

Modal Logic for Distributed Trust*

Niels Voorneveld

Peeter Laud

Cybernetica AS

March 24, 2026

Abstract

We propose a method for reasoning about trust in multi-agent systems, specifying a language for describing communication protocols and making trust assumptions and derivations. This is given an interpretation in a modal logic for describing the beliefs and communications of agents in a network. We define how information in the network can be shared via forwarding, and how trust between agents can be generalized to trust across networks.

We give specifications for the modal logic which can be readily adapted into a lambda calculus of proofs. We show that by nesting modalities, we can describe chains of communication between agents, and establish suitable notions of trust for such chains. We see how this can be applied to trust models in public key infrastructures, as well as other interaction protocols in distributed systems.

1 Introduction

When protocols are distributed across multiple parties, you must consider those parties in their proper context in order to make guarantees. Who has access to what information? Who is claiming that the information is correct. These are factors necessary for determining both whether something is correct, and whether someone has the necessary knowledge to learn about this correctness. Properties and guarantees we want to establish within distributed contexts involve claims regarding credentials, and correctness of distributed processes according to some specifications. Cooperation is needed to prove such guarantees, which fundamentally requires the sharing of statements and the consideration of other peoples statements.

There are two sides to cooperation. The first is to facilitate the building of strong claims by combining knowledge. This may not necessitate full transparency, and entities can choose what and with whom they share their claims, in order to cut costs or keep confidential information hidden. Such entities would like to learn the effectivity of sharing facts, and what others can do with them. They can prove that they can convince others, or that collaborators can use their facts effectively. Though the partial *hiding* of information is assumed, this paper is not about *privacy*. No guarantees are made that information cannot leak. It is simply specified in which contexts claims are officially shared, and proofs are built only on such specified assumptions. Still, some basic properties of privacy across communication channels can be asserted, and one can reason about consequences of leaked information.

The flipside to cooperation is *responsibility* and *accountability*. Sharing information is based on mutual trust, and entities making false claims can break this trust. As such, entities should take care not to share false facts, and careful entities should stick to their expertise. Entities are not just accountable for the claims they make, but also for the logical consequences of their claims. Implied lies are still lies. As such, it is important to specify exactly what information an entity has access to. This is another reason why the destination of claims matters, as it tells us who has received the information. Entities should not be allowed to ignore information shared with them. They should consider consequences of received claims, and be accountable for them. That said, it is important to distinguish between the claim as a whole, and the content of the claim. It is ok to accept that a claim has been made, yet to deny that the content of the claim is true. Whether or not to accept the contents of a claim as true is a matter of trust.

Trust is the cornerstone of this work, and has been studied from many points of view. If you trust a claim, you accept its content as true. Trust is formulated in a conditional way: it does not matter whether the claim has actually been made, trust states that if a certain claim is made by a certain entity, then it is true. Making trust conditional facilitates cooperation. Entities may declare their trust in others. This

*This research has been supported by Estonian Research Council, grant No. PRG1780.

can then be used by other entities, which could incorporate it in their model in order to make guarantees. For instance, a shop may declare trust in an authority in its ability to authenticate credentials. A buyer then knows it can prove their credentials to the shop via this authority. Such interacting trust relations can then be built up in order to prove guarantees of larger authentication networks and protocols. Note however, that if an entity receives a claim which they have previously declared as trustworthy, they are now accountable for the consequences of the claim. Hence, *declared trust*, which can be interpreted as *vouching* for the expertise of someone, requires one to accept accountability for a trusted entity’s statements regarding that area of expertise.

1.1 Constructive Modal Logic

Constructive logics have many useful applications in computer science and security information systems. Such logics force one to be more specific regarding how a property is satisfied. Proving a disjunction requires one to specify explicitly which of the two options are true (or give a decision procedure as such given the assumptions). Moreover, it disallows double negation elimination: “it cannot fail” is not a sufficient enough reason of why something is guaranteed to work.

In short, constructive logic requires one to build an explicit guarantee using a set of assumptions. This also increases accountability: the precise reason for the guarantee is embedded within the proof itself, allowing one to trace back which assumptions are used, and how a guarantee is satisfied. If the guarantee is later shown to be false, then it is easier to identify the entity that made a mistake.

To express both the claims being shared, and the internal perspectives of entities, we use modalities. Guided by our pursuit of accountable claims, we use modalities satisfying the axiom K and the necessity rule, which are considered standard axioms in epistemic (knowledge-based) logics. We moreover consider a collection of well-behaved axioms relating the many perspectives. Though many varieties can be considered, we start with a basic model of *belief* \mathcal{B} and *interaction* \mathcal{I} to mark internal and communicated knowledge.

A core building block of our logic is the consideration of modal *splitting axioms*, which declare that $\mathcal{M}A \Rightarrow \mathcal{N}\mathcal{R}A$ holds for any statement A . These express how information A from modality \mathcal{M} can be adopted by modality \mathcal{N} . This latter domain does not necessarily accept the statement as true. Instead, they consider it as flagged by modality \mathcal{R} , and can reason with this statement. A specific instance of such an axiom is an *awareness axiom* of the form $\mathcal{M}A \Rightarrow \mathcal{N}\mathcal{M}A$ for all A , where we say \mathcal{N} is aware of all statements A provable by \mathcal{M} . This would hold for instance if \mathcal{M} represents a communication channel which \mathcal{N} has access to.

Another aspect to consider is the *decidability* of the logic; does an algorithm exist which shows whether a guarantee is true given certain assumptions? Decidability is useful in order to increase accountability: it enables entities to compute and hence be aware of certain consequences of their claims. Checking for particularly significant consequences of one’s claims should be part of an entity’s due diligence. In order to facilitate decidability, we limit ourselves to a specified set of modal axioms, as well as avoiding the use of polymorphism and quantifiers.

1.2 Distributed Trust and Decentralized Identity

One of the main motivations for this paper is its applications to decentralized identities. We believe that the emergence of *self-sovereign identity* [45] acutely necessitates the ability to assign truth values to complex statements about and including trust. An entity’s sovereignty over their identity means the proliferation of identifiers that that entity intends to use in different contexts. These identifiers are presumably created by the entity themselves; some authority may or may not have bound them to identifiers in a more generally known namespace (e.g. the names of the residents of some country).

The statements that the various authorities make about an entity will be bound to the various identifiers created by that entity (as expressed by e.g. the “Privacy and minimal disclosure” principle of [38]). While it should generally be the responsibility of that entity to make sure that the claims they want to present together are bound to the same identifier, we expect the reality to not be so simple. When a relying party (i.e. a party that makes decisions on the basis of the claims it receives about an entity) receives a number of statements pertaining to a number of different identifiers, as well as some statements that relate these identifiers to each other (as per the “Delegation” principle of [38]), it becomes highly non-trivial to find out which of these claims can be combined with each other. This paper provides a system for relying parties to make such inferences.

The same question — what can be reasonably and confidently derived for which identifier? — is also something that the individual entities themselves want to find the answers to. Our logical approach

provides them with tools to make these inferences *from the point of view of some other party*, allowing them to forecast the decisions made by the relying parties. Indeed, the decisions can be derived for various parties, considering what we know about their trust relationships; these decisions may be combined with each other and considered from the point of view of different parties.

1.3 Misplaced Trust in Certificate Authorities

A motivating example we focus on is that of networks for public key validation. A certificate authority (CA) is an entity in such networks which makes publicly verifiable statements that bind public keys to entities. These statements are the *certificates*. A CA is trusted to be able to properly verify that an entity is in control of a public key. We consult CAs to find out an entity's public key. There are various mechanisms to find out whether an entity is a CA, e.g. a higher-level CA may issue a certificate to another entity stating that it is a (sub-)CA, too. We use such statements in hierarchical public-key infrastructures (PKI), where the knowledge of the public key of a *root* CA allows us to verify a *certificate chain* that ultimately binds a public key to an (end-)entity.

Of course, when an entity is declared to be a CA, then the declarer has to perform due diligence on that entity's ability to be a CA. Sometimes these declarations are made by CAs that should not be declaring other entities to be a CA. This creates confusion when verifying certificate chains, because the last certificate in that chain may be issued by a CA that should not be seen as a CA. Several incidents (see [36], subsections A.1.11, A.1.13, A.1.14, A.1.17) have taken place due to such misplaced trust. Even more (see subsections A.1.10 and A.1.25 in [36]) have happened due to other usage restrictions that have been absent in the issued certificate.

We employ our logic as a tool for specifying the function and duties of entities in distributed networks, like CAs in public key infrastructures. This way, they can be held accountable for mistakes made. Given a formal constructive proof of a property, if the property is later shown to be false, we can trace back the dependencies and find who is at fault, ensuring accountability. Moreover, definitions of trust can be fine-tuned to safeguard against certain errors and attacks. For instance, threshold trust can be employed to deal with issues of unresponsiveness or corruption of CAs.

1.4 Related Work

There exist a number of calculi and logics for defining and computing trust in distributed networks in general and PKIs in particular. The paper [44] considers a modal logic with belief modality, with the description of messages modeled by the underlying Kripke style semantics instead of explicitly described in the logic as we do. Non-modal approaches include a predicate calculus [19] studying a generalization of the usual graph description of PKI networks, and [6] which focuses on modeling the concurrent properties of such networks. Especially interesting are [34, 48] which both study trust statements and inferences which can be made using them, similar to our treatment of trust for PKI, which we describe within our logic. The goals of these calculi are similar to our formalisms — turning statements made by trusted entities into beliefs of users. Though they do not support *higher-order* reasoning, with messages referring to other messages, which our modal logic does consider.

We use a modal logic to describe and give meaning to trust in distributed systems. Modal logics have had many applications in computer science due to their flexibility and versatility, which shows their usefulness as a verification tool. Examples include modeling programming features such as variable scoping [42], and reasoning about cryptographic processes [20]. Using a modal logic to describe trust has been investigated in several different contexts. We ground our work on the definition of *cautious trust* by Liau [41], given in terms of belief and communication. This has been further investigated theoretically in a plethora of other works [16, 18, 40]. It has been shown that such trust definitions could be extended to consider probability [30, 39] and time [5].

We specifically look at *intuitionistic* modal logic [52], in order to keep proofs constructive and expressible by various lambda calculi [4, 9, 26] via the Curry-Howard correspondence. We maintain that calculi such as these are one of the requirements for a trust logic, due to the support it can give for the derivation of trust in software libraries.

1.5 Comparison to Authorization and Access Logics

The work has a strong connection to *authorization and access logics* (AL) [7, 8, 13, 22]. Though the basic motivation and foundations of the formalisms are distinct, they are not wholly incompatible. ALs are

about certain commands, and whether they were made by an authorized entity. This entails authenticating proofs of authority. In this paper, we do not focus on specific commands and control statements as such, but instead focus on the sharing and evaluating of information across distributed networks.

Many works mainly consider two primitives, a modality S_a for what a principal a *says*, and a primitive formula stating that a *speaks for* b . This logical operation already existed in foundational work [13], which focused on information regarding who owns which keys, and used in particular descriptions of what information people see, believe and control. Soon after, in [3], a more thorough investigation of the *says* constructor was done, studying a lattice structure and defining *control* as the logical affirmation: saying a statement makes that statement true.

Many different axiomatizations of S_a were considered, with order structure, consistency $S_a\perp \Rightarrow \perp$ (axiom D) and idempotency $S_aA \Leftrightarrow S_aS_aA$ occurring in [3] already. Consistently throughout further work [2, 21, 24] the axioms $A \Rightarrow S_aA$ and $S_aS_aA \Rightarrow S_aA$ were assumed, with some (mainly [2]) assuming even $S_aS_bA \Rightarrow S_bS_aA$. It is perhaps this last axiom which highlights a fundamental feature of authorization logic different from our approach: in ALs they are concerned about who has authority, not about who said something.

The axiom $A \Rightarrow S_aA$ signifies that whoever is reasoning about statements from a may complement any of a 's utterances with other things which are true. This comes from the perspective that a may want access or authority over something, and hence is ok with others completing their proof of access with additional details which may help them. This is in stark opposition to the aim in this paper, where a is accountable for their statements, and would not want others to put words into their mouth. Additionally, we would like a to be able to reason about consequences of their claims, and hold a precise model of how other people regard them. As such, they should only need to be accountable for knowledge they are aware of.

The paper [1] considers a variation in which $A \Rightarrow S_aA$ is replaced by $S_bA \Rightarrow S_aS_bA$, meaning one may adopt claims made by others. Utterances by b are considered as universally shared, which would make a aware of them. We would like to disallow this as well in general, as communications are not necessarily shared with everyone. Other work [31–33] allow $S_aA \Rightarrow S_aS_aA$ as we do when considering belief, though still have $S_aS_aA \Rightarrow S_aA$.

We drop axioms of the form $\mathcal{M}\mathcal{N}A \Rightarrow \mathcal{R}A$ mostly in order to keep decidability in the absence of S4's unit axioms. We predominantly focus on how information is shared and adopted, which entails axioms of the form $\mathcal{M}A \Rightarrow \mathcal{N}A$ and $\mathcal{M}A \Rightarrow \mathcal{N}\mathcal{R}A$. These include $\mathcal{B}_aA \Rightarrow \mathcal{B}_a\mathcal{B}_aA$, as well as agents adopting information they have access to, such as $\mathcal{I}_{a\rightarrow b}A \Rightarrow \mathcal{B}_a\mathcal{I}_{a\rightarrow b}A$. We moreover do not consider axiom D stating that $\mathcal{M}\perp \Rightarrow \perp$ [35, 49], since we do not wish the inconsistency of one party's statements to necessarily imply the inconsistency of the whole.

Certain foundational works consider lattice structures on their agents [2, 3] which give an axiomatic order on S_a , as well as meet and join axioms. We adopt the order structure to define how belief is shared, which allows us to model *broadcasting*. Other works have this order as dynamic statements in the logic, with derivable *speaks for* formulas. We do not consider those, as they are statements of authority, not necessarily trust. Some work formulating *can* statements [7, 8] does bear resemblance to our treatment of trust. There, if a says that b can do A , and if b does it, then a says A happens. Similarly, we say that if a claims that b 's opinion on A is valid (trust), and if b claims A , a is responsible for claim A .

The DKAL language [27, 29] considers a similar trust statement as primitive in a different context. In terms of treatment of trust itself, DKAL is perhaps the closest to our logic, with our derivations of validity (Sec. 4) having similarities with their operational semantics [29]. DKAL mostly focuses on parametric modalities satisfying axiom K and Necessity, and does not consider many further axioms, such as the connections between modalities we explore in this paper. Our treatment of multi-step information sharing (Sec. 3 and 5) is more fine-grained, with proper significance given to both the senders and the *intended* recipients of messages. We are also able to handle *disjunction*, which is useful to model certain protocols. But our greatest difference lies in semantics — the Kripke semantics of our logic gives us greater confidence in our axioms, inference rules, and proof search methods. DKAL instead has a translation to Liberal Datalog [11] for deriving truth. Their follow-up DKAL 2 [28] does have a Kripke model, but lacks some of the expressivity discussed above.

We consider our logic to be a specialization of former work on trust management systems [5, 34, 48] in the direction of distributed systems with defined information access. We focus in particular on how to specify protocols for sharing, forwarding, and trusting claims across networks. We do this by adopting and modifying the theory of belief and interaction from the past [41] and forming a decidable logic framework for specifying trust and interaction protocols.

1.6 Contributions and Paper Outline

Our first contribution is the formulation of a theory to reason about communications over distributed networks, enabling us to prove the existence of trust across chains of communications. We establish a set of axioms (Figure 1) for sharing and considering information, which have not been the focus of previous work, and discuss how to reason about indirect communication and trust in the logic. We also touch upon a calculus for building proofs in the logic.

The language for sharing and considering information allows us to state, what it means for one agent to trust another one with respect to the validity of a statement. Our calculus gives us some useful inference rules for validity in multi-agent systems (Figure 3).

We extend the trust with *levels*, allowing us to model examples such as public key infrastructures (PKI) that may involve several steps of delegation. Our language will be rich enough to express these levels, and to derive certain properties of them (Theorem 1) that turn to be rich enough to concisely describe the trust relations in networks of agents.

We start in Section 2 by discussing how we describe the relevant concepts with modalities and axioms. In Section 3 we show different ways of formulating indirect communications. We discuss reasoning with trust in Section 4, and formulate a formalism for sharing information and trust in a network in Section 5. We give further examples in Sections 6 and 7. We discuss two calculi in Section 10, and a Kripke model in 11, and give final remarks in Section 12.

Related papers: Since the writing of this article, two papers have been published focusing on and expanding on particular parts of this article. One paper [50] develops the decidability of a schema of intuitionistic modal logic, to which is partially covered here, and furthermore develops how to derive relevant consequences. Another paper [51] covers how you can reason about trust thresholds.

2 Modal Reasoning

We use a constructive modal logic to specify both the perspectives of different entities, as well as claims entities make to each other. We consider the general concept of *domain*, which subsumes specified entities, groups of entities, the public domain, and other entities identified by public keys. Let \mathbb{A} be a finite set of entities, domains or agents.

Following the traditional literature on trust, we consider two modalities which we will give a specified meaning:

- For each $a \in \mathbb{A}$, the modality \mathcal{B}_a signifies what a believes. In general the statement $\mathcal{B}_a A$ means that we can prove that given presumed assumptions of agent a , we can show that a should accept A as true. We can think of it as a can verify A given their own knowledge and expertise, and given what has been communicated to and shared with them.
- For each pair $a, b \in \mathbb{A}$, the modality \mathcal{I}_{a+b} signifies interaction between a and b . The statement $\mathcal{I}_{a+b} A$ expresses that through interaction between a and b , b has claimed and communicated to a enough information (claims) to imply that statement A is true.

There is a significant difference between what one believes and what one claims to be true. Firstly, in general one would not assume that $\mathcal{I}_{a+b} A \Rightarrow \mathcal{B}_b A$ as b may be more liberal with the truth than what they themselves truly believe and act upon. On the other hand, $\mathcal{B}_b A \Rightarrow \mathcal{I}_{a+b} A$ is not necessarily true as b may not want to communicate all that they know to a .

Entities are accountable for the logical consequences of their claims on top of the claims themselves. Similarly, we assume they can make logical inferences within their own model of belief. As such, we assume each modality \mathcal{M} satisfies the axiom K and the “necessity” inference rule.

- Axiom K states that $\mathcal{M} A \Rightarrow \mathcal{M}(A \Rightarrow B) \Rightarrow \mathcal{M} B$ for any formulas A and B , signifying that if we learn that both A and $A \Rightarrow B$ hold in belief or interaction expressed by \mathcal{M} , then the consequence B should also hold there. This axiom allows us to reason about inference that can be made within \mathcal{M} as an external observer.
- Necessity states that if formula A is undeniably true (a logical tautology), then so is $\mathcal{M} A$. This means everyone should accept obviously true logical statements; those which can be proven without further assumptions. This is weaker than saying $A \Rightarrow \mathcal{M} A$, which states that if A happens to be true (e.g. given as an optional assumption), then $\mathcal{M} A$ is true. The latter implies \mathcal{M} knows all

Axiom	Name
$\vdash \mathcal{M}A \Rightarrow \mathcal{M}(A \Rightarrow B) \Rightarrow \mathcal{M}B$	Axiom K
If $\vdash A$ then $\vdash \mathcal{M}A$	Necessity
$\vdash \mathcal{B}_a A \Rightarrow \mathcal{B}_a \mathcal{B}_a A$	Self awareness
$\vdash \mathcal{I}_{a \vdash b} A \Rightarrow \mathcal{B}_a \mathcal{I}_{a \vdash b} A$	Recipient awareness
$\vdash \mathcal{I}_{a \vdash b} A \Rightarrow \mathcal{I}_{a \vdash b} \mathcal{B}_b A$	Intent of claim
If $a \sqsubseteq b$ then $\vdash \mathcal{B}_a A \Rightarrow \mathcal{B}_b A$	Belief inheritance
If $a \sqsubseteq b$ then $\vdash \mathcal{I}_{a \vdash c} A \Rightarrow \mathcal{I}_{b \vdash c} A$	Receiver inheritance
If $a \sqsubseteq b$ then $\vdash \mathcal{I}_{c \vdash a} A \Rightarrow \mathcal{I}_{c \vdash b} A$	Sender inheritance
$\vdash \Box A \Rightarrow \mathcal{M} \Box A$	Public awareness
$\vdash \Box A \Rightarrow A$	Public verifiability

Figure 1: Axiom system. \mathcal{M} ranges over modalities, a, b over agents, and A, B over formulas.

details of the current situation, whereas necessity only states \mathcal{M} knows about facts which hold in all possible situations.

