

Urgency Annotations for Alternating Choices

Eren Keskin, Roland Meyer, and Sören van der Wall

TU Braunschweig, {e.keskin, roland.meyer, s.van-der-wall}@tu-bs.de

Abstract. We propose urgency programs, a new programming model with support for alternation, imperfect information, and recursion. The novelty are urgency annotations that decorate the (angelic and demonic) choice operators and control the order in which alternation is resolved. We study standard notions of contextual equivalence for urgency programs. Our first main result are fully abstract characterizations of these relations based on sound and complete axiomatizations. Our second main result settles their computability via a normal form construction. Notably, we show that the contextual preorder is $(2h - 1)$ -EXPTIME-complete for programs of maximal urgency h when the regular observable is given as an input resp. PTIME-complete when the regular observable is fixed. We designed urgency programs as a framework in which it is convenient to formulate and study verification and synthesis problems. We demonstrate this on a number of examples including the verification of concurrent and recursive programs and hyper model checking.

Keywords: Alternation, Imperfect Information, Contextual Equivalence, Full Abstraction, Axiomatization, Verification, Complexity.

1 Introduction

Algorithmic program verification may seem like a zoo of approaches that do not have much in common. Refinement is checked by establishing a relation between the program and a reference implementation [1], the model checking of linear-time properties [2] is formulated as language inclusion [3], and branching-time properties are reduced to games [4]. This variety makes it difficult to choose a verification approach when one is confronted with a new class of programs. Fortunately, there are master problems that provide guidance. When the new class of programs is concurrent, there is a good chance that the verification problem can be cast as coverability in well-structured transition systems [5,6,7]. When it is recursive, one would try to reduce it to higher-order model checking [8,9]. The master problem then informs us about how to implement the verification algorithm [10], from the symbolic representation of program states [11] to the search strategy [12]. But the master problems are no silver bullet. Well-structured transition systems are not good at modeling recursion, and the verification of branching-time properties is undecidable for them. Higher-order models, in turn, are not good at modeling concurrency.

We propose a new programming construct to capture verification tasks that are not handled well by the existing master problems. The key insight is that by

combining alternation [13] with imperfect information [14] one can model concurrency even in a sequential programming model. We combine the two by adding urgency annotations to the (angelic and demonic) choice operators. Choices are classically resolved in program order [13]. Urgency annotations are natural numbers that define when a choice has to be made: the higher the urgency, the sooner. The lower urgency choices remain unresolved in the program term until all higher urgency choices have been made. Choices of different urgency can thus be imagined as belonging to different components in the verification task, the program threads or the specification that has to judge the program behavior. To illustrate the order in which choices with urgency are made, consider the transition sequence

$$(a_l \frown a_r).(b \heartsuit c) \rightarrow (a_l \frown a_r).b \rightarrow a_l.b.$$

The demonic choice \frown of urgency 1 is resolved only after the angelic choice \heartsuit of urgency 2, although the demonic choice is written earlier in the program.

Urgency programs not only serve as a convenient backend to which to reduce verification tasks. By comparing the reductions, we also understand how the tasks are related. As a running example, we study the formulation of simulation and language inclusion as urgency programs. Let the systems of interest be \mathcal{A} and \mathcal{B} from Figure 1. We examine whether (a) \mathcal{A} is simulated by \mathcal{B} and (b) the language of \mathcal{A} is included in the language of \mathcal{B} . The transitions a_l and a_r carry the observation a . The programs p_{\lesssim} and p_{\subseteq} from Figure 1 encode the simulation resp. the inclusion problem. For now, just observe that the difference in these programs is merely a shift in urgency.




Fig. 1: Automata \mathcal{A} , \mathcal{B} and programs p_{\lesssim} , p_{\subseteq} for simulation resp. inclusion.

Urgency programs can only provide guidance for solving verification tasks if their verification problem is decidable and the decision procedure is easy to adapt to the new setting. To develop decision procedures for urgency programs, we follow a common approach: understand the contextual equivalence [15] of the programming model, then one understands how to summarize programs, and every algorithm will be a fixed point over these summaries. The notion of contextual equivalence depends on the level of detail at which one intends to observe the program behavior. We consider two standard definitions [16]:

$$\begin{aligned} p \simeq q & \quad \text{if } \forall O.\forall c[\bullet]. \quad c[p] \Downarrow O \quad \text{iff} \quad c[q] \Downarrow O \\ p \simeq_O q & \quad \text{if } \forall c[\bullet]. \quad c[p] \Downarrow O \quad \text{iff} \quad c[q] \Downarrow O. \end{aligned}$$

The former definition quantifies over observables O and contexts $c[\bullet]$ and is our notion of *contextual equivalence*. The latter fixes an observable, like termination or reaching an error, and we refer to it as *O -specialized contextual equivalence*. Due to the alternating choices in urgency programs, making an observation is defined in a game-theoretic way: $c[q] \Downarrow O$ means Eve has a winning strategy in the game arena $\llbracket c[q] \rrbracket$ when $O \subseteq \Sigma^*$ is the objective.

Our first contribution are full abstraction results: contextual equivalence and its specialized variant coincide with congruence relations that neither quantify over contexts nor observables. The congruences are defined axiomatically, and one may say we axiomatize the contextual equivalences. An important insight is that imperfect information distributes over perfect information. In the example,

$$(a_l \uparrow a_r).(b \Downarrow c) \simeq (a_l \uparrow a_r).b \Downarrow (a_l \uparrow a_r).c \simeq (a_l.b \uparrow a_r.b) \Downarrow (a_l.c \uparrow a_r.c) .$$

Our second contribution is to settle the complexity of the two contextual equivalences and their preorder variants. The main finding is that the specialized contextual preorder is $(2h - 1)$ -EXPTIME-complete for programs of maximal urgency h when the regular observable is given as an input resp. PTIME-complete when it is fixed. A consequence is indeed that the verification problem can be solved with the same complexity. To get the upper bound right, an important idea is to factorize the set of contexts. We equate contexts $c[\bullet]$ that have the same solution space: the same set of programs p so that $c[p] \Downarrow O$. The challenge is to handle the factorization algorithmically. We show how to represent solution spaces using the novel concept of characteristic terms.

For the lower bound, we reduce context-bounded multi-pushdown games [17] to the verification of urgency programs. This justifies our claim that urgency programs can capture concurrency in a rather natural way. To further demonstrate the usefulness of urgency programs, we show how to encode other popular verification problems, notably the recent hyper model checking [18,19].

Structure. After an introduction to urgency programs, we state the full abstraction results in Section 3, followed by the proofs of soundness, normalization, and completeness. In Section 7, we give the decidability and complexity results, followed by the upper and lower bound proofs in Sections 8 and 9.

2 Programming Model

Throughout the development, we fix a natural number $h > 0$ for the maximal urgency in programs, a finite alphabet Σ of terminal symbols with typical elements a, b, c , and a finite or infinite set of non-terminals N with elements A, B, C . Each non-terminal has a so-called defining program term given by $E : N \rightarrow \mathbf{T}$.

The set \mathbf{T} of *program terms* (of urgency up to h over Σ, N) is defined as

$$p ::= a \mid \text{skip} \mid \text{err} \mid A \mid p.p \mid \Psi P \mid \bigwedge P .$$

Terminal symbols represent program commands with visible behavior, `skip` is a command without visible effect, and `err` aborts the computation unsuccessfully.

Non-terminals model recursive functions, we have concatenation, and angelic (Ψ) as well as demonic (\bigwedge) choice of urgency $0 < u \leq h$. We use \oplus to mean Ψ or \bigwedge , infix notation for binary choices, and $\bigoplus p$ for $\bigoplus \{p\}$. An action is a term A or $\bigoplus P$. Terms that contain actions are called active. Terms that do not contain actions are called passive. Passive terms are words over $\Sigma \cup \{\text{skip}, \text{err}\}$, and we also call them word terms $w \in \mathbf{W}$. To avoid brackets, we let concatenation bind stronger than choices. Choices range over non-empty but possibly infinite sets of terms. Infinitary syntax requires care to make sure set and game-theoretic concepts remain sound. We moved the corresponding lemmas to Appendix A and B to keep the presentation light. The motivation for infinitary syntax will be given at the end of the section. We lift the notion of urgency from choices to program terms and define $\text{urg} : \mathbf{T} \rightarrow \mathbb{N}$ by

$$\begin{aligned} \text{urg}(\text{skip}) &= \text{urg}(\text{err}) = \text{urg}(a) = 0 & \text{urg}(A) &= h \\ \text{urg}(p.q) &= \max\{\text{urg}(p), \text{urg}(q)\} & \text{urg}(\bigoplus P) &= u . \end{aligned}$$

A *context* $c[\bullet]$ (of maximal urgency h over Σ and N) is a term that contains at most one occurrence of the fresh non-terminal \bullet . The set of all contexts is denoted by \mathbf{C} . The expression $c[p]$ refers to the term obtained from $c[\bullet]$ by replacing \bullet with the term p . A term q is a subterm of p if there is a context $c[\bullet]$ with $p = c[q]$. A subterm is called *outermost* if it is not enclosed by a choice. For example, p is outermost in $p.q$ but not in $(p \Psi q).r$. The *leading subterm* $\text{lead}(p)$ of a term p is defined as the outermost action with the highest urgency. If several outermost actions have this urgency, then the leftmost of them is leading. Passive terms do not have leading subterms. Where helpful, we will underline a subterm that contains the leading subterm. Note that $\underline{p}.q$ implies $\text{urg}(p) \geq \text{urg}(q)$ and $p.\underline{q}$ implies $\text{urg}(q) > \text{urg}(p)$. We denote the unique context enclosing the leading subterm $\text{lead}(p)$ in p by $\text{en}_p[\bullet] \in \mathbf{C}$. We have $p = \text{en}_p[\text{lead}(p)]$.

Example 1. Consider p_{\subseteq} from Figure 1. Its urgency is $\text{urg}(p_{\subseteq}) = 2$, the leading subterm is $b \Psi c$, and the enclosing context is $\text{en}_{p_{\subseteq}}[\bullet] = a.(a_l \bigwedge a_r).\bullet.(b \bigwedge c)$. Further, $a_l \bigwedge a_r$ is an outermost subterm of p_{\subseteq} and a_l is not.

2.1 Semantics

Given that our programming model has alternation, the operational semantics of a term is not a plain transition system but a *game arena* $\llbracket p \rrbracket = (\mathbf{T}, p, \text{own}, \rightarrow)$ in which positions are owned by player *Eve* or player *Adam*. The set of positions is the set of all terms. The initial position is the given term. The ownership assignment $\text{own} : \mathbf{T} \rightarrow \{\text{Eve}, \text{Adam}\}$ returns the owner of the leading subterm, $\text{own}(p) = \text{own}(\text{lead}(p))$. Adam owns the demonic choices, $\text{own}(\bigwedge P) = \text{Adam}$, and Eve owns the angelic choices, $\text{own}(\Psi P) = \text{Eve}$. We also give skip , err , terminals, and non-terminals to Eve. This has no influence on the semantics as there will be at most one move from these positions. The set of moves is defined as the smallest relation satisfying the following rules:

$$\frac{p \in P}{\bigoplus P \rightarrow p} \qquad \frac{}{A \rightarrow E(A)} \qquad \frac{\text{lead}(p) \rightarrow q}{p \rightarrow \text{en}_p[q]} .$$

A move always rewrites the leading subterm as illustrated by Figure 2. For a choice, it selects one alternative. For a non-terminal, it inserts the defining term. We define $\text{succ}(p) = \{q \mid p \rightarrow q\}$. It is worth noting that the game arena has perfect information. Imperfect information is modeled through choices, and they are eventually resolved.

The operational semantics is intensional in that it gives precise information about the program state at runtime. From a programming perspective, what matters is the result of a computation or, more generally, the *observable behavior* of the program. Due to Adam's influence, the observable behavior will rarely be a single word but rather a language $O \subseteq \Sigma^*$. We write $p \Downarrow O$ to mean that Eve can enforce termination and the result will be a word from O , no matter how Adam plays. We make this precise.

Our notion of observable behavior is based on concepts from game theory. We refer to a language $O \subseteq \Sigma^*$ as a reachability objective for the game arena $\llbracket p \rrbracket$. A play in this arena is a maximal (finite or infinite) sequence of positions $\pi = p_0, p_1, \dots$ that starts in the given term, $p_0 = p$, and respects the moves of the game arena, $p_i \rightarrow p_{i+1}$ for all i . If the play ends, the result is a word term $w \in \mathbf{W}$. We interpret it as an element of the monoid with zero $(\Sigma^* \cup \{\text{err}\}, \cdot, \text{skip}, \text{err})$. Here, skip is the unit, often denoted by ε , and err is the zero. We use $\stackrel{\cong}{\approx}$ to denote the monoid equality. Eve wins the play when w belongs to O , meaning there is $v \in O$ so that $w \stackrel{\cong}{\approx} v$ (write $w \in O$). Otherwise, Adam wins the play. In particular, Adam wins all infinite plays and all plays exhibiting $\text{err} \notin \Sigma$.

A positional strategy for Eve is a function $\sigma : \mathbf{T} \rightarrow \mathbf{T}$ so that $q \rightarrow \sigma(q)$ for all terms q owned by Eve that admit further rewriting. Since we are interested in reachability objectives, we can use positional strategies without loss of generality [20]. A play π is conform to σ if for all i with $\text{own}(p_i) = \text{Eve}$ and $\text{succ}(p_i) \neq \emptyset$ we have $\sigma(p_i) = p_{i+1}$. Eve wins objective O , if she has a strategy σ so as to win all plays that are conform to this strategy. This is what we denote by $p \Downarrow O$.

Example 2. We demonstrate how to encode simulation and inclusion between \mathcal{A} and \mathcal{B} from Figure 1. In p_{\prec} , Eve is tasked to find a violation of the simulation property, while Adam tries to prove it: whenever Eve takes a transition in \mathcal{A} , Adam tries to select a simulating transition right away in \mathcal{B} . The term p_{\subseteq} models inclusion: Eve selects a run in \mathcal{A} and Adam tries to come up with a run in \mathcal{B} with matching observations. We define an objective O so that $p_{\prec} \Downarrow O$ holds precisely when \mathcal{A} is simulated by \mathcal{B} and $p_{\subseteq} \Downarrow O$ holds precisely when \mathcal{A} 's language is included in \mathcal{B} 's. The objective has to make sure that the choices of transitions form paths in the automata and that the letters actually match. The latter is easy to check. The former can be guaranteed by interleaving $O_{\mathcal{A}}$ and $O_{\mathcal{B}}$ in an alternating fashion, so $a_1.a_2$ and $b_1.b_2$ yield $a_1.b_1.a_2.b_2$. Details are in Section 9.

Example 3. The encodings for simulation and inclusion only differ in the terms $(a_l \hat{\wedge} a_r).(b \check{\vee} c)$ and $(a_l \hat{\wedge} a_r).(b \check{\vee} c)$. Their semantics is given in Figure 2. Rectangular nodes are owned by Adam, circular ones by Eve, and for rectangular nodes with rounded corners the ownership does not matter. The objective is $O = \{a_l.c, a_r.b\}$. Indeed, \mathcal{A} 's language is included in \mathcal{B} 's, $(a_l \hat{\wedge} a_r).(b \check{\vee} c) \Downarrow O$, but \mathcal{A} is not simulated by \mathcal{B} , $(a_l \hat{\wedge} a_r).(b \check{\vee} c) \not\Downarrow O$.




Fig. 2: The game arenas $\llbracket (a_l \check{\vee} a_r).(b \hat{\wedge} c) \rrbracket$ and $\llbracket (a_l \check{\vee} a_r).(b \hat{\wedge}_2 c) \rrbracket$. O is gray.

2.2 Contextual Preorder

The notion of observable behavior is not compositional: we may have $p \Downarrow O$ if and only if $q \Downarrow O$ for all objectives O , yet the two terms behave differently when placed into a context. In our example, $(b \check{\vee} c) \Downarrow U$ if and only if $(b \check{\vee}_2 c) \Downarrow U$ for all objectives $U \subseteq \Sigma^*$. When inserting the terms into the context $c[\bullet] = (a_l \hat{\wedge} a_r).\bullet$, however, we have the difference discussed above. This is a common problem, and the way out is to consider the largest congruence that lives inside observational equivalence. It is more elegant to work with a precongruence and define the congruence of interest as a derived notion.

Definition 1. The contextual preorder $\preceq \subseteq \mathbf{T} \times \mathbf{T}$ is defined by $p \preceq q$, if

$$\forall O \subseteq \Sigma^*. \forall c[\bullet] \in \mathbf{C}. c[p] \Downarrow O \text{ implies } c[q] \Downarrow O.$$

The O -specialized contextual preorder $\preceq_O \subseteq \mathbf{T} \times \mathbf{T}$ is defined by fixing O and dropping the leading universal quantifier. The contextual equivalence is then $\simeq = \preceq \cap \succeq$, and the O -specialized contextual equivalence is $\simeq_O = \preceq_O \cap \succeq_O$.

In the example, $(b \check{\vee}_2 c) \preceq (b \check{\vee} c)$ and so by congruence $(a_l \hat{\wedge} a_r).(b \check{\vee}_2 c) \preceq (a_l \hat{\wedge} a_r).(b \check{\vee} c)$. The reverse does not hold, consider context $(a_l \hat{\wedge} a_r).\bullet$ and objective $\{a_l.c, a_r.b\}$ from above. Note that $p \preceq q$ implies $p \preceq_O q$ for all O .

With $O_{Term} = \Sigma^*$ as the objective, we can use the specialized contextual equivalence to study the termination behavior of programs. We can also introduce a letter loc so that $p \Downarrow O_{Reach}$ with $O_{Reach} = \Sigma^*.\text{loc}.\Sigma^*$ observes visits to a specific location. (Non-specialized) Contextual equivalence is more precise and takes into account all objectives. Both notions are also motivated by verification, where contextual equivalence gives information about which information can be abstracted away from an urgency term without an influence on the objective, similar to how bisimilarity preserves CTL* properties [21].

The motivation for an infinite set of non-terminals and infinitary terms is to model parameterized functions in a simple yet general way. The idea is to introduce a non-terminal for each instantiation of the function's formal parameters by actual values, inspired by value passing in process algebra [22]. Another more technical argument is that our normal form relies on infinitary syntax.

3 Full Abstraction

We define a precongruence $\sqsubseteq \subseteq \mathbf{T} \times \mathbf{T}$ on program terms that neither quantifies over contexts nor objectives but relates terms solely based on their syntactic

structure. The relation is defined through a set of axioms that should be understood as explaining the interplay between the operators in our programming model. The main finding is that this axiomatic precongruence coincides with the contextual preorder, and we say that we axiomatize (in a sound and complete way) the contextual preorder. This is our main theorem.

Theorem 1 (Full Abstraction 1). $p \sqsubseteq q$ if and only if $p \preceq q$.

We have a corresponding result for the O -specialized contextual preorder. In this setting, a complete axiomatization is considerably more difficult to obtain because, intuitively, we have to understand the concatenation behavior of language O . Our solution is partial in that we impose a side condition on the objective to obtain completeness: it should be right-separating, a notion we will define in a moment. Luckily, the objectives of interest O_{Term} and O_{Reach} are right-separating. It is always sound to reason with the O -specialized axiomatic precongruence.

Theorem 2 (Full Abstraction 2). $p \sqsubseteq_O q$ implies $p \preceq_O q$. If O is right-separating, then also $p \preceq_O q$ implies $p \sqsubseteq_O q$.

We understand the concatenation behavior of an objective with the help of the syntactic precongruence over the monoid $\Sigma^* \cup \{\text{err}\}$. It may relate terminal words to err in case they cannot be extended to a word from the objective.

Definition 2. The syntactic precongruence induced by O on $\Sigma^* \cup \{\text{err}\}$ is defined by $w \leq_O^s v$, if for all $x, y \in \Sigma^*$ we have $x.w.y \in O$ implies $x.v.y \in O$.

An objective is then right-separating, if the concatenation from left in the above definition is not needed to distinguish words. We define \leq_O^r on $\Sigma^* \cup \{\text{err}\}$ by $w \leq_O^r v$, if for all $y \in \Sigma^*$ we have $w.y \in O$ implies $v.y \in O$.

Definition 3 (and Lemma). Objective O is right-separating, if $\leq_O^s = \leq_O^r$. The objectives O_{Term} and O_{Reach} are right-separating.

Example 4. The syntactic congruence $\cong_O^s = \leq_O^s \cap \geq_O^s$ induced by the objective $O = \{a_l.c, a_r.b\}$ has classes $\Sigma^* \cup \{\text{err}\} / \cong_O^s = \{\text{[skip]}, [a_l], [a_r], [b], [c], [a_l.c, a_r.b]\}$ plus a class for the remaining words.

Intuitively, right-separating objectives allow us to evaluate the O -specialized contextual preorder by using contexts $\bullet.r$ that only append to the right. For arbitrary objectives, we have to consider contexts $s.\bullet.r$ and it is difficult to understand the interplay between high urgencies in r and low urgencies in s .

We now give the two axiomatizations and explain them on an intuitive level. Recall that a precongruence is a reflexive and transitive relation that is preserved when inserting related terms into the same context.

Definition 4. The axiomatic precongruence $\sqsubseteq \subseteq \mathbf{T} \times \mathbf{T}$ is the least precongruence satisfying the axioms in Figure 3 except (S). The O -specialized axiomatic precongruence \sqsubseteq_O on terms is the least precongruence satisfying all axioms in Figure 3. We use \equiv for $\sqsubseteq \cap \supseteq$ and \equiv_O for $\sqsubseteq_O \cap \supseteq_O$.

$$\begin{array}{ll}
\text{(a) Lattice} & \text{(b) Distributivity} \\
\text{(L1)} \frac{\forall i \in I. p_i \sqsubseteq q_i}{\bigoplus \{p_i \mid i \in I\} \sqsubseteq \bigoplus \{q_i \mid i \in I\}} & \text{(D1)} \frac{\text{urg}(p) < u}{p.(\bigoplus Q) \equiv \bigoplus \{p.q \mid q \in Q\}} \\
\text{(L2)} \frac{\forall_{i \in I} \bigwedge P_i \equiv \bigwedge_{f: I \rightarrow P_I} \bigvee \{f(i) \mid i \in I\}}{\bigvee_{i \in I} \bigwedge P_i \equiv \bigwedge_{f: I \rightarrow P_I} \bigvee \{f(i) \mid i \in I\}} & \text{(D2)} \frac{\text{urg}(p) \leq u}{(\bigoplus Q).p \equiv \bigoplus \{q.p \mid q \in Q\}} \\
\text{(L3)} \frac{\text{urg}(p) \leq u}{p \bigwedge (p \bigvee q) \equiv p \quad p \bigvee (p \bigwedge q) \equiv p} & \text{(c) Normalization} \\
\text{(L4)} \frac{}{\bigoplus_{i \in I} \bigoplus P_i \equiv \bigoplus \bigcup_{i \in I} P_i} & \text{(L5)} \frac{\text{urg}(p) \leq u}{p \sqsubseteq p \bigvee q} \quad \text{(d) Err} \quad \text{(e) Monoid} \\
\text{(FP)} \frac{}{A \equiv E(A)} & \text{(B1)} \frac{}{\text{err} \sqsubseteq p} \quad \text{(M)} \frac{w \stackrel{\cong}{=} v}{w \equiv v} \\
\text{(LFP)} \frac{\forall A \in N. E(A)\{N/p_N\} \sqsubseteq p_A}{B \sqsubseteq p_B} & \text{(g) Specialization} \\
\text{(S)} \frac{w \leq_O^s v}{w \sqsubseteq_O v}
\end{array}$$

Fig. 3: Axioms defining \sqsubseteq and \sqsubseteq_O .

With the axioms given in Figure 3a, the choice operators span a completely distributive lattice on each urgency. The monotonicity axiom (L1) is not covered by the precongruence but implements an infinite replacement. The axiom has a side condition that can be found Appendix A. Due to this axiom, nodes in our proof trees may have an infinite degree. Yet, every path is guaranteed to be finite. To see the premise in axiom (L5), consider $p = (a_l \bigwedge a_r)$ and context $\bullet.(b \bigvee c)$. Then in $p.(b \bigvee c)$ Eve wins while in $(p \bigvee p).(b \bigvee c)$ she loses, similar to Figure 2. For (L3), the reasoning is similar. As a consequence of the lattice axioms, one can derive the dual rules of (L2) and (L5). Distributivity (L2) states that the order of choices can be changed by considering all choice functions $f : I \rightarrow \bigcup_{i \in I} P_i$ with $f(i) \in P_i$ for all $i \in I$, denoted by $f : I \rightarrow P_I$.

The distributivity in (D1) captures the essence of imperfect information: concatenation to the left distributes over choice, provided the internals of the term are invisible as the choice has a higher urgency. The distributivity from the right in (D2) is similar but takes into account that the leading subterm for equal urgencies is leftmost. This clean interplay between imperfect information and choice came as a surprise and we consider these laws an important contribution.

A string with $\text{err} \notin \Sigma$ is the most disadvantageous term for Eve, because it belongs to no objective $O \subseteq \Sigma^*$.

The monoid axiom (M) refers to word terms $w, v \in \mathbf{W}$. We interpret them in the monoid $\Sigma^* \cup \{\text{err}\}$ and inherit the equality there, denoted by $\stackrel{\cong}{=}$ above. The equality strips brackets and skip, and interprets err as zero.

The normalization axiom (N) reflects the fact that only the outermost choice operator determines the urgency of a term. Towards soundness, note that once

the outer choice with urgency u is resolved, we are sure that the context to the left has urgency strictly smaller than u and the context to the right has urgency at most u . Hence, the inner choice is the next to be resolved, independent of whether its urgency is u or $v \geq u$.

The fixed-point axiom (FP) allows us to rewrite non-terminals to their defining terms. The axiom (LFP) allows us to rewrite non-terminals to a prefixed point, using Knaster and Tarski's characterization of least fixed points [23]. Here, we let p_N denote a vector of terms with one entry p_A per non-terminal $A \in N$, and use $\{N/p_N\}$ for the substitution of all non-terminals by these terms.

Recall that Axiom (S) only plays a role in the definition of the specialized axiomatic precongruence. The axiom depends on the objective O of interest, meaning it actually is a family of axioms. The axiom relates word terms $w, v \in \mathbf{W}$ as prescribed by the syntactic precongruence.

Example 5. We show that simulation implies inclusion in Example 2. Remember that Eve tries to refute the relation. We prove axiomatically

$$b \varphi c \stackrel{(L5)}{\sqsubseteq} (b \psi c) \varphi (b \psi c) = \varphi (b \psi c) \stackrel{(L3), (L4)}{\equiv} \bigwedge \varphi (b \psi c) \stackrel{(L3)}{\equiv} b \psi c.$$

By (L5), we have $b \sqsubseteq b \psi c$ and $c \sqsubseteq b \psi c$, which we can apply to subterms by congruence. The equality holds because choices range over sets. We apply (L3) with $p = q = \varphi(b \psi c)$, apply (L4) to flatten the choices, apply (L3) once more to remove the choices, and finally apply congruence. This yields $p \sqsubseteq p \approx$. With soundness of the axiomatization (Section 4), we obtain the desired $p \sqsubseteq p \approx$. With a normal form result (Section 5), we also show completeness (Section 6).

4 Soundness

Proposition 1 (Soundness). $p \sqsubseteq q$ implies $p \preceq q$, $p \sqsubseteq_O q$ implies $p \preceq_O q$.

This section is devoted to proving Proposition 1. Proving soundness is difficult because it requires us to reason over all contexts. If $p \sqsubseteq_O q$ is an axiom, then we need to show that $c[p] \Downarrow O$ implies $c[q] \Downarrow O$ for all $c[\bullet] \in \mathbf{C}$.

