

Temporal Logics for Hyperproperties

Michael R. Clarkson¹, Bernd Finkbeiner², Masoud Koleini¹,
Kristopher K. Micinski³, Markus N. Rabe², and César Sánchez⁴

¹George Washington University; ²Universität des Saarlandes;
³University of Maryland, College Park; ⁴IMDEA Software Institute

Abstract. Two new logics for verification of hyperproperties are proposed. Hyperproperties characterize security policies, such as noninterference, as a property of sets of computation paths. Standard temporal logics such as LTL, CTL, and CTL* can refer only to a single path at a time, hence cannot express many hyperproperties of interest. The logics proposed here, HyperLTL and HyperCTL*, add explicit and simultaneous quantification over multiple paths to LTL and to CTL*. This kind of quantification enables expression of hyperproperties. A model checking algorithm for the proposed logics is given. For a fragment of HyperLTL, a prototype model checker has been implemented.

1 Introduction

Trace properties, which developed out of an interest in proving the correctness of programs [32], characterize correct behavior as properties of individual execution traces. Although early verification techniques specialized in proving individual correctness properties of interest, such as mutual exclusion or termination, temporal logics soon emerged as a general, unifying framework for expressing and verifying trace properties. Practical model checking tools [11, 16, 28] based on those logics now enable automated verification of program correctness.

Verification of security is not directly possible with such tools, because some important security policies cannot be characterized as properties of individual execution traces [38]. Rather, they are properties of sets of execution traces, also known as *hyperproperties* [15]. Specialized verification techniques have been developed for particular hyperproperties [5, 27, 41, 43], as well as for *2-safety* properties [52], which are properties of pairs of execution traces. But a unifying program logic for expressing and verifying hyperproperties could enable automated verification of a wide range of security policies.

In this paper, we propose two such logics. Both are based, like hyperproperties, on examining more than one execution trace at a time. Our first logic, *HyperLTL*, generalizes linear-time temporal logic (LTL) [44]. LTL implicitly quantifies over only a single execution trace of a system, but HyperLTL allows explicit quantification over multiple execution traces simultaneously, as well as propositions that stipulate relationships among those traces. For example, HyperLTL can express information-flow policies such as *observational determinism* [37, 46, 61], which requires programs to behave as (deterministic) functions

from low-security inputs to low-security outputs. The following two programs do not satisfy observational determinism, because they leak the value of high-security variable h through low-security variable l , thus making the program behave nondeterministically from a low-security user’s perspective:

$$(1) \ l := h \qquad (2) \ \mathbf{if} \ h = 0 \ \mathbf{then} \ l := 1 \ \mathbf{else} \ l := 0$$

Other program logics could already express observational determinism or closely related policies [7, 30, 41]. Milushev and Clarke [40–42] have even proposed other logics for hyperproperties, which we discuss in Section 8. But HyperLTL provides a simple and unifying logic in which many information-flow security policies can be directly expressed.

Information-flow policies are not one-size-fits-all. Different policies might be needed depending on the power of the adversary. For example, the following program does not satisfy observational determinism, but the program might be acceptable if nondeterministic choices, denoted \square , are resolved such that the probability distribution on output value l is uniform:

$$(3) \ l := h \ \square \ l := 0 \ \square \ l := 1$$

On the other hand, if the adversary can influence the resolution of nondeterministic choices, program (3) could be exploited to leak information. Similarly, the following program does satisfy observational determinism, but the program might be unacceptable if adversaries can monitor execution time:

$$(4) \ \mathbf{while} \ h > 0 \ \mathbf{do} \ \{h := h - 1\}$$

In Section 3, we show how policies appropriate for the above programs, as well as other security policies, can be formalized in HyperLTL.

Our second logic, *HyperCTL**, generalizes a branching-time temporal logic, CTL* [18]. Although CTL* already has explicit trace quantifiers, only one trace is ever in scope at a given point in a formula (see Section 5.1), so CTL* cannot directly express hyperproperties. But HyperCTL* can, because it permits quantification over multiple execution traces simultaneously. HyperLTL and HyperCTL* enjoy a similar relationship to that of LTL and CTL*: HyperLTL is the syntactic fragment of HyperCTL* containing only formulas in *prenex* form—that is, formulas that begin exclusively with quantifiers and end with a quantifier-free formula. HyperCTL* is thus a strict generalization of HyperLTL. HyperCTL* also generalizes a related temporal logic, SecLTL [17], and subsumes epistemic temporal logic [19, 54] (see Section 5).

Having defined logics for hyperproperties, we investigate model checking of those logics. In Section 6, we show that for HyperCTL* the model checking problem is decidable by reducing it to the satisfiability problem for quantified propositional temporal logic (QPTL) [50]. Since HyperCTL* generalizes HyperLTL, we immediately obtain that the HyperLTL model checking problem is also decidable. We present a hierarchy of fragments, which allows us to precisely characterize the complexity of the model checking problem in the number quantifier

alternations. The lowest fragment, which disallows any quantifier alternation, can be checked by a space-efficient polynomial-time algorithm (NLOGSPACE in the number of states of the program).

We also prototype a model checker that can handle an important fragment of HyperLTL, including all the examples from Section 3. The prototype implements a new model checking algorithm based on a well-known LTL algorithm [58, 59] and on a self-composition construction [7, 52]. The complexity of our algorithm is exponential in the size of the program and doubly exponential in the size of the formula—impractical for real-world programs, but at least a demonstration that model checking of hyperproperties formulated in our logic is possible.

This paper contributes to theoretical and foundational aspects of security by:

- defining two new program logics for expressing hyperproperties,
- demonstrating that those logics are expressive enough to formulate important information-flow policies,
- proving that the model checking problem is decidable, and
- prototyping a new model checking algorithm and using it to verify security policies.

The rest of the paper is structured as follows. Section 2 defines the syntax and semantics of HyperLTL. Section 3 provides several example formulations of information-flow policies. Section 4 defines the syntax and semantics of HyperCTL*. Section 5 compares our two logics with other temporal and epistemic logics. Section 6 obtains a model checking algorithm for HyperCTL*. Section 7 describes our prototype model checker. Section 8 reviews related work, and Section 9 concludes.

2 HyperLTL

HyperLTL extends propositional linear-time temporal logic (LTL) [44] with explicit quantification over *traces*. A *trace* is an infinite sequence of sets of *atomic propositions*. Let AP denote the set of all atomic propositions. The set TR of all traces is therefore $(2^{\text{AP}})^\omega$.

We first define some notation for manipulating traces. Let $t \in \text{TR}$ be a trace. We use $t[i]$ to denote element i of t , where $i \in \mathbb{N}$. Hence, $t[0]$ is the first element of t . We write $t[0, i]$ to denote the prefix of t up to and including element i , and $t[i, \infty]$ to denote the infinite suffix of t beginning with element i .

Syntax. Let π be a *trace variable* from an infinite supply \mathcal{V} of trace variables. Formulas of HyperLTL are defined by the following grammar:

$$\begin{aligned} \psi &::= \exists \pi. \psi \mid \forall \pi. \psi \mid \varphi \\ \varphi &::= a_\pi \mid \neg \varphi \mid \varphi \vee \varphi \mid X \varphi \mid \varphi \text{ U } \varphi \end{aligned}$$

Connectives \exists and \forall are universal and existential trace quantifiers, read as “along some traces” and “along all traces.” For example, $\forall \pi_1. \forall \pi_2. \exists \pi_3. \psi$ means that for all traces π_1 and π_2 , there exists another trace π_3 , such that ψ holds on those three traces. (Since branching-time logics also have explicit path quantifiers, it

is natural to wonder why one of them does not suffice to formulate hyperproperties. Section 5.1 addresses that question.) A HyperLTL formula is *closed* if all occurrences of trace variables are bound by a trace quantifier.

