

Proof Search and Certificates for Evidential Transactions

Vivek Nigam¹, Giselle Reis²[0000–0002–5145–9829], Samar Rahmouni², and
Harald Ruess³

¹ Huawei Munich Research Center, Germany vivek.nigam@gmail.com

² Carnegie Mellon University, Qatar giselle@cmu.edu, srahmoun@andrew.cmu.edu

³ fortiss GmbH, Germany ruess@fortiss.org

Abstract. Attestation logics have been used for specifying systems with policies involving different principals. Cyberlogic is an attestation logic used for the specification of Evidential Transactions (ETs). In such transactions, evidence has to be provided supporting its validity with respect to given policies. For example, visa applicants may be required to demonstrate that they have sufficient funds to visit a foreign country. Such evidence can be expressed as a Cyberlogic proof, possibly combined with non-logical data (*e.g.*, a digitally signed document). A key issue is how to construct and communicate such evidence/proofs. It turns out that attestation modalities are challenging to use established proof-theoretic methods such as focusing. Our first contribution is the refinement of Cyberlogic proof theory with knowledge operators which can be used to represent knowledge bases local to one or more principals. Our second contribution is the identification of an executable fragment of Cyberlogic, called Cyberlogic programs, enabling the specification of ETs. Our third contribution is a sound and complete proof system for Cyberlogic programs enabling proof search similar to search in logic programming. Our final contribution is a proof certificate format for Cyberlogic programs inspired by Foundational Proof Certificates as a means to communicate evidence and check its validity.

Keywords: Attestation Logics · Proof Search · Sequent Calculus

1 Introduction

Attestation logics [3,15,22,16,7,6,29] have been used for the specification of policies of distributed systems, such as access control systems [3], distributed authorization policies [15,22], and evidential transactions (ETs) [16,6,7,7,29]. In these logics, one specifies policies involving attestation formulas of the form $K \triangleright F$, where K is a principal (or agent) in the system.

Cyberlogic is an attestation logic for ETs. In Cyberlogic, cryptographic keys K are identified with specific authorities, and attestations $K \triangleright A$ express the fact that principal K attests to statement A . For example, K may be a visa-granting authority and A the statement that the visa requester is authorized

to enter the specified country by the end of the year and at most once. An evidential transaction might issue a visa given that proof of sufficient funds has been provided in the form of a digital certificate whose validity can then be verified by customs authorities upon entry.

Formally, evidence in ETs can be expressed as a Cyberlogic proof. To carry out an ET, a Cyberlogic proof demonstrating policy compliance shall be produced and communicated. ETs therefore enable trust in, for example, distributed exchanges in electronic commerce, by enabling the exchange of various forms of *verifiable evidence*, such as evidence of funds in the visa example above.

The problem of producing attestation logic proofs (and proof objects) has not been given enough attention so far. Attestation logics have been formalized as Hilbert-style proof systems [3,16] that do not have the sub-formula property and therefore are not suitable for proof search. Other works on authorization logics [15,22] have proposed sequent calculi which do possess the sub-formula property. However, the search space is too great to enable efficient proof search.

The established proof-theoretic method for proof search is *focusing* [4,19]. Focusing distinguishes between inference rules that have “don’t know” and “don’t care” non-determinism to prune the proof search space. Interestingly, focused proof systems [8,19] provide a proof-theoretical justification for backward and forward-chaining, two proof-search strategies for Horn clauses (logic programs). Such justification, however, breaks when programs contain modalities, such as attestation modalities, *i.e.*, formulas of the form $K \multimap F$. This is because focusing is lost whenever any of these formulas is encountered and therefore, improvements to the search space because of focusing is not so significant for attestation logics.

Our main goal is the study of Cyberlogic’s proof theory in order to enable proof search (similar to the search involved in logic programming) and the generation of proof certificates for the communication of evidence in ETs.

Our first contribution, detailed in Section 2, is a Gentzen style proof system for Cyberlogic that admits cut elimination. A feature of the proof system is that it enables the combination of evidence represented as logical derivations as well as digital evidence, *e.g.*, signed hashes of documents, financial statements, medical records. The logic also includes a knowledge operator for sets of principals.

Our second contribution, detailed in Section 3, is the identification of a fragment of Cyberlogic, called Cyberlogic programs, akin to Horn clauses used in logic programming. This is motivated by the ongoing work on building distributed logic programming engines for ETs which extend existing engines [11] with attestations of the form $K \multimap A$.

Our third contribution, also detailed in Section 3, addresses the challenge of how to efficiently construct Cyberlogic program proofs. We propose a focused inspired proof system for Cyberlogic programs and prove that it is sound and complete in this fragment. This system enables more efficient proof search.

Our last contribution, detailed in Section 4, addresses the challenge of how to efficiently communicate evidence. We propose a proof certificate format for Cyberlogic programs inspired by Foundational Proof Certificates (FPCs) [10]. FPCs enable the reconstruction of proofs by using simple logic programs as

guides. This means that such certificates can elide parts that can be easily reconstructed or which one is willing to reconstruct.

2 Cyberlogic Proof Theory

Cyberlogic [29] is an intuitionistic modal logic which can be used for specifying ETs. The logic is parametrized by a finite set of principals $\mathcal{K} = \{K_1, \dots, K_n\}$, which are used in formulas as follows:

- $K_i : \triangleright F$: meaning that principal K_i attests the (Cyberlogic) formula F ;
- $\text{kb}_{\mathcal{Q}}F$, where $\mathcal{Q} \subseteq \mathcal{K}$: meaning that all principals in \mathcal{Q} know F , or, alternatively, that the combined knowledge of principals in \mathcal{Q} imply F ; and
- $\text{evidence}_{K_i}A$: standing for an external evidence signed by principal K_i .

External evidences are left unspecified since they fall outside the logical scope and depend on the ET being formalized. For example, $\text{evidence}_{K_i}A$ could be signed hashes of tickets, financial statements, medical records, etc. In Cyberlogic the evidence associated with an ET is a combination of a formal proof (in sequent calculus) and a collection of external evidences.

Cyberlogic formulas are constructed according to the following grammar:

$$F, G ::= A \mid F \wedge G \mid F \vee G \mid F \supset G \mid \top \mid \perp \mid K : \triangleright F \mid \text{kb}_{\mathcal{Q}}F \mid \forall x.F \mid \exists x.F$$

where A is an atom, $K \in \mathcal{K}$, and $\mathcal{Q} \subseteq \mathcal{K}$. The formula $K : \triangleright F$ is read as “principal K attests F ” and acts like the *says* modality in lax logics [14,27]. The formula $\text{kb}_{\mathcal{Q}}F$ is read as “principals in \mathcal{Q} know F ” and is inspired by the *knows* modality used in linear authorization logics [15,22]. Different from that logic, Cyberlogic allows the direct specification of knowledge shared by multiple principals, as illustrated in Example 1.

Cyberlogic sequents are of the shape $\Gamma \longrightarrow G$, where Γ is a multiset of formulas. The Cyberlogic proof system, $\text{CL}_{\mathcal{K}}$, is depicted in Figure 1. Rules for the intuitionistic connectives $\wedge, \vee, \supset, \forall, \exists$ are as in LJ [30]. The new rules are the ones involving assertions $K : \triangleright F$ and $\text{kb}_{\mathcal{Q}}$. Note that a “built-in” contraction of the main formula is needed on the left premise of \supset_l and the premise of \forall_l , as expected in intuitionistic logics. Also, the rule kb_l has an explicit contraction on the premise. These contractions are needed for cut admissibility (Theorem 2).

Rules \triangleright_l and \triangleright_r specify that \triangleright is a lax modality [27,22,25]. The intuition behind \triangleright_l is: if an assertion G of a principal K is provable using F , then it is also provable if K attests F . Rule \triangleright_r specifies that principals are rational, *i.e.*, they can always attest formulas that are derivable. Differently from existing systems with lax modalities, $\text{CL}_{\mathcal{K}}$ has the rule *ext*. This rule allows a proof of an attestation $K : \triangleright A$ to be completed whenever a principal provides evidence $\text{evidence}_K A$ for the claim A . This formalizes the intuition that principals may use digital evidence signed by their private key. We leave the definition of evidence unspecified as it depends on the intended ET specified.