On top of the general axioms, we assume some additional ones which signify who has access to what information. We predominantly consider *splitting* axioms of the form $\mathcal{M}X \Rightarrow \mathcal{N}\mathcal{R}X$. Awareness axioms state that agents down the line can reflect on and consider domains they have knowledge about:

- For any A , $\mathcal{B}_a A \Rightarrow \mathcal{B}_a \mathcal{B}_a A$.
- For any A , $\mathcal{I}_{a \vdash b} A \Rightarrow \mathcal{B}_a \mathcal{I}_{a \vdash b} A$.

So an agent is aware of their own internal viewpoint, as well as everything communicated to them.

Lastly we express the fact that claims made are intended as expressions of belief. For instance, if an agent a says some property is satisfied, then it is intended to mean that a claims *they believe* the property is satisfied. One may claim the latter is a weaker statement, but at the very least there exists an implication, which we express as a modal axiom.

- For any A , $\mathcal{I}_{a \vdash b} A \Rightarrow \mathcal{I}_{a \vdash b} \mathcal{B}_b A$.

2.1 Optional hierarchy of agents

We consider one more type of axiom we may add in order to increase the power of the logic. These are associated to the definition of *shared domains*, which combine several domains into one in order to evaluate consequences of their combined knowledge. We define a preorder \sqsubseteq on agents $a \sqsubseteq b$ denoting that b inherits all knowledge of a :

- If $a \sqsubseteq b$ then $\mathcal{B}_a A \Rightarrow \mathcal{B}_b A$ for any A .
- If $a \sqsubseteq c$ and $b \sqsubseteq d$ then $\mathcal{I}_{a \vdash b} A \Rightarrow \mathcal{I}_{c \vdash d} A$ for any A .

Consequently, $\mathcal{B}_a A \Rightarrow \mathcal{B}_b \mathcal{B}_c A$ if $a \sqsubseteq b$ and $a \sqsubseteq c$, and $\mathcal{I}_{a \vdash c} A \Rightarrow \mathcal{B}_b \mathcal{I}_{a \vdash c} A$ if $a \sqsubseteq b$.

As an example, we can take our domains to be subsets of entities, $\mathbb{A} = \mathcal{P}(\mathbb{A}')$, and let $S \sqsubseteq Z$ if S is a subset of Z . For $S \subseteq \mathbb{A}'$ we have that \mathcal{B}_S is the combined belief of all entities $a \in S$. We can for instance show that $\mathcal{B}_{\{a\}} A \Rightarrow \mathcal{B}_{\{b\}} (A \Rightarrow B) \Rightarrow \mathcal{B}_{\{a,b\}} (B)$. Though we have that $\mathcal{B}_{\{a\}} A \vee \mathcal{B}_{\{b\}} A \Rightarrow \mathcal{B}_{\{a,b\}} (A)$, the converse does not hold as we cannot prove $\mathcal{B}_{\{a\}} A \Rightarrow \mathcal{B}_{\{b\}} (A \Rightarrow B) \Rightarrow (B_{\{a\}}(B) \vee B_{\{b\}}(B))$.

We can also consider a special domain marking knowledge accessible to all. This is in line with the traditional necessity modality \Box which marks properties which are certainly true. We can incorporate this in our scheme by considering a special agent $\Omega \in \mathbb{A}$ for which $\Omega \sqsubseteq a$ for all other agents and assert the above axiom. This makes $\mathcal{I}_{\Omega \vdash a}$ into a kind of *broadcast* modality, as in what a makes public.

On the other side of the spectrum, we could consider the theoretical gathering of all perspectives, which can be used to check if there are any inconsistent opinions. This would be given by some agent Υ such that $b \sqsubseteq \Upsilon$ for any other agent b . In terms of our $\mathbb{A} = \mathcal{P}(\mathbb{A}')$ example, $\Omega = \emptyset$ and $\Upsilon = \mathbb{A}'$.

Last but not least, we could consider the \Box modality itself as well, to easily mark assumptions which we want all domains to adopt. We assume that $\Box A \Rightarrow \Box \Box A$, $\Box A \Rightarrow A$ and $\Box A \Rightarrow \mathcal{M}A$ for any other modality \mathcal{M} .

$$\begin{array}{c}
\frac{}{\Gamma, A \vdash A} \quad \frac{}{\Gamma \vdash \top} \quad \frac{\Gamma \vdash \perp}{\Gamma \vdash A} \quad \frac{\Gamma \vdash A \quad \Gamma \vdash A \Rightarrow B}{\Gamma \vdash B} \quad \frac{\Gamma, A \vdash B}{\Gamma \vdash A \Rightarrow B} \quad \frac{\Gamma \vdash A \wedge B}{\Gamma \vdash A} \quad \frac{\Gamma \vdash A \wedge B}{\Gamma \vdash B} \\
\frac{\Gamma \vdash A \quad \Gamma \vdash B}{\Gamma \vdash A \wedge B} \quad \frac{\Gamma \vdash A}{\Gamma \vdash A \vee B} \quad \frac{\Gamma \vdash B}{\Gamma \vdash A \vee B} \quad \frac{\Gamma \vdash A \vee B \quad \Gamma, A \vdash C \quad \Gamma, B \vdash C}{\Gamma \vdash C} \\
\frac{\Gamma \vdash \mathcal{M}A_1 \dots \Gamma \vdash \mathcal{M}A_n \quad A_1, \dots, A_n \vdash B}{\Gamma \vdash \mathcal{M}B} \quad \frac{\Gamma \vdash \mathcal{M}A \quad \mathcal{M} \Rightarrow l}{\Gamma \vdash l(A)}
\end{array}$$

Figure 2: Intuitionistic Multi-Modal K Logic

2.2 Logical Foundation

We define the set of formulas \mathbb{F} inductively as follows:

$$A, B := \top \mid \perp \mid t \mid \mathcal{M}A \mid A \Rightarrow B \mid A \wedge B \mid A \vee B$$

Note that by extension, we can formulate conjunction and disjunction over finite sets of formulas as well.

We consider a finite set of modalities \mathbb{M} satisfying axioms K and necessity, and other axioms given in the following form:

Definition 1. A modal unfolding axiom is given by a pair $(\mathcal{M} \Rightarrow l)$ where $\mathcal{M} \in \mathbb{M}$ is a modality, and $l \in \mathbb{M}^*$ is a list of modalities.

For each $l \in \mathbb{M}^*$ and formula A we define $l(A)$ inductively as $\varepsilon(A) = A$ and $(\mathcal{M} \cdot l)(A) = \mathcal{M}(l(A))$. The unfolding axiom $\mathcal{M} \Rightarrow l$ denotes the axiom $\mathcal{M}X \Rightarrow l(X)$.

Given some set of unfolding axioms Ax , we define our logic to be the intuitionistic multimodal logic, where each modality satisfies axiom K and necessity, and where furthermore the axioms from Ax are asserted. We denote this \mathcal{L}_{Ax} . In the rest of the paper, we simply take the sets of assertions Ax as discussed before, summarised in Figure 1.

There are many ways to set up such a logic. In Figure 2 is a setup for for a sequent calculus for the logic, with proof steps. A *sequent* is a pair $\Gamma \vdash A$ consisting of a set of formulas Γ denoting the assumptions, and a formula A denoting the consequent. The fractions in the figure, called *judgements*, tell us how to construct new true sequents out of known true sequents. The calculus is can be improved to better suit applications, and is mostly included here for reference. Later, in Figure 4, we formulate a Fitch-style lambda-calculus which forms a fragment of dependent modal type theory [26], and is more practical for optimizing proofs¹.

The rest of this subsection is there to consider *decidability* of the logic. That is, do we have an algorithm to determine whether or not a sequent $\Gamma \vdash A$ is provable. This is very useful in practice, as it allows us to check and verify the truth of properties given certain assumptions. This in turn allows us to hold entities accountable for their claims, as they should be aware of consequences of their claims and can be blamed if these uncover a lie.

2.3 Decidability

The main property is that the logic of Figure 2 given the modal axioms of Figure 1 given that \sqsubseteq is a preorder and \mathbb{A} is finite, forms a decidable logic. The rest of this subsection explores the nuances of what kind of axiomatic systems Ax gives rise to a decidable logic, and can be skipped. More details can be found in [50].

To get decidability, Ax needs to satisfy three additional properties.

- Ax is *reflexive* if $\mathcal{M} \Rightarrow \mathcal{M} \in \text{Ax}$ for each modality.
- Ax is *transitive* if for any $\mathcal{M} \Rightarrow l \cdot \mathcal{N} \cdot l' \in \text{Ax}$ and $\mathcal{N} \Rightarrow r \in \text{Ax}$, we have $\mathcal{M} \Rightarrow l \cdot r \cdot l' \in \text{Ax}$.
- Ax is *decomposable* if for any $\mathcal{M} \Rightarrow \mathcal{N} \cdot \mathcal{N}' \cdot l$ there is an $\mathcal{R} \in \mathbb{M}$ such that $\mathcal{M} \Rightarrow \mathcal{N} \cdot \mathcal{R} \in \text{Ax}$ and for any $\mathcal{M} \Rightarrow \mathcal{N} \cdot \mathcal{T} \cdot r \in \text{Ax}$, we have $\mathcal{R} \Rightarrow \mathcal{T} \cdot r \in \text{Ax}$.

¹We shall consider the Proofs=Programs Curry-Howard correspondence there. What we are mainly interested in in the formulation of that variant is to get the *substitutivity* property, which allows us to simplify proofs.

The first two expose the categorical structure of the axioms, whereas the third property asserts that any “unfolding” of a modality into three or more modalities can be factored into a composition of axioms of the form $\mathcal{M} \Rightarrow \mathcal{N} \cdot \mathcal{R}$, as well as such decompositions are universal allowing for easier proof search. We denote a choice of \mathcal{R} as given in the third point by $\mathcal{M} \ominus \mathcal{N}$, meaning that if $\mathcal{M} \Rightarrow \mathcal{N} \cdot \mathcal{N}' \cdot l$, then $\mathcal{M} \ominus \mathcal{N}$ exists, $\mathcal{M} \Rightarrow \mathcal{N} \cdot (\mathcal{M} \ominus \mathcal{N})$, and $\mathcal{M} \ominus \mathcal{N} \Rightarrow \mathcal{N}' \cdot l$.

The reflexive and transitive closure of Ax is denoted as $\widehat{\text{Ax}}$.

Lemma 1. *For any Ax , the logics \mathcal{L}_{Ax} and $\mathcal{L}_{\widehat{\text{Ax}}}$ are equivalent.*

Proof. We prove this by induction on the reflexive transitive closure. If $\mathcal{M} \Rightarrow l \in \widehat{\text{Ax}}$ then either:

- $\mathcal{M} \Rightarrow l \in \text{Ax}$ in which case we are done.
- $l = \mathcal{M}$, in which case the axiom asserts that $\mathcal{M}A \Rightarrow \mathcal{M}A$ for any formula A , which is trivially true.
- It came from the composition of $\mathcal{M} \Rightarrow l_1 \cdot \mathcal{N} \cdot l_2$ and $\mathcal{N} \Rightarrow r$, upon which we can apply the induction hypothesis. So for any formula A , we have that $\mathcal{M}(A) \Rightarrow l_1 \mathcal{N} l_2(A)$ and $\mathcal{N}(l_2 A) \Rightarrow r(l_2 A)$ are provable. Applying necessity on the latter, we get $l_1 \mathcal{N} l_2 A \Rightarrow l_1 r l_2 A$, and hence by intuitionistic reasoning ($A \Rightarrow B, B \Rightarrow C \vdash A \Rightarrow C$) we get that $\mathcal{M}A \Rightarrow l_1 r l_2 A$ is provable. □

Lemma 2. *If Ax is the system of axioms defined in Figure 1, then $\widehat{\text{Ax}}$ is decomposable.*

Proof. For \mathcal{B}_a , note that $\mathcal{B}_a \Rightarrow \mathcal{N} \cdot \mathcal{N}' \cdot l \in \widehat{\text{Ax}}$ means $\mathcal{N} \cdot \mathcal{N}' \cdot l = \mathcal{B}_{a_1} \cdot \mathcal{B}_{a_2} \cdot \dots \cdot \mathcal{B}_{a_n}$ for some list of at least two agents such that $a \sqsubseteq a_i$ for all $i \in \{1, \dots, n\}$ (without inheritance, all a_i are a). We simply take $\mathcal{B}_a \ominus \mathcal{B}_{a_1} = \mathcal{B}_a$, for which $\mathcal{B}_a \Rightarrow \mathcal{B}_{a_1} \cdot \mathcal{B}_a$ and $\mathcal{B}_a \Rightarrow \mathcal{B}_{a_2} \cdot \dots \cdot \mathcal{B}_{a_n}$.

For \mathcal{I}_{a+b} , axiom $\mathcal{I}_{a+b} \Rightarrow \mathcal{N} \cdot \mathcal{N}' \cdot l \in \widehat{\text{Ax}}$ means $\mathcal{N} \cdot \mathcal{N}' \cdot l = \mathcal{B}_{a_1} \cdot \dots \cdot \mathcal{B}_{a_n} \cdot \mathcal{I}_{a'+b'} \cdot \mathcal{B}_{b_1} \cdot \dots \cdot \mathcal{B}_{b_m}$ where $n+1+m \geq 2$, $a \sqsubseteq a_i$, $a \sqsubseteq a'$, $b \sqsubseteq b'$ and $b \sqsubseteq b_j$. If $n=0$ we simply take $\mathcal{I}_{a+b} \ominus \mathcal{I}_{a'+b'} = \mathcal{B}_b$, and if $n > 0$ we take $\mathcal{I}_{a+b} \ominus \mathcal{B}_{a_1} = \mathcal{I}_{a+b}$. □

3 Chains of Communications

Modalities let us mark different domains of knowledge, and are as such useful for describing distributed networks of trust. In such networks, entities make claims, which may be communicated throughout the network, passing through multiple entities. At each location, it can then be determined what is known and what can be derived.

We consider a canonical example, which we shall describe in a multitude of ways. The scenario is that a certain claim goes through multiple entities before getting to the targeted recipient. How is this claim sent, and what are the preconditions for the recipient to trust this claim?

We consider four entities in a chain of communications:

$$a \longrightarrow b \longrightarrow c \longrightarrow d$$

Here a will function as the origin of some claim A , and d as the intended recipient. However, suppose either a cannot send the claim directly to d , or d does not trust a directly. Whichever is the case, b and c are used as intermediaries.

Example 1 (A global view). The simplest model simply asserts that the intermediaries pass along the the claim as their own. This can be stated by $\mathcal{I}_{b+a}A \Rightarrow \mathcal{I}_{c+b}A$ and $\mathcal{I}_{c+b}A \Rightarrow \mathcal{I}_{d+c}A$. They say that if b receives claim A from a , it will send claim A to c , and similarly for c . The assertions compose into $\mathcal{I}_{b+a}A \Rightarrow \mathcal{I}_{d+c}A$, hence if a sends the claim it will end up at d .

The assertions in the above examples do require some trust by b and c , since they need to assert the truth of A . Secondly, note that we need access to all such assertions in order to make the final derivation. We shall address both these things in the next two examples.

Example 2 (Trust and Responsiveness). The statement $\mathcal{I}_{b+a}A \Rightarrow \mathcal{I}_{c+b}A$ requires b to repeat a claim to another, which effectively means b must state that A is true. Hence the assertion requires trust of b in a . Moreover, the statement asserts a response by b , the sending of information to c . We can break it up in the following components:

- Trust: $\mathcal{B}_b(\mathcal{I}_{b \leftarrow a} A \Rightarrow A)$, if b receives the claim A from a , then it is believed. Note that with axiom $\mathcal{I}_{b \leftarrow a} A \Rightarrow \mathcal{B}_b \mathcal{I}_{b \leftarrow a} A$, we may derive $\mathcal{I}_{b \leftarrow a} A \Rightarrow \mathcal{B}_a A$.
- Responsiveness: $\mathcal{B}_b(\mathcal{I}_{b \leftarrow a} A \wedge A) \Rightarrow \mathcal{I}_{c \leftarrow b} A$, if b receives the claim A and believes it, b will send it to c . We can split this up further into two statements.
 - $\mathcal{B}_b(\mathcal{I}_{b \leftarrow a} A \wedge A \Rightarrow \mathcal{I}_{c \leftarrow b} A)$ where b believes in their own commitment to sending the claim along, and
 - $\mathcal{B}_b \mathcal{I}_{c \leftarrow b} A \Rightarrow \mathcal{I}_{c \leftarrow b} A$ saying we trust b keeps their commitment².

Suppose we want to ensure that d can derive that the message is passed along. The assertions defined prior as stated are not accessible to d , by which we mean we have no reason for thinking d knows or believes in them. So if d does not have them as prior assumptions, the entities are required communicate these statements of trust and responsiveness to d themselves. For brevity, we write $\mathcal{I}_{a_1 \leftarrow a_2 \leftarrow a_3 \dots \leftarrow a_n}$ for $\mathcal{I}_{a_1 \leftarrow a_2} \mathcal{I}_{a_2 \leftarrow a_3} \dots \mathcal{I}_{a_{n-1} \leftarrow a_n}$.

Example 3 (Communicating Trust). Instead of stating $\mathcal{I}_{b \leftarrow a} A \Rightarrow \mathcal{I}_{c \leftarrow b} A$, we let b tell c that they will do this. We state that as $\mathcal{I}_{c \leftarrow b}(\mathcal{I}_{b \leftarrow a} A \Rightarrow A)$, meaning b tells c that if they receive claim A from a , then A is true. Agent b need not explicitly send claim A anymore, they can simply forward a received claim from a and thereby implying the truth of A . This consequence is derived by applying axiom K, getting $\mathcal{I}_{c \leftarrow b} \mathcal{I}_{b \leftarrow a} A \Rightarrow \mathcal{I}_{c \leftarrow b} A$.

With everything being forwarded to d , we get statements:

- $\mathcal{I}_{d \leftarrow c \leftarrow b}(\mathcal{I}_{b \leftarrow a} A \Rightarrow A)$, it is forwarded that b trusts a .
- $\mathcal{I}_{d \leftarrow c}(\mathcal{I}_{c \leftarrow b} A \Rightarrow A)$, it is communicated that c trusts b .
- $\mathcal{B}_d(\mathcal{I}_{d \leftarrow c} A \Rightarrow A)$, d trusts c .

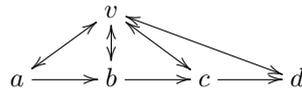
All composed together, we can derive $\mathcal{B}_d(\mathcal{I}_{d \leftarrow c \leftarrow b \leftarrow a} A \Rightarrow A)$, d trusts forwarded claims from a .

The communicated statements of the previous example can be split up into different parts as done in example 2, if more clarification is needed. There is still one disadvantage regarding the above approach: each link in the chain must trust the previous link regarding statement A , even though that previous link may not be an expert on A . This leaves the system open to a kind of injected claim attack: if any link starts claiming A , everyone down the line will start to accept this claim, even though the claim did not originate from a . Though this problem cannot be entirely prevented, as we need to assume at least some integrity from the links, there are three partial solutions we shall consider.

The first solution is to fine tune which claims are trusted. Instead of simply trusting A , we can only accept certain evidence of A . Even though a link may not be an expert on A , it could at least be trusted regarding other things, like communicated claims from other experts and verifying that such claims came from trusted individuals. This shall be the topic of Section 4. The second solution is to use multiple sources to double check claims, which we shall consider in Section 7 using *threshold trust*.

