Principal-Centric Reasoning in Constructive Authorization Logic

(Extended Abstract)

Deepak Garg
dg@cs.cmu.edu
Carnegie Mellon University

Abstract

We present an authorization logic that is quite similar to constructive modal S4. The logic assumes that principals are conceited in their beliefs. We describe the sequent calculus, Hilbert-style axiomatization, and Kripke semantics of the logic. A distinguishing characteristic of the sequent calculus is that hypothetical reasoning is relativized to beliefs of principals. We prove several meta-theoretic results including cut-elimination, and soundness and completeness for the Kripke semantics.

1 Introduction

Authorization refers to the act of deciding whether or not an agent making a request to perform an operation on a resource should be allowed to do so. For example, the agent may be a browser trying to read pages from a website. In that case, the site's web server may consult the browser's credentials and a `.htaccess` file to determine whether to send the pages or not. Such access control is pervasive in computer systems. As systems and their user environments evolve, policies used for access control may become complex and error prone. This suggests the need for formal mechanisms to represent, enforce, and analyze policies. Logic appears to be a useful mechanism for these purposes. Policies may be expressed as formulas in a suitably chosen logic. This has several merits. First, the logic's rigorous inference eliminates any ambiguity in meaning that may be inherent in a textual description of policies. Second, policies may be enforced end-to-end using generic logic-based mechanisms like proof-carrying authorization [8–10, 29]. Third, by writing policies in a logic, there is hope that the policies themselves can be checked for correctness against some given criteria.

Whereas first-order logic and sometimes propositional logic suffice to express many authorization policies, distributed systems pose a peculiar challenge: how do we express and combine policies of *different* agents and systems? This is often necessary since policies and the authorizations derived from them may vary from system to system. Policies on different systems may also interact to allow or deny access. To model such distributed policies, Abadi and others proposed logics with formulas of the form $K \mathsf{says} A$, where K is an agent or a system (abstractly called a principal) and A is a formula representing a policy [6, 28]. The intended meaning of the formula is that principal K states, or believes that policy A holds. From a logical perspective $K \mathsf{says} \cdot$ is a modality and the logic is an indexed modal logic with one modality for each principal. We call such a modal logic an *authorization logic*. In the past fifteen years there have been numerous proposals describing authorization logics that differ widely in the specific axioms (or inference rules) used for $K \mathsf{says} \cdot$ [2, 3, 8–10, 15, 18, 20, 24–26, 29, 30]. One emerging trend is the increased use of intuitionistic logics for authorization (e.g., [3, 16, 20, 24–26, 29, 37]) as opposed to classical logics.

This paper presents a new intuitionistic authorization logic called DTL_0. This logic is peculiar in a certain respect: it abandons the usual objectivity in reasoning from hypothesis, relativizing hypothetical reasoning to principals. The hypothetical judgment of the logic has the form $\Gamma \xrightarrow{K} A$, which means, up to a first approximation, that under the assumption that all beliefs of K are true, the hypotheses Γ imply A. Although this choice of binding hypothetical reasoning to principals may be unintuitive from a philosophical point of view, it seems attractive from the perspective of access control.

Our primary interest in developing DTL_0 is deployment in proof-carrying authorization [8–10, 29, 37]. Hence our main focus is DTL_0's proof-theory, especially the sequent calculus, which we describe in detail (Section 3). We prove several meta-theoretic properties of the sequent calculus, including cut-elimination (Section 3.1). We also present a Hilbert-style system for DTL_0 (Section 2), and sound and complete Kripke semantics (Section 4). The principal-centric reasoning of DTL_0 reflects in the Kripke semantics: worlds are explicitly associated with principals who may view them. This suggests that principals in DTL_0 may be related to nominals from hybrid logic [13, 14, 17]. We also show that DTL_0 is a generalization of constructive modal S4 [7, 33].

DTL_0 is a fragment of a larger authorization logic, DTL, which we are currently developing. The latter is quite broad, incorporating first-order quantifiers, explicit time for modeling time-bounded policies [20], and linearity for modeling consumable credentials [25]. Besides developing the logic's theory, a secondary goal of ongoing work is to understand how DTL_0 relates to existing authorization logics, through translations between them. The eventual objective of this line of work is more ambitious; we want to establish a common framework in which policies written in different logics may be combined. Initial efforts in this direction using (classical) modal S4 as foundation appeared in earlier work [24].

By itself, this paper makes two contributions. First, it presents a new authorization logic that explicitly relativizes hypothetical reasoning to principals, and describes the logic's proof theory. To the best of our understanding, such relativization is unique to our logic, at least in the context of authorization. A second, albeit minor contribution of the paper is sound and complete Kripke semantics, which are relatively rare for authorization logics (as opposed to their prevalence in modal logics). The only other examples we know of are Kripke semantics for authorization logics based on lax-like modalities [24], and those for an earlier authorization logic based on the modal logic K [6].

