

Capsule: A Protocol for Secure Collaborative Document Editing

Nadim Kobeissi

INRIA, Symbolic Software
nadim@symbolic.software

Abstract. Today’s global society strongly relies on collaborative document editing, which plays an increasingly large role in sensitive workflows. While other collaborative venues, such as secure messaging, have seen secure protocols being standardized and widely implemented, the same cannot be said for collaborative document editing. Popular tools such as Google Docs, Microsoft Office365 and Etherpad are used to collaboratively write reports and other documents which are frequently sensitive and confidential, in spite of the server having the ability to read and modify text undetected.

Capsule is the first formalized and formally verified protocol standard that addresses secure collaborative document editing. Capsule provides confidentiality and integrity on encrypted document data, while also guaranteeing the ephemeral identity of collaborators and preventing the server from adding new collaborators to the document. Capsule also, to an extent, prevents the server from serving different versions of the document being collaborated on.

In this paper, we provide a full protocol description of Capsule. We also provide formal verification results on the Capsule protocol in the symbolic model. Finally, we present a full software implementation of Capsule, which includes a novel formally verified signing primitive implementation.

Keywords: collaborative document editing, secure document editing protocol, formal verification, proverif, privacy enhancing technology

1 Introduction

Collaborative document editing software such as Google Docs and Microsoft Office365 has become indispensable in today’s work environment. In spite of no confidentiality guarantees, large working groups still find themselves depending on these services even for drafting sensitive documents. Peer pressure, a rush to agree on a working solution and the sheer lack of alternatives have created an ecosystem where a majority of sensitive drafting is currently conducted with no end-to-end encryption guarantees whatsoever. Google and Microsoft servers have full access to all documents being collaborated upon and so do any parties with read access to the internal databases used by these services.

Capsule aims to solve this problem by being the first formally specified and formally verified standard protocol for secure collaborative document editing. Capsule comes with security goals that are validated in the symbolic model using the ProVerif [1] automated protocol verifier. It is also presented with a full client and server implementation, published as open source software.

We note that Capsule does not target the much simpler problem of encrypting cloud documents that are at rest, such as files stored on Google Drive or Microsoft OneDrive. Rather, the Capsule Protocol targets files being actively, collaboratively edited in real time.

Capsule’s more involved use case scenario is responsible for providing a class of security guarantees vaguely resembling that of group secure messaging but without a need for forward secrecy. Another element that Capsule has to deal with is allowing the management of a document’s incremental history with a server that cannot read the document’s contents. This is accomplished by storing the document as *an authenticated hash chain of encrypted, signed diffs*: when a new participant joins a Capsule document, they pull the entire document hash chain, decrypting and parsing every diff until the entire document is reconstructed. Once within the document, each participant pushes and pulls encrypted diff blocks into this hash chain. Utilizing a hash chain is not only more efficient than more naïve approaches, but also makes it more difficult for the server to provide differing document histories to participants, without forking the hash chain entirely and with little payoff.

Capsule uses symmetric encryption to guarantee document confidentiality and integrity. Cryptographic signatures are used for authentication, although we employ only *ephemeral authentication* since identities are valid for the lifetime of the document. Symmetric primitives are also used for generating a type of proof that disallows the server from adding non-existent participants to the document, in spite of these fake participants already being unable to access any document plaintext.

1.1 Security Goals

Capsule aims to guarantee the following security goals. In the following, a *valid participant* is any entity with access to the Capsule collaborative document’s shared master secret. This master secret is generated upon document creation and is shared manually by the document creator in order to allow access to the document.

- **Participant List Integrity.** Only participants with access to the Capsule collaborative document’s shared master secret may appear as valid entries on the participant list. Illegitimate entries injected by the server or any other parties must be detectable by valid participants.
- **Confidentiality.** Any changes made to a collaborative document may only be viewed by valid participants.
- **Integrity.** Encrypted diffs that are appended to the Capsule document’s hash chain by a particular author cannot be tampered with by any other party.