$$\begin{array}{c}
\frac{}{\Gamma, A \longrightarrow A} \text{init} \quad \frac{\text{evidence}_{\mathcal{K}} A}{\Gamma \longrightarrow \mathcal{K} \triangleright A} \text{ext} \quad \frac{}{\Gamma \longrightarrow \top} \top_r \quad \frac{}{\Gamma, \perp \longrightarrow C} \perp_l \\
\frac{\Gamma, F_1, F_2 \longrightarrow G}{\Gamma, F_1 \wedge F_2 \longrightarrow G} \wedge_l \quad \frac{\Gamma \longrightarrow F_1 \quad \Gamma \longrightarrow F_2}{\Gamma \longrightarrow F_1 \wedge F_2} \wedge_r \quad \frac{\Gamma, F_1 \longrightarrow G \quad \Gamma, F_2 \longrightarrow G}{\Gamma, F_1 \vee F_2 \longrightarrow G} \vee_l \\
\frac{\Gamma, F_1 \supset F_2 \longrightarrow F_1 \quad \Gamma, F_2 \longrightarrow G}{\Gamma, F_1 \supset F_2 \longrightarrow G} \supset_l \quad \frac{\Gamma, F_1 \longrightarrow F_2}{\Gamma \longrightarrow F_1 \supset F_2} \supset_r \quad \frac{\Gamma \longrightarrow F_i}{\Gamma \longrightarrow F_1 \wedge F_2} \vee_{r_i} \\
\frac{\Gamma, \forall x.F, F[t/x] \longrightarrow G}{\Gamma, \forall x.F \longrightarrow G} \forall_l \quad \frac{\Gamma \longrightarrow F[\alpha/x]}{\Gamma \longrightarrow \forall x.F} \forall_r \quad \frac{\Gamma, F[\alpha/x] \longrightarrow G}{\Gamma, \exists x.F \longrightarrow G} \exists_l \quad \frac{\Gamma \longrightarrow F[t/x]}{\Gamma \longrightarrow \exists x.F} \exists_r \\
\frac{\Gamma, F \longrightarrow \mathcal{K} \triangleright G}{\Gamma, \mathcal{K} \triangleright F \longrightarrow \mathcal{K} \triangleright G} \triangleright_l \quad \frac{\Gamma \longrightarrow F}{\Gamma \longrightarrow \mathcal{K} \triangleright F} \triangleright_r \quad \frac{\Gamma, \text{kb}_{\mathcal{Q}} F, F \longrightarrow G}{\Gamma, \text{kb}_{\mathcal{Q}} F \longrightarrow G} \text{kb}_l \quad \frac{\Gamma \mid_{\mathcal{Q}} \longrightarrow F}{\Gamma \longrightarrow \text{kb}_{\mathcal{Q}} F} \text{kb}_r
\end{array}$$

Fig. 1. $CL_{\mathcal{K}}$ – Cyberlogic proof system for $\mathcal{K} = \{K_1, \dots, K_n\}$. Here A is an atomic formula, $\mathcal{Q} \subseteq \mathcal{K}$, and $\Gamma \mid_{\mathcal{Q}} = \{\text{kb}_{\mathcal{Q}'} F \mid \text{kb}_{\mathcal{Q}'} F \in \Gamma \wedge \mathcal{Q}' \subseteq \mathcal{Q}\}$. Moreover, in rules \exists_l and \forall_r , α is a fresh constant not appearing in Γ nor F .

Rules kb_l and kb_r refine Cyberlogic by enabling the collection of logical theories known by a set of principals. Such theories act as *knowledge bases*. Rule kb_l specifies that any common knowledge can be part of a knowledge base. The interesting rule is kb_r , which specifies that $\text{kb}_{\mathcal{Q}} F$ can only be proved using the local knowledge or evidence provided by principals in \mathcal{Q} . This is formally captured by restricting Γ in kb_r 's premise to the set $\Gamma \mid_{\mathcal{Q}} = \{\text{kb}_{\mathcal{Q}'} F \mid \text{kb}_{\mathcal{Q}'} F \in \Gamma \wedge \mathcal{Q}' \subseteq \mathcal{Q}\}$. This is a powerful construct that increases the expressiveness of Cyberlogic. In particular, it is straightforward to specify that certain assertions can be concluded from the shared knowledge of a set of principals.

Proposition 1. *The following sequents are provable in $CL_{\mathcal{K}}$ for all $\mathcal{K} \in \mathcal{K}$ and formulas F_1, F_2 . $F_1 \equiv F_2$ represents the sequents $(F_1 \longrightarrow F_2)$ and $(F_2 \longrightarrow F_1)$:*

1. $F \longrightarrow \mathcal{K} \triangleright F$
2. $\text{kb}_{\mathcal{Q}} F \longrightarrow F$
3. $\text{kb}_{\{\mathcal{K}\}} F \longrightarrow \mathcal{K} \triangleright F$
4. $\mathcal{K} \triangleright \mathcal{K} \triangleright F \equiv \mathcal{K} \triangleright F$
5. $\text{kb}_{\mathcal{Q}'} F \longrightarrow \text{kb}_{\mathcal{Q}} F$, if $\mathcal{Q}' \subseteq \mathcal{Q}$. In particular, $\text{kb}_{\mathcal{Q}_1} \text{kb}_{\mathcal{Q}_2} F \longrightarrow \text{kb}_{\mathcal{Q}_1 \cup \mathcal{Q}_2} F$.
6. $\text{kb}_{\mathcal{Q}_1} F \wedge \text{kb}_{\mathcal{Q}_2} F \longrightarrow \text{kb}_{\mathcal{Q}_1 \cup \mathcal{Q}_2} F$
7. $\mathcal{K} \triangleright (F_1 \wedge F_2) \equiv \mathcal{K} \triangleright F_1 \wedge \mathcal{K} \triangleright F_2$
8. $\text{kb}_{\mathcal{Q}} (F_1 \wedge F_2) \equiv \text{kb}_{\mathcal{Q}} F_1 \wedge \text{kb}_{\mathcal{Q}} F_2$
9. $(\mathcal{K} \triangleright F_1 \vee \mathcal{K} \triangleright F_2) \longrightarrow \mathcal{K} \triangleright (F_1 \vee F_2)$
10. $\text{kb}_{\mathcal{Q}} A \vee \text{kb}_{\mathcal{Q}} B \longrightarrow \text{kb}_{\mathcal{Q}} (A \vee B)$
11. $\mathcal{K} \triangleright (F_1 \supset F_2) \longrightarrow (\mathcal{K} \triangleright F_1 \supset \mathcal{K} \triangleright F_2)$
12. $\text{kb}_{\mathcal{Q}} (F_1 \supset F_2) \longrightarrow (\text{kb}_{\mathcal{Q}} F_1 \supset \mathcal{K} \triangleright F_2)$
13. $\mathcal{K} \triangleright (\nabla x.F) \equiv \nabla x. \mathcal{K} \triangleright F$, $\nabla \in \{\forall, \exists\}$
14. $\text{kb}_{\mathcal{Q}} (\nabla x.F) \equiv \nabla x. \text{kb}_{\mathcal{Q}} F$, $\nabla \in \{\forall, \exists\}$

Moreover, the following sequents are not provable if $\mathcal{K}_1 \neq \mathcal{K}_2$ and $\mathcal{Q}_1 \neq \mathcal{Q}_2$:

1. $\mathcal{K} \triangleright F \not\rightarrow F$
2. $F \not\rightarrow \text{kb}_{\mathcal{Q}} F$
3. $\mathcal{K} \triangleright F \not\rightarrow \text{kb}_{\{\mathcal{K}\}} F$
4. $\mathcal{K}_1 \triangleright (\mathcal{K}_2 \triangleright F) \not\rightarrow \mathcal{K}_2 \triangleright (\mathcal{K}_1 \triangleright F)$
5. $\text{kb}_{\mathcal{Q}_1} (\text{kb}_{\mathcal{Q}_2} F) \not\rightarrow \text{kb}_{\mathcal{Q}_2} (\text{kb}_{\mathcal{Q}_1} F)$
6. $\text{kb}_{\mathcal{Q}_1 \cup \mathcal{Q}_2} F \not\rightarrow \text{kb}_{\mathcal{Q}_i} F$, $i \in \{1, 2\}$
7. $\text{kb}_{\mathcal{Q}_1 \cup \mathcal{Q}_2} F \not\rightarrow \text{kb}_{\mathcal{Q}_1} F \wedge \text{kb}_{\mathcal{Q}_2} F$
8. $\text{kb}_{\mathcal{Q}} \mathcal{K} \triangleright A \not\rightarrow \mathcal{K} \triangleright \text{kb}_{\mathcal{Q}} A$
9. $\mathcal{K} \triangleright \text{kb}_{\mathcal{Q}} A \not\rightarrow \text{kb}_{\mathcal{Q}} \mathcal{K} \triangleright A$

In the remainder of the paper, we elide the set of principals \mathcal{K} whenever it can be deduced from the context.

Example 1. (Shared Knowledge) The ability to use kb with multiple principals allows the derivation of facts that depend on the combination of knowledge of multiple principals. Consider that principal K_1 knows A and $B \supset C$, and principal K_2 knows $A \supset B$, then the following sequent is provable in CL:

$$\text{kb}_{\{K_1\}}A, \text{kb}_{\{K_1\}}B \supset C, \text{kb}_{\{K_2\}}A \supset B \longrightarrow \text{kb}_{\{K_1, K_2\}}C$$

Remark 1. The original Cyberlogic paper [6] (and technical report [5]) proposed two kinds of attestations, $:\triangleright$ and \triangleright , to distinguish when an attestation is derived from a digital evidence or logical inferences. This combination, however, does not yield to a proof system with the cut-elimination property [28].

