INJECTING FINITENESS TO PROVE FINITE LINEAR TEMPORAL LOGIC COMPLETE

ERIC CAMPBELL AND MICHAEL GREENBERG

Abstract. Temporal logics over finite traces are not the same as temporal logics over potentially infinite traces [2, 6, 5]. Roşu first proved completeness for linear temporal logic on finite traces (LTL_f) with a novel coinductive axiom [13]. We offer a different proof, with fewer and more conventional axioms. Our proof is a direct adaptation of Kröger and Merz's proof [10]. The essence of our adaption is that we "inject" finiteness: that is, we alter the proof structure to ensure that models are finite.

§1. Introduction. Temporal logics have proven useful in a remarkable number of applications, in particular reasoning about reactive systems. To accommodate the nonterminating nature of such systems, temporal logics have used a possibly-infinite model of time. For nearly thirty years after Pnueli's seminal work [12], the prevailing wisdom held that proofs about infinite-trace temporal logics were sound for finite models of time. Researchers have recently overturned that conventional wisdom: some formulae are valid only in finite models [2, 6, 5].

Having realized that finite temporal logics differ from (possibly) infinite ones, we may wonder: how do these finite temporal logics behave? What are their model and proof theories like? Can we adapt existing metatheoretical techniques from infinite settings, or must we come up with new ones? Reworking the model theory of temporal logics for finite time is an uncomplicated exercise: the standard model is a (possibly infinite) sequence of valuations on primitive propositions; to consider only finite models, simply restrict the model to finite sequences of valuations. The proof theory is more challenging. In practice, it is sufficient to (a) add an axiom indicating that the end of time eventually comes, (b) add an axiom to say what happens when the end of time arrives, and (c) to relax (or strengthen) axioms from the infinite logic that may not hold in finite settings. For an example of (c), consider LTL_f. It normally holds that the next modality commutes with implication, i.e., $\circ(\phi \Rightarrow \psi) \Leftrightarrow (\circ\,\phi \Rightarrow \circ\,\psi)$; in a finite setting, we must relax the if-and-only-if to merely the left-to-right direction.

Once we settle on a set of axioms, what does a proof of deductive completeness look like? We believe that it is possible to adapt existing techniques for infinite temporal logics to finite ones *directly*. As evidence, we offer a proof of completeness for linear temporal logic over finite traces (LTL_f) with a a conventional structure: we define a graph of *positive-negative pairs* of formulae (PNPs), following Kröger and Merz's presentation [10]. The only change we make to their construction is that when we prove our satisfiability lemma—the core property relating the PNP graph to provability—we "inject" finiteness, adding a formula that guarantees a finite model.

We claim the following contributions:

- Evidence for the claim that the metatheory for infinite temporal logics readily adapts to finite temporal logics by means of *injecting finiteness* (Section 2 situates our work; Section 3 explains our model of finite time).
- A proof of deductive completeness for linear temporal logic on finite traces (LTL_f; Section 4) with fewer axioms than any prior proof [13].

§2. Related work. Pnueli [12] proved his temporal logic programs to be sound and complete over traces of "discrete systems" which may or may not be finite; Lichtenstein et al. [11] extended LTL with past-time operators and allowed more explicitly for the possibility of finite or infinite traces.

Baier and McIlraith were the first to observe that some formulae are only valid in infinite models, and so LTL_f and other 'truncated' finite temporal logics differ from their infinite originals [2]. Roşu [13] offers a translation from LTL_f to LTL that perserves satisfiability of formulae, but makes no claims about the inverse translation. De Giacomo and Vardi showed that satisfiability and validity were PSPACE-complete for these finite logics, relating LTL_f and linear dynamic logic (LDL_f) to other logics (potentially infinite LTL, FO[<], star-free regular expressions, MSO on finite traces) [6]; later, de Giacomo et al. were able to directly characterize when LTL_f and LDL_f formulae are sensitive to infiniteness [5]. De Giacomo and Vardi have also studied the synthesis problem for our logic of interest [7, 8]. Most recently, D'Antoni and Veanes offered a decision procedure for MSO on finite sequences, but without a deductive completeness result [4].

Roşu [13] was the first to show a deductive completeness result for a finite temporal logic: he showed LTL_f is deductively complete by replacing the induction axiom with a *coinduction* axiom coInd: if $\vdash \bullet\phi \Rightarrow \phi$

then $\vdash \Box\phi$.[1] He shows that coInd is equivalent to the combination of the conventional induction axiom Ind (if $\vdash (\phi \Rightarrow \bullet\phi)$ then $\vdash \phi \Rightarrow \Box\phi$) axiom and a finiteness axiom Fin, $\Diamond\bullet\bot$. Roşu proves consistency using "maximally consistent" worlds, i.e., in a greatest fixed-point style.

Our goal is to show that existing, conventional methods for infinite temporal logics suffice for proving that finite temporal logics are deductively complete. For LTL_f, we take the conventional inductive framing, extending Kröger and Merz's axioms with the axiom Fin : $\Diamond\bullet\bot$, i.e., $\Diamond\,\mathsf{end}$ (we call this axiom Finite). Surprisingly, we are able to prove completeness with only six temporal axioms—one fewer than Roşu's seven, though he conjectures his set is minimal. It turns out that some of his axioms are consequences of others—his necessitation axiom $\mathrm{N}_\Box$ can be proved from $\mathrm{N}_\bullet$ and coInd (via Ind; see Lemma 13). Our results for LTL_f show that a smaller axiom set exists. In fact, we could go still smaller: using Roşu's proofs, we can replace Finite and Induction with coInd, for only five axioms! We see our proof as offering a separate contribution, beyond shrinking the number of axioms needed. We follow Kröger and Merz's least fixed-point construction quite closely, adapting their proof from LTL to LTL_f by *injecting finiteness*. Our key idea is that existing techniques for infinite systems readily adapt to finite ones: we can reuse model theory which uses potentially infinite models so long as we can force the theory to work exclusively with finite models.

2.1. Applications. For de Giacaomo and Vardi, LTL_f is useful for AI planning applications [6, 5, 7, 8]. The second author first encountered finite temporal logics when designing Temporal NetKAT [3]. NetKAT is a specification language for network configurations [1] based on Kleene algebra with tests [9]. Temporal NetKAT extends NetKAT with the ability to write and analyze policies using past-time finite linear temporal logic, e.g., a packet may not arrive at the server unless it has previously been at the firewall. Our interest in the deductive completeness of LTL_f comes directly from the Temporal NetKAT work: the completeness result for Temporal NetKAT's equivalence relation relies on deductive completeness for LTL_f.

§3. Modeling finite time. Our logic, LTL_f, uses a finite model of time: *traces*. A trace over a fixed set of propositional variables is a (possibly infinite) sequence $(\eta_1, \eta_2, \dots)$ where η_i is a *valuation*, i.e., a function establishing the truth value ($\mathfrak{t}$ or $\mathfrak{f}$) for each propositional variable. We refer to each valuation as a 'state', with the intuition that each valuation represents a discrete moment in time. Formally:

[1] In Roşu's paper, empty circles mean "weak next" and filled ones mean "next", while we follow Kröger and Merz and do the reverse [10].

Definition 1 (Valuations and Kripke structures). Given a set of variables Var, a *valuation* is a function $\eta : \mathsf{Var} \to \{\mathfrak{t}, \mathfrak{f}\}$. A *Kripke structure* or a *trace* is a finite, non-empty sequence of valuations; we write $\mathsf{K}^n \in \mathsf{Model}_n$ to refer to a model with n valuations, i.e., $\mathsf{K}^n = (\eta_1, \ldots, \eta_n)$.

We particularly emphasize the finiteness of our Kripke structures, writing K^n and explicitly stating the number of valuations as a superscript each time. The number n is not directly accessible in our logic, though the size of models is observable (e.g., the LTL$_f$ formula $\circ\circ\circ\top$ will be satisfiable only in models with four or more steps). Our traces are not only finite, but they are necessarily non-empty—all LTL$_f$ formulae would trivially hold in empty models.

Suppose we have $\mathsf{Var} = \{x, y, z\}$. As a first example, the smallest possible model will be one with only one time step, $\mathsf{K}^1 = \eta_1$, where η_1 is a function from Var to the booleans, i.e., a subset of Var. As a more complex example, consider the following model K^4 with four time steps:

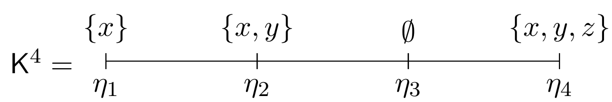

In the first state, x holds but y and z do not (i.e. $\eta_1(x) = \mathfrak{t}$, but $\eta_1(y) = \eta_1(z) = \mathfrak{f}$); then x and y hold; then no propositions hold; and then all primitive propositions hold.

Our logic uses Kripke structures to interpret formulae, defining a function $\mathsf{K}^n_i : \mathsf{Formula} \to \{\mathfrak{t}, \mathfrak{f}\}$. (Put another way: we define a function $\mathsf{interp} : \mathsf{Model}_n \times \{1, \ldots, n\} \times \mathsf{Formula} \to \{\mathfrak{t}, \mathfrak{f}\}$.) We lift this interpretation function to define validity and satisfiability.

Definition 2 (Semantic satisfiability and validity). For an interpretation function $\mathsf{K}^n_i : \mathsf{Formula} \to \{\mathfrak{t}, \mathfrak{f}\}$, we say for $\phi \in \mathsf{Formula}$:

- K^n *models* ϕ iff $\mathsf{K}^n_1(\phi) = \mathfrak{t}$;
- ϕ is *satisfiable* iff $\exists \mathsf{K}^n$ such that K^n models ϕ;
- $\mathsf{K}^n \models \phi$ (pronounced "K^n satisfies ϕ") iff $\forall 1 \leq i \leq n,\ \mathsf{K}^n_i(\phi) = \mathfrak{t}$;
- $\models \phi$ (pronounced "ϕ is valid") iff $\forall \mathsf{K}^n, \mathsf{K}^n \models \phi$; and
- $\mathcal{F} \models \phi$ for $\mathcal{F} \subseteq \mathsf{Formula}$ (pronounced "ϕ is valid under $\mathcal{F}$") iff $\forall \mathsf{K}^n$, if $\forall \psi \in \mathcal{F},\ \mathsf{K}^n \models \psi$ then $\mathsf{K}^n \models \phi$.

§4. LTL$_f$: linear temporal logic on finite traces. Linear temporal logic is a classical logic for reasoning on potentially infinite traces. The syntax of linear temporal logic on finite traces (LTL$_f$) is identical to that of its (potentially) infinite counterpart. We define LTL$_f$ as propositional logic with two temporal operators (Figure 1). The propositional fragment comprises: variables v from some fixed set of propositional variables Var; the false proposition, $\perp$; and implication, $\phi \Rightarrow \psi$. The temporal fragment comprises two operators: the *next modality*, written $\circ\,\phi$; and, *weak until*, written $\phi\ \mathcal{W}\ \psi$.

These core logical operators encode a more conventional looking logic (Figure 1), with the usual logical operators and an enriched set of temporal operators. Of these standard encodings, we remark on two in particular: end, the end of time, and $\bullet\,\phi$, the *weak next* modality. In the usual (potentially infinite) semantics, it is generally the case that $\circ\,\top$ holds, i.e., that the true proposition holds in the next state, i.e., that there *is* a next state. But at the end of time, there is no next state—and so $\circ\,\top$ ought not adhere. In every state *but* the last, we have $\circ\,\top$ as usual. We can therefore define $\mathsf{end} = \neg\circ\top$—if end holds, then we must be at the end of time. It's worth noting here that negation does not generally commute with the next modality[2]; observe that $\neg\circ\top$ holds only at the end of time, but $\circ\neg\top$ holds nowhere. In fact, $\neg\circ\neg\phi$ will hold at the end of time for every possible ϕ, and everywhere else, $\neg$ and $\circ$ commute, i.e. $\neg\circ\neg\phi$ holds if and only if $\circ\,\phi$ holds. Bearing these facts in mind, we define the *weak next* modality as $\bullet\,\phi = \neg\circ\neg\phi$. We say that weak next is *insensitive to the end of time* and $\circ\,\phi$ is *senstitive to the end of time.* To realize these intuitions, we must define our model.

The simple, standard model for LTL is a possibly-infinite *trace*; we will restrict ourselves to finite traces (Definition 1). Given a Kripke structure K^n, we assign a truth value to a proposition ϕ at time step $1 \le i \le n$ with the function $\mathsf{K}^n_i(\phi)$, defined as a fixpoint on formulae. The definitions for K^n_i in the propositional fragment are straightforward implementations of the conventional operations. The definitions for K^n_i in the temporal fragment also assign the usual meanings, being mindful of the end of time. When there is no next state, the formula $\circ\,\phi$ is necessarily false; when there is no next state, the formula $\phi\ \mathcal{W}\ \psi$ degenerates into $\phi \vee \psi$. Why? Suppose we are at the end of time; one of two cases adheres. Either we have ϕ until the end of time (which is now!), or we have ψ and have satisfied the until. We can verify our earlier intuitions about end and $\bullet\,\phi$. Observe that $\mathsf{K}^n_i(\mathsf{end}) = \mathsf{t}$ exactly when $i = n$; similarly, $\mathsf{K}^n_i(\bullet\,\top) = \mathsf{t}$ for all $1 \le i \le n$.

By way of example, consider K^4 from Section 3. We have $\mathsf{K}^4 \models y \Rightarrow x$, because $\mathsf{K}^4_i(y \Rightarrow x) = \mathsf{t}$ for all i, i.e., whenever y holds, so does x. Similarly, K^4 models $x\ \mathcal{W}\ y$ with the k in the existential equal to 2; we have K^4 models $z\ \mathcal{W}\ x$, too, but trivially with $k = 1$. The formula $\Box\, z$ doesn't hold in any state of K^4, but $\mathsf{K}^4 \models \Diamond\, z$. We lift K^n_i to satisfiability and validity in the usual way (Definition 2). We prove a semantic deduction theorem appropriate to our setting: rather than producing a bare implication, deduction produces an implication whose premise is under an 'always' modality.