To save space, proofs of theorems and many other results related to DTL_0 have been omitted from this extended abstract. These may be found in the full version of the paper that is available on the author's web page [23]. In addition to proofs and a description of some of the design choices, the full version contains a natural deduction system, a construction of canonical Kripke models, and sound and complete translations between DTL_0 and other modal logics, including several authorization logics and constructive multi-modal S4.

2 The logic DTL_0

DTL_0 extends propositional intuitionistic logic with a principal-indexed modality, $K \mathsf{\ says\ } A$. Principals, denoted K, are abstractions for users, programs, machines, and systems, that either create policies or request access to resources. We stipulate a fixed set of principals `Prin`, pre-ordered by a relation written $\succeq$. $K_1 \succeq K_2$ is read "principal K_1 is stronger than principal K_2", and entails that $K_1 \mathsf{\ says\ } A$ implies $K_2 \mathsf{\ says\ } A$ for every formula A. We assume that `Prin` has at least one maximum element, called the *local authority* (denoted ℓ).[1] The syntax of formulas in DTL_0 is shown below. P denotes atomic formulas.

$$A, B, C \ ::= \ P \mid A \wedge B \mid A \vee B \mid \top \mid \bot \mid A \supset B \mid K \mathsf{\ says\ } A$$

Axiomatic Proof-System. A Hilbert-style proof-system for DTL_0 consists of any axiomatization of propositional intuitionistic logic (elided here), and the following axioms and rules for $K \mathsf{\ says\ } A$. We write $\vdash A$ to mean that A is valid.

$$\frac{\vdash A}{\vdash K \mathsf{\ says\ } A} \qquad \text{(nec)}$$

$$\vdash (K \mathsf{\ says\ } (A \supset B)) \supset ((K \mathsf{\ says\ } A) \supset (K \mathsf{\ says\ } B)) \qquad \text{(K)}$$

$$\vdash (K \mathsf{\ says\ } A) \supset K \mathsf{\ says\ } K \mathsf{\ says\ } A \qquad \text{(4)}$$

$$\vdash K \mathsf{\ says\ } ((K \mathsf{\ says\ } A) \supset A) \qquad \text{(C)}$$

$$\vdash (K_1 \mathsf{\ says\ } A) \supset (K_2 \mathsf{\ says\ } A) \text{ if } K_1 \succeq K_2. \qquad \text{(S)}$$

(nec) and (K) are the usual necessitation rule and closure under consequence axiom for normal modal logics (see e.g., [12]). (4) is also standard from modal logics such as S4. (C)

[1] To the best of our understanding, the term *local authority* as used here was first introduced in the preview implementation of the language SecPAL [1].

is the characterizing axiom of DTL_0. It is characteristic of the doxastic logic of conceited reasoners (hence the name C) [35]. Intuitively, the axiom means that every principal says that all its statements are true. Although the propriety of this axiom in the context of doxastic reasoning has been questioned, it seems quite useful for authorization. The axiom (S) means that whenever principal K_1 believes a formula A, every weaker principal K_2 believes it as well.

The following properties may be established in DTL_0. $\not\vdash A$ means that A is not valid in the stated generality (although specific instances of A may be valid). $A \equiv B$ denotes $(A \supset B) \wedge (B \supset A)$.

$\vdash (\ell \mathsf{\ says\ } A) \supset (K \mathsf{\ says\ } A)$

$\vdash (K \mathsf{\ says\ } K \mathsf{\ says\ } A) \equiv (K \mathsf{\ says\ } A)$

$\not\vdash A \supset K \mathsf{\ says\ } A$

$\not\vdash (K \mathsf{\ says\ } A) \supset A$

$\vdash (K \mathsf{\ says\ } (A \wedge B)) \equiv ((K \mathsf{\ says\ } A) \wedge (K \mathsf{\ says\ } B))$

$\not\vdash (K \mathsf{\ says\ } (A \vee B)) \supset ((K \mathsf{\ says\ } A) \vee (K \mathsf{\ says\ } B))$

$\not\vdash \bot$

$\not\vdash (K \mathsf{\ says\ } A) \supset (K' \mathsf{\ says\ } (K \mathsf{\ says\ } A))$

The last property means that if a principal K states policy A, not every principal may believe this. In some cases, this may not be desirable, since some policies may be stated and *published* by K. If K publishes policy A, we may expect that $K' \mathsf{\ says\ } K \mathsf{\ says\ } A$. In DTL_0, published policies may be expressed using the defined connective $K \mathsf{\ publ\ } A = \ell \mathsf{\ says\ } K \mathsf{\ says\ } A$ (read "K publishes A"), which satisfies the following properties:

$\vdash (K \mathsf{\ publ\ } A) \supset K \mathsf{\ says\ } A.$

$\not\vdash (K \mathsf{\ says\ } A) \supset K \mathsf{\ publ\ } A.$

$\vdash (K \mathsf{\ publ\ } A) \supset K' \mathsf{\ says\ } (K \mathsf{\ publ\ } A).$

