THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

miTLS: Verifying Protocol Implementations against Real-World Attacks

Citation for published version:

Bhargavan, K, Fournet, C & Kohlweiss, M 2016, 'miTLS: Verifying Protocol Implementations against Real-World Attacks', *IEEE Security and Privacy*, vol. 14, no. 6, pp. 18-25. <https://doi.org/10.1109/MSP.2016.123>

Digital Object Identifier (DOI):

[10.1109/MSP.2016.123](https://doi.org/10.1109/MSP.2016.123)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

IEEE Security and Privacy

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

miTLS: Verifying Protocol Implementations Against Real-World Attacks

Karthikeyan Bhargavan karthikeyan.bhargavan@inria.fr, Cédric Fournet fournet@microsoft.com,
Markulf Kohlweiss markulf@microsoft.com

The TLS Internet Standard, previously known as SSL, is the default protocol for encrypting communications between clients and servers on the Web. Hence, TLS routinely protects our sensitive emails, health records, and payment information against network-based eavesdropping and tampering. For the last 20 years, the security of TLS has been analyzed in a variety of cryptographic and programming models, in order to establish strong formal guarantees for various configurations of the protocol. Yet, TLS deployments are still often found to be vulnerable to attacks, and still rely on security experts to fix the protocol implementations.

The miTLS project intends to solve this apparent contradiction between published proofs and real-world attacks, which reveals a gap between the theory and practice of TLS. To this end, we jointly develop a verified reference implementation and a cryptographic security proof that account for the low-level details of the protocol. The resulting formal development sheds light on recent attacks, yields security guarantees for typical usages of TLS, and informs the design of the next version of the protocol.

D.4.6.c Cryptographic controls < D.4.6 Security and Privacy Protection < D.4 Operating Systems < D Software/Software Engineering,

D.2.4.d Formal methods < D.2.4 Software/Program Verification < D.2 Software Engineering < D Software/Software Engineering

Both the Internet and cryptography took roots in military technology. One of the first uses of computers, pioneered by Alan Turing, was to decrypt German war-time communications; and the precursor of the Internet, the Arpanet, was designed for resilience in case of nuclear war. It was the invention of public-key cryptography by Diffie and Hellman that created the impetus for open academic research into cryptography, and eventually led to the ubiquitous use of encryption on the Internet.

The Secure Socket Layer (SSL) protocol, one of the first real-world deployments of public-key cryptography, was originally developed by Netscape, an early Internet

browser vendor, to provide secure channels for electronic commerce. One of its main designers was Elgamal, a student of Hellman. As SSL took over the Web, it was renamed Transport Layer Security (TLS) and documented as an open standard by the Internet Engineering Task Force (IETF). Over time, it has undergone major changes; its implementations currently feature five versions—SSL2, SSL3, TLS 1.0, TLS 1.1, and TLS 1.2—while the next version of the protocol is actively being discussed at the IETF.

TLS implements a network socket API on top of a reliable but insecure network. It consists of two main protocols: a *handshake* that establishes sessions between clients and servers, relying on public-key cryptography to compute shared session keys; and a *record layer* that uses those keys to encrypt and authenticate their communications. SSL2 initially supported a single handshake scheme, based on RSA encryption, and a few record encryption algorithms, such as RC2 and DES. SSL3 added Diffie-Hellman schemes to the handshake, and further encryption algorithms, such as RC4 and 3DES. Over time, many of these cryptographic constructions came under attack, and were supplemented with stronger alternatives.

Since the client and the server may support different sets of cryptographic algorithms, the handshake lets them negotiate a combination of algorithms, called a ciphersuite. Hence, any TLS client and server can inter-operate as long as they have at least one ciphersuite in common. Over time, the number of ciphersuites supported by TLS implementations has grown steadily. For example, the popular OpenSSL library now supports over a hundred ciphersuites.

Not all ciphersuites are equally strong. Like most commercial software during the cold war, SSL was subject to US export regulations that classified cryptography as a weapon. To comply with these regulations, all protocol versions up to TLS 1.0 included deliberately weakened encryption algorithms for use in US software, such as web browsers, exported to foreign countries. Cryptographers and security practitioners started a rebellion, dubbed the Crypto Wars¹, against this weakening of their work, and

eventually prevailed, but SSL and TLS implementations were still forced to support export-grade ciphersuites for interoperability.