The meta-theory of CL has been analysed using the L-framework [1], which uses rewriting logic to automatically derive structural proofs of sequent calculi properties [26]. The following lemma was used in the proofs of cut-elimination and invertibility.

Lemma 1. *If $\Gamma, K : \triangleright F \longrightarrow G$, then $\Gamma, F \longrightarrow G$.*

The proof proceeds by structural induction on the derivation of $\Gamma, K : \triangleright F \longrightarrow G$. The proof has been mechanically checked using the the L-framework with some few cases proved by hand.

As expected, $\supset_r, \wedge_r, \wedge_l, \vee_l, \forall_r, \exists_l$ are invertible whereas $\vee_r, \supset_l, \forall_l, \exists_r$ are not invertible. In addition, the rules $:\triangleright_l$ and kb_l are invertible whereas the $:\triangleright_r$ and kb_r are not invertible.

Lemma 2. *If $\Gamma, K : \triangleright F \longrightarrow K : \triangleright G$ then $\Gamma, F \longrightarrow K : \triangleright G$.*

This is a simple corollary of Lemma 1. Invertibility of kb_l is straightforward because of the contraction of the main formula.

Rules $:\triangleright_r$ and kb_r are not invertible. The counter examples are:

$$\begin{aligned} [:\triangleright_r] \quad & K : \triangleright a \longrightarrow K : \triangleright a \quad \text{but} \quad K : \triangleright a \not\rightarrow a \\ [\text{kb}_r] \quad & a, a \supset \text{kb}_K b \longrightarrow \text{kb}_K b \quad \text{but} \quad \not\rightarrow b \end{aligned}$$

Weakening is height perserving admissible in CL.

Theorem 1 (Identity expansion). *$F \longrightarrow F$ is provable in CL for any cyberlogic formula F .*

The proof is by structural induction on F .

Theorem 2 (Cut elimination). *If $\Gamma \longrightarrow F$ and $\Gamma, F \longrightarrow C$, then $\Gamma \longrightarrow C$.*

The proof proceeds by a nested induction on the structure of the proofs of $\Gamma \longrightarrow F$ and $\Gamma, F \longrightarrow C$, and the formula F . The noteworthy cases are the ones where cut needs to permute over kb rules. For kb_l , contraction of the main formula is needed, and the permutation over kb_r can be done only if cut is principal on the left (which is a lemma that can be proved). Details about these transformations are in Appendix A.1.

3 Cyberlogic Programs

Cyberlogic programs are fragment of CL which resembles Horn clauses in logic programming. Section 3.2 proposes a proof search operational semantics for cyberlogic programs and proves its soundness and completeness. The proof search discipline relies on ideas from focusing [4]. Focused proof systems for LJ [19] provide a proof theoretical justification of forward and backward chaining search. Each technique is enforced by the choice of polarity of atomic formulas: positive atoms lead to forward chaining and negative atoms lead to backward chaining. This correspondence, however, does not extend to cyberlogic due to attestation formulas $K :▷ A$ which cause focusing to be lost [22]. Consider the following example where the formula under focus is in brackets:

$$\frac{K_1 :▷ a \longrightarrow [K_1 :▷ a] \quad K_1 :▷ a, [K_2 :▷ b] \longrightarrow K_2 :▷ b}{K_1 :▷ a, [K_1 :▷ a \supset K_2 :▷ b] \longrightarrow K_2 :▷ b} \supset_l$$

In focused proof systems, forward chaining can be enforced by disallowing focus to be lost on the right formula in the left premise, *i.e.* $[K_1 :▷ a]$. However, if $:▷_r$ is applied to this sequent the premise would be $K_1 :▷ a \longrightarrow a$, which is not provable (see Proposition 1). In fact, $[K_1 :▷ a]$ must lose focus on the right for the proof to be completed. Therefore, if $:▷$ modalities are used in logic programs, other strategies for proof search need to be analysed.

3.1 Cyberlogic Program Syntax

Cyberlogic programs can be divided into goals, knowledge bases, common knowledge, and attestation clauses.

Goals (G) Cyberlogic programs are used to derive a goal G , defined as:

$$G ::= \top \mid K :▷ \text{kb}_{\mathcal{Q}}A \mid G_1 \wedge G_2 \mid \exists x.G$$

where A is an atomic formula. The restriction of $:▷ \text{kb}_{\mathcal{Q}}$ to atoms does not reduce the expressiveness of goals, given the equivalences in Proposition 1.

Knowledge Bases (\mathcal{B}): A knowledge base, written $\text{kb}_{\{K_i\}}\Gamma$, of a principal $K_i \in \mathcal{K}$ is a set of formulas Γ not containing the connectives $:▷$ or kb . Here, $\text{kb}_{\{K_i\}}\Gamma$ represents the set of formulas $\{\text{kb}_{\{K_i\}}F \mid F \in \Gamma\}$.

Intuitively, a knowledge base $\text{kb}_{\{K_i\}}\Gamma$ can be interpreted as K_i 's local knowledge. This means that K_i may use its own prover to derive new facts. For example, if Γ is a collection of Horn-clauses, then K_i may deploy a Prolog engine to derive some goal. Alternatively if Γ is a set of formulas in CNF form, then K_i may use resolution provers. The absence of modal connectives in knowledge bases has important impacts on the design of the proof certificate described in Section 4, as those may rely on existing certificates for different provers [10].

Common Knowledge (C): Common knowledge are knowledge bases that are known to all principals, written as $\text{kb}_\emptyset F$. Since $\emptyset \subseteq \mathcal{Q}$ for every \mathcal{Q} , these formulas remain in the context when applying kb_r . In this sense they contain first order formulas that may be used by all principals.

Attestation Formulas (D): Formulas of the form $K \vdash \text{kb}_\mathcal{Q} A$ are derived by attestation formulas of the form below where for all $1 \leq i \leq n$, $K_i \in \mathcal{K}$, $\mathcal{Q}_i \subseteq \mathcal{K}$, and A_1, \dots, A_n, A are atomic formulas and \mathbf{X} are bounded by universal quantifiers:

$$\begin{aligned} & \forall \mathbf{X}. (\text{kb}_{\mathcal{Q}_1}(K_1 \vdash A_1) \wedge \dots \wedge \text{kb}_{\mathcal{Q}_n}(K_n \vdash A_n) \wedge G \supset K \vdash (\text{kb}_\emptyset A)) \\ & \forall \mathbf{X}. (\text{kb}_{\mathcal{Q}_1}(K_1 \vdash A_1) \wedge \dots \wedge \text{kb}_{\mathcal{Q}_n}(K_n \vdash A_n) \wedge G \supset K \vdash (\text{kb}_{\{K\}} A)) \end{aligned}$$

Intuitively, an attestation formula belongs to a principal, namely K in the right-hand side of \supset . Such formulas derive K 's attestation of an atomic formula which is its own knowledge ($\text{kb}_{\{K\}} A$), or common knowledge ($\text{kb}_\emptyset A$). This means that K 's attestation formulas cannot derive knowledge belonging to other principals. Furthermore to derive an attestation, one can use the knowledge base of other principals, *i.e.* the formulas $\text{kb}_{\mathcal{Q}_i}(K_i \vdash A_i)$ or additional goals, *i.e.* G . Finally notice that $K \vdash (\text{kb}_\emptyset A)$ and $K \vdash (\text{kb}_{\{K\}} A)$ are attestation formulas themselves, where the left-hand side of \supset is empty (denoting \top).

The difference between formulas $K \vdash A$ and $K \vdash (\text{kb}_{\{K\}} A)$ is subtle. Note that the former can be derived using the evidence rule `ext`, while the latter cannot. $K \vdash (\text{kb}_{\{K\}} A)$ is K 's attestation that A follows from its local knowledge base. It is possible to specify that A can be derived from an external evidence, but this has to be made explicit by an attestation formula, *e.g.*, $\text{kb}_{\{K\}}(K \vdash A) \supset K \vdash (\text{kb}_{\{K\}} A)$. Note that this formula is not a tautology.

We are interested in proving goals from attestation formulas, knowledge bases, and common knowledge, which are formally represented by cyberlogic program sequents defined as follows.

Definition 1 (Cyberlogic Program Sequents (CPS)). A cyberlogic program sequent (CPS) is a sequent $\mathcal{C}, \mathcal{B}, \mathcal{D} \longrightarrow G$, where \mathcal{B} is a set of knowledge bases, \mathcal{C} is a set of common knowledge formulas, \mathcal{D} is a set of attestation formulas, and G is a goal formula.