$\vdash (K \mathsf{\ publ\ } A) \supset K' \mathsf{\ says\ } (K \mathsf{\ says\ } A).$

Example 2.1 (Policies in DTL_0)**.** We illustrate the use of DTL_0 for expressing authorization policies through a simple example. Suppose that the principal OAL (Online Academic Library) represents an online repository of scientific articles. Academics institutions (such as CMU) may buy corporate subscriptions that allow all their members to download articles from OAL. It is up to the subscribing institutions to tell OAL who their members are. Alice is an individual who wishes to download an article from OAL. Let the formula $\mathtt{downloadAlice}$ mean that Alice may download articles from OAL, and let $\mathtt{memberAliceCMU}$ mean that Alice is a member of CMU. Further, let us assume that CMU has a subscription at OAL. The following represent possible policies of the principals.

1. $\mathsf{OAL}\ \mathsf{says}\ ((\mathsf{CMU}\ \mathsf{says}\ \mathtt{memberAliceCMU}) \supset \mathtt{memberAliceCMU})$

2. $\mathsf{OAL}\ \mathsf{says}\ (\mathtt{memberAliceCMU} \supset \mathtt{downloadAlice})$

3. $\mathsf{CMU}\ \mathsf{publ}\ \mathtt{memberAliceCMU}$

The first policy, stated by OAL, means that if CMU says that Alice is its member, then this is the case. The second policy, also stated by OAL, means that if Alice is a member of CMU, then she may download articles. The third policy, stated and published by CMU, means that Alice is a member of CMU. It is easy to check that these three policies entail the formula $\mathsf{OAL}\ \mathsf{says}\ \mathtt{downloadAlice}$ in DTL_0, and that this would not be the case if we changed publ to says in the last policy.

3 Sequent Calculus

Now we describe a sequent calculus for DTL_0. Our presentation is inspired by earlier work on proof-theory for modal logics [25, 33]. Broadly, we follow Martin-Löf's judgmental method [31], and make a strong distinction between formulas and judgments. Judgments are the objects of knowledge, and are established through proofs. Formulas are subjects of judgments. For DTL_0, we use two basic (categorical) judgments: $A\ \mathsf{true}$, meaning that formula A is true, and $K\ \mathsf{claims}\ A$, meaning that principal K believes or claims that formula A is true. The two categorical judgments do not entail each other in general. $K\ \mathsf{says}\ A$ *internalizes* the judgment $K\ \mathsf{claims}\ A$ as a formula, allowing it to be combined with other connectives. In other words the judgments $(K\ \mathsf{says}\ A)\ \mathsf{true}$ and $K\ \mathsf{claims}\ A$ are equivalent.

To reason from hypothesis, we introduce hypothetical judgments (sequents) $\Gamma \xrightarrow{K} A\ \mathsf{true}$, informally meaning that principal K may reason from the hypothesis in Γ that A is true. Formally, the symbol Γ denotes a (possibly empty) multiset of categorical judgments, called the hypothesis or assumptions:

$$\Gamma \ ::= \ \cdot \ \mid \Gamma, A\ \mathsf{true} \ \mid \Gamma, K'\ \mathsf{claims}\ A$$

The principal K is called the *context* of the judgment. In context K, $K'\ \mathsf{claims}\ C$ entails $C\ \mathsf{true}$ if $K' \succeq K$. This is the only principle that distinguishes reasoning in one context from that in another. The formula A on the right of $\xrightarrow{K}$ is called the conclusion of the sequent.

The inference rules of the sequent calculus are shown in Figure 1. For brevity, we often elide the judgment name true, abbreviating $A\ \mathsf{true}$ to A. The notation $\Gamma|_K$ used in the rule

$$\frac{P \text{ atomic}}{\Gamma, P \xrightarrow{K} P}\text{init} \qquad \frac{\Gamma, K\ \mathsf{claims}\ A, A \xrightarrow{K'} C \qquad K \succeq K'}{\Gamma, K\ \mathsf{claims}\ A \xrightarrow{K'} C}\text{claims}$$

$$\frac{\Gamma|_K \xrightarrow{K} A}{\Gamma \xrightarrow{K'} K\ \mathsf{says}\ A}\mathsf{saysR} \qquad \frac{\Gamma, K\ \mathsf{says}\ A, K\ \mathsf{claims}\ A \xrightarrow{K'} C}{\Gamma, K\ \mathsf{says}\ A \xrightarrow{K'} C}\mathsf{saysL}$$

$$\frac{\Gamma \xrightarrow{K} A \qquad \Gamma \xrightarrow{K} B}{\Gamma \xrightarrow{K} A \wedge B}\wedge\text{R} \qquad \frac{\Gamma, A \wedge B, A, B \xrightarrow{K} C}{\Gamma, A \wedge B \xrightarrow{K} C}\wedge\text{L}$$