An atomic proposition a , where $a \in \text{AP}$, expresses some fact about states. Since formulas may refer to multiple traces, we need to disambiguate which trace the proposition refers to. So we annotate each occurrence of an atomic proposition with a trace variable π . Boolean connectives \neg and \vee have the usual classical meanings. Implication, conjunction, and bi-implication are defined as syntactic sugar: $\varphi_1 \rightarrow \varphi_2 \equiv \neg\varphi_1 \vee \varphi_2$, and $\varphi_1 \wedge \varphi_2 \equiv \neg(\neg\varphi_1 \vee \neg\varphi_2)$, and $\varphi_1 \leftrightarrow \varphi_2 \equiv (\varphi_1 \rightarrow \varphi_2) \wedge (\varphi_2 \rightarrow \varphi_1)$. True and false, written **true** and **false**, are defined as $a_\pi \vee \neg a_\pi$ and $\neg \text{true}$.

Temporal connective $X\varphi$ means that φ holds on the next state of every quantified trace. Likewise, $\varphi_1 U \varphi_2$ means that φ_2 will eventually hold of the states of all quantified traces that appear at the same index, and until then φ_1 holds. The other standard temporal connectives are defined as syntactic sugar: $F\varphi \equiv \text{true} U \varphi$, and $G\varphi \equiv \neg F \neg\varphi$, and $\varphi_1 W \varphi_2 \equiv (\varphi_1 U \varphi_2) \vee G\varphi_1$, and $\varphi_1 R \varphi_2 \equiv \neg(\neg\varphi_1 U \neg\varphi_2)$.

We also introduce syntactic sugar for comparing traces. Given a set P of atomic propositions, $\pi[0] =_P \pi'[0] \equiv \bigwedge_{a \in P} a_\pi \leftrightarrow a_{\pi'}$. That is, $\pi[0] =_P \pi'[0]$ holds whenever the first state in both π and π' agree on all the propositions in P . And $\pi =_P \pi' \equiv G(\pi[0] =_P \pi'[0])$, that is, all the positions of π and π' agree on P . The analogous definitions hold for \neq .

Semantics. The validity judgment for HyperLTL formulas is written $\Pi \models_T \psi$, where T is a set of traces, and $\Pi : \mathcal{V} \rightarrow \text{TR}$ is a *trace assignment* (i.e., a *valuation*), which is a partial function mapping trace variables to traces. Let $\Pi[\pi \mapsto t]$ denote the same function as Π , except that π is mapped to t . We write trace set T as a subscript on \models , because T propagates unchanged through the semantics; we omit T when it is clear from context. Validity is defined as follows:

$$\begin{array}{lll}
\Pi \models_T \exists\pi. \psi & \text{iff} & \text{there exists } t \in T : \Pi[\pi \mapsto t] \models_T \psi \\
\Pi \models_T \forall\pi. \psi & \text{iff} & \text{for all } t \in T : \Pi[\pi \mapsto t] \models_T \psi \\
\Pi \models_T a_\pi & \text{iff} & a \in \Pi(\pi)[0] \\
\Pi \models_T \neg\varphi & \text{iff} & \Pi \not\models_T \varphi \\
\Pi \models_T \varphi_1 \vee \varphi_2 & \text{iff} & \Pi \models_T \varphi_1 \text{ or } \Pi \models_T \varphi_2 \\
\Pi \models_T X\varphi & \text{iff} & \Pi[1, \infty] \models_T \varphi \\
\Pi \models_T \varphi_1 U \varphi_2 & \text{iff} & \text{there exists } i \geq 0 : \Pi[i, \infty] \models_T \varphi_2 \\
& & \text{and for all } 0 \leq j < i \text{ we have } \Pi[j, \infty] \models_T \varphi_1
\end{array}$$

Trace assignment *suffix* $\Pi[i, \infty]$ denotes the trace assignment $\Pi'(\pi) = \Pi(\pi)[i, \infty]$ for all π . If $\Pi \models_T \varphi$ holds for the empty assignment Π , then T *satisfies* φ .

We are interested in whether programs satisfy formulas, so we first derive a set T of traces from a program, first using *Kripke structures* as a unified representation of programs. A Kripke structure K is a tuple $(S, s_0, \delta, \text{AP}, L)$ comprising a set of *states* S , an *initial state* $s_0 \in S$, a transition function $\delta : S \rightarrow 2^S$, a set of *atomic propositions* AP , and a labeling function $L : S \rightarrow 2^{\text{AP}}$.

To ensure that all traces are infinite, we require that $\delta(s)$ is nonempty for every state s .

The set $\text{Traces}(K)$ of traces of K is the set of all sequences of labels produced by the state transitions of K starting from initial state. Formally, $\text{Traces}(K)$ contains trace t iff there exists a sequence $s_0s_1\dots$ of states, such that s_0 is the initial state, and for all $i \geq 0$, it holds that $s_{i+1} \in \delta(s_i)$; and $t[i] = L(s_i)$. A Kripke structure K *satisfies* φ , denoted by $K \models \varphi$, if $\text{Traces}(K)$ satisfies φ .

It will later be technically convenient to consider enlarging the set AP of atomic propositions permitted by a Kripke structure to a set AP' , such that $\text{AP} \subset \text{AP}'$. We extend $\text{Traces}(K)$ into the set of traces $\text{Traces}(K, \text{AP}')$ that is agnostic about whether each new proposition holds at each state. A trace $(P_0 \cup P'_0)(P_1 \cup P'_1)\dots \in \text{Traces}(K, \text{AP}')$ whenever $P_0P_1\dots \in \text{Traces}(K)$, and for all $i \geq 0$: $P'_i \subseteq \text{AP}' \setminus \text{AP}$. The final conjunct requires every possible set of new atomic propositions to be included in the traces.

3 Security Policies in HyperLTL

We now put HyperLTL into action by formulating several *information-flow security policies*, which stipulate how information may propagate from inputs to outputs. Information-flow is a very active field in security; see [20, 48] for surveys.

Noninterference. A program satisfies *noninterference* [23] when the outputs observed by low-security users are the same as they would be in the absence of inputs submitted by high-security users. Since its original definition, many variants with different execution models have been named “noninterference.” For clarity of our examples, we choose a simple *state-based* synchronous execution model in which atomic propositions of the traces contain the values of program variables, and in which progress of time corresponds to execution steps in the model. We also assume that the variables are partitioned into input and output variables, and into two security levels, *high* and *low*. (We could handle lattices of security levels by conjoining several formulas that stipulate noninterference between elements of the lattice.)

Noninference [38] is a variant of noninterference that can be stated in our simple system model. Noninference stipulates that, for all traces, the low-observable behavior must not change when all high inputs are replaced by a dummy input λ , that is, when the high input is removed. Noninference, a *liveness hyperproperty* [15], can be expressed in HyperLTL as follows:

$$\forall \pi. \exists \pi'. (\text{G } \lambda_{\pi'}) \wedge \pi =_L \pi' \tag{5}$$

where $\lambda_{\pi'}$ expresses that all of the high inputs in the current state of π' are λ , and $\pi =_L \pi'$ expresses that all low variables in π and π' have the same values.

Nondeterminism. Noninterference was introduced for use with deterministic programs. Nonetheless, nondeterminism naturally arises when program specifications abstract from implementation details, so many variants of noninterference

have been developed for nondeterministic programs. We formalize two variants here.