Example 2. (Local Computations) This example illustrates the use of `kb` to specify when parts of a derivation can be proved locally using a principal's knowledge. Consider that the following clause

$$\text{kb}_{\{K_1\}}(K_1 \vdash F_1) \wedge \text{kb}_{\{K_2\}}(K_2 \vdash F_2) \supset K \vdash \text{kb}_{\{K\}} G$$

specifies that for K to attest G , K_1 and K_2 have to attest F_1 and F_2 respectively, using *their own local theories, common knowledge, or evidence*. This means that computations carried out by K_1 and K_2 to derive their assertions $K_1 \vdash F_1$ and $K_2 \vdash F_2$ respectively, do not depend on other principals and therefore, the search for these derivations can be performed locally.

Example 3. (Levels of Trust) This example illustrates the use of **kb** to specify that some evidence should only be trusted if derived from trusted sources. Consider three principals $\mathcal{K} = \{K_T, K_U, K\}$ where K trusts evidence from K_T , but not all evidence from K_U . Then the following clause

$$\text{kb}_{\{K, K_T\}}(K : \triangleright \text{critical}(\text{ok})) \wedge \text{kb}_{\mathcal{K}}(K : \triangleright \text{nonCritical}(\text{ok})) \supset K : \triangleright \text{kb}_{\emptyset}(\text{all}(\text{ok}))$$

specifies that K can attest that everything is **ok** as a common knowledge if all the non-critical and critical elements are **ok**. However, the check of critical parts can only be performed by principals K trusts, namely K itself or K_T . Information from K_U 's knowledge bases cannot be used in the proof of **critical(ok)**.

Example 4. (Simplified Visa) Consider a visa issuing scenario where an applicant applies to a consulate (**cons**) for an entry visa. This is an example of an ET as, to obtain the visa, evidence has to be provided that, for example, the applicant has no crime records, or that they have sufficient funds. We illustrate how such an ET can be specified in Cyberlogic.

The formula below labelled **main** specifies conditions for a visa to be issued:

$$\begin{aligned} \mathbf{main}: & \forall \text{Id}. \forall \text{Doc}. \forall \text{V}. (\text{kb}_{\{\text{cons}\}}(\text{cons} : \triangleright \text{validVisit}(\text{Id}, \text{Doc})) \\ & \wedge \text{kb}_{\{\text{cons}\}}(\text{cons} : \triangleright \text{prepVisa}(\text{Id}, \text{V})) \\ & \wedge \text{cons} : \triangleright \text{kb}_{\{\text{cons}\}}(\text{suffFin}(\text{Doc})) \wedge \text{police} : \triangleright \text{kb}_{\{\text{police}\}}(\text{noCrimeRec}(\text{Id})) \\ & \supset \text{cons} : \triangleright \text{kb}_{\text{cons}}(\text{issVisa}(\text{Id}, \text{Doc}, \text{V}))) \end{aligned}$$

The transaction for **cons** issuing a visa V to an applicant Id requires **cons** to attest validity of Id 's visit by itself ($\text{validVisit}(\text{Id}, \text{Doc})$) and Id 's criminal record with the help of the **police** ($\text{noCrimeRec}(\text{Id})$). In addition, **cons** also needs to attest Id 's financial status ($\text{suffFin}(\text{Doc})$).

The following two clauses expand on how **cons** can attest $\text{suffFin}(\text{Doc})$: either via an employment contract or a bank statement.

$$\mathbf{cont}: \text{kb}_{\{\text{cons}\}}(\forall \text{Doc}. \forall \text{Cont}. (\text{empContract}(\text{Doc}, \text{Cont}) \wedge \text{valid}(\text{Cont}) \supset \text{suffFin}(\text{Doc})))$$

$$\begin{aligned} \mathbf{bankStmt}: & \forall \text{Doc}. \forall \text{Stmt}. (\text{kb}_{\{\text{cons}\}}(\text{cons} : \triangleright \text{bankStmt}(\text{Doc}, \text{Stmt})) \\ & \wedge \text{bank} : \triangleright \text{kb}_{\{\text{bank}\}}(\text{valid}(\text{Stmt})) \supset \text{cons} : \triangleright \text{kb}_{\{\text{cons}\}}(\text{suffFin}(\text{Doc}))) \end{aligned}$$

The formula labeled **cont** belongs to **cons**'s knowledge base. This means that **cons** can check the validity of an employment contract without evidence from other principals. For example, $\text{valid}(\text{Cont})$ may check the contract duration and salary. The formula labeled **bankStmt**, on the other hand, takes the bank statement **Stmt** from the given documents, **Doc**, and requires the **bank** to validate it using its knowledge base. This makes sense as Id 's financial records are sensitive and do not need to be disclosed to anyone else apart from her financial institute.

These clauses also illustrate the subtle difference between goal formulas $K : \triangleright \text{kb}_{\{K\}} F$ and knowledge base formulas $\text{kb}_{\{K\}} K : \triangleright F$. For example, in the **main** clause, the fact that applicant has come to their appointment at the consulate does not depend on other agents and that is why we use a knowledge base formula. The same applies to the visa preparation. On the other hand, the fact that applicant has sufficient funds may require evidence from other parties, *e.g.*, the applicant's **bank**. Therefore this is specified as a goal.

3.2 CPS Proof Search

Proof search of CPS can be divided into the following phases: goal decomposition, \triangleright_l application, attestation formula decomposition, $K \triangleright A$ decomposition, and first-order reasoning. We define a (focusing inspired) sequent calculus for the CPS fragment, called $\text{CL}_{\mathcal{P}}$ (Figure 2) for enforcing this proof search discipline. Sequents in $\text{CL}_{\mathcal{P}}$ have the following shape: $\Theta; \Lambda; \Delta \longrightarrow F$, where Θ contains kb formulas, Λ contains attestation formulas, Δ contains formulas of the form $K \triangleright \text{kb}_{\mathcal{Q}}A$, and F is either a goal formula, $\text{kb}_{\mathcal{Q}}(K \triangleright A)$, $K \triangleright A$ or A , where A is an atom. Moreover, the part of the sequent containing the formula that is being decomposed will be enclosed in square brackets. This will help distinguishing the phases mentioned above.

$$\begin{array}{c}
 \textbf{Goal decomposition} \\
 \frac{}{\Theta; \Lambda; \Delta \longrightarrow [\top]} \top_r \quad \frac{\Theta; \Lambda; \Delta \longrightarrow [G_1] \quad \Theta; \Lambda; \Delta \longrightarrow [G_2]}{\Theta; \Lambda; \Delta \longrightarrow [G_1 \wedge G_2]} \wedge_r \quad \frac{\Theta; \Lambda; \Delta \longrightarrow [G[t/x]]}{\Theta; \Lambda; \Delta \longrightarrow [\exists x.G]} \exists_r \\
 \\
 \frac{\Theta; \Lambda; [\Delta] \longrightarrow K \triangleright \text{kb}_{\mathcal{Q}}A}{\Theta; \Lambda; \Delta \longrightarrow [K \triangleright \text{kb}_{\mathcal{Q}}A]} G \Rightarrow \triangleright_l \quad \frac{\Theta \mid_{\mathcal{Q}}^* \longrightarrow A}{\Theta; \Lambda; \Delta \longrightarrow [K \triangleright \text{kb}_{\mathcal{Q}}A]} \triangleright_r + \text{kb}_r + \text{kb}_l \\
 \\
 \textbf{\(\(\triangleright_l\) application)} \\
 \frac{\Theta, \text{kb}_{\mathcal{Q}}A; \Lambda; [\Delta] \longrightarrow K \triangleright \text{kb}_{\mathcal{Q}'}A'}{\Theta; \Lambda; [\Delta, K \triangleright \text{kb}_{\mathcal{Q}}A] \longrightarrow K \triangleright \text{kb}_{\mathcal{Q}'}A'} \triangleright_l \\
 \\
 \frac{\Theta; [\Lambda]; \Delta^\dagger \longrightarrow K \triangleright \text{kb}_{\mathcal{Q}}A}{\Theta; \Lambda; [\Delta^\dagger] \longrightarrow K \triangleright \text{kb}_{\mathcal{Q}}A} \triangleright_l \Rightarrow \text{att} \quad \frac{\Theta; \Lambda; \Delta^\dagger \longrightarrow [K \triangleright \text{kb}_{\mathcal{Q}}A]}{\Theta; \Lambda; [\Delta^\dagger] \longrightarrow K \triangleright \text{kb}_{\mathcal{Q}}A} \triangleright_l \Rightarrow G \\
 \\
 \textbf{Attestation formula decomposition} \\
 \frac{\Theta; \Lambda; \Delta \longrightarrow [G\sigma] \quad \Theta; \Lambda; [\Delta, K \triangleright \text{kb}_{\mathcal{Q}}A\sigma] \longrightarrow K' \triangleright \text{kb}_{\mathcal{Q}'}A' \quad \Theta \mid_{\mathcal{Q}_1; \cdot; \cdot} \longrightarrow [K_1 \triangleright A_1\sigma] \quad \cdots \quad \Theta \mid_{\mathcal{Q}_n; \cdot; \cdot} \longrightarrow [K_n \triangleright A_n\sigma]}{\Theta; [\Lambda, \forall \mathbf{X}. (\text{kb}_{\mathcal{Q}_1}(K_1 \triangleright A_1) \wedge \cdots \wedge \text{kb}_{\mathcal{Q}_n}(K_n \triangleright A_n) \wedge G \supset K \triangleright \text{kb}_{\mathcal{Q}}A)]; \Delta \longrightarrow K' \triangleright \text{kb}_{\mathcal{Q}'}A'} \text{att} \\
 \\
 \textbf{\(\(\triangleright A\) decomposition)} \\
 \frac{\text{evidence}_{K}A}{\Theta; \cdot; \cdot \longrightarrow [K \triangleright A]} \text{ext} \quad \frac{\Theta^* \longrightarrow A}{\Theta; \cdot; \cdot \longrightarrow [K \triangleright A]} \triangleright_r + \text{kb}_l \\
 \\
 \textbf{First-order reasoning:} \\
 \text{All first-order rules from CL on } \Theta^* \longrightarrow A \text{ sequents}
 \end{array}$$