Many of the challenges in designing and deploying TLS securely were already apparent in the early days of the protocol. In particular, Bleichenbacher demonstrated a side-channel attack against the way RSA encryption was used in the SSL handshake, and Vaudenay discovered another side-channel attack on the way application data was encrypted in the record protocol.² Later versions of TLS continued to support those weak constructions, but mandated that implementations employ adequate countermeasures, triggering a series of increasingly sophisticated attacks and defences.

Besides cryptographic weaknesses, the SSL handshake protocol itself was shown to be vulnerable to logical flaws. The negotiation between strong and weak encryption had a protocol-level flaw in SSL2: a ciphersuite rollback (or downgrade) identified by Abadi³, enabling a network attacker to force a client and a server to use a weak export ciphersuite even though they both preferred a stronger ciphersuite. This flaw was fixed in SSL3, but a subsequent analysis revealed a more advanced downgrade attack⁴, enabling a network attacker to first force SSL3 clients and servers to use SSL2, and then exploit its known weaknesses. This was fixed by modifying the use of RSA encryption, which in turn enabled an improved Bleichenbacher-style side channel attack.

Hence, by the early 2000s, TLS was already caught in a cycle of attacks and fixes that continues to the current day. Formal foundations, to validate the protocol design and prevent any such attacks, became very attractive, and researchers from both the cryptographic and formal methods communities started applying various verification techniques to communications protocols.

Since the 1980s, cryptographers had been working on turning cryptography from an art into a science. The resulting theory is nowadays referred to as provable security. Conceived at the Theory of Computation group at MIT, it is concerned with reducing the difficulty of breaking cryptographic protocols to problems in complexity theory and mathematics. This approach resulted in ground-breaking works like those by the Turing award winners Shafi Goldwasser and Silvio Micali on probabilistic encryption and zero-knowledge proofs.

From the cryptographer's point of view, the TLS protocol is a combination of standard cryptographic constructions. Using compositional provable security techniques, one should be able to prove the security of each construction, and then put these proofs together to obtain a security theorem for TLS. In reality, composing proofs of various ad hoc parts of the protocol turned out to be hard, but over the last decade, cryptographers have successfully analyzed the security of many popular TLS ciphersuites. Their theorems confirm that, under some well-defined

implementation and mathematical assumptions, the cryptographic core of TLS is not vulnerable to attack.

From the programmer's point of view, protocols like TLS can be viewed as distributed processes that communicate across public channels and use cryptographic primitives as black boxes to protect their messages. The key analysis question is then whether the protocol, seen as a program, has logical flaws in its use of communications and cryptography, even if one assumes that the cryptographic building blocks are perfectly secure. For example, one may ask whether TLS admits ciphersuite or version downgrade attacks in the presence of an active network adversary.

The verification of concurrent and distributed processes has been investigated in a long line of research on programming language semantics, pioneered by other Turing award winners, Robin Milner and Tony Hoare. Their rigorous mathematical study of the meaning of programs allows us to formalize what it means for a program to keep a value secret, or for two programs to be equivalent. For simple cryptographic primitives, modelled as abstract mathematical functions, a message may hide a secret if it does not visibly depend on it; and two processes may be equivalent if they exchange similar-looking messages.

Most cryptographic algorithms hide secrets only computationally, meaning that given sufficient computational resources the secret can eventually be recovered. Their precise modelling requires complicated probabilistic definitions against restricted classes of adversaries. Instead, the semantics community proposed simpler symbolic approximations of cryptography to capture logical flaws, and developed tools to prove security (or find attacks) in their models. These techniques were used in a series of automated analyses of TLS, showing that the protocol is not vulnerable to logical attack, as long as the attacker is unable to break the cryptographic primitives.

Both the provable security and the symbolic verification of protocols were successful in their respective academic communities. As the former is more precise and the latter easier to automate, they are in principle complementary. However, the technical differences outlined above led to largely separate developments. As we will argue, this limited the impact that either of them had on the real-world security of TLS.

