is possible to perform a formal analysis of a protocol *during* the standardization process, hopefully strengthening the protocols that finally result in standards. To this end, I participated in a case study in which our group provided a formal analysis of the EAP-GPSK authentication protocol which was in the midst of the standardization process with the Internet Engineering Task Force, to be discussed below.

While well-designed cryptographic protocols ensure the establishment of secure communication channels, they do nothing to prevent the legitimate parties of the communication from misusing the information that is shared. For example, when paying for a meal at a restaurant, a customer will give her credit card to the waiter. The (physical) channel of handing the card to waiter can be secure, but this does not stop the waiter from recording the customer's credit card number.

More generally, several untrusting individuals or organizations may enter into a collaboration with a temporary and partial alignment of their interests. For example several competing companies may want to forecast the direction of the market in a collaborative way. Such a situation requires the companies to share some information that is quite sensitive. The decision about which pieces of information to share with whom will typically depend on the inhomogeneous trust relationships the organizations have with one another. A common way for companies to manage the trust relationship with their clients is to produce privacy policies which restrict the ways that the company is allowed to use the information it collects. The rest of my work has focused on developing an appropriate model in which to analyze confidentiality in such collaborative settings.

## 2 Analysis of EAP-GPSK

I begin by discussing the work I did while visiting the computer science department at Stanford University in the summer of 2007, hosted by Professor John C. Mitchell. I participated in a project with Professor Mitchell, his student Arnab Roy, and my advisor Professor Andre Scedrov, in which we analyzed a protocol undergoing the standardization process with the Internet Engineering Task Force (IETF), and that led to a publication in the Conference on Applied Cryptography and Network

Security [7]. The IETF is a large international community of network designers, vendors, researchers and operators with an interest in the evolution of the internet. They produce documents containing protocol standards, best current practices and various other informational documents. Participation is open to any interested individual. The protocol we analyzed is known as Extensible Authentication Protocol (EAP) Generalized Pre-Shared Key (GPSK). We performed our analysis at a critical time in the standardization process. We first looked at the fifth draft of the protocol specification [1]. At that time the protocol was sufficiently mature to undergo rigorous formal analysis, yet it had not yet been standardized or deployed in the real world.

The EAP-GPSK protocol is a lightweight, flexible authentication protocol designed to be used in small devices where memory and computation power are scarce. It is part of an ongoing IETF process to develop authentication methods for the EAP framework. We analyzed the protocol and found three weaknesses. The first was a client-side Denial-of-Service attack. This weakness allowed an intruder to flood the network with illegitimate messages overflowing the client's memory thereby blocking further progress of the protocol. The second weakness was an anomaly with the way in which keys are derived. Since cryptographic keys are the lynch pin of any cryptographic algorithm, if they are not created carefully, an intruder may be able to reverse engineer the process to learn the key. This would lead to a complete break of both secrecy and authentication, rendering the protocol useless. While we did not provide an explicit attack against the key-derivation function, its non-standard usage seemed to be a potential vector of attack. The third weakness was a ciphersuite downgrading attack. As part of EAP-GPSK the two users agree upon a ciphersuite, or list of cryptographic algorithms, that they both like. In this attack, the intruder could coax the client into choosing a weak ciphersuite. If it is weak enough then the intruder may be able to break it in real-time, again rendering the protocol useless.

Finding weaknesses is only the first part of analyzing a security protocol. In order for the analysis to be worthwhile we must show that the weaknesses can be fixed. Indeed, in discussions with the IETF working group in charge of this protocol, we offered possible solutions to the above weaknesses which were subsequently adopted in the next draft of the specification which has since completed the standardization process resulting in the final specification document, IETF RFC 5433 [2].

Our main goal in this analysis was to make a solid impact on the protocol's development in a timely manner. To that end, we used a variety of analysis methods at our disposal. First and foremost, we relied on past experience with cryptographic protocols to formulate the protocols in an abstract and formal way. Using this formalization we used the model checker Mur $\varphi$  to detect the Denial-of-Service attack. Finally, we used Protocol Composition Logic [3] to prove some of the stated properties of the fixed protocol.

### 3 Collaboration and Confidentiality

In addition to the work analyzing EAP-GPSK, which focuses on the initial stage of securing communication, I have also been engaged in a project focused on issues of confidentiality in collaborative systems. The work done in this direction [4, 5, 6] was completed jointly with Professor Max Kanovich of Queen Mary University, London and with my advisor Professor Andre Scedrov. The goal of this research is to develop and explore a formal, mathematical model of collaborative systems in which we can express and check certain properties about confidentiality. In order to motivate the most important details and results of the formalism, let us first consider one such system from the real

world that we would like to model.

Consider a patient who needs to have a medical test performed at a hospital. The process begins with the patient registering with the receptionist who anonymizes the patient with some ID number. A nurse then gets a test sample from the anonymized patient (*e.g.* draws blood, obtains DNA, etc.) to send off to the lab. A lab technician performs the test and sends the results to the patient's physician who then gives the appropriate diagnosis and/or prescription to the patient. The patient has some preferences regarding which agents learn (or can learn) certain combinations of data. For example, the lab technician must learn the result of the test and be able to connect it to his ID number. But the patient does not want the technician to connect the test result to his true identity. The patient will have similar but distinct preferences regarding the other agents involved in the process.