We first develop a proof technique for soundness that allows us to reduce the set of contexts we have to consider, and in a second step prove the axioms sound. To define the contexts that have to be considered, we introduce some terminology. We say that term p is *immediate* for context $c[\bullet] \in \mathbf{C}$, if $c[p] \in \mathbf{W}$ or $c[p]$ is active and p contains the leading subterm in $c[p]$, denoted by $c[\underline{p}]$. If this is not the case, we call p *paused* for $c[\bullet]$. As the names suggest, immediate terms get rewritten in the next move while paused terms do not. For example, term φP is immediate for $\bullet \bigwedge Q$ but paused for $\bullet \bigwedge Q$ and $\varphi\{\bullet, q\}$. In the last context, \bullet is enclosed by a choice. No term is immediate for such a context.

Lemma 1 (Proof Technique). *If $c[p] \Downarrow O$ implies $c[q] \Downarrow O$ for all contexts $c[\bullet] \in \mathbf{C}$ where at least one of p or q is immediate, then $p \preceq_O q$.*

Lemma 1 defines precisely the contexts we need to consider when proving the axioms sound. Its proof relies on the following observation: a paused term does not change the outcome of a move in the context (cf. Appendix C).

Lemma 2. *Consider terms p and q that are paused for context $c[\bullet]$. Then, $\text{own}(c[p]) = \text{own}(c[q])$ and there is a set of contexts $D \subseteq \mathbf{C}$ so that $\text{succ}(c[p]) = \{d[p] \mid d[\bullet] \in D\}$ and $\text{succ}(c[q]) = \{d[q] \mid d[\bullet] \in D\}$.*

We are now prepared to prove each axiom sound. Using Lemma 1, all proofs share a common approach: we fix an objective and pick a context that is immediate for at least one term in the axiom. Then, we unroll the game arena for a few moves until it reveals the winning implication we are after. We restrict our presentation to the soundness of Axiom (D1) and defer details to Appendix D. The proofs make use of two properties of immediate terms.

Lemma 3. *Let term p be immediate for context $c[\bullet]$. Then we have $\text{own}(c[p]) = \text{own}(p)$. If $\text{urg}(p) \leq \text{urg}(q)$, then also term q is immediate for $c[\bullet]$.*

Proof of Proposition 1 (Soundness). **Axiom (D1):** We consider binary choices owned by Eve, the generalization to choices over arbitrary sets and also Adam's case are similar. For $\text{urg}(p) < u$, the axiom says that $p.(q \Psi r) \sqsubseteq p.q \Psi p.r$ and vice versa. The goal is thus to show $p.(q \Psi r) \simeq p.q \Psi p.r$.

Consider an objective O and let $c[\bullet]$ be a context for which at least one of $p.(q \Psi r)$ or $p.q \Psi p.r$ is immediate. The urgencies of both terms are u . This means that not only one but actually both terms are immediate for $c[\bullet]$ and, moreover, the owner of $c[p.(q \Psi r)]$ and $c[p.q \Psi p.r]$ is Eve. The first moves in the game arenas are thus done by the same player and have the same result:



As a consequence, translating winning strategies becomes straightforward and the equivalence $c[p.(q \Psi r)] \Downarrow O$ iff $c[p.q \Psi p.r] \Downarrow O$ holds. \square

5 Normalization

As a first step towards completeness, we show that each term can be brought into a normal form using our axioms. Normalization is a standard approach in completeness proofs. The treatment of alternation and urgency is new.

The normal form eliminates non-terminals and orders the interplay between concatenation and choice: a normal form term is a tree of height $2h$ (with h the maximal urgency) that repeatedly alternates between Eve's and Adam's choices while decreasing the urgency. The leaves of the tree are terminal words, skip, or err. Inductively, we define $NF_0 = \Sigma^* \cup \{\text{err}\}$ and for $u > 0$,

$$ANF_u = \{\bigwedge P \mid \emptyset \neq P \subseteq NF_{u-1}\} \quad NF_u = \{\Psi P \mid \emptyset \neq P \subseteq ANF_u\}.$$

The base case terms are all owned by Eve. In an ANF_u term, Adam chooses over NF_{u-1} terms. In an NF_u term, Eve chooses over such ANF_u terms owned by Adam. The main result of this section is the following.

Proposition 2. *There is a function $\text{nf} : \mathbf{T} \rightarrow NF_h$ so that $\text{nf}(p) \equiv p$ for all p .*

Example 6. We illustrate the normal form computation on our running example. Let $h = 2$ and consider $(a_l \uparrow a_r).(b \Downarrow c)$.

$$\begin{aligned}
 (a_l \uparrow a_r).(b \Downarrow c) &\stackrel{\text{(D1)}}{\equiv} (a_l \uparrow a_r).b \Downarrow (a_l \uparrow a_r).c \\
 &\stackrel{\text{(D2)}}{\equiv} (a_l.b \uparrow a_r.b) \Downarrow (a_l.c \uparrow a_r.c) \\
 &\stackrel{\text{(L3)}, \text{(L4)}}{\equiv} [\Delta \Downarrow (a_l.b \uparrow a_r.b)] \Downarrow [\Delta \Downarrow (a_l.c \uparrow a_r.c)].
 \end{aligned}$$

Formally, we prove Proposition 2 in two steps. First, we compute a normal form term $\text{nf}(A) \equiv A$ for every non-terminal $A \in N$ in a fixed-point iteration. Then, we replace each non-terminal in p by its normal form. Finally, we normalize the remaining term to obtain $\text{nf}(p) \equiv p$. The two steps are formalized in the next lemma, details of the normalization can be found in Appendix E.

Lemma 4. *For all $p \in \mathbf{T}$ without non-terminals there is $\text{nf}(p) \in NF_h$ with $p \equiv \text{nf}(p)$. For all non-terminals $A \in N$, there is a $\text{nf}(A) \in NF_h$ with $\text{nf}(A) \equiv A$.*

6 Completeness

Proposition 3 (Completeness). *$p \preceq q$ implies $p \sqsubseteq q$. For a right-separating objective $O \subseteq \Sigma^*$, $p \preceq_O q$ implies $p \sqsubseteq_O q$.*

Interestingly, the completeness for the axiomatic precongruence follows from the completeness for the specialized version. To prove this implication, we consider objectives for which Axiom (S) does not add relations. Formally, we call an objective O *domain shattering*, if $\sqsubseteq = \sqsubseteq_O$.

Lemma 5. *The language $w.w^{\text{reverse}}$ is domain-shattering and right-separating.*

The argument for completeness is this. Let $O \subseteq \Sigma^*$ be the domain-shattering and right-separating objective from Lemma 5. We have

$$\preceq \stackrel{\text{Definition}}{\subseteq} \preceq_O \stackrel{\text{Prop. 3}}{\subseteq} \sqsubseteq_O \stackrel{\text{Shattering}}{=} \sqsubseteq \stackrel{\text{Soundness, Prop. 1}}{\subseteq} \preceq.$$

6.1 Completeness Proof, Specialized Case

To show Proposition 3, fix a right-separating $O \subseteq \Sigma^*$. With the normalization in Proposition 2 and soundness of the axiomatic precongruence in Proposition 1, it is sufficient to show completeness for terms in normal form. For $p, q \in NF_h$, we want to show that $p \not\sqsubseteq_O q$ implies $p \not\preceq q$ by giving a context that tells them

apart. This, however, is difficult as it requires us to understand precisely when the axiomatic congruence fails.

Our way out is to define a less flexible preorder that is easier to handle. The *domination preorder* \subseteq_O only relates normal form terms of the same urgency and owned by the same player. It is defined by induction on the urgency. For $w, v \in NF_0$, we have $w \subseteq_O v$ if $w \leq_O^s v$. For $u > 0$:

$$\bigwedge_u P \subseteq_O \bigwedge_u Q \text{ if } \forall s \in Q. \exists r \in P. r \subseteq_O s \quad \bigvee P \subseteq_O \bigvee Q \text{ if } \forall r \in P. \exists s \in Q. r \subseteq_O s$$

Proposition 3 holds with the following lemma, which we prove in Appendix F.

Lemma 6. *For a right-separating objective O , $p \preceq_O q$ implies $p \subseteq_O q$ and $p \subseteq_O q$ implies $p \preceq_O q$.*

7 Decidability and Complexity

We study the decidability and complexity of checking the contextual preorder and its specialized variant. To this end, we leave the setting of infinitary syntax and call p *finitary*, if it refers to a finite set of non-terminals (N, E) and all defining terms $E(A)$ and p itself are finite. We make the decision problem parametric in the relation R to be checked, and instantiate R with \simeq , \simeq_O , and \preceq_O :

h-DEC- R

Given: Finitary p, q over $\Sigma, (N, E)$ of urgency h .

Problem: Does $p R q$ hold?

The first finding is that already the contextual equivalence is undecidable. The proof is by a reduction from the equivalence problem for context-free languages and the result continues to hold if we fix an alphabet with at least two letters. A consequence is that the specialized contextual equivalence for domain-shattering objectives is also undecidable.

Proposition 4. *h-DEC- \simeq and h-DEC- \simeq_O with O domain-shattering are undecidable for every h .*

Recall that the language $w.w^{reverse}$ is domain-shattering, so already context-free objectives lead to undecidability.

Our main result in this section is that for regular objectives the specialized contextual preorder is decidable. We can also give the precise complexity, for which we measure the size of the input in the expected way as $|p| + |q| + |\Sigma| + |E|$. The size of the defining equations is $|E| = \sum_{A \in N} 1 + |E(A)|$. The size of a term is $|\text{skip}| = |\text{err}| = |a| = |A| = 1$, $|p.q| = 1 + |p| + |q|$, and $|\bigcirc^u P| = 1 + \sum_{p \in P} |p|$.

Theorem 3. *Let h be an urgency. For every regular objective $O \neq \emptyset$, the problem h-DEC- \preceq_O is PTIME-complete.*

It is worth noting that the result does not expect the objective to be right-separating. We can indeed decide the specialized contextual preorder \preceq_O for all regular objectives O . Moreover, the lower bound holds no matter the objective.

It is natural to define a variant $\text{h-DEC-}\preceq_*$ of the problem in which also the objective is part of the input and given as a deterministic finite automaton $(\Sigma, Q, i, \delta, \mathcal{F})$. In this case, we use $|O|$ to refer to $|\Sigma| + |Q|$. The following result shows the dramatic influence that the objective has on the complexity.

Theorem 4. $\text{h-DEC-}\preceq_*$ is $(2h - 1)$ – EXPTIME-complete.

A consequence is that we can also solve the problem of making an observation, denoted as $\text{h-DEC-}\Downarrow$ and defined with almost the same input as $\text{h-DEC-}\preceq_*$.

Corollary 1. $\text{h-DEC-}\Downarrow$ is $(2h - 1)$ – EXPTIME-complete.

8 Upper Bound

We prove the upper bounds claimed in Theorems 3 and 4 as follows.

Proposition 5. Given finitary terms p and q and a regular objective O as a DFA, deciding $p \preceq_O q$ can be done in time $(|p| + |q| + |E||N|) \cdot \exp_{2h-1}(\mathcal{O}(|O|^2))$.

The undecidability result in Proposition 4 shows that the normal form for the axiomatic congruence from Section 5 is insufficient as a basis for algorithms. The problem is that the normal form terms are typically infinite, and therefore difficult to handle computationally. The source of infinity can be found in the base case: already $NF_0 = \Sigma^* \cup \{\text{err}\}$ is infinite, and this propagates upwards. We realize that the O -specialized axiomatic congruence admits a more refined normal form that is guaranteed to yield finite terms (and finitely many of them). The key idea is to factorize NF_0 using Axiom (S).

We define the set of O -specialized normal form terms by induction on the urgency. The base case $SNF_0^O = \Sigma_{\text{err}}^* / \cong_O^s$ are classes of words in the syntactic congruence $\cong_O^s = \leq_O^s \cap \geq_O^s$ induced by O . For $u > 0$, the definition is

$$SANF_u^O = \{\bigwedge P \mid \emptyset \neq P \subseteq SNF_{u-1}^O\} \quad SNF_u^O = \{\bigvee P \mid \emptyset \neq P \subseteq SANF_u^O\}.$$

Note that since O is regular, the set SNF_0^O and so all $SANF_u^O$ and SNF_u^O are guaranteed to be finite [24]. Another aspect is that we change the alphabet to having \cong_O^s -congruence classes as letters. This can be fixed by working with a representative system: we represent every \cong_O^s -class by one of its elements.

We adapt the normalization process from Section 5 to compute a term in the O -specialized normal form. Only the base case changes, for the inductive cases we merely study the complexity. Interestingly, the overall normalization takes time $2h$ -fold exponential only in the size of the objective, using the common definition $\exp_0(x) = x$ and $\exp_{u+1}(x) = 2^{\exp_u(x)}$.

Lemma 7. Given a finitary term p and a regular objective $O \subseteq \Sigma^*$ as a DFA, we can compute $\text{nf}_O(p) \in SNF_h^O$ with $\text{nf}_O(p) \equiv_O p$ in time $(|p| + |E||N|) \cdot \exp_{2h-1}(\mathcal{O}(|SNF_0^O|))$. We have $|SNF_h^O| = \exp_{2h}(\mathcal{O}(|SNF_0^O|))$.

The result already allows us to decide the O -specialized contextual preorder as follows. Since we cannot assume the objective to be right-separating, the algorithm cannot rely on a full abstraction result. Instead, we have to evaluate $p \preceq_O q$ directly, by iterating over contexts. What makes this possible is the combination of our proof technique for soundness in Lemma 1 and the O -specialized normal form just introduced. With Lemma 1, we do not have to iterate over all contexts to show $p \preceq_O q$, but only over contexts of the form $r. \bullet .s$. With Lemma 7, the terms r and s can be normalized.

Corollary 2. *Let O be regular and p, q finitary. Then $p \preceq_O q$ iff for all $c[\bullet] = r. \bullet .s$ with $r \in SNF_{h-1}^O$, $s \in SNF_h^O$ we have $\text{nf}_O(c[p]) \Downarrow O$ implies $\text{nf}_O(c[q]) \Downarrow O$.*

The algorithm formulated in the corollary is slower than the promised upper bound by two exponents because SNF_h^O contains $\exp_{2h+1}(\mathcal{O}(|O|^2))$ many terms. To overcome the problem, the first step is to reduce the number of contexts that have to be considered. The idea is to factorize the contexts along their solution spaces. The solution space of a context is the set of terms p for which $c[p] \Downarrow O$ holds. When checking for $p \preceq_O q$, the job of a context $c[\bullet]$ is to disprove $p \preceq_O q$ by showing $c[p] \Downarrow O$ and $c[q] \not\Downarrow O$. Hence, when two contexts have the same solution space, it suffices to consider one of them.

What makes the solution space equivalence algorithmically interesting is that (i) it is coarse enough to save an exponent and (ii) we can directly compute with equivalence classes of contexts. The key insight behind both statements is that the solution space of a context can be represented in a convenient way: it is the \sqsubseteq_O -upward closure of a so-called *characteristic term*.

Definition 5. *Term p is characteristic for $c[\bullet]$ wrt. $O \subseteq \Sigma^*$, if for all $q \in \mathbf{T}$ we have $c[q] \Downarrow O$ if and only if $p \sqsubseteq_O q$.*

We will show that there are only $\exp_{2h}(\mathcal{O}(|O|^2))$ many characteristic terms for contexts of the form $r. \bullet .s$. Moreover, we can compute the characteristic terms directly, without building up the corresponding contexts. This saves an exponent in the complexity: we modify the algorithm in Corollary 2 to iterate through characteristic terms rather than contexts.

We save another exponent in the runtime of our algorithm by a more compact representation of the terms in O -specialized normal form. One exponent in the size of SNF_h^O is inherited from the base case, where already $|SNF_0^O| = \exp_1(\mathcal{O}(|O|^2))$. We use the fact that each class of the syntactic congruence can be represented by a function $Q \rightarrow Q$ between the states of the objective DFA [24]. The key idea is to see these functions as sets of state-pairs. We simulate a function, say for letter a , by letting Eve choose a state change (p, q) with $\delta(p, a) = q$. We thus represent the congruence class $[a] \in \Sigma_{\text{err}}^* / \cong_O^s$ by an angelic choice of urgency 1 over letters from the alphabet $Q \times Q$. The modification propagates to the normal form terms, and requires a simple modification of the objective. The new objective $\text{tr}(O)$ has a syntactic congruence with an exponentially smaller index, namely $|SNF_0^{\text{tr}(O)}| = \mathcal{O}(|O|^2)$.

Details on the compact term representation and on how to utilize characteristic terms are given in Appendix K.

9 Lower Bound and Applications

We show how to reduce well-known program verification tasks to the problem of making an observation. The first reduction justifies our complexity lower bound.

Multi-pushdown systems (MPDS). MPDS are a model for concurrent programs with recursive threads that synchronize through a shared, finite memory [25,26]. MPDS are Turing complete [27]. For bug hunting, it is therefore common to under-approximate their behavior by limiting the number of context switches [25]. A context is a sequence of transitions that operate only on one stack, and a context switch is the moment when this stack changes. We consider multi-pushdown games (MPDG), the alternating variant of MPDS. A *b-context-bounded 2-stack pushdown game (b-2PDG)* is an MPDG with two stacks.

Theorem 5. [17,26] *b-2PDG are $(b - 2) - \text{EXPTIME}$ -complete.*

Our lower bound is due to the following result that we prove in Appendix P.

Proposition 6. *Given a $(2h + 1)$ -2PDG PD, we can compute in polynomial time an alphabet Σ , a term p and non-terminals (N, E) over Σ all of maximal urgency h , and an objective DFA O so that Eve wins PD if and only if $p \Downarrow O$.*

Hyper model checking (HMC). Hyperproperties [18] are program properties that refer to sets of computations, as opposed to the classical linear-time properties [2] that refer to single computations. The corresponding logics [19] express hyperproperties by quantifying over the computations of the program \mathcal{K} , as in

$$\text{HP} = \forall w_1 \in L(\mathcal{K}). \exists w_2 \in L(\mathcal{K}) \dots \forall w_n \in L(\mathcal{K}). w_1 \oplus \dots \oplus w_n \in L_{\text{HP}} .$$

We use \oplus to denote the convolution of words. The convolution expects words of the same length (that is then chosen by the first quantifier) and forms their letter-by-letter product, meaning L_{HP} is a language over an alphabet Σ^n . The hyper model checking problem (HMC) asks whether a given program satisfies a given hyperproperty, denoted by $\mathcal{K} \models \text{HP}$. The standard way of solving HMC is to compose the program with L_{HP} , and then repeatedly with itself [28,29]. We show how to encode HMC into the problem of making an observation for urgency terms. To be precise, HMC is usually defined over infinite words, but we consider a finite-word version. The details are in Appendix M.

Theorem 6. *Given \mathcal{K} and HP, we can compute in polynomial time a term $p_{\mathcal{K}, \text{HP}}$ and an objective O so that $\mathcal{K} \models \text{HP}$ if and only if $p_{\mathcal{K}, \text{HP}} \Downarrow O$.*

Interestingly, our reduction to urgency terms even yields decidability results for HMC over recursive programs. As HMC generalizes inclusion, which is undecidable for context-free languages, the recursive case needs a restriction. We can work with a restriction similarly liberal as in the recent positive results [30,31], which form the decidability frontier for HMC.

Simulation, Inclusion, and Games with imperfect information. As explained in the introduction, urgency terms can express inclusion and simulation problems. We now give the details of the encoding, and also add imperfect information games to the picture. Let $\mathcal{A} = (Q_{\mathcal{A}}, \delta_{\mathcal{A}}, i_{\mathcal{A}}, \mathcal{F}_{\mathcal{A}})$ and $\mathcal{B} = (Q_{\mathcal{B}}, \delta_{\mathcal{B}}, i_{\mathcal{B}}, \mathcal{F}_{\mathcal{B}})$ be non-deterministic finite automata. Let p, q range over $Q_{\mathcal{A}}$ and r, s, t over $Q_{\mathcal{B}}$ with $s \neq t$. Imperfect information games (\mathcal{A}, hd) (for a definition, refer to [14]) are played on a finite automaton \mathcal{A} and hide the concrete states $p \in Q_{\mathcal{A}}$ from Adam and only present him an abstraction h , using a function $\text{hd} : Q \rightarrow H$. Such a game is played in rounds in which Adam sees the abstraction $h \in H$ of the current state p , chooses an action $a \in \Sigma$, and then Eve selects a transition from p for that action. Eve wins G when she can force the play from the initial state $i_{\mathcal{A}}$ into a final state in $\mathcal{F}_{\mathcal{A}}$.

We give the defining equations for the non-terminals:

$$\begin{aligned} A_p^{\subseteq} &= \forall \left\{ \left(\bigwedge_{(r,a,s) \in \delta_{\mathcal{B}}} (r, a, s) \right) \cdot A_q^{\subseteq} \mid (p, a, q) \in \delta_{\mathcal{A}} \right\} \cup \{\text{skip} \mid p \in \mathcal{F}_{\mathcal{A}}\} \\ A_p^{\supseteq} &= \forall \left\{ \left(\bigwedge_{(r,a,s) \in \delta_{\mathcal{B}}} (r, a, s) \right) \cdot A_q^{\supseteq} \mid (p, a, q) \in \delta_{\mathcal{A}} \right\} \cup \{\text{skip} \mid p \in \mathcal{F}_{\mathcal{A}}\} \\ A_h &= \bigwedge_{a \in \Sigma} \forall \left\{ A_{h'} \cdot \bigvee_{(p,a,q) \in \delta_{\mathcal{A}}, \text{hd}(q)=h'} (p, a, q) \mid h' \in H \right\} \cup \{\text{skip} \mid p \in \mathcal{F}_{\mathcal{A}}\} \end{aligned}$$

The only difference between the encodings of inclusion and simulation is the urgency of the angelic choice. In the case of inclusion, Adam only has to make a demonic choice when Eve has chosen the entire word. For simulation, the players take turns. For imperfect information games, Adam chooses the letter and Eve chooses the abstraction of the next state that will be made visible to Adam and the actual transition that is consistent with the letter and the abstraction. The point is that Adam should see the abstraction but not see the transition. We achieve this with a left-linear grammar where Eve's choice for the next transition will only be revealed when Adam has made all choices. By using urgency two, we could have avoided this trick and used a right-linear model. In all cases, the task of the objective O is to check that we have successive transitions.

Theorem 7. $A_{i_{\mathcal{A}}}^{\subseteq} \Downarrow O$ if and only if $\mathcal{L}(\mathcal{A}) \not\subseteq \mathcal{L}(\mathcal{B})$, $A_{i_{\mathcal{A}}}^{\supseteq} \Downarrow O$ if and only if $\mathcal{A} \not\preceq \mathcal{B}$, $A_{\text{hd}(i_{\mathcal{A}})} \Downarrow O$ if and only if Eve wins (\mathcal{A}, hd) .

It would strengthen the usefulness of urgency terms if the above reductions gave optimal complexity upper bounds. With what we have, this is true only for imperfect information games, which are EXPTIME-complete [14]. The inclusion problem is PSPACE-complete, but Corollary 1 only yields 3-EXPTIME upper bound. Simulation is PTIME-complete, but Corollary 1 yields an EXPTIME upper bound. The problem with simulation can be fixed with a different encoding that is also natural but does not show the similarity to language inclusion. There, we use the fact that the objective is fixed and a consequence of Theorem 3 that is similar to Corollary 1. In the following, we close the gap for language inclusion.

We identify natural subclasses of urgency terms and study their verification problems. In the term constructed for inclusion, for every urgency there is a

single player that makes a choice. This is not true for simulation nor for imperfect information games. We call this fragment the *weak urgency terms*. Linear grammars as used in $A_{i_A}^{\subseteq}$, $A_{i_A}^{\approx}$, and $p_{\mathcal{K}, \text{HP}}$ further reduce the complexity. The definition of weak and linear terms can be found in Appendix O.

Theorem 8. $\text{h-DEC-}\Downarrow$ for weak, linear terms is in $(\text{h} - 2)$ -EXPSPACE.

This yields the desired tight upper bound not only for inclusion checking, but also for the hyper model checking of non-recursive programs [32].

Procedure summaries. Procedure summaries [33,34,35] are a standard technique for the analysis of recursive programs. We now show how to derive procedure summaries from our axiomatization. The benefit is that the axiomatization guides the development of summaries, and thus reduces the moment of ingenuity required to come up with them.

Summaries are a closed-form representation of the call-return behavior of a recursive procedure. For functions, such a closed-form representation would capture the input-output behavior. In the context of linear-time model checking, the call-return behavior is the state change that the procedure may induce on the observer automaton. Procedure summaries are typically computed as a fixed point over a suitable domain. We now rediscover these domains as the normal forms for urgency terms. Our development covers the game version as well.

Pushdown games [34,35] extend recursive programs by alternation. We can assume reachability as the winning condition because parity can be reduced to reachability in polynomial time [36]. Let $\mathcal{P} = (Q_E, Q_A, \Sigma, \Gamma, \delta, \mathcal{F})$ be a pushdown game with transitions $\delta = (\delta_{\text{int}}, \delta_{\text{push}}, \delta_{\text{pop}})$. For convenience, we assume that final states in \mathcal{F} do not have outgoing transitions, but only loops that empty the stack. We solve two verification tasks: (a) whether Eve can force the play into a final state (b) while also producing a trace of observations from Σ accepted by a DFA \mathcal{A} . We model both with the same urgency grammar. The alphabet are again state changes (p, a, q) , this time of the pushdown, and the objective O checks whether they form a path. We define non-terminals A_x^p that express the fact that the current procedure is x and the current state is p . In the definition, \bigcirc is \bigvee if Eve owns the state, and \bigwedge if Adam owns the state, $a \in \Sigma$ is a letter, and $x, y \in \Gamma$ are stack symbols:

$$\begin{aligned} A_x^p &= \bigcirc \{ (p, a, q).A_x^q \mid (p, a, q) \in \delta_{\text{int}} \} & A_x &= \bigvee \{ A_x^p \mid p \in Q \} \\ &\cup \{ (p, a, q).A_y^q.A_x \mid (p, a, q, y) \in \delta_{\text{push}} \} \\ &\cup \{ (p, a, q) \mid (p, a, q, y) \in \delta_{\text{pop}} \} . \end{aligned}$$

We observe that the normal form $\text{nf}(A_x^p) = \bigvee_{i \in I} \bigwedge R_i$ is an Eve choice over a set I followed by an Adam choice over sets $R_i \subseteq Q \times \Sigma_{\text{err}}^* / \cong_{\mathcal{L}(\mathcal{A})}^s \times Q$. By definition, every element of R_i has p as the first component. Each $i \in I$ corresponds to a strategy σ that Eve can follow from state p with the top of stack being x until this symbol is popped. An element $(p, [w], q) \in R_i$ thus corresponds to a play through the pushdown according to σ that ends in state q after x has been popped, and

the play produces an observation $v \in [w]$. If $\mathcal{L}(\mathcal{A}) = \Sigma^*$, the middle part of our summaries becomes irrelevant. This is (up to parity tracking) the domain that Walukiewicz uses to solve pushdown games in EXPTIME [34], which matches our complexity. For other DFA, we obtain the summaries for pushdown games with inclusion objectives [35]. To be precise, we obtain an EXPTIME rather than a 2-EXPTIME upper bound because we assume the objective to be deterministic.