Theory vs Practice

Despite these theoretical successes, recent TLS versions have still been found vulnerable to practical attacks that rely on a combination of implementation bugs, cryptographic weaknesses, and protocol flaws. These attacks make it evident that the most advanced models of the provable security and verification communities still ignored many important implementation details. Such details include, for example, message formats, support for multiple protocol modes and algorithms for backward

compatibility, error handling exploitable as side channels, and signalling between the protocol and the application. Since these details affect the practical security of TLS, their omission limits the scope of theoretical statements.

It is worth reflecting on the cultural reasons for this gap between theory and practice. In their 2011 article on provable security, Degabriele, Paterson and Watson⁵ explain that a focus on principles can lead to simplistic or artificial models, and a neglect of implementation details. Interestingly, they notice a similar divide in the practical security community between specification writers and implementers. The former build in flexibility in specifications to allow for the competing interests of parties contributing to the development process. For example, specifications often avoid defining an API, and encourage implementations to accept a broad range of behaviours to support interoperability and backward compatibility. This flexibility can tempt cryptographers to interpret specifications in an overly abstract way that facilitates security analysis but misses real-world attacks that rely on implementation details.

Instead, we follow a model-attack-remodel cycle, informed by a dialog between practitioners and theoreticians. Concrete attack scenarios are invaluable for practice-oriented provable security: if they fall outside the security model, they encourage researchers to refine their model to better account for realistic threats. Conversely, model features that do not reflect any such scenario may point out simplifications.

Let us also mention a class of attacks often missed by practitioners and theoreticians alike. These attacks target the protocol design and evaluation process itself, sometimes directly, through the insertion of backdoors, or, more subtly, through influence on the culture in which designers operate. Juniper's VPN security hole is a recent example in this class. Besides awareness of the interests that some organizations may have in subverting Internet security, we believe that formal, open, practice-oriented protocol verification helps prevent such attacks.

A more technical challenge that prevents cryptographic analysis techniques to be applied to TLS deployments is that the protocol and its implementations have simply grown too complicated to be analysed by humans. Unsurprisingly, the highly-optimized C code of TLS implementations such as OpenSSL is amenable neither to cryptographic proofs nor to formal verification.

Cryptographic and symbolic models of TLS alike could not keep up with implementations and did not account for the details of the protocol as specified in the standard. Consequently, proofs of these models were likely to miss practical attacks on the protocol. Of course, they also missed attacks that exploited basic implementation flaws, such as incorrect certificate validation (GotoFail) or buffer overflows (HeartBleed).

In summary, we argue that the co-existence of proofs and attacks can be attributed to multiple gaps between verified models and real-world protocols:

- I. gaps between cryptographic models and standards;
- II. gaps between standards and implementations;
- III. gaps between individually secure ciphersuites and their insecure composition; and
- IV. gaps between APIs and application-level security.

In the rest of this article, we describe these gaps in more detail, and explain how we try to bridge them in the miTLS project.

miTLS: a verified reference implementation of TLS

By 2008, the theory and practice of TLS had largely diverged. To relate high-level specifications and low-level implementations (gaps I and II above), a group of researchers at the Microsoft Research-INRIA joint centre in Paris (including two authors of this paper) decided to build a reference implementation of the TLS 1.0 standard (RFC2246) in a style that enabled them to extract a formal model of the protocol directly from the code.⁶ By this approach, they ensured that the formal model was faithful to the standard and captured its low-level details. The model was then analyzed with a state-of-the-art protocol verifier, called ProVerif, to find both logical flaws in the protocol standard and implementation bugs in their code. Inasmuch as ProVerif did not find any flaws, they obtained high assurance in the security of their code against a large class of attacks.

This reference implementation, later dubbed miTLS (for Microsoft-INRIA TLS), was written in about 4000 lines of F# and the extracted symbolic models were among the largest to be automatically analyzed at the time, at the limits of verification technology. Symbolic tools like ProVerif are effective in automatically finding flaws without the need for any user intervention, but they do not necessarily scale well to large models. Verifying their TLS implementation for one protocol version and one cipher-suite took 3.5 hours and 4.5 GB of memory. Modelling other protocol modes was out of reach. Consequently, although they were able to find known attacks on early versions of SSL, they missed TLS renegotiation or Triple handshake attacks that were discovered later, because their models did not fully account for renegotiation.