This example contains all of the key elements of collaborative systems that we aim to capture with the formalism. It has numerous agents each with different abilities and access to different sets of data. They have a shared goal of providing the patient with the correct diagnosis and proper prescriptions. It also reflects the idea that the confidentiality preferences of the agents are heterogeneous and dependent on the specific context. That is, certain combinations of data is allowed to flow to some agents but not others, and these restrictions depend not only on the other agents involved but also on what actions they may have available to them. In another scenario, each agent might have their own list of confidentiality concerns.

A natural starting point when analyzing a situation such as the one above is to ask whether the agents have a way of collaborating to achieve their shared goal while simultaneously respecting the confidentiality concerns of all the participating agents. In other words, does respecting the confidentiality concerns of all the agents preclude the possibility of a successful collaboration? I will refer to this problem as the Collaborative Planning Problem with Confidentiality. Naturally, the answer to this question will depend on the specifics of the scenario. More subtly, the answer will also depend on the precise definitions of what it means for the confidentiality concerns to be respected. Our formal model allows us not only to ask these questions in a precise and rigorous way, but also to explore several possible interpretations for what it might mean to respect an agent's confidentiality. The main results of our work establish the decidability and complexity of the Collaborative Planning Problem with Confidentiality under a variety of technical modeling choices.

**Summary of the Formalism.** The formalism is a kind of state transition system in which each agent has some local configuration which is only accessible to that agent along with some publicly shared configuration accessible to all the agents. Every agent has a set of actions that describe the ways it can manipulate the system. We assume the actions are local in the sense that each agent's actions can only affect and depend upon the parts of the global configuration which are accessible to that agent. So, for example, to determine if agent  $A$  can apply an action  $a$ , it is sufficient to look at the local configuration of  $A$  and the shared public configuration. Similarly,  $A$ 's actions cannot change the local configuration of other agents. We assume that the system starts in some configuration  $W$  which specifies who has access to the different pieces of data.

In order to specify the goal of a collaboration, it is sufficient to specify the conditions which must be present in a configuration. Thus, in contract negotiations, a typical goal will be any contract which is signed by all parties. It is not necessary to specify the entire text of the contract at the outset. Also, the agents may agree on several goals, any of which could represent a successful outcome of the collaboration. The agents want to find a plan, or a sequence of actions, that take the system

from the initial configuration to one of the goal configurations. (We use the word plan as a result of the similarity of this problem to that of the classical planning problem in artificial intelligence.)

Finally, our formalism provides an explicit mechanism for representing the agents' confidentiality preferences. Namely, each agent has a *confidentiality policy* expressed as a set of "negative goals" which represent configurations that the agent views as undesirable or *critical*. For example, in the medical scenario above, the patient's policy would indicate that any configuration in which the lab technician knows the patient's name in conjunction with the test result should be marked as critical.

Thus the Collaborative Planning Problem with Confidentiality has four main inputs: the actions of the agents, the initial configuration of the system, the goal configuration(s), and the confidentiality policies of the agents. In addition, we offer three ways to define policy compliance with respect to the agents' critical configurations. The definitions may be applicable in different scenarios depending on the level of trust among the agents.

The first definition is *system compliance*. Roughly speaking, a system is compliant if the set of reachable configurations does not contain any critical configurations. If a system is known to be compliant then the agents can freely apply *any* sequence of actions without worrying about reaching a critical configuration. Such a definition is most appropriate when there is a low level of trust among the agents.

The second definition is *weak plan compliance*. A plan is weakly compliant if it does pass through any critical configurations. This definition is much weaker than the previous one. Under this definition the agents are content as long as their critical configurations are never actually reached. It does not matter if there was another course of action that could have produced a critical configuration. This definition is appropriate for situations in which the agents have a relatively high level of trust in one another.

The final definition is *plan compliance*. A plan is compliant if it is weakly compliant, and if at every point along the plan, no subset of agents can deviate from the plan to reach a critical configuration of one of the other agents. In other words, compliant plans are robust against arbitrary coalitions of agents. This definition is also useful for situations with a relatively low level of trust among the agents.

**Summary of Results** I can now summarize the main results of this part of my research. Our goal was to determine the decidability and complexity of the Collaborative Planning Problem under the various definitions of confidentiality. In addition, we showed the effect of a technical restriction on the agents' actions. Normally, an action can potentially grow or shrink the overall size of the global configuration, but our technical restriction forces the configuration size to remain constant. These restricted actions are called well-balanced. The intuitive reason for imposing such a restriction is because it more closely models a situation in which the memory resources are fixed and finite.