The third solution has a different flavour, with none of the links in the communication chain needing to trust each other. Instead, they all have access to a single trusted source which can be consulted and used to double check if a claim is correct. This source can for instance be some kind of internet server, or a certification authority.

Example 4 (Trusted Authority). Let v be a new entity everyone has access to.



The behaviour of v can be specified as follows: it can send out claims (e.g. certifications) and will confirm to others that they indeed made this claim.

This is given as follows: $\bigwedge_{x,y \in \mathbb{A}} \mathcal{I}_{v \leftarrow x \leftarrow y \leftarrow v} A \Rightarrow \mathcal{B}_v A \Rightarrow \mathcal{I}_{x \leftarrow v} A$, meaning if someone notifies v that they heard from someone else that v has claimed something, then if this is indeed true, then v confirms this claim.

This can then be used by b given two statements:

²One could argue this should be an axiom, though note it requires b to actively check and act on their commitments, hence it is reasonable to formulate it as a kind of trust in b instead of a separate axiom.

1. $\mathcal{I}_{b \leftarrow a \leftarrow v} A \Rightarrow \mathcal{I}_{v \leftarrow b \leftarrow a \leftarrow v} A \vee \mathcal{I}_{b \leftarrow v} A$, which says that if b receives from a a claim made by v , it will check this with v directly, unless it has already received this claim from v . Note that with the specification of v , this implies $\mathcal{I}_{b \leftarrow a \leftarrow v} A \Rightarrow \mathcal{B}_v A \Rightarrow \mathcal{I}_{b \leftarrow v} A$.
2. $\mathcal{I}_{b \leftarrow v} A \Rightarrow \mathcal{I}_{c \leftarrow b \leftarrow v} A$.

Supposing b , c and d have the first property, and a , b and c have the second property, we can derive $\mathcal{B}_v A \Rightarrow \mathcal{I}_{a \leftarrow v} A \Rightarrow \mathcal{I}_{d \leftarrow v} A$.

4 Trust

For cooperation, it is useful for entities to declare their trust in someone before that entity has made the relevant claim. This way, others can be assured that the entity can be convinced if needed, and how this can be done. We define trust in two stages. First we introduce some notation; for each list of modalities $l = \mathcal{M}_1, \dots, \mathcal{M}_n$ and formula A , *validity* of l concerning A is given by:

$$(\widehat{\mathcal{M}_1 \dots \mathcal{M}_n})A := \mathcal{M}_1 \dots \mathcal{M}_n A \Rightarrow A$$

For instance, validity $\widehat{\mathcal{I}_\alpha}$ for some $\alpha = a_1 \leftarrow a_2 \leftarrow \dots \leftarrow a_n$ is the assertion that if claim A is communicated through chain of communications α , then A is true. Hence validity depends on four aspects: the sender, the receiver, the route the message took, and the content of the claim.

The concept of *trust* is simply a statement of validity put into some context. Validity itself is the expression that “we”, meaning those considering the current set of assumptions, trust the claim made over some channel. We can also consider trust as either believed or claimed by a particular entity. E.g. $\mathcal{B}_b \widehat{\mathcal{I}_\alpha} A$ means b trusts claim A if communicated over α , and $\mathcal{I}_\beta \widehat{\mathcal{I}_\alpha}$ is this trust being proclaimed over the chain of communications β . Usually, trust occurs in the recipient of the considered claim, e.g. in $\mathcal{B}_b \widehat{\mathcal{I}_{b \leftarrow \alpha}} A$ or $\mathcal{I}_{\gamma \leftarrow b} \widehat{\mathcal{I}_{b \leftarrow \alpha}} A$.

Given our axioms, *cautious trust* $\mathcal{T}_{a,b} A := \mathcal{B}_a (\mathcal{I}_{a \leftarrow b} A \Rightarrow \mathcal{B}_b A \wedge \mathcal{B}_b A \Rightarrow A)$ as formulated by Liao [41] can be proven with the following validities:

$$\mathcal{B}_a \widehat{\mathcal{I}_{a \leftarrow b}} \mathcal{B}_b A \wedge \mathcal{B}_a \widehat{\mathcal{B}_b} A \Rightarrow \mathcal{T}_{a,b} A$$

$\widehat{\mathcal{I}_\alpha}$ satisfies two rules which helps reasoning about it:

- $A \Rightarrow \widehat{\mathcal{I}_\alpha} A$, if a statement A is already true, then communications making this claim are trivially valid.
- $\widehat{\mathcal{I}_\alpha} A \wedge \widehat{\mathcal{I}_\alpha} B \Rightarrow \widehat{\mathcal{I}_\alpha} (A \wedge B)$, a combined statement’s validity can be determined in parts. The converse does however not hold.

Lemma 3. *For any modality \mathcal{M} , and formulas A and B , $\widehat{\mathcal{M}}A \wedge \widehat{\mathcal{M}}(A \Rightarrow B) \Rightarrow \widehat{\mathcal{M}}B$ is not necessarily provable.*

Proof. We show that there is a formula A and a Kripke frame which does not satisfy this statement. Let W have two worlds, v and w , and let \leq be the identity relation. Let t and r be two tokens such that $S_t = \emptyset$ and $S_r = \{w\}$. Lastly, let $R_{\mathcal{M}}$ relate v with w , but not v with v .

Then neither v nor w model $\mathcal{M}t$, hence both model $\widehat{\mathcal{M}}t$. Neither v nor w model t , so both model $t \Rightarrow r$, and hence $\widehat{\mathcal{M}}(t \Rightarrow r)$. Since $v R_{\mathcal{M}} v$ does not hold, and $w \models r$, we have $v \models \mathcal{M}r$. However, since v does not model r , v does not model $\widehat{\mathcal{M}}r$. So v does not model $\widehat{\mathcal{M}}t \wedge \widehat{\mathcal{M}}(t \Rightarrow r) \Rightarrow \widehat{\mathcal{M}}r$. \square

Lemma 4.

$$\widehat{\mathcal{I}_{\alpha \leftarrow b}} A \wedge \mathcal{I}_{\alpha \leftarrow b} \widehat{\mathcal{I}_{b \leftarrow \gamma}} A \Rightarrow \widehat{\mathcal{I}_{\alpha \leftarrow b \leftarrow \gamma}} A$$

Proof. $\mathcal{I}_{\alpha \leftarrow b} \widehat{\mathcal{I}_{b \leftarrow \gamma}} A = \mathcal{I}_{\alpha \leftarrow b} (\mathcal{I}_{b \leftarrow \gamma} A \Rightarrow A)$ implies $\mathcal{I}_{\alpha \leftarrow b} \mathcal{I}_{b \leftarrow \gamma} A \Rightarrow \mathcal{I}_{\alpha \leftarrow b} A$ by axiom K, where $\mathcal{I}_{\alpha \leftarrow b} \mathcal{I}_{b \leftarrow \gamma} = \mathcal{I}_{\alpha \leftarrow b \leftarrow \gamma}$. Hence combined with $\widehat{\mathcal{I}_{\alpha \leftarrow b}} A = (\mathcal{I}_{\alpha \leftarrow b} A \Rightarrow A)$ we get $(\mathcal{I}_{\alpha \leftarrow b \leftarrow \gamma} A \Rightarrow A) = \widehat{\mathcal{I}_{\alpha \leftarrow b \leftarrow \gamma}} A$. \square

Here validity is determined in two stages:

- b communicates over α that they would trust claim A if coming over channel γ .
- Communication of A over α are trusted.
- As such, communications over γ through b and then over α are trusted.

We gather derivable properties of validities studied in this section in Figure 3 for reference.

Property	Name
$\vdash \widehat{\mathcal{M}}A \wedge \mathcal{M}A \Rightarrow A$	Validity definition
$\vdash A \Rightarrow \widehat{\mathcal{M}}A$	Validity unit
$\vdash \widehat{\mathcal{M}}A \wedge \widehat{\mathcal{M}}B \Rightarrow \widehat{\mathcal{M}}(A \wedge B)$	Merging validities
$\vdash \widehat{\mathcal{M}}\mathcal{N}A \wedge \widehat{\mathcal{N}}A \Rightarrow \widehat{\mathcal{M}}\mathcal{N}A$	Composition rule 1
$\vdash \mathcal{M}\widehat{\mathcal{N}}A \wedge \widehat{\mathcal{M}}A \Rightarrow \widehat{\mathcal{M}}\mathcal{N}A$	Composition rule 2
$\vdash \mathcal{M}\widehat{\mathcal{N}}A \wedge \widehat{\mathcal{M}}\widehat{\mathcal{N}}A \wedge \widehat{\mathcal{M}}\mathcal{N}A \Rightarrow \widehat{\mathcal{M}}\mathcal{N}A$	Composition rule 3
$\nmid \widehat{\mathcal{M}}A \wedge \widehat{\mathcal{M}}(A \Rightarrow B) \Rightarrow \widehat{\mathcal{M}}B$	No axiom K

Figure 3: Validity properties (Derivable in the logic)

4.1 Indirect trust

Let us look at a basic chaining of trust. We shall reason from the perspective of an agent a , though we could as easily have considered a communicating these derivations to other agents instead. The basic derivation done by applying axiom K to Lemma 4 is as follows:

$$\mathcal{B}_a \widehat{\mathcal{I}}_{a+b}A \wedge \mathcal{I}_{a+b} \widehat{\mathcal{I}}_{b+c}A \Rightarrow \mathcal{B}_a \widehat{\mathcal{I}}_{a+b+c}A$$

Here a reasons that they trust b regarding property A . If b then communicates they trust c regarding property A , a can infer trust in c if at least claims go through b . This is an effective way of chaining trust, which requires a to trust b regarding A without argument. Note however that we also have the following derivation:

$$\mathcal{B}_a \widehat{\mathcal{I}}_{a+b}A \wedge \mathcal{I}_{a+b}A \Rightarrow \mathcal{B}_a A$$

In certain situations, we may want to fine-tune and reduce the necessary trust required of a .

Let us instead suppose a accepts a certain proof from b , but does not want to simply accept the conclusion. The argument from b comes in two parts:

1. I (meaning b) trust c if they think A is true.
2. c has communicated to me that A is true.

The obvious conclusion of these two statements is that b believes A is true. Instead of simply accepting this conclusion, a can choose to only accept the arguments themselves. They can state the following:

$$\mathcal{B}_a \widehat{\mathcal{I}}_{a+b} \widehat{\mathcal{I}}_{b+c}A \wedge \mathcal{B}_a \widehat{\mathcal{I}}_{a+b} \mathcal{I}_{b+c}A$$

stating that they trust b regarding their trust in c , as well as trust b to not falsely forward messages from c . It should be noted that a can derive $\mathcal{B}_a \widehat{\mathcal{I}}_{a+b}(\widehat{\mathcal{I}}_{b+c}A \wedge \mathcal{I}_{b+c}A)$ but cannot derive $\mathcal{B}_a \widehat{\mathcal{I}}_{a+b}A$ as seen in Lemma 3. The arguments are accepted, but not the conclusion. Still, b 's declaration of trust allows a to gain trust as well:

$$\mathcal{B}_a \widehat{\mathcal{I}}_{a+b}(\widehat{\mathcal{I}}_{b+c}A \wedge \mathcal{I}_{b+c}A) \wedge \mathcal{I}_{a+b} \widehat{\mathcal{I}}_{b+c}A \Rightarrow \mathcal{B}_a \widehat{\mathcal{I}}_{a+b+c}A$$

The argument goes as follows: b has already communicated its trust to a . So if b then later mentions that c claimed A , the two claims by b can be combined into a single package $\mathcal{I}_{a+b}(\widehat{\mathcal{I}}_{b+c}A \wedge \mathcal{I}_{b+c}A)$ which is trusted by a , and hence $\mathcal{B}_a(\widehat{\mathcal{I}}_{b+c}A \wedge \mathcal{I}_{b+c}A)$. Having accepted the arguments then allows a to make the appropriate derivation, that $\mathcal{B}_a A$. However, if in this case b had communicated A directly, a would not have accepted it: a trusts b regarding identifying experts on A , but a does not trust b as an expert on A itself.

In general, suppose b has derived a statement B using assumptions A_1, \dots, A_n . People like a can trust b 's conclusion directly, as stated by $\mathcal{B}_a \widehat{\mathcal{I}}_{a+b}B$. Or a can instead require that b provide their arguments, and state $\mathcal{B}_a \widehat{\mathcal{I}}_{a+b}(A_1 \wedge \dots \wedge A_n)$ or $\mathcal{B}_a(\mathcal{I}_{a+b}A_1 \wedge \dots \wedge \widehat{\mathcal{I}}_{a+b}A_n)$. In the latter situation, b can only convince a by claiming all the arguments. There are still two things to be considered.

First note that the proof of B via A_1, \dots, A_n is never communicated. The proof is implied, and has to be done by a . Indeed, the burden of proof is naturally on the one who wants to trust, not on the entity providing the facts. One could imagine a asking b whether B is true, and b communicating $A_1 \wedge \dots \wedge A_n$ instead of B itself. Regardless of whether a trusts b regarding B , or regarding $A_1 \wedge \dots \wedge A_n$, a will be able to convince themselves of B . The fact that B is the conclusion is implied by the claims. Note that with decidability, a will be able to verify that the proof exists.

Secondly, it may not make sense for a to trust b regarding just a specific entity c . In most instances, the particular c in question may not be predetermined. We can state that a trusts b regarding any such entity, as can be formulated as:

$$\mathcal{B}_a \bigwedge_{c \in \mathbb{A}} \widehat{\mathcal{I}}_{a \rightarrow b}(\widehat{\mathcal{I}}_{b \rightarrow c} A \wedge \mathcal{I}_{b \rightarrow c} A)$$

We call this *higher-order trust*. We do not trust b directly regarding A , but do trust b to identify experts of A . These orders of trust are an important tool for describing trust infrastructures.

5 Forwarding Networks

We use our logic to describe trust in networks of entities. These can be used to describe public key infrastructures, as well as other multiparty networks of authentication. The core idea is to establish what the role of an entity in the network is, and what type of trust is required for productive cooperation with this entity. Agents may establish trust directly, or use intermediaries like certification authorities to establish trust.

We start by describing the language on a high level. We write $\text{trust}_n a, b(\mathbf{P})$ to say that a trusts b regarding predicate $\mathbf{P} : \mathbb{A} \rightarrow \mathbb{F}$, to a certain *order* n . Here, \mathbf{P} is some predicate on agents, describing some fundamental claim one wants to convey to others, like whether an agent owns a certain public key, or whether an agent has a certain certificate or diploma³. Most likely, we take as our set of tokens \mathbb{T} some statements that can refer to agents, e.g. with $\text{hasKey}(a) \in \mathbb{T}$ for each $a \in \mathbb{A}$, we can use hasKey as a predicate. The order n is a natural number saying whether there is direct trust between two agents, or whether there is trust between agents as intermediaries in bigger networks.

- $\text{trust}_0 a, b(\mathbf{P})$ means that if a receives a claim $\mathbf{P}(b)$ from b , then a will accept this claim $\mathbf{P}(b)$ as true.
- $\text{trust}_{n+1} a, b(\mathbf{P})$ means that if a receives a claim from b regarding the n -th order trust of b in \mathbf{P} , then a believes this trust is warranted.

We can look at some examples of different orders of trust, in the realm of public key infrastructures.

- Order 0 trust: This is the goal. Certification authorities record and communicate their order 0 trust regarding key ownership, and others can inherit this order 0 trust.
- Order 1 trust: This is the order of trust one has in a certification authority. With order 1 trust in the CA, one can inherit their order 0 trust.
- Order 2 trust: This is the trust one has in authorities which validate CAs, like domain trusted lists. With order 2 trust in a domain trusted list, one can inherit order 1 trust in the CA which is listed.
- Order 3 trust: This is trust in entities which claim to have order 2 trust. For instance, trust in a government backing some domain trusted list.

In order to establish trust across networks, claims will need to be forwarded and checked by intermediary agents. We have to specify the allowed paths of communication.

Definition 2. A *forwarding network* is a set of non-empty, non-repeating lists of agents $S \subseteq \mathbb{A}^*$, such that:

1. If $(a, \beta, c) \in S$ then $(\beta, c) \in S$ and $(a, \beta) \in S$ (in other words, S is closed under taking sublists).
2. If $(\alpha, a, \beta, b, \gamma) \in S$ and $(a, \beta', b) \in S$ then $(\alpha, a, \beta', b, \gamma) \in S$.

If $a, \alpha, b \in S$, we say that α is a defined path from b to a . Any segment of the path can be replaced by an alternative route, if specified by S . Note that α can be empty, in which case there is a direct path.

We write $\mathcal{N}_S(a, b, c)$ if there is a path $a, \alpha, b, \beta, c \in S$.

Lemma 5. Given a forwarding network $S \subseteq \mathbb{A}^*$, then \mathcal{N}_S has the following properties:

- If $\mathcal{N}_S(a, b, d) \& \mathcal{N}_S(b, c, d)$ then $\mathcal{N}_S(a, b, c) \& \mathcal{N}_S(a, c, d)$.
- If $\mathcal{N}_S(a, b, c) \& \mathcal{N}_S(a, c, d)$ then $\mathcal{N}_S(a, b, d) \& \mathcal{N}_S(b, c, d)$.

³ \mathbf{P} need not depend on a , and can be some other primitive statement.

5.1 Forwarding Modality

It is assumed that messages will be passed along across every specified path. This is independent of whether the message itself is trusted. As we have seen before, we can establish commitments for forwarding messages which do not require additional trust from the network members, except for trusting they will pass along the message. If you forward a message, you are not necessarily claiming that the content is true. We define the collection of all forwardings with a *composite modality*. For $a, b \in \mathbb{A}$, let

$$\mathcal{J}_{a+b}A := \bigwedge \{ \mathcal{I}_{a+\gamma+b}A \mid \gamma \in \mathbb{A}^* \text{ such that } (a, \gamma, b) \in S \}$$

Note that $\mathcal{J}_{a+a}A = \top$ by definition. The composite modality behaves like a modality itself:

Lemma 6. *The following properties hold:*

- \mathcal{J}_{a+b} satisfies axiom K and necessity.
- $\mathcal{J}_{a+b}A \implies \mathcal{B}_a\mathcal{J}_{a+b}A$.
- $\mathcal{J}_{a+b}A \implies \mathcal{J}_{a+b}\mathcal{B}_bA$.

Last but certainly not least, we have the additional property which allows us to factor these modalities further.

Lemma 7. *If $\mathcal{N}_S(a, b, c)$ then $\vdash \mathcal{J}_{a+c}A \implies \mathcal{J}_{a+b}\mathcal{J}_{b+c}A$.*

Proof. Assume $\mathcal{J}_{a+c}A$, hence for any $a, \alpha, c \in S$ we have $\mathcal{I}_{a+\alpha+c}A$. To prove $\mathcal{J}_{a+b}\mathcal{J}_{b+c}A$ we prove $\mathcal{I}_{a+\beta+b}\mathcal{J}_{b+c}A$ for each $(a, \beta, b) \in S$. $\mathcal{I}_{a+\beta+b}\mathcal{J}_{b+c}A$ in turn can be shown by proving $\mathcal{I}_{a+\beta+b}\mathcal{I}_{b+\gamma+c}A$ for each $(b, \gamma, c) \in S$ and applying axiom K, using that there are only finitely many possible γ -s.

So suppose $(a, \beta, b) \in S$ and $(b, \gamma, c) \in S$, and suppose $\mathcal{N}_S(a, b, c)$, meaning $(a, \beta', b, \gamma', c) \in S$ for some β' and γ' . By applying property 2 of forwarding networks twice, $(a, \beta, b, \gamma, c) \in S$. Since $\mathcal{J}_{a+c}A$ implies any path from c to a , it implies $\mathcal{I}_{a+\beta+b+\gamma+c}A$ which by definition is $\mathcal{I}_{a+\beta+b}\mathcal{I}_{b+\gamma+c}A$, what we need. \square

Basically, since \mathcal{J}_{a+c} attempts all paths of communication from c to a , these include all paths through agent b . We again use $\widehat{\mathcal{J}}_{a+b}$ to express validity of such communications:

$$\widehat{\mathcal{J}}_{a+b}A := \mathcal{J}_{a+b}A \Rightarrow A$$

Note that this considers all possible paths, so such validity can be obtained if validity exists across any one path.

