



Expressive Logical Combinators for Free

Pierre Genevès, Nabil Layaïda, Alan Schmitt

► To cite this version:

Pierre Genevès, Nabil Layaïda, Alan Schmitt. Expressive Logical Combinators for Free. 2013. hal-00868724v1

HAL Id: hal-00868724

<https://inria.hal.science/hal-00868724v1>

Preprint submitted on 1 Oct 2013 (v1), last revised 6 May 2015 (v4)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Expressive Logical Combinators for Free

Pierre Genevès¹, Nabil Layaïda², and Alan Schmitt²

¹ CNRS

² Inria

Abstract. A popular technique in the static analysis for query languages relies on the construction of compilers that effectively translate queries into logical formulas. These formulas are then solved for satisfiability using an off-the-shelf satisfiability solver. A critical aspect in this approach is the size of the obtained logical formula, since it constitutes a factor that affects the combined complexity of the global approach.

We show that a whole class of logical combinators (or “macros”) can be used as an intermediate language between the query language and the logical language. Those logical combinators provide an exponential gain in succinctness over the corresponding explicit logical representation, yet preserving the typical exponential time complexity of the subsequent logical decision procedure.

This opens the way for solving a wide range of problems such as satisfiability and containment for expressive query languages in exponential-time even though their direct formulation into the underlying logic results in an exponential blowup of the formula size, yielding an incorrectly presumed two-exponential time complexity.

We illustrate this from a very practical point of view on a few examples such as numerical occurrence constraints and tree frontier properties which are concrete problems found in the XML world.

1 Introduction

Modal logics have recently been increasingly used as target languages for compiling high level query languages such as XPath [5,9], SPARQL [2], and even CSS selectors [6], as well as policy languages such as XACML [8]. The common goal is to characterize these languages in terms of expressivity and complexity, and in particular to build adapted and effective static analyzers.

Such an approach requires the construction of efficient compilers that translate queries into logical formulas. Those formulas are then solved for satisfiability using an off-the-shelf satisfiability solver such as the ones found in [14,11,5,13,10,4]. A critical aspect is then the size of the obtained logical formula, since it is a factor that affects the combined complexity³ of the global approach.

³ In the context of problem-solving by reduction to logical satisfiability, combined complexity considers both the complexity of the translation of the problem into logic (taking into account any potential blow-up in size induced by the change in representation), and the complexity of testing satisfiability of the logical formulation.

We show that a whole class of logical combinators (or “macros”) can be used as an intermediate language between the query language and the logical language. Those logical combinators provide an exponential gain in succinctness over the corresponding explicit logical representation, yet preserving the typical exponential time complexity [14,11,5,13] of the subsequent logical decision procedure.

This opens the way for solving a wide range of problems such as query satisfiability and query containment in exponential-time even though their direct formulation into the underlying logic results in an exponential blowup of the formula size, yielding an incorrectly presumed two-exponential time complexity.

Specifically, two essential steps are involved in the reduction of a problem to logical satisfiability: (1) the translation of the initial problem into a logical formula, and (2) the actual satisfiability check of the formula. Traditionally, the complexity of the satisfiability test is stated in terms of the size of the formula, thus every duplication of sub-formulas during the first step may affect the combined complexity and severely impact the practical applicability of the entire approach. Interestingly, we observe that a common form of μ -calculus sub-formula duplication has a very limited impact on combined complexity in existing implementations, such as [14,11,5,13]. The reason lies in the fact that satisfiability-testing algorithms can operate directly on a Hintikka-set-like representation of formulas composed of atomic propositions and modal sub-formulas. In this setting, we prove that the time complexity actually depends on the number of *distinct* atomic propositions and modal sub-formulas. This makes explicit a notion of truth-status sharing for identical sub-formulas not exhibited in the analysis of the time complexity of such algorithms.

We develop this idea in the context of the μ -calculus, whose expressive power subsumes the ones of many modal logics. More specifically we develop this idea using the alternation-free μ -calculus with converse modalities whose models are finite trees, following [5]. Trees are encoded in binary, without loss of generality [3], through the “first-child” and “next-sibling” modalities, respectively noted $\langle 1 \rangle$ and $\langle 2 \rangle$ (see Figure 1).