As discussed above, a limitation of symbolic approaches is that they assume that the underlying cryptographic building blocks are perfect, and hence miss attacks. Some semi-automated tools, such as CryptoVerif, can analyze protocols in a more precise computational model of cryptography, but similarly do not scale up to large models. They applied CryptoVerif to core fragments of their TLS implementation, but were not able to analyze the full protocol using this tool. For example, they did not model

features like compression or the details of Cipher Block Chaining, and hence they missed subsequent vulnerabilities like BEAST and CRIME.

For the next version of miTLS⁷, we wanted to use a proof technique that could handle multiple versions and features of the protocol at the same time, and would rely on standard computational assumptions for the underlying cryptographic constructions. To this end, we switched to a verification method based on refinement types (to be explained shortly), originally designed for symbolic protocol analysis by Bhargavan, Fournet and Gordon,⁸ and then extended to modular computational proofs by Fournet, Kohlweiss, and Strub.⁹

Figure 1: miTLS verification architecture

Refinement types allow programmers to annotate each function with logical formulas. These annotations can capture program invariants, cryptographic assumptions, protocol events, and many security guarantees. To verify that a program meets its type annotations, the developer runs a type-checker that automatically verifies the program with the aid of an external SMT solver to discharge logical proof obligations. Crucially, type-checking is compositional, in the sense that each function can be independently verified, assuming that all previous functions also meet their type annotations. Consequently, the time for type-checking a large program is more-or-less linear in its size, and can be controlled by writing additional intermediate annotations.

The miTLS implementation currently supports TLS 1.0, 1.1, and 1.2, with multiple handshake and record modes. It also fully supports session resumption and renegotiation. The code is written in about 5000 lines of code, and is split into a sequence of modules, each of which equipped with a refinement-type interface. The verification approach is depicted in Figure 1. For modules containing protocol code, the interface represents the target security goals we wish to verify. For modules implementing cryptographic primitives, the interface represents the idealized functionality of the primitive, according to some standard cryptographic security assumption.

The top-level security guarantees for miTLS are stated in terms of a secure channel interface presented by TLS to

the application. This interface guarantees that application data sent on a connection between a miTLS client and a miTLS server is kept confidential, as long as the connection uses strong cryptographic algorithms and the long-term private keys of the two peers are unknown to the attacker. Moreover, the interface guarantees that the stream of application data received at one end is a prefix of the stream sent by the other. The security proof relies on type-checking each module, after applying a series of game-based transformations on the core cryptographic modules to replace the concrete algorithms by their ideal functionality. By this approach, we are able to verify the full miTLS implementation, module by module, under precise computational security assumptions. The total time for verification is under 20 minutes.

While it is valuable to have a security theorem for a reference implementation of TLS 1.2, the impact of miTLS is perhaps better evaluated in terms of the parts of the protocol design we were unable to prove, or where we had to make special cryptographic assumptions. These corner cases resulted in the discovery of weaknesses in the protocol and attacks on its real-world usage, discussed next.

Application Interface (API) and its Security Goals

Many problems stem from a mismatch between the security properties expected by applications using TLS and the actual guarantees provided by TLS (gap IV). The TLS standard does not specify an application interface (API) and so each implementation is free to implement its own. Application developers are expected to understand these APIs in detail and to use them in the right way to achieve their security goals. For example, some TLS libraries expect applications to validate the certificate presented by the server, and thus developers who wrongly assume that the library will do it for them become vulnerable to server-impersonation attacks. More generally, many attacks appear when building application-level authentication on top of TLS.

Consider an application that uses TLS to establish a secure channel where the client is initially unauthenticated. The application then runs an authentication protocol on top of TLS that allows the user to present a credential to the server. In this setting, the client expects that its use of TLS guarantees that the credential will only be presented at the target server; and a server that receives the credential over TLS may expect that the user intended to authenticate to it. However, as demonstrated by the attack outlined in Figure 2, these expectations are ill placed. We follow cryptographic tradition and refer to the client as Alice, the server as Bob, and the attacker as Eve. If Alice is willing to use the same credential (say, an X.509 certificate) with both Bob and Eve, then Eve can impersonate Alice at Bob, by forwarding Alice's credential (say, her signature over some authentication message) over his own

channel with Bob. Even if Alice is careful and uses her credential only with Bob, a sophisticated attacker may impersonate Bob to mount a Man-in-the-Middle attack (MitM), by operating a phishing web-site, obtaining mis-issued certificates, or compromising the server key. Such credential forwarding attacks can only be prevented if Alice not only authenticates herself, but also her channel, e.g., by signing a unique identifier extracted from the TLS connection. Then, if Bob compares these channel identifiers he can detect the attack.