[2]This is not true in the possibly-infinite semantics, where $\models \neg\circ\phi \Leftrightarrow \circ\neg\phi$

Syntax

$$\begin{array}{llcl} \text{Variables} & v \in \mathsf{Var} & & \\ \text{LTL}_f \text{ formulae} & \phi, \psi \in \text{LTL}_f & ::= & v \mid \bot \mid \phi \Rightarrow \psi \\ & & & \mid \circ\,\phi \mid \phi\,\mathcal{W}\,\psi \end{array}$$

Encodings

$$\begin{array}{rclrcl} \neg\phi & = & \phi \Rightarrow \bot & \top & = & \neg\bot \\ \phi \vee \psi & = & \neg\phi \Rightarrow \psi & \phi \wedge \psi & = & \neg(\neg\phi \vee \neg\psi) \\ \mathsf{end} & = & \neg\circ\top & \bullet\,\phi & = & \neg\circ\neg\phi \\ \Box\,\phi & = & \phi\,\mathcal{W}\,\bot & \Diamond\,\phi & = & \neg\Box\neg\phi \\ \phi\,\mathcal{U}\,\psi & = & \phi\,\mathcal{W}\,\psi \wedge \Diamond\,\psi & & & \end{array}$$

Semantics $\boxed{\mathsf{K}^n_i : \text{LTL}_f \to \{\mathfrak{t}, \mathfrak{f}\}}$

$$\mathsf{K}^n_i(v) = \eta_i(v) \tag{1}$$

$$\mathsf{K}^n_i(\bot) = \mathfrak{f} \tag{2}$$

$$\mathsf{K}^n_i(\phi \Rightarrow \psi) = \begin{cases} \mathfrak{t} & \mathsf{K}^n_i(\phi) = \mathfrak{f} \text{ or } \mathsf{K}^n_i(\psi) = \mathfrak{t} \\ \mathfrak{f} & \text{otherwise} \end{cases} \tag{3}$$

$$\mathsf{K}^n_i(\circ\,\phi) = \begin{cases} \mathsf{K}^n_{i+1}(\phi) & i < n \\ \mathfrak{f} & i = n \end{cases} \tag{4}$$

$$\mathsf{K}^n_i(\phi\,\mathcal{W}\,\psi) = \begin{cases} \mathfrak{t} & \forall i \le j \le n,\ \mathsf{K}^n_j(\phi) = \mathfrak{t} \text{ or} \\ & \exists i \le k \le n,\ \mathsf{K}^n_k(\psi) = \mathfrak{t} \text{ and} \\ & \forall i \le j < k,\ \mathsf{K}^n_j(\phi) = \mathfrak{t} \\ \mathfrak{f} & \text{otherwise} \end{cases} \tag{5}$$

FIGURE 1. LTL_f syntax and semantics

THEOREM 3 (Semantic deduction). $\mathcal{F} \cup \{\phi\} \models \psi$ *iff* $\mathcal{F} \vdash \Box\,\phi \Rightarrow \psi$.

PROOF. We prove each direction separately.

($\Rightarrow$) Suppose $\mathcal{F} \cup \{\phi\} \models \psi$. Let K^n be given such that $\mathsf{K}^n \models \chi$ for all $\chi \in \mathcal{F}$. We will show that $\mathsf{K}^n_i(\Box\,\phi \Rightarrow \psi)$ for all i.

Let an i be given. If $\mathsf{K}^n_i(\Box\,\phi) = \mathfrak{f}$, we are done immediately—so suppose $\mathsf{K}^n_j(\Box\,\phi) = \mathfrak{t}$ for all $j \ge i$. It remains to be seen that $\mathsf{K}^n(\psi) = \mathfrak{t}$ for all $j \ge i$.

Let j be given. We extract a smaller Kripke structure, $\mathsf{K}'^{n-i} = (\eta_i, \ldots, \eta_{n-i})$. By definition, $\mathsf{K}'^{n-i}_k = \mathsf{K}^n_{k-1+i}$ for all $1 \le k \le m-i$. Then, our assumption that $\mathsf{K}^n \models \chi$ for all $\chi \in \mathcal{F} \cup \{\phi\}$ implies

Proof theory $\boxed{\vdash \subseteq 2^{\mathrm{LTL}_f} \times \mathrm{LTL}_f}$

Axioms

all propositional tautologies	TAUT
$\vdash \bullet(\phi \Rightarrow \psi) \Leftrightarrow (\bullet\,\phi \Rightarrow \bullet\,\psi)$	WKNEXTDISTR
$\vdash \mathsf{end} \Rightarrow \neg \circ \phi$	ENDNEXTCONTRA
$\vdash \Diamond\,\mathsf{end}$	FINITE
$\vdash \phi\,\mathcal{W}\,\psi \Leftrightarrow \psi \vee (\phi \wedge \bullet(\phi\,\mathcal{W}\,\psi))$	WKUNTILUNROLL
$\dfrac{\vdash \phi}{\vdash \bullet\,\phi}$	WKNEXTSTEP
$\dfrac{\vdash \phi \Rightarrow \psi \quad \vdash \phi \Rightarrow \bullet\,\phi}{\vdash \phi \Rightarrow \Box\,\psi}$	INDUCTION

Consequences

$\vdash \neg(\circ\top \wedge \circ\bot)$	*Lemma* 6
$\vdash \neg \circ \phi \Leftrightarrow \mathsf{end} \vee \circ\neg\phi$	*Lemma* 7
$\vdash \bullet\,\phi \Leftrightarrow \circ\,\phi \vee \mathsf{end}$	*Lemma* 8
$\vdash \neg\mathsf{end} \wedge \bullet\neg\phi \Rightarrow \neg \bullet \phi$	*Lemma* 9
$\vdash \bullet(\phi \wedge \psi) \Leftrightarrow \bullet\,\phi \wedge \bullet\,\psi$	*Lemma* 10
$\vdash \neg \bullet \phi \Rightarrow \bullet\neg\phi$	*Lemma* 11
$\vdash \Box\,\phi \Leftrightarrow \phi \wedge \bullet\Box\,\phi$	*Lemma* 12
$\dfrac{\vdash \phi}{\vdash \Box\,\phi}$	*Lemma* 13

$\mathcal{F} \vdash \phi$ iff assuming $\vdash \psi$ for each $\psi \in \mathcal{F}$ we have $\vdash \phi$

FIGURE 2. LTL_f proof theory

$\mathsf{K}'^{n-i} \models \chi$ for all $\chi \in \mathcal{F} \cup \phi$. We already assumed that $\mathcal{F} \vdash \{\phi\} \models \psi$, so we can conclude that $\mathsf{K}'^{n-i} \models \psi$.

Hence, K'^{n-i} assigns $\mathfrak{t}$ to $\Box\,\phi \Rightarrow \psi$, and so $\mathsf{K}^n_i(\Box\,\phi \Rightarrow \psi) = \mathfrak{t}$.

($\Leftarrow$) Suppose $\mathcal{F} \models \Box\,\phi \Rightarrow \psi$. Let K^n be given such that $\mathsf{K}^n \models \chi$ for all $\chi \in \mathcal{F} \cup \{\phi\}$. We must show that $\mathsf{K}^n_i(\psi)$ for all i. Since $\mathsf{K}^n \models \chi$ for all $\chi \in \mathcal{F}$, then $\mathsf{K}^n_i(\Box\,\phi \Rightarrow \psi) = \mathfrak{t}$ by assumption. Furthermore, we know that $\mathsf{K}^n \models \phi$, i.e., $\mathsf{K}^n_j(\phi) = \mathfrak{t}$ for all $i \leq j \leq n$. Then by definition of the interpretation function, $\mathsf{K}^n_i(\Box\,\phi) = \mathfrak{t}$, so the implication in the assumption cannot hold vacuously: conclude $\mathsf{K}^n_i(\psi) = \mathfrak{t}$ as desired. $\dashv$

For our axioms (Figure 2), we adapt Kröger and Merz's presentation [10]. Two axioms are new: FINITE says that time will eventually end; ENDNEXTCONTRA says that at the end of time, there is no next state. Other axioms are lightly adapted: wherever one would ordinarily use the (strong) next modality, $\circ\,\phi$, we instead use weak next, $\bullet\,\phi$. Changing these axioms to use strong next would be unsound in finite models. We can, however, characterize the relationship between the next modality, negation, and the end of time ("Consequences" in Figure 2 and Section 4.2).

Roşu proves completeness with a slightly different set of axioms, replacing FINITE and INDUCTION with a single *coinduction* axiom he calls COIND:

$$\frac{\vdash \bullet\,\phi \Rightarrow \phi}{\vdash \Box\,\phi}$$

He proves that COIND is equivalent to the conjunction of FINITE and INDUCTION, so it does not particularly matter which axioms we choose. In order to emphasize how little must change to make our logic finite, we keep our presentation as close to Kröger and Merz's as possible[3].

4.1. Soundness. Proving that our axioms are sound is, as usual, relatively straightforward: we verify each axiom in turn.

THEOREM 4 (LTL_f soundness). *If* $\vdash \phi$ *then* $\models \phi$.

PROOF. By induction on the derivation of $\vdash \phi$. Our proof refers to the various cases in the definition of the model (Figure 1).

(TAUT) As for propositional logic.

(WKNEXTDISTR) We have $\vdash \bullet(\phi \Rightarrow \psi) \Leftrightarrow (\bullet\,\phi \Rightarrow \bullet\,\psi)$. To show validity in the model, let K^n be given. We show that K^n assigns true to the left-hand side iff it assigns true to the right-hand side.

$$\begin{array}{ll}
\multicolumn{2}{l}{\mathsf{K}^n \models \bullet(\phi \Rightarrow \psi)} \\
\text{iff} & \forall 1 \le i \le n,\ \mathsf{K}^n_i(\bullet(\phi \Rightarrow \psi)) = \mathfrak{t} \\
\text{iff} & \forall 1 \le i \le n-1,\ \mathsf{K}^n_{i+1}(\phi \Rightarrow \psi) = \mathfrak{t} \\
\text{iff} & \forall 1 \le i \le n-1,\ \mathsf{K}^n_{i+1}(\phi) = \mathfrak{f} \text{ or } \mathsf{K}^n_{i+1}(\psi) = \mathfrak{t} \\
\text{iff} & \forall 1 \le i \le n,\ \mathsf{K}^n_i(\bullet\,\phi) = \mathfrak{f} \text{ or } \mathsf{K}^n_i(\bullet\,\psi) = \mathfrak{t} \\
 & \qquad\qquad\qquad \text{where } \mathsf{K}^n_n(\bullet\,\psi) = \mathfrak{t} \text{ trivially} \\
\text{iff} & \forall 1 \le i \le n,\ \mathsf{K}^n_i(\bullet\,\phi \Rightarrow \bullet\,\psi) \\
\text{iff} & \mathsf{K}^n \models \bullet\,\phi \Rightarrow \bullet\,\psi
\end{array}$$

By unfolding the encodings of logical operators, we can derive that $\mathsf{K}^n \models \bullet(\phi \Rightarrow \psi) \Leftrightarrow (\bullet\,\phi \Rightarrow \bullet\,\psi)$.

[3] They use a less-expressive syntax, omitting $\mathcal{W}$ and $\mathcal{U}$. We extend their methodology to include these operators.

(EndNextContra) We have $\vdash \mathsf{end} \Rightarrow \neg \circ \phi$; let K^n be given to show $\mathsf{K}^n \models \mathsf{end} \Rightarrow \neg \circ \phi$, i.e., that K^n assigns $\mathfrak{t}$ to the given formula at each $1 \leq i \leq n$. Let i be given.

We have $\mathsf{K}^n_i(\mathsf{end} \Rightarrow \neg \circ \phi)$. There are two cases: $i < n$ and $i = n$.

- ($i < n$) We have $\mathsf{K}^n_i(\mathsf{end}) = \mathsf{K}^n_i(\neg \circ \top)$. Since $i < n$, then $\mathsf{K}^n_i(\circ \top) = \mathsf{K}^n_{i+1}(\top) = \mathsf{K}^n_{i+1}(\neg \bot) = \mathfrak{t}$, and so $\mathsf{K}^n_i(\mathsf{end}) = \mathfrak{f}$. The implication is thus vacuous: $\mathsf{K}^n_i(\mathsf{end} \Rightarrow \neg \circ \phi) = \mathfrak{t}$, by the first clause of case (3).
- ($i = n$) Since $i = n$, we have $\mathsf{K}^n_n(\circ \phi) = \mathfrak{f}$, so $\mathsf{K}^n_n(\neg \circ \phi) = \mathfrak{t}$. Therefore $\mathsf{K}^n_i(\mathsf{end} \Rightarrow \neg \circ \phi) = \mathfrak{t}$, by the second clause of case (3).

(Finite) We have $\vdash \Diamond \mathsf{end}$; let K^n be given to show $\mathsf{K}^n \models \Diamond \mathsf{end}$, i.e., that for all $1 \leq i \leq n$, we have $\mathsf{K}^n_i(\Diamond \mathsf{end}) = \mathfrak{t}$. Let i be given.

Unfolding our encodings, we must show that:

$$\mathsf{K}^n_i(\Diamond \mathsf{end}) = \mathsf{K}^n_i(\neg \Box \neg \mathsf{end}) = \mathsf{K}^n_i(\neg(\neg \mathsf{end}\ \mathcal{W}\ \bot)) = \mathfrak{t}$$

By cases (3) and (2), it will suffice to show that $\mathsf{K}^n_i(\neg \mathsf{end}\ \mathcal{W}\ \bot) = \mathfrak{f}$. By case (5), there are two ways for the weak-until to be assigned true; we will show that neither adheres. First, observe that $\mathsf{K}^n_k(\bot) = \mathfrak{f}$ for all $1 \leq k \leq n$, so there is no k to satisfy the second clause of case (5). Next, observe that when $j = n$, where $\mathsf{K}^n_j(\neg \mathsf{end}) = \mathfrak{f}$, so the first clause of case (5) cannot be satisfied. Since neither case holds, we find $\mathsf{K}^n_i(\neg \mathsf{end}\ \mathcal{W}\ \bot) = \mathfrak{f}$.

(WkUntilUnroll) We have $\vdash \phi\ \mathcal{W}\ \psi \Leftrightarrow \psi \vee (\phi \wedge \bullet(\phi\ \mathcal{W}\ \psi))$; let K^n be given to show that $\mathsf{K}^n \models \phi\ \mathcal{W}\ \psi \Leftrightarrow \psi \vee (\phi \wedge \bullet(\phi\ \mathcal{W}\ \psi))$, i.e., that for all $1 \leq i \leq n$, the left-hand side of our formula is assigned true by K^n_i iff the right-hand side is. Let an i be given; we prove each side independently.