10 Related Work

Urgency annotations are related to priorities in process algebra [37]. The key difference is that priorities preserve the program order while urgencies do not. This out-of-order execution brings a new form of unbounded memory to the semantics (the unresolved choices) that our axiomatization explains how to handle. Also alternation [38] is not common in process algebra. Indeed, we have not found a study of angelic and demonic choice from the perspective of contextual equivalence. It is the special case of our work when the urgency is $h = 1$.

Our axiomatization is related to Salomaa’s work on language equivalence [39] and Milner’s on bisimilarity [40]. The former reference is particularly close: our Axioms (D1) and (D2) generalize Salomaa’s Axioms (A4) and (A5) to higher urgencies. Also our context lemma for proving soundness has relatives in process algebra [41]. Characteristic objects (here, contexts and terms) have appeared as early as [42]. Our contribution is to adapt the general idea to our setting.

We showed that urgency programs can capture hyper model checking over regular languages. Hyper model checking remains decidable if one language is context-free and the others are regular [43]. We can model this fragment with urgency programs as well. Urgency programs combine recursion and imperfect information. This is interesting as the canonical pushdown games with imperfect information [44] are known to be undecidable even under strong restrictions [45].

The goal of effective denotational semantics [46,47] is to solve verification problems with tailor-made denotational semantics. Given a specification O , the task is to derive a denotational semantics \mathcal{D}_O such that $\mathcal{D}_O(p)$ answers the question of whether program p satisfies specification O . Satisfying a specification coincides with our notion of making an observation, $p \Downarrow O$. The link is close. In Appendix Q.1, we show how our axiomatization of the specialized contextual preorder induces a denotational semantics that is effective for finitary programs. This shows that the semantic domain can be obtained in a systematic way. The landmark result in this field is an effective denotational semantics that captures the higher-order model-checking problem [48]. It would be interesting to consider it from the axiomatic point of view.

11 Conclusion and Future Work

We presented urgency annotations for alternating choices as a new programming construct and studied the standard notions of contextual equivalence for the resulting programs. We gave sound and complete axiomatizations and settled the

complexity. Our findings can be used to obtain new algorithms for verification and synthesis tasks, as we demonstrated on examples.

The next step is to extend urgency programs to infinite words. The challenge is to find the right semantics. If we define the semantics via infinite unrollings, then it is unclear how to ever switch to choices with lower urgency. Instead, it seems appropriate to work with programs that contain an ω -operator and can be rewritten a finite number of times. This, however, calls for a different set of algebraic techniques [49].

References

1. R. Milner, “An algebraic definition of simulation between programs,” in *IJCAI*. Kaufmann, 1971, pp. 481–489.
2. A. Pnueli, “The temporal logic of programs,” in *FOCS*. IEEE, 1977, pp. 46–57.
3. M. Y. Vardi and P. Wolper, “An automata-theoretic approach to automatic program verification,” in *LICS*. IEEE, 1986, pp. 332–344.
4. E. A. Emerson and C. S. Jutla, “Tree automata, mu-calculus and determinacy,” in *FOCS*. IEEE, 1991, pp. 368–377.
5. A. Finkel, “A generalization of the procedure of Karp and Miller to well structured transition systems,” in *ICALP*, ser. LNCS, vol. 267. Springer, 1987, pp. 499–508.
6. P. A. Abdulla and B. Jonsson, “Verifying programs with unreliable channels,” in *LICS*. IEEE, 1993, pp. 160–170.
7. A. Finkel and P. Schnoebelen, “Well-structured transition systems everywhere!” *TCS*, vol. 256, no. 1-2, pp. 63–92, 2001.
8. C. L. Ong, “On model-checking trees generated by higher-order recursion schemes,” in *LICS*. IEEE, 2006, pp. 81–90.
9. N. Kobayashi and C. L. Ong, “A type system equivalent to the modal mu-calculus model checking of higher-order recursion schemes,” in *LICS*. IEEE, 2009, pp. 179–188.
10. N. Kobayashi, R. Sato, and H. Unno, “Predicate abstraction and CEGAR for higher-order model checking,” in *PLDI*. ACM, 2011, pp. 222–233.
11. A. Finkel and J. Goubault-Larrecq, “Forward analysis for WSTS, part I: completions,” in *STACS*, ser. LIPIcs, vol. 3. Dagstuhl, 2009, pp. 433–444.
12. J. Kloos, R. Majumdar, F. Niksic, and R. Piskac, “Incremental, inductive coverability,” in *CAV*, ser. LNCS, vol. 8044. Springer, 2013, pp. 158–173.
13. A. K. Chandra, D. Kozen, and L. J. Stockmeyer, “Alternation,” *JACM*, vol. 28, no. 1, pp. 114–133, 1981.
14. J. H. Reif, “The complexity of two-player games of incomplete information,” *JCSS*, vol. 29, no. 2, pp. 274–301, 1984.
15. R. Milner, “Fully abstract models of typed λ -calculi,” *TCS*, vol. 4, no. 1, pp. 1–22, 1977.
16. G. Jaber and A. S. Murawski, “Complete trace models of state and control,” in *ESOP*, ser. LNCS, vol. 12648. Springer, 2021, pp. 348–374.
17. A. Seth, “Games on multi-stack pushdown systems,” in *LFCS*, ser. LNCS, vol. 5407. Springer, 2009, pp. 395–408.
18. M. R. Clarkson and F. B. Schneider, “Hyperproperties,” *JCS*, vol. 18, no. 6, pp. 1157–1210, 2010.
19. M. R. Clarkson, B. Finkbeiner, M. Koleini, K. K. Micinski, M. N. Rabe, and C. Sánchez, “Temporal logics for hyperproperties,” in *POST*, ser. LNCS, vol. 8414. Springer, 2014, pp. 265–284.
20. D. A. Martin, “Borel determinacy,” *AMATH*, vol. 102, no. 2, pp. 363–371, 1975.
21. M. Hennessy and R. Milner, “Algebraic laws for nondeterminism and concurrency,” *JACM*, vol. 32, no. 1, pp. 137–161, 1985.
22. J. A. Bergstra, A. Ponse, and S. A. Smolka, Eds., *Handbook of Process Algebra*. Elsevier, 2001.
23. G. Birkhoff, *Lattice Theory*. AMS, 1967.
24. M. O. Rabin and D. S. Scott, “Finite automata and their decision problems,” *IBM J. Res. Dev.*, vol. 3, no. 2, pp. 114–125, 1959.

25. S. Qadeer and J. Rehof, "Context-bounded model checking of concurrent software," in *TACAS*, ser. LNCS, vol. 3440. Springer, 2005, pp. 93–107.
26. R. Meyer and S. van der Wall, "On the complexity of multi-pushdown games," in *FSTTCS*, ser. LIPIcs, vol. 182. Dagstuhl, 2020, pp. 52:1–52:35.
27. G. Ramalingam, "Context-sensitive synchronization-sensitive analysis is undecidable," *TOPLAS*, vol. 22, no. 2, pp. 416–430, 2000.
28. G. Barthe, P. R. D'Argenio, and T. Rezk, "Secure information flow by self-composition," in *CSFW*. IEEE, 2004, pp. 100–114.
29. B. Finkbeiner, M. N. Rabe, and C. Sánchez, "Algorithms for model checking HyperLTL and HyperCTL*," in *CAV*, ser. LNCS, vol. 9206. Springer, 2015, pp. 30–48.
30. J. O. Gutsfeld, M. Müller-Olm, and C. Ohrem, "Deciding asynchronous hyperproperties for recursive programs," *CoRR*, vol. abs/2201.12859, 2022.
31. A. Bajwa, M. Zhang, R. Chadha, and M. Viswanathan, "Stack-aware hyperproperties," in *TACAS*, ser. LNCS, vol. 13993. Springer, 2023, pp. 308–325.
32. M. N. Rabe, "A temporal logic approach to information-flow control," Ph.D. dissertation, Saarland University, 2016.
33. M. Sharir and A. Pnueli, *Two approaches to interprocedural data flow analysis*. New York, NY: New York Univ. Comput. Sci. Dept., 1978.
34. I. Walukiewicz, "Pushdown processes: Games and model-checking," *IC*, vol. 164, no. 2, pp. 234–263, 2001.
35. L. Holík, R. Meyer, and S. Muskalla, "Summaries for context-free games," in *FSTTCS*, ser. LIPIcs, vol. 65. Dagstuhl, 2016, pp. 41:1–41:16.
36. M. Hague, R. Meyer, S. Muskalla, and M. Zimmermann, "Parity to safety in polynomial time for pushdown and collapsible pushdown systems," in *MFCS*, ser. LIPIcs, vol. 117. Dagstuhl, 2018, pp. 57:1–57:15.
37. R. Cleaveland, G. Lüttgen, and V. Natarajan, "Priority in process algebra," in *Handbook of Process Algebra*. Elsevier, 2001, pp. 711–765.
38. A. K. Chandra and L. J. Stockmeyer, "Alternation," in *FOCS*. IEEE, 1976, pp. 98–108.
39. A. Salomaa, "Two complete axiom systems for the algebra of regular events," *JACM*, vol. 13, no. 1, pp. 158–169, 1966.
40. R. Milner, "A complete inference system for a class of regular behaviours," *JCSS*, vol. 28, no. 3, pp. 439–466, 1984.
41. A. Pitts and I. Stark, "Operational reasoning for functions with local state," in *Higher Order Operational Techniques in Semantics*. CUP, 1998, pp. 227–273.
42. R. McNaughton and S. A. Papert, *Counter-Free Automata*. MIT Press, 1971.
43. A. Pommellet and T. Touili, "Model-checking HyperLTL for pushdown systems," in *SPIN*, ser. LNCS, vol. 10869. Springer, 2018, pp. 133–152.
44. B. Aminof, A. Legay, A. Murano, O. Serre, and M. Y. Vardi, "Pushdown module checking with imperfect information," *IC*, vol. 223, pp. 1–17, 2013.
45. L. Bozzelli, "New results on pushdown module checking with imperfect information," in *GandALF*, ser. EPTCS, vol. 54, 2011, pp. 162–177.
46. K. Aehlig, "A finite semantics of simply-typed lambda terms for infinite runs of automata," *LMCS*, vol. 3, no. 3, 2007.
47. K. Terui, "Semantic evaluation, intersection types and complexity of simply typed lambda calculus," in *RTA*, ser. LIPIcs, vol. 15. Dagstuhl, 2012, pp. 323–338.
48. S. Salvati and I. Walukiewicz, "A model for behavioural properties of higher-order programs," in *CSL*, ser. LIPIcs, vol. 41. Dagstuhl, 2015, pp. 229–243.
49. T. Wilke, "An algebraic theory for regular languages of finite and infinite words," *Algebra and Computation*, vol. 3, no. 4, pp. 447–489, 1993.

50. L. M. Goldschlager, “The monotone and planar circuit value problems are log space complete for P,” *ACM SIGACT News*, vol. 9, no. 2, pp. 25–29, 1977.
51. O. M. Committee. Openssl. [Online]. Available: [openssl.org](https://www.openssl.org)
52. P. C. Kocher, “Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems,” in *CRYPTO*, ser. LNCS, vol. 1109. Springer, 1996, pp. 104–113.
53. D. Brumley and D. Boneh, “Remote timing attacks are practical,” in *USENIX Security*. USENIX Association, 2003.
54. W. J. Savitch, “Relationships between nondeterministic and deterministic tape complexities,” *Journal of Computer and System Sciences*, vol. 4, no. 2, pp. 177–192, 1970.

A Handling Infinitary Syntax

For the development of the section, we call a term an **h**-term when no choice subterm has urgency higher than **h** and the expansion of each non-terminal is an **h**-term. Infinitary syntax requires care to make sure that set theoretic concepts remain sound. The problem lies on the unbounded nature of the choice operator. If we were to allow choices to range over *arbitrary sets* of terms, the class of all terms would no longer form a set. To see this, suppose that the class of terms with unbounded choice were a set \mathbf{T} . Then, we could build the term $\bigvee \mathbf{T} \in \mathbf{T}$, which contradicts the Axiom of Regularity. But for our development in the paper, it is a convenience to have \mathbf{T} as a set and not be distracted by subtle differences of classes and sets. For this reason, we impose a restriction on the terms, that does not hinder our developments. It is defined over the structural depth of a term, which is

$$\begin{aligned} \text{depth}(A) &= \text{depth}(a) = \text{depth}(\text{err}) = \text{depth}(\text{skip}) = 0 \\ \text{depth}(p.q) &= \max\{\text{depth}(p) + 1, \text{depth}(q) + 1\} \\ \text{depth}(\bigoplus P) &= \sup\{\text{depth}(p') + 1 \mid p' \in P\}. \end{aligned}$$

The notion of structural depth lifts naturally to defining assignments (N, E) . We let $\text{depth}(E)$ be the smallest limit ordinal strictly greater than $\text{depth}(E(A))$ for all $A \in N$. So, whenever the paper mentions the set of all terms \mathbf{T} , we refer, in fact, to restriction $\mathbf{T}_E = \{p \mid \text{depth}(p) < \text{depth}(E)\}$. The set of permitted contexts \mathbf{C}_E is defined in the same way, as contexts are defined as terms built from $(N \uplus \{\bullet\}, E)$. We let $\mathbf{C}_E = \{c[\bullet] \mid \text{depth}(c[\bullet]) < \text{depth}(E)\}$, and drop the subscript from this set as well. This restriction indeed does not hinder our development. All definitions we apply to terms can be expressed as context free replacements and the set \mathbf{T}_E is closed under such replacements. The only exceptions build Axioms (L1) and (LFP), for which we present side conditions so to stay in the restricted set \mathbf{T}_E .

Lemma 8. *Let (N, E) be a defining assignment to **h**-terms. For all $p \in \mathbf{T}_E$ and all $c[\bullet] \in \mathbf{C}_E$, $c[p] \in \mathbf{T}_E$.*

Proof. Proof is by an induction on the structure of $c[\bullet]$. Let $p \in \mathbf{T}_E$ and let $\alpha = \text{depth}(E)$. Note that $\text{depth}(p) < \alpha$. The case $c[\bullet] = p \in \mathbf{T}_E$ is clear, since $c[\bullet] = c[p]$. For the case $c[\bullet] = \bullet$, $\text{depth}(c[p]) = \text{depth}(p) < \alpha$.

For the concatenative inductive case, let $c[\bullet] = q.d[\bullet]$. The case $c[\bullet] = d[\bullet].q$ is analogous. Since $\text{depth}(c[\bullet]) < \alpha$, we also have $\text{depth}(d[\bullet]) < \alpha$ and $\text{depth}(q) < \alpha$. Applying the induction hypothesis yields $\text{depth}(d[p]) < \alpha$. And since α is a limit ordinal, also $\text{depth}(q) + 1 < \alpha$ and $\text{depth}(c[p]) + 1 < \alpha$.

For the choice inductive case, let $c[\bullet] = \bigoplus\{d[\bullet]\} \cup Q$. We have $\text{depth}(c[p]) = \sup(\{\text{depth}(d[p]) + 1\} \cup \{\text{depth}(q) \mid q \in Q\})$. This is equal to the maximum of $\text{depth}(d[p]) + 1$ and $\sup\{\text{depth}(r) + 1 \mid r \in Q\}$. We know that $\text{depth}(c[\bullet]) < \alpha$. So, $\sup\{\text{depth}(r) + 1 \mid r \in Q\} \leq \text{depth}(c[\bullet]) < \alpha$. Per definition, we also have $\text{depth}(d[\bullet]) < \alpha$. We apply the induction hypothesis to get $\text{depth}(d[p]) < \alpha$ and due to α being a limit ordinal, $\text{depth}(c[p]) < \alpha$. \square

An important implication of Lemma 8 is that the successor relation is well defined for \mathbf{T}_E . All rewriting has the form $c[\bigoplus P] \rightarrow c[p]$ with $p \in P$ or $c[A] \rightarrow c[E(A)]$ for some context $c[\bullet]$ with $\text{depth}(c[\bullet]) < \text{depth}(E)$. Then, the successors have $\text{depth}(c[p]) < \text{depth}(c[\bigoplus P]) < \text{depth}(E)$ for all $p \in P$, and $\text{depth}(E(A)) < \text{depth}(E)$ for all $A \in N$.

A.1 Axiom Side Conditions on Depth

The restriction to \mathbf{T}_E also restricts the axiom system to terms only in \mathbf{T}_E . This affects Axioms (L1) and (FP).

For Axiom (L1), note that $\bigoplus \mathbf{T}_E \not\subseteq \mathbf{T}_E$.

$$(L1) \frac{\forall i \in I. p_i \sqsubseteq q_i}{\bigoplus \{p_i \mid i \in I\} \sqsubseteq \bigoplus \{q_i \mid i \in I\}}$$

So to keep terms in \mathbf{T}_E , we impose the following side condition to Axiom (L1): $\sup\{\text{depth}(p_i) + 1 \mid i \in I\} < \text{depth}(E)$ and $\sup\{\text{depth}(q_i) + 1 \mid i \in I\} < \text{depth}(E)$. This results in $\bigoplus \{p_i \mid i \in I\} \in \mathbf{T}_E$ and $\bigoplus \{q_i \mid i \in I\} \in \mathbf{T}_E$.

For Axiom (FP), note that \mathbf{T}_E is closed under finite substitutions, but not under infinite substitutions. However, N is infinite, so for arbitrary p_N , it is not guaranteed that $E(A)\{N/p_N\} \in \mathbf{T}_E$.

$$(LFP) \frac{\forall A \in N. E(A)\{N/p_N\} \sqsubseteq p_A}{B \sqsubseteq p_B}$$

We need to ensure that the substituted term belongs to \mathbf{T}_E . Thus, we require $\text{depth}(E(A)\{N/p_N\}) < \text{depth}(E)$ for all $A \in N$. In fact, the requirement is already implicitly stated by the axiom. The precondition requires $E(A)\{N/p_N\} \sqsubseteq p_B$ and $\sqsubseteq \subseteq \mathbf{T}_E \times \mathbf{T}_E$. So, $\text{depth}(E(A)\{N/p_N\}) < \text{depth}(E)$ is already required by the axiom implicitly.

B Strategy Tree Bounds

Let $G = (V, v, \text{own}, E)$ be a game arena with reachability objective $O \subset V$ and $\sigma : V \rightarrow V$ be a winning strategy for Eve from starting position $v \in V$. Consider the subgraph T of (V, E) reachable from v , where for all $v \in V$ owned by Eve only the successor $\sigma(v)$ is part of T . Since σ is a winning strategy, all paths from v in T must reach the objective in finitely many steps, i.e. T is a tree and every branch is finite. We prove in a second, that this is sufficient to obtain an ordinal α to bound the depth of T . If Eve has a strategy tree with depth α , we say that Eve wins in α turns.

The existence of uniform positional strategies results in an important property. If Eve wins from $v \in V$ in α turns, then there must be a $w \in E(v)$ from which Eve wins in $\gamma < \beta$ turns, if $\text{own}(v) = \text{Eve}$. If $\text{own}(v) = \text{Adam}$, Eve wins from all $w \in E(v)$ in $\alpha_w < \alpha$ turns.

It remains to prove that the absence of an infinite branch is sufficient to obtain α . Let $T = (V, E)$ be a directed tree and α_V be the smallest ordinal with $|V| \leq |\alpha_V|$.

Definition 6. A function $\text{depth} : V \rightarrow \alpha_V$ is called *depth assignment* when $\text{depth}(v) < \text{depth}(w)$ for all $(v, w) \in E$.

Note that $E(v) = \emptyset$ implies $\text{depth}(v) = \sup(\emptyset) = 0$.

Lemma 9. If a directed tree $T = (V, E)$ has no infinite path then T has a depth assignment $\text{depth} : V \rightarrow \alpha_V$.

Proof. We prove the contraposition. Let $T = (V, E)$ be a directed tree with root $r \in V$. We use $T_v = (V_v, E_v)$ for the subtree rooted in $v \in V$. The key insight is that for all $v \in V$, where T_v has no depth assignment, there must be $w \in E_v(v)$ where also T_w has no depth assignment. To see the validity of this statement, suppose the existence of $v \in V$ where T_v has no depth assignments, while all $w \in E(v)$, T_w have depth assignments depth_w . Note that all V_w are disjoint. Then $\text{depth} : V_v \rightarrow \alpha_{V_v}$ is a depth assignment, where $\text{depth}(u) = \text{depth}_w(u)$ if $u \in V_w$, and $\text{depth}(v) = \sup\{\text{depth}(w) + 1 \mid w \in E_v(v)\}$. The depth property is satisfied for all $u \in V_w$ and also for $v \in V_v$. So $\text{depth} : V_v \rightarrow \alpha_{|V_v|}$ is a depth assignment, which contradicts the assumption.

Let $T = (V, E)$ have no depth assignment. We inductively construct an infinite sequence of nodes (v_0, v_1, \dots) with $v_i, v_{i+1} \in E$ and so that T_{v_i} has no depth assignment. The root is $v_0 = r$ and to extend (v_0, \dots, v_n) , where T_{v_n} has no depth assignment, we choose any v_{n+1} such that $T_{v_{n+1}}$ has no depth assignment either, which we have shown to exist. \square

C Soundness: Proof Technique

Proof of Lemma 1. Assume Eve wins O from $c[p]$. Then she does so in at most β -many moves, where β is an ordinal that is guaranteed to exist by results in Appendix B. To be clear, even for transfinite β , Eve wins each play after a finite number of moves. The ordinal β limits the size of the game arena reachable from $c[p]$ when she plays according to her strategy. We show $c[q] \Downarrow O$ by transfinite induction on β . The base case is simple, yet instructive. If $\beta = 0$ then $c[p] \in O$, meaning $c[p] \in \mathbf{W}$. Then p is immediate for context $c[\bullet]$, and the premise yields $c[q] \Downarrow O$.

In the inductive case, it will make no difference whether β is a limit ordinal or a successor ordinal, so we will not distinguish the two. We have $c[p] \Downarrow O$. If p or q is immediate for $c[\bullet]$, then the premise of the lemma already tells us $c[q] \Downarrow O$. Therefore, assume both terms are paused for $c[\bullet]$. Intuitively, we will see that Eve can copy her strategy from $c[p]$ to $c[q]$ (until the inserted term becomes immediate). Since both terms are paused for $c[\bullet]$, we can apply Lemma 2. It shows that, after insertion, the owner is the same, $\text{own}(c[p]) = \text{own}(c[q])$, and also gives a set of contexts $D \subseteq \mathbf{C}$ capturing the successors. Let $\text{own}(c[p]) = \text{own}(c[q]) = \text{Eve}$. Then, there must be a context $d[\bullet] \in D$ so that Eve wins $d[p]$

in $\beta' < \beta$ moves. By the induction hypothesis, $d[q] \Downarrow O$. Moreover, Eve can play $c[q] \rightarrow d[q]$ and win. If $\text{own}(c[p]) = \text{own}(c[q]) = \text{Adam}$, then for all $d[\bullet] \in D$, Eve must win $d[p]$ in $\beta' < \beta$ turns. By the induction hypothesis, we get $d[q] \Downarrow O$ for all $d[\bullet] \in D$. This is exactly $\text{succ}(c[q])$, so we have $c[q] \Downarrow O$ as well. \square

D Soundness: Missing Axiom Proofs

Proofs of the remaining axioms. In all of the soundness proofs, we conclude with Lemma 1 to generalize from contexts for which at least one side of the conclusion is immediate, to all contexts. To avoid repetition, this conclusion is omitted.

Most axioms have a direct proof. Only Axioms (L1) and (LFP) require a simultaneous induction on the proof structure. The induction is kept implicit.

Axiom (LFP): The proof makes heavy use of substitution. For the sake of readability, we use q^N for $q\{N/p_N\}$. We also extend this notation to contexts and write $d^N[\bullet]$ for the context obtained from $d[\bullet] \in \mathbf{C}$ by replacing all non-terminals $A \in N$ (different from \bullet) by p_A .

Assuming $E(A)^N \sqsubseteq p_A$ holds for all $A \in N$, the axiom yields $B \sqsubseteq p_B$. We proceed by an (outer) induction on the ordinal height of proof trees. The induction hypothesis yields $E(A)^N \preceq p_A$ for all $A \in N$. We have to show $B \preceq p_B$. Consider an objective O and a context $d[\bullet]$ for which p_B or B is immediate. We prove that $d[B] \Downarrow O$ implies $d[p_B] \Downarrow O$ with a detour. If Eve wins O from B , she does so in β -many moves, with β an ordinal, Appendix B. We apply transfinite induction on β to establish the following more general statement. For all contexts $c[\bullet]$ and all terms p , if Eve wins O from $c[p]$ in β moves, then she wins O from $c[p^N]$. Letting $c[\bullet] = d[\bullet]$ and $p = B$ gives us the desired conclusion.

The base case $\beta = 0$ is trivial: the term $c[p]$ must be a word term and hence $c[p] = c[p^N]$. Before moving on with the inductive step, we make a remark. Since the urgency of non-terminals is maximal, we have $\text{urg}(A) \geq \text{urg}(p)$ for all non-terminals A and all terms p . Moreover, the urgency of a term is monotonic in the urgency of its subterms, so in particular $\text{urg}(q) \geq \text{urg}(q^N)$ holds.

In the inductive step, let p be a term and $c[\bullet]$ be a context so that Eve wins O from $c[p]$ in β moves. We first consider the case that p is paused for $c[\bullet]$. Since $\text{urg}(p) \geq \text{urg}(p^N)$, Lemma 3 tells us that p^N is also paused for $c[\bullet]$. Similar to the proof of Lemma 1, we can use the induction hypothesis to argue that Eve wins O from $c[p^N]$. It is worth noting that in the paused case the move will change the surrounding context, which is why we strengthened the inductive statement to universally quantify over contexts.

Assume p is immediate for $c[\bullet]$. Let $q = \text{lead}(p)$ be the leading subterm and recall that $p = \text{en}_p[q]$, the term can be written as the unique context enclosing the leading subterm with the leading subterm inserted. The substitution distributes to all subterms and we also have $p^N = \text{en}(p)^N[q^N]$. Showing that Eve wins from $c[p^N]$ thus means to show that she wins from $c[\text{en}(p)^N[q^N]]$. Since q is leading in p , it is a choice or a non-terminal. We begin with the choice, $q = \textcircled{u}Q$. We argue that q^N must be immediate for $c[\text{en}(p)^N[\bullet]]$. To see this, note that p is immediate for $c[\bullet]$ and so q is immediate for $c[\text{en}_p[\bullet]]$. The substitution

distributes over the choice and we have $(\bigoplus Q)^N = \bigoplus \{r^N \mid r \in Q\}$. This shows $\text{urg}(q) = \text{urg}(q^N)$. For the outermost actions in $c[\text{en}_p[\bullet]]$, the substitution can only lower the urgencies.


The fact that q^N is immediate for $c[\text{en}(p)^N[\bullet]]$ yields

$$\text{succ}(c[\text{en}(p)^N[\underline{q^N}]])) = \{c[\text{en}(p)^N[r^N] \mid q \rightarrow r\} = \{c[s^N] \mid p \rightarrow s\}.$$

We also have $\text{succ}(c[\underline{p}]) = \{c[s] \mid p \rightarrow s\}$. Similar to the proof of Lemma 1, Eve wins from $c[s]$ in $\beta' < \beta$ turns for all/one $s \in \text{succ}(p)$, depending on the owner of $c[p]$. For every successor s , the induction hypothesis tells us that $c[s] \Downarrow O$ implies $c[s^N] \Downarrow O$. Since the owner of both $c[p]$ and $c[p^N]$ is the owner of the choice \bigcirc , Eve can copy her strategy.