- **Ephemeral Authentication.** The identity keys which are generated for the lifetime of the Capsule document session must successfully identify the participant that owns them and disallow impersonation. Public identity keys must benefit from Trust on First Use logic and be verifiable out-of-band.
- **Transcript Consistency.** A malicious server cannot selectively omit changes to the collaborative document short of entirely forking the hash chain. Furthermore, in no scenario may a malicious server maliciously *append* changes to the collaborative document.

1.2 Threat Model

Capsule assumes the following threat model, which is fairly standard for protocols aiming to provide end-to-end security:

- **Untrusted Network.** We assume that an attacker controls the network and so can intercept, tamper with and inject network messages.
- **Malicious Server.** We assume that a Capsule server may be potentially interested in misleading users with regards to the document’s history and in the apparent list of identities participating in a collaborative document editing session.

1.3 Related Work

To the best of our knowledge, the only prior existing work regarding collaborative document encryption is CryptPad [2], an open source web client. CryptPad shares similarities with Capsule especially in that both use a hash chain of encrypted diffs in order to manage document collaboration and to reconstruct the document. However, CryptPad adopts a more relaxed threat model of an “honest but curious cloud server” and does not appear to guard against a server interfering with the document’s list of participants or its history. Meanwhile, Capsule explicitly guards against a server injecting false participants by requiring a certain proof from all participants. CryptPad’s software implementation is also limited within a web browser and unlike Capsule’s, does not employ formally verified cryptographic primitives.

2 Primitives

In designing the Capsule protocol, we wanted to focus on obtaining the smallest attack surface possible on an architectural cryptographic level. This is why we only use a single hash function, BLAKE2s [3] for all symmetric operations: hashing, key derivation, message authentication and encryption. Similarly, we use only one primitive, Ed25519 [4] for cryptographic signatures over an elliptic curve field. Capsule does not require any Diffie-Hellman type primitive.

We expose the following cryptographic functions:

- **Hashing.** $\text{BLAKE2}(x) \rightarrow y$.
- **Hash-Based Key Derivation.** $\text{BKDF}(k, \text{salt}) \rightarrow (k_0, k_1, k_2, k_3)$. BLAKE2s is also used for key derivation, producing four 16-byte outputs.
- **Encryption and Decryption.** BLAKE2 is used for encryption and decryption. We hash the encryption key over a counter and a nonce to generate a keystream. BLAKE2 is then used as a keyed hash with a MAC key to generate an HMAC value over the ciphertext.
 - **Encryption.** $\text{BENC}(ek, mk, p) \rightarrow (\text{ENC}_p, n)$.
 - **Decryption.** $\text{BDEC}(ek, mk, \text{ENC}_p, n) \rightarrow p$.
- **Signatures.** Ed25519 is chosen due to its speed and minimal input validation requirements.
 - **Key Generation.** $\text{EDGEN}(sk) \rightarrow pk$.
 - **Signing.** $\text{EDSIGN}(sk, x) \rightarrow \text{SIG}_x$.
 - **Signature Verification.** $\text{EDVERIF}(pk, x, sig_x) \rightarrow \{\top, \perp\}$.

3 Protocol Description

During the lifetime of a Capsule collaborative document, there are only two subprotocols that a participant has to follow. The first is the key generation subprotocol which produces the necessary key material to conduct the session. The second is the hash chain [5] protocol, which we call *DiffChain*. By pushing and pulling encrypted diffs to the diff chain, participants can reconstruct the document upon joining the session and then begin exchanging modifications.

3.1 Key Material

The following key material is necessary for all participants.

Client Key Materials A client A owns the following key material in relationship to collaborative document V :

- $V_k \xleftarrow{R} \{0, 1\}^{128}$, a randomly generated master secret, selected by the document creator or otherwise obtained from the document creator. Acts as the lifetime token for legitimately joining a collaborative document.
- $(V_{ek}, V_{mk}, V_{sp}, V_{id}) \leftarrow \text{BKDF}(V_k, \text{CAPSULECORP})$. A set of symmetric subkeys used for encryption, message authentication, proof of being a legitimate participant and to derive the document ID used by the server for bookkeeping.
- $AV_{sk} \xleftarrow{R} \{0, 1\}^{256}$, a signing private key.
- $(AV_{sk}, AV_{pk}) \leftarrow \text{EDGEN}(AV_{sk})$.
- $AV_{pv} \leftarrow \text{BLAKE2MAC}(V_{sp}, A || V_{id} || AV_{pk})$