A (nondeterministic) program satisfies *observational determinism* [61] if every pair of traces with the same initial low observation remain indistinguishable for low users. That is, the program appears to be deterministic to low users. Programs that satisfy observational determinism are immune to *refinement attacks* [61], because observational determinism is preserved under refinement. Observational determinism, a *safety hyperproperty* [15], can be expressed in HyperLTL as follows:

$$\forall \pi. \forall \pi'. \pi[0] =_{L, \text{in}} \pi'[0] \rightarrow \pi =_{L, \text{out}} \pi' \quad (6)$$

where $\pi =_{L, \text{in}} \pi'$ and $\pi =_{L, \text{out}} \pi'$ express that both traces agree on the low input and low output variables, respectively.

Generalized noninterference (GNI) [35] permits nondeterminism in the low-observable behavior, but stipulates that low-security outputs may not be altered by the injection of high-security inputs. Like noninterference, GNI was originally formulated for event-based systems, but it can also be formulated for state-based systems [38]. GNI is a liveness hyperproperty and can be expressed as follows:

$$\forall \pi. \forall \pi'. \exists \pi''. \pi =_{H, \text{in}} \pi'' \wedge \pi' =_L \pi'' \quad (7)$$

The trace π'' in (7) is an *interleaving* of the high inputs of the first trace and the low inputs and outputs of the second trace. Other security policies based on interleavings, such as *restrictiveness* [36], *separability* [38], and *forward correctness* [39] can similarly be expressed in HyperLTL.

Declassification. Some programs need to reveal secret information to fulfill functional requirements. For example, a password checker must reveal whether the entered password is correct or not. The noninterference policies we have examined so far prohibit such behavior. More flexible security policies have been designed to permit *declassification* of information; see [49] for a survey.

With HyperLTL, we easily specify customized declassification policies. For example, suppose that a system inputs a password in its initial state, then declassifies whether that password is correct in the next state. The following policy (a safety hyperproperty) stipulates that leaking the correctness of the password is permitted, but that otherwise observational determinism must hold:

$$\forall \pi. \forall \pi'. (\pi[0] =_{L, \text{in}} \pi'[0] \wedge X(pw_\pi \leftrightarrow pw_{\pi'})) \rightarrow \pi =_{L, \text{out}} \pi' \quad (8)$$

where atomic proposition pw expresses that the entered password is correct.

Quantitative noninterference. Quantitative information-flow policies [12, 14, 24, 31] permit leakage of information at restricted rates. One way to measure leakage is with *min-entropy* [51], which quantifies the amount of information an attacker can gain given the answer to a single guess about the secret. The *bounding problem* [60] for min-entropy is to determine whether that amount is bounded

from above by a constant n . Assume that the program whose leakage is being quantified is deterministic, and assume that the secret input to that program is uniformly distributed. The bounding problem then reduces to determining that there is no tuple of $2^n + 1$ low-distinguishable traces [51,60] (a safety hyperproperty). We can express that as follows:

$$\neg \exists \pi_0. \dots \exists \pi_{2^n}. \left(\bigwedge_i \pi_i =_{L,\text{in}} \pi_0 \right) \wedge \bigwedge_{i \neq j} \pi_i \neq_{L,\text{out}} \pi_j \quad (9)$$

The initial negation can be pushed inside to obtain a proper HyperLTL formula.

Quantitative flow and entropy naturally bring to mind probabilistic systems. We haven't yet explored extending our logics to enable specification of policies that involve probabilities. Perhaps techniques previously used with epistemic logic [25] could be adapted; we leave this as future work.

Event-based systems. Our examples above use a synchronous state-based execution model. Many formulations of security policies, including the original formulation of noninterference [23], instead use an *event-based* system model, in which input and output events are not synchronized and have no relation to time. HyperLTL can express policies for asynchronous execution models, too. For example, HyperLTL can express the original definition of noninterference [23] and observational determinism; the companion technical report [13] shows how. The key idea is to allow the system to stutter and to quantify over all stuttered versions of the executions. We characterize the correct *synchronization* of a pair of traces as having updates to low variables only at the same positions. We then add an additional antecedent to the policy formula to require that only those pairs of traces that are synchronized correctly need to fulfill the security condition.

4 HyperCTL*

HyperLTL was derived from LTL by extending the models of formulas from single traces to sets of traces. However, like LTL, HyperLTL is restricted to linear time and cannot express branching-time properties (e.g., all states that succeed the current state satisfy some proposition). We show now that a branching-time logic for hyperproperties could be derived from a branching-time logic for trace properties, such as CTL* [18]. We call this logic HyperCTL*. The key idea is again to use sets instead of singletons as the models of formulas.

Syntax. HyperCTL* generalizes HyperLTL by allowing quantifiers to appear anywhere within a formula. Quantification in HyperCTL* is over *paths* through a Kripke structure. A path p is an infinite sequence of pairs of a state and a set of atomic propositions. Hence, a path differs from a trace by including a state of the Kripke structure in each element. Formally, $p \in (S \times 2^{\text{AP}})^\omega$, where S is the states of the Kripke structure. As with traces, $p[i]$ denotes the element i of

p , and $p[i, \infty]$ denotes the suffix of p beginning with element i . We also define a new notation: let $p(i)$ be the state in element i of p .

In HyperCTL*, π is a *path variable* and $\exists\pi$ is a *path quantifier*. Formulas of HyperCTL* are defined by the following grammar:

$$\varphi ::= a_\pi \mid \neg\varphi \mid \varphi \vee \varphi \mid X\varphi \mid \varphi U \varphi \mid \exists\pi. \varphi$$

We introduce all the syntactic sugar for derived logical operators, as for HyperLTL. The universal quantifier can now be defined as syntactic sugar, too: $\forall\pi. \varphi \equiv \neg\exists\pi. \neg\varphi$. A HyperCTL* formula is *closed* if all occurrences of some path variable π are in the scope of a path quantifier. A HyperCTL* *specification* is a Boolean combination of closed HyperCTL* formulas each beginning with a quantifier (or its negation).

Semantics. The validity judgment for HyperCTL* formulas is written $\Pi \models_K \varphi$, where K is a Kripke structure, and $\Pi : \mathcal{V} \rightarrow (S \times 2^{\text{AP}})^\omega$ is a *path assignment*, which is a partial function mapping path variables to paths. We write K as a subscript on \models , because K propagates unchanged through the semantics; we omit K when it is clear from context. Validity is defined as follows:

$$\begin{aligned} \Pi \models_K a_\pi & \quad \text{iff} \quad a \in L(\Pi(\pi)(0)) \\ \Pi \models_K \neg\varphi & \quad \text{iff} \quad \Pi \not\models_K \varphi \\ \Pi \models_K \varphi_1 \vee \varphi_2 & \quad \text{iff} \quad \Pi \models_K \varphi_1 \text{ or } \Pi \models_K \varphi_2 \\ \Pi \models_K X\varphi & \quad \text{iff} \quad \Pi[1, \infty] \models_K \varphi \\ \Pi \models_K \varphi_1 U \varphi_2 & \quad \text{iff} \quad \text{there exists } i \geq 0 : \Pi[i, \infty] \models_K \varphi_2 \\ & \quad \text{and for all } 0 \leq j < i \text{ we have } \Pi[j, \infty] \models_K \varphi_1 \\ \Pi \models_K \exists\pi. \varphi & \quad \text{iff} \quad \text{there exists } p \in \text{Paths}(K, \Pi(\pi')(0)) : \Pi[\pi \mapsto p] \models_K \varphi \end{aligned}$$

In the clause for existential quantification, π' denotes the path variable most recently added to Π (i.e., closest in scope to π). If Π is empty, let $\Pi(\pi')(0)$ be the initial state of K . It would be straightforward but tedious to further formalize this notation, so we omit the details. That clause uses another new notation, $\text{Paths}(K, s)$, which is the set of paths produced by Kripke structure K beginning from state s . Formally, $\text{Paths}(K, s)$ contains path p , where $p = (s_0, P_0)(s_1, P_1) \dots$ and $P_i \in 2^{\text{AP}}$, iff there exists a sequence $s_0 s_1 \dots$ of states, such that s_0 is s , and for all $i \geq 0$, it holds that $s_{i+1} \in \delta(s_i)$ and $P_i = L(s_i)$.