- If $(a, \alpha, b) \in S$, then $\widehat{\mathcal{I}}_{a+\alpha+b}A \implies \widehat{\mathcal{J}}_{a+b}A$.

Lastly note that if no path from b to a exists, then $\widehat{\mathcal{J}}_{a+b}A \equiv (\top \Rightarrow A) \equiv A$, hence trust across non existing paths only happens when the claim in question is true.

5.2 Orders of trust

Direct trust from a in b can be defined as $\mathcal{B}_a\widehat{\mathcal{I}}_{a+b}\mathbf{P}(b)$, meaning a trusts b regarding claim $\mathbf{P}(b)$. Agent b is effectively claiming that predicate \mathbf{P} holds for them, for instance saying: “I own public key k ”, or “I trust c ”.

We consider the following two properties fundamental to the development. If $\mathcal{N}_S(a, b, c)$ then:

- $\mathcal{J}_{a+b}\widehat{\mathcal{J}}_{b+c}A \wedge \widehat{\mathcal{J}}_{a+b}A \implies \widehat{\mathcal{J}}_{a+c}A$
- $\mathcal{J}_{a+b}\widehat{\mathcal{J}}_{b+c}A \wedge \widehat{\mathcal{J}}_{a+b}\widehat{\mathcal{J}}_{b+c}A \wedge \widehat{\mathcal{J}}_{a+b}\mathcal{J}_{b+c}A \implies \widehat{\mathcal{J}}_{a+c}A$

Note that the proof of the latter statement first deduces $\widehat{\mathcal{J}}_{a+b}A$ from the pieces of evidence $\widehat{\mathcal{J}}_{a+b}\widehat{\mathcal{J}}_{b+c}A$ and $\widehat{\mathcal{J}}_{a+b}\mathcal{J}_{b+c}A$, and then applies the former statement.

In the above statements, A can either be a basic statement $\mathbf{P}(c)$, or itself a statement necessary for composing trust further. As such, we need sequences of communication and validity claims, whose length depends on the order of trust we want to express. Using \mathbf{P} as a basis, we define a set of higher-order statements $\mathbf{C}_a^n(\mathbf{P})$ inductively as follows:

$$\begin{aligned} \mathbf{C}_a^0(\mathbf{P}) &:= \{\mathbf{P}(a)\} \\ \mathbf{C}_a^{n+1}(\mathbf{P}) &:= \{\widehat{\mathcal{J}}_{a+b}A, \mathcal{J}_{a+b}A \mid b \in \mathbb{A}, A \in \mathbf{C}_b^n(\mathbf{P})\} \end{aligned}$$

Using the fact that each set is finite, we define higher order trust as a formula using conjunctions.

Definition 3. We define n -th order validity of a in b as:

$$\text{val}_n(a+b)(P) := \bigwedge \{ \widehat{\mathcal{J}}_{a+b} A \mid A \in \mathbf{C}_b^n(P) \}$$

We define n -th order trust of a in b as:

$$\text{trust}_n(a+b)(P) := \mathcal{B}_a \text{val}_n(a+b)(P)$$

For example, $\text{trust}_0(a+b)(P) \equiv \mathcal{B}_a \widehat{\mathcal{J}}_{a+b} P(b)$ and $\text{trust}_1(a+b)(P) \equiv \bigwedge_c (\mathcal{B}_a \widehat{\mathcal{J}}_{a+b} \widehat{\mathcal{J}}_{b+c} P(c) \wedge \mathcal{B}_a \widehat{\mathcal{J}}_{a+b} \mathcal{J}_{b+c} P(c))$.

Let us establish some needed lemmas provable in the logic.

Lemma 8. If $\mathcal{N}_S(a, b, c)$ then $\vdash \text{val}_{n+1}(a+b)(P) \Rightarrow \widehat{\mathcal{J}}_{a+b}(\text{val}_n(b+c)(P))$.

Proof. Note that $\widehat{\mathcal{J}}_{a+b} A \wedge \widehat{\mathcal{J}}_{a+b} B \Rightarrow \widehat{\mathcal{J}}_{a+b}(A \wedge B)$ since $\mathcal{J}_{a+b}(A \wedge B) \Rightarrow \mathcal{J}_{a+b} A \wedge \mathcal{J}_{a+b} B$. These implications can be generalised to conjunctions over finite sets. As a consequence,

$$\begin{aligned} \bigwedge_{A \in \mathbf{C}_c^n(P)} \widehat{\mathcal{J}}_{a+b} \widehat{\mathcal{J}}_{b+c} A &= \bigwedge_{A \in \mathbf{C}_c^n(P)} (\mathcal{J}_{a+b} \widehat{\mathcal{J}}_{b+c} A \Rightarrow \widehat{\mathcal{J}}_{b+c} A) \\ \Rightarrow (\widehat{\mathcal{J}}_{a+b} \text{val}_n(b+c)(P) \Rightarrow \text{val}_n(b+c)(P)) \\ &= \widehat{\mathcal{J}}_{a+b}(\text{val}_n(b+c)(P)) \end{aligned} \quad \square$$

We can compose validity in the following way.

Lemma 9. If $\mathcal{N}_S(a, b, c)$, the following properties hold:

1. $\vdash \text{val}_{n+1}(a+b)(P) \wedge \text{val}_n(b+c)(P) \Rightarrow \text{val}_n(a+c)(P)$.
2. $\vdash \text{val}_{n+1}(a+b)(P) \wedge \mathcal{J}_{a+b} \text{val}_n(b+c)(P) \Rightarrow \text{val}_n(a+c)(P)$.
3. $\vdash \text{val}_n(a+b)(P) \wedge \mathcal{J}_{a+b} \text{val}_n(b+c)(P) \Rightarrow \text{val}_n(a+c)(P)$.
4. $\vdash \text{trust}_{n+1}(a+b)(P) \wedge \mathcal{J}_{a+b} \text{val}_n(b+c)(P) \Rightarrow \text{trust}_n(a+c)(P)$.
5. $\vdash \text{trust}_n(a+b)(P) \wedge \mathcal{J}_{a+b} \text{val}_n(b+c)(P) \Rightarrow \text{trust}_n(a+c)(P)$.

Proof. Let $A \in \mathbf{C}_c^n(P)$, we want to show that $\widehat{\mathcal{J}}_{a+c} A$ in the first three cases, which means $\mathcal{J}_{a+c} A \Rightarrow A$. Since $\mathcal{J}_{a+c} A \Rightarrow \mathcal{J}_{a+b} \mathcal{J}_{b+c} A$ it is sufficient to prove A from $\mathcal{J}_{a+b} \mathcal{J}_{b+c} A$. So assume $\mathcal{J}_{a+b} \mathcal{J}_{b+c} A$,

1. $\text{val}_{n+1}(a+b)(P)$ implies $\widehat{\mathcal{J}}_{a+b} \mathcal{J}_{b+c} A$, hence $\mathcal{J}_{b+c} A$. From $\text{val}_n(b+c)(P)$ we have $\widehat{\mathcal{J}}_{b+c} A$, hence A .
2. Like before we can derive $\mathcal{J}_{b+c} A$. $\text{val}_{n+1}(a+b)(P)$ also implies $\widehat{\mathcal{J}}_{a+b} \widehat{\mathcal{J}}_{b+c} A$, and from $\mathcal{J}_{a+b}(\text{val}_n(b+c)(P))$ we have $\mathcal{J}_{a+b} \widehat{\mathcal{J}}_{b+c} A$, hence $\widehat{\mathcal{J}}_{b+c} A$, which together with $\mathcal{J}_{b+c} A$ makes A .
3. From $\mathcal{J}_{a+b}(\text{val}_n(b+c)(P))$ we have $\mathcal{J}_{a+b} \widehat{\mathcal{J}}_{b+c} A$, hence with $\mathcal{J}_{a+b} \mathcal{J}_{b+c} A$ we get $\mathcal{J}_{a+b} A$. From $\text{val}_n(a+b)(P)$ we get $\widehat{\mathcal{J}}_{a+b} A$, and hence A .
4. Apply axiom K and $\mathcal{J}_{a+b} A \Rightarrow \mathcal{B}_a \mathcal{J}_{a+b} A$ to property 2.
5. Apply axiom K and $\mathcal{J}_{a+b} A \Rightarrow \mathcal{B}_a \mathcal{J}_{a+b} A$ to property 3.

□

Suppose we have some forwarding network S ,

Definition 4. We say that a trust statement $\text{trust}_n(a+b)(P)$, is *shared over* S if $\text{trust}_n(a+b)(P)$ and $\mathcal{J}_{c+b} \text{val}_n(a+b)(P)$ hold for any $c \in \mathbb{A}$ such that $\mathcal{N}_S(c, a, b)$.

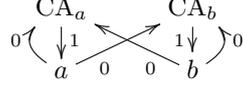
Sharing trust over a network involves sending your claim of trust to all relevant parties. The forwarding network S specifies everyone who can use your claims and may depend on them, and as such gives a specification of where to send your claims.

Theorem 1. Given a forwarding network S and $\mathcal{N}_S(a, b, c)$,

- If both $\text{trust}_n(a+b)(P)$ and $\text{trust}_n(b+c)(P)$ are shared over S , then $\text{trust}_n(a+c)(P)$ is shared over S .
- If both $\text{trust}_n(a+b)(P)$ and $\text{trust}_{n+1}(b+c)(P)$ are shared over S , then $\text{trust}_n(a+c)(P)$ is shared over S .

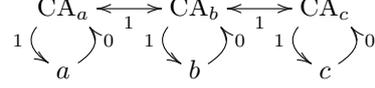
Proof. Suppose $\mathcal{N}_S(a, b, c)$, and both $\text{trust}_n(a+b)(P)$ and $\text{trust}_n(b+c)(P)$ are shared over S . Then it holds that $\text{trust}_n(a+b)(P)$ and $\mathcal{J}_{a+b} \text{val}_n(b+c)(P)$, so by property 5 of Lemma 9, $\text{trust}_n(a+c)(P)$. Supposing moreover $\mathcal{N}_S(d, a, c)$, then by Lemma 5, $\mathcal{N}_S(d, b, c)$ and $\mathcal{N}_S(d, a, b)$, so $\mathcal{J}_{d+a} \text{val}_n(a+b)(P)$ and $\mathcal{J}_{d+b} \text{val}_n(b+c)(P)$, the latter implying $\mathcal{J}_{d+a} \mathcal{J}_{a+b} \text{val}_n(b+c)(P)$. With axiom K extended to \mathcal{J} by Lemma 6, and by property 3 of Lemma 9, we get $\mathcal{J}_{d+a} \text{val}_n(a+c)(P)$. We conclude that $\text{trust}_n(a+c)(P)$ is shared over S . The second property proven similarly, using Properties 4 and 2 of Lemma 9 instead. □

Personal CA Here, everyone has their own CA to consult and check trustworthiness of other agents.



We can derive $(a \overset{0}{\dashrightarrow} b)$ and $(b \overset{0}{\dashrightarrow} a)$.

Mesh PKI Here, everyone has their own CA they can use to advocate validity of their claims. These CAs are connected in a mesh, where each CA trusts other CAs in their ability to authenticate keys. A user then only needs to trust their own CA in being able to authenticate others' claims.



Note that not all CAs need to be directly connected. As long as there is a chain of trust between each of the CAs, trust can be derived.

7 Threshold Trust

There are situations in which a CA is considered unreliable. This can happen for two reasons:

- CAs may be incorrect, as in they are lying or simply wrong about the validity of a claim.
- CAs may be unresponsive regarding trust, as in they have not verified or communicated whether they consider someone trustworthy.⁴

In both cases, we can use multiple CAs to mitigate the issue. In the first situation, we can confirm with multiple CAs whether an entity is trusted, before accepting the entity's claims. In the second situation, we attempt to consult multiple CAs hoping at least some assert their trust in the entity.

Consider Lemma 9 again, and see that when $\mathcal{N}_S(a, i, b)$:

$$\text{trust}_1(a \dashrightarrow i)(P) \wedge \mathcal{I}_{a \dashrightarrow i}(\text{val}_0(i \dashrightarrow b)(P)) \implies (\text{trust}_0(a \dashrightarrow b)(P))$$

We see that trust composition requires two statements; the trust in the CA given by $\text{trust}_0(a \dashrightarrow i)(P)$, and the sharing of validation by the CA given by $\mathcal{I}_{a \dashrightarrow i}(\text{val}_0(i \dashrightarrow b)(P))$. If we think a CA is dishonest, we lack the former statement, and if we think a CA is unresponsive, we lack the latter statement.

Consider the following situation. We have two users a and b , and two CAs named i and j . We assume that trust of a in b gets delegated to both i and j , asserting $\mathcal{N}_S(a, i, b)$ and $\mathcal{N}_S(a, j, b)$: We get the following two results:

- $\text{trust}_1(a \dashrightarrow i)(P) \vee \text{trust}_1(a \dashrightarrow j)(P)$ and $\mathcal{I}_{a \dashrightarrow i}(\text{val}_0(i \dashrightarrow b)(P)) \wedge \mathcal{I}_{a \dashrightarrow j}(\text{val}_0(j \dashrightarrow b)(P))$ implies $\text{trust}_0(a \dashrightarrow b)(P)$. If only one of the two CAs is trusted by a , then both need to assert trust in b to guarantee that trust can be derived.
- $\text{trust}_1(a \dashrightarrow i)(P) \wedge \text{trust}_1(a \dashrightarrow j)(P)$ and $\mathcal{I}_{a \dashrightarrow i}(\text{val}_0(i \dashrightarrow b)(P)) \vee \mathcal{I}_{a \dashrightarrow j}(\text{val}_0(j \dashrightarrow b)(P))$ implies $\text{trust}_0(a \dashrightarrow b)(P)$. If both CAs are trusted, only one needs to assert trust in b .

The above two examples resolve problems created by potential attacks on the network. Possible incorrectness of a CA may be due to an attacker impersonating the CA, and in the first example trust can still be derived if one CA gets corrupted this way. Possible unresponsiveness may be due to an attacker removing a CA's published certificates, and in the second example trust can still be derived if one CA gets disrupted.

In cases where both incorrectness and unresponsiveness are possibilities, we would need to consult at least three CAs to validate a claim. We get *threshold PKI*, with the simplest being a *2-out-of-3* threshold. Given formulas A, B, C , define:

$$\text{2of3}(A, B, C) := (A \wedge B) \vee (A \wedge C) \vee (B \wedge C)$$

Note that $\text{2of3}(A, B, C) \equiv (A \vee B) \wedge (A \vee C) \wedge (B \vee C)$.

Consider again agents a and b , together with three CAs i, j and k , and let $S = \{aib, ajb, akb\}$, hence $\mathcal{N}_S(a, i, b)$, $\mathcal{N}_S(a, j, b)$, and $\mathcal{N}_S(a, k, b)$. We get threshold trust:

⁴Note that the CA may still be responsive regarding forwarding of claims, as that is dealt with separately and can be independently verified.

Lemma 10. *If $\mathcal{N}_S(a, i, b)$, $\mathcal{N}_S(a, j, b)$, and $\mathcal{N}_S(a, k, b)$, then:*
 $2of\mathcal{B}(\text{trust}_1(a+i)(P), \text{trust}_1(a+j)(P), \text{trust}_1(a+k)(P)) \wedge$
 $2of\mathcal{B}(\mathcal{I}_{a+i}\text{val}_0(i+b)(P), \mathcal{I}_{a+j}\text{val}_0(j+b)(P), \mathcal{I}_{a+k}\text{val}_0(k+b)(P))$
 $\implies \text{trust}_0(a+b)(P)$

In [51], a logic is considered in which the trust thresholds are directly baked into the modalities themselves.

8 Wish Modality

We have seen how agents can request information. As discussed before, we have to be careful when using the communication modality, as any message is considered to be a claim made by the sender. As is, there is no direct way to formally communicate facts which are not yet believed without having to lie, except by forwarding other people's claims. This is quite inconvenient in networks based on trust.

A solution is to use a *wish modality* \mathcal{W}_a , which tells us what an agent a wants to be true. The modality fits in the framework of this paper, since it satisfies axiom K; if agent a desires A to be true, and desires B to be true, then one should be able to reason that a wants both of them to be true. In cases where the agent wants either of the two to be true, but not necessarily both, we instead have to write $\mathcal{W}_a(A \vee B)$.

Most commonly, a wish can be used to send a request to another agent: $\mathcal{I}_{a,b}\mathcal{W}_bA$ means b tells a it wishes A to be true. There are several examples of useful wishes one can utter:

- With $\mathcal{I}_{a,b}\mathcal{W}_bt$, b tells a it likes some principal statement t to be true, e.g. key ownership. If a has the ability to make t true, this is a direct request from b to do so.
- With $\mathcal{I}_{a,b}\mathcal{W}_b\mathcal{I}_{c,a}A$, b requests a to send the claim A to agent c . For instance, a is some authority which is requested to send signed credentials to some retailer c .
- $\mathcal{I}_{a,b}\mathcal{W}_b\mathcal{B}_cA$ is a request of b to a to help convince c of a certain fact.

As an example, suppose t is the token expressing that a owns a certain product or key. Suppose we have a store c which can make t true, as specified by $\mathcal{B}_ct \Rightarrow t$. A customer a would like t to be true and sends a request to c , $\mathcal{I}_{c,a}\mathcal{W}_at$. Now, c is willing to make t true if a has the proper credentials, which is expressed by some other token u . So c sends $\mathcal{I}_{a,c}(u \Rightarrow t)$.

Now a needs to convince c of the truth of u . They can do so in several ways. If c does not trust a directly (at least not regarding u), we need an intermediate authority b which c trusts. Either a sends a request to b to validate and send the credentials to c using the above request protocol. Alternatively, a can let c know that b can vouch for the truth of the credentials, and c can request validation from b independently.

One last example is using the wish modality to request information. Suppose agent b wants to know whether A is true or not. In this case, it would make the statement $\mathcal{W}_a((A \Rightarrow \mathcal{B}_aA) \wedge (\neg A \Rightarrow \mathcal{B}_a(\neg A)))$, and send this to someone who may know. Of course, without prearranged obligations or guarantees, requests could be ignored.

9 Key Ownership and Proxy Trust

One of the main motivations of this paper is to describe networks for authenticating ownership of keys. We could consider such keys as entities in their own right, used as aliases to send further messages. Even if one does not know the owner of a key, one may still reason about the belief and interactions of the owner, as represented by the key. In authorization logics, one may use statements of ownership to translate authority from an entity to their owned keys. Here we can similarly use key ownership here to translate trust in an entity to trust in their owned keys.

We take as set of tokens $\mathbb{T} = \{\mathcal{C}_{a,b} \mid a, b \in \mathbb{A}\}$, where $\mathcal{C}_{a,b}$ states a is controlled by b . We consider this as expressing a reflexive and transitive relationship, with $\forall_a \mathcal{C}_{a,a}$ and $\forall_{a,b,c}(\mathcal{C}_{a,b} \Rightarrow \mathcal{C}_{b,c} \Rightarrow \mathcal{C}_{a,c})$ as additional public assumptions. Unlike the delegation network \mathbb{N} , ownership of keys is not public knowledge, and each entity would need to verify on their own whether a certain key is owned by a certain entity.

In public key infrastructures, entities make claims about ownership of keys. One can trust a person directly when they claim to control a key, or one may trust a CA to validate someones claims of ownership

of keys. This is modeled by instantiating the P in trust statements $\text{trust}_n(a+b)(P)$ with $P(c) = \mathcal{C}_{d,c}$ for different choices of $d \in \mathbb{A}$. This implements the claim: I have control of agent d .