$$\frac{\Gamma \xrightarrow{K} A}{\Gamma \xrightarrow{K} A \vee B}\vee\text{R}_1 \qquad \frac{\Gamma \xrightarrow{K} B}{\Gamma \xrightarrow{K} A \vee B}\vee\text{R}_2 \qquad \frac{\Gamma, A \vee B, A \xrightarrow{K} C \qquad \Gamma, A \vee B, B \xrightarrow{K} C}{\Gamma, A \vee B \xrightarrow{K} C}\vee\text{L}$$

$$\frac{}{\Gamma \xrightarrow{K} \top}\top\text{R} \qquad \frac{}{\Gamma, \bot \xrightarrow{K} C}\bot\text{L}$$

$$\frac{\Gamma, A \xrightarrow{K} B}{\Gamma \xrightarrow{K} A \supset B}\supset\text{R} \qquad \frac{\Gamma, A \supset B \xrightarrow{K} A \qquad \Gamma, A \supset B, B \xrightarrow{K} C}{\Gamma, A \supset B \xrightarrow{K} C}\supset\text{L}$$

Figure 1: Sequent calculus for DTL_0

(saysR) stands for the multiset $\{(K'\ \mathsf{claims}\ C) \in \Gamma \mid K' \succeq K\}$. If we assume that formula A is true, we should certainly be able to conclude that A is true. For atomic formulas, this may be established by the rule (init); for others we prove it as a theorem (see Theorem 3.2).

The rules (claims), (saysR), and (saysL) characterize DTL_0. Read from the conclusion to the premises, the rule (claims) states that whenever we assume $K\ \mathsf{claims}\ A$, we are also justified in assuming that A is true, if we are reasoning in a context K' such that $K \succeq K'$. The rule (saysR) means that $K\ \mathsf{says}\ A$ may be established in any context if we can prove in context K that A is true using only assumptions $K''\ \mathsf{claims}\ C$ for $K'' \succeq K$. Observe that this is the only rule that changes the context of the sequent. The rule (saysL) captures the idea that $K\ \mathsf{says}\ A$ internalizes $K\ \mathsf{claims}\ A$: if we assume that $K\ \mathsf{says}\ A$ is true, then we may also assume $K\ \mathsf{claims}\ A$.

The rules for the connectives $\wedge$, $\vee$, $\top$, $\bot$, and $\supset$ are standard, except for a context which is associated with each sequent. We elide a description of these standard rules, and turn to the meta-theoretic properties of the sequent calculus.

3.1 Meta-Theory

Meta-theoretic properties, such as cut-elimination, are important from our perspective because proof-carrying authorization (our intended application) is heavily based in proof-checking, and proof-construction. Besides, meta-theoretic properties also imply that the inference rules of the logic fit well with each other, increasing faith in the logic's good foundation. Cut-elimination also means that all proofs can be normalized. Normalization is sometimes useful for auditing proofs of authorization.

Formally, the cut-elimination theorem states that adding a cut rule to a sequent calculus does not make more judgments provable. This is an easy consequence of the following theorem.

Theorem 3.1 (Admissibility of Cut)**.** *The following hold for the sequent calculus of Figure 1.*

1. $\Gamma \xrightarrow{K} A$ *and* $\Gamma, A \xrightarrow{K} C$ *imply that* $\Gamma \xrightarrow{K} C$.
2. $\Gamma|_K \xrightarrow{K} A$ *and* $\Gamma, K\ \textsf{claims}\ A \xrightarrow{K'} C$ *imply that* $\Gamma \xrightarrow{K'} C$.

Proof (Outline). Both statements can be proved simultaneously by lexicographic induction, first on the size of the cut formula A, and then on the size of the two given derivations, as in earlier work [32]. □

The logical dual of the cut-elimination theorem is identity, which states that whenever A $\textsf{true}$ is assumed as a hypothesis, we may conclude it. The following theorem captures this generalization of the (init) rule.

Theorem 3.2 (Identity)**.** *For each formula* A, $\Gamma, A \xrightarrow{K} A$.

Proof (Outline). By induction on A. □

Another theorem of interest for DTL_0 is subsumption, which states that contexts lower in the order $\succeq$ allow more provable formulas.

Theorem 3.3 (Subsumption)**.** *If* $\Gamma \xrightarrow{K} A$ *and* $K \succeq K'$*, then* $\Gamma \xrightarrow{K'} A$.

Proof (Outline). By induction on the given derivation of $\Gamma \xrightarrow{K} A$. □

Finally, we prove equivalence of the sequent calculus and the Hilbert-style system.

Theorem 3.4 (Equivalence)**.** $\Gamma \xrightarrow{K} A$ *if and only if* $\vdash K\ \textsf{says}\ (\Gamma \supset A)$.