Fig. 1. N-ary to Binary Tree Encoding

In this setting, an elegant way of building a μ -calculus formula is to apply a combinator to another formula. For instance, $\mathbf{split}(X) = \langle 1 \rangle X \wedge \langle 2 \rangle X$ is a combinator that generates a formula such that the input formula must hold in

both successors of the current node. Although X is duplicated, the increase of the size of the lean generated from $\text{split}(\varphi)$ when compared to the one generated from φ is only a small constant, independent of φ .

The paper is organized as follows. We introduce the logical formulas and combinators in §2, state our main result in §3, and give several examples in §4.

2 Basic Logical Formulas and Combinators

We recall the syntax of the logic of [5], used to prove properties about finite binary trees. We consider a set AP of atomic propositions, representing the tree node names, which includes a special reserved name ϵ ; a set Var of variables, used in fixpoints; and a set Prog = $\{1, 2, \bar{1}, \bar{2}\}$ of programs, to describe navigation in a tree. Program 1 navigates to the first child (left successor in Figure 1), program 2 navigates to the next sibling (right successor), program $\bar{1}$ to the parent (predecessor to the right, if it exists), and program $\bar{2}$ to the previous sibling (predecessor to the left, if it exists). We let $\bar{\bar{a}} = a$ for any $a \in \text{Prog}$. A logical formula is defined using the following syntax.

- \top , \perp , σ or $\neg\sigma$ for all $\sigma \in \text{AP} \setminus \{\epsilon\}$;
- x for all $x \in \text{Var}$;
- $\varphi_1 \vee \varphi_2$ or $\varphi_1 \wedge \varphi_2$ where φ_1 and φ_2 are logical formulas;
- $\langle a \rangle \varphi$ or $\neg \langle a \rangle \top$ where $a \in \text{Prog}$ and φ is a logical formula;
- $\mu x. \varphi$ where $x \in \text{Var}$ and φ is a logical formula.

We now give an intuition of the interpretation of formulas in the setting of finite trees. A formula is *satisfiable* if there exists a tree such that a node of this tree is *selected* by the formula. The truth formula \top selects every node whereas \perp selects none. The σ formula selects every node whose name is σ whereas $\neg\sigma$ selects the nodes with other names (the node name ϵ is used to represent names of nodes not occurring in the formula). Formula conjunction and disjunction correspond to set intersection and union, respectively. A formula $\langle a \rangle \varphi$ selects a node if the node reached following a is selected by formula φ . Formula $\neg \langle a \rangle \top$ selects a node if there is no node reachable through a . Finally, a fixpoint $\mu x. \varphi$ is interpreted as the smallest fixpoint (the intersection of every pre-fixpoint).

A formula is *closed* if every occurrence of a variable x is bound by an enclosing μx . We write $\varphi \prec \psi$ if φ is a sub-formula of ψ , and $\varphi \not\prec \psi$ if it is not.

A *combinator* F is a formula with zero or more occurrences of a placeholder, written X , possibly negated ($\neg X$). We write $F\{\varphi/X\}$ for the combinator F where every instance of X has been replaced by the closed formula φ .⁴ We often write $F(X)$ to make clear the name of the placeholder, and $F(\varphi)$ for $F\{\varphi/X\}$.

We consider formulas in negation normal form. The negation of a formula or combinator, written \bar{F} , is defined in Figure 2. The negation of a modality is a disjunction: the modality is false either because there is no node in that direction,

⁴ In case of a negated placeholder, we replace $\neg X$ with $\bar{\varphi}$, the negation normal form of φ (see Figure 2).

or because the node in that direction does not satisfy the sub-formula. Following [5], the greatest and smallest fixpoint coincide (provided a simple restriction on formulas, namely for cycle-free formulas using sets of finite trees as models: see [5] for details). Every combinator presented here fulfill this restriction. However, but this work could also be done in a setting where the smallest and greatest fixpoints differ, and in this case one defines $\overline{\mu x.F}$ as $\nu x.\overline{F}$.