Fig. 2. $\text{CL}_{\mathcal{P}}$ – Sequent calculus for cyberlogic programs. A , A' and A_i are atoms, Δ^\dagger is such that for all $K' \triangleright \text{kb}_{\mathcal{Q}'}A' \in \Delta^\dagger$, $K' \neq K$, and $\Theta^* = \{F \mid \text{kb}_{\mathcal{Q}}F \in \Theta\}$.

Lemma 3. *The kb_r rules permutes down every left rule in the CPS fragment.*

Proof. First we note that, in the CPS fragment, \wedge, \vee, \forall , and kb formulas on the left do not have kb modalities as subformulas. We look at the case of kb_l , as the others follow a similar argument.

Since F is not a kb formula, then $F \notin (\Gamma, \text{kb}_{\mathcal{Q}'}F, F) \mid_{\mathcal{Q}}$. Therefore we can conclude that $(\Gamma, \text{kb}_{\mathcal{Q}'}F, F) \mid_{\mathcal{Q}} = (\Gamma, \text{kb}_{\mathcal{Q}'}F) \mid_{\mathcal{Q}}$ and the permutation is:

$$\frac{\frac{(\Gamma, \text{kb}_{\mathcal{Q}'}F, F) \mid_{\mathcal{Q}} \xrightarrow{\varphi} G}{\Gamma, \text{kb}_{\mathcal{Q}'}F, F \xrightarrow{\varphi} \text{kb}_{\mathcal{Q}}G} \text{kb}_r}{\Gamma, \text{kb}_{\mathcal{Q}'}F \xrightarrow{\varphi} \text{kb}_{\mathcal{Q}}G} \text{kb}_l}{\Gamma, \text{kb}_{\mathcal{Q}'}F \xrightarrow{\varphi} \text{kb}_{\mathcal{Q}}G} \text{kb}_r \rightsquigarrow \frac{(\Gamma, \text{kb}_{\mathcal{Q}'}F) \mid_{\mathcal{Q}} \xrightarrow{\varphi} G}{\Gamma, \text{kb}_{\mathcal{Q}'}F \xrightarrow{\varphi} \text{kb}_{\mathcal{Q}}G} \text{kb}_r$$

The case for \triangleright_l holds vacuously, as it is impossible to have \triangleright_l immediately below kb_r since the former requires the right formula to be of the shape $\text{K} \triangleright$.

The remaining case is \supset_l . Observe that in the CPS fragment, the formula F_2 in $F_1 \supset F_2$ is of the form $\text{K} \triangleright \text{kb}_{\mathcal{Q}'}A$. Therefore, $(\Gamma, F_2) \mid_{\mathcal{Q}} = \Gamma \mid_{\mathcal{Q}}$. Also, $(\Gamma, F_1 \supset F_2) \mid_{\mathcal{Q}} = \Gamma \mid_{\mathcal{Q}}$. Thus the permutation is:

$$\frac{\frac{\Gamma \xrightarrow{\varphi} F_1 \quad \frac{(\Gamma, F_2) \mid_{\mathcal{Q}} \xrightarrow{\varphi} G}{\Gamma, F_2 \xrightarrow{\varphi} \text{kb}_{\mathcal{Q}}G} \text{kb}_r}{\Gamma, F_1 \supset F_2 \xrightarrow{\varphi} \text{kb}_{\mathcal{Q}}G} \supset_l}{\Gamma, F_1 \supset F_2 \xrightarrow{\varphi} \text{kb}_{\mathcal{Q}}G} \supset_l}{\Gamma, F_1 \supset F_2 \xrightarrow{\varphi} \text{kb}_{\mathcal{Q}}G} \supset_l \rightsquigarrow \frac{(\Gamma, F_1 \supset F_2) \mid_{\mathcal{Q}} \xrightarrow{\varphi} G}{\Gamma, F_1 \supset F_2 \xrightarrow{\varphi} \text{kb}_{\mathcal{Q}}G} \text{kb}_r$$

□

Notice that it is crucial for attestation formulas to have a \triangleright modality formula on the consequent, otherwise Lemma 3 would not hold. As seen below, this lemma is key to proving completeness of the proof search procedure for CPS.

Theorem 3 (Soundness and completeness of $\text{CL}_{\mathcal{P}}$). $\Theta; \Lambda; \Delta \longrightarrow [F]$ in $\text{CL}_{\mathcal{P}}$ if and only if $\Theta, \Lambda, \Delta \longrightarrow F$ in CL

Proof. Soundness is straightforward: a proof in $\text{CL}_{\mathcal{P}}$ can be transformed into a proof in CL by using the same logical rules (possibly expanded – e.g. att becomes a sequence of $\forall_l + \supset_l + \wedge_r + \text{kb}_r$) and skipping the phase transition rules \Rightarrow (which only change the syntax of the sequent, but not its content).

Completeness is achieved by reasoning about invertibility and permutability of inference rules in the specific case of CPS. We argue that each phase can be performed in the proposed order.

Goal decomposition The goal formula can be eagerly decomposed until becoming $\text{K} \triangleright \text{kb}_{\mathcal{Q}}A$ before applying other rules because: \top_r and \wedge_r are invertible, and in the absence of \forall_r and \exists_l, \exists_r permutes down every rule. Once the right side formula is $\text{K} \triangleright \text{kb}_{\mathcal{Q}}A$, there are two options to continue: (1) change to \triangleright_l application phase, or (2) apply rules $\triangleright_r + \text{kb}_r + \text{kb}_l$ in Figure 1.

The first case is discussed below. In the second case, we need to argue that kb_r may be applied immediately above \triangleright_r . Once \triangleright_r is applied, we could choose

a formula from the context to continue with. However, kb_r permutes down all left rules for the CPS fragment, as shown in Lemma 3. Therefore any proof that continues with a formula in Θ , Λ , or Δ above $:\triangleright_r$ can be transformed into a proof where kb_r is applied immediately above $:\triangleright_r$. Since kb_l is invertible, it can be applied to exhaustion safely.

$:\triangleright_l$ application After eagerly decomposing the goal, $:\triangleright_l$ can be applied to exhaustion since it is an invertible rule (Lemma 2).

Attestation formula decomposition This phase contains only one rule, namely **att**, which encompasses \forall_l , \supset_l , \wedge_r , and kb_r . The quantifier rule can always be delayed until its subformula is needed, and \wedge_r is an invertible rule, therefore these can be chained together without loss of completeness. Due to Lemma 3, the application of kb_r can be permuted down for the CPS fragment and thus it is safe to apply the rule as soon as possible.

The two top premises of **att** force the proof search to go back to applying invertible rules, which does not break completeness.

$\text{K}:\triangleright A$ decomposition Once this state is reached, Θ is left with **kb** formulas whose subformulas are in first-order logic (i.e., no modalities). In this case, one can either close the proof with an external evidence, or apply $:\triangleright_r + \text{kb}_l$ to release the atom on the right side. The eager application of kb_l is justified due to its invertibility. It can also be delayed until this point because it permutes up \supset_l and $:\triangleright_r$ in CL, and it permutes up kb_r in the CPS fragment (Lemma 3).

First-order reasoning From this point onwards, there are no modalities in the sequent so it will be proved using only first-order reasoning. □

4 Proof Certificates

Cyberlogic programs may be used to derive facts about attestation (goals), using pure logical reasoning (knowledge bases), principal delegation (attestation formulas), and external evidence. Once a goal is derived, evidence shall be available so that any interested party can verify that the proof is correct. Verifiable evidence means that entities do not need to trust each other's proof producing process, as long as they can check the proofs using their own trusted processes.