Proof (Outline). In each direction by induction on the given derivations. For proving the "only-if" clause, we have to generalize the Hilbert-style system to allow hypothesis and prove the deduction theorem. This is done in the usual way. □

Observe that there is no equivalent of $\vdash B$ in the sequent calculus unless B has the form K says A. In this sense, the above theorem actually *embeds* the sequent calculus into the Hilbert-style system. While it is possible to recover the entire Hilbert-style system in the sequent calculus by adding non-indexed hypothetical judgments $\Gamma \longrightarrow A$, this extension seems uninteresting for authorization policies, and we omit it here.

4 Kripke Semantics for DTL$_0$

Next we describe sound and complete Kripke semantics for DTL$_0$. Although not directly applicable to policies, Kripke semantics are an invaluable tool for proving properties of the logic (e.g., [4, 24]). There is also hope that Kripke countermodels can be used as proofs of *failure*, in case an authorization does not succeed. Our presentation of Kripke semantics is inspired by work on the modal logic constructive S4 [7], and also uses some ideas from work on Kripke semantics of lax logic [22, 24].

The distinguishing characteristic of our Kripke semantics are *views* [24]. With each world w, we associate a set of principals $\theta(w)$ to whom the world is said to be visible. Our correctness property is that $\cdot \xrightarrow{K} A$ if and only if *each world visible to K satisfies A*.[2] In this manner, views allow us to distinguish reasoning in one context from that in another. If $K \succeq K'$ then we require that any world visible to K' also be visible to K. This ensures that context K validates fewer formulas than context K', and captures the subsumption principle (Theorem 3.3).

We model falsehood by explicitly specifying in each frame a set F of worlds where $\bot$ holds. These worlds are called fallible worlds [21, 22, 36]. We say that $w \models \bot$ iff $w \in F$. To model intuitionistic implication, we use a pre-order $\leq$ between worlds (as usual) and say that $w \models A \supset B$ iff for all w', $w \leq w'$ and $w' \models A$ imply $w' \models B$. Finally, to model the modality says, we use a principal-indexed binary relation $\sqsubseteq_K$ between worlds and define:

$$w \models K \text{ } \mathsf{says} \text{ } A \text{ iff either } w \in F \text{ or for all } w', w'', \; w \leq w' \sqsubseteq_K w'' \text{ implies } w'' \models A.$$

The clause $w \in F$ in the above definition is required to validate $\bot \supset K$ says A. The remaining definition is a generalization of satisfaction for $\Box A$ from Kripke semantics of constructive S4 [7]. To validate axiom (4), we stipulate that $\sqsubseteq_K ; \leq$ be a subset of $\sqsubseteq_K$.[3]

Both the use of a pre-order to model intuitionistic implication, and the use of different binary relations to model each modality are standard in modal logic. The novelty here is the interaction of these relations with views. We require that $\leq$ preserve views, i.e., if $w \leq w'$ and w be visible to K, then w' also be visible to K. We also require that whenever

[2] Throughout this section we use the sequent calculus of DTL$_0$ to state correctness properties. Use of the sequent calculus as opposed to the axiomatic system is partly a matter of personal taste and partly a matter of technical convenience.

[3] We believe that this condition can be weakened to $(\sqsubseteq_K ; \leq) \subseteq (\leq ; \sqsubseteq_K)$ without affecting the correctness of the Kripke semantics, but have not verified that this is the case.

$w \sqsubseteq_K w'$, w' be visible to K. For example, in the definition of $w \models K\ \mathsf{says}\ A$ above, w'' would be visible to K. By forcing these restrictions, we ensure that the semantics of all connectives except $K\ \mathsf{says}\ \cdot$ can be defined without changing views. On the other hand, the semantics of $K\ \mathsf{says}\ \cdot$ shift the reasoning to worlds that are visible to K. This subtle interaction between views and binary relations captures the exact meaning of formulas in DTL$_0$.

Definition 4.1 (Kripke Models)**.** A Kripke model M for DTL$_0$ is a tuple $(W, \theta, \leq, (\sqsubseteq_K)^{K \in \mathtt{Prin}}, \rho, F)$, where

- W is a non-empty set of worlds (worlds are denoted w).
- $\theta : W \mapsto 2^{\mathtt{Prin}}$ is a *view function* that maps each world w to a set of principals. If $K \in \theta(w)$, we say that w is visible to K, else w is said to be invisible to K. We often write W^K for the set $\{w \in W \mid K \in \theta(w)\}$. We require that:

  (View-closure) $K \in \theta(w)$ and $K' \succeq K$ imply $K' \in \theta(w)$.
- $\leq$ is a pre-order on W called the *implication relation*. We require that:

  (Imp-mon) $w \leq w'$ imply $\theta(w) \subseteq \theta(w')$.
- For each K, $\sqsubseteq_K$ is a subset of $W \times W^K$ called the *modality relation*. We require that:

  (Mod-refl) If $w \in W^K$, then $w \sqsubseteq_K w$.