Like with Traces in HyperLTL, we define $\text{Paths}(K, s, \text{AP}')$ as follows: We have $(s_0, P_0 \cup P'_0)(s_1, P_1 \cup P'_1) \dots \in \text{Paths}(K, s, \text{AP}')$ iff $(s_0, P_0)(s_1, P_1) \dots \in \text{Paths}(K, s)$, and for all $i \geq 0$, it holds that $P'_i \subseteq \text{AP}' \setminus \text{AP}$.

We say that a Kripke structure K *satisfies* a HyperCTL* specification φ , denoted by $K \models \varphi$, if $\Pi \models_K \varphi$ holds true for the empty assignment. The *model checking problem* for HyperCTL* is to decide whether a given Kripke structure satisfies a given HyperCTL* specification.

HyperCTL vs. HyperLTL.* LTL can be characterized as the fragment of CTL* containing formulas of the form $A\varphi$, where A is the CTL* universal path quantifier and φ contains no quantifiers. Formula $A\varphi$ is satisfied in CTL* by a Kripke structure iff φ is satisfied in LTL by the traces of the Kripke structure.

A similar relationship holds between HyperLTL and HyperCTL*: HyperLTL can be characterized as the fragment of HyperCTL* containing formulas in *prenex* form—that is, a series of quantifiers followed by a quantifier-free formula. A formula φ in prenex form is satisfied in HyperCTL* by a Kripke structure iff φ is satisfied in HyperLTL by the traces of the Kripke structure. HyperCTL* is a strict generalization of HyperLTL, which extends HyperLTL with the capability to use quantified formulas as subformulas in the scope of temporal operators. For example, consider the program $(l := 0 \parallel l := 0) \parallel (l := 1 \parallel l := 1)$. A low-observer can infer which branch of the center-most nondeterministic choice is taken, but not which branch is taken next. This is expressed by HyperCTL* formula $\forall\pi. X \forall\pi'. X (l_\pi \leftrightarrow l_{\pi'})$. There is no equivalent HyperLTL formula.

As we show in Subsection 5.3, the temporal logic SecLTL [17] can be encoded in HyperCTL*, but not in HyperLTL. This provides further examples that distinguish HyperLTL and HyperCTL*.

5 Related Logics

We now examine the expressiveness of HyperLTL and HyperCTL* compared to several existing temporal logics: LTL, CTL*, QPTL, ETL, and SecLTL. There are many other logics that we could compare to in future work; some of those are discussed in Section 8.

5.1 Temporal Logics

HyperCTL* is an extension of CTL* and therefore subsumes LTL, CTL, and CTL*. Likewise, HyperLTL subsumes LTL. But temporal logics LTL, CTL, and CTL* cannot express information-flow policies. LTL formulas express properties of individual execution paths. All of the noninterference properties of Section 3 are properties of sets of execution paths [15, 38]. Explicit path quantification does enable their formulation in HyperLTL.

Even though CTL and CTL* have explicit path quantifiers, information-flow security policies, such as observational determinism (6), cannot be expressed with them. Consider the following fragment of CTL* semantics:

$$\begin{aligned} s \models A\varphi & \text{ iff for all } p \in \text{Paths}(K, s) : p \models \varphi \\ p \models \Phi & \text{ iff } p(0) \models \Phi \end{aligned}$$

Path formulas φ are modeled by paths p , and *state* formulas Φ are modeled by states s . State formula $A\varphi$ holds at state s when all paths proceeding from s satisfy φ . Any state formula Φ can be treated as a path formula, in which case Φ holds of the path iff Φ holds in the first state on that path. Using this semantics, consider the meaning of $AA\varphi$, which is the form of observational determinism (6):

$$\begin{aligned} s \models AA\varphi & \\ = \text{for all } p \in \text{Paths}(K, s) : p \models A\varphi & \\ = \text{for all } p \in \text{Paths}(K, s) \text{ and } p' \in \text{Paths}(K, s) : p' \models \varphi & \end{aligned}$$

Note how the meaning of $\text{AA } \varphi$ is ultimately determined by the meaning of φ , where φ is modeled by the single path p' . Path p is ignored in determining the meaning of φ ; the second universal path quantifier causes p to “leave scope.” Hence φ cannot express correlations between p and p' , as observational determinism requires. So CTL^* path quantifiers do not suffice to express information-flow policies. Neither do CTL path quantifiers, because CTL is a sub-logic of CTL^* . In fact, even the modal μ -calculus does not suffice to express some information-flow properties [2].

By using the self-composition construction [7, 52], it is possible to express relational noninterference in CTL [7] and observational determinism in CTL^* [30]. Those approaches resemble HyperCTL^* , but HyperCTL^* formulas express policies directly over the original system, rather than over a self-composed system. Furthermore, the self-composition approach does not seem capable of expressing policies that require both universal and existential quantifiers over infinite executions, like noninterference (5) and generalized noninterference (7). It is straightforward to express such policies in our logics.

QPTL. Quantified propositional temporal logic (QPTL) [50] extends LTL with quantification over propositions, whereas HyperLTL extends LTL with quantification over traces. Quantification over traces is more powerful than quantification over propositions, as we now show.

QPTL formulas are generated by the following grammar, where $a \in \text{AP}$:

$$\psi ::= a \mid \neg\psi \mid \psi \vee \psi \mid X\psi \mid F\psi \mid \exists a. \psi$$

All QPTL connectives have the same semantics as in LTL, except for propositional quantification:

$$p \models \exists a. \psi \quad \text{iff} \quad \text{there exists } p' \in (2^{\text{AP}})^\omega : p =_{\text{AP} \setminus a} p' \text{ and } p' \models \psi.$$

Theorem 1. *HyperLTL subsumes QPTL, but QPTL does not subsume HyperLTL.*

Proof sketch. To express a QPTL formula in HyperLTL , rewrite the formula to prenex form, and rename all bound propositions with unique fresh names from a set AP' . These propositions act as free variables, which are unconstrained because they do not occur in the Kripke structure. Replace each propositional quantification $\exists a$ in the QPTL formula by a path quantification $\exists \pi_a$ in the HyperLTL formula. And replace each occurrence of a by a_{π_a} . The result is a HyperLTL formula that holds iff the original QPTL formula holds.

But not all HyperLTL formulas can be expressed in QPTL. For example, QPTL cannot express properties that require the existence of paths, such as $\exists \pi. X a_\pi$. \square

In Section 6, we exploit the relationship between HyperLTL and QPTL to obtain a model checking algorithm for HyperLTL .

5.2 Epistemic Logics

HyperLTL and HyperCTL* express information-flow policies by explicit quantification over multiple traces or paths. Epistemic temporal logic has also been used to express such policies [4, 10, 26, 55] by implicit quantification over traces or paths with the *knowledge* connective K of epistemic logic [19]. We do not yet know which is more powerful, particularly for information-flow policies. But we do know that HyperLTL subsumes a common epistemic temporal logic.