- ($\Rightarrow$) We have $\mathsf{K}^n_i(\phi\ \mathcal{W}\ \psi) = \mathfrak{t}$ iff $\forall i \leq j \leq n,\ \mathsf{K}^n_j(\phi) = \mathfrak{t}$ or $\exists i \leq k \leq n,\ \mathsf{K}^n_k(\psi) = \mathfrak{t}$ and $\forall i \leq j < k,\ \mathsf{K}^n_j(\phi) = \mathfrak{t}$. We go by cases.
  - (ϕ always holds) In this case, $\mathsf{K}^n_i(\phi) = \mathfrak{t}$ and $\mathsf{K}^n_{i+1}(\phi\ \mathcal{W}\ \psi) = \mathfrak{t}$, so $\mathsf{K}^n_i(\psi \vee (\phi \wedge \bullet(\phi\ \mathcal{W}\ \psi))) = \mathfrak{t}$.
  - (ϕ holds until ψ eventually holds) If $k = i$, then $\mathsf{K}^n_i(\psi) = \mathfrak{t}$ implies $\mathsf{K}^n_i(\psi \vee (\phi \wedge \bullet(\phi\ \mathcal{W}\ \psi))) = \mathfrak{t}$.
    If not, then $\mathsf{K}^n_{i+1}(\phi\ \mathcal{W}\ \psi)$ holds by the second clause of (5), using our same k. Therefore, $\mathsf{K}^n_i(\bullet(\phi\ \mathcal{W}\ \psi)) = \mathfrak{t}$ along with $\mathsf{K}^n_i(\phi) = \mathfrak{t}$ (since j can be i), so $\mathsf{K}^n_i(\psi \vee (\phi \wedge \bullet(\phi\ \mathcal{W}\ \psi))) = \mathfrak{t}$.
- ($\Leftarrow$) We have $\mathsf{K}^n_i(\psi \vee (\phi \wedge \bullet(\phi\ \mathcal{W}\ \psi))) = \mathfrak{t}$; we must show $\mathsf{K}^n_i(\phi\ \mathcal{W}\ \psi) = \mathfrak{t}$. We go by cases on which side of the disjunction holds.
  - ($\mathsf{K}^n_i(\psi) = \mathfrak{t}$) In this case, $k = i$ witnesses the second clause of case (5) with $k = i$.
  - ($\mathsf{K}^n_i(\phi) = \mathsf{K}^n_i(\bullet(\phi\ \mathcal{W}\ \psi)) = \mathfrak{t}$) If $i = n$, we are done by the first clause of case (5). because $\forall i + 1 \leq j \leq n,\ \mathsf{K}^n_j(\phi) = \mathfrak{t}$, then $\mathsf{K}^n_i(\phi) = \mathfrak{t}$ completes first clause of case (5). Otherwise,

$\mathsf{K}^n_{i+1}(\phi\,\mathcal{W}\,\psi) = \mathfrak{t}$, because there is some $i+1 \leq k \leq n$ such that $\mathsf{K}^n_k(\psi) = \mathfrak{t}$ and $\forall i+1 \leq j < k,\ \mathsf{K}^n_j(\phi) = \mathfrak{t}$. Since we also have $\mathsf{K}^n_i(\phi) = \mathfrak{t}$, k witnesses the the second clause of case (5).

(WkNextStep) We have $\vdash \bullet\,\phi$; as our IH on $\vdash \phi$, we have $\models \phi$, i.e., $\forall \mathsf{K}^n \forall 1 \leq i \leq n, \mathsf{K}^n_i(\phi) = \mathfrak{t}$. Let a K^n be given to show $\mathsf{K}^n \models \bullet\,\phi$, i.e., $\forall 1 \leq i \leq n,\ \mathsf{K}^n_i(\bullet\,\phi) = \mathfrak{t}$. Let i be given; we go by cases on $i = n$.

$(i = n)$ Since $\mathsf{K}^n_n(\circ\,\neg\phi) = \mathfrak{f}$, conclude $\mathsf{K}^n_n(\bullet\,\phi) = \mathsf{K}^n_n(\neg\circ\neg\phi) = \mathfrak{t}$.

$(i < n)$ By definition (4), $\mathsf{K}^n_i(\circ\,\neg\phi) = \mathsf{K}^n_{i+1}(\neg\phi)$. By the IH, we know that $\mathsf{K}^n_{i+1}(\phi) = \mathfrak{t}$, so $\mathsf{K}^n_{i+1}(\neg\phi) = \mathfrak{f}$, and $\mathsf{K}^n_i(\circ\,\neg\phi) = \mathfrak{f}$. Conclude that $\mathsf{K}^n_i(\bullet\,\phi) = \mathsf{K}^n_i(\neg\circ\neg\phi) = \mathfrak{t}$.

(Induction) We have $\vdash \phi \Rightarrow \Box\,\psi$; as our IHs, we have (1) $\models \phi \Rightarrow \psi$ and (2) $\models \phi \Rightarrow \bullet\,\phi$, i.e., every Kripke structure K^n assigns $\mathfrak{t}$ to those formulae at every index. Let K^n be given to show that $\forall 1 \leq i \leq n,\ \mathsf{K}^n_i(\phi \Rightarrow \Box\,\psi) = \mathfrak{t}$. If $\mathsf{K}^n_i(\phi) = \mathfrak{f}$, the implication vacuously holds; instead consider the case where $\mathsf{K}^n_i(\phi) = \mathfrak{t}$. Let $k = n - i$; we go by induction on k to show $\mathsf{K}^n_i(\Box\,\psi) = \mathfrak{t}$.

$(k = 0)$ Here $n = i$. We have $\mathsf{K}^n_n(\phi) = \mathfrak{t}$; by outer IH (1), it must be that $\mathsf{K}^n_n(\psi) = \mathfrak{t}$, which means that $\mathsf{K}^n_n(\Box\,\psi) = \mathsf{K}^n_n(\psi\,\mathcal{W}\,\bot) = \mathfrak{t}$ by the first clause of case (5).

$(k = k' + 1)$ Here $i < n$. We must show that $\mathsf{K}^n_i(\Box\,\psi) = \mathfrak{t}$.
Since $\mathsf{K}^n_i(\phi) = \mathfrak{t}$, the outer IHs immediately give $\mathsf{K}^n_i(\psi) =_{\text{IH (1)}} \mathfrak{t}$ (outer IH on $\vdash \phi \Rightarrow \psi$) and $\mathsf{K}^n_i(\bullet\,\phi) = \mathfrak{t}$ (outer IH on $\vdash \phi \Rightarrow \bullet\,\phi$). Futhermore, we have $i < n$, so unfolding the definition of $\bullet\,\phi$ gives $\mathsf{K}^n_{i+1}(\neg\phi) = \mathfrak{f}$, or equivalently $\mathsf{K}^n_{i+1}(\phi) = \mathfrak{t}$. Conclude $\mathsf{K}^n_{i+1}(\Box\,\psi) = \mathsf{K}^n_{i+1}(\psi\,\mathcal{W}\,\bot) = \mathfrak{t}$ by the inner IH.
Since $\mathsf{K}^n_j \models \neg\bot$, the above conclusion that $\mathsf{K}^n_{i+1}(\psi\,\mathcal{W}\,\bot) = \mathfrak{t}$ must hold because $\forall i+1 \leq j \leq n,\ \mathsf{K}^n_j(\psi) = \mathfrak{t}$ (as per case (5)). Since $\mathsf{K}^n_i(\psi) = \mathfrak{t}$ as well, we can see that $\forall i \leq j \leq n\ \mathsf{K}^n_j(\psi) = \mathfrak{t}$; conclude that $\mathsf{K}^n_i(\Box\,\psi) = \mathsf{K}^n_i(\psi\,\mathcal{W}\,\bot) = \mathfrak{t}$. ⊣

We can also prove a deduction theorem for our proof theory analogous to Theorem 3.

Theorem 5 (Deduction). *$\mathcal{F} \cup \{\phi\} \vdash \psi$ iff $\mathcal{F} \vdash \Box\,\phi \Rightarrow \psi$.*

Proof. From left to right, by induction on the derivation:

(ψ an axiom or $\psi \in \mathcal{F}$) We have $\mathcal{F} \vdash \Box\,\phi \Rightarrow \psi$ by Taut.

(WkNextStep) We have $\psi = \bullet\,\chi$. WkNextStep concludes $\bullet\,\chi$ from $\mathcal{F} \cup \{\phi\} \vdash \chi$, which gives the IH of $\mathcal{F} \vdash \Box\,\phi \Rightarrow \chi$. Applying WkNextStep to the IH, we have $\mathcal{F} \vdash \bullet(\Box\,\phi \Rightarrow \chi)$; by WkNextDistr, we have $\mathcal{F} \vdash \bullet\,\Box\,\phi \Rightarrow \bullet\,\chi$. By Lemma 12, we know that $\vdash \Box\,\phi \Rightarrow \phi \wedge \bullet\,\Box\,\phi$, so by Taut we have $\mathcal{F} \vdash \Box\,\phi \Rightarrow \psi$.

(Induction) We have $\psi = \chi \Rightarrow \Box\,\rho$, with $\mathcal{F} \cup \{\phi\} \vdash \chi \Rightarrow \rho$ and $F \cup \{\phi\} \vdash \chi \Rightarrow \bullet\,\chi$. By the IH, we know that $\mathcal{F} \vdash \Box\,\phi \Rightarrow \chi \Rightarrow \rho$ and $\mathcal{F} \vdash \Box\,\phi \Rightarrow \chi \Rightarrow \bullet\,\chi$. By Taut, we have $\mathcal{F} \vdash \Box\,\phi \wedge \chi \Rightarrow \rho$ and

$\mathcal{F} \vdash \Box\,\phi \wedge \chi \Rightarrow \bullet\,\chi$, so by INDUCTION we have $\mathcal{F} \vdash \Box\,\phi \wedge \chi \Rightarrow \Box\,\rho$. By TAUT, we find $\mathcal{F} \vdash \Box\,\phi \Rightarrow \chi \Rightarrow \Box\,\rho$.

From right to left, we have $\mathcal{F} \vdash \Box\,\phi \Rightarrow \psi$ and must show $\mathcal{F} \cup \{\phi\} \vdash \psi$. By TAUT, we have $\mathcal{F} \cup \{\phi\} \vdash \Box\,\phi \Rightarrow \psi$. We must prove that $\mathcal{F} \cup \{\phi\} \vdash \Box\,\phi$, which will give $\mathcal{F} \cup \{\phi\} \vdash \psi$ via TAUT.

By WKNEXTSTEP and TAUT, we know that $\mathcal{F} \cup \{\phi\} \vdash \phi \Rightarrow \bullet\,\phi$. We therefore have by induction that $\mathcal{F} \cup \{\phi\} \vdash \phi \Rightarrow \Box\,\phi$. By TAUT, we can conclude $\mathcal{F} \cup \{\phi\} \vdash \phi$ and subsequently $\mathcal{F} \cup \{\phi\} \vdash \Box\,\phi$. ⊣

4.2. Consequences. Before proceeding to the proof of completeness, we prove a variety of properties in LTL$_f$ that will be necessary for the proof: characterizations of the modality (Lemma 6, 8, and 12) and distributivity over connectives (Lemmas 7, 9, 10, and 11). We also derive Roşu's necessitation axiom, $\mathrm{N}_\Box$ (Lemma 13).

Together Lemma 9 and Lemma 11 completely characterize the relationship between the weak next modality and negation: the former can pull a negation out of a weak next modality when not at the end; the latter can push one in whether or not the end has arrived. The latter lemma comes later in our development to account for dependencies in its proof.

LEMMA 6 (Modal consistency). $\vdash \neg(\circ\,\top \wedge \circ\,\bot)$

PROOF. Suppose for a contradiction that $\vdash \circ\,\top \wedge \circ\,\bot$. We have $\vdash \top$ by TAUT, so $\vdash \bullet\,\top$ by WKNEXTSTEP. But $\bullet\,\top$ desugars to $\neg \circ \neg\top$, i.e., $\neg \circ \bot$—a contradiction. ⊣

LEMMA 7 (Negation of next). $\vdash \neg \circ \phi \Leftrightarrow \mathsf{end} \vee \circ \neg\phi$

PROOF. We prove each direction separately.

($\Rightarrow$) We have $\vdash \mathsf{end} \vee \neg\mathsf{end}$ by the law of the excluded middle (TAUT). If end holds, we are done. Otherwise, suppose $\vdash \neg \circ \phi \wedge \neg\mathsf{end}$; we must show $\vdash \circ\neg\phi$. By resugaring and TAUT, we have $\vdash \bullet\neg\phi$; by WKNEXTDISTR, we have $\vdash \neg \bullet \phi$; by desugaring, we have $\vdash \circ\neg\phi$.

($\Leftarrow$) By the law of the excluded middle, we have $\vdash \mathsf{end} \vee \neg\mathsf{end}$ (TAUT). If end holds, then we have $\vdash \neg \circ \phi$ by ENDNEXTCONTRA immediately.

So we have $\vdash \neg\mathsf{end} \wedge \circ \neg\phi$ and we must show $\vdash \neg \circ \phi$. By resugaring and TAUT, we have $\vdash \neg \bullet \phi$; by WKNEXTDISTR, we have $\vdash \bullet \neg\phi$. By desugaring and TAUT, we have $\vdash \neg \circ \phi$ as desired. ⊣

LEMMA 8 (Weak next/next equivalence). $\vdash \bullet\,\phi \Leftrightarrow \circ\,\phi \vee \mathsf{end}$

PROOF. We prove each direction separately.

($\Rightarrow$) Suppose $\vdash \bullet\,\phi$. $\bullet\,\phi$ desugars to $\neg \circ \neg\phi$. By Lemma 7, we have $\mathsf{end} \vee \circ \neg\neg\phi$. If end holds, we are done immediately by TAUT. Otherwise, we have $\circ \neg\neg\phi$, which gives us $\circ\,\phi$ by TAUT, as well.

($\Leftarrow$) Suppose $\vdash \circ\,\phi \vee \mathsf{end}$. By the law of the excluded middle, we have $\mathsf{end} \vee \neg\mathsf{end}$ (TAUT). If end holds, then we are done immediately by ENDNEXTCONTRA.