  (Mod-trans) $\sqsubseteq_K$ be transitive.

  (Mod-closure) $w \sqsubseteq_K w'$ and $K' \succeq K$ imply $w \sqsubseteq_{K'} w'$

  (Commutativity) If $w \sqsubseteq_K w' \leq w''$, then $w \sqsubseteq_K w''$.
- $\rho : W \mapsto 2^{\mathtt{AtomicFormulas}}$ is a *valuation function* that maps each world to the set of atomic formulas that hold in it. We require that:

  (Rho-her) $P \in \rho(w)$ and $w \leq w'$ imply $P \in \rho(w')$.
- $F \subseteq W$ is the set of *fallible worlds.* We require that:

  (F-her) $w \in F$ and $w \leq w'$ imply $w' \in F$.

  (F-univ) $w \in F$ imply $P \in \rho(w)$

Definition 4.2 (Satisfaction)**.** Given a model $M = (W, \theta, \leq, (\sqsubseteq_K)^{K \in \mathtt{Prin}}, \rho, F)$, and a world $w \in W$, the satisfaction relation $w \models A$ (world w satisfies formula A) is defined by induction on A as follows.

$w \models P$ iff $P \in \rho(w)$.

$w \models A \wedge B$ iff $w \models A$ and $w \models B$.

$w \models A \vee B$ iff $w \models A$ or $w \models B$.

$w \models \top$.

$w \models \bot$ iff $w \in F$.

$w \models A \supset B$ iff for all w', $w \leq w'$ and $w' \models A$ imply $w' \models B$.

$w \models K$ $\textsf{says}$ A iff either $w \in F$ or for all w', w'', $w \leq w' \sqsubseteq_K w''$ implies $w'' \models A$.

We say that a principal K validates A in model M (written $M \models^K A$) if for each world $w \in W^K$ in M, it is the case that $w \models A$. The Kripke semantics defined above are sound and complete in the following sense.

Theorem 4.3 (Soundness and Completeness). $\cdot \xrightarrow{K} A$ *in the sequent calculus if and only if for each Kripke model M, $M \models^K A$.*

Soundness ("only if" direction) follows by an induction on the given sequent calculus proof. We must generalize the statement to allow non-empty hypotheses. The proof of completeness ("if" direction) uses a canonical model construction, which we omit here. The construction generalizes Alechina et al's construction of canonical models for constructive S4 [7].

DTL$_0$ as a Generalization of Constructive S4

In the special case where there is only one principal (say ℓ), DTL$_0$ reduces to the modal logic constructive S4. The sole modality ℓ $\textsf{says}$ A behaves exactly like the necessitation modality $\Box A$. The sequent calculus of Figure 1 reduces to a judgmental sequent calculus for constructive S4 (e.g., [25]). Similarly, the Kripke semantics reduce to those of constructive S4 described by Alechina et al [7], with the exception that our treatment of falsehood uses fallible worlds explicitly, and that DTL$_0$ does not have a modality corresponding to $\Diamond$. The following theorem is straightforward.

Theorem 4.4. *In the special case where there is only one principal ℓ, the following are equivalent:*

1. *$\vdash A$ treating ℓ $\textsf{says}$ $\cdot$ as the $\Box$ modality from constructive S4.*

2. *$\xrightarrow{\ell} A$ in the sequent calculus of Figure 1.*

3. *$\vdash \ell$ $\textsf{says}$ A in the axiomatic system of Section 2.*

Even though DTL_0 reduces to constructive S4 when there is only one principal, it is very different from the multi-modal constructive S4 obtained by taking independent S4 $\Box$ modalities (i.e., the logic S4 $\otimes$ S4... $\otimes$ S4). For example, the latter logic validates $(K \mathsf{\ says\ } K' \mathsf{\ says\ } A) \supset K' \mathsf{\ says\ } A$, which DTL_0 does not. In earlier work, we described the use of this logic for modeling *knowledge* in authorization policies [25].

5 Related Work

Many authorization logics have been proposed in the past, all of which contain the modality $K \mathsf{\ says\ } A$ [2, 3, 8–10, 15, 18, 20, 24–26, 29, 30]. The axioms and rules used in these logics differ widely. The particular combination of rules used in DTL_0 appears to be novel. Perhaps most closely related to DTL_0 is a proposal by Abadi in a survey paper [2], where the axiom $(K \mathsf{\ says\ } A) \supset (K' \mathsf{\ says\ } K \mathsf{\ says\ } A)$ is suggested. says with this axiom behaves very much like the defined connective publ in DTL_0. In a recent paper, Abadi studies connections between many possible axiomatizations of says, as well as higher level policy constructs such as delegation and control [4].

Also related to DTL_0 is work on languages for authorization (e.g., [11, 19, 27, 34]), most notably the languages Soutei and Binder [19, 34]. Our use of the term "context" is borrowed from the latter. Binder was also one of the earliest languages to explicitly define a notion of exporting policies from one context to another, which is very similar to publication of policies illustrated in Section 2. The pre-order $\succeq$ on principals draws on ideas from the Dependency Core Calculus [3, 5], where the modal indices are elements of a lattice.