Given a cyberlogic program sequent of the shape: $\Theta; \Lambda; \Delta \longrightarrow G$ one could take its full sequent calculus proof in $\text{CL}_{\mathcal{P}}$ as evidence. If the interested parties know the calculus, checking validity of proofs reduces to checking the valid application of each rule. However, these proofs are too fine grained, and contain many uninteresting details that can be easily inferred. Proof certificates elide such details, and keep only the crucial steps for proof reconstruction.

Proof certificates for cyberlogic are defined inspired by λ -terms and *foundational proof certificates* [9,21] (FPC). FPC is a framework for checking proofs in different formalisms using a small trusted kernel. The proposed kernels are the sequent calculus focused systems LKF and LJF [19] for LK and LJ respectively, augmented with predicates for guiding proof search [10]. The definition

$$\begin{array}{c}
\frac{}{\text{top} : \Theta; A; \Delta \rightarrow [\top]} \top_r \quad \frac{\Xi : \Theta; A; \Delta \rightarrow [G[t/x]]}{\Xi : \Theta; A; \Delta \rightarrow [\exists x.G]} \exists_r \\
\\
\frac{\Xi_1 : \Theta; A; \Delta \rightarrow [G_1] \quad \Xi_2 : \Theta; A; \Delta \rightarrow [G_2]}{\text{split}(\Xi_1, \Xi_2) : \Theta; A; \Delta \rightarrow [G_1 \wedge G_2]} \wedge_r \\
\\
\frac{\Xi : \Theta; A; [\Delta] \rightarrow K : \triangleright \text{kb}_{\mathcal{Q}} A}{\text{toSays}_{\mathcal{L}}(\Xi) : \Theta; A; \Delta \rightarrow [K : \triangleright \text{kb}_{\mathcal{Q}} A]} G \Rightarrow : \triangleright_l \quad \frac{\Psi : \Theta \mid_{\mathcal{Q}}^* \rightarrow A}{\text{fol}(\Psi) : \Theta; A; \Delta \rightarrow [K : \triangleright \text{kb}_{\mathcal{Q}} A]} : \triangleright_r + \text{kb}_r + \text{kb}_l \\
\\
\frac{\Xi : \Theta, \text{kb}_{\mathcal{Q}} A; A; [\Delta] \rightarrow K : \triangleright \text{kb}_{\mathcal{Q}'} A'}{\Xi : \Theta; A; [\Delta, K : \triangleright \text{kb}_{\mathcal{Q}} A] \rightarrow K : \triangleright \text{kb}_{\mathcal{Q}'} A'} : \triangleright_l \\
\\
\frac{\Xi : \Theta; [A]; \Delta^\dagger \rightarrow K : \triangleright \text{kb}_{\mathcal{Q}} A}{\text{toAtt}(\Xi) : \Theta; A; [\Delta^\dagger] \rightarrow K : \triangleright \text{kb}_{\mathcal{Q}} A} : \triangleright_l \Rightarrow \text{att} \quad \frac{\Xi : \Theta; A; \Delta^\dagger \rightarrow [K : \triangleright \text{kb}_{\mathcal{Q}} A]}{\text{toGoal}(\Xi) : \Theta; A; [\Delta^\dagger] \rightarrow K : \triangleright \text{kb}_{\mathcal{Q}} A} : \triangleright_l \Rightarrow G \\
\\
\frac{\Xi' : \Theta; A; \Delta \rightarrow [G\sigma] \quad \Xi'' : \Theta; A; [\Delta, K : \triangleright \text{kb}_{\mathcal{Q}} A\sigma] \rightarrow K' : \triangleright \text{kb}_{\mathcal{Q}'} A'}{\Xi_1 : \Theta \mid_{\mathcal{Q}_1}; \cdot \rightarrow [K_1 : \triangleright A_1\sigma] \quad \cdots \quad \Xi_n : \Theta \mid_{\mathcal{Q}_n}; \cdot \rightarrow [K_n : \triangleright A_n\sigma]} \text{att} \\
\text{att}(i, \sigma, [\Xi_1, \dots, \Xi_n], \Xi', \Xi'') : \\
\Theta; [A, i : \forall X. (\text{kb}_{\mathcal{Q}_1}(K_1 : \triangleright A_1) \wedge \cdots \wedge \text{kb}_{\mathcal{Q}_n}(K_n : \triangleright A_n) \wedge G \supset K : \triangleright \text{kb}_{\mathcal{Q}} A)]; \Delta \rightarrow K' : \triangleright \text{kb}_{\mathcal{Q}'} A' \\
\\
\frac{\text{evidence}_{\mathcal{K}}(E, A)}{\text{ext}(E) : \Theta; \cdot; \cdot \rightarrow [K : \triangleright A]} \text{ext} \quad \frac{\Psi : \Theta^* \rightarrow A}{\text{fol}(\Psi) : \Theta; \cdot; \cdot \rightarrow [K : \triangleright A]} : \triangleright_r + \text{kb}_l
\end{array}$$

Fig. 3. $\text{CL}_{\mathcal{P}}^{\mathcal{Q}}$ – $\text{CL}_{\mathcal{P}}$ kernel for verifying $\text{CL}_{\mathcal{P}}$ proof certificates of Cyberlogic programs. Δ^\dagger is such that for all $K' : \triangleright \text{kb}'_{\mathcal{Q}'} A' \in \Delta^\dagger$, $K' \neq K$ and $\Theta^* = \{F \mid \text{kb}_{\mathcal{Q}} F \in \Theta\}$.

of proof certificates for a proof system \mathcal{S} relies on two parts: (1) a translation of \mathcal{S} 's formulas into LKF or LJF formulas; and (2) a correspondence of \mathcal{S} proofs (or proof steps) to LKF or LJF proof steps. Given these two elements, a proof certificate for a proof of F in \mathcal{S} consists of a predicate which guides a proof of F 's translation in LKF or LJF. The following proof formats can be checked in FPC: resolution, λ -terms, Horn clauses, Frege proofs, matings, tableaux, etc.

Defining LKF or LJF FPCs for cyberlogic is challenging due to the modalities $:\triangleright$ and kb , and digital evidences. LKF has been used to check proofs in modal logics [20], but the translation of modal formulas into LK formulas used the modalities' semantic definition. Instead, we propose a modular $\text{CL}_{\mathcal{P}}$ kernel which allows facts derived from knowledge bases or external evidence to be checked by the appropriate engine or entity.

The $\text{CL}_{\mathcal{P}}$ kernel $\text{CL}_{\mathcal{P}}^{\mathcal{Q}}$ (Figure 3) is constructed by augmenting sequents with a certificate Ξ (a term indicating how the proof must proceed) and indices for the

formulas in Λ . A certificate for a proof of $\Theta; \Lambda; \Delta \longrightarrow G$ is $\Xi : \Theta; \Lambda_I; \Delta \longrightarrow G$, where Ξ is a term built from the predicates used in CL_P^a , and Λ_I is a mapping from indices to formulas in Λ . The indices are used in Ξ . The checking of a cyberlogic sequent $\Theta; \Lambda; \Delta \longrightarrow G$ with certificate Ξ starts from the sequent $\Xi : \Theta; \Lambda_I; \Delta \longrightarrow [G]$. Certificates denoted by the letter Ψ can represent proofs in other formalisms and may be checked by another engine. The predicates in Ξ are used for the following purposes during a derivation in CL_P^a .

First of all, they indicate how the proof should continue when there are multiple choices. For example, if the sequent is of the form $\Theta; \Lambda; \Delta \longrightarrow [K : \triangleright \text{kb}_Q A]$, then Ξ must be one of $\text{toSays}_L(-)$ or $\text{fol}(-)$, indicating whether to work on \triangleright modalities on the left, or finish the proof with first-order reasoning, respectively.

Secondly, certificates relay information at the appropriate moment. For example, $\text{split}(-, -)$ contains the certificates for each of the branches on a splitting rule, and $\text{ext}(-)$ includes an external evidence for proposition A . Note that there is no certificate for \exists_R since these can be instantiated with meta-variables, and unification can be verified when the proof is completed.

The certificate for rule **att** is more interesting. It includes the index i of the attestation formula to be decomposed, the substitution σ for the \forall quantifier, and certificates for each premise. Note that each Ξ_1, \dots, Ξ_n must be $\text{ext}(-)$ or $\text{fol}(-)$.

Example 5. Consider Example 4, and let the indices of the formulas be their labels: **main**, **cont**, and **bankStmt**. The certificate for a proof that alice can get a visa is $\Xi : \text{cont}; \text{main}, \text{bankStmt}; \cdot \longrightarrow \text{cons} : \triangleright \text{kb}_{\{\text{cons}\}} \text{issVisa}(\text{alice}, \text{doc}, \text{visa})$. Where Ξ is:

$$\text{att}(\text{main}, \{\text{id} \mapsto \text{alice}, \text{Doc} \mapsto \text{doc}, \forall \mapsto \text{visa}\}, [\text{fol}(\Psi_{\text{validVisit}}), \text{fol}(\Psi_{\text{prepVisa}})], \Xi_G, \Xi_0)$$

The certificates $\Psi_{\text{validVisit}}$ and Ψ_{prepVisa} are first-order logic proof certificates from derivations using the consulate's own knowledge base.