There are therefore six computational problems for which one can determine the decidability and complexity: The Collaborative Planning Problem with Confidentiality with respect to each of the three definitions of confidentiality for general actions as well as for systems with well-balanced actions. Figure 1 summarizes the results for five of the six problems. I am currently investigating the sixth problem, and I am confident this result will appear in my PhD thesis.

In addition to these decidability and complexity results, we also explored correspondences between our formalism and other well-established formalisms. In particular, we show that goal reachability in our formalism is equivalent to derivability of certain sequents in affine logic, a variant

	System Compliance	Weak Plan Compliance	Plan Compliance
Well-balanced Actions	PSPACE	PSPACE	—
(Possibly) Un-balanced Actions	EXSPACE	Undecidable	Undecidable

Figure 1: Summary of the published complexity results.

of linear logic that allows the rule of weakening. Similarly, the EXSPACE result is proved via a correspondence between the reachability problem in our formalism and the covering problem for Petri nets. Such correspondences serve to demonstrate the connection between our decidability and complexity results and similar results stated in other formalisms. They give our formalism a solid foundation in well-established work.

## 4 Possible Directions for Future Work

Since EAP-GPSK has completed the standardization process, there is not an obvious way to extend that particular analysis. However, new protocols are being developed every day, and they could all benefit from a sound, formal analysis. While some automated tools exist for verifying the correctness of cryptographic protocols, they are not yet usable by non-experts. A program which can automatically verify security properties of a protocol that is also usable by industry professionals would be a very valuable contribution.

The research regarding confidentiality in collaborative settings leads naturally to a wide variety of new questions that can drive further research in the area. To begin with, the formalism can be enriched in a number of ways. For instance, it is common for cryptographic protocols to require the agents to produce fresh, random values. Indeed, many algorithms require the use of random values to work properly. The formalism could be enriched by including a mechanism to model the generation of such fresh values. We could also relax the locality of the actions and allow actions to produce facts directly in other agents' local configurations. This could model more direct communication channels.

It would also be natural to introduce policies that can more accurately represent the kinds of privacy policies that are already implemented by many organizations. Namely, policies tend to place restrictions and obligations on actions instead of on the data itself. In addition, it could be useful to explicitly model time in the formalism in order to reason about time-sensitive policies. Perhaps using a synchronous version of this system would be appropriate for such purposes. One could continue to explore the decidability and complexity of the Collaborative Planning Problem with Confidentiality under various combinations of these features.

The system could be extended to capture the notion of data provenance, which is useful for making decisions based on the integrity of the data. If the agents have local policies about which pieces of data they can trust, then the problem of reaching the goal could be more complex. Keeping track of data's origin is also a possible way to model auditable systems. If confidentiality is violated, an audit might be able to reveal the responsible agent.

The agents in the system would be better served if they had *local* algorithms that allowed them to reason about confidentiality properties of the system. Concepts from mechanism design may be useful for such purposes. It is possible that introducing utility functions to the formalism could lead to some new avenues for analysis. For example, the definition of plan compliance is somewhat reminiscent of a Nash equilibrium in game theory. Using utility functions might shed more light on this connection.

Another interesting avenue of research might be in trying to implement such a system and apply it to several real-world case studies. Given the high complexity of the problems, it would be wise to start with small examples, but some experience with actual systems could provide insight into appropriate ways to restrict the system in order to increase the efficiency of the algorithms. Such investigations may also introduce more natural questions that arise in this collaborative context.

## References

- [1] T. Clancy and H. Tschofenig. EAP Generalized Pre-Shared Key (EAP-GPSK), (work in progress). April 2007. <http://www.ietf.org/internet-drafts/draft-ietf-emu-eap-gpsk-05.txt>.
- [2] T. Clancy and H. Tschofenig. Extensible Authentication Protocol - Generalized Pre-Shared Key (EAP-GPSK) Method, RFC 5433, February 2009. <http://tools.ietf.org/html/rfc5433>.
- [3] A. Datta, A. Derek, J. C. Mitchell, and A. Roy. Protocol Composition Logic (PCL). *Electron. Notes Theor. Comput. Sci.*, 172:311–358, 2007.
- [4] M. Kanovich, P. Rowe, and A. Scedrov. Collaborative planning with confidentiality. *Submitted*. [http://www.math.upenn.edu/~rowep/CPWC\\_full.pdf](http://www.math.upenn.edu/~rowep/CPWC_full.pdf).
- [5] M. Kanovich, P. Rowe, and A. Scedrov. Policy compliance in collaborative systems. *Submitted*. [http://www.math.upenn.edu/~rowep/Policy\\_Compliance.pdf](http://www.math.upenn.edu/~rowep/Policy_Compliance.pdf).
- [6] M. Kanovich, P. Rowe, and A. Scedrov. Collaborative Planning with Privacy. In *A. Sebelfeld, ed., "20-th IEEE Computer Security Foundations Symposium (CSF)"*, pages 265–278. IEEE Computer Security Press, July 2007.
- [7] J. C. Mitchell, A. Roy, P. Rowe, and A. Scedrov. Analysis of eap-gpsk authentication protocol. In *ACNS*, pages 309–327, 2008.