Among the above values, AV_{pv} is called a *proof of participation value*: Since the server does not know V_{sp} , it cannot generate a xV_{pv} for any possible participant x . Validating this value, therefore, protects the collaborative document editing session from containing illegitimate participants who were not given access to V_k :

Furthermore, we make the following observation: if all participants accept ciphertext signed only by identities with a valid xV_{pv} value for any participant x , then the authentication component of the BENC becomes redundant.

Proof. Assuming that BLAKE2MAC is a one-way function, that V_k is secret and that AV_{pk} is authenticated as a public signing identity belonging to A :

- If value AV_{pv} is validated, this proves that generator A must possess V_{sp} .
- Given that V_k is secret, ownership of V_{sp} in turn proves that A must possess V_k .
- V_k maps into $(V_{ek}, V_{mk}, V_{sp}, V_{id})$. Therefore, proving ownership of V_{sp} under a secret V_k is equivalent to proving ownership of $(V_{ek}, V_{mk}, V_{sp}, V_{id})$. This already satisfies the original purpose of validating the proof of participation value. \square
- A cannot produce AV_{pv} without also being able to produce (V_{ek}, V_{mk}) . Since signature verification disallows tampering, any unauthenticated ciphertext encrypted under V_{ek} and signed under A is equivalent to a ciphertext encrypted under V_{ek} , authenticated using V_{mk} and signed under A . \square

This second observation allows potentially replacing BENC with a unauthenticated encryption cipher, which could lead to improved performance.

Server Key Materials A server S obtains access to the following key materials in relationship to collaborative document V :

- V_{id} , used as the server-side identifier for the document.
- AV_{pk} as Alice’s ephemeral identity. Optionally, the server can also use this to validate DiffChain blocks sent in by Alice for commitment.
- AV_{pv} , used by Alice to prove that her identity is being served as a legitimate participant to the collaborative document.

3.2 Session Setup

Suppose Alice (A) wants to create a new collaborative document. She generates the keys mentioned in §3.1 and then communicates the following key material to server S :

$$A \longrightarrow (V_{id}, AV_{pk}, AV_{pv}) \longrightarrow S$$

The server creates the new document identified server-side as V_{id} and stores the triple (A, AV_{pk}, AV_{pv}) as the first participant to the document.

Session setup is now complete. In order to invite a fellow collaborator into the document, Alice simply shares V_k . Alice can choose to encode V_k as a string or QR code to make sharing easier. Once Bob obtains V_k , he too can immediately generate all necessary key material to join the document, encrypt and authenticate diff information and prove to the rest of the participants that his participation is legitimate.

When Bob (B) joins the document, the following occurs (the last message is repeated for every existing participant:)

$$\begin{aligned} B &\longrightarrow (V_{id}, BV_{pk}, BV_{pv}) \longrightarrow S \\ B &\longleftarrow (V_{id}, AV_{pk}, AV_{pv}) \longleftarrow S \end{aligned}$$

3.3 Managing Collaborative Document History with DiffChain

As a protocol, Capsule needs to provide a fast, efficient method to allow for the continuous update of a document by an unbounded number of participants and for the constant synchronization of this document between the participants. Aside from homomorphic encryption, which currently does not appear to be ready for use in this kind of real-world system, the other clear potential solution seemed to be an encrypted, append-only authenticated log of diff information. We construct such a data structure and call it DiffChain.

Pulling a DiffChain to reconstruct a collaborative document upon joining it is very similar to pulling the Bitcoin blockchain and parsing it in order to reconstruct the currency's current value and activity by going through all transactions since the original block. A DiffChain block contains the following data:

- **Encrypted Diff.** Alice generates this encrypted diff using $\text{BENC}(V_{ek}, V_{mk}, \text{diff})$.
- **Hash of Previous Block.** This is a standard element in hash chain constructions.
- **ID of Current Block.** This is conventionally a UNIX timestamp. It is appended by the server.
- **ID of Previous Block.** This allows for faster lookups.
- **Signature.** All of the above fields are included in a single signature, with the exception of the ID of the current block since it is appended by the server.