Beside being examples of claims we want to validate, ownership of keys can have additional effects on the modalities. We can for instance add the following axioms to Δ :

- $\forall_{a,b}(\mathcal{C}_{a,b} \Rightarrow \mathcal{B}_b \mathcal{C}_{a,b})$.
- $\forall_{a,b}(\mathcal{C}_{a,b} \Rightarrow (\mathcal{B}_a \mathbf{X} \Leftrightarrow \mathcal{B}_b \mathbf{X}))$.
- $\forall_{a,b}(\mathcal{C}_{a,b} \Rightarrow \forall_c(\mathcal{I}_{a,c} \mathbf{X} \Rightarrow \mathcal{I}_{b,c} \mathbf{X}))$.
- $\forall_{a,b}(\mathcal{C}_{a,b} \Rightarrow \forall_c(\mathcal{I}_{c,a} \mathbf{X} \Rightarrow \mathcal{I}_{c,b} \mathbf{X}))$

These statements say the following; agents know which keys they control. If agent a is controlled by agent b , then a and b believe the same things as they are identifiers (alter-egos) of the same entity. Moreover, messages sent to and from a can be considered as sent to and from b respectively.

Using the axioms above, we can *transfer* trust. If you trust an agent, then you can trust any agent controlled by them. This result will allow people to act using a digital identity they control, and still be trusted.

Lemma 11. $\vdash_{\Delta} \forall_{b,c}(\mathcal{C}_{b,c} \Rightarrow \forall_a(\mathcal{V}_{a,c} \mathbf{X} \Rightarrow \mathcal{V}_{a,b} \mathbf{X}))$.

Proof. Suppose $\mathcal{C}_{b,c}$ and $\mathcal{V}_{a,c} A$ hold for some keys a, b, c and statement A . We prove $\mathcal{V}_{a,b} A$.

If $\mathcal{I}_{a,b} A$, then since $\mathcal{C}_{b,c}$ we know by the new axiom that the message was actually sent by c , so $\mathcal{I}_{a,c} A$. By validity $\mathcal{V}_{a,c} A$, we know that c believes it: $\mathcal{B}_c A$. Again by $\mathcal{C}_{b,c}$, we conclude that b believes it too: $\mathcal{B}_b A$.

If $\mathcal{B}_b A$, then by $\mathcal{C}_{b,c}$ and the new axioms, $\mathcal{B}_c A$. Using reliability from $\mathcal{V}_{a,c} A$ we can derive A . \square

Corollary 1. $\vdash_{\Delta} \forall_{a,b,c} \mathcal{B}_a \mathcal{C}_{b,c} \Rightarrow (\mathcal{T}_{a,c} \mathbf{X} \Rightarrow \mathcal{T}_{a,b} \mathbf{X})$.

Proof. Suppose $\mathcal{B}_a \mathcal{C}_{b,c}$ and $\mathcal{T}_{a,c} A = \mathcal{B}_a \mathcal{V}_{a,c} A$ holds for some keys a, b, c and statement A . Then $\mathcal{T}_{a,b} A = \mathcal{B}_a \mathcal{V}_{a,b} A$ is true by lifting Lemma 11 to perspective \mathcal{B}_a using axiom (K). \square

10 Logic Calculi

There are many different equivalent ways of formulating a calculus whose power matches our logic. We highlight mainly three different calculi:

1. A basic intuitionistic sequent calculus for provability, as formulated in Figure 2.
2. A lambda calculus following the Fitch-style proof systems, as used in for instant dependent modal logics. [15]
3. A cut-free sequent calculus for decidable proof search.

We shall briefly describe a version of the second, which is suitable for expressing proofs as terms in a lambda calculus. Note that the logic used in this paper is less powerful than those considered in recent works, e.g. as used in *dependent modal type theory* in [10, 26]. This is in order to facilitate easier proof searching, as is useful when imposing accountability of implied statements, as well as to not unnecessarily introduce unneeded complexity. The logic can of course be extended if needed, possibly sacrificing decidability.

We define *modal contexts* as follows:

$$\Gamma, \Delta := \varepsilon \mid \Gamma, x : A \mid \Gamma, \{\mathcal{M}\}$$

These may contain assumptions in the form of statements A with associated name x for proofs. Additionally, we have modalities $\{\mathcal{M}\}$ expressing a modal shift: e.g. the context $x : A, \{\mathcal{B}_a\}, y : B$ can be read as, supposing A is true, and we consider the perspective of a , where we additionally suppose a believes in B .

Let $[-]$ be the map from contexts to \mathbb{M}^* defined as: $[\varepsilon] = \varepsilon$, $[\Gamma, x : A] = \Gamma$ and $[\Gamma, \{\mathcal{M}\}] = [\Gamma], \mathcal{M}$. See Figure 4 for the Fitch-style lambda calculus. Comparing this to the literature, here we rename **mod** to **lock**, and **unmod** to **key**. Furthermore, we do not apply our additional modal axioms on the contexts directly, instead folding those into the **key** operations. Hence, once one unlocks a term, the modal

$$\begin{array}{c}
\frac{[\Delta] = \varepsilon}{\Gamma, x : A, \Delta \vdash \text{var}(x) : A} \quad \frac{}{\Gamma \vdash * : \top} \quad \frac{\Gamma \vdash P : \perp}{\Gamma \vdash \text{abs}_A(P) : A} \quad \frac{\Gamma, \{\mathcal{M}\} \vdash P : A}{\Gamma \vdash \text{lock}_{\mathcal{M}}(P) : \mathcal{M}A} \\
\frac{\Gamma \vdash P : \mathcal{M}A \quad \mathcal{M} \Rightarrow l \quad [\Delta] = l}{\Gamma, \Delta \vdash \text{key}_{\mathcal{M} \Rightarrow l}(P) : A} \quad \frac{\Gamma \vdash P : A \quad \Gamma \vdash Q : A \Rightarrow B}{\Gamma \vdash P \cdot Q : B} \quad \frac{\Gamma, x : A \vdash P : B}{\Gamma \vdash \lambda x : A.P : A \Rightarrow B} \\
\frac{\Gamma \vdash P : A_1 \wedge A_2}{\Gamma \vdash \pi_i(P) : A_i} \quad \frac{\Gamma \vdash P : A_1 \quad \Gamma \vdash Q : A_2}{\Gamma \vdash (P, Q) : A_1 \wedge A_2} \\
\frac{\Gamma, x : A \vdash P : C \quad \Gamma, y : B \vdash Q : C \quad \Gamma \vdash R : A \vee B}{\Gamma \vdash \text{case}_C(R)\{\nu_1(x) \mapsto P, \nu_2(y) \mapsto Q\} : C} \quad \frac{\Gamma \vdash P : A_i}{\Gamma \vdash \nu_i(P) : A_1 \vee A_2}
\end{array}$$

Figure 4: Fitch-style lambda calculus for our logic

context gets fixed. Shifting modal contexts now becomes a meta operation on terms, similar to variable substitutions.

As an example, a possible lambda term for the sequent $x : \mathcal{B}_a(\mathcal{I}_{a+b}A \Rightarrow A) \vdash \mathcal{I}_{a+b}A \Rightarrow \mathcal{B}_aA$ is $\lambda y : \mathcal{I}_{a+b}A. \text{lock}_{\mathcal{B}_a}(\text{key}_{\mathcal{B}_a \Rightarrow \mathcal{B}_a}(\text{var}(x)) \cdot \text{lock}_{\mathcal{I}_{a+b}}(\text{key}_{\mathcal{I}_{a+b} \Rightarrow \mathcal{B}_a \cdot \mathcal{I}_{a+b}}(\text{var}(y))))$.

10.1 Meta constructions

The intuitionistic modal logic forms a foundation for expressing a plethora of concepts. We can construct a variety of other operations using the above tools. For instance, with \perp marking absurdity, we can define negation $\neg A$ as $\neg A = A \Rightarrow \perp$ following the standard intuitionistic tradition.

Another useful thing we can define is conjunction and disjunction over finite sets of formulas. Given a finite subset of $S \subseteq_{\text{fin}} \mathbb{F}$, we define the following:

- $\bigwedge S$ is the conjunction over S , which holds precisely when all $A \in S$ hold.
- $\bigvee S$ is the disjunction over S , which holds precisely if some $A \in S$ holds.

These are defined inductively, iterating binary conjunctions and disjunctions. In particular, $\bigwedge \emptyset \equiv \top$ and $\bigvee \emptyset \equiv \perp$.

When describing distributed networks, we consider the additional finite set of agents \mathbb{A} . Both modalities and tokens may refer to specific agents, marking their perspectives and primitive statements of interest. As such, formulas can refer to specific agents, and the referenced agent could be used as a parameter. Given such a parametrized formula $f : \mathbb{A} \rightarrow \mathbb{F}$, and given a subset $S \subseteq \mathbb{A}$, we write:

- $\forall_{x \in S} f(x)$ for the universal quantification, equivalent to $\bigwedge \{f(a) \mid a \in S\}$.
- $\exists_{x \in S} f(x)$ for the existential quantification, equivalent to $\bigvee \{f(a) \mid a \in S\}$.

We will write $\forall_x f(x)$ and $\exists_x f(x)$ in case $S = \mathbb{A}$, and $\forall_{P(x)} f(x)$ and $\exists_{P(x)} f(x)$ if $S = \{a \in \mathbb{A} \mid P(a)\}$ with P some predicate on agents.

These quantifiers could be expressed using formulas polymorphic over agent variables as well. However, we would like to avoid such unnecessary language generalisations at this point, focusing instead on the underlying theory of trust.

10.2 Notes on Decidable Proof Search

Let us briefly address how the decomposability property on axioms facilitates easier proof search. Note that given decomposability, we can reduce any proof to one only using $\text{key}_{\mathcal{M} \Rightarrow \mathcal{N}}$ and $\text{key}_{\mathcal{M} \Rightarrow \mathcal{N} \cdot \mathcal{M} \ominus \mathcal{N}}$ operators. Supposing we want to prove $\mathcal{N}B$ with a context containing $\mathcal{M}A$, then we know it is sufficient to use the formula in context as $\mathcal{N}A$ and/or $\mathcal{N}(\mathcal{M} \ominus \mathcal{N})A$ depending which axioms are available.

More concretely, proof search may happen in a cut-free variant of the logic extending the usual calculi and their cut elimination proofs [43]. These calculi use contexts excluding the $\{\mathcal{M}\}$ modal locks, instead defining for each modality \mathcal{N} an explicit operation \mathcal{N}^{-1} on contexts where $\mathcal{N}^{-1}(\Gamma)$ contains A for any $\mathcal{M}A \in \Gamma$ such that $\mathcal{M} \Rightarrow \mathcal{N}$, and $(\mathcal{M} \ominus \mathcal{N})A$ for any $\mathcal{M}A \in \Gamma$ such that $\mathcal{M} \Rightarrow \mathcal{N} \cdot \mathcal{M} \ominus \mathcal{N}$. We then handle modalities with the single rule:

$$\frac{\mathcal{M}^{-1}(\Gamma) \vdash A}{\Gamma \vdash \mathcal{M}A}$$

We then establish that the associated calculus has cut elimination, and note that we can associate a subformula order establishing that a proof search must terminate. We get that if \mathbb{M} is finite and Ax is a reflexive, transitive and decomposable set of axioms, the associated calculus has decidable proof search. Decidability may moreover be derived by proving our axioms as an instance of those systems specified an proven to be decidable in [23], though this has not been verified.

This need not mean that the proof search is practical. Moreover, this is sensitive to the chosen axioms; note the need for decomposability and the lack of axioms of the form $\mathcal{M}\mathcal{N}A \Rightarrow \mathcal{R}A$ in our formalism.

11 Kripke Model

The proof system introduced in the previous section connects our logic to a *categorical model* [26], describing explicitly the handling and manipulation of proofs. We shall furthermore consider a *Kripke model* for establishing soundness and giving an interpretation of the formulas in terms of possible worlds which describe who knows and says what.

For the Kripke semantics, we consider a traditional variant [12, 17] of intuitionistic modal logic, and also briefly consider Simpson's [47] version.

Definition 5. A *modal Kripke frame* on (\mathbb{T}, \mathbb{M}) consists of a quadruple (W, \leq, S_-, R_-) where W is a set of worlds, \leq is a preorder on W , S associates to each token $t \in \mathbb{T}$ a subset $S_t \subseteq W$ on W and R associates to each modality $\mathcal{M} \in \mathbb{M}$ a binary endorelation $R_{\mathcal{M}} \subseteq W^2$ on W such that:

1. If $v \leq w$ and $v \in S_t$, then $w \in S_t$.
2. If $v \leq w$ and $wR_{\mathcal{M}}v'$, then there is a $v' \in W$ such that $vR_{\mathcal{M}}v'$ and $v' \leq w'$.

At their core, Kripke frames consider possible worlds W , where each world not only describes which fundamental statements are true with subsets S_t , but also what is communicated and believed. The relation $R_{\mathcal{B}_a}$ for instance describes for each v all possible worlds w agent a thinks they might be in. So if $vR_{\mathcal{B}_a}w \in S_t$, $vR_{\mathcal{B}_a}k \notin S_t$, then in world v agent a is unsure of whether statement t is true, hence \mathcal{B}_at is not true in v .

Last but not least, \leq implements an interpretation of intuitionistic logic. In this situation, statements are not simply true or false. Some statements either *become* true or proven, and some statements *remain* false or unproven. As such, we can think of \leq as a progression of time; the more agents communicate with each other, the more they learn. This progression may be nondeterministic, as multiple distinct worlds may spawn from one world.

Note that in some definitions of Kripke frames, the second condition is removed and instead this property is baked into the interpretation of modal formulas. In our case, some conditions on R cannot be avoided regardless, since they need to accommodate axioms like $\mathcal{I}_{a+b}A \Rightarrow \mathcal{B}_a\mathcal{I}_{a+b}A$. Let Ax be a set of unfolding modal axioms on \mathbb{M} .

Definition 6. A modal Kripke frame (W, \leq, S, R) conforms to Ax if $\mathcal{M} \Rightarrow \mathcal{N}_1 \dots \mathcal{N}_n \in \text{Ax}$ implies $R_{\mathcal{N}_1}; \dots; R_{\mathcal{N}_n} \subseteq R_{\mathcal{M}}$.

Given a Kripke frame, we define *satisfiability* of a formula A in a world $w \in W$ inductively on formulas:

- $w \models \top$, $w \not\models \perp$, and $w \models t$ iff $w \in S_t$
- $w \models \mathcal{M}A$ iff $\forall v.(wR_{\mathcal{M}}v \implies v \models A)$
- $w \models A \Rightarrow B$ iff $\forall v.((w \leq v \wedge v \models A) \implies v \models B)$
- $w \models A \wedge B$ iff $w \models A$ and $w \models B$
- $w \models A \vee B$ iff $w \models A$ or $w \models B$

Some immediate properties:

- If $w \models A$ and $w \leq v$ then $v \models A$.
- If the frame conforms to Ax and $\mathcal{M} \Rightarrow \mathcal{N}_1 \dots \mathcal{N}_n \in \text{Ax}$, then $w \models \mathcal{M}A \Rightarrow \mathcal{N}_1 \dots \mathcal{N}_n A$ for each $A \in \mathbb{F}$ and $w \in W$.

Theorem 2. If $\vdash A$ is provable, then $w \models A$ for any modal Kripke frame (W, \leq, S, R) and world $w \in W$.

12 Conclusions

We end this paper with some final considerations.

12.1 Regarding Privacy

In this paper, we have not considered *privacy*. We have considered responsiveness based on consent, for instance with the wish modality, where an authority has promised to share some fact after being asked to do so. However, this does not prevent the same authority from sharing this fact without consent of the owner of the information.

Privacy specifically considers statements about entities *not* believing or receiving messages about some sensitive piece of data. Adding privacy may be an interesting next step, which would involve focusing more on negative statements, and would potentially necessitate adding further modalities.

12.2 Universal Knowledge

Both the modalities \mathcal{B}_Ω and \Box describe a kind of public knowledge. The main difference is that we have $\Box X \Rightarrow X$, but not for \mathcal{B}_Ω . For instance if we can prove $\mathcal{B}_\Omega(A_1 \wedge \dots \wedge A_n \Rightarrow B)$ we know that everyone (except us) can prove $A_1, \dots, A_n \vdash B$.

The \Box can be used to make universally held assumptions. This particular use of the \Box modality goes back to early proof theory by Gentzen [25], and is inspired by the work on *dual-context modal logic* developed in [37]. In the latter, a second context is used to describe formulas which are marked by the \Box modality, and an equivalence is established between dual-context modal logic and modal logic with \Box . We could further generalize the logic of this paper to use \Box for marking public statements. In order for this new modality to describe public knowledge, we need to add the following extra assumptions to the logic:

- $\Box A \Rightarrow A$.
- $\Box A \Rightarrow \mathcal{M}\Box A$ for any modality \mathcal{M} .

As a result, \Box satisfies the usual K4 axioms from modal logic.

Using the \Box modality would allow one to reason about public statements internally, an agent may reason about consequences of releasing public statements. However, it seems counter intuitive to think of public statements as anything but universally known, and hence we keep the fact whether a statement is public completely external in this paper.

One could also generalise the public modality further. It might be useful to instead specify different scopes of knowledge and groups of entities, using a preorder on modalities: $\mathcal{M} \leq \mathcal{N}$ if \mathcal{N} is aware of everything known by the perspective \mathcal{M} . This can be expressed with axioms:

- $\mathcal{M}A \Rightarrow \mathcal{N}A$
- $\mathcal{M}A \Rightarrow \mathcal{N}\mathcal{M}A$.

An exploration of such systems is subject to future research.

12.3 Evaluating Risk of Trust

Trust may not always result in certainty. There is always a risk when building ones guarantees based on statements from other entities. As such, it is worthwhile to determine some risk related to a proof. Though this may be done by generalizing to probabilistic modal logic, a lot could already be achieved by simply analyzing proofs within the current logic. The fact of the matter is, that assuming sufficient trust, certain guarantees can be made. The risk is the trust assumptions themselves.

As such, one can collect different proofs of the same guarantee, and analyze the sets of trust assumptions necessary to make the proof. Then one could associate a certain level of risk to each assumption, depending on who needs to be trusted to get the result. This could be a specific risk (e.g. 5 percent), or some unspecified constant risk ε . For instance, in the 2-3 threshold example, associating a 5 percent risk to trusting each of the three CAs gets us a total failure risk of only 0.725 percent, which is a significantly reduced risk.