If $\circ\,\phi \wedge \neg\mathsf{end}$ holds, then we can show $\vdash \bullet\,\phi$ by desugaring to $\circ\,\neg\phi \Rightarrow \bot$. Suppose, for a contradiction, that $\circ\,\neg\phi$. We have $\circ\,\phi$ and $\circ\,\neg\phi$, so $\circ\,\bot$. But from $\neg\mathsf{end}$, we have $\circ\,\top$—and by Lemma 6 we have a contradiction. ⊣

Lemma 9 (Weak next negation before the end).

$$\vdash \neg\mathsf{end} \wedge \bullet\,\neg\phi \Rightarrow \neg\bullet\,\phi$$

Proof. We have $\neg\mathsf{end} \wedge \bullet\,\neg\phi$. By unrolling syntax, we have $\neg\mathsf{end} \wedge \neg\circ\neg\neg\phi$. By Taut, we have $\neg\mathsf{end} \wedge \neg\circ\,\phi$. By Lemma 8, $\circ\,\phi \Rightarrow \bullet\,\phi$, so we have $\neg\bullet\,\phi$. ⊣

Lemma 10 (Weak next distributes over conjunction).

$$\vdash \bullet(\phi \wedge \psi) \Leftrightarrow \bullet\,\phi \wedge \bullet\,\psi.$$

Proof. From $\bullet(\phi \Rightarrow \psi)$, we have $\bullet(\neg(\phi \Rightarrow \neg\psi))$ by Taut. By LEM, we have $\mathsf{end} \vee \neg\mathsf{end}$; by the definition of $\bullet$, we can refactor our formula into $\bullet(\neg(\phi \Rightarrow \neg\psi)) \wedge \mathsf{end} \vee \neg\circ(\phi \Rightarrow \neg\psi) \wedge \neg\mathsf{end}$, also by Taut. By EndNextContra, we have $\mathsf{end} \Rightarrow \neg\circ\neg\neg(\phi \Rightarrow \neg\psi)$, so we can eliminate the left-hand disjunct by Taut, to find $\mathsf{end} \vee \neg\circ(\phi \Rightarrow \neg\psi) \wedge \neg\mathsf{end}$. By Lemma 8, we can weaken our next modality to find $\mathsf{end} \vee \neg\circ(\phi \Rightarrow \neg\psi) \wedge \neg\mathsf{end}$. By WkNextDistr, we distribute the modality across the implication and we have $\mathsf{end} \vee \neg(\bullet\,\phi \Rightarrow \bullet\,\neg\psi) \wedge \neg\mathsf{end}$. Using Lemma 9, we can change $\bullet\,\neg\psi$ to $\neg\bullet\,\psi$, because we have $\neg\mathsf{end}$ on that branch; we now have $\mathsf{end} \vee \neg(\bullet\,\phi \Rightarrow \neg\bullet\,\psi) \wedge \neg\mathsf{end}$. By Taut, we can rearrange the implication to find $\mathsf{end} \vee (\bullet\,\phi \wedge \bullet\,\psi) \wedge \neg\mathsf{end}$. By EndNextContra, we can introduce $\bullet\,\phi$ and $\bullet\,\psi$ on the left-hand disjunct, to find $\bullet\,\phi \wedge \bullet\,\psi \wedge \mathsf{end} \vee \bullet\,\phi \wedge \bullet\,\psi \wedge \neg\mathsf{end}$. Finally, Taut allows us to rearrange our term to $\bullet\,\phi \wedge \bullet\,\psi \wedge (\mathsf{end} \vee \neg\mathsf{end})$, where the rightmost conjunct falls out and we find $\bullet\,\phi \wedge \bullet\,\psi$. ⊣

Lemma 11 (Weak next negation). $\vdash \neg\bullet\,\phi \Rightarrow \bullet\,\neg\phi$

Proof. By definition, $\neg\bullet\,\phi$ is equivalent to $\neg\neg\circ\neg\phi$. By Taut, we eliminate the double negation to find $\circ\,\neg\phi$. By Lemma 7 from right to left, we have $\neg\circ\,\phi$. By Taut , we reintroduce an inner double negation, to find $\neg\circ\neg\neg\phi$. By definition, we have $\bullet\,\neg\phi$. ⊣

Lemma 12 (Always unrolling). $\vdash \Box\,\phi \Leftrightarrow \phi \wedge \bullet\Box\,\phi$

Proof. Desugaring, we must show $\vdash (\phi\,\mathcal{W}\,\bot) \Leftrightarrow \phi \wedge \bullet(\phi\,\mathcal{W}\,\bot)$. By WkUntilUnroll, $\vdash (\phi\,\mathcal{W}\,\bot) \Leftrightarrow \bot \vee (\phi \wedge \bullet(\phi\,\mathcal{W}\,\bot))$. By Taut we can eliminate the $\bot$ case of the disjunction on the right. ⊣

Lemma 13 (Necessitation). *If* $\vdash \phi$ *then* $\vdash \Box\,\phi$.

Proof. We apply Induction with $\phi = \top$ and $\psi = \phi$ to show that $\vdash \top \Rightarrow \Box\,\phi$, i.e., $\vdash \Box\,\phi$ by Taut.

We must prove both premises: $\vdash \top \Rightarrow \phi$ and $\vdash \top \Rightarrow \bullet\,\top$.

Since we've assumed $\vdash \phi$, we have the first premise by Taut.

By Taut and WkNextStep, we have $\vdash \bullet\,\top$, and so $\vdash \top \Rightarrow \bullet\,\top$ by Taut. ⊣

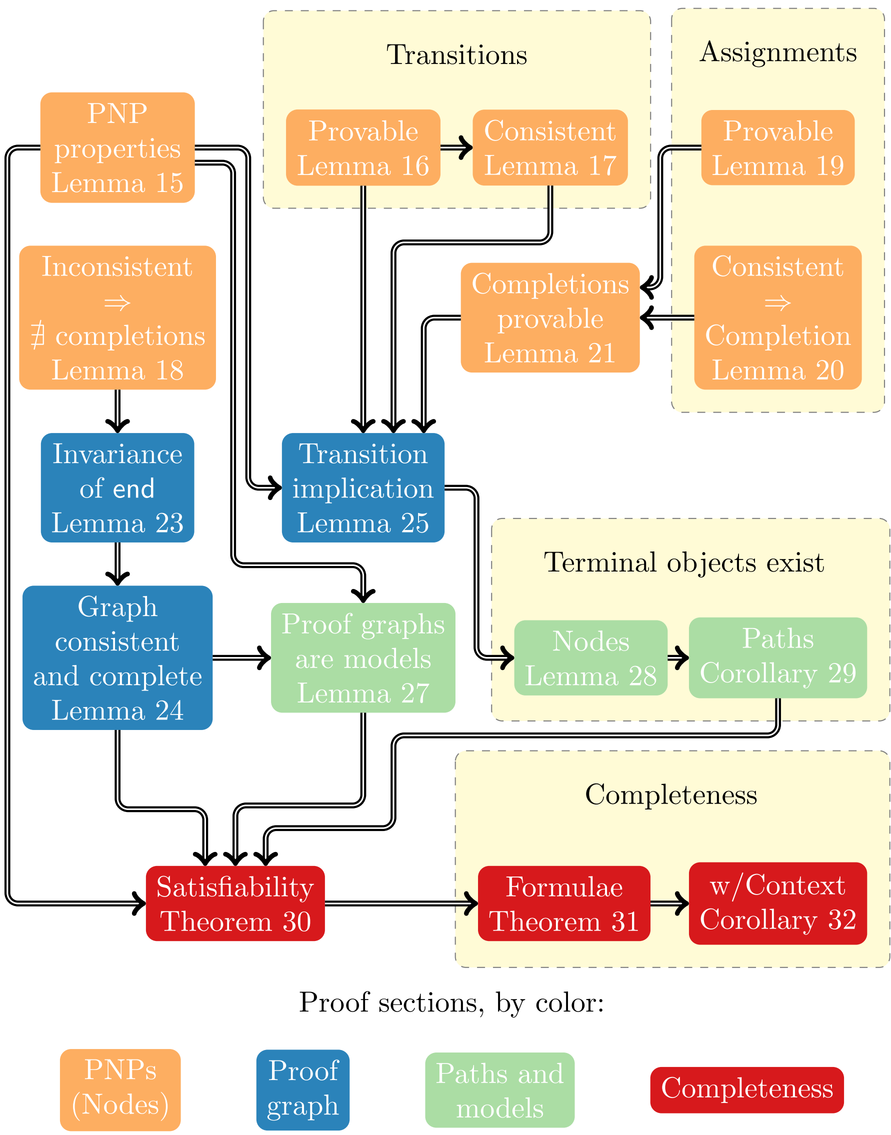

FIGURE 3. Completeness for LTL_f

4.3. Completeness. To show deductive completeness for LTL_f, we must find that if $\models \phi$ then $\vdash \phi$. To do so we will construct a graph that does two things at once: first, paths from the root of the graph to a terminal state correspond to Kripke structures which ϕ satisfies; second, consistency properties in the graph relate to the provability of the underlying formula ϕ.

Our construction follows the standard least-fixed point approach found in Kröger and Merz's book [10]: we constructs a graph whose nodes assign truth values to each subformula of our formula of interest, ϕ, by putting each subformula in either the true, "positive" set or in the false, "negative" set. Our completeness proof ultimately constructs a graph with the assumption that $\nvdash \neg\neg\phi$, showing that the negated graph has no models—the law of the excluded middle yields $\vdash \phi$ (since LTL$_f$'s propositional core is classical). Our proof itself is classical, using the law of the excluded middle to define the proof graph (comps) and prove some of its properties (Lemma 27).

What about the 'f' in LTL$_f$? Nothing described so far differs in any way from the Henkin-Hasenjaeger graph approach used by Kröger and Merz [10]. Kröger and Merz's graphs were always finite, but their notion of satisfying paths forces paths to be infinite. We restrict our attention to terminating paths: paths where not only is our formula of interest satisfied, but so is $\Diamond\,\mathsf{end}$. To ensure such paths exist, we *inject* $\Diamond\,\mathsf{end}$ when we create the graph.

The proof follows the following structure (Figure 3): we define the nodes of the graph (Definition 14); we define the edge relation on the graph (Figure 4) and show that it maps appropriately to time steps in the proof theory (Lemma 21 finds a consistent successor; Lemma 16 shows the successor is a state in our graph); we show that the graph structure results in a finite structure with appropriate consistency properties (Lemma 25); we define which paths in the graph represent our Kripke structure of interest (Lemma 27 shows that our graph's transitions correspond to the semantics; Lemma 29 guarantees that we have appropriate finite models). The final proof comes in two parts: we show that consistent graphs correspond to satisfiable formulae (Theorem 30), which is enough to show completeness (Theorem 31).

Definition 14 (PNP). A *positive-negative pair* (PNP) $\mathcal{P}$ is a pair of finite sets of formulae $(\mathsf{pos}(\mathcal{P}),\ \mathsf{neg}(\mathcal{P}))$. We refer to the collected formulas of $\mathcal{P}$ as $\mathcal{F}_{\mathcal{P}} = \mathsf{pos}(\mathcal{P}) \cup \mathsf{neg}(\mathcal{P})$; we call the set of all PNPs PNP.

We write the *literal interpretation* of a PNP $\mathcal{P}$ as:

$$\widehat{\mathcal{P}} = \bigwedge_{\phi \in \mathsf{pos}(\mathcal{P})} \phi \wedge \bigwedge_{\psi \in \mathsf{neg}(\mathcal{P})} \neg\psi.$$

We say $\mathcal{P}$ is *inconsistent* if $\vdash \neg\widehat{\mathcal{P}}$; conversely, $\mathcal{P}$ is *consistent* when it is not the case that $\vdash \neg\widehat{\mathcal{P}}$.

Positive-negative pairs will form the states of our proof graph, where each state will ultimately be a collection of formulae that hold (or not) in a given moment in time. Before we can even begin constructing the graph, we show that they adequately characterize a moment in time: that is, they are without contradiction, can be 'saturated' with all of the formulae of interest, and respect the general rules of our logic. Readers may be familiar with 'atoms', but PNPs are themselves not atoms; *complete* PNPs are more or less atoms (Figure 4).

LEMMA 15 (PNP properties). *For all consistent PNPs* $\mathcal{P}$*:*

1. $\mathsf{pos}(\mathcal{P}) \cap \mathsf{neg}(\mathcal{P}) = \emptyset$*;*
2. *For all* ϕ*, either* $(\mathsf{pos}(\mathcal{P}) \cup \{\phi\},\ \mathsf{neg}(P))$ *or* $(\mathsf{pos}(\mathcal{P}),\ \mathsf{neg}(P) \cup \{\phi\})$ *is consistent;*
3. $\bot \notin \mathsf{pos}(\mathcal{P})$*;*
4. *if* $\{\phi, \psi, \phi \Rightarrow \psi\} \subseteq \mathcal{F}_{\mathcal{P}}$*, then* $\phi \Rightarrow \psi \in \mathsf{pos}(\mathcal{P})$ *iff* $\phi \in \mathsf{neg}(\mathcal{P})$ *or* $\psi \in \mathsf{pos}(P)$*;*
5. *if* $\vdash \phi \Rightarrow \psi$ *and* $\phi \in \mathsf{pos}(\mathcal{P})$ *and* $\psi \in \mathcal{F}_{\mathcal{P}}$*, then* $\psi \in \mathsf{pos}(\mathcal{P})$*.*

PROOF. Let a given PNP $\mathcal{P}$ be consistent. We show each case by reasoning based on whether each formula is assigned to the positive or the negative set of $\mathcal{P}$, deriving contradictions as appropriate.

1. Suppose for a contradiction that $\phi \in \mathsf{pos}(\mathcal{P}) \cap \mathsf{neg}(\mathcal{P})$. We have $\vdash \neg(\phi \wedge \neg\phi)$ by TAUT, but $\vdash \widehat{\mathcal{P}} \Rightarrow \phi \wedge \neg\phi$, and so $\vdash \neg\widehat{\mathcal{P}}$ by TAUT—making $\mathcal{P}$ inconsistent, a contradiction.
2. If $\phi \in \mathsf{pos}(\mathcal{P})$ or $\phi \in \mathsf{neg}(\mathcal{P})$ already, we are done; we already know by (1) that $\phi \notin \mathsf{pos}(\mathcal{P}) \cap \mathsf{neg}(\mathcal{P})$.