Credential forwarding attacks and their countermeasures have appeared multiple times in TLS applications. They were first discussed in the context of tunnelled compound authentication protocols for network access. They then reappeared in the context of user-authenticated TLS renegotiation as commonly used on the web. In response to these attacks, a variety of channel identifiers were defined for TLS and exposed within the APIs of various implementations. Compound authentication protocols used the TLS session key (called master secret) as an identifier for binding application-level credentials. TLS renegotiation used the protocol transcript of the previous handshake as a connection identifier.

We implemented these countermeasures in miTLS and tried to prove that applications using miTLS are not vulnerable to credential forwarding, but we failed. Instead, we discovered several counterexamples. A malicious server is able to synchronize the TLS session keys on two different connections, one from Alice and one to Bob, so that the channel identifiers on both connections are the same, hence defeating the compound authentication countermeasure. Then by running a second TLS connection that uses session resumption, the server can also synchronize the protocol transcripts on these connections, hence defeating the TLS renegotiation countermeasure. In fact, such channel synchronization attacks break all known credential forwarding protections over TLS by exploiting a misunderstanding of the TLS API; the protocol does not guarantee unique channel identifiers.

This class of channel synchronization attacks was called the Triple Handshake¹⁰, since it requires a sequence of up to three runs of TLS before the attack succeeds against TLS client authentication. Although it has been present in the TLS protocol since SSL3, it escaped previous analyses because they did not consider sequences of TLS connections, and they did not model credential forwarding as a threat. In response to these attacks, we helped the TLS working

group to standardize a new protocol-level fix called the extended master secret that systematically protects all compound authentication protocols.

Implementing Negotiation

In addition to designing an API, a second major challenge for a TLS implementation is that it needs to handle a variety of protocol versions, extensions, authentication modes, and ciphersuites at the same time. While the TLS standard describes each mode in isolation, it does not always specify how an implementation should compose them (gap III). In particular, the protocol state machine is left unspecified and each implementation can design its own. In miTLS, we define and verify our own state machine. Our type-based proofs rely on careful invariants that require that the current protocol state is consistent with the desired protocol mode, and that the transcripts and signature formats for different modes are disjoint. Considering the effort that was required to prove our own state machine correct, we then tested other implementations to see if they implemented the TLS standard correctly, and to our surprise, many of them failed this test, resulting in subtle attacks.¹¹

Some implementations failed to correctly implement the composition of the handshake and record protocols and allowed application data to be sent unencrypted, before the handshake was complete. Other implementations failed to correctly compose regular RSA ciphersuites with export RSA ciphersuites, allowing a downgrade attack, called FREAK, whereby a MitM attacker could fool a TLS client into accepting export-grade 512-bit RSA keys even though it wanted to use regular RSA. In all, by testing other open-source TLS libraries against miTLS, we found dozens of state machine bugs across all major TLS implementations, including four that could be exploited for real-world attacks.

Another attack on TLS negotiation, found by a large group of researchers including one of the current authors, relies on a protocol flaw rather than an implementation bug.¹² In Logjam, the server supports both regular Diffie-Hellman (DH) groups as well as export-grade 512-bit DH groups. The client does not support export-grade DH, but it allows the server to pick the group. As depicted in Figure 3, this situation leads to an MitM attack. The attacker tampers with the protocol messages to fool the server into thinking that the client only supports export-grade DH. So, the server sends the export-grade group to the client who thinks this is the server's regular group and accepts it. This kind of attacks is sometimes called a cross-protocol attack since it involves confusions between two different protocols (DH and export DH). It is enabled by a protocol flaw in TLS: the server's signature format for export DH ciphersuites is indistinguishable from its signature for regular DH. Hence, the attacker can successfully downgrade the connection to use export DH even though the client does not realize it. To complete the attack, the attacker still needs to solve the discrete log problem for the export DH group, which is well within reach of modern processing power.