Responsible pushing and pulling. As a rule, clients should not push any new blocks before obtaining confirmation from the server that their local diff state is current. This will help avoid merge conflicts.

Diffing algorithm. The Capsule standard library comes equipped with an efficient diff generation algorithm that features a JSON-compatible syntax for compatibility. Independent implementations are free to adapt their own diff payload representations, as this does not strictly affect the security or operation of the cryptographic protocol itself.

4 Symbolic Verification with ProVerif

We describe Capsule using the applied-pi calculus and, using the ProVerif symbolic protocol verifier, verify that the Capsule protocol meets its security goals under its proposed threat model.

In the symbolic model, cryptographic primitives are logically perfect black boxes: hash functions are one-way mappings, encryption functions are permutations, random values are unique and perfectly indistinguishable. Processes are defined and can then be executed as a parallel sequence on an unbounded network. In this setting, a Dolev-Yao [6] active attacker exists and attempts to reconstruct inputs and outputs in such a way as to reach goals defined by queries and events.

4.1 Capsule Processes in ProVerif

In ProVerif, we describe a single public channel, `pub`, intending to represent regular Internet network exchanges. Then, two types of client processes are described. We illustrate them here in an abbreviated format.¹

The top-level process consists of an unbounded parallel executions of the two processes below in the context of an active attacker. Three identities are used: Alice, Bob, Carol and a compromised identity Mallory for which the secret identity key is leaked on the network. The symmetric master secret V_k is never leaked.

Writer Client Process The writer process running under identity x joins a document V using a pre-shared V_k and writes a block on the DiffChain.

$$\begin{aligned} \text{WRITERCLIENT}(V_k, x, x_{sk}, \text{diff}) = & \\ & x_{pk} \leftarrow \text{EDGEN}(x_{sk}) \\ & (V_{ek}, V_{mk}, V_{sp}, V_{id}) \leftarrow \text{BKDF}(V_k) \\ & x_{vp} \leftarrow \text{BLAKE2MAC}(V_{sp}, (x \| V_{id} \| x_{pk})) \\ & \text{event ProofSent}(x, x_{vp}) \\ & \text{out}(\text{pub}, (V_{id}, x, x_{pk}, x_{vp})) \\ & \text{ENC}_{\text{diff}} \leftarrow \text{BENC}(V_{ek}, V_{mk}, \text{diff}) \\ & \text{SIG}_{\text{ENC}_{\text{diff}}} \leftarrow \text{EDSIGN}(x_{pk}, (x \| \text{ENC}_{\text{diff}} \| V_{id})) \\ & \text{event Pushed}(V_{id}, x, x_{pk}, \text{ENC}_{\text{diff}}, \text{SIG}_{\text{ENC}_{\text{diff}}}) \\ & \text{out}(\text{pub}, (V_{id}, x, x_{pk}, \text{ENC}_{\text{diff}}, \text{SIG}_{\text{ENC}_{\text{diff}}})) \end{aligned}$$

Reader Client Process The reader process running under participant identity z comes equipped with x_{fp} value for each participant x whose public identity

¹ Capsule models and reference implementations will be available at <https://symbolic.software/capsule/>.

keys have been authenticated out of band, using the generated fingerprint:

$$x_{fp} \leftarrow \text{BLAKE2MAC}(x_{pk}, x)$$

$$\begin{aligned} \text{READERCLIENT}(V_k, z, x_{fp}) = & \\ & (V_{ek}, V_{mk}, V_{sp}, V_{id}) \leftarrow \text{BKDF}(V_k) \\ & \text{in}(\text{pub}, (=V_{id}, x, x_{pk}, x_{vp})) \\ & x_{fp}^c \leftarrow \text{BLAKE2MAC}(x_{pk}, x) \\ & x_{vp}^c \leftarrow \text{BLAKE2MAC}(V_{sp}, (x \| V_{id} \| x_{pk})) \\ \text{if } ((x_{fp}^c = x_{fp}) \wedge (x_{vp}^c = x_{vp})) \text{ then } (& \\ & \text{event ProofVerified}(x, z, x_{vp}) \\ & \text{in}(\text{pub}, =V_{id}, =x, =x_{pk}, \text{ENC}_{\text{diff}}, \text{ENC}_{\text{SIG}_{\text{diff}}}) \\ \text{if } (\text{EDVERIF}(x_{pk}, \text{ENC}_{\text{diff}}, \text{SIG}_{\text{ENC}_{\text{diff}}})) \text{ then } (& \\ & \text{diff} \leftarrow \text{BDEC}(V_{ek}, V_{mk}, \text{ENC}_{\text{diff}}) \\ & \text{event Pulled}(V_{id}, x, x_{pk}, \text{ENC}_{\text{diff}}, \text{SIG}_{\text{ENC}_{\text{diff}}})) \end{aligned}$$

4.2 Security Goals in the Symbolic Model

We apply ourselves on the security goals described in §1.1 in order to model and verify participant list integrity, confidentiality, integrity and ephemeral authentication in ProVerif.²

Participant List Integrity We assert that if a proof of participation value x_{vp} is verified by participant y , then there must exist an identity x which has issued x_{vp} . As shown in §3.1, this means that whoever asserts identity x must also possess V_{sp} :

$$\text{event ProofVerified}(x, y, x_{vp}) \implies \text{event ProofIssued}(x, x_{vp})$$

Confidentiality We simply query for attacker access to `diff`:

$$\text{query attacker}(\text{diff})$$

Integrity and Ephemeral Authentication We assert that if a DiffChain block $(\text{ENC}_{\text{diff}}, \text{SIG}_{\text{ENC}_{\text{diff}}})$ was received by y as part of document V , then this block will successfully authenticate and decrypt as originating from identity x if and only if it was pushed by x for document V .

$$\begin{aligned} \text{event Pulled}(V_{id}, x, y, x_{pk}, \text{ENC}_{\text{diff}}, \text{SIG}_{\text{ENC}_{\text{diff}}}) \implies \\ \text{event Pushed}(V_{id}, x, x_{pk}, \text{ENC}_{\text{diff}}, \text{SIG}_{\text{ENC}_{\text{diff}}}) \end{aligned}$$

² Transcript consistency, while not explicitly modeled in ProVerif, is obtained through the hash chain structure of DiffChains, as described in §3.3.

4.3 Results Under a Dolev-Yao Attacker

All queries mentioned above complete successfully. By tweaking the model, we are also able to quickly test for the event where an ephemeral identity is not verified out-of-band by the other participants. Since all encryption is symmetric and no asymmetric key agreement ever occurs in the Capsule protocol, the outcome of this is minor and does not result in the compromise of confidential information.

5 Software Implementation

Capsule is provided with a client/server reference implementation called Capsulib. It is meant to allow for quick deployment and testing of the Capsule protocol in production.

5.1 Capsulib: a Capsule Client/Server Implementation

Capsulib Server is a Node.js application that uses Redis as a database backend but is otherwise self-contained. It achieves low latency by communicating over WebSockets. Its main purpose is to hold a DiffChain record for every document and to facilitate the exchange of encrypted blocks.

Capsulib Client is meant to be executed within an Electron runtime. It provides a user interface, the Capsulib cryptographic library and the full Capsule client protocol implementation.

5.2 Formally Verified Cryptographic Primitives in WebAssembly

While Capsulib is written in JavaScript, much of the Capsulib cryptographic primitives library is automatically generated as WebAssembly (WASM) [7] code. We use a toolchain that provides strong verification guarantees, hence ruling out potentially catastrophic bugs in the cryptographic layer.

Algorithms such as Ed25519 are specified in F* [8], an ML-like programming language with support for program verification via the use of a dependent type system, effect annotations and SMT-based automation. Our specifications are executable, which allows us to run them against the RFC test vectors to ensure their correctness.