Define ETL (epistemic temporal logic) to be LTL with the addition of K under its perfect recall semantics [4, 19, 54]. The model of an ETL formula is a pair (K, Agts) of a Kripke structure K and a set Agts of equivalence relations on AP , called the *agents*; each relation models the knowledge of an agent. (*Interpreted systems*, rather than Kripke structures, are often used to model ETL formulas [19, 54]. Interpreted systems differ in style but can be translated to our formulation.) In the *asynchronous* semantics of ETL, $K_A\psi$ holds on state i of trace $t \in \text{Traces}(K)$, denoted $t, i \models K_A\varphi$, iff

$$\text{for all } t' \in \text{Traces}(K) : t[0, i] \approx_A t'[0, i] \text{ implies } t', i \models \varphi,$$

where \approx_A denotes stutter-equivalence on finite traces with respect to A . In the *synchronous* semantics of ETL, stutter-equivalence is replaced by stepwise-equivalence.

The following two theorems show that HyperLTL subsumes ETL:

Theorem 2. *In the synchronous semantics, for every ETL formula ψ and every set Agts of agents, there exists a HyperLTL formula φ such that for all Kripke structures K , we have $(K, \text{Agts}) \models \psi$ iff $K \models \varphi$.*

Theorem 3. *In the asynchronous semantics, for every ETL formula ψ and every set Agts of agents, there exists a HyperLTL formula φ such that for all asynchronous Kripke structures K , we have $(K, \text{Agts}) \models \psi$ iff $K \models \varphi$.*

Proofs of both theorems appear in the companion technical report [13]. Theorem 3 requires an additional assumption that K is an *asynchronous* Kripke structure, i.e. that it can always stutter in its current state and that it is indicated in an atomic proposition whether the last state was a stuttering step.

HyperLTL and ETL have the same worst-case complexity for model checking, which is non-elementary. But, as we show in Section 6, the complexity of our model checking algorithm on the information-flow policies of Section 3 is much better—only NLOGSPACE (for observational determinism, declassification, and quantitative noninterference for a fixed number of bits) or PSPACE (for noninterference and generalized noninterference) in the size of the system. For those policies in NLOGSPACE, that complexity, unsurprisingly, is as good as algorithms based on self-composition [7]. This ability to use a general-purpose, efficient HyperLTL model checking algorithm for information flow seems to be an improvement over encodings of information flow in ETL.

5.3 SecLTL

SecLTL [17] extends temporal logic with the *hide* modality \mathcal{H} , which allows to express information flow properties such as noninterference [23]. The semantics of SecLTL is defined in terms of labeled transition system, where the edges are labeled with valuations of the set of variables. The formula $\mathcal{H}_{H,O}\varphi$ specifies that the current valuations of a subset H of the input variables I are kept secret from an attacker who may observe variables in O until the release condition φ becomes true. The semantics is formalized in terms of a set of *alternative paths* to which the *main path* is compared:

$$\text{AltPaths}(p, H) = \{p' \in \text{Paths}(K_M, p[0]) \mid p[1] =_{I \setminus H} p'[1] \text{ and } p[2, \infty] =_I p'[2, \infty]\}$$

where K_M is the equivalent Kripke structure for the labeled transition system M (we will explain the translation later in this section.) A path p satisfies the SecLTL formula $\mathcal{H}_{H,O}\varphi$, denoted by $p \models \mathcal{H}_{H,O}\varphi$, iff

$$\forall p' \in \text{AltPaths}(p, H). (p =_O p', \text{ or there exists } i \geq 0 : \\ p[i, \infty] \models_K \varphi \text{ and } p[1, i-1] =_O p'[1, i-1])$$

A labeled transition system M satisfies a SecLTL formula ψ , denoted by $M \models \psi$, if every path p starting in the initial state satisfies ψ .

SecLTL can express properties like the dynamic creation of secrets discussed in Section 4, which cannot be expressed by HyperLTL. However, SecLTL is subsumed by HyperCTL*. To encode the hide modality in HyperCTL*, we first translate M into a Kripke structure K_M , whose states are labeled with the valuation of the variables on the edge leading into the state. The initial state is labeled with the empty set. In the modified system, $L(p[1])$ corresponds to the current labels. We encode $\mathcal{H}_{H,O}\varphi$ as the following HyperCTL* formula:

$$\forall \pi'. \pi[1] =_{I \setminus H} \pi'[1] \wedge X (\pi[1] =_O \pi'[1] \text{ W } (\pi[1] \neq_I \pi'[1] \vee \varphi))$$

Theorem 4. *For every SecLTL formula ψ and transition system M , there is a HyperCTL* formula φ such that $M \models \psi$ iff $K_M \models \varphi$.*

The model checking problem for SecLTL is PSPACE-hard in the size of the Kripke structure [17]. The encoding of SecLTL specifications in HyperCTL* implies that the model checking problem for HyperCTL* is also PSPACE-hard (for a fixed specification of alternation depth ≥ 1), as claimed in Theorem 6.

6 Model Checking and Satisfiability

In this section we exploit the connection between HyperCTL* and QPTL to obtain a model checking algorithm for HyperCTL* and study its complexity. We identify a hierarchy of fragments of HyperCTL* characterized by the number of quantifier alternations. This hierarchy allows us to give a precise characterization of the complexity of the model checking problem. The fragment of formulas with quantifier alternation depth 0 includes already many formulas of interest and our result provides an NLOGSPACE algorithm in the size of the Kripke structure.

Definition 1 (Alternation Depth). A *HyperCTL** formula φ in NNF has alternation depth 0 plus the highest number of alternations from existential to universal and universal to existential quantifiers along any of the paths of the formula's syntax tree starting in the root. Occurrences of U and R count as an additional alternation.

Theorem 5. The model checking problem for *HyperCTL** specifications φ with alternation depth k on a Kripke structure K is complete for $\text{NSPACE}(g_c(k, |\varphi|))$ and it is in $\text{NSPACE}(g_c(k - 1, |K|))$ for some $c > 0$.

The function $g_c(x, y)$ denotes a tower of exponentials of height x with argument y : $g_c(0, y) = y$ and $g_c(x, y) = c^{g_c(x-1, y)}$. $\text{NSPACE}(g_c(x, y))$ denotes the class of languages accepted by a Turing machine bounded in space by $O(g_c(x, y))$. Abusing notation, we define $g_c(-1, y) = \log y$ and $\text{NSPACE}(\log y) = \text{NLOGSPACE}$ in y .

Proof. Both directions, the lower bound and the upper bound, are based on the complexity of the satisfiability problem for QPTL formulas φ in prenex normal form and with alternation depth k , which is complete for $\text{NSPACE}(g(k, |\varphi|))$ [50].

For the upper bound on the *HyperCTL** model checking complexity, we first translate until operators $\psi U \psi'$ as $\exists t. t \wedge G(t \rightarrow \psi' \vee (\psi \wedge X t)) \wedge \neg G t$. Let $\psi(K, \text{AP}')$ encode a Kripke structure K , where $K = (S, s_0, \delta, \text{AP}, L)$, as a QPTL formula (cf. [33]) using the set of atomic propositions AP' , which must contain atomic propositions replacing those of AP and additional atomic propositions to describe the states S . The formula $\psi(K, \text{AP}')$ is linear in $|K|$ and does not require additional quantifiers.