Towards TLS 1.3

At the time when SSL was first designed, there was a real enthusiasm and sense of purpose to deploy practical cryptographic protocols. Often one and the same person

worked on and understood both the cryptography and the implementation—to the extent possible at the time. Since then, advances in cryptographic theory and analysis have greatly improved our understanding of when protocols achieve their security goals and when they fail to do so. However, typically this analysis is performed either on toy protocols, or in retrospect on partial aspects of a protocol specification. Implementers still primarily follow a fix-attempt-fix cycle. This cycle, however, only gets worse as protocols grow in complexity. We believe that the only way

forward is through an active collaboration between theoreticians and practitioners.

A promising development in this direction is the standardization effort behind the upcoming TLS 1.3 protocol, which fixes many weaknesses in TLS 1.2 and, at the same time, promises improved performance. From the early stages of its design, the TLS working group has invited and encouraged the participation of academic researchers, who have responded with significant numbers. Not only was the design of the cryptographic core of TLS 1.3 strongly influenced by the OPTLS protocol by Krawczyk and Wee¹³, but we now have multiple published security proofs for different draft versions of the protocol even before it has been standardized. Such careful cryptographic analysis for a new standard is unprecedented at the IETF. As a result of this process, many attacks and weaknesses were detected and removed from early drafts, resulting in a simpler and more secure protocol.

Our main contribution to the standardization effort is a new version of miTLS that implements TLS 1.3, but also supports older versions for backwards compatibility. Since mainstream TLS implementations will continue to support such older versions for the foreseeable future, we were especially concerned with the potential for version downgrade attacks that might nullify the security advantages of TLS 1.3.

TLS 1.3 signs all exchanged messages to prevent MitM attacks like Logjam that rely on tampering with handshake messages for downgrade attacks. However, we discovered that by downgrading the protocol version to TLS 1.2, the attacker can force the server to use the weaker TLS 1.2 signature that does not cover all messages, hence re-enabling such tampering attacks. The problem is that, in older versions of TLS, clients cannot verify the maximum supported server version until the end of the protocol, by when it is too late.

That this downgrade attack went unnoticed until Draft 10 of TLS 1.3 is an example for the many intricacies and pitfalls of practical protocol security. Once detected and brought to the attention of the IETF, we helped develop a verified countermeasure depicted in Figure 4,¹⁴ that is peculiar but simple: shorten the server nonce, which is signed in TLS 1.2, and use some of its bytes to encode the server's highest supported version number.

The future of verified implementations

The miTLS approach necessarily involves multi-disciplinary teams of cryptography, programming semantics, tooling and verification experts, as well as generalists knowledgeable of real-world security concerns and system performance. We require implementations to be written in a programming language with a well-defined formal semantics so that protocol properties devised by theorists can be verified using sound automated tools on code co-developed with practitioners.

Even verified implementations have to rely on cryptographic assumptions, the accuracy of the security model, and the correctness of proofs and verification tools. To ensure that our modelling assumptions do not miss concrete attacks, we advocate a comprehensive penetration testing regime that uses the miTLS codebase to find and implement attacks on miTLS and other TLS implementations. Such attacks can be on cryptographic primitives, on the TLS protocol level, but also on the HTTPS ecosystem and even against the soundness of our verification tools. Our goal is to use a combination of verification and testing to span and evaluate all four of these levels in order to reduce the trusted computing base for TLS applications.

Verification alone is not enough to ensure that a TLS implementation will be widely used. Real-world implementations have to be performant. A key challenge for future work on miTLS is to extend our verification techniques so that they can handle the programming idioms used in high-performance code. As our tools improve, we anticipate that the feature and performance gap between verified and unverified protocol implementations will vanish.

Karthikeyan Bhargavan karthikeyan.bhargavan@inria.fr,

Karthikeyan Bhargavan is a researcher at INRIA, the French national lab for computer science. He is based in Paris where he leads a team called Prosecco ("programming securely with cryptography") and is the principal investigator of an ERC consolidator grant CIRCUS on provably secure implementations of cryptographic web applications. Karthik was trained at IIT New Delhi and the University of Pennsylvania. Before coming to Paris in 2009, he worked as a researcher at Microsoft Research lab in Cambridge, England. His publications and CV are available from <http://prosecco.inria.fr/personal/karthik>

Cédric Fournet fournet@microsoft.com,

Cédric Fournet leads the Constructive Security group at the Microsoft Research lab in Cambridge, UK. He is interested in security, privacy, cryptography, programming languages, and formal verification. He is currently working on a verified TLS/HTTPS protocol stack and techniques for outsourcing computations with strong security and privacy guarantees. Cédric graduated from Ecole Polytechnique and Ecole Nationale des Ponts et Chaussées, and completed a PhD at INRIA in France. See also <https://www.microsoft.com/en-us/research/people/fournet>.