References

- [1] Martín Abadi. Logic in access control. In *Proceedings of the 18th Annual IEEE Symposium on Logic in Computer Science*, LICS '03, page 228, USA, 2003. IEEE Computer Society.
- [2] Martín Abadi. Access control in a core calculus of dependency. *SIGPLAN Not.*, 41(9):263–273, September 2006.
- [3] Martín Abadi, Michael Burrows, Butler Lampson, and Gordon Plotkin. A calculus for access control in distributed systems. *ACM Trans. Program. Lang. Syst.*, 15(4):706–734, September 1993.
- [4] Matteo Acclavio, Davide Catta, and Federico Olimpieri. Canonicity in modal lambda calculus. *ArXiv*, abs/2304.05465, 2023.
- [5] Alessandro Aldini. A calculus for trust and reputation systems. In Jianying Zhou, Nurit Gal-Oz, Jie Zhang, and Ehud Gudes, editors, *Trust Management VIII*, pages 173–188, Berlin, Heidelberg, 2014. Springer Berlin Heidelberg.
- [6] Benjamin Aziz and Geoff Hamilton. Modelling and analysis of PKI-based systems using process calculi. *International Journal of Foundations of Computer Science*, 18(03):593–618, 2007.
- [7] Moritz Becker, Cedric Fournet, and Andrew Gordon. Design and semantics of a decentralized authorization language. In *20th IEEE Computer Security Foundations Symposium (CSF'07)*, pages 3–15, 2007.
- [8] Moritz Y. Becker. Information flow in trust management systems. *J. Comput. Secur.*, 20(6):677–708, November 2012.
- [9] Gianluigi Bellin, Valeria De Paiva, and Eike Ritter. Extended curry-howard correspondence for a basic constructive modal logic. In *Proceedings of the 2nd Workshop on Methods for Modalities*, Amsterdam, Netherlands, November 2001.
- [10] Lars Birkedal, Ranald Clouston, Bassel Manna, Rasmus Ejlers Møgelberg, Andrew M. Pitts, and Bas Spitters. Modal dependent type theory and dependent right adjoints. *Mathematical Structures in Computer Science*, 30(2):118–138, 2020.
- [11] Andreas Blass and Yuri Gurevich. Two forms of one useful logic: Existential fixed point logic and liberal datalog. *Bull. EATCS*, 95:164–182, 2008.
- [12] Milan Božić and Kosta Došen. Models for normal intuitionistic modal logics. *Studia Logica*, 43(3):217–245, 1984.
- [13] Michael Burrows, Martin Abadi, and Roger Needham. A logic of authentication. *ACM Trans. Comput. Syst.*, 8(1):18–36, February 1990.
- [14] CEF eDelivery. Cef edelivery building block v1.2. trust models guidance, May 2018.
- [15] Ranald Clouston. Fitch-style modal lambda calculi. In Christel Baier and Ugo Dal Lago, editors, *Foundations of Software Science and Computation Structures*, pages 258–275, Cham, 2018. Springer International Publishing.
- [16] Mehdi Dastani, Andreas Herzig, Joris Hulstijn, and Leendert van der Torre. Inferring trust. In João Leite and Paolo Torroni, editors, *Computational Logic in Multi-Agent Systems*, pages 144–160, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.
- [17] Kosta Dosen. Models for stronger normal intuitionistic modal logics. *Studia Logica*, 44:39–70, 1985.
- [18] Besik Dundua and Levan Uridia. Trust and belief , interrelation. In *Proceedings of the Third Workshop on Agreement Technologies, Bahia Blanca, Argentina, November 1, 2010*, 2010.
- [19] H. El Bakkali and B.I. Kaitouni. A predicate calculus logic for the PKI trust model analysis. In *Proceedings IEEE International Symposium on Network Computing and Applications. NCA 2001*, pages 368–371, 2001.

- [20] Ulrik Frendrup, Hans Hüttel, and Jesper Nyholm Jensen. Modal logics for cryptographic processes. *Electronic Notes in Theoretical Computer Science*, 68(2):124–141, 2002. EXPRESS’02, 9th International Workshop on Expressiveness in Concurrency (Satellite Workshop of CONCUR 2002).
- [21] Deepak Garg and Martín Abadi. A modal deconstruction of access control logics. In Roberto Amadio, editor, *Foundations of Software Science and Computational Structures*, pages 216–230, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg.
- [22] Deepak Garg, Lujo Bauer, Kevin D. Bowers, Frank Pfenning, and Michael K. Reiter. A linear logic of authorization and knowledge. In *Proceedings of the 11th European Conference on Research in Computer Security*, ESORICS’06, page 297–312, Berlin, Heidelberg, 2006. Springer-Verlag.
- [23] Deepak Garg, Valerio Genovese, and Sara Negri. Countermodels from sequent calculi in multi-modal logics. In *2012 27th Annual IEEE Symposium on Logic in Computer Science*, pages 315–324, 2012.
- [24] Deepak Garg and F. Pfenning. Non-interference in constructive authorization logic. In *19th IEEE Computer Security Foundations Workshop (CSFW’06)*, pages 11 pp.–296, 2006.
- [25] G. Gentzen. Untersuchungen über das logische Schließen I. *Mathematische Zeitschrift*, 39:176–210, 1935.
- [26] Daniel Gratzer, G. A. Kavvos, Andreas Nuyts, and Lars Birkedal. Multimodal dependent type theory. In *Proceedings of the 35th Annual ACM/IEEE Symposium on Logic in Computer Science*, LICS ’20, page 492–506, New York, NY, USA, 2020. Association for Computing Machinery.
- [27] Yuri Gurevich and Itay Neeman. Dkal: Distributed-knowledge authorization language. In *2008 21st IEEE Computer Security Foundations Symposium*, pages 149–162, 2008.
- [28] Yuri Gurevich and Itay Neeman. DKAL 2 - A Simplified and Improved Authorization Language. Technical Report MSR-TR-2009-11, Microsoft Research, 2009.
- [29] Yuri Gurevich and Arnab Roy. Operational semantics for dkal: Application and analysis. In Simone Fischer-Hübner, Costas Lambrinoudakis, and Günther Pernul, editors, *Trust, Privacy and Security in Digital Business*, pages 149–158, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg.
- [30] Andreas Herzig. Modal probability, belief, and actions. *Fundamenta Informaticae*, 57(2-4):323–344, 2003.
- [31] Andrew K. Hirsch and Michael R. Clarkson. Belief semantics of authorization logic. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *2013 ACM SIGSAC Conference on Computer and Communications Security, CCS’13, Berlin, Germany, November 4-8, 2013*, pages 561–572. ACM, 2013.
- [32] Andrew K. Hirsch, Pedro H. Azevedo de Amorim, Ethan Cecchetti, Ross Tate, and Owen Arden. First-order logic for flow-limited authorization. In *33rd IEEE Computer Security Foundations Symposium, CSF 2020, Boston, MA, USA, June 22-26, 2020*, pages 123–138. IEEE, 2020.
- [33] Jinwei Hu, Yan Zhang, Ruixuan Li, and Zhengding Lu. A logic for authorization provenance. In *Proceedings of the 5th ACM Symposium on Information, Computer and Communications Security, ASIACCS ’10*, page 238–249, New York, NY, USA, 2010. Association for Computing Machinery.
- [34] Jingwei Huang and David Nicol. A calculus of trust and its application to PKI and identity management. In *Proceedings of the 8th Symposium on Identity and Trust on the Internet, IDtrust ’09*, page 23–37, New York, NY, USA, 2009. Association for Computing Machinery.
- [35] Zeinab Iranmanesh, Morteza Amini, and Rasool Jalili. A logic for multi-domain authorization considering administrators. In *2008 IEEE Workshop on Policies for Distributed Systems and Networks*, pages 189–196, 2008.
- [36] Aleksander Kamenik, Peeter Laud, Alisa Pankova, Triin Siil, and Nikita Snetkov. SPOF2.3 - eID infrastruktuuri usaldusmodel (Trust Model for eID Infrastructure). Research report D-16-158, Cybernetica AS, October 2022. <https://www.ria.ee/sites/default/files/documents/2022-11/SPOF-2.3-eid-infrastruktuuri-usaldusmodel-28.10.2022.pdf>.

- [37] G. A. Kavvos. Dual-context calculi for modal logic. In *2017 32nd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–12, 2017.
- [38] Line Kofoed and Sankarshan Mukopadhyay. Principles of SSI V3. <https://sovrin.org/principles-of-ssi/>, 2022.
- [39] Reto Kohlas, Jacek Jonczyk, and Rolf Haenni. A trust evaluation method based on logic and probability theory. In Yücel Karabulut, John Mitchell, Peter Herrmann, and Christian Damsgaard Jensen, editors, *Trust Management II*, pages 17–32, Boston, MA, 2008. Springer US.
- [40] Christopher Leturc and Grégory Bonnet. A normal modal logic for trust in the sincerity. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '18*, page 175–183, Richland, SC, 2018. International Foundation for Autonomous Agents and Multiagent Systems.
- [41] Churn-Jung Liau. Belief, information acquisition, and trust in multi-agent systems—a modal logic formulation. *Artificial Intelligence*, 149(1):31–60, 2003.
- [42] Aleksandar Nanevski, Frank Pfenning, and Brigitte Pientka. Contextual modal type theory. *ACM Trans. Comput. Logic*, 9(3), June 2008.
- [43] Frank Pfenning. Structural cut elimination: I. intuitionistic and classical logic. *Information and Computation*, 157(1):84–141, 2000.
- [44] P.V. Rangan. An axiomatic basis of trust in distributed systems. In *Proceedings. 1988 IEEE Symposium on Security and Privacy*, pages 204–211, 1988.
- [45] Drummond Reed, Riëks Joosten, and Oskar van Deventer. The basic building blocks of SSI. In Alex Preukschat and Drummond Reed, editors, *Self-Sovereign Identity: Decentralized Digital Identity and Verifiable Credentials*, pages 21–38. Manning Publications Co., 2021.
- [46] Helena Rifà-Pous and Jordi Herrera-Joancomartí. An interdomain PKI model based on trust lists. In Javier López, Pierangela Samarati, and Josep L. Ferrer, editors, *Public Key Infrastructure, 4th European PKI Workshop: Theory and Practice, EuroPKI 2007, Palma de Mallorca, Spain, June 28-30, 2007, Proceedings*, volume 4582 of *Lecture Notes in Computer Science*, pages 49–64. Springer, 2007.
- [47] Alex K. Simpson. *The proof theory and semantics of intuitionistic modal logic*. PhD thesis, University of Edinburgh, UK, 1994.
- [48] Munindar P. Singh. Trust as dependence: a logical approach. In *Adaptive Agents and Multi-Agent Systems*, 2011.
- [49] Emin Gün Sirer, Willem de Bruijn, Patrick Reynolds, Alan Shieh, Kevin Walsh, Dan Williams, and Fred B. Schneider. Logical attestation: an authorization architecture for trustworthy computing. In *Proceedings of the Twenty-Third ACM Symposium on Operating Systems Principles, SOSP '11*, page 249–264, New York, NY, USA, 2011. Association for Computing Machinery.
- [50] Niels Voorneveld. Forward proof search for intuitionistic multimodal k logics. In Gian Luca Pozzato and Tarmo Uustalu, editors, *Automated Reasoning with Analytic Tableaux and Related Methods*, pages 335–353, Cham, 2026. Springer Nature Switzerland.
- [51] Niels Voorneveld. Threshold trust logic. In Raimundas Matulevičius, Liina Kamm, and Mubashar Iqbal, editors, *Secure IT Systems, 30th Nordic Conference, NordSec 2025, Tartu, Estonia, November 12-13, 2025, Proceedings*, volume 16325 of *Lecture Notes in Computer Science*, March 2026.
- [52] Frank Wolter and Michael Zakharyashev. Intuitionistic modal logic. In Andrea Cantini, Ettore Casari, and Pierluigi Minari, editors, *Logic and Foundations of Mathematics: Selected Contributed Papers of the Tenth International Congress of Logic, Methodology and Philosophy of Science, Florence, August 1995*, pages 227–238. Springer Netherlands, Dordrecht, 1999.

A Operational Semantics for Modal Lambda Calculus

In this appendix, we expand on the semantics of our lambda calculus. For $l \in \mathbb{M}^*$, we inductively define $\Gamma, \{l\}$ as:

- $\Gamma, \{\} = \Gamma$,
- $\Gamma, \{\mathcal{M}, l\} = \Gamma, \{\mathcal{M}\}, \{l\}$.

First of all, we have weakening in two directions:

- If $\Gamma, \Delta \vdash P : A$ then $\Gamma, x : A, \Delta \vdash P : A$.
- If $\Gamma \vdash P : A$ then $\{\mathcal{M}\}, \Gamma \vdash P : A$.

We write $l \ni r_1 \mid \dots \mid r_n$ if $l = \mathcal{M}_1 \cdot \dots \cdot \mathcal{M}_n$ and $\mathcal{M}_i \ni r_i$ for each i . Given $[\Gamma] \ni l_1 \mid \dots \mid l_n$ we can define the substitution $\Gamma \langle l_1 \mid \dots \mid l_n \rangle$ as: $\varepsilon \langle \rangle = \varepsilon$, and $(\Gamma, x : A) \langle l_1 \mid \dots \mid l_n \rangle = \Gamma \langle l_1 \mid \dots \mid l_n \rangle, x : A$, and $(\Gamma, \{\mathcal{M}\}) \langle l_1 \mid \dots \mid l_n \rangle = \Gamma \langle l_1 \mid \dots \mid l_{n-1} \rangle, \{\mathcal{M}\}$. Then $[\Gamma \langle l_1 \mid \dots \mid l_n \rangle] = l_1, \dots, l_n$. This modal substitution can be extended to terms as well, allowing for context shifting.

Lemma 12. *If $[\Gamma] \ni l_1 \mid \dots \mid l_n$ and $\Gamma \vdash P : A$ then $\Gamma \langle l_1 \mid \dots \mid l_n \rangle \vdash P \langle l_1 \mid \dots \mid l_n \rangle : A$ for some $P \langle l_1 \mid \dots \mid l_n \rangle$.*

Proof. We define $P \langle l_1 \mid \dots \mid l_n \rangle$ by induction:

- $\text{var}(x) \langle l_1 \mid \dots \mid l_n \rangle = \text{var}(x)$.
- $* \langle l_1 \mid \dots \mid l_n \rangle = *$.
- $\text{abs}_A(P) \langle l_1 \mid \dots \mid l_n \rangle = \text{abs}_A(P \langle l_1 \mid \dots \mid l_n \rangle)$.
- $\text{lock}_{\mathcal{M}}(P) \langle l_1 \mid \dots \mid l_n \rangle = \text{lock}_{\mathcal{M}}(P \langle l_1 \mid \dots \mid l_n \mid \mathcal{M} \rangle)$.
- Suppose $\Gamma \vdash P : \mathcal{M}A$, $\mathcal{M} \ni r$, $[\Delta] = r$ and $[\Gamma, \Delta] \ni l_1 \mid \dots \mid l_n$. We can divide the last statement into $[\Gamma] \ni l_1 \mid \dots \mid l_i$ and $[\Delta] = r \ni l_{i+1} \mid \dots \mid l_n$. By transitivity, $\mathcal{M} \ni l_{i+1}, \dots, l_n$, so we can define $\text{key}_{\mathcal{M} \ni r}(P) \langle l_1 \mid \dots \mid l_n \rangle = \text{key}_{\mathcal{M} \ni l_{i+1}, \dots, l_n}(P \langle l_1 \mid \dots \mid l_i \rangle)$.

□

We can also perform partial modal substitutions. For instance, if $[\Gamma] = \mathcal{M}_1, \dots, \mathcal{M}_n, [\Delta]$ and $[\Delta] \ni l_1 \mid \dots \mid l_m$, then we can write $(-)\langle l_1 \mid \dots \mid l_n \rangle$ as applied to Γ and terms in context Γ as a shorthand for $(-)\langle \mathcal{M}_1 \mid \dots \mid \mathcal{M}_n \mid l_1 \mid \dots \mid l_n \rangle$.

Similarly, we can define variable substitution. Given $\Gamma, x : A, \Delta \vdash P : B$ and $\Gamma \vdash Q : A$ there is a $P[Q/x]$ such that $\Gamma, \Delta \vdash P[Q/x] : B$. This implements a cut rule. We get the following normalisation procedure, which tidies up the terms and hence the proofs.

- $\text{key}_{\mathcal{M} \ni l}(\text{lock}_{\mathcal{M}}(P)) \rightsquigarrow P \langle l \rangle$.
- $(\lambda x : A. P) \cdot Q \rightsquigarrow P[Q/x]$.
- $\pi_i((P_1, P_2)) \rightsquigarrow P_i$.
- $\text{case}_C(\nu_i(R))\{\nu_1(x_1) \mapsto P_1, \nu_2(x_2) \mapsto P_2\} \rightsquigarrow P_i[R/x_i]$.
- $\text{abs}_A(\text{abs}_{\perp}(P)) \rightsquigarrow \text{abs}_A(P)$.
- $\text{abs}_{A \Rightarrow B}(P) \cdot Q \rightsquigarrow \text{abs}_B(P)$.
- $\pi_i(\text{abs}_{A_1 \wedge A_2}(P)) \rightsquigarrow \text{abs}_{A_i}(P)$.
- $\text{case}_C(\text{abs}_{A \vee B}(R)\{\dots\}) \rightsquigarrow \text{abs}_C(R)$.

B Proof of Decidability

Suppose we have a finite set of modalities \mathbb{M} and a reflexive, transitive and decomposable set of axioms Ax . We shall write $\mathcal{M} \Rightarrow l$ to mean that $\mathcal{M} \Rightarrow l \in \text{Ax}$. We take the partial map $\ominus : \mathbb{M}^2 \rightarrow \mathbb{M}$, where $\mathcal{M} \ominus \mathcal{N}$ is defined as a witness to the decomposability property. So $\mathcal{M} \Rightarrow \mathcal{N} \cdot (\mathcal{M} \ominus \mathcal{N})$ if and only if $\mathcal{M} \ominus \mathcal{N}$ is defined, and if $\mathcal{M} \Rightarrow \mathcal{N} \cdot l$ with l having at least one element, then $\mathcal{M} \ominus \mathcal{N}$ is defined and $\mathcal{M} \ominus \mathcal{N} \Rightarrow l$. We write $\mathcal{M} \ominus \mathcal{N} \downarrow$ if it is defined, and $\mathcal{M} \ominus \mathcal{N} \uparrow$ if it is not.

A *flat context* Γ is given by a list of formulas. In other words, it is a context without any $\{\mathcal{M}\}$ in it. For any modality \mathcal{M} , we define a function $(_) \ominus \mathcal{M}$ on flat contexts (which is written as $\mathcal{M}^{-1}(_)$ in the main body of the paper), where:

- $(_) \ominus \mathcal{N} = (_)$.
- $(\Gamma, \mathcal{M}B) \ominus \mathcal{N} = \begin{cases} \Gamma \ominus \mathcal{N}, B, (\mathcal{M} \ominus \mathcal{N})B & \text{if } \mathcal{M} \Rightarrow \mathcal{N}, \text{ and } \mathcal{M} \ominus \mathcal{N} \downarrow \\ \Gamma \ominus \mathcal{N}, B & \text{if } \mathcal{M} \Rightarrow \mathcal{N}, \text{ and } \mathcal{M} \ominus \mathcal{N} \uparrow \\ \Gamma \ominus \mathcal{N}, (\mathcal{M} \ominus \mathcal{N})B & \text{if } \neg \mathcal{M} \Rightarrow \mathcal{N}, \text{ and } \mathcal{M} \ominus \mathcal{N} \downarrow \\ \Gamma \ominus \mathcal{N} & \text{if } \neg \mathcal{M} \Rightarrow \mathcal{N}, \text{ and } \mathcal{M} \ominus \mathcal{N} \uparrow \end{cases}$
- $(\Gamma, B) \ominus \mathcal{N} = \Gamma \ominus \mathcal{N}$ for any other B .

We write $\Gamma \subseteq \Delta$ if any formula of Γ is in Δ . We write $\Gamma \sqsubseteq \Delta$ if any formula of Γ is either in Δ , or of the form $\mathcal{M}C$ with Δ containing $\mathcal{N}C$ for some \mathcal{N} such that $\mathcal{N} \Rightarrow \mathcal{M}$. This is capturing the fact that a formula $\mathcal{N}C$ is more general than $\mathcal{M}C$ if $\mathcal{N} \Rightarrow \mathcal{M}$. Hence in $\Gamma \sqsubseteq \Delta$, Δ is a stronger set of assumptions than Γ . By transitivity of Ax , \sqsubseteq is transitive.

Lemma 13. *If $\Gamma \sqsubseteq \Delta$, then $\Gamma \ominus \mathcal{M} \sqsubseteq \Delta \ominus \mathcal{M}$*

Proof. For $C \in (\Gamma \ominus \mathcal{M})$, either $C = C'$, $\mathcal{N}C' \in \Gamma$ with $\mathcal{N} \Rightarrow \mathcal{M}$, or $C = (\mathcal{N} \ominus \mathcal{M})C'$ and $\mathcal{N}C' \in \Gamma$. Regardless of the case, Δ has $\mathcal{R}C'$ such that $\mathcal{R} \Rightarrow \mathcal{N}$ (note that \mathcal{R} could be \mathcal{N}).