It remains to consider the case that q is a non-terminal A . Then, $c[\text{en}_p[A]]$ can only be played into $c[\text{en}_p[E(A)]]$ and Eve wins from this position in $\beta' < \beta$ moves. By the hypothesis, $c[\text{en}_p[E(A)]^N]$ is also won by Eve. We can write this term as $c[\text{en}(p)^N[E(A)^N]]$. The hypothesis of the outer induction yields $E(A)^N \preceq p_A$. Therefore, Eve must also win from $c[\text{en}(p)^N[p_A]] = c[\text{en}_p[A]^N] = c[p^N]$.

Axiom (N): We only show the case $\forall \bigoplus Q \simeq \forall \bigvee Q$ with $v < u$. Like in the previous proof, both terms are immediate for the context $c[\bullet]$ of interest. The key is to note that the inner choice cannot be resolved until the outer choice has been made. The game arenas are:



As before, translation of strategies is straightforward.

Axiom (B1): To show $\perp \preceq p$, note that Eve never wins from a term $c[\perp]$ and hence the implication is trivial.

Axiom (L1): Assume $p_i \sqsubseteq q_i$ for all $i \in I$ for which also $p_i \preceq q_i$ holds. Further let $\bigoplus \{p_i \mid i \in I\}$ and $\bigoplus \{q_i \mid i \in I\}$ be valid terms (i.e. they satisfy the depth constraints from Appendix A). Acquire $c[\bullet]$ for which one of these terms are immediate. Since $\text{urg}(\bigoplus \{p_i \mid i \in I\}) = \text{urg}(\bigoplus \{q_i \mid i \in I\}) = u$, Lemma 3 tells us that both terms are immediate for $c[\bullet]$. Fix an objective $O \subseteq \Sigma^*$. Let $c[\bigoplus \{p_i \mid i \in I\}] \Downarrow O$. We have that $\text{succ}(c[\bigoplus \{p_i \mid i \in I\}]) = \{c[p_i] \mid i \in I\}$. Then, for some $i \in I$ ($\bigcirc = \vee$) [for all $i \in I$ ($\bigcirc = \wedge$)] holds $c[p_i] \Downarrow O$. Since $p_i \preceq q_i$, also $c[q_i] \Downarrow O$. Thus, $c[\bigoplus \{q_i \mid i \in I\}] \Downarrow O$.

Axiom (L4): Let $\bigoplus_{i \in I} \bigoplus P_i$ and $\bigoplus \bigcup_{i \in I} P_i$ be terms. Let $c[\bullet]$ be a context where one, and by Lemma 3 both, terms are immediate. After one move from $c[\bigoplus_{i \in I} \bigoplus P_i]$, the resulting term is always of the form $c[\bigoplus P_i]$ for some $i \in I$. Since $\text{urg}(\bigoplus P_i) = u$, Lemma 3 states that this term is also immediate for $c[\bullet]$. Then $\text{succ}(c[\bigoplus P_i]) = \{c[p] \mid p \in P_i\}$. This position is owned by the same player that owns the initial position, $c[\bigoplus_{i \in I} \bigoplus P_i]$ and the position $c[\bigoplus \bigcup_{i \in I} P_i]$. Then, in two moves, this player reaches $\bigcup_{i \in I} \{c[p] \mid p \in P_i\} = \{c[p] \mid p \in \bigcup_{i \in I} P_i\}$ from

$c[\bigoplus_{i \in I} \bigoplus P_i]$. We also have $\text{succ}(c[\bigoplus_{i \in I} P_i]) = \{c[p] \mid p \in \bigcup_{i \in I} P_i\}$. So it is straightforward to lift the strategies from one term to the other easily under any objective.

Axiom (L5): Let $p, q \in \mathbf{T}$ with $\text{urg}(p) \leq u$. Acquire a context $c[\bullet]$ for which one of p or $p \Psi q$ is immediate. If p is immediate for $c[\bullet]$, since $\text{urg}(p) \leq u = \text{urg}(p \Psi q)$, $p \Psi q$ is also immediate for $c[\bullet]$. Then, in any case, $p \Psi q$ is immediate for $c[\bullet]$. Fix an objective $O \subseteq \Sigma^*$ and let $c[p] \Downarrow O$. Since $c[p \Psi q]$, Eve can choose p in the inserted choice to reach $c[p]$ and win, i.e. $c[p \Psi q] \Downarrow O$.

Axiom (L3): Let $p, q \in \mathbf{T}$ and let $\text{urg}(p) \leq u$. We will only show $p \simeq p \Psi (p \Delta q)$. The proof of the dual statement is analogous. Fix an objective $O \subseteq \Sigma^*$. Acquire a context $c[\bullet]$ for which one of p or $p \Psi (p \Delta q)$ be immediate. Similarly to (L5), $p \Psi (p \Delta q)$ is guaranteed to be immediate. Let $c[p] \Downarrow O$. Then Eve can play $c[p \Psi (p \Delta q)] \rightarrow c[p]$ and win, so $c[p \Psi (p \Delta q)] \Downarrow O$ as well. Let $c[p] \Downarrow O$ not hold. Then Adam has a winning strategy from $c[p]$, since reachability games are determined. If Eve were to play $c[p \Psi (p \Delta q)] \rightarrow c[p]$, Adam would win from this position. If Eve were to instead play $c[p \Psi (p \Delta q)] \rightarrow c[p \Delta q]$, Adam can play $c[p \Delta q] \rightarrow c[p]$ and win. Thus $c[p \Psi (p \Delta q)] \Downarrow O$ does not hold.

Axiom (L2): Let $\bigwedge_{i \in I} \Psi P_i \in \mathbf{T}$. Per Axiom of choice, $\{f \mid f : I \rightarrow P_I\} \neq \emptyset$ and $\bigvee_{f: I \rightarrow P_I} \bigwedge_{i \in I} \{f(i) \mid i \in I\}$ is well defined. Let $\bigvee_{f: I \rightarrow P_I} \bigwedge_{i \in I} \{f(i) \mid i \in I\} \in \mathbf{T}$. Fix an objective $O \subseteq \Sigma^*$ and acquire a context $c[\bullet]$ for which one of $\bigvee_{f: I \rightarrow P_I} \bigwedge_{i \in I} \{f(i) \mid i \in I\}$ and $\bigwedge_{i \in I} \Psi P_i$ is immediate. Since both terms have urgency u , Lemma 3 states that both terms are immediate. Per definition, we have $c[\bigwedge_{i \in I} \Psi P_i] \Downarrow O$ if and only if for all $i \in I$, there is a $p \in P_i$ with $c[p] \Downarrow O$. Since we assume axiom of choice, we can apply Skolemization to get the following equivalent statement: There is a $f : I \rightarrow P_I$ where for all $i \in I$, $c[f(i)] \Downarrow O$. But this is equivalent to $c[\bigvee_{f: I \rightarrow P_I} \bigwedge_{i \in I} \{f(i) \mid i \in I\}] \Downarrow O$. Then $c[\bigvee_{f: I \rightarrow P_I} \bigwedge_{i \in I} \{f(i) \mid i \in I\}] \Downarrow O$ if and only if $c[\bigwedge_{i \in I} \Psi P_i] \Downarrow O$.

Axiom (D2): Let $p, \bigoplus Q \in \mathbf{T}$ where $\text{urg}(p) \leq u$. Acquire a context for which at least one, and per Lemma 3 both, of $(\bigoplus Q).p$ and $\bigoplus\{q.p \mid q \in Q\}$ are immediate. We have the owner $\text{own}(c[(\bigoplus Q).p]) = \text{own}(c[\bigoplus\{q.p \mid q \in Q\}])$. It follows that we have $\text{succ}(c[(\bigoplus Q).p]) = \{c[q.p] \mid q \in Q\}$ due to $\text{succ}((\bigoplus Q).p) = \{q.p \mid q \in Q\}$. Further, $\text{succ}(c[\bigoplus\{q.p \mid q \in Q\}]) = \{c[q.p] \mid q \in Q\}$ as well. So under any objective $O \subseteq \Sigma^*$, lifting the strategies from one term to the other is straightforward.

Axiom (M): Let $w, v \in \mathbf{W}$ with $w \stackrel{*}{\simeq} v$. Acquire a context $c[\bullet]$ that is immediate for one of w or v . Since these are both word terms, the only way one of these terms can be immediate is if $c[\bullet]$ is a concatenation of terminals and \bullet . Then, we have $c[w] \stackrel{*}{\simeq} c[v]$. For an objective $O \subseteq \Sigma^*$, we also see that $c[w] \Downarrow O$ if and only if $c[v] \Downarrow O$.

Axiom (S): Let $O \subseteq \Sigma^*$ be an objective. Let $w, v \in \Sigma_{\text{err}}^*$ be word terms with $w \sqsubseteq_O v$ due to $w \leq_O^s v$. Let $c[\bullet]$ be a context for which one of the words is immediate. As the words have urgency zero, the context must be a word as well: $c[\bullet] = x \bullet y$ (ignoring the bracketing) for $x, y \in \Sigma_{\text{err}}^*$. Assume Eve wins O from $c[w] = x.w.y$. Then $x.w.y \in O$. By definition of \leq_O^s , we get $x.v.y \in O$. So Eve wins O from $c[v]$ as well.

Axiom (FP): Let $A \in N$. Acquire a context $c[\bullet]$ for which at least one of A or $E(A)$ is immediate. Since $\text{urg}(A) = \mathbf{h} \geq E(A)$, Lemma 3 tells us that A is guaranteed to be immediate. The term $c[\underline{A}]$ has exactly one successor, $c[E(A)]$. Then under any objective $O \subseteq \Sigma^*$, $c[\underline{A}] \Downarrow O$ if and only if $c[E(A)] \Downarrow O$. \square

E Normalization

We rely on proof rules that can be derived from the axioms. These follow from the axioms in Figure 3. Before we address normalization, we take a brief detour to prove the utilized proof rules correct. These will also be used in Appendix I.

$$\begin{aligned} \text{(REP)} \quad & \frac{\forall A \in N. p_A \sqsubseteq A}{q\{N/p_N\} \sqsubseteq q} & \text{(L6)} \quad & \frac{\text{urg}(p) \leq \mathbf{u}}{p \equiv \bigoplus p} \\ \text{(DN)} \quad & \frac{\mathbf{v} < \mathbf{u}}{\bigwedge \bigoplus Q \equiv \bigwedge \bigoplus \mathbf{v} Q} \end{aligned}$$

Proof sketch. For (REP), we use (L2) and the congruence rule inductively on the subterms. Utilizing (L3) twice yields (L6). (DN) follows from

$$\begin{aligned} \bigwedge_{i \in I} \bigoplus Q_i & \stackrel{\text{(L6)}}{\equiv} \bigoplus \bigwedge_{i \in I} \bigoplus Q_i \stackrel{\text{(L2)}}{\equiv} \bigwedge_{i \in I} \bigoplus Q_i \\ & \stackrel{\text{(N)}}{\equiv} \bigwedge_{i \in I} \bigoplus \bigoplus \mathbf{v} Q_i \stackrel{\text{(L6)}}{\equiv} \bigwedge_{i \in I} \bigoplus \mathbf{v} Q_i \end{aligned}$$

\square

E.1 Proof of Lemma 4, Part 1

The function $\text{nf}(p)$ is defined by induction. For p being a terminal or skip, err , we use (L6) to introduce a sequence of $2\mathbf{h}$ choices over singleton sets and arrive at a term $\text{nf}(p) \equiv p$ in normal form. For a concatenation or choice, we recursively normalize the operands and then invoke specialized functions that rely on the operands being normalized:

$$\begin{aligned} \text{nf}(p.q) &= \text{nfconc}(\text{nf}(p).\text{nf}(q)) \\ \text{nf}(\bigoplus P) &= \text{nfchoice}(\bigoplus \{\text{nf}(p) \mid p \in P\}) . \end{aligned}$$

Lemma 10. *Let $R \subseteq NF_{\mathbf{h}}$ and $r = \bigoplus R$. We can find $\text{nfchoice}(r) \in NF_{\mathbf{h}}$ with $\text{nfchoice}(r) \equiv r$.*

Proof. Consider the proof of Lemma 12. \square

Lemma 11. *For $p, q \in NF_{\mathbf{h}}$ we can find $\text{nfconc}(p.q) \in NF_{\mathbf{h}}$ with $\text{nfconc}(p.q) \equiv p.q$.*

Proof. We strengthen the statement and show that for all urgencies u , if we have normal form terms $p, q \in NF_u$, then we can obtain a normal form in NF_u . We proceed by induction on u . For the base case $u = 0$, (M) yields $p.q \stackrel{\cong}{=} r \in NF_0$.

For the inductive case, let $u > 0$. Then $p = \Psi_{i \in I} \bigwedge_{\mathfrak{u}} P_i$ and $q = \Psi_{j \in J} \bigwedge_{\mathfrak{u}} Q_j$. So we can write:

$$\begin{aligned} p.q &= (\Psi_{i \in I} \bigwedge_{\mathfrak{u}} P_i) \cdot (\Psi_{j \in J} \bigwedge_{\mathfrak{u}} Q_j) \\ &\stackrel{2 \times (\text{D2})}{\equiv} \Psi_{i \in I} \bigwedge_{p' \in P_i} p' \cdot (\Psi_{j \in J} \bigwedge_{\mathfrak{u}} Q_j) \\ &\stackrel{2 \times (\text{D1})}{\equiv} \Psi_{i \in I} \bigwedge_{p' \in P_i} \Psi_{j \in J} \bigwedge_{q' \in Q_j} p' \cdot q' . \end{aligned}$$

We have $p', q' \in NF_{u-1}$. We apply I.H. to obtain $\text{nfconc}(p'.q') \in NF_{u-1}$

$$\stackrel{\text{I.H.}}{\equiv} \Psi_{i \in I} \bigwedge_{p' \in P_i} \Psi_{j \in J} \bigwedge_{q' \in Q_j} \text{nfconc}(p'.q') .$$

The term is not in normal form due to the two layers of u choices. We apply Lemma 10 to obtain a normal form. \square

E.2 Proof of Lemma 4, Part 2

We proceed by a transfinite Kleene iteration. By induction on α , we construct normal form terms $A^{(\alpha)} \in NF_{\mathfrak{h}}$ for all non-terminals A and all ordinals α . We write $p^{(\gamma)}$ to denote $p\{N/N^{(\gamma)}\}$, where $N^{(\gamma)}$ refers to a vector of terms that has $A^{(\gamma)}$ as its A component. We let $A^{(0)} = \text{nf}(\text{err})$ and for all $\alpha > 0$:

$$A^{(\alpha)} = \text{nf}(\Psi\{E(A)^{(\beta)} \mid \beta < \alpha\}) .$$

In both cases, we rely on the normalization from Appendix E.1.

We claim that $A^{(\beta)} \sqsubseteq A^{(\alpha)} \sqsubseteq A$ for all ordinals $\beta < \alpha$. By definition, the alternatives available to Eve in $A^{(\beta)}$ are contained in the alternatives available to her in $A^{(\alpha)}$. The former precongurence $A^{(\beta)} \sqsubseteq A^{(\alpha)}$ then follows from a simple application of Axioms (L1), (L4) and (L5).

With a transfinite induction we show that for all ordinals α and for all non-terminals A we have $A^{(\alpha)} \sqsubseteq A$. For the base case, we already have $A^{(0)} \equiv \text{err} \sqsubseteq A$. For the inductive case, let α be an ordinal so that for all ordinals $\beta < \alpha$ and all non-terminals A we have $A^{(\beta)} \sqsubseteq A$. Using (REP), we see that $E(A)^{(\beta)} \sqsubseteq E(A)$. So we can apply (L1) to get

$$\Psi\{E(A)^{(\beta)} \mid \beta < \alpha\} \sqsubseteq \Psi E(A) \stackrel{(\text{L6})}{\equiv} E(A) \stackrel{(\text{FP})}{\equiv} A .$$

It remains to show that $A \sqsubseteq A^{(\gamma)}$ for some ordinal γ . The largest \sqsubseteq -chain of strictly increasing elements in $NF_{\mathfrak{h}}$ has size $|ANF_{\mathfrak{h}}|$. Then the largest such chain

in $N \rightarrow NF_h$ has size $|N||ANF_h|$. As a chain forms a well-ordered set, there is an ordinal γ having at least this size. This means we have $A^{(\gamma)} \equiv A^{(\gamma+1)}$ for all non-terminals, and

$$\begin{aligned} E(A)^{(\gamma+1)} &\stackrel{\text{(L5)}}{\sqsubseteq} E(A)^{(\gamma)} \Psi \Psi \{E(A)^{(\beta)} \mid \beta < \gamma\} \\ &\stackrel{\text{(L4)}}{\equiv} \Psi \{E(A)^{(\beta)} \mid \beta < \gamma + 1\} \equiv A^{(\gamma+1)} \equiv A^{(\gamma)}. \end{aligned}$$

Applying (LFP), we get $A \sqsubseteq A^{(\gamma)}$ for all $A \in N$.

F Completeness: Missing Proof

Proof. We focus on the former implication, the latter is simple. The implication trivially holds for terms p that are minimal in the domination preorder. We will need information about the shape of these minimal terms. In urgency zero, minimal are all terms $w \leq_O^s \text{err}$, meaning there is no chance to extend w to a word in O . For higher urgencies $u > 0$, terms ΨP are minimal where all elements in P are minimal. Terms $\bigwedge P$ are minimal where an element in P is minimal.

To show the implication for terms p and q where p is not minimal, we use *characteristic contexts*. Given a normal form term p , we construct a context $c_p[\bullet]$ so that for all normal form terms q of the same urgency as p and owned by the same player we have:

$$c_p[q] \Downarrow O \quad \text{iff} \quad p \subseteq_O q. \quad (1)$$

The implication from $p \preceq_O q$ to $p \subseteq_O q$ indeed follows. Since $p \subseteq_O p$, we obtain $c_p[p] \Downarrow O$ by Equivalence (1). The assumption $p \preceq_O q$ now yields $c_p[q] \Downarrow O$. Hence, again by Equivalence (1), we have $p \subseteq_O q$. For the maximal urgency, we need a special treatment.

It remains to give the construction of $c_p[\bullet]$. It will have the shape $\bullet.t_p$ where t_p is again in normal form. We proceed by induction on the urgency u and need a special case for the maximal urgency h .

Base Case $p \in NF_0$. We define

$$t_p = \bigwedge \{y \in \Sigma^* \mid p.y \in O\}.$$

The set is non-empty as p is not minimal. Moreover, for every normal form term q of urgency zero, Equivalence (1) holds by the definition of \leq_O^s .

Inductive Case $p = \Psi P \in NF_u$, $u < h$. We define

$$t_p = \bigwedge_{u+1} \{t_r \mid r \in P \text{ not minimal}\}.$$

Note that we can increase the urgency because $u < h$. A non-minimal r is guaranteed to exist in P by the assumption that ΨP itself is not minimal. To

prove Equivalence (1), we consider $\Psi Q \in NF_u$ and argue as follows:

$$\begin{aligned}
& (\Psi Q).(\bigwedge_{u+1} \{t_r \mid r \in P \text{ not minimal}\}) \Downarrow O \\
& \text{iff } \forall r \in P \text{ not minimal. } \exists s \in Q. s.t_r \Downarrow O \\
\{\text{I.H.}\} & \text{iff } \forall r \in P \text{ not minimal. } \exists s \in Q. r \subseteq_O s \\
& \text{iff } \forall r \in P. \exists s \in Q. r \subseteq_O s \\
& \text{iff } \Psi P \subseteq_O \Psi Q .
\end{aligned}$$

The inductive case for $p = \bigwedge P \in ANF_u$, $u \leq h$, is similar.

Special Case NF_h : Since the urgency $h + 1$ is not allowed, we are not able to construct the context in the way we did above. Instead, we show that for $p, q \in NF_h$ with $p \not\subseteq_O q$, there is a term $r \in ANF_h$ where $p.t_r$ is won by Eve and $q.t_r$ is won by Adam. This yields $p \not\subseteq q$.

Let $p = \Psi P$ and $q = \Psi Q$ both in NF_h with $p \not\subseteq_O q$. By definition of the domination preorder, there is $r \in P$ so that for all $s \in Q$ we have $r \not\subseteq_O s$. The term r cannot be minimal, and hence t_r is guaranteed to exist. We claim that Eve wins O from $p.t_r$ while Adam wins O from $q.t_r$. Note that the leading terms are $\underline{p}.t_r$ resp. $\underline{q}.t_r$. To see that Eve wins from $\underline{p}.t_r$, let her choose $r \in P \subseteq ANF_h$ to reach $r.t_r$. Since $r \subseteq_O r$, Equivalence (1) yields $r.t_r \Downarrow O$. To see that Adam wins from $\underline{q}.t_r$, let Eve choose $s \in Q \subseteq ANF_h$. As $r \not\subseteq_O s$, Equivalence (1) implies that Eve loses from $s.t_r$. \square

G Completeness: Characteristic Context

The missing inductive case is $p = \bigwedge P \in ANF_u$. We set

$$t_p = \Psi\{t_r \mid r \in P\}.$$

Since p is not minimal, no element $r \in P$ is minimal and hence the t_r are guaranteed to exist. To see Equivalence (1), consider $\bigwedge Q \in ANF_u$:

$$\begin{aligned}
& (\bigwedge Q).(\Psi\{t_r \mid r \in P\}) \Downarrow O \\
& \text{iff } \forall s \in Q. \exists r \in P. s.t_r \Downarrow O \\
\{\text{I.H.}\} & \text{iff } \forall s \in Q. \exists r \in P. r \subseteq_O s. \\
& \text{iff } \bigwedge P \subseteq_O \bigwedge Q.
\end{aligned}$$

H Completeness: Domain Shattering Objectives

We show Lemma 5, namely there are right-separating and domain-shattering objectives. For $|\Sigma| > 1$, we take

$$O = \{w.w^{\text{reverse}} \mid w \in \Sigma^*\} .$$

For $\Sigma = \{a\}$, we let $O = \{a^{(n^2)} \mid n \in \mathbb{N}\}$. It is known that these objectives have a syntactic congruence with singleton classes. Even the syntactic precongruence is the reflexive relation, $\leq_O^s = \cong_O^s$.

Let $O = \{a^{(n^2)} \mid n \in \mathbb{N}\}$ if $\Sigma = \{a\}$ and $O = \{w.w^{reverse} \mid w \in \Sigma^*\}$ if $|\Sigma| > 1$. We show that for all $w, v \in \mathbf{W}$ with $w \not\cong v$, we can find a $x \in \Sigma^+$ with $w.x \in O$ and $v.x \notin O$. This proves that O is right separating. In particular, it also showcases that Axiom (S) cannot relate terms w, v unless $w \cong v$. So it also proves that O is domain-shattering. This also shows for all $w \in \Sigma^+$ the existence of x with $w.x \in O$ to establish the right separability of w and \perp .

Case $|\Sigma| = 1$: Let $w = a^i$ and $v = a^j$ with $i \neq j$. Then set $x = a^{t^2-i}$ with $t = i + j + 1$. Obviously, $w.x \in O$. Suppose $a^j.a^{t^2-i} \in O$. Then, $j + t^2 - i = k^2$ for some $k \in \mathbb{N}$. So $j - i = t^2 - k^2 = (t - k) \cdot (t + k)$. $i \neq j$ implies $|t - k| \geq 1$. Note that $k, i, j \geq 0$ so we get the contradiction $|j - i| < i + j < t + k < |t - k| \cdot |t + k| = |j - i|$.

Case $|\Sigma| > 1$: Let $w = w_0w_1 \dots w_n$ and $v = v_0v_1 \dots v_m$. Let $a \neq w_m$ be a terminal. We set $x = a.a.w^{reverse}$, where $a \neq w_{|w|}$ if $|w| > |v|$, $a \neq v_{|w|}$ if $|v| > |w|$, and a arbitrary if $|w| = |v|$. Suppose $v.a.a.w^{reverse} \in O$. Then, $v.a.a.w^{reverse} = w.a.a.v^{reverse}$. If $|w| = |v|$, $w = v$ is a contradiction. Otherwise, if $|v| < |w|$, we have $(v.a)_{|v|} = a \neq w_{|v|} = (w.a)_{|v|}$ for a contradiction. Similarly, if $|w| < |v|$, the contradiction is $(w.a)_{|w|} = a \neq v_{|w|} = (v.a)_{|w|}$.

I Specialized Normal Form

We now provide a sketch of the normalization algorithm with the complexity $(|p| + |E||N|) \cdot \exp_{2h-1}^O$ as given in Lemma 7.

We use methods from Section 5 with slight modifications. To avoid repetition, we only note the modifications that need to be made, instead of giving a full algorithm. The section takes the form of constructing modified normalization functions $\text{nf}_O(\cdot)$, $\text{nfchoice}_O(\cdot)$, and $\text{nfconc}_O(\cdot)$. We provide lemmas that reference those in Section 5. The $\text{nfchoice}_O(\cdot)$ and $\text{nfconc}_O(\cdot)$ calls made by $\text{nf}_O(\cdot)$ are kept the same. Note that implementing $\text{nfchoice}_O(\cdot)$ and $\text{nfconc}_O(\cdot)$ with time complexity $|p| \cdot \exp_{2h-1}^O$ for the input term p means that $\text{nf}_O(q)$ is also constructed in $|q| \cdot \exp_{2h-1}^O$ time, if q has no non-terminals. For the rest of the section, fix a regular objective $O \subseteq \Sigma^*$, a finite set of non-terminal symbols N , and a finitary defining assignment $E : N \rightarrow \mathbf{T}$.

Lemma 12 (Modifies Lemma 10). *Let $P \subseteq \text{SNF}_h^O$ and $r = \bigoplus P$. We can find $\text{nfchoice}_O(r) \in \text{SNF}_h^O$ with $\text{nfchoice}_O(r) \equiv_O r$ in $|r| \cdot \exp_{2h-1}^O$ time.*

Proof. By induction on h . Note that for normalforms in SNF_u^O also have maximal urgency u . In any case of u we need to normalize a term $r = \bigoplus_{i \in I} \Psi P_i$. The harder case is $\bigcirc = \bigwedge$. We apply (N) if $u < h$ and (L2) to obtain $\Psi_{f: I \rightarrow P_I} \bigwedge_{i \in I} f(i)$ and utilize (L4) to combine the $\bigwedge_{i \in I}$ with the $f(i) \in \text{SANF}_u^O$:

$$\text{nfchoice}_O(r) = \Psi_{f: I \rightarrow P_I} \bigwedge_{i \in I} Q_f$$

¹ We use \exp_u^O for $\exp_u(\mathcal{O}(|\text{SNF}_0^O|))$. Usages of \leq and $=$ are to be understood by means of \in or \subseteq .

where $Q_f = \bigcup\{Q \mid f(i) = \bigwedge Q, i \in I\}$. Naively, applying Axiom (L2) requires us to enumerate all the choice functions $f : I \rightarrow P_I$. However, we do not need to account for all of them, because after applying (L4) for a single f , we know that $\bigwedge Q_f \in SANF_u^O$. Knowing this means that the distribution considers way more functions f (namely $\prod_{i \in I} |P_i| \geq 2^{|SANF_u^O|}$) than there can be sets Q_f ($|SANF_u^O|$). Instead, we can use (L4) to split the application of (L2) into $|I| - 1$ many single applications of (L2) and keeping the size of the intermediary results bound by $|SANF_u^O| = \exp_{2u-1}^O$: In the case of $J \subseteq I$ with $|J| = 2$, i.e. constant, the number of functions $f : J \rightarrow P_J$ is bound by $|P_1| \cdot |P_2| \leq |SANF_u^O|^2$. For such an f (with binary co-domain) the union Q_f can be computed in time $|SANF_{u-1}^O|^2$ as long as the terms $f(j)$ are in normalform $f(j) \in SANF_u^O$. In case of $h = u$, this will be the case. Otherwise, $u < h$ and we apply the induction hypothesis to obtain normalforms equivalent to $f(1)$ and $f(2)$ before computing Q_f . That way, we compute no more than $|I| \leq |r|$ many (L4) for splitting, $|I|$ many (L2) enumerating $|SANF_{u-1}^O|^2 \leq \exp_{2u-1}^O$ functions f each, and for each f we apply (L4) to create the union Q_f in time $|SANF_u^O|^2 \leq \exp_{2u-2}^O$. Together (for $h = u$):

$$\underbrace{|I|}_{(L4)} + |I| \cdot \underbrace{|SANF_u^O|^2}_{(L2)} \cdot \underbrace{|SANF_{u-1}^O|^2}_{(L4), Q_f} \leq |r| \cdot \exp_{2u-1}^O$$

In case of $u < h$, the right summand changes to

$$\begin{aligned} |I| \cdot \underbrace{|SANF_u^O|^2}_{(L2)} \cdot \underbrace{(|SANF_{u-1}^O|^2 + 2|SANF_{h-1}^O| \cdot \exp_{2h-3}^O)}_{(L4), Q_f, \text{I.H.}} \\ \leq |r| \cdot \exp_{2h-1}^O \end{aligned}$$

Adding singleton choice operators (L6) to reobtain a term in SNF_h^O costs close to no time. \square

Lemma 13 (Modifies Lemma 11). *For any $p, q \in SNF_h^O$, we can find $\text{nfconc}_O(p, q) \in SNF_h^O$ with $\text{nfconc}_O(p, q) \equiv_O p \cdot q$ in \exp_{2h-1}^O time.*

Proof. The proof in Section 5 proceeds by an induction on urgency. We change the base case to $h = 1$ and for both, base and inductive case, apply (D1) and (D2) as in the proof of Lemma 11. The application of (D1) creates at most $|p|$ copies of q , and (D2) creates another $|q|$ copies of each p' . Call the resulting term s with $|s| \leq 2 \cdot |p| \cdot |q|$.