Our Kripke semantics, as well as the completeness proof, are based on those of Alechina et al's work [7] for constructive S4. View functions were used earlier by the author and Abadi to describe semantics of authorization logics with lax-like modalities [24]. Fallible worlds have been used in the past to explain intuitionistic logic [21, 36], and also in semantics of lax logic [22]. It also appears to us that DTL_0 may be closely related to intuitionistic hybrid logics, and especially to the work of Chadha and others [17], but further investigation is needed to make an explicit connection. The presentation of the sequent calculus for DTL_0 is inspired by Pfenning and Davies' work on constructive S4 [33], and more directly by earlier work of the author and others [25].

6 Conclusion

We have presented a new constructive authorization logic, which explicitly relativizes hypothetical reasoning to the policies of individual principals. We have described the proof-theory and Kripke semantics of the logic. In ongoing work, we are considering extensions of the logic with first-order connectives, explicit time, and linearity to model other policy motifs. We are also translating existing authorization logics and languages for writing authorization policies to DTL$_0$, with the goal of understanding relations between the different formalisms.

There are several other avenues for future work. For instance, there seem to be strong connections between DTL$_0$ and hybrid logics. A useful generalization of DTL$_0$ would be to internalize the pre-order $\succeq$ as a formula. Such an extension would allow us to model delegation, along lines of the "speaks for" connective present in some authorization logics [3, 6, 24, 28]. Although the proof-theory of such an extension is relatively straightforward, it would be interesting to see its effects on Kripke semantics.

Acknowledgment. This work was partly supported by the Air Force Research Laboratory under grant no. FA87500720028. The author wishes to acknowledge Frank Pfenning for discussions and feedback on the logic and the paper, and Martín Abadi for feedback on the logic.

References

[1] SecPAL Preview release for .NET, 2006. http://research.microsoft.com/projects/SecPAL/.

[2] Martín Abadi. Logic in access control. In *Proceedings of the 18th Annual Symposium on Logic in Computer Science (LICS'03)*, pages 228–233, June 2003.

[3] Martín Abadi. Access control in a core calculus of dependency. *Electronic Notes in Theoretical Computer Science*, 172:5–31, April 2007. *Computation, Meaning, and Logic: Articles dedicated to Gordon Plotkin.*

[4] Martín Abadi. Variations in access control logic, 2008. Personal communication.

[5] Martín Abadi, Anindya Banerjee, Nevin Heintze, and Jon G. Riecke. A core calculus of dependency. In *Conference Record of the 26th Sympoisum on Principles Of Programming Languages (POPL'99)*, pages 147–160, San Antonio, Texas, January 1999. ACM Press.

[6] Martín Abadi, Michael Burrows, Butler Lampson, and Gordon Plotkin. A calculus for access control in distributed systems. *ACM Transactions on Programming Languages and Systems*, 15(4):706–734, 1993.

[7] Natasha Alechina, Michael Mendler, Valeria de Paiva, and Eike Ritter. Categorical and Kripke semantics for constructive S4 modal logic. In *CSL '01: Proceedings of the 15th International Workshop on Computer Science Logic*, pages 292–307, London, UK, 2001. Springer-Verlag.

[8] Andrew W. Appel and Edward W. Felten. Proof-carrying authentication. In G. Tsudik, editor, *Proceedings of the 6th ACM Conference on Computer and Communications Security*, pages 52–62, Singapore, November 1999. ACM Press.

[9] Lujo Bauer. *Access Control for the Web via Proof-Carrying Authorization.* PhD thesis, Princeton University, November 2003.

[10] Lujo Bauer, Scott Garriss, Jonathan M. McCune, Michael K. Reiter, Jason Rouse, and Peter Rutenbar. Device-enabled authorization in the Grey system. In *Information Security: 8th International Conference (ISC '05)*, Lecture Notes in Computer Science, pages 431–445, September 2005.

[11] Moritz Y. Becker, Cédric Fournet, and Andrew D. Gordon. Design and semantics of a decentralized authorization language. In *20th IEEE Computer Security Foundations Symposium*, pages 3–15, 2007.

[12] P. Blackburn, J. van Benthem, and F. Wolter. *Handbook of Modal Logic.* Elsevier B. V., 2007.

[13] Patrick Blackburn. Representation, reasoning, and relational structures: A hybrid logic manifesto. *Logic Journal of IGPL*, 8(3):339–365, 2000.

[14] Torben Braüner and Valeria de Paiva. Towards constructive hybrid logic. In *Electronic Proceedings of Methods for Modalities 3 (M4M3)*, 2003.

[15] J. G. Cederquist, R. Corin, M. A. C. Dekker, S. Etalle, J. I. den Hartog, and G. Lenzini. Audit-based compliance control. *International Journal of Information Security*, 6(2):133–151, 2007.