The cryptographic algorithms themselves are written in Low*, a subset of F* that enjoys an optimized compilation scheme to low-level targets. The Low* implementation of our algorithms is shown to match the original F* specification, which ensures functional correctness and memory safety.

Based on this, we rule out both incorrect math and memory errors such as buffer overflows. In addition to function correctness and memory safety, we enforce basic side-channel resistance, by restricting secrets to a very limited set of operations. In effect, this ensures that secret-manipulating code is branch-free and cannot access memory using a secret index. This rules out, respectively, classic timing and cache attacks.

Once verified, Low* programs can be compiled to low-level targets using the Kremlin [9] compiler, which currently supports C and WASM as backends. Compared to Emscripten, the Kremlin compiler offers a much smaller trusted computing base, smaller resulting WASM files and almost no dependencies on untrusted code, such as a `libc` implementation.

The cryptographic primitives that we compile into WebAssembly using this process are originally part of the HACLS* [10] cryptographic library, specified in F*.

6 Conclusion

In this paper, we have presented Capsule, the first formally verified protocol for secure collaborative document editing. We believe Capsule will address an important need in the space of privacy enhancing technologies. We have provided a full description of the Capsule protocol and symbolic verification results for its security goals under the specified threat model.

Finally, we also provide Capsulib, a reference implementation of Capsule based on modern web technologies. While Capsulib is built for web use, it is also the first software project to employ the HACLS* formally verified cryptographic library by translating it into WASM. By adopting such technologies, Capsule protocol is able to be presented complete with a practical implementation that follows best practices and achieves sound practical security in deployment.

Acknowledgments

We would like to thank Karthikeyan Bhargavan, Denis Merigoux and Jonathan Protzenko for various feedback and assistance. We also thank Vitalik Buterin for encouragement to pursue this work.

References

1. Vincent Cheval and Bruno Blanchet. Proving more observational equivalences with ProVerif. In *International Conference on Principles of Security and Trust*, pages 226–246. Springer, 2013. 2
2. CryptPad: Zero Knowledge, Collaborative Real Time Editing, 2016. <https://cryptpad.fr/what-is-cryptpad.html>. 3
3. Jean-Philippe Aumasson, Samuel Neves, Zooko Wilcox-O’Hearn, and Christian Winnerlein. BLAKE2: simpler, smaller, fast as MD5. In *International Conference on Applied Cryptography and Network Security*, pages 119–135. Springer, 2013. 3
4. Daniel J Bernstein, Niels Duif, Tanja Lange, Peter Schwabe, and Bo-Yin Yang. High-speed high-security signatures. *Journal of Cryptographic Engineering*, 2(2):77–89, 2012. 3
5. Leslie Lamport. Password authentication with insecure communication. *Communications of the ACM*, 24(11):770–772, 1981. 4
6. Paul Syverson, Catherine Meadows, and Iliano Cervesato. Dolev-Yao is no better than machiavelli. Technical report, Naval Research Lab, Washington DC Center for High Assurance Computing Systems, 2000. 7
7. WebAssembly Core Specification, 2018. <https://webassembly.github.io/spec/core/bikeshed/index.html>. 9
8. Jonathan Protzenko, Jean-Karim Zinzindohoué, Aseem Rastogi, Tahina Ramananandro, Peng Wang, Santiago Zanella-Béguelin, Antoine Delignat-Lavaud, Cătălin Hrițcu, Karthikeyan Bhargavan, Cédric Fournet, et al. Verified low-level programming embedded in F. *Proceedings of the ACM on Programming Languages*, 1(ICFP):17, 2017. 9
9. Peng Wang, Karthikeyan Bhargavan, Jean-Karim Zinzindohoué, Abhishek Anand, Cédric Fournet, Bryan Parno, Jonathan Protzenko, Aseem Rastogi, and Nikhil Swamy. Extracting from F* to C: a progress report. 10
10. Jean-Karim Zinzindohoué, Karthikeyan Bhargavan, Jonathan Protzenko, and Benjamin Beurdouche. HACL*: A verified modern cryptographic library. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 1789–1806. ACM, 2017. 10