Certificate Ξ_0 corresponds to **att**'s premise where the conclusion of **main** is added to the context. This branch can be closed by removing the modalities, so $\Xi_0 = \text{toGoal}(\text{fol}(\text{id}))$, where **id** is a first-order logic directive to close the proof.

Certificate Ξ_G guides the proof of the new goal:

$$\text{cons} : \triangleright \text{kb}_{\{\text{cons}\}}(\text{suffin}(\text{doc})) \wedge \text{police} : \triangleright \text{kb}_{\{\text{police}\}}(\text{noCrimeRec}(\text{alice}))$$

and thus $\Xi_G = \text{split}(\Xi_{\text{fin}}, \Xi_{\text{crime}})$. Ξ_{fin} depends on how **cons** decides to check for sufficient funds. It could rely on the **bank** and use the attestation formula **bankStmt**, in which case Ξ_{fin} has the shape

$$\text{toSays}_L(\text{toAtt}(\text{att}(\text{bankStmt}, -, -, -, -)))$$

Or it could use **cont** from its knowledge base, in which case Ξ_{fin} would be $\text{fol}(-)$.

5 Related Work

Attestation logics have been proposed for the specification of policies of several distributed systems [15,22,16,6,29,3]. We have been inspired by some of this work

in the design of Cyberlogic. Actually, Cyberlogic was proposed some decades ago [29,6], but until now its proof theory had not been carefully investigated. In particular, there were no statements on cut-elimination. Additionally, we have been inspired by the previous works on authorization logics [15,22,16] to extend Cyberlogic with knowledge operators.

The main contribution of our work is the study of proof search and proof certificates for attestation logics with knowledge operators.

In previous work [15] in intuitionistic authorization logic, knowledge was restricted to one principal. As demonstrated in Example 1, allowing for multiple principal knowledge databases ensures collaboration in reasoning.

Proof search for attestation logics is not adequately addressed in the literature. Either the proposed proof systems are Hilbert-style [3,2,18] which do not enjoy the sub-formula property and therefore are not suitable for proof search, or they are sequent calculus proof system, but not focused proof systems [15,22,29,6,17]. [15] only speculates that logic programming languages can be used to carry out proof search for fragments of attestation logic. We confirm this speculation with the definition of Cyberlogic programs.

Our main inspiration for proof certificate is the work on foundational proof certificates [10]. However, the existing work did not consider proof certificates for attestation logics. Closer to our objective is the work of Libal and Volpe [20], which define proof certificates for modal logics by encoding (the semantics of) these logics in LKF. Our work instead proposes proof certificates directly in Cyberlogic. This means that we are able to capitalize on rules, such as attestation rules, to build more compact certificates. Another difference is that our proof certificates may contain (pointers to) extra-logical evidence.

Cyberlogic has been formalized in Coq [12], encoding evidential transactions for Schengen Visa applications. Our approach is different in that it lays a proof theoretic foundation to Cyberlogic. In particular, proof search is formally justified as well as the representation of Cyberlogic proofs as FPCs.

Logic programming engines, such as ETB [11], have been proposed for programming ETs. However, these engines do not (yet) support attestations, such as $K \triangleright F$, local knowledge, such as $\text{kb}_Q F$, nor the use of digital certificates. We believe that this work can greatly profit from the foundations laid by this paper.

Finally, works [16,7] propose the use of evidence for authorization. Specifically, [17] show that a fragment of their system is decidable in linear time. It would be interesting to investigate how this fragment relates to Cyberlogic programs, and whether proof certificates as defined in this work can be applied to the decidable fragment. This is left for future work.

6 Conclusions

This paper lays the proof-theoretic foundations for Cyberlogic, an attestation logic for evidential transactions, and refine Cyberlogic with epistemic modalities. We identify a fragment of Cyberlogic, Cyberlogic programs, and propose a proof system similar to focused proof systems for enabling sound and complete proof

search. The necessary permutations for completeness rely on the careful interplay between attestation, \triangleright , and knowledge modalities, $\text{kb}_{\mathcal{Q}}$. We then propose a concise proof certificate format for proofs of Cyberlogic programs.

This paper is the first step for a framework enabling evidential transactions that we are currently implementing. In particular, we are extending Distributed Datalog engines available in [11] to support Cyberlogic. Moreover, we are integrating such engines with PKI infrastructure, available in, for example, Distributed Ledger Technologies. This means that evidence, both in the form of digital evidence and logical derivations in the form of FPCs, can be stored and audited through the Ledger Technologies.

We are currently investigating extensions to Cyberlogic programs to include other modalities, such as temporal and epistemic [24,13] while still preserving its good proof search properties.

We have also started to study conditions for when two attestation rules can be introduced in any order. If two clauses can be introduced in any order, then they can also be introduced in parallel. Therefore, this would provide proof-theoretic justification for proof search optimization. This could be used, for example, for proposing refinements to dependency graphs used for evaluating distributed logic programming [23] which take principals into account. These results will impact the maintenance of evidential transactions, whose applications can have important consequences to, *e.g.*, certification in automotive and avionics domains.

Acknowledgment: We would like to thank Dian Balta, Natarajan Shankar and Tewodros Beyene for useful discussions and valuable feedback on earlier versions of this paper.

References

1. L-framework. <https://carlosolarte.github.io/L-framework/>, accessed on 03-01-2021
2. Abadi, M., Burrows, M., Lampson, B., Plotkin, G.: A calculus for access control in distributed systems. vol. 15, pp. 1–23 (01 1995). https://doi.org/10.1007/3-540-46766-1_1
3. Abadi, M.: Logic in access control. In: 18th IEEE Symposium on Logic in Computer Science (LICS) Proceedings. p. 228. IEEE Computer Society (2003). <https://doi.org/10.1109/LICS.2003.1210062>
4. Andreoli, J.M.: Logic Programming with Focusing Proofs in Linear Logic. *Journal of Logic and Computation* **2**(3), 297–347 (1992). <https://doi.org/10.1093/logcom/2.3.297>
5. Bernat, V.: First-Order Cyberlogic Hereditary Harrop Logic. Tech. rep., SRI International (2006), <http://www.lsv.ens-cachan.fr/Publis/PAPERS/PS/Bernat-cyberlogic1.ps>
6. Bernat, V., Ruess, H., Shankar, N.: First-order Cyberlogic. Technical Report CSL-SRI-04-03, SRI International Computer Science Laboratory (2004)
7. Blass, A., Gurevich, Y., Moskal, M., Neeman, I.: Evidential authorization. In: Nanz, S. (ed.) *The Future of Software Engineering*. pp. 73–99. Springer (2010). https://doi.org/10.1007/978-3-642-15187-3_5