[16] Andrew Gordon Cédric Fournet and Sergio Maffeis. A type discipline for authorization in distributed systems. In *CSF '07: Proceedings of the 20th IEEE Computer Security Foundations Symposium*, pages 31–48. IEEE Computer Society, 2007.

[17] Rohit Chadha, Damiano Macedonio, and Vladimiro Sassone. A hybrid intuitionistic logic: Semantics and decidability. *Journal of Logic and Computation*, 16:27–59(33), February 2006.

[18] Jason Crampton, George Loizou, and Greg O' Shea. A logic of access control. *The Computer Journal*, 44(1):137–149, 2001.

[19] John DeTreville. Binder, a logic-based security language. In M. Abadi and S. Bellovin, editors, *Proceedings of the 2002 Symposium on Security and Privacy (S&P'02)*, pages 105–113, Berkeley, California, May 2002. IEEE Computer Society Press.

[20] Henry DeYoung, Deepak Garg, and Frank Pfenning. An authorization logic with explicit time. In *Proceedings of the 21st IEEE Computer Security Foundations Symposium (CSF-21)*, Pittsburgh, Pennsylvania, June 2008. IEEE Computer Society Press. To appear. Extended version available as Technical Report CMU-CS-07-166.

[21] M. Dummett. *Elements of Intuitionism.* Oxford University Press, 1977.

[22] M. Fairtlough and M.V. Mendler. Propositional lax logic. *Information and Computation*, 137(1):1–33, August 1997.

[23] Deepak Garg. Principal-centric reasoning in constructive authorization logic (full version), 2008. Available electronically from http://www.cs.cmu.edu/~dg.

[24] Deepak Garg and Martín Abadi. A modal deconstruction of access control logics. In *Proceedings of the 11th International Conference on Foundations of Software Science and Computation Structures (FoSSaCS 2008)*, pages 216–230, Budapest, Hungary, April 2008.

[25] Deepak Garg, Lujo Bauer, Kevin Bowers, Frank Pfenning, and Michael Reiter. A linear logic of affirmation and knowledge. In D. Gollman, J. Meier, and A. Sabelfeld, editors, *Proceedings of the 11th European Symposium on Research in Computer Security (ESORICS '06)*, pages 297–312, Hamburg, Germany, September 2006. Springer LNCS 4189.

[26] Deepak Garg and Frank Pfenning. Non-interference in constructive authorization logic. In J. Guttman, editor, *Proceedings of the 19th Computer Security Foundations Workshop (CSFW '06)*, pages 283–293, Venice, Italy, July 2006. IEEE Computer Society Press.

[27] Yuri Gurevich and Itay Neeman. DKAL: Distributed-knowledge authorization language. In *Proceedings of the 21st IEEE Symposium on Computer Security Foundations (CSF-21)*, 2008. To appear.

[28] Butler Lampson, Martín Abadi, Michael Burrows, and Edward Wobber. Authentication in distributed systems: Theory and practice. *ACM Transactions on Computer Systems*, 10(4):265–310, November 1992.

[29] Chris Lesniewski-Laas, Bryan Ford, Jacob Strauss, Robert Morris, and M. Frans Kaashoek. Alpaca: Extensible authorization for distributed services. In *Proceedings of the 14th ACM Conference on Computer and Communications Security (CCS-2007)*, Alexandria, VA, October 2007.

[30] Ninghui Li, Benjamin N. Grosof, and Joan Feigenbaum. Delegation logic: A logic-based approach to distributed authorization. *ACM Transactions on Information and Systems Security*, 6(1):128–171, 2003.

[31] Per Martin-Löf. On the meanings of the logical constants and the justifications of the logical laws. *Nordic Journal of Philosophical Logic*, 1(1):11–60, 1996.

[32] Frank Pfenning. Structural cut elimination I. Intuitionistic and classical logic. *Information and Computation*, 157(1/2):84–141, March 2000.

[33] Frank Pfenning and Rowan Davies. A judgmental reconstruction of modal logic. *Mathematical Structures in Computer Science*, 11:511–540, 2001.

[34] Andrew Pimlott and Oleg Kiselyov. Soutei, a logic-based trust-management system. In *Proceedings of the Eighth International Symposium on Functional and Logic Programming (FLOPS 2006)*, pages 130–145, 2006.

[35] Raymond M. Smullyan. *Forever Undecided*. Oxford University Press, 1988.

[36] A. S. Troelstra and D. Van Dalen. *Constructivism in Mathematics: Volume 2*. Elsevier Science Publishing Company, 1988.

[37] Jeffrey A. Vaughan, Limin Jia, Karl Mazurak, and Steve Zdancewic. Evidence-based audit. In *Proceedings of the 21st IEEE Symposium on Computer Security Foundations (CSF-21)*, 2008. To appear.