In the base case, p' and q' belong to SNF_0^O . We employ (FP) and find an $x \in SNF_0^O$ with $x \cong_O^s w.v$. Finding the syntactical congruence class of $w.v$ can be done in time $(\exp_0^O)^k$ for some fixed $k \in \mathbb{N}$ by naively checking.

In the inductive case, we apply the induction hypothesis to normalize the concatenative subterms $p'.q'$ into SNF_h^O . A term in SNF_h^O term is bound by size $|SANF_h^O| \leq \exp_{2h-1}^O$. After applying the induction hypothesis to each pair in s , the resulting term t has size $|t| \leq 2 \cdot |p| \cdot |q| \cdot \exp_{2h-1}^O = \exp_{2h-1}^O$. In Section 5 the layers of choices are resolved by invoking Lemma 10. In our case, we employ the modified Lemma 12 for another $|t| \cdot \exp_{2h-1}^O \leq (\exp_{2h-1}^O)^2 = \exp_{2h-1}^O$. \square

Lemma 14 (Modifies Lemma 4). *For any $A \in N$, we can find $\text{nf}(A) \in \text{SNF}_h^O$ with $\text{nf}(A) \equiv_O A$ in $|E| \cdot |N| \cdot \exp_{2^{h-1}}^O$ time.*

Proof Sketch. We construct an increasing chain of SNF_h^O terms using the same least fixed point construction from Section 5. The constructed term for component A has size at most $|E(A)| \cdot (\exp_{2^{h-1}}^O)^2$. Per Lemmas 12 and 13, this term is normalized in $|E(A)| \cdot (\exp_{2^{h-1}}^O)^3 = |E(A)| \cdot \exp_{2^{h-1}}^O$ time. So normalizing all components takes $|E| \cdot \exp_{2^{h-1}}^O$ time. The chain converges in at most $|N| \cdot |\text{SNF}_h^O| = |N| \cdot \exp_{2^{h-1}}^O$ iterations. So the normalization process takes $|E| \cdot |N| \cdot \exp_{2^{h-1}}^O$ time. \square

We complete the normalization function analogously to Section 5. To calculate $\text{nf}_O(p)$ for any finitary $p \in \mathbf{T}$ with non-terminals, the algorithm first constructs $\text{nf}_O(A)$ for all $A \in N$. This takes $|E| \cdot |N| \cdot \exp_{2^{h-1}}^O$ time. Then, we let $\text{nf}_O(p) = \text{nf}_O(p\{N/p_N\})$ where $p_A = \text{nf}_O(A)$. The normalization takes $|p| \cdot \exp_{2^{h-1}}^O$ time per Lemmas 12 and 13. This concludes the construction of the algorithm and thus proves Lemma 7.

J Decidability and Complexity

Proof of Proposition 4. Consider two CFG $G_i = (N_i, E_i, S_i), i \in \{0, 1\}$ over some alphabet Σ and with disjoint sets of non-terminals $N_1 \cap N_2 = \emptyset$. Note that each non-terminal has a single production rule of the shape $E_i(A) = w_1 | \dots | w_n$, $w_i \in (N \cup \Sigma)^*$. We reproduce this non-deterministic structure by yielding the choice to Eve. We construct the program-term grammar $(N_1 \cup N_2, E)$, where E is defined via $E(A) = \bigvee\{w_1, \dots, w_n\}$ for each production rule of above shape. Kleene iteration yields the normal form

$$\text{nf}(S_i) = \bigvee\{\bigwedge w \mid w \in L(G_i)\} \simeq \bigvee L(G_i).$$

We utilize $\preceq_O = \preceq$ by choosing the shattering objective $O = \{w.w^{\text{reverse}} \mid w \in \Sigma^*\}$. Since \preceq_O is the reflexivity relation on $\Sigma^* \times \Sigma^*$, the domination pre-order yields $\bigvee L =_O \bigvee M$ if and only if $L = M$ for any $L, M \subseteq \Sigma^*$. Using Lemma 6, Theorem 2, and that O is shattering and right-separating, $L(G_1) = L(G_2)$ if and only if $S_1 \simeq \bigvee L(G_1) \simeq \bigvee L(G_2) \simeq S_2$. \square

Proof of Theorem 3. We show that $\text{h-DEC-}\preceq_O$ is PTIME-hard (wrt. log-space reductions) by sketching out a reduction from the Monotonic Circuit Value Problem [50]. The problem consists of assignments of boolean variables $P_{i \leq n}$ to \vee/\wedge clauses built out of the variables $P_{j < i}$, to **true**, or to **false**. The input is accepted if and only if the variable P_n evaluates to true. Let non-empty $O \subseteq \Sigma^*$ and let $w \in O$. For each boolean variable P_i , the log-space Turing Machine outputs a non-terminal P_i and a defining equation $E(P_i)$. The machine outputs $E(P_i) = w$ if $P_i = \text{true}$, and it outputs $E(P_i) = \text{err}$ if $P_i = \text{false}$. For $P_i = P_{j_0} \wedge \dots \wedge P_{j_k}$, the machine outputs $E(P_i) = \bigwedge\{P_{j_l} \mid l \leq k\}$. For $P_i = P_{j_0} \vee \dots \vee P_{j_k}$, it outputs

$\bigvee\{P_{j_l} \mid l \leq k\}$. The reduced problem instance asks whether $w \preceq_O P_n$. By induction on i we show $P_i \equiv_O w$ if P_i evaluates to true, and $P_i \equiv_O w$ and if P_i evaluates to false. Since \equiv_O sound, and $w \not\preceq_O \text{err}$ (with the witnessing context $c[\bullet] = \bullet$) this proof suffices.

For the base case, we have $P_0 = \text{true}$ or $P_0 = \text{false}$. After applying Axiom (FP), we get $P \equiv_O w$ and $P \equiv_O \text{err}$ respectively. For the inductive case, we have $P_i = \text{true}$, $P_i = \text{false}$, $P_i = P_{j_0} \wedge \dots \wedge P_{j_k}$, or $P_i = P_{j_0} \vee \dots \vee P_{j_k}$. The first two sub-cases are already handled in the base case. The latter two cases are dual, so we only handle the \wedge -case. Per Axiom (FP), we have $P_i \equiv_O \bigwedge\{P_{j_l} \mid l \leq k\}$. The variable P_i evaluates to true, if and only if P_{j_l} evaluates to true for all $l \leq k$. Let all variables evaluate to true. Induction hypothesis delivers us $P_{j_l} \equiv_O w$ and thus $P_i \equiv_O \bigwedge\{w\}$. The rule Axiom (L6) tells us $\bigwedge\{w\} \equiv_O w$. W.l.o.g. let a variable P_{j_0} evaluate to false. then induction hypothesis delivers us $P_{j_0} \equiv_O \text{err}$ and $P_{j_l} \equiv_O w$ or $P_{j_l} \equiv_O \text{err}$ for all $0 < l \leq k$. We have $P_i \equiv_O \bigwedge\{\text{err}\}$ or $P_i \equiv_O \bigwedge\{\text{err}, w\}$. In the former case we have $P_i \equiv_O \text{err}$ per Axiom (L6). In the latter case, the lattice axioms tell us that $\bigwedge\{w, \text{err}\} \equiv_O \text{err}$, and thus $P_i \equiv_O \text{err}$. \square

K Upper Bound: Compact Terms Representation and Characteristic Terms

K.1 Compact Term Representation

Recall that the idea behind our compact term representation is to translate syntactic congruence classes $[a] : Q \rightarrow Q$ into angelic choices over state changes. With a representative system for the syntactic congruence, we can assume the source alphabet is Σ rather than $\Sigma_{\text{err}}^* / \cong^s$. The translation of term p over Σ is the term $\text{tr}(p)$ over the alphabet Q^2 defined by

$$\begin{aligned} \text{tr}(a) &= \bigvee\{(p, q) \mid \delta(p, a) = q\} & \text{tr}(\text{err}) &= \text{err} \\ \text{tr}(\bigoplus P) &= \bigoplus\{\text{tr}(p) \mid p \in P\} & \text{tr}(\text{skip}) &= \text{skip} \\ \text{tr}(p.q) &= \text{tr}(p).\text{tr}(q) & \text{tr}(A) &= A, \end{aligned}$$

and we also translate the equations, $E_{\text{tr}}(A) = \text{tr}(E(A))$. The new alphabet calls for a translation of the objective. The DFA $\text{tr}(O)$ is obtained from O by modifying the transitions and adding a failure state \perp . We let $\text{tr}(\delta)(p, (q, r)) = r$ if $p = q$ and $\text{tr}(\delta)(p, (q, r)) = \perp$ if $p \neq q$. The accepting and initial states remain the same. The translation is faithful when it comes to our notion of observable behavior.

Lemma 15. $p \Downarrow O$ if and only if $\text{tr}(p) \Downarrow \text{tr}(O)$.

By combining Lemma 15 and Corollary 2, we can decide the specialized contextual preorder $p \preceq_O q$ by checking whether $c[\text{tr}(p)] \Downarrow \text{tr}(O)$ implies $c[\text{tr}(q)] \Downarrow \text{tr}(O)$ for all contexts $\text{tr}(s).\bullet.\text{tr}(t)$ with $s \in \text{SNF}_{h-1}^O$ and $t \in \text{SNF}_h^O$. The problem with this algorithm is that, for complexity reasons, we cannot work with an explicit translation of terms in specialized normal form. To overcome the problem, an attempt would be to generalize the above set of contexts to all $s.\bullet.t$ with

$s \in SNF_{h-1}^{\text{tr}(O)}$ and $t \in SNF_h^{\text{tr}(O)}$. However, the set $SNF_h^{\text{tr}(O)}$ contains terms over the new alphabet Q^2 that do not result from a translation of an SNF_h^O term. Unfortunately, these extra contexts may, incorrectly so, disprove the specialized contextual preorder of interest.

Our solution is to come up with a direct construction for the image of SNF_h^O under $\text{nf}_{\text{tr}(O)} \circ \text{tr}$, the translation followed by a normalization. This is the appropriate subset of $SNF_h^{\text{tr}(O)}$ over which we should form contexts. The idea behind the construction is to explicitly translate the urgency 1 terms in specialized normal form, and build up the higher orders in the standard way. We define the *translated O-specialized normal form* terms (with $u > 1$) by

$$\begin{aligned} TNF_0^{\text{tr}(O)} &= SNF_0^{\text{tr}(O)} & TNF_1^{\text{tr}(O)} &= \text{nf}_{\text{tr}(O)}(\text{tr}(SNF_1^O)) \\ TANF_u^{\text{tr}(O)} &= \{\bigwedge P \mid \emptyset \neq P \subseteq TNF_{u-1}^{\text{tr}(O)}\} \\ TNF_u^{\text{tr}(O)} &= \{\bigvee P \mid \emptyset \neq P \subseteq TANF_u^{\text{tr}(O)}\}. \end{aligned}$$

The set of contexts we should iterate over is thus

$$\mathbf{C}_{\text{tr}(O)} = \{r \bullet s \mid r \in TNF_{h-1}^{\text{tr}(O)} \text{ and } s \in TNF_h^{\text{tr}(O)}\}.$$

The argumentation leads to the following algorithm for checking the specialized contextual preorder.

Proposition 7. *Let O be regular and p, q finitary. Then $p \preceq_O q$ iff for all $c[\bullet] \in \mathbf{C}_{\text{tr}(O)}$, we have that $c[\text{tr}(p)] \Downarrow \text{tr}(O)$ implies $c[\text{tr}(q)] \Downarrow \text{tr}(O)$.*

The benefit over Corollary 2 is that the translated normal form terms are in $\text{tr}(O)$ -specialized normal form, $TNF_h^{\text{tr}(O)} \subseteq SNF_h^{\text{tr}(O)}$, so we inherit the bound.

Lemma 16. $|TNF_h^{\text{tr}(O)}| \leq \exp_{2h}(\mathcal{O}(|Q|^2))$.

K.2 Characteristic Terms

With Proposition 7, we need to iterate over $2h$ -exponentially many contexts. We now eliminate another exponent by factorizing the contexts with the help of characteristic terms. Recall that term p is characteristic for context $c[\bullet]$ wrt. $\text{tr}(O)$, if its $\sqsubseteq_{\text{tr}(O)}$ -upward closure is the solution space of the context: for all q we have $p \sqsubseteq_{\text{tr}(O)} q$ if and only if $c[q] \Downarrow \text{tr}(O)$.

For the contexts $c[\bullet] \in \mathbf{C}_{\text{tr}(O)}$ we just defined, giving Adam a choice over the solution space yields a characteristic term:

$$\chi(c[\bullet]) = \bigwedge_h \{p \in SNF_{h-1}^{\text{tr}(O)} \mid c[p] \Downarrow \text{tr}(O)\}. \quad (2)$$

To see that the term is characteristic indeed, we rely on the domination preorder introduced in the completeness proof.

Lemma 17. *Term $\chi(c[\bullet])$ is characteristic for $c[\bullet] \in \mathbf{C}_{\text{tr}(O)}$ wrt. $\text{tr}(O)$ and can be computed in time $\exp_{2h-1}(\mathcal{O}(|Q|^2))$.*

Proof. To prove that $\chi(c[\bullet])$ is characteristic, consider term $p = \bigvee_{i \in I} \bigwedge P_i$ in $\text{tr}(O)$ -specialized normal form. As p is immediate for $c[\bullet]$, we get $c[p] \Downarrow \text{tr}(O)$ if and only if there is an index $i \in I$ so that for all $q \in P_i$ we have $c[q] \Downarrow \text{tr}(O)$. This can be shown to be equivalent to the domination preorder $\bigvee \bigwedge \{p \in \text{SNF}_{h-1}^{\text{tr}(O)} \mid c[p] \Downarrow \text{tr}(O)\} \sqsubseteq_{\text{tr}(O)} \bigvee_{i \in I} \bigwedge P_i$. For the terms at hand, this domination preorder is equivalent to $\bigwedge \{p \in \text{SNF}_{h-1}^{\text{tr}(O)} \mid c[p] \Downarrow \text{tr}(O)\} \sqsubseteq_{\text{tr}(O)} \bigvee_{i \in I} \bigwedge P_i$, even if the objective is not right-separating.

To compute the characteristic term, we have to check, for every term $p \in \text{SNF}_{h-1}^{\text{tr}(O)}$, whether $c[p] \Downarrow \text{tr}(O)$ holds. Such a check requires a normalization of $c[p]$, followed by a polynomial-time evaluation of the resulting term. The normalization takes time $|c[p]| \cdot \exp_{2h-1}(\mathcal{O}(|Q|^2))$, Lemma 7. The dominating factor in $|c[p]|$ is the size of the $\text{TNF}_h^{\text{tr}(O)}$ term in the context, which is bounded by $\exp_{2h-1}(\mathcal{O}(|Q|^2))$. There are $\exp_{2h-2}(\mathcal{O}(|Q|^2))$ terms p we have to go through. The overall runtime is thus bounded by $\exp_{2h-1}(\mathcal{O}(|Q|^2))$. \square

Let $\chi(\mathbf{C}_{\text{tr}(O)}) = \{\chi(c[\bullet]) \mid c[\bullet] \in \mathbf{C}_{\text{tr}(O)}\}$ denote the set of characteristic terms. As these terms belong to $\text{SANF}_h^{\text{tr}(O)}$, we inherit the following bound.

Lemma 18. $|\chi(\mathbf{C}_{\text{tr}(O)})| \leq \exp_{2h-1}(\mathcal{O}(|O|^2))$.

Compared to Lemma 16, there are exponentially fewer characteristic terms than contexts. To decide $p \preceq_O q$, we thus intend to iterate over all $x \in \chi(\mathbf{C}_{\text{tr}(O)})$ and check whether $x \preceq_{\text{tr}(O)} p$ implies $x \preceq_{\text{tr}(O)} q$. We will use the domination preorder for these checks.

Lemma 19. *Given $p, q \in \text{SNF}_h^{\text{tr}(O)}$, we can decide $p \sqsubseteq_{\text{tr}(O)} q$ in time $|p| \cdot |q|$.*

There is a last obstacle: we do not know the characteristic terms, like we did not know the translated normal form terms above. Going through all contexts and determining the characteristic terms is prohibitively expensive. Generalizing from $\chi(\mathbf{C}_{\text{tr}(O)})$ to $\text{SANF}_h^{\text{tr}(O)}$ is incorrect. The way out is to give a direct construction of the characteristic terms.

The key insight is that the characteristic terms satisfy the following equation, where we have $s \in \text{SNF}_{h-1}^{\text{tr}(O)}$, $T \subseteq \text{SANF}_h^{\text{tr}(O)}$, and $c[\bullet] = s \bullet \bigvee T$:

$$\chi(c[\bullet]) \equiv_{\text{tr}(O)} \bigwedge \{\chi(s \bullet t) \mid t \in T\} . \quad (3)$$

The equation follows from Equation (2), Appendix L. The impact of Equation (3) may not be immediate: we still have to make sure to construct the characteristic term for every set $T \subseteq \text{SANF}_h^{\text{tr}(O)}$. What the equation does is to give us an inductive formulation of the characteristic terms which allows us to compute the set of all characteristic terms in a fixed point. We first construct the characteristic terms for singleton sets $|T| = 1$. Then we conjoin the characteristic terms as prescribed by Equation (3) to obtain the characteristic terms for sets of size $|T| \leq 2$. We repeat the latter conjunction until we reach a fixed point. Throughout the process, we work up to $\equiv_{\text{tr}(O)}$. With Lemma 18, the sets we compute with have size at most $\exp_{2h-1}(\mathcal{O}(|O|^2))$. Moreover, we are guaranteed to reach the fixed

point after at most $\exp_{2h-1}(\mathcal{O}(|O|^2))$ steps. To state the correctness, define for $P \subseteq SNF_{h-1}^{\text{tr}(O)}$ and $Q \subseteq SANF_h^{\text{tr}(O)}$:

$$\mathbb{X}(P, Q) = \bigcup_{s \in P, T \subseteq Q} \{\chi(s \bullet \Psi T)\}.$$

Lemma 20. $\chi(\mathbf{C}_{\text{tr}(O)}) = \mathbb{X}(TNF_{h-1}^{\text{tr}(O)}, TANF_h^{\text{tr}(O)})$. The set can be computed in time $\exp_{2h-1}(\mathcal{O}(|O|^2))$.

The following proposition yields the overall algorithm.

Proposition 8. Let O be regular and p, q finitary. Then $p \preceq_O q$ iff for all $x \in \mathbb{X}(TNF_{h-1}^{\text{tr}(O)}, TANF_h^{\text{tr}(O)})$, we have $x \subseteq_{\text{tr}(O)} \text{nf}_{\text{tr}(O)}(p)$ implies $x \subseteq_{\text{tr}(O)} \text{nf}_{\text{tr}(O)}(q)$.

The domination preorder is sound for checking the axiomatic preorder even for objectives that fail to be right-separating because we have characteristic terms on the left. The time for computing the characteristic terms is given in Lemma 20. The normalization is Lemma 7, and we make use of the fact that the syntactic congruence of $\text{tr}(O)$ has size quadratic in $|Q|$. By Lemma 19, the comparison takes quadratic time. This concludes the proof of Proposition 5.

L Upper Bound: More Details on Compact Term Representation and Characteristic Terms

In this section, we handle the omitted proofs from Section 8. Namely, we prove Proposition 7, Lemma 15, the special case $h = 1$ in Lemma 20 and we justify Equation (3).

Reasoning for Equation (3): To see Equation (3), note that by Equation (2) the characteristic term $\chi(c[\bullet])$ is an urgency- h choice owned by Adam over the set of $SNF_{h-1}^{\text{tr}(O)}$ solutions of $c[\bullet]$. This set of solutions is the union of the $SNF_{h-1}^{\text{tr}(O)}$ solutions for $s \bullet t$ with $t \in T$. The reason is that the choice over T in the context has a higher urgency than the inserted term. We can thus stratify Adam's choice into a choice over $t \in T$ followed by a choice of the $SNF_{h-1}^{\text{tr}(O)}$ solutions for $s \bullet t$. Again by Equation (2), this is precisely the right-hand side of Equation (3).

Proofs of Lemma 15 and Proposition 7: Proofs of these statements require us to observe the inner workings of terms. To do so cleanly, we extend the relation $\stackrel{*}{\sim}$ to terms. The relation $\stackrel{*}{\sim}$ on terms is the smallest equivalence relation that contains $\stackrel{*}{\sim}$ on \mathbf{W} , and the equalities $p.(q.r) \stackrel{*}{\sim} (p.q).r$, $p.\text{skip} \stackrel{*}{\sim} p$, $\text{skip}.p \stackrel{*}{\sim} p$, $p.\text{err}.q \stackrel{*}{\sim} \text{err}$ for all $p, q, r \in \mathbf{T}$. Note that this is different than Axiom (M), which allows us to apply $\stackrel{*}{\sim}$ to words enclosed by arbitrary contexts. We also define $\text{hd} : \mathbf{T} \rightarrow \mathbf{W}$ and $\text{tl} : \mathbf{T} \rightarrow \mathbf{T}$. We let $\text{hd}(p) = w \in \mathbf{W}$ be the concatenation of outermost terminals (including any skip and err) that appear in p before the leftmost outermost action. If there are no such terms, $\text{hd}(p) = \text{skip}$. We let $\text{tl}(p) = q$ be the concatenation of the remaining outermost actions and commands. If no

outmost action exists in the term, we let $\text{tl}(p) = \text{skip}$. Note that $p \cong \text{hd}(p).\text{tl}(p)$ for all $p \in \mathbf{T}$. Finally, we define a set of quasi-runs QRun_w for all $w \in \Sigma^* \cup \{\text{err}\}$. We employ the notation QRun_w for a word term $w \in \mathbf{W}$, as a shorthand for QRun_v where v is the monoid element from $\Sigma^* \cup \{\text{err}\}$. We let $\text{QRun}_{\text{skip}} = \{\text{skip}\}$ and $\text{QRun}_{\text{err}} = \{\text{err}\}$. For a word $w = a_0 \dots a_{k-1} \in \Sigma^+$, the set QRun_w contains all terms of the form $(i, p_1).(p_1, p_2) \dots (p_{n-1}, p_n).\text{tr}(a_n) \dots \text{tr}(a_{k-1})$, where $i \xrightarrow{a_0} p_1 \dots p_{n-1} \xrightarrow{a_{n-1}} p_n$ is a run in the DFA for O . So an element of QRun_w runs the DFA on the prefix of w up to n and has undetermined transitions in form of the terms $\text{tr}(a_i)$ for the remainder of w .

We find it useful to prove a stronger version of Lemma 15.

Lemma 21. *Let $p \in \mathbf{T}$, $r \in \text{QRun}_{\text{hd}(p)}$ and $O \subseteq \Sigma^*$. Then for all $q \cong r.\text{tr}(\text{tl}(p))$, $p \Downarrow O$ if and only if $q \Downarrow \text{tr}(O)$.*

Assuming Lemma 21 we show Proposition 7.

Proof of Proposition 7. Let $O \subseteq \Sigma^*$. Per Corollary 2, we know that $p \preceq_O q$ holds if and only if $s.p.t \Downarrow O$ implies $s.q.t \Downarrow O$ for all $s \in \text{SNF}_{h-1}^O$ and $t \in \text{SNF}_h^O$. We apply Lemma 15 to see that this is equivalent to the statement $\text{tr}(s.p.t) \Downarrow \text{tr}(O)$ implies $\text{tr}(s.q.t) \Downarrow \text{tr}(O)$ for all $s \in \text{SNF}_{h-1}^O$ and $t \in \text{SNF}_h^O$. Per definition, we have $\text{tr}(s.p.t) = \text{tr}(s).\text{tr}(p).\text{tr}(t)$ and $\text{tr}(s.q.t) = \text{tr}(s).\text{tr}(q).\text{tr}(t)$. For the moment, assume $\text{TNF}_u^{\text{tr}(O)} = \{\text{tr}(r) \mid r \in \text{SNF}_u^O\}$ (up to $\equiv_{\text{tr}(O)}$) for all $u \geq 1$ without proof. For $h > 1$, this makes the previous statement equivalent with the desired $s.\text{tr}(p).t \Downarrow O$ implies $s.\text{tr}(q).t \Downarrow O$ for all $s \in \text{TNF}_{h-1}^{\text{tr}(O)}$ and $t \in \text{TNF}_h^{\text{tr}(O)}$.

Now let $h = 1$. We show that for all $s \in \text{SNF}_0^O$ and $t \in \text{SNF}_1^O$, there are $s' \in \text{TNF}_0^{\text{tr}(O)}$ and $t' \in \text{TNF}_1^{\text{tr}(O)}$ where $s.r.t \Downarrow O$ if and only if $s'.\text{tr}(r).t' \Downarrow \text{tr}(O)$ for all $r \in \text{SNF}_1^O$. We can also conversely find $s \in \text{SNF}_0^O$ and $t \in \text{SNF}_1^O$ for all $s' \in \text{TNF}_0^{\text{tr}(O)}$ and $t' \in \text{TNF}_1^{\text{tr}(O)}$ with the same property. Then, the statement $s.p.t \Downarrow O$ implies $s.q.t \Downarrow O$ for all $s \in \text{SNF}_0^O$, $t \in \text{SNF}_1^O$ is equivalent to the statement $s'.\text{tr}(p).t' \Downarrow \text{tr}(O)$ implies $s'.\text{tr}(q).t' \Downarrow \text{tr}(O)$ for all $s' \in \text{TNF}_0^{\text{tr}(O)}$, $t' \in \text{TNF}_1^{\text{tr}(O)}$. Note that, we have $\text{hd}(s.p.t) = s$ and $\text{tl}(s.p.t) = p.t$. Per Lemma 21, we know that for all $r_s \in \text{QRun}_s$, $s.p.t \Downarrow O$ if and only if $r_s.\text{tr}(p.t) \Downarrow O$. We can let $s' = r_s = (i, p_1).(p_1, p_2) \dots (p_{n-1}, p_n) \in \text{TNF}_0^{\text{tr}(O)} \cap \text{QRun}_s$ where the DFA runs on s from i to p_n and $t' = \text{tr}(t)$.