8. Chaudhuri, K., Pfenning, F., Price, G.: A Logical Characterization of Forward and Backward Chaining in the Inverse Method. In: Automated Reasoning, Third International Joint Conference, IJCAR, Proceedings. pp. 97–111 (2006). https://doi.org/10.1007/11814771_9
9. Chihani, Z., Miller, D., Renaud, F.: Foundational Proof Certificates in First-Order Logic. In: Bonacina, M.P. (ed.) Automated Deduction – CADE-24. pp. 162–177. Springer Berlin Heidelberg (2013)
10. Chihani, Z., Miller, D., Renaud, F.: A Semantic Framework for Proof Evidence. *J. Autom. Reasoning* **59**(3), 287–330 (2017). <https://doi.org/10.1007/s10817-016-9380-6>
11. Cruanes, S., Hamon, G., Owre, S., Shankar, N.: Tool Integration with the Evidential Tool Bus. In: Verification, Model Checking, and Abstract Interpretation, 14th International Conference, VMCAI. Proceedings. pp. 275–294 (2013). https://doi.org/10.1007/978-3-642-35873-9_18
12. Dargaye, Z., Kirchner, F., Tucci-Piergiovanni, S., Gürcan, O.: Towards Secure and Trusted-by-Design Smart Contracts. In: JFLA (2018)
13. DeYoung, H., Garg, D., Pfenning, F.: An authorization logic with explicit time. In: Proceedings of the 21st IEEE Computer Security Foundations Symposium, CSF. pp. 133–145. IEEE Computer Society (2008). <https://doi.org/10.1109/CSF.2008.15>
14. Fairtlough, M., Mendler, M.: Propositional Lax Logic. *Inf. Comput.* **137**(1), 1–33 (1997). <https://doi.org/10.1006/inco.1997.2627>
15. Garg, D., Bauer, L., Bowers, K.D., Pfenning, F., Reiter, M.K.: A Linear Logic of Authorization and Knowledge. In: Computer Security - ESORICS 2006, 11th European Symposium on Research in Computer Security, Proceedings. pp. 297–312 (2006). https://doi.org/10.1007/11863908_19
16. Gurevich, Y., Neeman, I.: DKAL: Distributed-Knowledge Authorization Language. Tech. Rep. MSR-TR-2008-09 (January 2008), <https://www.microsoft.com/en-us/research/publication/191tr-dkal-distributed-knowledge-authorization-language/>
17. Gurevich, Y., Neeman, I.: Dkal 2 - a simplified and improved authorization language. Tech. Rep. MSR-TR-2009-11 (February 2009), <https://www.microsoft.com/en-us/research/publication/200-dkal-2-a-simplified-and-improved-authorization-language/>
18. Gurevich, Y., Neeman, I.: Logic of infons: The propositional case. *ACM Trans. Comput. Log.* **12**, 9 (01 2011)
19. Liang, C., Miller, D.: Focusing and polarization in linear, intuitionistic, and classical logics. *Theor. Comput. Sci.* **410**(46), 4747–4768 (2009). <https://doi.org/10.1016/j.tcs.2009.07.041>
20. Libal, T., Volpe, M.: A general proof certification framework for modal logic. *Math. Struct. Comput. Sci.* **29**(8), 1344–1378 (2019). <https://doi.org/10.1017/S0960129518000440>
21. Miller, D.: Foundational Proof Certificates. In: Delahaye, D., Paleo, B.W. (eds.) All about Proofs, Proofs for All, All about Proofs, Proofs for All, vol. Mathematical Logic and Foundations, 55, pp. 150–163. College Publications (2015), <https://hal.inria.fr/hal-01239733>
22. Nigam, V.: A framework for linear authorization logics. *Theor. Comput. Sci.* **536**, 21–41 (2014). <https://doi.org/10.1016/j.tcs.2014.02.018>
23. Nigam, V., Jia, L., Loo, B.T., Scedrov, A.: Maintaining distributed logic programs incrementally. *Computer Languages, Systems & Structures* **38**(2), 158–180 (2012). <https://doi.org/10.1016/j.cl.2012.02.001>

24. Nigam, V., Olarte, C., Pimentel, E.: A General Proof System for Modalities in Concurrent Constraint Programming. In: CONCUR. LNCS, vol. 8052, pp. 410–424. Springer (2013)
25. Nigam, V., Pimentel, E., Reis, G.: An extended framework for specifying and reasoning about proof systems. J. Log. Comput. **26**(2), 539–576 (2016). <https://doi.org/10.1093/logcom/exu029>
26. Olarte, C., Pimentel, E., Rocha, C.: Proving Structural Properties of Sequent Systems in Rewriting Logic. In: Rusu, V. (ed.) Rewriting Logic and Its Applications. pp. 115–135. Springer International Publishing (2018)
27. Pfenning, F., Davies, R.: A judgmental reconstruction of modal logic. Mathematical Structures in Computer Science **11**(4), 511–540 (2001). <https://doi.org/10.1017/S0960129501003322>
28. Reis, G.: Observations about the proof theory of cyberlogic. <http://www.gisellereis.com/papers/cyberlogic-report.pdf> (2019)
29. Ruess, H., Shankar, N.: Introducing Cyberlogic. In: Proceedings of the 3rd Annual High Confidence Software and Systems Conference (2003)
30. Troelstra, A.S., Schwichtenberg, H.: Basic Proof Theory. Cambridge University Press (1996)

A Proofs

A.1 Cut-elimination

Proof. (Sketch) The proof follows the usual Gentzen strategy of reducing the cuts’ grade and rank. The interesting cases are rank reduction over kb rules.

In the case of kb_l , contraction of the main formula is needed for the permutation to work. If this was not the case, we could not conclude $\Gamma, A \rightarrow G$ from $\Gamma, \text{kb}_Q A \rightarrow G$. The transformations are:

$$\begin{array}{c}
 \frac{\frac{\Gamma, \text{kb}_Q A, A \xrightarrow{\varphi_1} C}{\Gamma, \text{kb}_Q A \xrightarrow{\varphi_1} C} \text{kb}_l \quad \frac{\Gamma, \text{kb}_Q A, C \xrightarrow{\varphi_2} G}{\Gamma, \text{kb}_Q A \xrightarrow{\varphi_2} G} \text{cut}}{\Gamma, \text{kb}_Q A \xrightarrow{\varphi_1} G} \text{cut} \quad \rightsquigarrow \quad \frac{\frac{\Gamma, \text{kb}_Q A, A \xrightarrow{\varphi_1} C \quad \frac{\Gamma, \text{kb}_Q A, A, C \xrightarrow{\varphi_2 + \text{weakening}} G}{\Gamma, \text{kb}_Q A, A, C \xrightarrow{\varphi_2} G} \text{cut}}{\Gamma, \text{kb}_Q A, A \xrightarrow{\varphi_1} G} \text{kb}_l}{\Gamma, \text{kb}_Q A \xrightarrow{\varphi_1} G} \text{cut} \\
 \\
 \frac{\frac{\Gamma, \text{kb}_Q A \xrightarrow{\varphi_1} C \quad \frac{\Gamma, \text{kb}_Q A, A, C \xrightarrow{\varphi_2} G}{\Gamma, \text{kb}_Q A, C \xrightarrow{\varphi_2} G} \text{kb}_l}{\Gamma, \text{kb}_Q A \xrightarrow{\varphi_1} G} \text{cut}}{\Gamma, \text{kb}_Q A \xrightarrow{\varphi_1} G} \text{cut} \quad \rightsquigarrow \quad \frac{\frac{\Gamma, \text{kb}_Q A, A \xrightarrow{\varphi_1 + \text{weakening}} C \quad \frac{\Gamma, \text{kb}_Q A, A, C \xrightarrow{\varphi_2} G}{\Gamma, \text{kb}_Q A, A, C \xrightarrow{\varphi_2} G} \text{cut}}{\Gamma, \text{kb}_Q A, A \xrightarrow{\varphi_1} G} \text{kb}_l}{\Gamma, \text{kb}_Q A \xrightarrow{\varphi_1} G} \text{cut}
 \end{array}$$

The other interesting case is when we need to permute a cut over a kb_r rule on the right branch:

$$\frac{\frac{\Gamma \xrightarrow{\varphi_1} C \quad \frac{(\Gamma, C) \upharpoonright_{Q_i} \xrightarrow{\varphi_2} G}{\overline{\Gamma}, C \xrightarrow{\varphi_2} \text{kb}_{Q_i} G} \text{kb}_r}{\Gamma \xrightarrow{\varphi_1} \text{kb}_{Q_i} G} \text{cut}}{\Gamma \xrightarrow{\varphi_1} \text{kb}_{Q_i} G} \text{cut}$$

There are two cases to consider:

1. $C \equiv \text{kb}_{\mathcal{Q}_j} C'$ and $\mathcal{Q}_i \preceq \mathcal{Q}_j$: in this case, we can permute the cut over rules on φ_1 (left rules except $:\triangleright_L$, which is never applicable) until it is principal. This lemma can be proved by case analysis. At this point, the premise on the left branch will be $\Gamma \mid_{\mathcal{Q}_j} \rightarrow C'$. Then kb_R can be applied to the end-sequent, resulting in:

$$\frac{\frac{\Gamma \mid_{\mathcal{Q}_i} \xrightarrow{\varphi'_1} \text{kb}_{\mathcal{Q}_j} C' \quad \Gamma \mid_{\mathcal{Q}_i, \text{kb}_{\mathcal{Q}_j} C'} \rightarrow G}{\Gamma \mid_{\mathcal{Q}_i} \rightarrow G} \text{ cut}}{\Gamma \rightarrow \text{kb}_{\mathcal{Q}_i} G} \text{ kb}_r$$

The proof φ'_2 is exactly φ_2 , since $(\Gamma, \text{kb}_{\mathcal{Q}_j} C') \mid_{\mathcal{Q}_i} \equiv \Gamma \mid_{\mathcal{Q}_i, \text{kb}_{\mathcal{Q}_j} C'}$ when $\mathcal{Q}_i \preceq \mathcal{Q}_j$. The proof φ'_1 is obtained from the proof of $\Gamma \mid_{\mathcal{Q}_j} \rightarrow C'$, since $\Gamma \mid_{\mathcal{Q}_j} \subseteq \Gamma \mid_{\mathcal{Q}_i}$ when $\mathcal{K}_i \preceq \mathcal{Q}_j$.

2. $C \not\equiv \text{kb}_{\mathcal{Q}_j} C'$ or $\mathcal{Q}_i \not\preceq \mathcal{Q}_j$: in this case $C \notin (\Gamma, C) \mid_{\mathcal{Q}_i}$, so kb_r can be applied directly to the end-sequent, and the cut can be removed.

□