Appendix

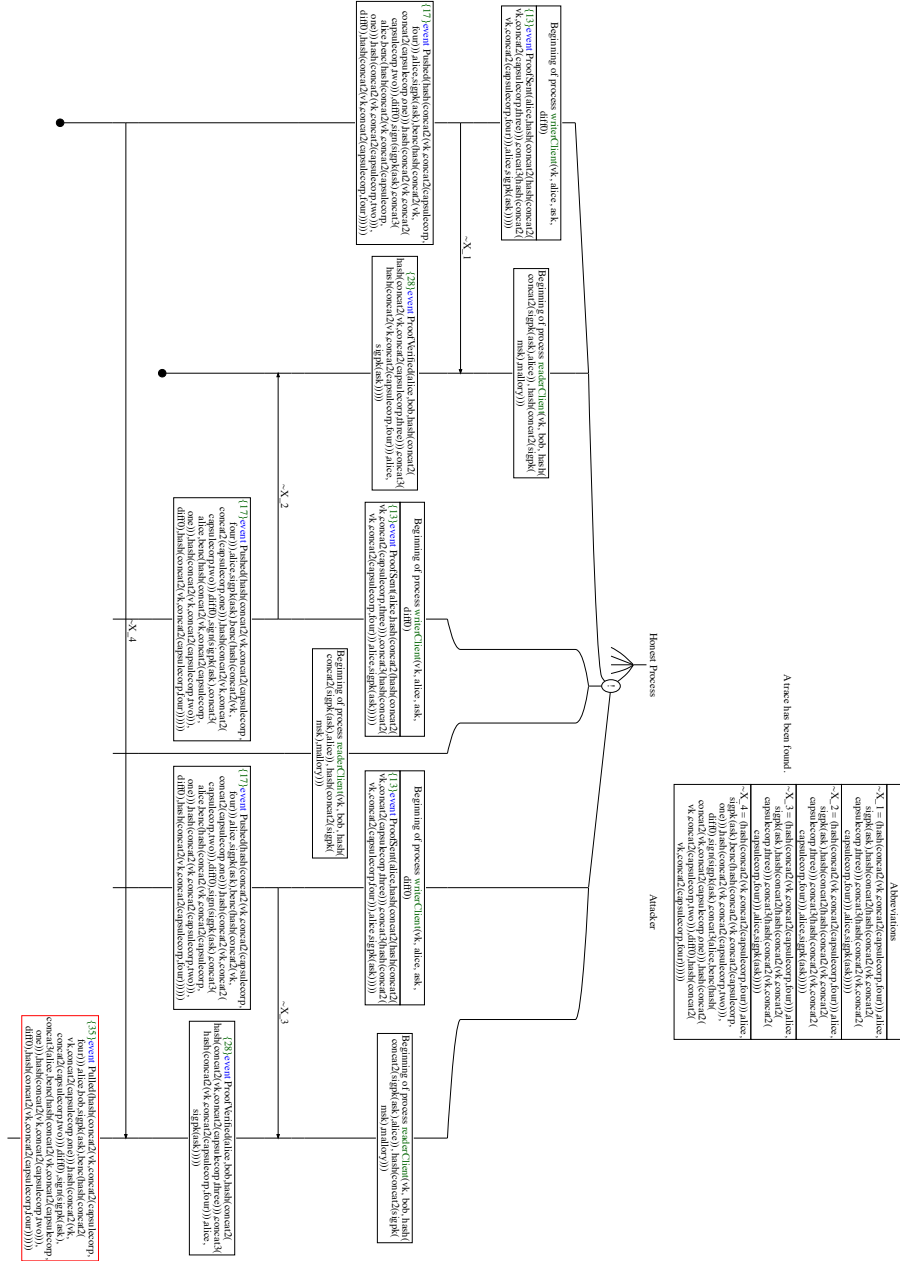


Fig. 1. A ProVerif trace demonstrating a full run of the Capsule protocol, as modeled. The Pulled event is reachable, indicating that a correct protocol run has been completed.