Finally, we show our assumption $\text{TNF}_u^{\text{tr}(O)} = \{\text{tr}(r) \mid r \in \text{SNF}_u^O\}$ (up to $\equiv_{\text{tr}(O)}$) for all $u \geq 1$. The inclusion $\text{TNF}_u^{\text{tr}(O)} \subseteq \{\text{tr}(r) \mid r \in \text{SNF}_u^O\}$ follows from the fact that $\text{tr}(\bigoplus P) = \bigoplus \{\text{tr}(p) \mid p \in P\}$ along with the definitions of $\text{TNF}_u^{\text{tr}(O)}$ and $\text{TANF}_u^{\text{tr}(O)}$. The inclusion $\{\text{tr}(r) \mid r \in \text{SNF}_u^O\} \subseteq \text{TNF}_u^{\text{tr}(O)}$ is proven by induction on u . For the base case, we have $\text{TNF}_1^{\text{tr}(O)} = \{\text{nf}(\text{tr}(r)) \mid r \in \text{SNF}_1^O\}$ per definition. For the inductive case, the cases SNF_u^O and SANF_u^O are analogous, so we only handle one. Letting $\Psi P \in \text{SNF}_u^O$, we get $\text{nf}(\text{tr}(\Psi P)) = \text{nf}(\Psi \{\text{nf}(\text{tr}(r))\} \mid r \in P)$. Induction hypothesis tells us that $\text{nf}(\text{tr}(r)) \in \text{TANF}_u^{\text{tr}(O)}$ for all $r \in P$, so $\Psi \{\text{nf}(\text{tr}(r))\} \in \text{TNF}_u^{\text{tr}(O)}$. \square

The proof of Lemma 21 is more involved and relies on the following Lemmas 22 and 23, which we prove first.

Lemma 22. *Let $p, q \in \mathbf{T}$ with $p \stackrel{\cong}{\approx} q$. Then $\text{own}(p) = \text{own}(q)$ and Eve wins from p in β moves if and only if Eve wins from q in β moves.*

Proof Sketch. The definition of ownership already implies $\text{own}(p) = \text{own}(q)$ for all $p, q \in \mathbf{T}$ with $p \stackrel{\cong}{\approx} q$. The remainder of the statement is proven by transfinite induction on β . The base case follows from the monoid evaluation of word terms. We sketch out the inductive case. W.l.o.g. we can only handle one direction of the implication. Let $p \stackrel{\cong}{\approx} q$ and let Eve win from p in β moves. In both terms, the same i -th concatenation operand (ignoring the bracketing) will be leading. Then, the successor sets are equal up to $\stackrel{\cong}{\approx}$. If $\text{own}(p) = \text{own}(q) = \text{Eve}$, then Eve wins from at least one $p' \in \text{succ}(p)$ in $\gamma < \beta$ moves. Here, we can apply the induction hypothesis to lift the strategy to q . The case $\text{own}(p) = \text{own}(q) = \text{Adam}$ is dual. \square

We call $p \in \mathbf{T}$ *headless* if $\text{hd}(p) = \text{skip}$. For headless terms, $\text{tr}(\cdot)$ and $\text{succ}(\cdot)$ commute.

Lemma 23. *For headless $p \in \mathbf{T}$, $\text{succ}(\text{tr}(p)) = \text{tr}(\text{succ}(p))$.*

Proof. The proof is by (transfinite) structural induction on $p \in \mathbf{T}$. For the base case, we let $p = A \in N$. This is the only base case, because p is only headless if $p = A \in N$. Per definition, we have $\text{succ}(\text{tr}(A)) = \{\text{tr}(E(A))\} = \text{tr}(\text{succ}(A))$.

The first inductive case is $p = \bigoplus P$. Per definition, $\text{tr}(\bigoplus P) = \bigoplus \{\text{tr}(p) \mid p \in P\}$. And thus, $\text{succ}(\bigoplus \{\text{tr}(p) \mid p \in P\}) = \{\text{tr}(p) \mid p \in P\} = \text{tr}(\text{succ}(\bigoplus P))$.

The second inductive case is $p = q.r$. If q is not headless, then p would not be headless. So we deduce that q must be headless. A headless term must contain an outermost choice or a non-terminal, since terms can not be empty. So, $\text{urg}(q) \geq 1$. We observe that for all $s \in \mathbf{T}$, $\text{urg}(s) > 0$ implies $\text{urg}(\text{tr}(s)) = \text{urg}(s)$ and $\text{urg}(s) = 0$ implies $\text{urg}(\text{tr}(s)) \leq 1$. This is clear from the construction of $\text{tr}(\cdot)$: Only the urgencies of urgency 0 subterms change. These get replaced by urgency 1 terms, unless they are *err* or *skip*.

The first case is $\underline{q}.r$. So, we have $\text{urg}(q) \geq \text{urg}(r)$. When $\text{urg}(r) = 0$, then $\text{urg}(\text{tr}(q)) \geq 1 \geq \text{urg}(\text{tr}(r))$. Otherwise for q to be leading, $\text{urg}(q) \geq 1$ must hold and

$$\text{urg}(\text{tr}(q)) = \text{urg}(q) \geq \text{urg}(r) = \text{urg}(\text{tr}(r)).$$

In either case we have $\text{tr}(q).\text{tr}(r)$. By I.H., $\text{succ}(\text{tr}(q)) = \text{tr}(\text{succ}(q))$. This results in $\text{succ}(\text{tr}(q.r)) = \text{succ}(\text{tr}(q).\text{tr}(r)) = \text{tr}(\text{succ}(q).\text{tr}(r)) = \text{tr}(\text{succ}(q.r))$.

Now let $\underline{q}.\underline{r}$. Since the term q is headless, r has urgency $\text{urg}(r) > \text{urg}(q) \geq 1$. Let $r = s.\underline{x}.t$ for some action x . If we are strict, we also need to handle the cases where r equals to $s.\underline{x}$, $\underline{x}.t$, and \underline{x} . We omit them to avoid repetition. Since $\text{urg}(x) = \text{urg}(r) > 1$, the term x can only be a choice $\bigoplus P$ or a non-terminal A . Both of these terms are headless, so we apply the induction hypothesis and obtain $\text{succ}(\text{tr}(x)) = \text{tr}(\text{succ}(x))$. Because $\text{urg}(x) > \text{urg}(q.s)$, the leading subterm is

$\text{tr}(q).\text{tr}(s).\text{tr}(x).\text{tr}(t)$. Indeed, $\text{urg}(\text{tr}(q.s)) \leq \max(1, \text{urg}(q.s))$ and $\text{urg}(\text{tr}(q.s)) < \text{urg}(\text{tr}(x))$. Finally, we derive

$$\begin{aligned} \text{succ}(\text{tr}(q.\underline{x})) &= \text{succ}(\text{tr}(q.s.\underline{x}.t)) \\ &= \text{tr}(q.s).\text{succ}(\text{tr}(x)).\text{tr}(t) \\ &= \text{tr}(q.s).\text{tr}(\text{succ}(x)).\text{tr}(t) \\ &= \text{tr}(\text{succ}(q.s.\underline{x}.t)) \\ &= \text{tr}(\text{succ}(q.\underline{x})) \end{aligned}$$

□

Proof of Lemma 21. We show both directions by an induction on the number of moves Eve needs to win.

Forward Direction: Let $p \in \mathbf{T}$ and $r \in \text{QRun}_{\text{hd}(p)}$. For the base case, let Eve win O from p in 0 moves. Then, $\text{hd}(p) = p \in \mathbf{W}$ and $\text{tl}(p) = \text{skip}$. Because Eve wins, p can not contain err . Then, $r \in \text{QRun}_p$ must be of the form $(i, p_1).(p_1, p_2) \dots (p_{n-1}, p_n).\text{tr}(a_n) \dots \text{tr}(a_{k-1})$. We know that the word $w = a_0 \dots a_n \dots a_{k-1}$ that corresponds to $p \stackrel{*}{\approx} w$ has a run on the DFA for O . Per definition of QRun_p , the DFA runs w from i to p_n in $n - 1$ steps, and Eve can choose the remaining transitions to reach $p_k \in \mathcal{F}$.

For the inductive case, let Eve reach O from p in β moves. Per Lemma 22, this also holds from $w.q$, where $w = \text{hd}(p)$ and $q = \text{tl}(p)$. We show that $r \Downarrow \text{tr}(O)$ for some $r_w \in \text{QRun}_w$ and $r \stackrel{*}{\approx} r_w.\text{tr}(q)$. Per Lemma 22, showing this for one such r suffices. We first observe that Eve has a strategy to reach $r'_w.\text{tr}(q)$ where $r'_w \in \text{QRun}_w$. In case of $\text{urg}(q) \geq 2$ we have $r_w = r'_w$ and the leading subterm $r_w.\text{tr}(q)$ due to $\text{urg}(r_w) \leq 1 < \text{urg}(\text{tr}(q))$. Otherwise, $\text{urg}(\text{tr}(q)) = \text{urg}(q) = 1$ (it can't be 0 because $q = \text{tl}(p)$). If $r_w \in \text{QRun}_w \cap \mathbf{W}$, we already have $r_w.\text{tr}(q)$. So let $r_w \in \text{QRun}_w \setminus \mathbf{W}$, i.e. $\text{urg}(r_w) = 1$. Eve resolves each term of the form $\bigvee \{(q, r) \mid r \in \delta(q, a)\}$ for some $a \in \Sigma$ and extend the determined part of the run in QRun_w . We know that $w \not\approx \text{err}$, because Eve wins $w.q$. So, per definition of QRun_w , Eve can find fitting transitions. Exhaustive application of this strategy rewrites r_w to $r'_w \in \text{QRun}_w \cap \mathbf{W}$ for the desired $r'_w.\text{tr}(q)$.

We now show $r'_w.\text{tr}(q) \Downarrow \text{tr}(O)$. Let $\text{own}(w.q) = \text{Eve}$. The case $\text{own}(w.q) = \text{Adam}$ is dual. Remember $p = w.q$ and $w = \text{hd}(p)$, so we have $w.\underline{q}$. Then, there is a $w.q' \in \text{succ}(w.q)$ from which Eve wins in $\gamma < \beta$ moves. Since $w.q' \stackrel{*}{\approx} w.\text{hd}(q').\text{tl}(q')$, Lemma 22 tells us that Eve also wins from $w.\text{hd}(q').\text{tl}(q')$. We have $r'_w.\text{tr}(\text{hd}(q')) \in \text{QRun}_{w.\text{hd}(q')} = \text{QRun}_{\text{hd}(w.q')}$, so we can apply the induction hypothesis to see that Eve wins from $r'_w.\text{tr}(\text{hd}(q')).\text{tr}(\text{tl}(q'))$ in γ turns. The fact $r'_w.\text{tr}(\text{hd}(q')).\text{tr}(\text{tl}(q')) \stackrel{*}{\approx} r'_w.\text{tr}(q')$ and Lemma 22 imply that $r'_w.\text{tr}(q') \Downarrow \text{tr}(O)$. Since q is headless, Lemma 23 tells us $\text{succ}(\text{tr}(p)) = \text{tr}(\text{succ}(p))$ and thus $r'_w.\text{tr}(q') \in \text{succ}(r'_w.\text{tr}(q))$. Then, we also have $r'_w.\text{tr}(q) \Downarrow \text{tr}(O)$. This concludes this direction of the proof.

Backward Direction: Let $p \in \mathbf{T}$, $w = \text{hd}(p)$, $q = \text{tl}(p)$ and $O \subseteq \Sigma^*$. For the base case, let Eve win $r_w.\text{tr}(q)$ in 0 moves. Then, $\text{tr}(q)$ is a command and this is only possible if $q = \text{skip}$. Since $r_w \in \text{tr}(O)$, the run r_w is accepting in the DFA for O . So, $w = p \in O$.

For the inductive case, let Eve win $r_w.\text{tr}(q)$ in β moves for some $r_w \in \text{QRun}_w$. The first case is $r_w.\text{tr}(q)$. Then, there is a term $\text{tr}(a)$ in r_w , so Eve can simply choose the corresponding transition in the DFA to extend r_w to $r'_w \in \text{QRun}_w$ per definition of QRun_w . Eve then wins from $r'_w.\text{tr}(q)$ in $\gamma < \beta$ moves. We can apply the induction hypothesis to see that Eve wins from $w.q$, and per Lemma 22 from p . The second case is $r_w.\text{tr}(q)$. Let $\text{own}(r_w.\text{tr}(q)) = \text{Eve}$. Again, the case $\text{own}(r_w.\text{tr}(q)) = \text{Adam}$ is dual. Since q is headless, we have $\text{succ}(\text{tr}(q)) = \text{tr}(\text{succ}(q))$ per Lemma 23. Then, Eve wins from some $r_w.\text{tr}(q')$ where $q' \in \text{succ}(q)$ in $\gamma < \beta$ moves. Write $r_w.\text{tr}(q') \stackrel{\cong}{=} r_w.\text{tr}(\text{hd}(q')).\text{tr}(\text{tl}(q'))$. As in the previous direction, we have $r_w.\text{tr}(\text{hd}(q')) \in \text{QRun}_{w.\text{hd}(q')}$. Per definition, $\text{tl}(q')$ is headless. We can apply the induction hypothesis to get that Eve wins from $w.\text{hd}(q').\text{tl}(q') \stackrel{\cong}{=} w.q' \in \text{succ}(w.q)$. So Eve wins from $w.q \stackrel{\cong}{=} p$ as well, completing the proof. \square

Proof of Lemma 20 for $h = 1$:

Preliminary Facts: Before moving on to the computation of $\chi(\mathbf{C}_{\text{tr}(O)})$ for $h = 1$, we study $\text{SNF}_0^{\text{tr}(O)}$ more closely.

Lemma 24. *The syntactic monoid of $\text{tr}(O)$ is $\text{SNF}_0^{\text{tr}(O)} = \mathcal{Q}^2 \cup \{\text{err}, \text{skip}\}$. Furthermore, the $(p, q).(q, r) \equiv_{\text{tr}(O)} (p, r)$ and $(p, q).(s, r) \equiv_{\text{tr}(O)} \text{err}$ hold for all $p, q, r, s \in \mathcal{Q}$ with $q \neq s$.*

Proof. First we show that $\mathcal{Q}^2 \cup \{\text{err}, \text{skip}\} \subseteq \text{SNF}_0^{\text{tr}(O)}$, i.e. that these elements are pairwise not equivalent. It is clear that for any $(p, q) \in \mathcal{Q}^2$, $(i, p).(p, q).(q, f) \Downarrow O$, so we get $(p, q) \not\stackrel{\cong}{=}_{\text{tr}(O)} \text{err}$ and thus $(p, q) \not\stackrel{\cong}{=}_{\text{tr}(O)} \text{err}$ per soundness. It is also clear to see that for any $(p, q) \neq (s, r) \in \mathcal{Q}^2$, we have $(i, p).(p, q).(q, f) \Downarrow \text{tr}(O)$ but $(i, p).(s, r).(q, f) \not\Downarrow \text{tr}(O)$ fails to hold since $p \neq s$ or $q \neq r$. skip is the neutral element of the syntactic monoid and unless $|\mathcal{Q}| = 1$ it is different from any $(p, q) \in \mathcal{Q}^2$.

Now, we show $\text{SNF}_0^{\text{tr}(O)} \subseteq \mathcal{Q}^2 \cup \{\text{err}, \text{skip}\}$. These are exactly the terminal symbols, so it will suffice to show that the set $\mathcal{Q}^2 \cup \{\text{err}\}$ is closed under concatenation up to $\equiv_{\text{tr}(O)}$. Concatenations that involve err or skip elements are reduced by Axiom (M) to an element from $\mathcal{Q}^2 \cup \{\text{err}, \text{skip}\}$.

It remains to show the claimed congruences. For that, let $p, q, r, s \in \mathcal{Q}$ with $q \neq s$ and let $c[\bullet] = s.\bullet.t$ be a concatenative context with $s, t \in \mathbf{W}$. The first observation is that $c[(p, q).(r, s)]$ is losing for Eve, since she can not form a continuous run, i.e. she can not reach $\text{tr}(O)$. Then, per Axiom (S) we get $(p, q).(s, r) \equiv_{\text{tr}(O)} \text{err}$. An accepting, continuous run needs to have only three properties. Namely, it must start from i , it must not have discontinuities, and it must end at some $f \in \mathcal{F}$. But, $(p, q).(q, r)$ and (p, r) are both continuous while starting and ending at the same states. So we see that $c[(p, q).(q, r)]$ is an accepting run if and only if $c[(p, r)]$ is an accepting run. Thus, $(p, q).(q, r) \equiv_{\text{tr}(O)} (p, r)$. \square

We define $\delta[w] = \{(p, q) \in \mathcal{Q}^2 \mid p \xrightarrow{v} q, w \stackrel{\cong}{=} v \in \Sigma^+\}$ for all $w \in \mathbf{W}$ without err that are not skip .

Lemma 25. $\text{tr}(w) \equiv_{\text{tr}(O)} \bigvee \delta[w]$.

Proof. This is done by induction on $|w|$, the number of symbols contained in w . For the base case, we have $|w| = 1$. Then, $\text{tr}(w) = \bigvee \{(p, q) \mid q \in \delta(p, w)\}$. For the inductive case, let $w.v \in \mathbf{W}$ unequal to err or skip by $\stackrel{\#}{\equiv}$. Then,

$$\begin{aligned} \text{tr}(w.v) &= \text{tr}(w).\text{tr}(v) \\ &\equiv_{\text{tr}(O)} \bigvee \delta[w].\bigvee \delta[v] \end{aligned}$$

Applying (D1) and (D2) and flattening with (L4):

$$\begin{aligned} &\equiv_{\text{tr}(O)} \bigvee \{(p, q).(s, r) \mid \\ &\quad (p, q) \in \delta[w], (s, r) \in \delta[v]\} \end{aligned}$$

Using Lemma 24 and (L5) to remove inconsistent runs:

$$\begin{aligned} &\equiv_{\text{tr}(O)} \{(p, q).(q, s) \mid \\ &\quad (p, q) \in \delta[w], (q, r) \in \delta[v]\} \\ &\equiv_{\text{tr}(O)} \bigvee \delta[w.v] \end{aligned}$$

□

L.1 Iterating the Characteristic Terms

Recall that we have constructed $TNF_1^{\text{tr}(O)}$ to be the image of SNF_1^O under $\text{nf}_{\text{tr}(O)} \circ \text{tr}$. This poses a problem for $h = 1$. An Adam choice alone can not represent the image of a SNF_1^O term under normalization. Namely, for some $\bigvee_{i \in I} \bigwedge P_i \in SNF_1^{\text{tr}(O)}$, we might not have any $p \in \mathbf{T}$ with $\text{tr}(p) \equiv_{\text{tr}(O)} \bigwedge P_i$. So, if we naively apply the construction for $h = 1$, we are forced to iterate over a substantial subset of $SNF_1^{\text{tr}(O)}$. This results in a $\exp_2(\mathcal{O}(|O|^2))$ time complexity. For this reason, we compute $\chi(\mathbf{C}_{\text{tr}(O)})$ directly, by exploiting the Myhill-Nerode right-precongruence on the states of the DFA.

Definition 7. For any $p, q \in \mathcal{Q}$, we have $p \leq_N q$ if and only if for all $w \in \Sigma^*$, $p \xrightarrow{w} f \in \mathcal{F}$ implies $q \xrightarrow{w} f'$ for some $f' \in \mathcal{F}$.

We extend this to pairs of states $(p, p'), (q, q') \in \mathcal{Q}^2$. We let $(p, p') \leq_N (q, q')$ if and only if $p = q$ and $p' \leq_N q'$. We call a state pair (p, q) dead, if there is no run from q to a state in \mathcal{F} in the DFA for O . For some $c[\bullet] \in \mathbf{C}_{\text{tr}(O)}$, we call $\mathcal{S}(c[\bullet]) = \{p \in SNF_0^{\text{tr}(O)} \mid c[p] \Downarrow O\}$ the solution space of $c[\bullet]$. Towards computing $\chi(\mathbf{C}_{\text{tr}(O)})$, we show two important facts that expose the relationship between the extended \leq_N and solution spaces of contexts.

First, we show that there is a context $c[\bullet] \in \mathbf{C}_{\text{tr}(O)}$ for each not-dead $(p, q) \in \mathcal{Q}^2$, where $\mathcal{S}(c[\bullet])$ is the \leq_N -upward closure of (p, q) .

Lemma 26. For all $(p, q) \in \mathcal{Q}^2$ that is not dead, there is a $c[\bullet] \in \mathbf{C}_{\text{tr}(O)}$ with $\mathcal{S}(c[\bullet]) = \{(p', q') \in \mathcal{Q}^2 \mid (p, q) \leq_N (p', q')\}$.

Proof. We claim that $c[\bullet] = (q_0, p) \cdot \bullet \cdot \bigwedge_{w \in P} \text{tr}(w)$ where $P = \{v \mid q \xrightarrow{v} q_f \in \mathcal{F}\}$. Note that $P \neq \emptyset$, because (p, q) is not dead. For $(s, t) \in Q^2$ with $p \neq r$, $c[(s, t)]$ is losing for Eve, since $(q_0, p).(s, t) \equiv_O \text{err}$ per Lemma 24. Now assume $(p, s) \in Q^2$. We use Lemma 25 and get $c[(p, s)] \equiv_{\text{tr}(O)} (q_0, p).(p, s) \cdot \bigwedge_{w \in P} \forall \text{Rch}_w$. Then, Eve wins $c[(p, s)]$ if and only if Eve wins $(q_0, p).(p, s) \cdot \forall \text{Rch}_w$ for all $w \in P$. This is equivalent to w having a run from s to some $q_f \in \mathcal{F}$ for all $w \in P$. This is the definition of $q \leq_N s$. \square

Then, we show that for all $c[\bullet] \in \mathbf{C}_{\text{tr}(O)}$, the solution spaces are upward closed.

Lemma 27. *Let $c[\bullet] \in \mathbf{C}_{\text{tr}(O)}$ and $(p, q), (p', q') \in Q$ with $(p, q) \leq_N (p', q')$. Then, $(p, q) \in \mathcal{S}(c[\bullet])$ implies $(p', q') \in \mathcal{S}(c[\bullet])$.*

Proof of Lemma 27. W.l.o.g. let $(p, q), (p', q') \in Q^2$ with $(p, q) \leq_N (p', q')$ and $c[\bullet] = (s, t) \cdot \bullet \cdot \forall_{i \in I} \bigwedge_{w \in P_i} \forall \text{Rch}_w$. If $s \neq q_0$ or $t \neq p$ no matter which (p, r) is inserted to this context, Eve can never derive a continuous run. Now let $c[\bullet] = (q_0, p) \cdot \bullet \cdot \forall_{i \in I} \bigwedge_{w \in P_i} \forall \text{Rch}_w$ and assume that Eve wins from $c[(p, q)]$. Then, there is an $i \in I$ where for all $w \in P$, Eve wins $(q_0, p).(p, q) \cdot \forall \text{Rch}_w$. So, there is an $i \in I$ where for all $w \in P_i$, there is a run $q \xrightarrow{w} q_f$ for some $q_f \in \mathcal{F}$. The latter part of the statement and $(p, q) \leq_N (p', q')$ implies that there also is a run $q' \xrightarrow{w} q'_f$ for some $q'_f \in \mathcal{F}$. So, from $c[(p, q')]$, Eve plays the same $i \in I$. Let Adam play some $w \in P_i$. The resulting term is $(q_0, p).(p, q') \cdot \forall \text{Rch}_w$. So Eve chooses $(q', f') \in \text{Rch}_w$ to reach $(q_0, p).(p, q').(q', f')$ and win. \square

By applying Equation (3) along with Lemma 26 and Lemma 27, we see that a set $X \subseteq \text{SNF}_0^{\text{tr}(O)}$ is a solution space if and only if it is an upward closure that does not contain dead pairs. We can iterate through all left-sides $p \in Q$ and right-sides $Q' \subseteq Q$ and build the upward closures, therefore solution spaces, in $\exp_1(\mathcal{O}(|O|^2))$ time. Recall that for $c[\bullet] \in \mathbf{C}_{\text{tr}(O)}$, we had $\chi(c[\bullet]) = \bigwedge \mathcal{S}(c[\bullet])$ (Equation (2)). So, we can build $\{\bigwedge \mathcal{S}(c[\bullet]) \mid c[\bullet] \in \mathbf{C}_{\text{tr}(O)}\} = \{\chi(c[\bullet]) \mid c[\bullet] \in \mathbf{C}_{\text{tr}(O)}\} = \chi(\mathbf{C}_{\text{tr}(O)})$ in $\exp_1(\mathcal{O}(|O|^2))$ time.

M Hyperproperties and Urgency

Hyperproperties emerged as a unifying approach to information flow and security properties, which cannot be stated as classical safety or liveness properties on a single trace. A hyperproperty relates traces and is formulated over a set of traces rather than a single trace. A novel approach are logics like HyperLTL [29] to describe hyperproperties. We will stay more general and define an n -trace hyperproperty to be any DFA \mathcal{A} over $(\Sigma^n)^*$. Note that we consider finite traces rather than infinite ones. In particular, a hyperproperty takes about the relation of traces and does so in a highly synchronized manner: A set of traces is accepted by \mathcal{A} ensures that they all share the exact same length. To model different traces of the same system so that the resulting observations are synchronized is

a common problem with hyperproperties. We will not tackle this issue here, but assume that the traces of a system \mathcal{K} are partitioned into sets of same length. That means that the set of traces $L(\mathcal{K}) \subseteq \Sigma^*$ is partitioned into sets of traces $\text{tr}_l(\mathcal{K}) \subseteq \Sigma^l$ of length $l \in \mathbb{N}$.

Definition 8. *Let \mathcal{K} be a system (Kripke structure) with trace-set $L(\mathcal{K}) \subseteq \Sigma^*$. An n -trace hyperproperty \mathcal{A} is satisfied by \mathcal{K} , $\mathcal{K} \models \mathcal{A}$, if there is $l \in \mathbb{N}$ such that*

$$\exists w_1 \in \text{tr}_l(\mathcal{K}) \forall w_2 \in \text{tr}_l(\mathcal{K}) \dots Q w_n \in \text{tr}_l(\mathcal{K}). \prod_{i=1}^n w_i \in \mathcal{L}(\mathcal{A})$$

The definition might be non-intuitive at first glance: Only for a single length l , the set $\text{tr}_l(\mathcal{K})$ has to satisfy the hyperproperty. Further, every hyperproperty begins to quantify with an \exists quantifier. This makes it impossible to formulate hyperproperties from the $\forall\exists$ -fragment of HyperLTL. However, the ability to deterministically decide hyperproperties of the above shape is already sufficient to decide any hyperproperty, including ones that start with \forall , for which we check its negation.