   So ϕ does not already occur in $\mathcal{P}$. Suppose, for a contradiction, that adding ϕ to either set is inconsistent, i.e. both $\vdash \neg(\widehat{\mathcal{P}} \wedge \phi)$ and $\vdash \neg(\widehat{\mathcal{P}} \wedge \neg phi)$. By TAUT, that would imply that $\vdash \neg\widehat{\mathcal{P}} \wedge (\phi \vee \neg\phi)$, which is the same as simply $\vdash \neg\widehat{\mathcal{P}}$—a contradiction.
3. Suppose for a contradiction that $\bot \in \mathsf{pos}(\mathcal{P})$; by TAUT we have $\vdash \widehat{\mathcal{P}} \Rightarrow \bot$, which is syntactic sugar for $\vdash \neg\widehat{\mathcal{P}}$—a contradiction.
4. Suppose $\{\phi, \psi, \phi \Rightarrow \psi\} \subseteq \mathcal{F}_{\mathcal{P}}$.

   ($\phi \Rightarrow \psi \in \mathsf{pos}(\mathcal{P})$)) We must show that $\phi \in \mathsf{pos}(\mathcal{P})$ or that $\psi \in \mathsf{neg}(\mathcal{P})$. Suppose, for a contradiction, that neither is in the appropriate set; we then have $\vdash \widehat{\mathcal{P}} \Rightarrow (\phi \Rightarrow \psi) \wedge \phi \wedge \neg\psi$; by TAUT, we can then conclude $\vdash \neg\widehat{\mathcal{P}}$—a contradiction.

   ($\phi \in \mathsf{neg}(\mathcal{P})$ or $\psi \in \mathsf{pos}(\mathcal{P})$) We must show that $\phi \Rightarrow \psi \in \mathsf{pos}(\mathcal{P})$. Suppose, for a contradiction, that its not the case that $\phi \Rightarrow \psi \in \mathsf{pos}(\mathcal{P})$. Since $\phi \Rightarrow \psi \in \mathcal{F}_{\mathcal{P}}$, then $\phi \Rightarrow \psi \in \mathsf{neg}(\mathcal{P})$. We have either $\vdash \widehat{\mathcal{P}} \Rightarrow \neg(\phi \Rightarrow \psi) \wedge \neg\phi$ or $\vdash \widehat{\mathcal{P}} \Rightarrow \neg(\phi \Rightarrow \psi) \wedge \psi$. By TAUT, we can convert $\neg(\phi \Rightarrow \psi)$ into $\phi \wedge \neg\psi$—and either way we can find by TAUT that $\vdash \neg\widehat{\mathcal{P}}$, a contradiction.
5. Suppose $\vdash \phi \Rightarrow \psi$ and $\psi \in \mathsf{pos}(\mathcal{P})$ with $\phi \in \mathcal{F}_{\mathcal{P}}$. We must show that $\phi \in \mathsf{pos}(\mathcal{P})$. Suppose, for a contradiction, that $\phi \in \mathsf{neg}(\mathcal{P})$. We then have $\vdash \widehat{\mathcal{P}} \Rightarrow (\phi \Rightarrow \psi) \wedge \psi \wedge \neg\phi$; by TAUT, we can then find $\vdash \neg\widehat{\mathcal{P}}$, which is a contradiction. $\dashv$

Transition functions $\boxed{\sigma_i^\bullet : \mathsf{PNP} \to 2^{\mathrm{LTL}_f}}$ $\boxed{\sigma : \mathsf{PNP} \to \mathsf{PNP}}$

$$\begin{aligned}
\sigma_1^+(\mathcal{P}) &= \{\phi \mid \circ\,\phi \in \mathsf{pos}(\mathcal{P})\} \\
\sigma_2^+(\mathcal{P}) &= \{\phi\,\mathcal{W}\,\psi \mid \phi\,\mathcal{W}\,\psi \in \mathsf{pos}(\mathcal{P}), \psi \in \mathsf{neg}(\mathcal{P})\} \\
\sigma_3^-(\mathcal{P}) &= \{\phi \mid \circ\,\phi \in \mathsf{neg}(\mathcal{P})\} \\
\sigma_4^-(\mathcal{P}) &= \{\phi\,\mathcal{W}\,\psi \mid \phi\,\mathcal{W}\,\psi \in \mathsf{neg}(\mathcal{P}),\ \phi \in \mathsf{pos}(\mathcal{P})\} \\
\sigma(\mathcal{P}) &= (\sigma_1^+(\mathcal{P}) \cup \sigma_2^+(\mathcal{P}),\ \sigma_3^-(\mathcal{P}) \cup \sigma_4^-(\mathcal{P}))
\end{aligned}$$

$$\begin{array}{rcl rcl}
\tau(v) &=& \{v\} & \tau(\bot) &=& \{\bot\} \\
\tau(\phi \Rightarrow \psi) &=& \{\phi \Rightarrow \psi\} \cup \tau(\phi) \cup \tau(\psi) & \tau(\circ\,\phi) &=& \{\circ\,\phi\} \\
\tau(\phi\,\mathcal{W}\,\psi) &=& \{\phi\,\mathcal{W}\,\psi\} \cup \tau(\phi) \cup \tau(\psi) & & & \\
\tau(\mathcal{F}) &=& \bigcup_{\phi \in \mathcal{F}} \tau(\phi) & \tau(\mathcal{P}) &=& \tau(\mathcal{F}_{\mathcal{P}})
\end{array}$$

Extensions, completions, and possible assignments

$\boxed{\preceq\, \subseteq \mathsf{PNP} \times \mathsf{PNP}}$ $\boxed{\mathsf{comps} : 2^{\mathrm{LTL}_f} \to 2^{\mathsf{PNP}}}$ $\boxed{\mathsf{comps} : \mathsf{PNP} \to 2^{\mathsf{PNP}}}$

$\boxed{\mathsf{assigns} : 2^{\mathrm{LTL}_f} \to 2^{\mathsf{PNP}}}$

$$\mathcal{P} \preceq \mathcal{Q} \text{ iff } \mathsf{pos}(\mathcal{P}) \subseteq \mathsf{pos}(\mathcal{Q}) \text{ and } \mathsf{neg}(\mathcal{P}) \subseteq \mathsf{neg}(\mathcal{Q})$$
$$\mathsf{assigns}(\mathcal{F}) = \{\mathcal{P} \mid \mathcal{F}_{\mathcal{P}} = \tau(\mathcal{F})\}$$
$$\mathsf{comps}(\mathcal{P}) = \{\mathcal{Q} \mid \mathcal{F}_{\mathcal{Q}} = \tau(\mathcal{P}),\ \mathcal{P} \preceq \mathcal{Q},\ \mathcal{Q} \text{ consistent}\}$$

FIGURE 4. Step and closure functions; extensions and completions

We write $\mathcal{P} \preceq \mathcal{Q}$ (read "$\mathcal{P}$ is extended by $\mathcal{Q}$" or "$\mathcal{Q}$ extends $\mathcal{P}$") when $\mathcal{Q}$'s positive and negative sets subsume $\mathcal{P}$'s (Figure 4). We say $\mathcal{P}$ is complete when $\mathcal{F}_{\mathcal{P}} = \tau(\mathcal{P})$. We say a complete PNP $\mathcal{Q}$ is a *completion* of $\mathcal{P}$ when $\mathcal{P} \preceq \mathcal{Q}$ and $\mathcal{Q}$ is consistent and complete. We define the set of all consistent completions of a given PNP $\mathcal{P}$ with $\mathsf{comps}(\mathcal{P})$. Our goal is to generate successors states to build a graph of PNPs; to do so, we define two functions: a step function σ and a closure function τ (Figure 4). The step function σ takes a PNP and generates those formulae which must hold in the next step, thereby characterizing the transitions in our graph. The closure function τ takes a PNP and produces all of its subterms that are relevant for the current state, i.e., it doesn't go under the next modality. The set of completions, comps, is not a constructive set, since we have (as yet) no way to determine whether a given PNP is consistent or not.

First, we show that each PNP implies its successor (Lemma 16); next, consistent PNPs produce consistent successors (Lemma 17).

LEMMA 16 (Transitions are provable). *For all* $\mathcal{P} \in \mathsf{PNP}$*, we have* $\vdash \widehat{\mathcal{P}} \Rightarrow \bullet\,\widehat{\sigma(\mathcal{P})}$.

PROOF. Unfolding the definition of $\widehat{\mathcal{P}}$, we must show

$$\vdash \left[\bigwedge_{\phi\in\mathsf{pos}(\mathcal{P})} \phi \wedge \bigwedge_{\psi\in\mathsf{neg}(\mathcal{P})} \psi\right] \Rightarrow \bullet \left[\bigwedge_{\phi\in\mathsf{pos}(\sigma(\mathcal{P}))} \phi \wedge \bigwedge_{\psi\in\mathsf{neg}(\sigma(\mathcal{P}))} \psi\right].$$

By cases on the clauses of σ, we show that $\widehat{\mathcal{P}}$ implies each of the parts of $\widehat{\sigma(\mathcal{P})}$, tying the cases together by TAUT and Lemma 10:

(σ_1^+) Suppose $\phi \in \sigma_1^+(\mathcal{P})$ because $\circ\,\phi \in \mathcal{P}$. We have $\vdash \widehat{\mathcal{P}} \Rightarrow \bullet\,\phi$ because $\circ\,\phi \Rightarrow \bullet\,\phi$ by Lemma 8.

(σ_2^+) Suppose $\phi\ \mathcal{W}\ \psi \in \sigma_2^+(\mathcal{P})$ because $\vdash \phi\ \mathcal{W}\ \psi \in \mathsf{pos}(\mathcal{P})$ and $\psi \in \mathsf{neg}(\mathcal{P})$. Then, $\vdash \widehat{\mathcal{P}} \Rightarrow \neg\psi \wedge \phi\ \mathcal{W}\ \psi$. We have $\widehat{\mathcal{P}} \Rightarrow \bullet\,\phi\ \mathcal{W}\ \psi$ by WKUNTILUNROLL and TAUT.

(σ_3^-) Suppose $\phi \in \sigma_3^-(\mathcal{P})$ because $\circ\,\phi \in \mathsf{neg}(\mathcal{P})$. We have $\vdash \widehat{\mathcal{P}} \Rightarrow \bullet\,\neg\phi$ because $\neg\circ\,\phi \Rightarrow \bullet\,\neg\phi$ by Lemmas 7 and 8.

(σ_4^-) Suppose $\phi\ \mathcal{W}\ \psi \in \sigma_4^-(\mathcal{P})$ because $\phi\ \mathcal{W}\ \psi \in \mathsf{neg}(\mathcal{P})$ and $\phi \in \mathsf{pos}(\mathcal{P})$. We have:

$$\begin{array}{rll}
& \vdash \widehat{\mathcal{P}} \Rightarrow \neg(\phi\ \mathcal{W}\ \psi) \wedge \phi & \\
\text{iff} & \vdash \widehat{\mathcal{P}} \Rightarrow \neg(\psi \vee (\phi \wedge \bullet(\phi\ \mathcal{W}\ \psi))) \wedge \phi & \textsc{WkUntilUnroll} \\
\text{iff} & \vdash \widehat{\mathcal{P}} \Rightarrow \neg\psi \wedge \neg(\phi \wedge \bullet(\phi\ \mathcal{W}\ \psi)) \wedge \phi & \textsc{Taut} \\
\text{iff} & \vdash \widehat{\mathcal{P}} \Rightarrow \neg\psi \wedge (\neg\phi \vee \neg\,\bullet(\phi\ \mathcal{W}\ \psi)) \wedge \phi & \textsc{Taut} \\
\text{iff} & \vdash \widehat{\mathcal{P}} \Rightarrow \neg\psi \wedge \neg\,\bullet(\phi\ \mathcal{W}\ \psi) & \textsc{Taut} \\
\text{implies} & \vdash \widehat{\mathcal{P}} \Rightarrow \neg\psi \wedge \bullet\,\neg(\phi\ \mathcal{W}\ \psi) \wedge \phi & \textit{Lemma } 11 \\
\text{iff} & \vdash \widehat{\mathcal{P}} \Rightarrow \bullet\,\neg(\phi\ \mathcal{W}\ \psi) & \textsc{Taut}
\end{array}$$ ⊣

LEMMA 17 (Transitions are consistent). *For all consistent PNPs* $\mathcal{P}$*, if* $\vdash \widehat{\mathcal{P}} \Rightarrow \neg\mathsf{end}$ *then* $\sigma(\mathcal{P})$ *is consistent.*

PROOF. Let $\mathcal{P}$ be a consistent PNP such that $\vdash \widehat{\mathcal{P}} \Rightarrow \neg\mathsf{end}$. Assume for the sake of contradiction, that $\vdash \neg\widehat{\sigma(\mathcal{P})}$. Then we can write, by WKNEXT, $\vdash \bullet\neg\widehat{\sigma(\mathcal{P})}$, which is equivalent to $\vdash \neg\circ\widehat{\sigma(\mathcal{P})}$. By Lemma 16, $\vdash \widehat{\mathcal{P}} \Rightarrow \bullet(\widehat{\sigma(\mathcal{P})}) \wedge \neg\mathsf{end}$, which by Lemma 8 and TAUT gives $\vdash \widehat{\mathcal{P}} \Rightarrow \circ(\widehat{\sigma(\mathcal{P})})$. Now we can derive $\vdash \widehat{\mathcal{P}} \Rightarrow \bot$, or equivalently $\vdash \neg\widehat{\mathcal{P}}$—a contradiction. Conclude that $\sigma(\mathcal{P})$ is consistent. ⊣

Having established the fundamental properties of our successors, we must *complete* them: each PNP state needs to be 'saturated' to include all formulae of interest from the previous state. A consistent PNP will have many such possible completions—and we prove as much below—but we first observe that an inconsistent PNP one will have none.

LEMMA 18 (Inconsistent PNPs have no completions). *If a PNP $\mathcal{P}$ is inconsistent, then* $\mathsf{comps}(\mathcal{P}) = \emptyset$.

PROOF. Let $\mathcal{P}$ be given; suppose, for a contradiction, that there exists $\mathcal{Q} \in \mathsf{comps}(\mathcal{P})$, i.e, $\mathcal{F}_{\mathcal{Q}} = \tau(\mathcal{P})$ and $\mathcal{P} \preceq \mathcal{Q}$ and $\mathcal{Q}$ is consistent. We have $\vdash \widehat{\mathcal{Q}} \Rightarrow \widehat{\mathcal{P}}$ by TAUT, because $\mathcal{Q}$ is an extension of $\mathcal{P}$. But we know that $\vdash \neg\widehat{\mathcal{P}}$, so it must be the case that $\vdash \neg\widehat{\mathcal{Q}}$—which would mean that $\mathcal{Q}$ was inconsistent, a contradiction. ⊣

In order to fully define our graph, we must show that not only are successors of PNPs provable, so are their completions. We do so in two steps: first, we show that there is always *some* provable assignment of propositions in each set of formulas.