- If $C = C'$, and $\mathcal{N}C' \in \Gamma$ with $\mathcal{N} \Rightarrow \mathcal{M}$. By transitivity, $\mathcal{R} \Rightarrow \mathcal{M}$, hence $C = C' \in (\Delta \ominus \mathcal{M})$.
- If $C = (\mathcal{N} \ominus \mathcal{M})C'$, and $\mathcal{N}C' \in \Gamma$, then $\mathcal{R} \Rightarrow \mathcal{M} \cdot (\mathcal{N} \ominus \mathcal{M})$, hence $\mathcal{R} \ominus \mathcal{M}$ is defined and $\mathcal{R} \ominus \mathcal{M} \Rightarrow \mathcal{N} \ominus \mathcal{M}$ by definition. Hence $(\mathcal{R} \ominus \mathcal{M})C' \in (\Delta \ominus \mathcal{M})$, which covers for $C = (\mathcal{N} \ominus \mathcal{M})C'$.

□

Lemma 14. *If $\mathcal{M} \Rightarrow \mathcal{N}$, then $(\Gamma \ominus \mathcal{M}) \sqsubseteq (\Gamma \ominus \mathcal{N})$.*

Proof. For $C \in (\Gamma \ominus \mathcal{M})$, then we have two cases:

- If $C = C'$, $\mathcal{R}C' \in \Gamma$ with $\mathcal{R} \Rightarrow \mathcal{M}$. By transitivity, $\mathcal{R} \Rightarrow \mathcal{N}$, hence $C = C' \in (\Gamma \ominus \mathcal{N})$.
- If $C = (\mathcal{R} \ominus \mathcal{M})C'$ and $\mathcal{R}C' \in \Gamma$. By transitivity, $\mathcal{R} \Rightarrow \mathcal{N} \cdot (\mathcal{R} \ominus \mathcal{M})$, hence $(\mathcal{R} \ominus \mathcal{N})$ is defined and $(\mathcal{R} \ominus \mathcal{N}) \Rightarrow (\mathcal{R} \ominus \mathcal{M})$. So $(\mathcal{R} \ominus \mathcal{N})C' \in (\Gamma \ominus \mathcal{N})$ which covers for $C = (\mathcal{R} \ominus \mathcal{M})C'$.

□

Lemma 15. *If $\mathcal{M} \ominus \mathcal{N}$ is defined, then $(\Gamma \ominus \mathcal{M}) \sqsubseteq ((\Gamma \ominus \mathcal{N}) \ominus (\mathcal{M} \ominus \mathcal{N}))$.*

Proof. For $C \in (\Gamma \ominus \mathcal{M})$, then we have two cases:

- If $C = C'$, $\mathcal{R}C' \in \Gamma$ with $\mathcal{R} \Rightarrow \mathcal{M}$. By transitivity, $\mathcal{R} \Rightarrow \mathcal{N} \cdot \mathcal{M} \ominus \mathcal{N}$, hence $\mathcal{R} \ominus \mathcal{N}$ exists and $(\mathcal{R} \ominus \mathcal{N}) \Rightarrow (\mathcal{M} \ominus \mathcal{N})$. So, $(\mathcal{R} \ominus \mathcal{N})C' \in ((\Gamma \ominus \mathcal{N}) \ominus (\mathcal{M} \ominus \mathcal{N}))$.
- If $C = (\mathcal{R} \ominus \mathcal{M})C'$ and $\mathcal{R}C' \in \Gamma$. By transitivity, $\mathcal{R} \Rightarrow \mathcal{N} \cdot \mathcal{M} \ominus \mathcal{N} \cdot \mathcal{R} \ominus \mathcal{M}$. Hence $\mathcal{R} \ominus \mathcal{N}$ exists and $(\mathcal{R} \ominus \mathcal{N}) \Rightarrow \mathcal{M} \ominus \mathcal{N} \cdot \mathcal{R} \ominus \mathcal{M}$. Hence $(\mathcal{R} \ominus \mathcal{N}) \ominus (\mathcal{M} \ominus \mathcal{N})$ exists, and $(\mathcal{R} \ominus \mathcal{N}) \ominus (\mathcal{M} \ominus \mathcal{N}) \Rightarrow \mathcal{R} \ominus \mathcal{M}$. So $(\mathcal{R} \ominus \mathcal{N})C' \in ((\Gamma \ominus \mathcal{N}) \ominus (\mathcal{M} \ominus \mathcal{N}))$ and $((\mathcal{R} \ominus \mathcal{N}) \ominus (\mathcal{M} \ominus \mathcal{N}))C' \in ((\Gamma \ominus \mathcal{N}) \ominus (\mathcal{M} \ominus \mathcal{N}))$ which covers for $C = (\mathcal{R} \ominus \mathcal{M})C'$.

□

$$\begin{array}{c}
\frac{}{\Gamma, t, \Delta \vdash t} (\text{Var}) \quad \frac{\Gamma, A \vdash B}{\Gamma \vdash A \Rightarrow B} (\text{ImpR}) \quad \frac{\Gamma, A \Rightarrow B \vdash A \quad \Gamma, A \Rightarrow B, B \vdash D}{\Gamma, A \Rightarrow B \vdash D} (\text{ImpL}) \\
\frac{\Gamma \ominus \mathcal{M} \vdash A}{\Gamma \vdash \mathcal{M}A} (\text{ModR}) \quad \frac{}{\Gamma \vdash \top} (\text{TopR}) \quad \frac{}{\Gamma, \perp, \Delta \vdash A} (\text{BotL}) \\
\frac{\Gamma, A_1 \wedge A_2, A_1 \vdash B}{\Gamma, A_1 \wedge A_2 \vdash B} (\text{AndL1}) \quad \frac{\Gamma, A_1 \wedge A_2, A_2 \vdash B}{\Gamma, A_1 \wedge A_2 \vdash B} (\text{AndL2}) \quad \frac{\Gamma \vdash B_1 \quad \Gamma \vdash B_2}{\Gamma \vdash B_1 \wedge B_2} (\text{AndR}) \\
\frac{\Gamma, A_1 \vee A_2, A_1 \vdash B \quad \Gamma, A_1 \vee A_2, A_2 \vdash B}{\Gamma, A_1 \vee A_2 \vdash B} (\text{OrL}) \quad \frac{\Gamma \vdash B_1}{\Gamma \vdash B_1 \vee B_2} (\text{OrR1}) \quad \frac{\Gamma \vdash B_2}{\Gamma \vdash B_1 \vee B_2} (\text{OrR2})
\end{array}$$

Figure 5: Decidable Sequent Calculus

B.1 The sequent calculus

Figure 5 gives the sequent calculus, showing proof rules to determine which sequent of the form $\Gamma \vdash A$, with Γ a flat context, are provable.

Proposition 1 (Structural weakening). *Suppose \mathcal{D} gives a proof of $\Gamma \vdash C$, and $\Gamma \sqsubseteq \Delta$, then there is a proof \mathcal{D}' of $\Delta \vdash C$, where \mathcal{D}' has the same shape of \mathcal{D} .*

Proof. Can be done by induction on \mathcal{D} . In the ModR rule, we use Lemma 13 in order to make the inductive call. Note in particular that the (Var) rule only targets tokens, and if Γ has token t , then Δ has token t as well. \square

We shall freely apply the above lemma to modify contexts accordingly, including swapping and copying formulas.

Proposition 2 (Identity theorem). *For any formula A , the sequent $\Gamma, A \vdash A$ is provable.*

Proof. Proven by induction on A :

- If $A = t$, $A = \perp$ or $A = \top$, it is directly provable by (Var), (BotL) and (TopR) respectively.

- If $A = A_1 \Rightarrow A_2$:

$$\frac{\frac{\Gamma, A_1 \Rightarrow A_2, A_1 \vdash A_1}{\Gamma, A_1 \Rightarrow A_2, A_1 \vdash A_2} (\text{IH}) \quad \frac{\Gamma, A_1 \Rightarrow A_2, A_1, A_2 \vdash A_2}{\Gamma, A_1 \Rightarrow A_2, A_1 \vdash A_2} (\text{ImpL})}{\Gamma, A_1 \Rightarrow A_2 \vdash A_1 \Rightarrow A_2} (\text{ImpR})$$

- If $A = \mathcal{M}B$,

$$\frac{\frac{\Gamma \ominus \mathcal{M}, B, (\mathcal{M} \ominus \mathcal{M})B \vdash B}{\Gamma, \mathcal{M}B \vdash \mathcal{M}B} (\text{IH})}{\Gamma, \mathcal{M}B \vdash \mathcal{M}B} (\text{Mod})$$

Leave out $(\mathcal{M} \ominus \mathcal{M})B$ if $(\mathcal{M} \ominus \mathcal{M})$ is undefined. $\mathcal{M} \Rightarrow \mathcal{M}$ holds by reflexivity, hence $(\Gamma, \mathcal{M}B) \ominus \mathcal{M}$ contains B .

- The \wedge and \vee cases are standard. \square

Proposition 3 (Cut elimination). *Given a proof \mathcal{D} of $\Gamma \vdash A$, and a proof \mathcal{E} of $\Gamma, A \vdash B$, then we can construct a proof \mathcal{F} of $\Gamma \vdash B$.*

Proof. We use Pfenning's structural cut elimination proof [43] as a basis. We do induction on: Size of A , size of \mathcal{E} , size of \mathcal{D} , in that order. We consider the size of $\mathcal{M}A$ and $\mathcal{N}A$ for any two modalities \mathcal{M} and \mathcal{N} to be the same.

We focus on the non-standard new case our calculus includes: Both \mathcal{D} and \mathcal{E} end with ModR. Case.

$$\mathcal{D} = \frac{\frac{\mathcal{D}_1}{\Gamma \ominus \mathcal{M} \vdash A} (\text{d})}{\Gamma \vdash \mathcal{M}A} (\text{ModR})$$

$$\mathcal{E} = \frac{\frac{\mathcal{E}_1}{\Gamma \ominus \mathcal{N}, A, (\mathcal{M} \ominus \mathcal{N})A \vdash B} \text{ (e1)}}{\Gamma, \mathcal{M}A \vdash \mathcal{N}B} \text{ (ModR)}$$

Note that the $(\mathcal{M} \ominus \mathcal{N})A$ and A under \mathcal{E}_1 exist depending on whether $\mathcal{M} \ominus \mathcal{N}$ exists and $\mathcal{M} \Rightarrow \mathcal{N}$ holds. We shall inductively cut these two formulas. If the formula to be cut does not exist, the respective cut can be left out.

To cut $(\mathcal{M} \ominus \mathcal{N})A$, we use Lemma 15 to note that $(\Gamma \ominus \mathcal{M}) \sqsubseteq ((\Gamma \ominus \mathcal{N}) \ominus (\mathcal{M} \ominus \mathcal{N}))$, and by Lemma 13, $((\Gamma \ominus \mathcal{N}) \ominus (\mathcal{M} \ominus \mathcal{N})) \sqsubseteq ((\Gamma \ominus \mathcal{N}, A) \ominus (\mathcal{M} \ominus \mathcal{N}))$ (regardless of whether this A is there), hence by structural weakening \mathcal{D}_1 gives a proof of $(\Gamma \ominus \mathcal{N}, A) \ominus (\mathcal{M} \ominus \mathcal{N}) \vdash A$.

To cut A , then since A is there by the prerequisite that $\mathcal{M} \Rightarrow \mathcal{N}$, we can use Lemma 14 to see that $(\Gamma \ominus \mathcal{M}) \sqsubseteq (\Gamma \ominus \mathcal{N})$. Hence by structural theorem, \mathcal{D}_1 gives a proof of $\Gamma \ominus \mathcal{N} \vdash A$.

We use these versions of \mathcal{D}_1 to perform the cuts in the following way, $\mathcal{F} =$

$$\frac{\frac{\mathcal{D}'_1}{\Gamma \ominus \mathcal{N} \vdash A} \quad \frac{\frac{\frac{\mathcal{D}'_1}{\Gamma \ominus \mathcal{N}, A \vdash (\mathcal{M} \ominus \mathcal{N})A} \quad \frac{\frac{\mathcal{E}_1}{\Gamma \ominus \mathcal{N}, A, (\mathcal{M} \ominus \mathcal{N})A \vdash B}}{\Gamma \ominus \mathcal{N}, A \vdash B} \text{ (IH-}\mathcal{E}\text{)}}{\Gamma \ominus \mathcal{N} \vdash A} \text{ (IH-A)}}{\Gamma \ominus \mathcal{N} \vdash B} \quad \frac{\Gamma \ominus \mathcal{N} \vdash B}{\Gamma \vdash \mathcal{N}B} \text{ (IH-A)}}{\Gamma \vdash \mathcal{N}B} \text{ (IH-}\mathcal{E}\text{)}$$

□

It should be noted that in the above proof, the fact that Ax is decomposable is used in a fundamental way. $\Gamma \ominus \mathcal{N}$ can for each $\mathcal{M}A$ only spin up A and $(\mathcal{M} \ominus \mathcal{N})A$ at most. We can freely cut the A , as a reduced formula (A compared to $\mathcal{M}A$) can be used as a basis for the induction hypothesis, allowing us to liberally change the needed proofs as the size of formula is the dominant (first) inductive argument in the cut elimination proof. Cutting $(\mathcal{M} \ominus \mathcal{N})A$ is more difficult: we ensure that for determining termination of the cut elimination proof, we consider the size of $(\mathcal{M} \ominus \mathcal{N})A$ to be the same as the size of $\mathcal{M}A$. This is ok, as they both have the same number of *proper* subformulas, and hence the measure for checking termination on the formula to be cut does not increase. Cutting $\mathcal{M} \ominus \mathcal{N})A$ can be done using a reduced proof \mathcal{E} and same size proof \mathcal{D} .

Decomposability also ensures that we need only consider one such modal formula of the same size as $\mathcal{M}A$. Cutting multiple is problematic, as it makes the termination argument more difficult (maybe impossible), so we let $(\mathcal{M} \ominus \mathcal{N})A$ be sufficient for implying all other modal formulas $\mathcal{R}A$ which may be needed. So if we have multiple \mathcal{R} for which $\mathcal{M} \Rightarrow \mathcal{N} \cdot \mathcal{R}$, the single formula $(\mathcal{M} \ominus \mathcal{N})A$ covers all of them.

B.2 Proving Decidability

The decidability proof follows the usual recipe. Consider again the subformula property \leq of formulas, which is extended to include $A \leq \mathcal{M}A$ and $\mathcal{M}A \leq \mathcal{N}B$ if $A \leq B$. Since we only have a finite number of modalities \mathbb{M} , we get that for any formula A there are only finitely many B such that $B \leq A$. We observe that for any $A \in (\Gamma \ominus \mathcal{M})$ there is a $B \in \Gamma$ such that $A \leq B$. As a consequence, we can show by induction on proofs that suppose \mathcal{D} proves $\Gamma \vdash A$, and $\Delta \vdash B$ appears somewhere in \mathcal{D} , then for any $C \in \Delta, B$ there is a $D \in \Gamma, A$ such that $C \leq D$.

Given structural weakening shown in Proposition 1, we see that whenever a $\Delta \vdash A$ appears above a $\Gamma \vdash A$ in a proof, where $\Delta \subseteq \Gamma$, then we have a redundancy: the proof of $\Delta \vdash A$ could replace the proof of $\Gamma \vdash A$, reducing the proof tree. We say that \mathcal{D} is reduced if there are no $\Delta \vdash A$ appearing above $\Gamma \vdash A$ in \mathcal{D} such that $\Delta \subseteq \Gamma$. By the above remark, if there is a proof of a sequent, there is a reduced proof of that sequent.

For each sequent $\Gamma \vdash A$, there is a bound n such that any reduced proof of $\Gamma \vdash A$ has at most height n . This is because by the subformula property, any $\Delta \vdash B$ in a proof of $\Gamma \vdash A$ can only contain formulas drawn from the set of subformulas of Γ, A , of which there are finitely many (note there are only finitely many modalities): we say there are m such formulas. So suppose S is a set of sequents $\Delta \vdash B$ possibly appearing in proofs of $\Gamma \vdash A$, such that no two sequents are eachother weakenings, then S has at most $m \cdot 2^m$ elements (m possible consequents, and 2^m possible ordered lists of non-repeating assumptions). So any branch in a reduced proof has at most $m \cdot 2^m$ giving a bound on the height of reduced trees.

Since each proof steps makes at most two branches, the bound on the height gives a bound on the size of reduced proofs. Hence, up to weakening, there are only a finite number of reduced proof trees we need to check to see if $\Gamma \vdash A$ has a proof. For a theoretical argument of decidability, we simply check

all possible reduced proofs to find if one works. More practical algorithms involve more targeted proof searches.

B.3 Extra: Proving axioms

We can show that the sequent calculus correctly adapts \mathcal{L}_{Ax} , as it has the same rules as the standard sequent calculus for intuitionistic logics on the non-modal side. Let us prove the axioms:

- For axiom K, note that $(\mathcal{M}A, \mathcal{M}(A \Rightarrow B)) \ominus \mathcal{M}$ by reflexivity of Ax includes A and $A \Rightarrow B$, which proves B . So $(\mathcal{M}A, \mathcal{M}(A \Rightarrow B)) \ominus \mathcal{M} \vdash B$ is provable, and hence by the ModR rule, $\mathcal{M}A, \mathcal{M}(A \Rightarrow B) \vdash \mathcal{M}B$. Axiom K is then constructed by using the ImpR rule twice.
- For necessity, note that with have a proof of $\vdash A$, then since $(\cdot \ominus \mathcal{M}) = \cdot$, this is also a proof of $(\cdot \ominus \mathcal{M}) \vdash A$. Hence by the ModR rule, $\vdash \mathcal{M}A$ is provable.
- Suppose $\mathcal{M} \Rightarrow \mathcal{N} \in \text{Ax}$, then $(\mathcal{M}A \ominus \mathcal{N})$ has A . So $(\mathcal{M}A \ominus \mathcal{N}) \vdash A$ is provable by the identity theorem, and hence $\mathcal{M}A \vdash \mathcal{N}A$ is provable by ModR, and $\vdash \mathcal{M}A \Rightarrow \mathcal{N}A$ by ImpR.
- Suppose $\mathcal{M} \Rightarrow \mathcal{N} \cdot \mathcal{R} \in \text{Ax}$, then $(\mathcal{M}A \ominus \mathcal{N})$ has $(\mathcal{M} \ominus \mathcal{N})A$ and $\mathcal{M} \ominus \mathcal{N} \Rightarrow \mathcal{R} \in \text{Ax}$ by decomposability. So $(\mathcal{M}A \ominus \mathcal{N}) \ominus \mathcal{R}$ has A , and we can prove $(\mathcal{M}A \ominus \mathcal{N}) \ominus \mathcal{R} \vdash A$, and apply ModR twice to get $\mathcal{M}A \vdash \mathcal{N}\mathcal{R}A$, proving $\vdash \mathcal{M}A \Rightarrow \mathcal{N}\mathcal{R}A$ with ImpR.
- We generalize the previous case, showing that by induction on the length of l , if $\mathcal{M} \Rightarrow l \in \text{Ax}$ and $\mathcal{M}A \in \Gamma$, then $\Gamma \vdash l(A)$. Base case for length of l being 1 (or even 2) has been covered above. Suppose $l = \mathcal{N}, l'$, then $\mathcal{M} \ominus \mathcal{N}$ exists and $\mathcal{M} \ominus \mathcal{N} \Rightarrow l' \in \text{Ax}$. So supposing $\mathcal{M}A \in \Gamma$, then $(\mathcal{M} \ominus \mathcal{N})A \in (\Gamma \ominus \mathcal{N})$, and by induction hypothesis, $(\Gamma \ominus \mathcal{N}) \vdash l'(A)$. Using ModR, we can conclude that $\Gamma \vdash \mathcal{N}l'(A)$ where $\mathcal{N}l' = l$. This finishes the induction.

We conclude that if $\mathcal{M} \Rightarrow l \in \text{Ax}$, then $\mathcal{M}A \vdash l(A)$ and hence $\vdash \mathcal{M}A \Rightarrow l(A)$ by ImpR.