*HyperCTL** path quantifiers $\exists \pi. \varphi$ and $\forall \pi. \varphi$ are then encoded as $\exists \text{AP}_\pi. \psi(K, \text{AP}_\pi) \wedge \varphi_{\text{AP}_\pi}$ and $\forall \text{AP}_\pi. \psi(K, \text{AP}_\pi) \rightarrow \varphi_{\text{AP}_\pi}$, where AP_π is a set of *fresh* atomic propositions including a copy of AP and additional atomic propositions to describe the states S . The formula φ_{AP_π} is obtained from φ by replacing all atomic propositions referring to path π by their copies in AP_π . Atomic propositions in the formula that are not in AP (i.e. their interpretation is not fixed in K) need to be added to the sets AP_π accordingly.

For the lower bound, we reduce the satisfiability problem for a given QPTL formula φ in prenex normal form to a model checking problem $K \models \varphi'$ of *HyperCTL**. We assume, without loss of generality, that φ is closed (if a free proposition occurs in φ , we bind it with an existential quantifier) and each quantifier in φ introduces a different proposition.

The Kripke structure K consists of two states $S = \{s_0, s_1\}$, is fully connected $\delta(s) = S$ for all $s \in S$, and has a single atomic proposition $\text{AP} = \{p\}$. The states are labeled as follows: $L(s_0) = \emptyset$ and $L(s_1) = \{p\}$. Essentially, paths in K can encode all sequences of valuations of a variable in QPTL. To obtain the *HyperCTL** formula, we now simply replace every quantifier in the QPTL formula with a path quantifier. The only technical problem left is that quantification in QPTL allows to choose freely the value of p in the current state, while path quantification in *HyperCTL** only allows the path to differ in the next state. We solve the issue by shifting the propositions using a next operator. \square

Lower bounds in $|K|$. An NLOGSPACE *lower* bound in the size of the Kripke structure for fixed specifications with alternation depth 0 follows from the non-emptiness problem of non-deterministic Büchi automata. For alternation depth 1 and more we can derive PSPACE hardness in the size of the Kripke structure from the encoding of the logic SecLTL into HyperCTL* (see Subsection 5.3).

The result can easily be transferred to HyperLTL, since in the SecLTL formula that is used to prove PSPACE hardness, the Hide operator does not occur in the scope of temporal operators and hence the translation yields a HyperLTL formula.

Theorem 6. *For HyperLTL formulas the model checking problem is hard for PSPACE in the size of the system.*

A Remark on Efficiency The use of the standard encoding of the until operator in QPTL with an additional quantifier shown above is, in certain cases, wasteful. The satisfiability of QPTL formulas can be checked with an automata-theoretic construction, where we first transform the formula into prenex normal form, then generate a nondeterministic Büchi automaton for the quantifier-free part of the formula, and finally apply projection and complementation to handle the existential and universal quantifiers. In this way, each quantifier alternation, including the alternation introduced by the encoding of the until operators, causes an exponential blow-up. However, if an until operator occurs in the quantifier-free part, the standard transformation of LTL formulas to nondeterministic Büchi automata handle this until operator without requiring a quantifier elimination, resulting in an exponential speedup.

Using this insight, the model checking complexity for many of the formulas presented above and in Section 3 can be reduced by one exponent. Additionally, the complexity with respect to the size of the system reduces to NLOGSPACE for HyperCTL* formulas where the leading quantifiers are all of the same type and are followed by some quantifier-free formula which may contain until operators without restriction. Observational determinism and the declassification policy discussed in Section 3 are examples for specifications in this fragment. This insight was used for the prototype implementation described in Section 7 and it avoids an additional complementation step for noninference (5).

Satisfiability. The positive result regarding the model checking problem for HyperCTL* does not carry over to the satisfiability problem. The *finite-state satisfiability problem* consists of the existence of a finite model, while the *general satisfiability problem* asks for the existence of a possibly infinite model.

Theorem 7. *For HyperCTL*, finite-state satisfiability is hard for Σ_1^0 and general satisfiability is hard for Σ_1^1 .*

In the proof, located in the companion technical report [13], we reduce the LTL synthesis problem of distributed systems to the satisfiability problem of HyperCTL*.

7 Prototype Model Checker

The results of the previous section yield a model checking algorithm for all of HyperCTL*. But most of our information-flow policy examples do not require the full expressiveness of HyperCTL*. In fact, we have been able to implement a prototype model checker for an expressive fragment of the logic mostly using off-the-shelf components.

Define *HyperLTL₂* as the fragment of HyperLTL (and of HyperCTL*) in which the series of quantifiers at the beginning of a formula may involve at most one alternation. Every formula in HyperLTL₂ thus may begin with at most two (whence the name) kinds of quantifiers—a sequence of \forall 's followed by a sequence of \exists 's, or vice-versa. For example, $\exists\pi.\psi$ and $\forall\pi_1.\forall\pi_2.\exists\pi_3.\psi$ are allowed, but $\forall\pi_1.\exists\pi_2.\forall\pi_3.\psi$ is not. HyperLTL₂ suffices to express all the security policies formulated in Section 3. (Another logic for hyperproperties, \mathcal{L}_μ^k [41], similarly restricts fixpoint operator alternations with no apparent loss in expressivity for security policies.)

Our model checking algorithm for HyperLTL₂, detailed in the companion technical report [13], is based on algorithms for LTL model checking [21, 22, 57]. Those LTL algorithms determine whether a Kripke structure satisfies an LTL formula by performing various automata constructions and by checking language containment. Our algorithm likewise uses automata constructions and language containment, as well as self composition [7, 52] and a new *projection* construction.

We prototyped this algorithm in about 3,000 lines of OCaml code. Our prototype accepts as input a Kripke structure and a HyperLTL₂ formula, then constructs the automata required by our algorithm, and outputs a countermodel if the formula does not hold of the structure. For automata complementation, our prototype outsources to GOAL [53], an interactive tool for manipulating Büchi automata. We have used the prototype to verify noninterference (5), observational determinism (6), and generalized noninterference (7) for small Kripke structures (up to 10 states); running times were about 10 seconds or less.

Since our algorithm uses automata complementation, the worst-case running time is exponential in the size of the Kripke structure's state space and doubly exponential in the formula size. So as one might expect, our prototype currently does not scale to medium-sized Kripke structures (up to 1,000 states). But our purpose in building this prototype was to demonstrate a proof-of-concept for model checking of hyperproperties. We conjecture that practical symbolic model checking algorithms, such as BMC and IC3, could be used to scale up our approach to real-world systems.

8 Related Work

McLean [38] formalizes security policies as closure with respect to *selective interleaving functions*. He shows that trace properties cannot express security policies such as noninterference and average response time, because those are not properties of single execution traces. Mantel [34] formalizes security policies with *basic security predicates*, which stipulate *closure conditions* for trace sets.

Clarkson and Schneider [15] introduce *hyperproperties*, a framework for expressing security policies. Hyperproperties are sets of trace sets, and are able to formalize security properties such as noninterference, generalized noninterference, observational determinism and average response time. Clarkson and Schneider use second-order logic to formulate hyperproperties. That logic isn't verifiable, in general, because it cannot be effectively and completely axiomatized. Fragments of it, such as HyperLTL and HyperCTL*, can be verified.

Alur et al. [2] show that modal μ -calculus is insufficient to express all *opacity* policies [9], which prohibit observers from discerning the truth of a predicate. (Alur et al. [2] actually write “secrecy” rather than “opacity.”) Simplifying definitions slightly, a trace property P is *opaque* iff for all paths p of a system, there exists another path p' of that system, such that p and p' are low-equivalent, and exactly one of p and p' satisfies P . Noninference (5) is an opacity policy [47] that HyperLTL can express.