LEMMA 19 (Assignments are provable). $\vdash \bigvee_{\mathcal{P} \in \mathsf{assigns}(\mathcal{F})} \widehat{\mathcal{P}}$

PROOF. By induction on the size of $\mathcal{F}$.

($|\mathcal{F}| = 0$) We have $\vdash \top$ by TAUT.

($|\mathcal{F}| = n+1$) Let $\phi \in \mathcal{F}$ be a maximal formula, i.e., $\phi \notin \tau(\mathcal{F} - \{\phi\})$.We have $\mathsf{assigns}(\mathcal{F}) = \{\mathcal{P} \mid \mathcal{P}' \in \mathsf{assigns}(\mathcal{F}'),\ \mathcal{F}_{\mathcal{P}} = \mathcal{F}_{\mathcal{P}'} \cup \tau(\phi)\}$, i.e., each formula in $\tau(\phi)$ not already assigned in $\mathcal{P}'$ is put in either the positive or negative set of $\mathcal{P}$. That is, we take each formula in $\mathcal{P}'$ and conjoin $\psi \vee \neg\psi$ for each $\psi \in \tau(\phi)$. We know by the IH that $\vdash \bigvee_{\mathcal{P}' \in \mathsf{assigns}(\mathcal{F}')} \widehat{\mathcal{P}'}$, so by TAUT we have $\vdash \bigvee_{\mathcal{P} \in \mathsf{assigns}(\mathcal{F})} \widehat{\mathcal{P}}$. ⊣

Having established that there are consistent assignments, we can show that conditionally provable assignments are in fact completions.

LEMMA 20 (Consistent assignments are completions).
For all consistent PNPs $\mathcal{P}$ and for all $\mathcal{Q} \in \mathsf{assigns}(\widehat{\mathcal{P}})$, if $\vdash \widehat{\mathcal{P}} \Rightarrow \widehat{\mathcal{Q}}$ then $\mathcal{Q} \in \mathsf{comps}(\mathcal{P})$.

PROOF. Let $\mathcal{P}$ and $\mathcal{Q} \in \mathsf{assigns}(\widehat{\mathcal{P}})$ be given such that $\vdash \widehat{\mathcal{P}} \Rightarrow \widehat{\mathcal{Q}}$.
Suppose for a contradiction that $\mathcal{Q} \notin \mathsf{comps}(\mathcal{P})$. It must be the case that either $\mathcal{Q}$ does not extend $\mathcal{P}$ or $\mathcal{Q}$ is inconsistent—we will show that both cases are contradictory.

If $\mathcal{P} \not\preceq \mathcal{Q}$, then there exists some formula ϕ such that $\phi \in \mathsf{pos}(\mathcal{P})$ and $\phi \in \mathsf{neg}(\mathcal{Q})$ or vice versa. Then, $\vdash \widehat{\mathcal{P}} \Rightarrow \phi$ and $\vdash \widehat{\mathcal{Q}} \Rightarrow \neg\phi$. Then $\vdash \widehat{\mathcal{P}} \wedge \widehat{\mathcal{Q}} \Rightarrow \phi \wedge \neg\phi$, which by TAUT means $\vdash \neg(\widehat{\mathcal{P}} \wedge \widehat{\mathcal{Q}})$, or equivalently, that $\vdash \widehat{\mathcal{P}} \Rightarrow \neg\widehat{\mathcal{Q}}$. When combined with the assumption that $\vdash \widehat{\mathcal{P}} \Rightarrow \widehat{\mathcal{Q}}$, via TAUT, we can derive $\vdash \neg\widehat{\mathcal{P}}$—a contradiction with $\mathcal{P}$'s consistency.

If, on the other hand $\mathcal{Q}$ is inconsistent, then we can see from $\vdash \widehat{\mathcal{P}} \Rightarrow \widehat{\mathcal{Q}}$ that $\vdash \neg\widehat{\mathcal{Q}}$—and by TAUT, it must be that $\vdash \neg\widehat{\mathcal{P}}$, which contradicts $\mathcal{P}$'s consistency. ⊣

Combining the last two proofs we find that consistent completions are provable.

LEMMA 21 (Consistent completions are provable).
For all consistent PNPs $\mathcal{P}$, *we have* $\vdash \widehat{\mathcal{P}} \Rightarrow \bigvee_{\mathcal{Q} \in \mathsf{comps}(\mathcal{P})} \widehat{\mathcal{Q}}$.

PROOF. By Lemma 19, we have $\vdash \bigvee_{\mathcal{Q} \in \mathsf{assigns}(\widehat{\mathcal{P}})} \widehat{\mathcal{Q}}$. By TAUT, we have $\vdash \widehat{\mathcal{P}} \Rightarrow \bigvee_{\mathcal{Q} \in \mathsf{assigns}(\widehat{\mathcal{P}})} \widehat{\mathcal{Q}}$. By Lemma 20, we know that we only need to keep those $\mathcal{Q} \in \mathsf{assigns}(\widehat{\mathcal{P}})$ which are also in $\mathsf{comps}(\mathcal{P})$, and so we have $\vdash \widehat{\mathcal{P}} \Rightarrow \bigvee_{\mathcal{Q} \in \mathsf{comps}(\mathcal{P})} \widehat{\mathcal{Q}}$ as desired. ⊣

Having established the fundamental properties of consistent completions, we set about defining the structure on which we build our proof. We show that, starting from a PNP formed from a given formula, we can construct a graph where nodes are PNPs and a node $\mathcal{P}$'s successors are consistent completions of $\sigma(\mathcal{P})$.

DEFINITION 22 (Proof graphs). For a consistent and complete PNP $\mathcal{P}$ (i.e., where $\mathcal{F}_{\mathcal{P}} = \tau(\mathcal{P})$ and it is not the case that $\vdash \neg\widehat{\mathcal{P}}$), we define a *proof graph* $\mathcal{G}_{\mathcal{P}}$ as follows: (a) $\mathcal{P}$ is the *root* of $\mathcal{G}_{\mathcal{P}}$; (b) $\mathcal{P}$ has an edge to the root of $\mathcal{G}_{\mathcal{Q}}$ for each $\mathcal{Q} \in \mathsf{comps}(\sigma(\mathcal{P}))$. The $\mathcal{Q} \in V(\mathcal{G}_{\mathcal{P}})$ are those PNPs reachable from $\mathcal{P}$.

Since $\mathcal{P}$ is composed of a finite number of formulae, the set of all subsets of $\tau(\mathcal{P})$ is finite, as are any assignments of those subsets to PNPs. Hence the number of nodes in the proof graph must be finite.[4]

Our innovation in adapting the completeness proof to finite time is *finiteness injection*, where we make sure that $\Diamond\,\mathsf{end}$ is in the positive set of the root of the proof graph we construct to show completeness. After injecting finiteness, every node of the proof graph will either have end in its positive set (and no successors) or all of its successors have $\Diamond\,\mathsf{end}$ in their positive set.

Every lemma we prove, from here to the final completeness result, will have some premise concerning the end of time: by only working with PNPs with $\Diamond\,\mathsf{end}$ in the positive set, we guarantee that time eventually ends.

LEMMA 23 (end injection is invariant). *If* $\mathcal{P}$ *is a consistent and complete PNP with* $\Diamond\,\mathsf{end} \in \mathsf{pos}(\mathcal{P})$*, then either:*

- $\mathsf{end} \in \mathsf{pos}(\mathcal{P})$ *and* $\mathcal{P}$ *has no successors (i.e.,* $\mathsf{comps}(\sigma(\mathcal{P})) = \emptyset$*), or*
- $\mathsf{end} \in \mathsf{neg}(\mathcal{P})$ *and for all* $\mathcal{Q} \in \mathsf{comps}(\sigma(\mathcal{P}))$*, we have* $\Diamond\,\mathsf{end} \in \mathsf{pos}(\mathcal{Q})$*.*

PROOF. Recall that $\Diamond\,\mathsf{end}$ desugars to $\neg(\neg\neg \circ \top \;\mathcal{W}\; \bot)$. Since $\mathcal{P}$ is complete, we know that $\mathsf{end} \in \mathcal{F}_{\mathcal{P}}$. By cases on where end appears:

[4] Confusingly, Kröger and Merz [10] call this graph an "infinite tree" in their proof of completeness for potentially infinite LTL, even though it turns out to be finite in that setting, as well.

($\mathsf{end} \in \mathsf{pos}(\mathcal{P})$) If end (i.e., $\neg \circ \top$) is in $\mathsf{pos}(\mathcal{P})$ and $\mathcal{P}$ is consistent, it must be the case that $\circ \top \in \mathsf{neg}(\mathcal{P})$ by Lemma 15. We therefore have that $\top \in \sigma_3^-(\mathcal{P})$, so $\vdash \widehat{\sigma(\mathcal{P})} \Rightarrow \neg\top$, i.e., $\sigma(\mathcal{P})$ is inconsistent—and therefore $\mathsf{comps}(\sigma(\mathcal{P})) = \emptyset$, because there are no consistent completions of an inconsistent PNP (Lemma 18).

($\mathsf{end} \in \mathsf{neg}(\mathcal{P})$) If end (i.e., $\neg \circ \top$) is in $\mathsf{neg}(\mathcal{P})$ and $\mathcal{P}$ is consistent, then it must be the case that $\circ \top \in \mathsf{pos}(\mathcal{P})$. We must have have $\Diamond\,\mathsf{end} \in \sigma_2^+(\mathcal{P})$, which means $\Diamond\,\mathsf{end} \in \mathsf{pos}(\sigma(\mathcal{P}))$. It must be therefore be the case that $\Diamond\,\mathsf{end} \in \mathsf{pos}(\mathcal{Q})$ for all $\mathcal{Q} \in \mathsf{comps}(\sigma(\mathcal{P}))$, since each such $\mathcal{Q}$ must be an extension of $\sigma(\mathcal{P})$. $\dashv$

We can go further, showing that $\Diamond\,\mathsf{end}$ is in fact in *every* node's positive set.

LEMMA 24 (Proof graphs are consistent). *For all consistent and complete PNPs $\mathcal{P}$, every node $\mathcal{Q} \in \mathcal{G}_\mathcal{P}$ is consistent and complete. If $\Diamond\,\mathsf{end} \in \mathsf{pos}(\mathcal{P})$, then $\Diamond\,\mathsf{end} \in \mathsf{pos}(\mathcal{Q})$.*

PROOF. By induction on the length of the shortest path from $\mathcal{P}$ to $\mathcal{Q}$ in $\mathcal{G}_\mathcal{P}$. When $n = 0$, we have $\mathcal{Q} = \mathcal{P}$, so we have $\mathcal{P}$'s completeness and consistency by assumption; the second implication is immediate.

When $n = n' + 1$, we have some path $\mathcal{P}, \mathcal{P}_2, \mathcal{P}_3, \ldots, \mathcal{P}_{n'}, \mathcal{Q}$. We know that $\mathcal{P}_{n'}$ is complete and consistent; we must show that $\mathcal{Q}$ is complete and consistent. By construction, we know that $\mathcal{Q} \in \mathsf{comps}(\sigma(\mathcal{P}_{n'}))$, so $\mathcal{Q}$ must be consistent and complete by definition. By the IH, we know that $\Diamond\,\mathsf{end} \in \mathsf{pos}(\mathcal{P}_{n'})$, so by Lemma 23, we find the same for $\mathcal{Q}$. $\dashv$

Each node has the potential for successors: for each node $\mathcal{Q} \in \mathcal{G}_P$, we can prove that $\widehat{\mathcal{Q}}$ implies the disjunction of every other node's literal interpretation.

LEMMA 25 (Step implication). *For all consistent and complete PNPs $\mathcal{P}$ where $\Diamond\,\mathsf{end} \in \mathsf{pos}(\mathcal{P})$ then $\vdash \bigvee_{\mathcal{Q} \in \mathcal{G}_\mathcal{P}} \widehat{\mathcal{Q}} \Rightarrow \bullet \bigvee_{\mathcal{Q} \in \mathcal{G}_\mathcal{P}} \widehat{\mathcal{Q}}$.*

PROOF. By Lemma 16, we know that $\vdash \widehat{\mathcal{Q}} \Rightarrow \bullet \widehat{\sigma(\mathcal{Q})}$. By Lemma 24, we know that $\mathcal{Q}$ is consistent and complete and $\Diamond\,\mathsf{end} \in \mathsf{pos}(\mathcal{Q})$. Since $\mathcal{Q}$ is complete, we know that $\mathsf{end} \in \mathcal{F}_\mathcal{Q}$. We go by cases on where end occurs:

($\mathsf{end} \in \mathsf{pos}(Q)$) By Lemma 23, we know that $\mathsf{comps}(\sigma(\mathcal{Q})) = \emptyset$, so we must find $\vdash \widehat{\mathcal{Q}} \Rightarrow \bullet \bot$. Since $\vdash \widehat{\mathcal{Q}} \Rightarrow \mathsf{end}$, we are done by Lemma 8 with $\phi = \bot$.

($\mathsf{end} \in \mathsf{neg}(Q)$) We have $\vdash \widehat{\mathcal{Q}} \Rightarrow \neg\mathsf{end}$, so by Lemma 17 we know that $\sigma(Q)$ is consistent. We therefore have $\vdash \widehat{\sigma(\mathcal{Q})} \Rightarrow \bigvee_{Q' \in \mathsf{comps}(\sigma(\mathcal{Q}))} \widehat{\mathcal{Q}}'$ by Lemma 21.

Since $\vdash \widehat{\mathcal{Q}} \Rightarrow \bullet \bigvee_{\mathcal{Q}' \in \mathsf{comps}(\sigma(\mathcal{Q}))} \widehat{\mathcal{Q}}'$, we can show that $\vdash \widehat{\mathcal{Q}} \Rightarrow \bullet \bigvee_{\mathcal{Q}' \in \mathcal{G}_{\mathcal{P}}} \widehat{\mathcal{Q}}'$, because $\mathsf{comps}(\sigma(\mathcal{Q})) \subseteq \mathcal{G}_{\mathcal{P}}$ by definition. Since we find this for each $\mathcal{Q}$, we conclude $\vdash \bigvee_{\mathcal{Q} \in \mathcal{G}_{\mathcal{P}}} \widehat{\mathcal{Q}} \Rightarrow \bullet \bigvee_{\mathcal{Q} \in \mathcal{G}_{\mathcal{P}}} \widehat{\mathcal{Q}}$. ⊣

We have so far established that the proof graph $\mathcal{G}_{\mathcal{P}}$ is rooted at $\mathcal{P}$, preserves any finiteness we may inject, and has provable successors. We are nearly done: we show that our proof graph corresponds to a Kripke structure which models $\mathcal{P}$.