We show how to model check a hyperproperty by translating the system \mathcal{K} and the hyperproperty \mathcal{A} into a term p and an objective O , such that \mathcal{A} is satisfied by \mathcal{K} if and only if $p \Downarrow O$. Intuitively, we model the quantifiers in Definition 8 by the players. Each quantifier is resolved in one level of urgency throughout the whole term. Only then, and in knowledge of the resolution of previous quantifiers, a player resolves the next quantifier. Formally, we define the terms p_i for $i \in [1, n+1]$ over the alphabet $\Sigma = \delta$ by induction:

$$p_{n+1} = \text{skip} \quad p_i = \begin{cases} \bigvee_{n-i+1} \delta.p_{i+1} & i \text{ odd} \\ \bigwedge_{n-i+1} \delta.p_{i+1} & i \text{ even} \end{cases}$$

The final term p is a single non-terminal $p = A$ with $E(A) = \text{skip} \Psi p_1.A$. This lets Eve choose the first trace and while doing so, she also fixed the length of the considered traces.

For the objective O , we want to capture the hyperproperty \mathcal{A} , which already is a DFA. For the transformation to O we only need to transform the input alphabet from δ^n to δ and make sure the players choose actual traces of the system \mathcal{K} . Technically, for a word $w \in (\Sigma^n)^*$, we obtain $w_i \in \Sigma^*$ by restricting w to the i th component. We define the flattening flat of a word in $(\Sigma^n)^*$ to Σ^* inductively by $\text{flat}(\text{skip}) = \text{skip}$ and $\text{flat}((t_1, \dots, t_n).v) = t_1 \dots t_n.\text{flat}(v)$. We set the set of correct runs on component i by $O_i = \{w \in (\Sigma^n)^* \mid w_i \in L(\mathcal{K})\}$.

$$O = \text{flat}(\mathcal{L}(\mathcal{A})) \cap \bigcap_{i \text{ odd}} O_i \cup \bigcup_{i \text{ even}} \overline{O_i}$$

Note that $|O|$ is linear in $|\mathcal{A}| + |\mathcal{K}|$.

Theorem 9. *$\mathcal{K} \models \mathcal{A}$ if and only if $p \Downarrow O$.*

Proof. By construction, $p \Downarrow O$ if and only if there is $l \in \mathbb{N}$ with $(p_1)^l \Downarrow O$, because unrolling A an infinite number of times yields the win to Adam.

Next we inspect the terms $p(w_1, \dots, w_m)$ that can occur whenever urgency $n - m$ is next to be resolved. Here, w_i are from $\text{tr}_l(\mathcal{K})$ and for odd i are chosen corresponding to the \exists -quantifiers. By construction, $p(w_1, \dots, w_m)$ has shape

$$t_{1,1}t_{2,1} \dots t_{m,1}p_{m+1} \dots t_{1,l}t_{2,l} \dots t_{m,l}p_{m+1}.$$

where $w_i = t_{i,1} \dots t_{i,l}$.

We prove the statement: Eve wins $p(w_1, \dots, w_m) \Downarrow O$ if and only if m is even and there is $w_{m+1} \in \text{tr}_l(\mathcal{K})$ with $p(w_1, \dots, w_m, w_{m+1}) \Downarrow O$, or m is odd and for all $w_{m+1} \in \text{tr}_l(\mathcal{K})$ holds $p(w_1, \dots, w_m, w_{m+1}) \Downarrow O$.

If m is even, then Eve can choose a sequence $w_{m+1} \in \delta^l$ to transform $p(w_1, \dots, w_m)$ into $p(w_1, \dots, w_m, w_{m+1})$. If she chooses a word outside $\text{tr}_l(\mathcal{K})$ she loses due to O_{m+1} . Thus, she wins if and only if there is $w_{m+1} \in \text{tr}_l(\mathcal{K})$ she can choose and $p(w_1, \dots, w_m, w_{m+1}) \Downarrow O$.

The case of m odd is similar.

Finally, $n = m$ makes $p_{m+1} = \text{skip}$, so $p(w_1, \dots, w_n) \Downarrow O$ if and only if $p(w_1, \dots, w_n) \in \text{flat}(\mathcal{L}(\mathcal{A}))$ by definition. \square

The complexity of checking an n -trace hyperproperty $\mathcal{K} \models \mathcal{A}$ using this approach is $(2n - 1) - \text{EXPTIME}$ bounded (Corollary 1). This is still far from optimal, considering that n -trace hyperproperties can be decided in $(n - 1) - \text{EXSPACE}$ [29]. The overhead stems from the shape of normalforms, which allow for Adam and Eve choices in each layer of urgencies, while our approach produces only one type of non-determinism for each urgency. We intend to investigate this special case in a separate study.

N Model Checking Hyperproperties for Recursive Programs

Checking hyperproperties on pushdown systems is known to be undecidable in the general case [43]. But showing it undecidable does not satisfy the desire to check recursive programs against hyperproperties. Indeed, a first approach was proposed directly in [43], where one type of quantifier has to work with finite state approximations instead of the actual system. We take a different route to restrict the general setting. To motivate the restriction, consider the algorithm in Listing 1.1 for multiplication of high-bit numbers.

```

1 bit[] KM(bit[] x, bit[] y) {
2     if (len(x) < 64) return (int) x * (int) y;
3
4     mid = len(x) / 2;
5     x1 = x[:mid];
6     y1 = y[:mid];
7     x2 = x[mid+1:];
8     y2 = y[mid+1:];

```

```

9
10   z2 = KM(x1, y1);
11   z0 = KM(x2, y2);
12   z1 = KM(x1 + x2, y1 + y2) - z2 - z0;
13
14   return (z2 << 2mid) + (z1 << mid) + z0;
15 }

```

Listing 1.1: Karatsuba multiplication.

The algorithm performs a high-bit multiplication of two numbers x and y . Instead of naively multiplying 64-bit windows of the bit streams, the algorithm recursively splits the inputs in half. It performs 3 multiplications on the half-sized integers and adds them together for the final result. The actual algorithm needs some more care for bit-overflows, but we shall ignore it here.

Algorithms like Listing 1.1 are used by security protocols like `OpenSSL` [51]. A known weakness of security protocols are timing-based attacks [52,53]. In these attacks, one does not run the program once to derive a secret. Instead, we measure its execution time over multiple runs with different controllable inputs and deduct secret values used by the program from the measured execution times. So, safety from timing attacks can be obtained by requiring a 2-trace hyperproperty on `KM`: For all traces w, v of `KM`, their execution time does not differ. We do not go into modelling details of how to phrase this property as a DFA \mathcal{A} . But we make one crucial discovery when we want to compare multiple runs of the Karatsuba algorithm: Its recursion depth is only dependent on the length of x and y , a parameter that is very often public knowledge. Even more important, runs of different recursion depth have close to no chance for a similar execution time. The previously stated hyperproperty would very likely not be satisfied. But when the recursion depth parameter is usually known, we instead want to ask for a different hyperproperty to hold: For all recursion depths d , and for all traces w, v with recursion depth d , their execution time does not differ. More generally speaking, the traces w and v agree on their *recursive structure*. We utilize this observation and restrict our check for a hyperproperty \mathcal{A} to sets of traces that agree on their recursive structure.

N.1 Recursive Programs

We consider the following simple language, where we abstract away from the actual commands and focus on the recursion principle of the language.

$$\begin{aligned}
 c &::= a \mid \mathbf{f}() \mid c.c \mid \mathbf{if} \ e \ c:c \\
 \mathbf{f} &::= \mathbf{f}() \rightarrow c.\mathbf{return}
 \end{aligned}$$

We have commands $a \in \Sigma$ to modify the state, function calls $\mathbf{f}()$ and concatenation. Parameters and return values are passed to/from \mathbf{f} by state-manipulation. A program $\mathcal{P} = (\mathcal{F}, E)$ is a set of function symbols \mathcal{F} with a distinguished initial function $\mathbf{main} \in \mathcal{F}$ and the function definitions E .

N.2 Semantics and Recursion Structure

We assume the domain of booleans $\mathcal{D} = [\mathcal{V} \rightarrow \mathbb{B}]$. The semantics are kept arbitrary for expressions $\llbracket e \rrbracket : \mathcal{D} \rightarrow \mathbb{B}$ and commands $\llbracket a \rrbracket : \mathcal{D} \rightarrow \mathcal{D}$. Semantics of function calls and concatenation are as expected. **if**-statements branch left on non-zero evaluation and right on zero. For the further development, we also require a special command **assume** $e \in \Sigma$. Its semantics operate on a fresh variable η and are tuned to observe whether an **assume** command failed.

$$\llbracket \mathbf{assume} \ e \rrbracket(d) = \begin{cases} d & \llbracket e \rrbracket(d) = \mathbf{true} \\ d[\eta \mapsto \mathbf{true}] & \llbracket e \rrbracket(d) = \mathbf{false} \end{cases}$$

A *branching* of \mathcal{P} is a finite tree t (a prefix closed subset of \mathbb{N}^*) with a labelling **cmd**. The labelling assigns nodes of t commands a , function calls \mathbf{f} , and **return** statements. The root node is labeled with $\mathbf{cmd}(\varepsilon) = \mathbf{main}$. Nodes tn labelled with $\mathbf{cmd}(tn) = a$ or **return** are leaves. A $\mathbf{cmd}(tn) = \mathbf{f}(n)$ -labelled node has children corresponding to one branch of their body c . The set of branches $\mathbf{br}(c)$ is

$$\begin{aligned} \mathbf{br}(a) &= \{a\} & \mathbf{br}(\mathbf{f}()) &= \{\mathbf{f}\} & \mathbf{br}(c_1.c_2) &= \mathbf{br}(c_1).\mathbf{br}(c_2) \\ \mathbf{br}(\mathbf{if} \ e \ c_1 : c_2) &= \mathbf{assume} \ e.\mathbf{br}(c_1) \cup \mathbf{assume} \ !e.\mathbf{br}(c_2) \end{aligned}$$

So if $w \in \mathbf{br}(c)$ is the chosen branch, then tn has $|w|$ children labelled by the corresponding a , \mathbf{f} , or **return** in w .

A *trace* of \mathcal{P} is a pair $(\mathbf{cmd}, \mathbf{st})$ where \mathbf{cmd} is a branching with domain t and $\mathbf{st} : t \rightarrow \mathcal{D}$ is a state assignment. Every node tn is labelled by $\mathbf{st}(tn) \in \mathcal{D}$, the state of the program before execution of $\mathbf{cmd}(tn)$, $\mathbf{st}(\varepsilon) \in \mathcal{D}$ is the input state. Consider node tn with $\mathbf{cmd}(tn) = \mathbf{f}$, so the children are one branch of its defining body. Its leftmost child $tn.0$ inherits the state, $\mathbf{st}(tn.0) = \mathbf{st}(tn)$. The rightmost child $tn.m$ has label $\mathbf{cmd}(tn.m) = \mathbf{return}$ and yields back $\llbracket \mathbf{f} \rrbracket(\mathbf{st}(tn)) = \mathbf{st}(tn.m)$. Intermediate children $tn.i$ just carry out their semantics to the next node by $\mathbf{st}(tn.i+1) = \llbracket \mathbf{cmd}(tn.i) \rrbracket(\mathbf{st}(tn.i))$.

A *recursion structure* r for \mathcal{P} is a tree labelling $r : t \rightarrow \mathcal{F}$. The internal of a tree t is the set $\mathbf{int}(t) = \{tn \mid tn.0 \in t\}$.

Definition 9. A *branching* \mathbf{cmd} has *recursion structure* r if $\mathbf{int}(\mathbf{dom} \ \mathbf{cmd}) = \mathbf{dom} \ r$ and \mathbf{cmd} and r coincide on $\mathbf{dom} \ r$.

The recursive structure thus identifies runs that differ only in \mathbf{st} and the actual commands executed, but the function calls are exactly the same. We say $(\mathbf{cmd}, \mathbf{st})$ is an r -trace when \mathbf{cmd} has recursion structure r . A recursion structure r is proper, if there is an r -trace of \mathcal{P} .

N.3 Checking Hyperproperties for similar Recursion Structure

We define the yield $\mathbf{yld}(t)$ of a tree as usual. The set of trace-observations with recursion structure r is the set $\mathbf{tr}_r(\mathcal{P}) = \{\mathbf{yld}(\mathbf{cmd}) \mid (\mathbf{cmd}, \mathbf{st}) \text{ is an } r\text{-trace of } \mathcal{P}\} \subseteq \Sigma^*$.

Definition 10. *Let \mathcal{A} be an n -trace hyperproperty. A program \mathcal{P} satisfies \mathcal{A} , $\mathcal{P} \models \mathcal{A}$, if there is a proper recursion structure r with*

$$\exists w_1 \in \text{tr}_r(\mathcal{P}) \forall w_2 \in \text{tr}_r(\mathcal{P}) \dots Q w_n \in \text{tr}_r(\mathcal{P}). \prod_{i=1}^n w_i \in \mathcal{L}(\mathcal{A})$$

To present our approach, we assume that programs are prefix-branching. We call a program \mathcal{P} *prefix-branching* when function calls are never succeeded by commands $a \in \Sigma$. Formally, we assume the program code stems from the following, slightly different grammar:

$$\begin{aligned} f & ::= \mathbf{f}() \mid f.f \\ c & ::= a \mid c.c \\ pb\text{-}c & ::= f \mid c.pb\text{-}c \mid \text{if } e \text{ } pb\text{-}c : pb\text{-}c \\ \mathbf{f} & ::= \mathbf{f}() \text{ -> } pb\text{-}c.\text{return} \end{aligned}$$

Note, that this is not only a presentational decision. While every program \mathcal{P} has an equivalent representation in prefix-branching form, the translation incurs more functions and thus, fixing a recursion structure may fix more behavior (e.g. branching beyond function calls) than desired. This means that we actually restrict our class of programs by enforcing prefix-branching. We do not exactly know, how impactful the restriction is for practical code, but our example (Listing 1.1) exerts the desired structure (up to parameter passing).

Theorem 10. *$\mathcal{P} \models \mathcal{A}$ is decidable for prefix-branching \mathcal{P} .*

Similar as for the finite state case, we will translate \mathcal{P} and \mathcal{A} into a term p and an objective O . The intent is to first fix a recursion structure, and then replay the call of a function \mathbf{f} with function body c by the rewriting of a non-terminal \mathbf{f} into a term $p(\mathbf{f})$. Term $p(\mathbf{f})$ basically chooses the branch through c . For prefix-branching, a branch is a word from $\Sigma^* \mathcal{F}^*$. With a fixed recursive structure also the sequence of function calls in c is fixed. For a fixed sequence of function calls $f = \mathbf{f}_1 \dots \mathbf{f}_m$ the set of available prefixes is captured by $\text{br}_f(c) \subseteq \Sigma^*$ with

$$\text{br}_f(c).f = \text{br}(c) \cap \Sigma^*.f.$$

We assume that all command prefixes (the Σ^* part) in $\text{br}_f(c)$ share the same length, $\text{br}_f(c) \subseteq \Sigma^l$ for some $l \in \mathbb{N}$. Similar to the finite state case this is a modelling issue for synchronization towards \mathcal{A} . It can be achieved by introducing skips to branches too short. Since $\text{br}(c)$ is finite, there is only a finite set of occurring function call sequences $f(c) = \{f \mid \text{br}_f(c) \neq \emptyset\}$. Finally, we define $p(\mathbf{f})$ in a term game (\mathcal{F}, E) with $E(\mathbf{f}) = p(\mathbf{f})$.

$$\begin{aligned} p_{n+1}^{c,f}(w_1, \dots, w_n) &= \text{flat}(w_1, \dots, w_n) \\ p_i^{c,f}(w_1, \dots, w_{i-1}) &= \begin{cases} \bigvee_{w_i \in \text{br}_f(c)} p_{i+1}^{c,f}(w_1, \dots, w_i) & i \text{ odd} \\ \bigwedge_{w_i \in \text{br}_f(c)} p_{i+1}^{c,f}(w_1, \dots, w_i) & i \text{ even} \end{cases} \\ p(\mathbf{f}) &= \bigvee_{f \in f(c)} p_1^{c,f}().f \end{aligned}$$

It remains to model the objective O . As before, we construct sets $O_i \subseteq (\Sigma^n)^*$ for $i \in [1, n]$.

$$O = \text{flat}(\mathcal{L}(\mathcal{A})) \cap \bigcap_{i \text{ odd}} O_i \cup \bigcup_{i \text{ even}} \overline{O_i}$$

$$O_i = \{w \in (\Sigma^n)^* \mid \llbracket w_i \rrbracket(\eta) = \text{true}\}$$

A command sequence $w \in \Sigma^*$ operates on the specific finite domain \mathcal{D} . These language is regular because \mathcal{D} is finite.

Theorem 11. $\text{main} \Downarrow O$ if and only if $\mathcal{P} \models \mathcal{A}$.

Proof Sketch. Notice that the set of all reachable word terms from $p_1^{c,f}$ is exactly $\text{flat}(\text{br}_f(c)^n)$. So by construction, playing from **main** until the term has no more non-terminals yields a term p_r for some recursion structure r , where the reachable word terms form precisely the set $\text{flat}(\text{br}_r(\mathcal{P})^n)$, where $\text{br}_r(\mathcal{P}) = \{\text{yld}(\text{cmd}) \mid \text{cmd} \text{ is an } r\text{-branching}\}$. As before, the alternative is that there is an infinite rewriting of non-terminals, in which case Eve loses. Thus, $\text{main} \Downarrow O$ if and only if there is r such that $p_r \Downarrow O$.

Next we again inspect the terms $p(w_1, \dots, w_m)$ that can occur whenever urgency $n - m$ is next to be resolved (at first, we have $p_r = p()$ with $m = 0$). This time, w_i are from $\text{tr}_r(\mathcal{P})$ and for odd i are chosen corresponding to the \exists -quantifiers. By construction, $p(w_1, \dots, w_m)$ has shape

$$p_{m+1}^{c^1, f^1}(w_1^1, \dots, w_m^1) \cdot \dots \cdot p_{m+1}^{c^l, f^l}(w_1^l, \dots, w_m^l),$$

where $c^1 \dots c^l$ is the depth-first left to right traversal of r .

We prove: $p(w_1, \dots, w_m) \Downarrow O$ if and only if m is even and there is $w_{m+1} \in \text{tr}_r(\mathcal{P})$ with $p(w_1, \dots, w_m, w_{m+1}) \Downarrow O$, or m is odd and for all $w_{m+1} \in \text{tr}_r(\mathcal{P})$ we have $p(w_1, \dots, w_m, w_{m+1}) \Downarrow O$.

If m is even, then Eve can choose a sequence of branches $w_{m+1}^i \in \text{br}_{f^i}(c^i)$ which together form the yield of a branching $w_{m+1} = w_{m+1}^1 \dots w_{m+1}^m \in \text{br}_r(\mathcal{P})$. She does so by transforming $p(w_1, \dots, w_m)$ into $p(w_1, \dots, w_m, w_{m+1})$. If she chooses a branching which has no trace $w_{m+1} \notin \text{tr}_r(\mathcal{P})$ she loses due to O_{m+1} . Thus, she wins if and only if there is $w_{m+1} \in \text{tr}_r(\mathcal{P})$ to choose and we have $p(w_1, \dots, w_m, w_{m+1}) \Downarrow O$. The case of m odd is similar.

The final case, where $n = m$ holds, implies that the immediate terms are $p_{m+1}^{c^i, f^i}(w_1^i, \dots, w_m^i) = \text{flat}(w_1^i, \dots, w_m^i)$ and by definition of \Downarrow for word terms, $p(w_1, \dots, w_n) \Downarrow O$ if and only if $p(w_1, \dots, w_n) \in \text{flat}(\mathcal{L}(\mathcal{A}))$. \square

O Specialized Decision Algorithm for Weak Terms with Linear Grammars

In this section, we show that a certain fragment, weak linear grammars/terms, admits more efficient decision algorithms. For maximal urgency h , the complexity goes down from $2h - 1 - \text{EXPTIME}$ to $(h - 2) - \text{EXSPACE}$. In the following, we

first build up to the definition of weak linearity. Then, we show how observations in linear terms with weak linear grammars can be decided in $(h-2)$ -EXPSPACE.

The weak fragment of urgency imposes a restriction on available choice operators per urgency. But as Axiom (N) suggests, choice operators can potentially combine to mimic the effect of disallowed choice operators. For example, let \Downarrow be disallowed, however let \Uparrow and \Downarrow be both allowed. Then, we observe

$$\Uparrow (p \Downarrow q) \stackrel{(N)}{\equiv} \Uparrow (p \Downarrow q) \equiv (p \Downarrow q)$$

We avoid this effect by further restricting the structure of weak terms. To this end, we define the internal urgency of a term. Intuitively, the internal urgency $\text{inurg}(p)$ of a term p is the largest urgency that appears in a choice operator.

$$\text{inurg}(p \oplus q) = \max\{\text{inurg}(p), \text{inurg}(q), u\} \quad \text{inurg}(p.q) = \max\{\text{inurg}(p), \text{inurg}(q)\}$$

The internal urgency of non-terminals requires a more subtle definition. This is because for $A \in N$, even if the term $E(A)$ does not contain choice operators from a certain urgency, $E(A)$ may invoke some other B for which $E(B)$ might recreate that urgency. To obtain a sound definition, we have to consider $E(A)$ for all $A \in N$ at once. We define $\text{inurg}(A)$ for each $A \in N$ to be the smallest assignment that satisfies $\text{inurg}(A) = \text{inurg}(E(A))$, along with the properties of $\text{inurg}(\cdot)$ defined above. Intuitively, this means that non-terminals with internal urgency u can only use choice operators of urgency u and non-terminals with internal urgency at most u . With this at hand, we define urgency-alignment. Intuitively, urgency-alignment of p requires that the urgency of all terms derivable not greater than $\text{urg}(p)$. Formally, we use an inductive definition. Terms $p = A$ or $p \in \Sigma$ are always urgency-aligned. A term $p = q.r$ is urgency-aligned if q and r are both urgency aligned, and a term $p = q \oplus r$ is urgency aligned if both q and r are urgency-aligned with $\text{inurg}(q), \text{inurg}(r) \leq u$. A grammar is urgency-aligned if $E(A)$ is urgency-aligned for all $A \in N$.

We are now ready to define weakness and linearity. We say that a term p is *weak*, if it, and its grammar are urgency-aligned, with p and $E(A)$ for all $A \in N$ only containing choice operators from $\{\Psi \mid u \text{ even}\} \cup \{\Uparrow \mid u \text{ odd}\}$. Weak terms are normalized into (O -specialized) weak normal forms WNF_u in $(u-1)$ -EXPTIME Lemma 28. The base case $WNF_0 = SNF_0^O$ stays the same. For positive urgencies, odd u and even v :

$$WNF_u = \{\Uparrow P \mid \emptyset \neq P \subseteq WNF_{u-1}\} \quad WNF_v = \{\Psi P \mid \emptyset \neq P \subseteq WNF_{v-1}\}$$

Linearity restricts the recursive behaviour of the term and the grammar. The syntactic construction of a *linear* term l is given below, where p represents an arbitrary term that does not contain non-terminals. We call a grammar (N, E) linear if $E(A)$ is linear for all $A \in N$.

$$l ::= A \mid p \mid l \oplus l \mid l.p \mid p.l$$

O.1 Weak Normal Form

We present a slightly modified normalization algorithm that takes advantage of the restricted expressiveness of weak grammars and weak terms. More specifically, we modify the functions $\text{nfchoice}_O(\cdot)$ and $\text{nfconc}_O(\cdot)$ to this end. The modified functions return terms with urgencies no larger than the urgency of the input and preserve weakness. The proofs remain largely the same. The normalization process of non-terminals remains the same in relation to $\text{nfchoice}_O(\cdot)$ and $\text{nfconc}_O(\cdot)$. In this case, we extract upper bounds by employing the lattice height of WNF_h . We prove the lemma below in this section. Note that for $h \geq 2$, the power \cdot^5 gets absorbed by $\mathcal{O}(\cdot)$.

Lemma 28 (Modifies Lemma 7). *Given a weak term p , a weak grammar (N, E) , and a regular objective $O \subseteq \Sigma^*$ given as a DFA, we can compute $\text{nf}_O(p) \in WNF_h$ with $\text{nf}_O(p) \equiv_O p$ in time $(|p| + |E||N|) \cdot \exp_{h-1}(\mathcal{O}(|SNF_0^O|^5))$.*

We first handle $\text{nfchoice}_O(\cdot)$.

Lemma 29 (Modifies Lemma 10). *Let $P \subseteq WNF_u$, and $r = \bigoplus P$. We can construct its normalform $\text{nfchoice}_O(r) \in WNF_u$ with $\text{nfchoice}_O(r) \equiv_O r$ in time $\mathcal{O}(|r|) \leq \exp_{h-1}(\mathcal{O}(|SNF_0^O|))$.*

Proof. We know for each $p' \in P$ we have $p' = \bigoplus Q_{p'}$. We simply construct the union $r' = \bigoplus \bigcup_{p' \in P} Q_{p'}$ in time $\mathcal{O}(|r|) \leq \exp_{h-1}(\mathcal{O}(|SNF_0^O|))$. We observe $r \equiv r'$ by (L4). This concludes the proof. \square

Lemma 30 (Modifies Lemma 11). *For any $p, q \in WNF_u$, we can find $\text{nfconc}_O(p.q) \in WNF_u$ with $\text{nfconc}_O(p.q) \equiv_O p.q$ in time $\exp_{h-1}(\mathcal{O}(|SNF_0^O|^4))$.*

Proof. The base case $h = 1$ is handled similarly to the inductive case, so we only show the inductive case. We first use the distribution axioms on the concatenation.

$$\begin{aligned}
 p.q &= (\bigoplus P).(\bigoplus Q) \\
 &\stackrel{\text{(D2)}}{\equiv} \bigoplus_{p' \in P} p'.(\bigoplus Q) \\
 &\stackrel{\text{(D1)}}{\equiv} \bigoplus_{p' \in P} \bigoplus_{q' \in Q} p'.q'.
 \end{aligned}$$

The resulting term allows for an application of the induction hypothesis on the subterms with lower urgency for $h > 1$.

$$\stackrel{\text{I.H.}}{\equiv} \bigoplus_{p' \in P} \bigoplus_{q' \in Q} \text{nfconc}_O(p'.q').$$

The distribution causes the term to jump to at most the square of its previous size. In any case, there are at most $\exp_{h-1}(\mathcal{O}(|SNF_0^O|))^2$ choice operands

of lower urgency. For $h = 1$, the concatenation is resolved by monoid multiplication which takes $\mathcal{O}(|SNF_0^O|)^2$ time. This results in an overall complexity of $\exp_{h-1}(\mathcal{O}(|SNF_0^O|))^4$. For $h > 1$ we employ the induction hypothesis and see that the normalization of each subterm takes $\exp_{h-2}(\mathcal{O}(|SNF_0^O|))^2$ time. So in total we again need at most $\exp_{h-1}(\mathcal{O}(|SNF_0^O|))^4$ steps. \square

Lemma 28 follows as a corollary of two observations. First, each iteration of the least fixed point calculation will require at most $|E| \cdot \exp_{h-1}(\mathcal{O}(|SNF_0^O|))^4$ steps for normalization. Second, we only need $|N| \cdot |WNF_{h-1}|$ iterations towards a fixed point, since the height of the lattice $WNF_h^{|N|}$ is bounded by this number. Replacing the non-terminals in p by their normalizations, the term grows by at most a factor of $\exp_{h-1}(\mathcal{O}(|SNF_0^O|))$. This is then normalized in $|p| \cdot \exp_{h-1}(\mathcal{O}(|SNF_0^O|))^5$. Resulting in an overall time complexity of at most $(|p| + |E||N|) \cdot \exp_{h-1}(\mathcal{O}(|SNF_0^O|)^5)$.

O.2 Deciding the winner for weak linear terms

The main claim of the section is given in the lemma below.