Huisman et al. reduce observational-determinism properties to properties in CTL* [30] and in modal μ -calculus [29] on a self-composed system. Barthe et al. use self composition to verify observational determinism [7] and noninterference [6] on terminating programs. Van der Meyden and Zhang [56] reduce a broader class of information-flow policies to safety properties on a self-composed system expressible in standard linear and branching time logics, and use model checking to verify noninterference policies. Their methodology requires customized model checking algorithms for each security policy, whereas this work proposes a single algorithm for all policies.

Balliu et al. [4] use a linear-time temporal epistemic logic to specify many declassification policies derived from noninterference. Their definition of noninterference, however, seems to be that of observational determinism (6). They do not consider any information-flow policies involving existential quantification, such as noninference. They also do not consider systems that accept inputs after execution has begun. Halpern and O'Neill [26] use a similar temporal epistemic logic to specify *secrecy* policies, which subsume many definitions of noninterference; they do not pursue model checking algorithms.

Alur et al. [1] discuss branching-time logics with *path equivalences* that are also able to express certain security properties. The authors introduce operators that resemble the knowledge operator of epistemic logics. As the logics build on branching-time logics they are not subsumed by HyperLTL. The relationship to HyperCTL* is still open.

Milushev and Clarke [40–42] propose three logics for hyperproperties:

- *Holistic hyperproperty logic* \mathcal{HL} , which is based on coinductive predicates over streams. Holistic hyperproperties “talk about whole traces at once; their specifications tend to be straightforward, but they are difficult to reason about, exemplified by the fact that no general approach to verifying such hyperproperties exists” [40]. HyperLTL and HyperCTL* are logics that talk about whole traces at once, too; and they have straightforward specifications as well as a general approach to verification.

- *Incremental hyperproperty logic* \mathcal{IL} is a fragment of *least fixed-point logic* [8]. There is a manual verification methodology for \mathcal{IL} [40], but no automated decision procedure.
- Another incremental hyperproperty logic \mathcal{IL}_μ^k , a fragment of polyadic modal μ -calculus [3] that permits at most one quantifier alternation (a greatest fixed-point followed by a least fixed-point). There is an automated model checking technique [41] for \mathcal{IL}_μ^k based on *parity games*. That technique has been prototyped and applied to a few programs.

All these logics suffice to express security policies such as noninterference and generalized noninterference. Like our logics, the exact expressive limitation is still an open problem.

As the preceding discussion makes clear, the expressiveness of HyperLTL and HyperCTL* versus several other logics is an open question. It’s possible that some of those logics will turn out to be more expressive or more efficiently verifiable than HyperLTL or HyperCTL*. It’s also possible that it will turn out to be simply a matter of taste which style of logic is more suitable for hyperproperties. The purpose of this paper was to explore one design option: a familiar syntax, based on widely-used temporal logics, that can straightforwardly express well-known hyperproperties.

9 Concluding Remarks

In designing a logic for hyperproperties, starting with HyperLTL was natural, because hyperproperties are sets of trace sets, and LTL uses trace sets to model programs. From HyperLTL, the extension to HyperCTL* was also natural: we simply removed the restrictions on where quantifiers could appear. The curtailment to HyperLTL₂ was also natural, because it was the fragment needed to express information-flow security policies. HyperLTL₂ permits up to one quantifier alternation, but what about hyperproperties with more? We do not yet know of any security policies that are examples. As Rogers [45] writes, “The human mind seems limited in its ability to understand and visualize beyond four or five alternations of quantifier. Indeed, it can be argued that the inventions... of mathematics are devices for assisting the mind in dealing with one or two additional alternations of quantifier.” For practical purposes, we might not need to go much higher than one quantifier alternation.

Acknowledgements. Fred B. Schneider suggested the name “HyperLTL.” We thank him, Rance Cleaveland, Rayna Dimitrova, Dexter Kozen, José Meseguer, and Moshe Vardi for discussions about this work. Adam Hinz worked on an early prototype of the model checker. This work was supported in part by AFOSR grant FA9550-12-1-0334, NSF grant CNS-1064997, the German Research Foundation (DFG) under the project SpAGAT within the Priority Program 1496 “Reliably Secure Software Systems — RS³,” and Spanish Project “TIN2012-39391-C04-01 STRONGSOFT.”

References

1. R. Alur, P. Cerný, and S. Chaudhuri. Model checking on trees with path equivalences. In *Proc. Tools and Algorithms for the Construction and Analysis of Systems*, pages 664–678, Mar. 2007.
2. R. Alur, P. Černý, and S. Zdancewic. Preserving secrecy under refinement. In *Proc. International Colloquium on Automata, Languages and Programming*, pages 107–118, 2006.
3. H. R. Andersen. A polyadic modal mu-calculus. Technical Report 1994-145, Technical University of Denmark (DTU), 1994.
4. M. Balliu, M. Dam, and G. L. Guernic. Epistemic temporal logic for information flow security. In *Proc. Workshop on Programming Languages and Analysis for Security*, June 2011.
5. A. Banerjee and D. A. Naumann. Stack-based access control and secure information flow. *Journal of Functional Programming*, 15(2):131–177, 2005.
6. G. Barthe, J. M. Crespo, and C. Kunz. Beyond 2-safety: asymmetric product programs for relational program verification. In *Proc. Logical Foundations of Computer Science*, pages 29–43, Jan. 2013.
7. G. Barthe, P. R. D’Argenio, and T. Rezk. Secure information flow by self-composition. In *Proc. IEEE Computer Security Foundations Workshop*, pages 100–114, June 2004.
8. J. Bradfield and C. Stirling. Modal mu-calculi. In *Handbook of Modal Logic*, pages 721–756. Elsevier, Amsterdam, The Netherlands, 2007.
9. J. Bryans, M. Koutny, L. Mazaré, and P. Y. A. Ryan. Opacity generalised to transition systems. In *Proc. Formal Aspects in Security and Trust*, pages 81–95, July 2005.
10. R. Chadha, S. Delaune, and S. Kremer. Epistemic logic for the applied pi calculus. In *Proc. Formal Techniques for Distributed Systems*, pages 182–197, June 2009.
11. A. Cimatti, E. M. Clarke, F. Giunchiglia, and M. Roveri. NuSMV: A new symbolic model verifier. In *Proc. Computer Aided Verification*, pages 495–499, July 1999.
12. D. Clark, S. Hunt, and P. Malacaria. Quantified interference for a while language. *Electronic Notes in Theoretical Computer Science*, 112:149–166, Jan. 2005.
13. M. R. Clarkson, B. Finkbeiner, M. Koleini, K. K. Micinski, M. N. Rabe, and C. Sánchez. Temporal logics for hyperproperties, Jan. 2014. Available at <http://arxiv.org/abs/1401.4492>.
14. M. R. Clarkson, A. C. Myers, and F. B. Schneider. Quantifying information flow with beliefs. *Journal of Computer Security*, 17(5):655–701, Oct. 2009.
15. M. R. Clarkson and F. B. Schneider. Hyperproperties. *Journal of Computer Security*, 18(6):1157–1210, 2010.
16. B. Cook, E. Koskinen, and M. Vardi. Temporal property verification as a program analysis task. In *Proc. Computer Aided Verification*, pages 333–348, July 2011.
17. R. Dimitrova, B. Finkbeiner, M. Kovács, M. N. Rabe, and H. Seidl. Model checking information flow in reactive systems. In *Proc. International Conference on Verification, Model Checking, and Abstract Interpretation*, pages 169–185, Jan. 2012.
18. E. A. Emerson and J. Y. Halpern. “Sometimes” and “not never” revisited: On branching versus linear time temporal logic. *Journal of the ACM*, 33(1):151–178, Jan. 1986.
19. R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning About Knowledge*. MIT Press, Cambridge, 1995.