Definition 26 (Terminal nodes and paths). A node $\mathcal{Z} \in \mathcal{G}_{\mathcal{P}}$ is *terminal* when $\circ \top \in \mathsf{neg}(\mathcal{Z})$. A path $\mathcal{P}_1, \ldots, \mathcal{P}_n$ is *terminal* when $\mathcal{P}_n$ is terminal.

Lemma 27 (Proof graphs are models). *For all consistent and complete PNPs $\mathcal{P}$, if $\mathcal{P}_1, \mathcal{P}_1, \mathcal{P}_2, \ldots, \mathcal{P}_n$ is a terminal path in $\mathcal{G}_{\mathcal{P}}$, then for all i:*

1. *For all formulae ϕ, if $\circ\,\phi \in \mathcal{F}_{\mathcal{P}_i}$ then $\circ\,\phi \in \mathsf{pos}(\mathcal{P}_i)$ iff $\phi \in \mathsf{pos}(\mathcal{P}_{i+1})$.*
2. *For all formulae ϕ and ψ, if $\phi \mathcal{W} \psi \in \mathcal{F}_{\mathcal{P}_i}$ then $\phi \mathcal{W} \psi \in \mathsf{pos}(\mathcal{P}_i)$ iff either $\phi \in \mathsf{pos}(P_j)$ for all $j \geq i$ or there is some $k \geq i$ such that $\psi \in \mathsf{pos}(\mathcal{P}_k)$ and $\forall i \leq j < k,\ \phi \in \mathsf{pos}(\mathcal{P}_j)$.*

Proof.

1. We have $\mathcal{P}_{i+1} \in \mathsf{comps}(\sigma(\mathcal{P}_i))$ by definition. We go by cases:
   ($\circ\,\phi \in \mathsf{pos}(\mathcal{P}_i)$) We have $\phi \in \mathsf{pos}(\sigma(\mathcal{P}_i))$, and so all consistent completions have ϕ in the positive set—in particular, $\mathcal{P}_{i+1}$.
   ($\phi \in \mathsf{pos}(\mathcal{P}_{i+1})$) Since $\circ\,\phi \in \mathcal{F}_{P_i}$, it must be the case that $\circ\,\phi$ is in one of $\mathsf{pos}(\mathcal{P}_i)$ or $\mathsf{neg}(\mathcal{P}_i)$. In the former case, we are done immediately. Suppose, for a contradiction, that $\circ\,\phi \in \mathsf{neg}(\mathcal{P}_i)$. Since $\mathcal{P}_{i+1}$ is a completion of $\sigma(\mathcal{P}_i)$, it must be that $\mathsf{neg}(\sigma(\mathcal{P}_i)) \subseteq \mathsf{neg}(\mathcal{P}_{i+1})$. Since $\circ\,\phi \in \mathsf{neg}(\mathcal{P}_i)$, we must have $\phi \in \mathsf{neg}(\sigma(\mathcal{P}_i))$, so $\phi \in \mathsf{neg}(\mathcal{P}_{i+1})$. But we have $\phi \in \mathsf{pos}(\mathcal{P}_{i+1})$ by assumption—and we have contradicted the consistency of $\mathcal{P}_{i+1}$ (Lemma 24).
2. We have $\mathcal{P}_{j+1} \in \mathsf{comps}(\sigma(\mathcal{P}_j))$ for all j, by definition. Further, we know that $\phi \mathcal{W} \psi \in \mathsf{pos}(\mathcal{P}_i)$ implies that $\{\phi, \psi\} \subseteq \mathcal{F}_{\mathcal{P}_i}$. We go by cases on where $\phi \mathcal{W} \psi$ occurs in $\mathcal{F}_{\mathcal{P}_i}$:
   ($\phi \mathcal{W} \psi \in \mathsf{pos}(\mathcal{P}_i)$) We must show that either $\phi \in \mathsf{pos}(P_j)$ for all $j \geq i$ or there is some $k \geq i$ such that $\psi \in \mathsf{pos}(\mathcal{P}_k)$ and $\forall i \leq j < k,\ \phi \in \mathsf{pos}(\mathcal{P}_j)$.
   We show (for all $\mathcal{P}$ on the path) that if $\phi \mathcal{W} \psi \in \mathsf{pos}(\mathcal{P})$, then either $\psi \in \mathsf{pos}(\mathcal{P})$ or $\phi \in \mathsf{pos}(\mathcal{P})$ and for all $\mathcal{Q} \in \mathsf{comps}(\sigma(\mathcal{P}))$, we have $\phi \mathcal{W} \psi \in \mathsf{pos}(\mathcal{Q})$.
   Since $\phi \mathcal{W} \psi \in \mathsf{pos}(\mathcal{P})$, by WkUntilUnroll we know that $\vdash \widehat{\mathcal{P}} \Rightarrow \psi \vee \phi \wedge \bullet(\phi \mathcal{W} \psi)$. Since ψ and ϕ are both in $\mathcal{F}_{\mathcal{P}}$, we can simply inspect $\mathcal{P}$. If $\psi \in \mathsf{pos}(\mathcal{P})$, we are done. So suppose $\psi \in \mathsf{neg}(\mathcal{P})$. We must therefore have $\phi \in \mathsf{pos}(\mathcal{P})$. By the definition

of σ_2^+, we have $\phi \mathcal{W} \psi \in \sigma(\mathcal{P})$, and so any $\mathcal{Q} \in \mathsf{comps}(\sigma(\mathcal{P}))$ must also have $\phi \mathcal{W} \psi \in \mathsf{pos}(\mathcal{Q})$.

We strengthen the inductive hypothesis, showing that for the remainder of the terminal path $\mathcal{P}_i \ldots \mathcal{P}_{i+n}$ either $\{\phi, \phi \mathcal{W} \psi\} \subseteq \mathsf{pos}(\mathcal{P}_j)$ for all $i \leq j \leq n$, or there exists a $k \geq i$ such that $\psi \in \mathsf{pos}(\mathcal{P}_k)$ and $\{\phi, \phi \mathcal{W} \psi\} \in \mathsf{pos}(\mathcal{P}_j)$ for all $i \leq j < k$. We go by induction on n.

$(n = 0)$ By the above, we either have $\psi \in \mathsf{pos}(\mathcal{P}_i)$ (and so $k = i$) or $\phi \in \mathsf{pos}(\mathcal{P}_i)$ (and then the path ends).

$(n = n' + 1)$ We know the path from $\mathcal{P}_i$ to $\mathcal{P}_{i+n}$ has either ϕ in every positive set or eventually ψ occurs after ϕs. In the latter case, we can simply reuse the k from the inductive hypothesis.

In the former case, we know $\{\phi, \phi \mathcal{W} \psi\} \subseteq \mathsf{pos}(\mathcal{P}_{n'})$, so by the above we can find that either $\psi \in \mathcal{P}_{n'}$ or since $\mathcal{P}_n \in \mathsf{comps}(\sigma(\mathcal{P}_{n'}))$ has $\phi \mathcal{W} \psi \in \mathsf{pos}(\mathcal{P}_n)$. By the above again, we can find that either $\psi \in \mathsf{pos}(\mathcal{P}_n)$ (and so $k = n$) or $\phi \in \mathsf{pos}(\mathcal{P}_n)$ (and we have $\phi \in \mathsf{pos}(\mathcal{P}_j)$ for all $j \geq i$).

$(\phi \mathcal{W} \psi \notin \mathsf{pos}(\mathcal{P}_i))$ We have $\phi \mathcal{W} \psi \in \mathsf{neg}(\mathcal{P}_i)$, so we must show that it is not the case that either $\phi \in \mathsf{pos}(P_j)$ for all $j \geq i$ or there is some $k \geq i$ such that $\psi \in \mathsf{pos}(\mathcal{P}_k)$ and $\forall i \leq j < k,\ \phi \in \mathsf{pos}(\mathcal{P}_j)$. We show that all paths out of $\mathcal{P}_i$ have ϕ in the positive set for zero or more transitions, but eventually neither ϕ nor ψ holds.

First, we show that if $\phi \mathcal{W} \psi \in \mathsf{neg}(\mathcal{P})$, then (a) $\psi \in \mathsf{neg}(\mathcal{P})$ and (b) either $\phi \in \mathsf{neg}(\mathcal{P})$ or $\phi \in \mathsf{pos}(\mathcal{P})$ and $\forall \mathcal{Q} \in \mathsf{comps}(\sigma(\mathcal{P}))$, we have $\phi \mathcal{W} \psi \in \mathsf{neg}(\mathcal{Q})$.

Since $\phi \mathcal{W} \psi \in \mathsf{neg}(\mathcal{P})$, we have $\vdash \widehat{\mathcal{P}} \Rightarrow \neg(\psi \vee \phi \wedge \bullet(\phi \mathcal{W} \psi))$ by WkUntilUnroll. By Taut we have $\vdash \widehat{\mathcal{P}} \Rightarrow \neg\psi \wedge (\neg phi \vee \neg \bullet(\phi \mathcal{W} \psi)$; by desugaring and Taut we have $\vdash \widehat{\mathcal{P}} \Rightarrow \neg\psi \wedge (\neg\phi \vee \circ\neg(\phi \mathcal{W} \psi)$.

To have $\mathcal{P}$ consistent, it must be that $\psi \in \mathsf{neg}(\mathcal{P})$. If $\phi \in \mathsf{neg}(\mathcal{P})$, we are done—we have satisfied (a) and (b). Suppose $\phi \in \mathsf{pos}(\mathcal{P})$. By the definition σ_4^-, we now have $\phi \mathcal{W} \psi \in \mathsf{neg}(\sigma(\mathcal{P}))$, so it must be the case that for any completion $\mathcal{Q}$, we have $\phi \mathcal{W} \psi \in \mathsf{neg}(\mathcal{Q})$.

Now, finally, suppose $\phi \mathcal{W} \psi \in \mathsf{neg}(\mathcal{P}_i)$. For $\mathcal{P}_i$ to be consistent, it will be the case that $\vdash \widehat{\mathcal{P}}_i \Rightarrow \widehat{\mathcal{P}}_{i+1} \Rightarrow \ldots \Rightarrow \widehat{\mathcal{P}}_n$ (since no node is terminal until $\mathcal{P}_n$). One such node must have $\phi \in \mathsf{neg}(\mathcal{P}_i)$: apply the reasoning above to see that no node can have $\psi \in \mathsf{pos}(\mathcal{P}_i)$ and, furthermore, if $\phi \in \mathsf{pos}(\mathcal{P}_i)$ then $\phi \mathcal{W} \psi \in \mathsf{neg}(\mathcal{P}_{i+1})$. If it does not happen before the terminal node, the last one has no successor, so WkUntilUnroll shows that necessarily $\phi \in \mathsf{neg}(\mathcal{P}_n)$. $\dashv$

Here we slightly depart from Kröger and Merz's presentation: since their models can be infinite, they must make sure that their paths are able to in some sense 'fulfill' temporal predicates. We, on the other hand, know that all of our paths will be finite, so our reasoning can be simpler. First, there must exist some terminal node.

LEMMA 28 (Injected finiteness guarantees terminal nodes).
For all consistent and complete PNPs $\mathcal{P}$, if $\Diamond\,\mathsf{end} \in \mathsf{pos}(\mathcal{P})$ then there is a terminal node $\mathcal{Z} \in \mathcal{G}_{\mathcal{P}}$.

PROOF. Suppose, for the sake of a contradiction, that $\circ\top \in \mathsf{pos}(Q)$ for all $\mathcal{Q} \in \mathcal{G}_{\mathcal{P}}$.

We have assumed $\Diamond\,\mathsf{end} \in \mathsf{pos}(\mathcal{P})$; by desugaring, $\Diamond\,\mathsf{end}$ amounts to $\neg\Box\neg\mathsf{end}$, i.e., $\neg(\neg\mathsf{end}\;\mathcal{W}\;\bot)$, i.e., $\neg(\circ\top\;\mathcal{W}\;\bot)$.

We have $\vdash \mathcal{Q} \Rightarrow \circ\top$ for each $\mathcal{Q} \in \mathcal{G}_{\mathcal{P}}$ by assumption, so by TAUT, we have $\vdash \bigvee_{\mathcal{Q}\in\mathcal{G}_{\mathcal{P}}} \widehat{\mathcal{Q}} \Rightarrow \circ\top$.

We have $\vdash \bigvee_{\mathcal{Q}\in\mathcal{G}_{\mathcal{P}}} \widehat{\mathcal{Q}} \Rightarrow \bullet\bigvee_{\mathcal{Q}\in\mathcal{G}_{\mathcal{P}}}$ by Lemma 25. So by INDUCTION, $\vdash \bigvee_{\mathcal{Q}\in\mathcal{G}_{\mathcal{P}}} \Rightarrow \Box\circ\top$, i.e, $\circ\top\;\mathcal{W}\;\bot$. Since $\mathcal{P} \in \mathcal{G}_{\mathcal{P}}$, we know $\vdash \widehat{\mathcal{P}} \Rightarrow \circ\top\;\mathcal{W}\;\bot$ by Lemma 25 again. But $\Diamond\,\mathsf{end} \in \mathsf{pos}(\mathcal{P})$ means that $\vdash \widehat{\mathcal{P}} \Rightarrow \Diamond\,\mathsf{end}$, so $\vdash \widehat{\mathcal{P}} \Rightarrow \neg(\circ\top\;\mathcal{W}\;\bot)$, as well! It must then be the case that $\vdash \neg\widehat{\mathcal{P}}$, which contradicts our assumption that $\mathcal{P}$ is consistent.

We therefore conclude that there must exist some node $\mathcal{Z} \in \mathcal{G}_{\mathcal{P}}$ such that $\circ\top \in \mathsf{neg}(\mathcal{Z})$. ⊣

Since our proof graph is constructed to be connected, the existence of a terminal node implies the existence of a terminal path.