Lemma 31. *Let (N, E) be a weak linear grammar with maximal urgency h , p a weak linear term, and $O \subseteq \Sigma^*$ a regular objective. Then $p \Downarrow O$ can be decided in $(\max\{|E|^2, |p|^2\} \cdot |E| \cdot |N| \cdot \exp_{h-2}(\mathcal{O}(|SNF_0^O|)))^4$ space.*

We fix a weak linear grammar (N, E) with maximal urgency h , and a regular objective O for the rest of this section. A naive normalization of p into WNF_h gives us a complexity of $(h - 1) - \text{EXPTIME}$ Lemma 28. Here, the size of the normal form terms is the main bottleneck, which can grow up to $\exp_{h-1}(\mathcal{O}(|O|))$. To circumvent this problem, the algorithm avoids a complete normalization. Instead, it aims to non-deterministically guess a small-size portion $q \in WNF_{h-1}$ of the normalization, which suffices to show $p \Downarrow O$. The existence of this is guaranteed by the lemma below. Because of it admits a direct check, we use the domination preorder \sqsubseteq_O instead of \sqsubseteq_O or \preceq_O .

Lemma 32. *Let p be a weak term. Then, there is a $q \in WNF_{h-1}$ with $\Psi q \sqsubseteq_O \text{nf}_O(p)$, if h is even and $\wedge q \supseteq_O \text{nf}_O(p)$ if h is odd, where $p \Downarrow O$ if and only if $q \Downarrow O$.*

Proof. Assume h is even. The case with odd h is dual. Let p be a weak term. Then, there is a term $p' \in WNF_h$ with $\text{nf}_O(p) = p'$. The definition of WNF_h tells us $p' = \Psi P$ for some $P \subseteq WNF_{h-1}$. We observe that $p' \Downarrow O$ holds if and only if there is a $r \in P$ with $r \Downarrow O$. Using the definition of \sqsubseteq_O , we also get $\Psi r \sqsubseteq_O p'$ for any $r \in P$. Consider the case of $p' \Downarrow O$, thus there is $p'' \in P \subseteq WNF_{h-1}$ and $p'' \sqsubseteq_O p'$ as well as $p'' \Downarrow O$ holds. Similarly, if $p' \not\Downarrow O$, any element $p'' \in P \subseteq WNF_{h-1}$ has $p'' \not\Downarrow O$ and $p'' \sqsubseteq_O p'$. \square

It remains to show that this fragment can be effectively guessed. This is stated by the lemma below. The normalization process tells us that for $k =$

$|N| \cdot |WNF_{h-1}|$, where $|WNF_{h-1}| \leq \exp_h(|O|)$, we get $p^{(k)} \equiv p$. Here, $p^{(i)}$ is p with non-terminals replaced by the i 'th iteration towards a fixed point as discussed in Appendix O.1, the syntax is borrowed from Appendix E.2. Initializing $i = k$ in the lemma below yields the desired $(h - 2) - \text{EXPSPACE}$ decidability result. To achieve this complexity, the algorithm avoids maintaining a call stack.

Lemma 33. *Let p be a weak linear term with urgency h , let $t \in WNF_{h-1}$ and let $i \in \mathbb{N}$. If p does not contain non-terminals, $t \subseteq_O \text{nf}_O(p)$ can be checked in space*

$$(|p|^2 \cdot (|E| \cdot |N| \cdot \exp_{h-2}(\mathcal{O}(|SNF_0^O|)))^2)^2.$$

If conversely, p contains non-terminals, $t \subseteq_O p^{(i)}$ can be checked in space

$$(\mathcal{O}(\log(i)) + \max\{|E|^2, |p|^2\} \cdot (|E| \cdot |N| \cdot \exp_{h-2}(\mathcal{O}(|SNF_0^O|)))^2)^2.$$

Proof. We only handle the case where h is even, the case with an odd h is dual. Our algorithm is a non-deterministic space bounded algorithm. We show that the check is possible in non-deterministic $|p|^2 \cdot \text{Norm}_O^2$ space if p does not contain non-terminals, and non-deterministic $\mathcal{O}(\log(i)) + \max\{|p|^2, |E|^2\} \cdot \text{Norm}_O^2$ space if p contains non-terminals, where $\text{Norm}_O = |E| \cdot |N| \cdot \exp_{h-2}(\mathcal{O}(|O|))$. These complexities can be translated into deterministic space by Savitch's Theorem [54]. This yields the deterministic complexities $(|p|^2 \cdot \text{Norm}_O^2)^2$ without non-terminals and $(\mathcal{O}(\log(i)) + \max\{|p|^2, |E|^2\} \cdot \text{Norm}_O^2)^2$ with non-terminals.

We proceed by an outer induction on $i \in \mathbb{N}$ and an inner structural induction on p . The base case is $p \in \Sigma \cup \{\text{err}, \text{skip}\}$ or $p \in N$. The check $t \subseteq_O \text{nf}_O(p')$ for $p' \in \Sigma \cup \{\text{err}, \text{skip}\}$ can be done in Norm_O time, and thus in the same amount of space, by employing the usual normalization procedure followed by evaluating the relation. This also covers the case $p \in N$ as $A^{(0)} = \text{err}$ for all $A \in N$. The inductive case for the inner induction is very similar between the outer base case $i = 0$ and the outer inductive case moving from $i - 1$ to i , so we do not handle it separately.

We continue with the outer inductive case. For the base case of the inner induction, where $i \in \mathbb{N}$, we have $p \in N$ or $p \in \Sigma \cup \{\text{err}, \text{skip}\}$. The case $p \in \Sigma \cup \{\text{err}, \text{skip}\}$ is handled exactly as in the case we initially handled. Let $p = A \in N$. Then, $p^{(i)} = E(A)^{(i-1)}$. Because (N, E) is a linear grammar, $E(A)$ is a linear term. We call the induction hypothesis to examine $t \subseteq_O \text{nf}_O(E(A)^{(i-1)})$ and returns its result. This is possible in non-deterministic $\mathcal{O}(\log(i - 1)) + \max\{|E(A)|^2, |E|^2\} \cdot \text{Norm}_O^2$ space by the induction hypothesis. Note that we do not need to compute beyond calling the induction hypothesis, so we do not need additional space for it (i.e. we are a tail recursion).

For the inner inductive case, let p be a weak linear term. If we have $\text{inurg}(p) < h$, all non-terminals that appear in p must also have internal urgency of at most $h - 1$. These non-terminals can only refer to non-terminals of internal urgency at most $h - 1$. Then we can apply the normalization procedure, limited only to these non-terminals, to construct $p' = \text{nf}_O(p^{(i)})$. This takes $\mathcal{O}(\log(i)) + |p| \cdot \text{Norm}_O$ time,

i.e. the same amount of space. The former term stems from the counter used to stop at the i -th LFP iteration, and the latter term stems from the normalization.

Now let $p = q \Psi r$. We have $\text{nf}_O(p^{(i)}) = \text{nf}_O(q^{(i)}) \Psi \text{nf}_O(r^{(i)})$. Even if p does not contain non-terminals, we still handle $p^{(i)}$ instead of the equal term p to unify our notation. We write $q' = \text{nf}_O(q^{(i)})$, $r' = \text{nf}_O(r^{(i)})$, and also let $q' = \Psi Q$ and $r' = \Psi R$, where $Q, R \subseteq \text{WNF}_{h-1}$. We stress that the algorithm will not explicitly construct these terms. We have

$$\text{nf}_O(p^{(i)}) = (\Psi Q) \Psi (\Psi R) =_O \Psi(Q \cup R).$$

The final equality $=_O$, i.e. $(\subseteq_O \cap \supseteq_O)$, follows from the definition of \subseteq_O , which ignores bracketing among choices. By the definition of \subseteq_O , we observe that $t \subseteq_O p^{(i)}$ holds if and only if there is a $s \in Q$ or $s \in R$ with $t \subseteq_O s$. The conditions $s \in Q$ and $s \in R$ are respectively equivalent to $s \subseteq_O q$ and $s \subseteq_O r$. This means that a $s \in \{q, r\}$ exists with $t \subseteq_O s$ if and only if $t \subseteq_O p^{(i)}$ holds. To verify the existence of such terms, algorithm first guesses a term $s \in \{q, r\}$. Then, it removes the rest of the term from the memory and checks $t \subseteq_O s^{(i)}$. This can be done in non-deterministic $\mathcal{O}(\log(i)) + \max\{|q|^2, |r|^2, |E|^2\} \cdot \text{Norm}_O^2$ space by the induction hypothesis. This yields the desired complexity. If p does not contain non-terminals, we have $s^{(i)} = s$ for all $s \in \{q, r\}$ and the check is possible in $\max\{|q|^2, |r|^2\} \cdot \text{Norm}_O^2$ space.

For the remaining inductive case, we have $p = q.r$. We similarly write $q' = \text{nf}_O(q^{(i)}) = \Psi Q$ and $r' = \text{nf}_O(r^{(i)}) = \Psi R$ with $Q, R \subseteq \text{WNF}_{h-1}$. We get

$$\text{nf}_O(p^{(i)}) = \text{nfconc}_O((\Psi Q).(\Psi R)) = \Psi\{\text{nf}_O(q''.r'') \mid q'' \in Q, r'' \in R\}.$$

Note that this distribution is only possible because all terms $q'' \in Q$ and $r'' \in R$ have urgency $h-1$. Same as before, we observe $t \subseteq_O \text{nf}_O(p)$ if and only if there are $q'', r'' \in \text{WNF}_{h-1}$ with $t \subseteq_O \text{nf}_O(q''.r'')$, $q'' \subseteq_O \text{nf}_O(q^{(i)})$, and $r'' \subseteq_O \text{nf}_O(r^{(i)})$. If no such q'' and r'' can be found, this implies that $t \subseteq_O p^{(i)}$ does not hold.

Assume p contains no non-terminals. Then $q^{(i)} = q$ and $r^{(i)} = r$. The algorithm proceeds as follows. It first guesses $q'', r'' \in \text{WNF}_h$ and verifies $t \subseteq_O \text{nf}_O(q''.r'')$ by the usual normalization procedure, which takes at most $(|q''| + |r''|) \cdot \text{Norm}_O \leq \text{Norm}_O^2$ space. If this fails, the algorithm returns false. Then it removes all information except $q'', q, r'',$ and r . The space required for this bounded from above by $4 \cdot |p| \cdot \text{Norm}_O$. It proceeds by checking $q'' \subseteq_O \text{nf}_O(q)$ and $r'' \subseteq_O \text{nf}_O(r)$, which is possible in $\max\{|q|^2, |r|^2\} \cdot \text{Norm}_O^2$ by the induction hypothesis. In total, the algorithm requires $4 \cdot |p| \cdot \text{Norm}_O + \max\{|q|^2, |r|^2\} \cdot \text{Norm}_O^2$ non-deterministic space. We have

$$\begin{aligned} 4 \cdot |p| \cdot \text{Norm}_O + \max\{|q|^2, |r|^2\} \cdot \text{Norm}_O^2 &= 4 \cdot |p| \cdot \text{Norm}_O + 3 \cdot \max\{|q|^2, |r|^2\} \cdot \text{Norm}_O^2 \\ &\leq (4 \cdot |p| + 3 \cdot (|p| - 1)^2) \cdot \text{Norm}_O^2 \\ &= (4 \cdot |p| + 3 \cdot |p|^2 - 6 \cdot |p| + 3) \cdot \text{Norm}_O^2 \\ &= (3 \cdot |p|^2 - 2 \cdot |p| + 3) \cdot \text{Norm}_O^2 \\ &\leq 3(|p|^2) \cdot \text{Norm}_O^2 \end{aligned}$$

Now conversely assume that p contains non-terminals. Because p is linear, one of q or r does not contain non-terminals. W.l.o.g. let q be this term. The algorithm first guesses $q'', r'' \in WNF_{h-1}$ and verifies $t \subseteq_O \text{nf}_O(q''.r'')$ using the usual normalization process in Norm_O^2 space. If this fails, the algorithm returns false. Similarly to the case without non-terminals, it then removes all information except $q'', q, r'',$ and r . Following this, it verifies $q'' \subseteq_O q$. As before, if the check fails, it returns false. The term q does not contain non-terminals, so the check takes an additional space of only $|q|^2 \cdot \text{Norm}_O^2$ by the induction hypothesis. The same calculation as in the no non-terminal case yields an upper bound $|p|^2 \cdot \text{Norm}_O^2$ on total memory used to complete this step. Upon completing the step, the algorithm erases q'' and q from memory. Then it checks $r'' \subseteq_O r$ and returns the result of the check. This check is possible in $\mathcal{O}(\log(i)) + \max\{|r|^2, |E|^2\} \cdot \text{Norm}_O^2$ non-deterministic space by the induction hypothesis. Note that the algorithm does not store any additional information besides the information required to perform this check. This yields the desired space complexity. \square

By slightly modifying the approach in Lemma 15, we can also get a complexity bound wrt. to $|O|$. Instead of letting Eve resolve the transitions in the DFA for $|O|$, we let Adam resolve them. We abuse the notation and refer to this translation with the symbol $\text{tr}(\cdot)$ for the rest of this section.

$$\begin{aligned} \text{tr}(a) &= \bigwedge \{(p, q) \mid \delta(p, a) = q\} & \text{tr}(\text{err}) &= \text{err} \\ \text{tr}(\bigoplus P) &= \bigoplus \{\text{tr}(p) \mid p \in P\} & \text{tr}(\text{skip}) &= \text{skip} \\ \text{tr}(p.q) &= \text{tr}(p).\text{tr}(q) & \text{tr}(A) &= A, \end{aligned}$$

The translation for the non-terminals also remains the same $E_{\text{tr}}(A) = \text{tr}(E(A))$. The transitions of the translated objective $\text{tr}(O)$ remain the same, however it accepts \perp as well as the original final states. This is because the task of choosing the correct transition is now assigned to Adam, and a failure leads to an Eve win. The translation only introduces \bigwedge at the lowest urgency level and does not modify the urgencies. Thus, a weak linear term p and a weak linear grammar (N, E) , the translated $\text{tr}(p)$ and $(N, \text{tr}(E))$ both remain weak linear. The Lemma Lemma 15 holds for this type of translation as well.

Lemma 34. $p \Downarrow O$ if and only if $\text{tr}(p) \Downarrow \text{tr}(O)$.

Proof Sketch. Letting Adam simulate the translations of the automaton does not change the winner. Determinicity still only allows for one choice of transition per $\text{tr}(a)$ that does not lead to his loss. The proof follows the same steps as those of Lemma Lemma 15. \square

The complexity wrt. $|O|$ follows.

Lemma 35. Let (N, E) be a weak linear grammar with maximal urgency h , p a weak linear term, and $O \subseteq \Sigma^*$ a regular objective. Then $p \Downarrow O$ can be decided in $(\max\{|E|, |p|\} \cdot |E| \cdot |N| \cdot \exp_{h-2}(\mathcal{O}(|O|^2)))^4$ space.

P Lower Bound

We prove the lower bound given in Theorem 4 with a reduction from context-bounded multi-pushdown games, a concurrent programming model the complexity of which is well-understood [25,26]. The proof of the lower bound given in Theorem 3 can be found in Appendix J.

P.1 Multi-Pushdown Games

We introduce multi-pushdown games trimmed to our needs. A *b-context-bounded 2-stack pushdown game* (*b-2PDG*) is a tuple $(Q, E, p_0, F, \Gamma, \delta)$ consisting of a finite set of states Q , a set of states $E \subseteq Q$ owned by Eve, an initial state p_0 , a set of goal states $F \subseteq Q$, a stack alphabet Γ , and a set of transitions $\delta \subseteq Q \times Op \times Q$. Transitions are annotated by a stack operation from $Op = \Gamma \times \Gamma^{\leq 2} \cup \{\text{nx}\}$. With (a, w) , we pop a from and push w to the active stack. With nx , we change the active stack, called a context switch. We assume there is a bottom of stack symbol $\$ \in \Gamma$ that is never removed.

The semantics of a *b-2PDG* is a game arena $(CF, \rightarrow, \text{own})$ with a reachability objective CF_F for Eve. The positions are configuration from $CF = Q \times [0, b] \times \Gamma^* \times \Gamma^*$. A configuration (p, k, s_1, s_0) stores the current state p , the number of context switches k that have occurred so far, and the contents of the two stacks. Stack s_0 is active after an even number of context switches, stack s_1 is active when k is odd. The owner and moves are defined as expected, there are no context switches beyond b , and we assume there are no deadlocks. The objective is $CF_F = F \times \{\mathbf{b}\} \times \Gamma^* \times \Gamma^*$, meaning we reach a goal state and have exhausted the context switches. Plays, strategies, and winning are defined like for urgency programs. The task is to decide whether Eve has a strategy to win from $(p_0, 0, \$, \$)$.

Theorem 12. [26] *b-2PDG are $(b - 2) - \text{EXPTIME}$ -complete.*

P.2 Reduction

The reduction is in two steps, we first reduce 2PDG to the problem of making an observation:

Proposition 9. *Repetition of Proposition 6 Given a $(2h+1)$ -2PDG PD, we can compute in poly time p over Σ and (N, E) of maximal urgency h and an objective DFA O so that Eve wins PD if and only if $p \Downarrow O$.*

We now reduce the problem of making an observation to the specialized contextual equivalence. Indeed, $p \Downarrow O$ is the same as to check $\chi_O(\bullet) \preceq_O p$, where $\chi_O(\bullet)$ is the characteristic term of the empty context formed for objective O . The problem is that the characteristic term may be exponential. We utilize the trick from Appendix K.1.

Lemma 36. *Given p over Σ and (N, E) , and objective O , we can compute in poly time $\chi_{\text{tr}(O)}(\bullet)$, $\text{tr}(p)$, and $\text{tr}(O)$ so that $p \Downarrow O$ if and only if $\chi(\bullet) \preceq_{\text{tr}(O)} \text{tr}(p)$.*

We sketch the proof of Proposition 6. We encode positions $(p, k, x \dots \$, y \dots \$)$ of the 2PDG as urgency terms

$$H^w . \underbrace{g.x^u \dots \$^u . @}_{s_1} . \underbrace{h.y^v \dots \$^v . @}_{s_0} .$$

Stack symbols $x \in \Gamma$ are represented by terms x^u of urgency u . Terminal $@$ marks the end of a stack content encoding. The terms g and h represent the history of the play. Finally, H^w implements context switches. The construction controls w , u , and v so that the top of the active stack is leading.

The top of the active stack allows the game to proceed as

$$\dots g . \underbrace{\left(\bigvee_{p \in Q} \bigcirc_{t \in \delta_{p,x}} \langle t \rangle^u \right)}_{\text{Rewritten from } x^u} \dots \$^u . @ \dots \rightarrow \dots g . \langle t \rangle_x^u \dots \$^u . @ \dots$$

Eve selects the current state $p \in Q$. Then the player owning this state selects the next transition. We use $\delta_{p,x} \subseteq \delta$ to denote the set of transitions from state p with top of stack symbol x . The set is non-empty because the 2PDG does not deadlock. The term $\langle t \rangle^u$ of the chosen transition contains terminals which join history g to record the state change and the urgency u . The objective O is a product DFA that reads the terminals for each urgency separately and enforces consistency with the 2PDA transitions.

Push/pop operations modify the active stack encoding in the expected way. For context switches, the leading term must swap the stack. To implement this, we use a decrement process on the now no longer active stack. We define stack symbols x^u as $\langle \rightarrow st \rangle^u . r \Psi \langle \rightarrow nx \rangle^u . x^{u-1}$, where r is the choice of the next transition explained above. The decrement process relies on the alternative $\langle \rightarrow nx \rangle^u . x^{u-1}$, which replaces x^u by x^{u-1} . A snapshot of the decrement process is

$$\dots g' . \langle nx \rangle \dots \overbrace{\langle \rightarrow nx \rangle . x_{i-1}^{u-1} . x_i^u}^{\text{Urgency } u-1} \dots \$^u @ \dots$$

Progression of the leading term \rightarrow

The terminals $\langle \rightarrow st \rangle$, $\langle \rightarrow nx \rangle$, $\langle nx \rangle$, and $@$ allow the objective to check the decrement process for correctness.

Each urgency simulates two contexts. Since we do not need to access the odd stack before the first context switch, we only generate this stack when it is first accessed. This allows us to simulate three contexts with the maximal urgency. In total, the construction simulates $2h + 1$ contexts with urgency h .

Q Lower Bound Details

Construction, Objective: The set of terminal symbols Σ consists of assignments $x := val$ and assertions $x =, val$ of the variables $f_u \in QU\{-\}$, $s_u \in QU\{-\}$,

and $c_u \in \{\text{nx}, \text{sty}\}$ for each $0 < u \leq h$, along with a variable $gn \in \{+, -\}$. In the parts of the play, where the urgency of the term is $0 < u \leq h$, variable f_u will keep track of the first MPDG state, the variable s_u will keep track of the latest MPDG state, and $c_u \in \{\text{nx}, \text{sty}\}$ will be used to enforce the correctness of context switches. The variable $gn \in \{+, -\}$ keeps track of whether the game has generated the second stack by making the first context switch. The objective DFA \mathcal{A} processes the updates and assertions on the values of these variables. For each $0 < u \leq h$, the DFA also keeps an assertion failure flag $err_u \in \{\perp, \top\}$, that records whether there has been assertion failure for f_u , s_u , c_u , and gn . If an assertion failure happens for one of these variables, then err_u is irrevocably set to \perp . In the initial state i , we have $gn = f_u = s_u = -$, $c_u = \text{st}$, and $err_u = \top$ for all $0 < u \leq h$. The DFA \mathcal{A} accepts if and only if $s_1 \in F$, there are no assertion errors ($err_u = \top$ for all $0 < u \leq h$), and the latest states are consistent with the first states ($s_{u+1} = f_u$ for all $0 < u < h$).

Construction, Assignments: We now move on to the construction of the defining assignments. Each stack symbol is represented by a different term for each urgency. The set of non-terminals is $N = \{x_{NT}^u \mid x \in \Gamma, 0 < u \leq h\}$. The representation of an individual stack symbol for urgency u , wraps the corresponding non-terminal in a unary choice with urgency u . Formally the representing term is the singleton choice $x^u = \forall x_{NT}^u$ for all $0 < u \leq h$. This ensures that the u -representation of a stack symbol has urgency u (remember that non-terminals have highest urgency). Furthermore, the term that represent the stack symbol must be the leftmost action. This allows a concatenation of terms that represent stack symbols to act like one stack in the MPDG. The defining assignments $E : N \rightarrow \mathbf{T}$ are laid out below for all $x_{NT}^u \in N$. We use helper terms to simplify the representation. For all $u \leq h$, $v < h$, $p, q \in Q$, $x \in \Gamma$, and $w \in x^{\leq 2}$ we have:

$$\begin{aligned} E(x_{NT}^u) &= (c_u =! \text{nx}.x^{u-1}) \Psi (c_u =! \text{st}.Pop_x^u) \\ Pop_x^u &= \forall_{p \in Q} s_u =! p. \langle \delta_{p,x} \rangle^u \quad \langle \delta_{p,x} \rangle^u = \bigoplus_{t \in \delta_{p,x}} \langle t \rangle_x^u \\ \langle (p, x, w, q) \rangle_x^u &= s_u := q.w_0^u \dots w_n^u \\ \langle (p, \text{nx}, q) \rangle_x^h &= (gn =! -.gn := +.\$^h. @.x^h) \Psi \\ &\quad (gn =! +.s_h := q.c_h := \text{nx}.x^h) \\ \langle (p, \text{nx}, q) \rangle_x^v &= s_v := q.c_v := \text{nx}.x^v \end{aligned}$$

$$@ = c_1 := \text{st} \dots c_h := \text{st}$$

$$H^u = H^{u-1}. \forall_{p \in Q} f_u := p.s_u := p$$

The initial term for the game is simply $H^{h-1}.\underline{\$^h}.@$. The terminals $\langle \rightarrow \text{st} \rangle^u$ and $\langle \rightarrow \text{nx} \rangle^u$ used in the main paper refer to the assertions $c_u =! \text{st}$ and $c_u =! \text{nx}$. At context switches in urgency h , Eve also needs to “guess” whether the second

stack has been generated. In the case where it has not yet been generated, the correct choice generates it. In the case where it has already been generated, the correct choice triggers a context switch in the usual way.

Q.1 Denotational Semantics

We show how to define a denotational semantics based on our axiomatization. What we find interesting is that, with the axiomatization at hand, the denotational semantics is a derived construct: the semantic domain and the interpretation of function symbols are induced by the axiomatization, yet the semantics is guaranteed to be fully abstract wrt. contextual equivalence resp. its specialized variant. The creativity that is saved in the definition of the semantics of course had to be invested up front when coming up with the axiomatization. We found it easier to study an axiomatization than a denotational semantics, because the problem is narrowed down to understanding the interplay between operators as opposed to coming up with a freely chosen semantic domain. We recall the basics of denotational semantics before turning to the details.

A *denotational semantics* for our programming language is a pair $((\mathcal{D}, \subseteq), \mathcal{I})$ consisting of a complete partial order (\mathcal{D}, \subseteq) of semantic elements and an interpretation $\mathcal{I} : F \rightarrow \mathcal{D}^\omega \rightarrow \mathcal{D}$ that assigns to each function symbol $f \in F$ in our language a monotonic function $f^\mathcal{I} : \mathcal{D}^{ar(f)} \rightarrow \mathcal{D}$ of the expected arity. The function symbols F are Σ , $\{\text{skip}, \text{err}, \cdot\}$, and choices of arbitrary arity with urgency 1 to h . We lift the interpretation to all terms p and assign them an element $\mathcal{D}(p) \in \mathcal{D}$, called the denotational semantics of the term. For recursion-free terms, the lifting is purely compositional:

$$\mathcal{D}(a) = a^\mathcal{I} \quad \mathcal{D}(p.q) = \mathcal{D}(p) \cdot^\mathcal{I} \mathcal{D}(q),$$

and similar for the other function symbols. For the non-terminals (N, E) , this allows us to understand the defining equations as a monotonic function

$$E^\mathcal{I} : (N \rightarrow \mathcal{D}) \rightarrow N \rightarrow \mathcal{D}.$$

The least solution of this function is the denotational semantics of the non-terminals: $\mathcal{D}(A) = [lfp.E^\mathcal{I}](A)$ for every $A \in N$. This is the missing case to define the semantics of arbitrary program terms again in a compositional way.

We focus on the denotational semantics induced by the axiomatic congruence. The development for the O -specialized axiomatic congruence with O right-separating is the same. If O is not right-separating, we cannot give a guarantee that the resulting semantics will be fully abstract. The *denotational semantics induced by \equiv* is $((\mathcal{D}_\equiv, \subseteq), \mathcal{I}_\equiv)$. The set of semantic elements is $\mathcal{D}_\equiv = \mathbf{T}/\equiv$, we factorize the set of terms along the axiomatic congruence. The complete partial order on these congruence classes is the one given by the axiomatic precongruence. It is guaranteed to be well-defined due to the precongruence. It is guaranteed to stabilize in an ordinal by the fact that chains are well-ordered sets. The interpretation of the function symbols is as expected:

$$a^{\mathcal{I}_\equiv} = [a]_\equiv \quad [p]_\equiv \cdot^{\mathcal{I}_\equiv} [q]_\equiv = [p.q]_\equiv.$$

Well-definedness holds because \equiv is a congruence, monotonicity holds because \sqsubseteq is a precongruence. The semantics is fully abstract wrt. contextual equivalence, $\mathcal{D}(p) = \mathcal{D}(q)$ iff $p \preceq q$, which is merely a reformulation of Theorem 1. We can define other fully abstract semantics by introducing representative systems on the congruence classes, for example based on normal forms.