20. R. Focardi and R. Gorrieri. Classification of security properties (Part I: Information flow). In *International School on Foundations of Security Analysis and Design: Tutorial Lectures*, pages 331–396, 2000.
21. P. Gastin and D. Oddoux. Fast LTL to Büchi automata translation. In *Proc. Computer Aided Verification*, pages 53–65, July 2001.
22. R. Gerth, D. Peled, M. Y. Vardi, and P. Wolper. Simple on-the-fly automatic verification of linear temporal logic. In *Proc. Protocol Specification, Testing and Verification*, pages 3–18, June 1995.
23. J. A. Goguen and J. Meseguer. Security policies and security models. In *Proc. IEEE Symposium on Security and Privacy*, pages 11–20, Apr. 1982.
24. J. W. Gray, III. Toward a mathematical foundation for information flow security. In *Proc. IEEE Symposium on Security and Privacy*, pages 210–34, May 1991.
25. J. W. Gray, III and P. F. Syverson. A logical approach to multilevel security of probabilistic systems. *Distributed Computing*, 11(2):73–90, 1998.
26. J. Y. Halpern and K. R. O’Neill. Secrecy in multiagent systems. *ACM Transactions on Information and System Security*, 12(1):5:1–47, Oct. 2008.
27. C. Hammer and G. Snelling. Flow-sensitive, context-sensitive, and object-sensitive information flow control based on program dependence graphs. *International Journal of Information Security*, 8(6):399–422, Dec. 2009.
28. G. J. Holzmann. The model checker SPIN. *IEEE Transactions on Software Engineering*, 23:279–295, 1997.
29. M. Huisman and H. Blondeel. Model-checking secure information flow for multi-threaded programs. In *Proc. Theory of Security and Applications*, pages 148–165, Mar. 2011.
30. M. Huisman, P. Worah, and K. Sunesen. A temporal logic characterisation of observational determinism. In *Proc. IEEE Computer Security Foundations Workshop*, pages 3–15, July 2006.
31. B. Köpf and D. Basin. An information-theoretic model for adaptive side-channel attacks. In *Proc. ACM Conference on Computer and Communications Security*, pages 286–296, Oct. 2007.
32. L. Lamport. Proving the correctness of multiprocess programs. *IEEE Transactions on Software Engineering*, 3(2):125–143, 1977.
33. Z. Manna and A. Pnueli. *The Temporal Logic of Reactive and Concurrent Systems: Specification*. Springer-Verlag, New York, 1992.
34. H. Mantel. Possibilistic definitions of security—an assembly kit. In *Proc. IEEE Computer Security Foundations Workshop*, pages 185–199, July 2000.
35. D. McCullough. Noninterference and the composability of security properties. In *Proc. IEEE Symposium on Security and Privacy*, pages 177–186, Apr. 1988.
36. D. McCullough. A hookup theorem for multilevel security. *Proc. IEEE Transactions on Software Engineering*, 16(6):563–568, June 1990.
37. J. McLean. Proving noninterference and functional correctness using traces. *Journal of Computer Security*, 1(1):37–58, 1992.
38. J. McLean. A general theory of composition for trace sets closed under selective interleaving functions. In *Proc. IEEE Symposium on Security and Privacy*, pages 79–93, Apr. 1994.
39. J. K. Millen. Unwinding forward correctness. In *Proc. IEEE Computer Security Foundations Workshop*, pages 2–10, June 1994.
40. D. Milushev and D. Clarke. Towards incrementalization of holistic hyperproperties. In *Proc. Conference on Principles of Security and Trust*, pages 329–348, 2012.
41. D. Milushev and D. Clarke. Incremental hyperproperty model checking via games. In *Proc. Nordic Conference on Secure IT Systems*, pages 247–262, Oct. 2013.

42. D. V. Milushev. *Reasoning about Hyperproperties*. PhD thesis, Katholieke Universiteit Leuven, June 2013.
43. A. C. Myers. JFlow: Practical mostly-static information flow control. In *Proc. ACM Symposium on Principles of Programming Languages*, pages 228–241, 1999.
44. A. Pnueli. The temporal logic of programs. In *Proc. Foundations of Computer Science*, pages 46–57, Sept. 1977.
45. H. Rogers. *Theory of Recursive Functions and Effective Computability*. MIT Press, Cambridge, Massachusetts, 1987.
46. A. W. Roscoe. CSP and determinism in security modelling. In *Proc. IEEE Symposium on Security and Privacy*, pages 114–127, May 1995.
47. P. Y. A. Ryan and T. Peacock. Opacity—further insights on an information flow property. Technical Report CS-TR-958, Newcastle University, Apr. 2006.
48. A. Sabelfeld and A. C. Myers. Language-based information-flow security. *IEEE Journal on Selected Areas in Communications*, 21(1):5–19, Jan. 2003.
49. A. Sabelfeld and D. Sands. Dimensions and principles of declassification. In *Proc. IEEE Computer Security Foundations Workshop*, pages 255–269, 2005.
50. A. P. Sistla, M. Y. Vardi, and P. Wolper. The complementation problem for Büchi automata with applications to temporal logic. *Theoretical Computer Science*, 49:217–237, 1987.
51. G. Smith. On the foundations of quantitative information flow. In *Proc. Conference on Foundations of Software Science and Computation Structures*, pages 288–302, Mar. 2009.
52. T. Terauchi and A. Aiken. Secure information flow as a safety problem. In *Proc. Static Analysis*, pages 352–367, Sept. 2005.
53. Y. Tsay, Y. Chen, M. Tsai, K. Wu, and W. Chan. GOAL: a graphical tool for manipulating Büchi automata and temporal formulae. In *Proc. Tools and Algorithms for the Construction and Analysis of Systems*, pages 466–471, Mar. 2007.
54. R. van der Meyden. Axioms for knowledge and time in distributed systems with perfect recall. In *Proc. IEEE Symposium on Logic in Computer Science*, pages 448–457, 1993.
55. R. van der Meyden and T. Wilke. Preservation of epistemic properties in security protocol implementations. In *Proc. ACM Conference on Theoretical Aspects of Rationality and Knowledge*, pages 212–221, 2007.
56. R. van der Meyden and C. Zhang. Algorithmic verification of noninterference properties. *Electronic Notes in Theoretical Computer Science*, 168:61–75, Feb. 2007.
57. M. Y. Vardi. An automata-theoretic approach to linear temporal logic. In *Proc. Banff Higher Order Workshop*, pages 238–266, Aug. 1996.
58. M. Y. Vardi and P. Wolper. Reasoning about infinite computations. *Information and Computation*, 115(1):1–37, 1994.
59. P. Wolper. Constructing automata from temporal logic formulas: A tutorial. In *Lectures on Formal Methods and Performance Analysis, First EEF/Euro Summer School on Trends in Computer Science*, pages 261–277. Springer, 2000.
60. H. Yasuoka and T. Terauchi. On bounding problems of quantitative information flow. In *Proc. European Symposium on Research in Computer Security*, pages 357–372, Sept. 2010.
61. S. Zdancewic and A. C. Myers. Observational determinism for concurrent program security. In *Proc. IEEE Computer Security Foundations Workshop*, pages 29–43, June 2003.