Markulf Kohlweiss markulf@microsoft.com

Markulf Kohlweiss is a researcher at Microsoft Research Cambridge in the Programming Principles and Tools group. He did his PhD at the COSIC (Computer Security and Industrial Cryptography) group at the K.U. Leuven, and his master thesis at IBM Research Zurich. Dr. Kohlweiss' research focus is on privacy-enhancing cryptography and formal reasoning about cryptographic protocols. See also <https://www.microsoft.com/en-us/research/people/markulf>.

1 https://en.wikipedia.org/wiki/Crypto_Wars

2. Serge Vaudenay. 2002. Security Flaws Induced by CBC Padding - Applications to SSL, IPSEC, WTLS (Serge Vaudenay), In EUROCRYPT 2002

3 Prudent Engineering Practice for Cryptographic Protocols (M. Abadi and R. Needham), In IEEE Transactions on Software Engineering, 22(1):2–15, January 1996.

4 Analysis of the SSL 3.0 protocol (David Wagner and Bruce Schneier). In USENIX Workshop on Electronic Commerce Proceedings, 1996: 29-40

5 J. P. Degabriele, K. Paterson and G. Watson, "Provable Security in the Real World," in IEEE Security & Privacy, vol. 9, no. 3, pp. 33-41, May-June 2011.

6 Cryptographically verified implementations for TLS (Karthikeyan Bhargavan, Cédric Fournet, Ricardo Corin,

Eugen Zalinescu), In ACM Conference on Computer and Communications Security, 2008

7 Implementing TLS with Verified Cryptographic Security (Karthikeyan Bhargavan, Cédric Fournet, Markulf Kohlweiss, Alfredo Pironti, Pierre-Yves Strub), In IEEE Symposium on Security & Privacy (Oakland), 2013.

8 Modular verification of security protocol code by typing (Karthikeyan Bhargavan, Cédric Fournet, Andrew D. Gordon). POPL 2010: 445-456

9 Modular code-based cryptographic verification (Cédric Fournet, Markulf Kohlweiss, Pierre-Yves Strub), In ACM Conference on Computer and Communications Security 2011: 341-350

10 Triple Handshakes and Cookie Cutters: Breaking and Fixing Authentication over TLS (Karthikeyan Bhargavan, Antoine Delignat-Lavaud, Cédric Fournet, Alfredo Pironti, and Pierre-Yves Strub). In IEEE Symposium on Security and Privacy (Oakland), 2014

11 A Messy State of the Union: Taming the Composite State Machines of TLS (Benjamin Beurdouche,

Karthikeyan Bhargavan, Antoine Delignat-Lavaud, Cédric Fournet, Markulf Kohlweiss, Alfredo Pironti, Pierre-Yves Strub, Jean Karim Zinzindohoue), In IEEE Symposium on Security & Privacy 2015 (Oakland), 2015

12 Imperfect Forward Secrecy: How Diffie-Hellman Fails in Practice (David Adrian, Karthikeyan Bhargavan, Zakir Durumeric, Pierrick Gaudry, Matthew Green, J. Alex Halderman, Nadia Heninger, Drew Springall, Emmanuel Thomé, Luke Valenta, Benjamin VanderSloot, Eric Wustrow, Santiago Zanella-Béguelin, and Paul Zimmermann), In ACM Conference on Computer and Communications Security (CCS'15), 2015

13 The OPTLS Protocol and TLS 1.3, Hugo Krawczyk Hockett Wee, October 9, 2015, Cryptology ePrint Archive, eprint.iacr.org/2015/978.pdf

14 Downgrade Resilience in Key-Exchange Protocols (Karthikeyan Bhargavan, Christina Brzuska, Cédric Fournet, Matthew Green, and Markulf Kohlweiss and Santiago Zanella-Béguelin) In IEEE Symposium on Security and Privacy 2016 (Oakland), 2016.