C Equivalence between Calculi

Given two flat contexts Γ and Δ , we say that $\Gamma \vdash \Delta$ is provable if $\Gamma \vdash B$ is provable for any $B \in \Delta$.

Lemma 16. *If $\Gamma \vdash \Delta$ and $\Delta \vdash \Phi$ are provable, then $\Gamma \vdash \Phi$ is provable.*

Proof. Starting with $\Gamma, \Delta \vdash C$ for some $C \in \Phi$, use a sequence of sequents $\Gamma, B_1, \dots, B_{i-1} \vdash B_i$ to cut away Δ , until getting to $\Gamma \vdash C$. \square

Consider a modal context Γ , which may include modalities $\{\mathcal{M}\}$. Hence Γ is either flat (has no modalities), or of the form $\Gamma', \{\mathcal{M}\}, \Gamma''$ with Γ' flat and Γ'' a modal context.

Definition 7. Given context Γ and flat context Δ , then $\Gamma \vdash \Delta$ is *constructable* if, by induction on Γ :

- If Γ is flat, then $\Gamma \vdash \Delta$ is constructable if it is provable.
- If $\Gamma = \Gamma_0, \{\mathcal{M}\}, \Gamma_1$ with Γ_0 flat, then $\Gamma \vdash \Delta$ is constructable if there is a flat context Φ such that:
 - $\Gamma_0 \vdash \Phi$ is provable.
 - $\Phi \ominus \mathcal{M}, \Gamma_1 \vdash \Delta$ is constructable.

If $\Gamma \vdash \Phi$ is constructable, and $\Phi, \Psi \vdash \Delta$ is constructable, then $\Gamma, \Psi \vdash \Delta$ is constructible.

Lemma 17 (Weakening of Constructability). *If $\Phi \subseteq \Psi$ and $\Gamma, \Phi, \Delta \vdash \Omega$ is constructable, then $\Gamma, \Psi, \Delta \vdash \Omega$ is constructable.*

Lemma 18. *If $\Gamma \vdash \Delta_0$ and $\Gamma \vdash \Delta_1$ are constructable, then $\Gamma \vdash \Delta_0, \Delta_1$ is constructable.*

Proof. Prove is done by induction on Γ .

If Γ is flat, then $\Gamma \vdash \Delta_0$ and $\Gamma \vdash \Delta_1$ are provable, hence $\Gamma \vdash \Delta_0, \Delta_1$ is provable.

If $\Gamma = \Gamma_0, \{\mathcal{M}\}, \Gamma_1$ with Γ_0 flat, then there are Φ_0 and Φ_1 such that $\Gamma_0 \vdash \Phi_0$ and $\Gamma_0 \vdash \Phi_1$ are provable and hence $\Gamma_0 \vdash \Phi_0, \Phi_1$ is provable, and both $\Phi_0 \ominus \mathcal{M}, \Gamma_1 \vdash \Delta_0$ and $\Phi_1 \ominus \mathcal{M}, \Gamma_1 \vdash \Delta_1$ are constructive. Note that $(\Phi_0, \Phi_1) \ominus \mathcal{M} = (\Phi_0 \ominus \mathcal{M}), (\Phi_1 \ominus \mathcal{M})$, hence by weakening $(\Phi_0, \Phi_1) \ominus \mathcal{M}, \Gamma_1 \vdash \Delta_0$ and $(\Phi_0, \Phi_1) \ominus \mathcal{M}, \Gamma_1 \vdash \Delta_1$ are constructive. By induction hypothesis, $(\Phi_0, \Phi_1) \ominus \mathcal{M}, \Gamma_1 \vdash \Delta_0, \Delta_1$. \square

Theorem 3. *If there is a term $\Gamma \vdash P : A$, then $\Gamma \vdash A$ is constructable.*

Proof. Induction on t :

$t \cdot r$ with $t : A \Rightarrow B$ and $r : A$, then by induction hypothesis, $\Gamma \vdash A \Rightarrow B$ and $\Gamma \vdash A$ are constructable, hence by Lemma 18, $\Gamma \vdash A \Rightarrow B, A$ is constructable. We show that $A \Rightarrow B, A \vdash B$ is provable, to construct $\Gamma \vdash B$.

$$\frac{\frac{}{A \Rightarrow B, A \vdash A} \text{(Var)}}{A \Rightarrow B, A \vdash B} \text{(ImpL)} \quad \frac{\frac{}{A \Rightarrow B, A, B \vdash B} \text{(Var)}}{A \Rightarrow B, A \vdash B} \text{(ImpR)}$$

$\lambda A.(t)$ with $t : B$, then by induction hypothesis $\Gamma, A \vdash B$ is constructable. Take the tail-end provable $\Gamma, A \vdash B$ and apply (ImpR) to construct $\Gamma \vdash A \Rightarrow B$.

$\text{key}_{\mathcal{M} \Rightarrow \alpha}(t)$ with $\Gamma \vdash t : \mathcal{M}A$. By induction hypothesis, $\Gamma \vdash \mathcal{M}A$ is constructable. We do case distinction on α :

- If $\alpha = \mathcal{N}$ then $\mathcal{M} \Rightarrow \mathcal{N}$. $\mathcal{M}A, \{\mathcal{M}\} \vdash A$ is constructible since $\mathcal{M}A \vdash \mathcal{M}A$ and $\mathcal{M}A \ominus \mathcal{N} \vdash A$ is provable by (Var). Composing $\mathcal{M}A, \{\mathcal{N}\} \vdash A$ with $\Gamma, \{\mathcal{N}\} \vdash \mathcal{M}A$, we construct $\Gamma, \{\mathcal{N}\} \vdash A$.
- If $\alpha = \mathcal{N}_1, \dots, \mathcal{N}_n$ with $n > 1$, we know $\mathcal{M} \Rightarrow \mathcal{N}_1 \dots \mathcal{N}_n$, and hence can define $\mathcal{R}_1 = \mathcal{M}$, $\mathcal{R}_{i+1} = \mathcal{R}_i \ominus \mathcal{N}_i$ up to \mathcal{R}_n for which we have $\mathcal{R}_n \Rightarrow \mathcal{N}_n$. Then $\mathcal{M}A, \{\mathcal{N}_1\} \vdash \mathcal{R}_1A$, and $\mathcal{R}_iA, \{\mathcal{N}_i\} \vdash \mathcal{R}_{i+1}A$ up to $\mathcal{R}_n, \{\mathcal{N}_n\} \vdash A$ are all constructable, composing into $\Gamma, \{\mathcal{N}_1\}, \dots, \{\mathcal{N}_n\} \vdash A$.

$\text{lock}_{\mathcal{M}}(t)$ with $\Gamma, \{\mathcal{M}\} \vdash t : A$. By induction hypothesis, we end with a proof of $\Gamma \ominus \mathcal{M} \vdash A$ with constructable $\Gamma \vdash \Gamma$. This gives us a proof of $\Gamma \vdash \mathcal{M}A$ using (ModR). Composing $\Gamma \vdash \Gamma$ and $\Gamma \vdash \mathcal{M}A$ we construct $\Gamma \vdash \mathcal{M}A$.

Other cases are simpler. □

Lemma 19. *For any term $(\Gamma \ominus \mathcal{M}), \Delta \vdash t : A$, there is a term $\Gamma, \{\mathcal{M}\}, \Delta \vdash r : A$.*

Proof. $(\Gamma \ominus \mathcal{M}), \Delta \vdash t : A$ can be weakened to $\Gamma, \{\mathcal{M}\}, (\Gamma \ominus \mathcal{M}), \Delta \vdash t : A$. We reason that for any formula B from $(\Gamma \ominus \mathcal{M})$ there is a lambda term $\Gamma, \{\mathcal{M}\} \vdash r_B : B$, and hence we can use substitutions to get a term $\Gamma, \{\mathcal{M}\}, \Gamma \vdash r : A$. Suppose $B \in (\Gamma \ominus \mathcal{M})$, then either:

- $x : \mathcal{N}B \in \Gamma$ for some x , and $\mathcal{N} \Rightarrow \mathcal{M}$. Then $\Gamma, \{\mathcal{M}\} \vdash \text{key}_{\mathcal{N} \Rightarrow \mathcal{M}}(\text{var}(x)) : B$.
- $B = (\mathcal{N} \ominus \mathcal{M})B'$ and $x : \mathcal{N}B' \in \Gamma$. Then $\Gamma, \{\mathcal{M}\} \vdash \text{lock}_{\mathcal{N} \ominus \mathcal{M}}(\text{key}_{\mathcal{N} \Rightarrow \mathcal{M}, \mathcal{N} \ominus \mathcal{M}}(\text{var}(x))) : (\mathcal{N} \ominus \mathcal{M})B'$. □

Theorem 4. *If $\Gamma \vdash A$ is provable, there is a lambda term t such that $\Gamma \vdash t : A$*

Proof. By induction on proofs. Most cases are standard, we shall just consider the modal case.

(ModR) Proving $\Gamma \vdash \mathcal{M}A$, using a proof of $\Gamma \ominus \mathcal{M} \vdash A$. By induction hypothesis, we have a term $\Gamma \ominus \mathcal{M} \vdash t : A$, and we can apply Lemma 19 to get $\Gamma, \{\mathcal{M}\} \vdash r : A$ and hence $\Gamma \vdash \text{lock}_{\mathcal{M}}(s) : A$. □

Theorem 5. *If $\Gamma \vdash A$ is constructable, there is a lambda term t such that $\Gamma \vdash t : A$*

Proof. By induction on Γ .

If $\Gamma = \Gamma$, then $\Gamma \vdash A$ is provable, and we use the previous theorem.

If $\Gamma = \Gamma, \{\mathcal{M}\}, \Gamma'$, with provable $\Gamma \vdash \Delta$, and constructable $\Delta \ominus \mathcal{M}, \Gamma' \vdash A$. We use the previous theorem and the induction hypothesis to find terms $\Gamma \vdash t_D : D$ for each $D \in \Delta$, and $\Delta \ominus \mathcal{M}, \Gamma' \vdash r : A$. Applying Lemma 19 to the latter, we get $\Delta, \{\mathcal{M}\}, \Gamma' \vdash s : A$. Substituting each t_D into s we get the desired term $\Gamma, \{\mathcal{M}\}, \Gamma' \vdash t' : A$. □

We can conclude that $\Gamma \vdash A$ is constructible if and only if there is a lambda term t such that $\Gamma \vdash t : A$. If in particular Γ is flat, then $\Gamma \vdash A$ is provable in the sequent calculus if and only if there is a lambda term t such that $\Gamma \vdash t : A$. Hence by decidability, there is an algorithm which determines for each sequent $\Gamma \vdash A$ with Γ a flat context, whether a lambda term exists. This can be extended to any modal context as follows.

We define the map sending pairs $(\Gamma \vdash A)$ of modal context and formulas to a formula $\langle \Gamma, A \rangle$ as follows:

- $\langle \vdash A \rangle := A$,
- $\langle (\Gamma, x : B) \vdash A \rangle := \langle \Gamma \vdash B \Rightarrow A \rangle$,

- $\langle\langle\Gamma, \{\mathcal{M}\}\rangle \vdash A\rangle := \langle\Gamma \vdash \mathcal{M}A\rangle$.

By induction, $\langle(x : B, \Gamma) \vdash A\rangle = B \Rightarrow \langle\Gamma \vdash A\rangle$ and $\langle(\{\mathcal{M}\}, \Gamma) \vdash A\rangle = \mathcal{M}\langle\Gamma \vdash A\rangle$.

Lemma 20. *For each Γ , Δ and A , the sequence $\Gamma, \Delta \vdash A$ has a lambda term if and only if $\Gamma \vdash \langle\Delta \vdash A\rangle$ has a lambda term.*

Proof. We prove both directions of the implication by induction on the length of Δ . The base case is simple, since $\langle\vdash A\rangle = A$. First from left to right:

- If $\Delta = \Delta', x : B$ and $\Gamma, \Delta', x : B \vdash t : A$. Then $\Gamma, \Delta' \vdash \lambda x : B. t : B \Rightarrow A$ which by induction hypothesis gives us a term $\Gamma \vdash t' : \langle\Delta' \vdash B \Rightarrow A\rangle$, where $\langle\Delta' \vdash B \Rightarrow A\rangle = \langle\Delta', x : B \vdash A\rangle$.
- If $\Delta = \Delta', \{\mathcal{M}\}$ and $\Gamma, \Delta, \{\mathcal{M}\} \vdash t : A$. Then $\Gamma, \Delta' \vdash \text{lock}_{\mathcal{M}}(t) : \mathcal{M}A$ which by induction hypothesis gives us a term $\Gamma \vdash t' : \langle\Delta' \vdash \mathcal{M}A\rangle$, where $\langle\Delta' \vdash \mathcal{M}A\rangle = \langle\Delta', \{\mathcal{M}\} \vdash A\rangle$.

The converse:

- If $\Delta = x : B, \Delta'$ and $\Gamma \vdash t : \langle x : B, \Delta' \vdash A \rangle$, then since $\langle x : B, \Delta' \vdash A \rangle = B \Rightarrow \langle \Delta' \vdash A \rangle$ we have $\Gamma, x : B \vdash t \cdot \text{var}(x) : \langle \Delta' \vdash A \rangle$. This gives us by induction hypothesis a term $\Gamma, x : B, \Delta' \vdash t' : A$ as desired.
- If $\Delta = \{\mathcal{M}\}, \Delta'$ and $\Gamma \vdash t : \langle \{\mathcal{M}\}, \Delta' \vdash A \rangle$, then since $\langle \{\mathcal{M}\}, \Delta' \vdash A \rangle = \mathcal{M}\langle \Delta' \vdash A \rangle$ we have $\Gamma, \{\mathcal{M}\} \vdash \text{key}_{\mathcal{M} \Rightarrow \mathcal{M}}(t) : \langle \Delta' \vdash A \rangle$. This gives us by induction hypothesis a term $\Gamma, \{\mathcal{M}\}, \Delta' \vdash t' : A$ as desired.

□

We conclude that $\Gamma \vdash A$ has a lambda term if and only if $\vdash \langle \Gamma \vdash A \rangle$ has a lambda term if and only if it has a proof in the decidable sequent calculus. Hence it is decidable whether $\Gamma \vdash A$ has a lambda term.

D Additional Proofs

Proof of Lemma 5. If $\mathcal{N}_S(a, b, d)$ then $(a, \alpha, b, \beta, d) \in S$ for some possibly empty $\alpha, \beta \in \mathbb{A}^*$. If $\mathcal{N}_S(b, c, d)$ then $(b, \gamma, c, \delta, d) \in S$ for some $\gamma, \delta \in \mathbb{A}^*$. By property 2 of a forwarding network, we can replace β in (a, α, b, β, d) with (γ, c, δ) , creating $(a, \alpha, b, \gamma, c, \delta, d) \in S$. This shows that $\mathcal{N}_S(a, c, d)$ as desired. Moreover, using property 1 of forwarding networks, we can show that $(b, \gamma, c, \delta, d) \in S$ as well, and hence $\mathcal{N}_S(b, c, d)$. The second property has a similar proof. □

Proof of Lemma 6. First note that if \mathcal{M} and \mathcal{N} satisfy axiom K and necessity, then $\mathcal{M}\mathcal{N}$ satisfies axiom K and necessity, since firstly $\vdash A$ implies $\vdash \mathcal{N}A$ implies $\vdash \mathcal{M}\mathcal{N}A$. Secondly, $\vdash \mathcal{N}A \wedge \mathcal{N}(A \Rightarrow B) \Rightarrow \mathcal{N}B$ so $\mathcal{M}\mathcal{N}A \wedge \mathcal{M}\mathcal{N}(A \Rightarrow B) \Rightarrow \mathcal{M}\mathcal{N}B$ by necessity and axiom K for \mathcal{M} . We conclude that \mathcal{I}_α satisfies axiom K and necessity for any appropriate α .

Proving Axiom K, suppose $\mathcal{J}_{a \vdash b}A$ and $\mathcal{J}_{a \vdash b}(A \Rightarrow B)$, we want to show that $\mathcal{J}_{a \vdash b}B$. Let $\gamma \in \mathbb{A}^*$ such that $(a, \gamma, b) \in S$, then $\mathcal{I}_{a \vdash \gamma \vdash b}A$ and $\mathcal{I}_{a \vdash \gamma \vdash b}(A \Rightarrow B)$, hence by axiom K on $\mathcal{I}_{a \vdash \gamma \vdash b}$ we get $\mathcal{I}_{a \vdash \gamma \vdash b}B$. This is for all relevant γ , hence $\mathcal{J}_{a \vdash b}B$. Proving necessity, if $\vdash A$, then $\vdash \mathcal{I}_{a \vdash \gamma \vdash b}A$ for all γ such that $(a, \gamma, b) \in S$, hence $\vdash \mathcal{J}_{a \vdash b}A$.

For $\mathcal{J}_{a \vdash b}A \Rightarrow \mathcal{B}_a \mathcal{J}_{a \vdash b}A$ and $\mathcal{J}_{a \vdash b}A \Rightarrow \mathcal{J}_{a \vdash b} \mathcal{B}_b A$, simply note: 1) that $\mathcal{I}_{a \vdash \gamma \vdash b}$ starts with some $\mathcal{I}_{a \vdash a'}$ for which $\mathcal{I}_{a \vdash a'}B \Rightarrow \mathcal{B}_a \mathcal{I}_{a \vdash a'}B$ for any B , hence $\mathcal{I}_{a \vdash \gamma \vdash b}A \Rightarrow \mathcal{B}_a \mathcal{I}_{a \vdash \gamma \vdash b}A$. 2) $\mathcal{I}_{a \vdash \gamma \vdash b}$ ends with some $\mathcal{I}_{b' \vdash b}$ for which $\mathcal{I}_{b' \vdash b}A \Rightarrow \mathcal{I}_{b' \vdash b} \mathcal{B}_b A$. Using axiom K on the rest of $\mathcal{I}_{a \vdash \gamma \vdash b}$ (which may be empty), we get $\mathcal{I}_{a \vdash \gamma \vdash b}A \Rightarrow \mathcal{I}_{a \vdash \gamma \vdash b} \mathcal{B}_b A$. □

Lemma 21. S_G , the set of shortest paths on a finite graph is a forwarding network.

Proof. Note first that a shortest path between two vertices is given by a non-empty list of distinct vertices. The shortest path from a to a is simply given by (a) . We prove the three additional properties:

1. Suppose given a shortest path $(a, b_1 \dots, b_n, c)$ from c to a , then $(a, b_1 \dots, b_n)$ has to be a shortest path from b_n to a ; if not, there would be a shorter path from c to a . Hence $(a, b_1 \dots, b_n) \in S_G$, and by a similar argument, $(b_1 \dots, b_n, c) \in S$.

2. Suppose $(a_1, \dots, a_n, b, c_1, \dots, c_m, d, e_1, \dots, e_k)$ is a shortest path, we leave ambiguous which elements are the source and target, since n and/or k could be zero. This is a path of length $n+m+k+1$. We know that (b, c_1, \dots, c_m, d) is the shortest path from d to b . Taking some other shortest path $(b, c'_1, \dots, c'_p, d)$ from d to b , we note that it has to have the same length, hence $p = m$. Hence $(a_1, \dots, a_n, b, c'_1, \dots, c'_p, d, e_1, \dots, e_k)$ is also a path of length $n + m + k + 1$. Note in particular that $(a_1, \dots, a_n, b, c'_1, \dots, c'_p, d, e_1, \dots, e_k)$ cannot have repeats, since otherwise we could cut out the cycle and create a shorter path, contradicting that that the shortest path in of length $n + m + k + 1$. We conclude that $(a_1, \dots, a_n, b, c'_1, \dots, c'_p, d, e_1, \dots, e_k)$ is also a shortest path, and hence in S_G . □

Lemma 22. *If \mathcal{G} is acyclic, S'_G is a forwarding network.*

Proof. This works since paths are closed under subpaths (property 1), and the result of replacing subpaths by alternative subpaths (property 2) cannot create a path without repeats. □