COROLLARY 29 (Injected finiteness guarantees terminal paths).
For all consistent and complete PNPs $\mathcal{P}$, if $\Diamond\,\mathsf{end} \in \mathsf{pos}(\mathcal{P})$ then there is a terminal path $\mathcal{P}, \mathcal{P}_2, \ldots, \mathcal{P}_{n-1}, \mathcal{Z} \in \mathcal{G}_{\mathcal{P}}$.

PROOF. By Lemma 28, there exists some terminal node $\mathcal{Z} \in \mathcal{G}_{\mathcal{P}}$. Since $\mathcal{G}_{\mathcal{P}}$ is constructed by iterating comps and σ on $\mathcal{P}$, there must exist some $\mathcal{P}_{n-1}$ such that $\mathcal{Z} \in \mathsf{comps}(\sigma(\mathcal{P}_{n-1}))$, and some $\mathcal{P}_{n-2}$ such that $\mathcal{P}_{n-1} \in \mathsf{comps}(\sigma(\mathcal{P}_{n-2}))$ and so on back to $\mathcal{P}$—yielding a path. ⊣

We can now prove the key lemma: consistent PNPs are satisfiable. To do this we show that a proof graph for a consistent PNP $\mathcal{P}$ induces a Kripke structure modeling $\mathcal{P}$'s literal interpretation, $\widehat{\mathcal{P}}$. The proof actually considers a version of $\mathcal{P}$ with $\Diamond\,\mathsf{end}$ (the FINITE axiom) injected into the positive set—we inject finiteness to make sure we're building an appropriate, finite model.

THEOREM 30 (LTL$_f$ satisfiability). *If $\mathcal{P}$ is a consistent PNP, then $\widehat{\mathcal{P}}$ is satisfiable.*

PROOF. Let $\mathcal{P}' = (\{\Diamond \mathsf{end}\} \cup \mathsf{pos}(\mathcal{P}), \mathsf{neg}(\mathcal{P}))$. If $\mathcal{P}$ is consistent, then so is $\mathcal{P}'$. (If not, it must be because $\vdash \widehat{\mathcal{P}} \Rightarrow \neg \Diamond \mathsf{end}$; by TAUT and FINITE, we have $\vdash \widehat{\mathcal{P}} \Rightarrow \Diamond \mathsf{end}$, and so $\vdash \neg\widehat{\mathcal{P}}$ and $\mathcal{P}$ is not consistent.)

To show that $\widehat{\mathcal{P}}$ is satisfiable, we will use the terminal path from Corollary 29 to construct a Kripke structure. Suppose our terminal path is of the form $\mathcal{P}, \mathcal{P}_2, \ldots, \mathcal{P}_n$; let $\mathsf{K}^n = (\eta_1, \ldots, \eta_n)$ where we define:

$$\eta_i(v) = \begin{cases} \mathfrak{t} & v \in \mathsf{pos}(\mathcal{P}_i) \\ \mathfrak{f} & \text{otherwise} \end{cases}$$

We must show that $\mathsf{K}^n_1(\widehat{\mathcal{P}}) = \mathfrak{t}$; it will suffice to show that $\mathsf{K}^n_1(\widehat{\mathcal{P}'}) = \mathfrak{t}$. To do so, we prove generally that for all $\phi \in \mathcal{F}_{P'}$, we have $\mathsf{K}^n_i(\phi) = \mathit{true}$ iff $phi \in \mathsf{pos}(\mathcal{P}_i)$. We go by induction on ϕ; throughout, we rely on the fact that every node is consistent and complete (Lemma 24).

- ($\phi = v$) By the definition of K^n and η_i.
- ($\phi = \bot$) By definition, we have $\mathsf{K}^n_i(\bot) = \mathfrak{f}$ and $\bot \notin \mathsf{pos}(\mathcal{P}_i)$ by Lemma 15.
- ($\phi = \phi \Rightarrow \psi$) Let an i be given. We know $\mathcal{P}_i$ is a consistent and complete PNP, so $\{\phi, \psi\} \in F_{\mathcal{P}_i}$. By the IH, we have $\mathsf{K}^n_i(\phi) = \mathfrak{t}$ iff $\phi \in \mathsf{pos}(\mathcal{P}_i)$ and similarly for ψ. We have $\mathsf{K}^n_i(\phi \Rightarrow \psi) = \mathfrak{t}$ iff $\mathsf{K}^n_i(\phi) = \mathfrak{f}$ or $\mathsf{K}^n_i(\psi) = \mathfrak{t}$ iff $\phi \in \mathsf{neg}(\mathcal{P}_i)$ or $\psi \in \mathsf{pos}(\mathcal{P}_i)$ iff $\phi \Rightarrow \psi \in \mathsf{pos}(\mathcal{P}_i)$(again by Lemma 15).
- ($\phi = \circ\, \phi$) Let an i be given. We have $\mathsf{K}^n_i(\circ\, \phi) = \mathfrak{t}$ iff $i > n$ and $\mathsf{K}^n_{i+1}(\phi)\mathfrak{t}$ iff in $\mathsf{K}^n_{i+1}(\phi) = \mathfrak{t}$ iff $\phi \in \mathsf{pos}(\mathcal{P}_{i+1})$ (by the IH) iff $\circ\, \phi \in \mathsf{pos}(\mathcal{P}_i)$ (since $\circ\, \phi \in \mathcal{F}_{\mathcal{P}'}$, we can apply Lemma 27).
- ($\phi = \phi\, \mathcal{W}\, \psi$) We have $\mathsf{K}^n_i(\phi\ \mathcal{W}\ \psi) = \mathfrak{t}$ iff either for all $i \leq j \leq n, \mathsf{K}^n_j(\phi) = \mathfrak{t}$ or there exists a $i \leq k \leq n$ such that $\mathsf{K}^n_k(\psi) = \mathfrak{t}$ and for all $i \leq j < k$ we have $\mathsf{K}^n_j(\phi) = \mathfrak{t}$. By the IH, those hold iff formulae are in appropriate positive sets; by Lemma 27, those formulae are in appropriate positive sets iff $\phi\, \mathcal{W}\, \psi$ is in the appropriate positive set.

At this point, we can see that $\mathsf{K}^n_1(\widehat{\mathcal{P}}) = \mathfrak{t}$. ⊣

Finally, we can show completeness. The proof is the usual one, where we to find a proof of ϕ we try to see if $\neg\phi$ is satisfiable—if not, then the PNP for $\neg\phi$ will be inconsistent, and so $\vdash \neg\neg\phi$, which yields $\vdash \phi$.

THEOREM 31 (LTL$_f$ completeness). *If* $\models \phi$ *then* $\vdash \phi$.

PROOF. If $\models \phi$, then for all Kripke structures K^n, we have $\mathsf{K}^n_i(\phi) = \mathfrak{t}$ for all i. Conversely, it must also be the case that $\mathsf{K}^n_i(\neg\phi) = \mathfrak{f}$ for all i, and so $\neg\phi$ is unsatisfiable. In other words, the PNP $(\emptyset, \{\phi\})$ is unsatisfiable. By the contrapositive of Theorem 30, it must be the case that $(\emptyset, \{\phi\})$ is inconsistent, i.e., $\vdash \neg\neg\phi$. By TAUT, we can conclude that $\vdash \phi$. ⊣

We extend the proof of completeness to allow for assumptions in the usual way.

Corollary 32 (LTL_f completeness, with contexts). *If* $\mathcal{F} \models \phi$ *then* $\mathcal{F} \vdash \phi$.

Proof. By induction on the size of $\mathcal{F}$.

($|F| = 0$) By Theorem 31.
($|F| = n+1$) We have $\{\phi_1, \ldots, \phi_{n+1}\} \models \psi$. By Theorem 3, we have $\{\phi_1, \ldots, \phi_n\} \models \Box \phi_{n+1} \Rightarrow \psi$. By the IH, we have $\{\phi_1, \ldots, \phi_n\} \vdash \Box \phi_{n+1} \Rightarrow \psi$. By Theorem 5, we have $\{\phi_1, \ldots, \phi_{n+1}\} \vdash \psi$. ⊣

§5. Decision procedure. We have implemented a satisfiability decision procedure for LTL_f.[5] Our method is based Kröger and Merz's tableau-based decision procedure [10]. Kroğer and Merz generate tableaux where the states are PNPs; they proceed to unfold propositional and then temporal formulae while checking for *closedness*. If a certain kind of path exists in the resulting graph, then the formula is satisfiable—we can use that path to generate a Kripke structure.

The closed nodes of their tableaux are inductively defined as those which are manifestly contradictory (e.g., $\bot \in \mathsf{pos}(\mathcal{P})$ or $\mathsf{pos}(\mathcal{P}) \cap \mathsf{neg}(\mathcal{P}) \neq \emptyset$), those where all of their successors are contradictory (e.g. $\bot \mathcal{W} \bot \in \mathsf{pos}(P)$ isn't obviously contradictory, but both of its temporal successors are), and those where a negated temporal formula is never actually falsified (e.g., if $\Box \phi \in \mathsf{neg}(\mathcal{P})$ and we are generating an infinite Kripke structure, we had better falsify ϕ at some point). The third criterion is a critical one: Kröger and Merz, by default, generate infinite paths in their tableaux, which correspond to infinite Kripke structures. If they were to drop their third criterion, they would find infinite paths where, say, $\neg \Box \phi$ is meant to hold but ϕ is never falsified. Such "dishonest" infinite paths must be carefully avoided.

Our decision procedure diverges slightly from theirs. First, we generalize their always-based approach to include weak until. Next, we simplify their approach to exclude the third condition on paths. Since we deal with finite models of time, we'll never consider infinite paths—and so we avoid the issue of dishonest infinite paths wholesale.

Our simplified notion of closedness means we can implement a more efficient algorithm. While Kröger and Merz need to keep the tableau around in order to identify the "honest" strongly connected components of the tableau, we need not do so. We can perform a perfectly ordinary graph search without having to keep the whole tableau in memory. (We do have to keep the *states* of the tableau in memory, though.) To be clear: we claim no asymptotic advantage, and our algorithm remains exponential; rather, our implementation is simpler.

[5] https://github.com/ericthewry/ltlf-decide

§6. Discussion. We have studied a finite temporal logic for linear time: LTL_f. We were able to adapt techniques for *infinite* temporal logics to show deductive completeness in a finite setting. We are by no means the first to prove completeness for LTL_f, but we do so (a) in direct analogy to existing methods and (b) improving on Roşu's axioms [13]. The proof of deductive completeness calls for only minor changes to the proof with potentially infinite time: we *inject finiteness* by inserting $\Diamond\,\mathsf{end}$ into our proof graphs, allowing us to directly adapt methods from an infinite logic; injecting finiteness simplifies the selection of the path used to generate the Kripke structure in the satisfiability proof (Lemma 28 and Corollary 29). We believe that the technique is general, and will adapt to other temporal logics; we offer this proof as evidence.

To be clear, we claim that the proof of completeness for a 'finitized' logic is relatively straightforward *once you find the right axioms.* We can offer only limited guidance on finding the right axioms. Finite temporal logics should have an axiom saying that time is, indeed, finite; some sort of axiom will be needed to establish the meaning of temporal modalities at the end of time (e.g., FINITE); when porting axioms from the infinite logic, one must be careful to check that the axioms are sound at the end of time (e.g., ENDNEXTCONTRA), when temporal modalities may change in meaning (e.g., changing distribution over implication to use the weak next modality, as in WKNEXTDISTR).

Acknowledgments. The comments of anonymous FoSSaCS reviewers helped improve this work.

REFERENCES

[1] CAROLYN JANE ANDERSON, NATE FOSTER, ARJUN GUHA, JEAN-BAPTISTE JEANNIN, DEXTER KOZEN, COLE SCHLESINGER, and DAVID WALKER, *Netkat: Semantic foundations for networks*, ***Acm sigplan notices***, vol. 49, ACM, 2014, pp. 113–126.

[2] JORGE A. BAIER and SHEILA A. MCILRAITH, *Planning with first-order temporally extended goals using heuristic search*, ***Proceedings of the 21st national conference on artificial intelligence - volume 1***, AAAI'06, AAAI Press, 2006, pp. 788–795.

[3] RYAN BECKETT, MICHAEL GREENBERG, and DAVID WALKER, *Temporal netkat*, ***ACM SIGPLAN Notices SIGPLAN Not.***, vol. 51 (2016), no. 6, p. 386401.

[4] LORIS D'ANTONI and MARGUS VEANES, *Monadic second-order logic on finite sequences*, ***Proceedings of the 44th acm sigplan symposium on principles of programming languages*** (New York, NY, USA), POPL 2017, ACM, 2017, pp. 232–245.

[5] GIUSEPPE DE GIACOMO, RICCARDO DE MASELLIS, and MARCO MONTALI, *Reasoning on ltl on finite traces: Insensitivity to infiniteness.*, ***Aaai***, Citeseer, 2014, pp. 1027–1033.

[6] Giuseppe De Giacomo and Moshe Y Vardi, *Linear temporal logic and linear dynamic logic on finite traces*, ***Ijcai'13 proceedings of the twenty-third international joint conference on artificial intelligence***, Association for Computing Machinery, 2013, pp. 854–860.

[7] ———, *Synthesis for ltl and ldl on finite traces*, ***Proc. of ijcai***, 2015.

[8] Giuseppe De Giacomo and Moshe Y. Vardi, *Ltlf and ldlf synthesis under partial observability*, ***Proceedings of the twenty-fifth international joint conference on artificial intelligence***, IJCAI'16, AAAI Press, 2016, pp. 1044–1050.

[9] Dexter Kozen, *Kleene algebra with tests*, ***ACM Trans. Program. Lang. Syst.***, vol. 19 (1997), no. 3, pp. 427–443.

[10] Fred Kroger and Stephan Merz, ***Temporal logic and state systems***, Springer, 2008.

[11] Orna Lichtenstein, Amir Pnueli, and Lenore Zuck, *The glory of the past*, ***Workshop on logic of programs***, Springer, 1985, pp. 196–218.

[12] A. Pnueli, *The temporal logic of programs*, ***Foundations of computer science, 1977., 18th annual symposium on***, Oct 1977, pp. 46–57.

[13] Grigore Roşu, *Finite-trace linear temporal logic: Coinductive completeness*, ***International conference on runtime verification***, Springer, 2016, pp. 333–350.

CORNELL UNIVERSITY
ITHACA, NY, USA
E-mail: ehc86@cornell.edu

POMONA COLLEGE
CLAREMONT, CA, USA
E-mail: michael@cs.pomona.edu