$$\begin{array}{ll}
\overline{\top} = \perp & \overline{\sigma} = \neg\sigma \\
\overline{F \wedge G} = \overline{F} \vee \overline{G} & \overline{\perp} = \top \\
\overline{\neg\sigma} = \sigma & \overline{F \vee G} = \overline{F} \wedge \overline{G} \\
\overline{X} = \neg X & \overline{x} = x \\
\overline{\langle a \rangle \top} = \neg \langle a \rangle \top & \overline{\neg X} = X \\
\overline{\mu x.F} = \mu x.\overline{F} & \overline{\neg \langle a \rangle \top} = \langle a \rangle \top \\
\overline{\langle a \rangle F \vee \neg \langle a \rangle \top} = \langle a \rangle \overline{F} & \overline{\langle a \rangle F} = \langle a \rangle \overline{F} \vee \neg \langle a \rangle \top \quad F \neq \top
\end{array}$$

Fig. 2. Negation Normal Form

3 Deciding Combined Formulas

3.1 The Lean

Following [12,5,13], we define the lean of F as follows, with $\overline{\mathcal{L}}_F(F)$ defined in Figure 3 (the environment F is the set of already unfolded fixpoints). The main difference with the usual approaches is that we close the lean under negation.

$$\text{Lean}(F) = \{ \langle a \rangle \top \mid a \in \{1, 2, \overline{1}, \overline{2}\} \} \cup \{ \emptyset \} \cup \overline{\mathcal{L}}_\emptyset(F)$$

3.2 The Factorization Power of the Lean

We can now state the main theorem of this paper: the lean size is not impacted by the duplication of sub-formulas. We write $|S|$ for the size of a set S .

Theorem 31 *Let F be a combinator and φ a closed formula. We have $|\text{Lean}(F\{\varphi/x\})| \leq |\text{Lean}(F)| + |\text{Lean}(\varphi)|$.*

The theorem is a direct consequence of Lemma A8 which is proved in the appendix.

We now give some intuition about this result, through a simple example. Recall the `split`(X) combinator defined as $\langle 1 \rangle X \wedge \langle 2 \rangle X$. Since the elements of

$$\begin{aligned}
\bar{\mathcal{L}}_I(\top) &= \bar{\mathcal{L}}_I(\perp) = \bar{\mathcal{L}}_I(x) = \bar{\mathcal{L}}_I(X) = \bar{\mathcal{L}}_I(\neg X) = \emptyset \\
\bar{\mathcal{L}}_I(\sigma) &= \bar{\mathcal{L}}_I(\neg\sigma) = \{\sigma\} \\
\bar{\mathcal{L}}_I(F \vee G) &= \bar{\mathcal{L}}_I(F \wedge G) = \bar{\mathcal{L}}_I(F) \cup \bar{\mathcal{L}}_I(G) \\
\bar{\mathcal{L}}_I(\langle a \rangle F) &= \{\langle a \rangle F, \langle a \rangle \bar{F}, \langle a \rangle \top\} \cup \bar{\mathcal{L}}_I(F) \\
\bar{\mathcal{L}}_I(\neg \langle a \rangle \top) &= \{\langle a \rangle \top\} \\
\bar{\mathcal{L}}_I(\mu x.F) &= \bar{\mathcal{L}}_I(F) \quad \text{if } x \notin F \\
\bar{\mathcal{L}}_I(\mu x.F) &= \emptyset \quad \text{if } \mu x.F \in \Gamma \text{ or } \mu x.\bar{F} \in \Gamma \\
\bar{\mathcal{L}}_I(\mu x.F) &= \bar{\mathcal{L}}_{\Gamma \cup \{\mu x.F\}}(F\{\mu x.F/x\}) \quad \text{otherwise}
\end{aligned}$$

Fig. 3. Negation Closed Lean

the lean are either atomic propositions (node names) and modalities, the lean of $\text{split}(\varphi)$ includes the lean of φ and four new elements: $\langle 1 \rangle \varphi$, $\langle 1 \rangle \bar{\varphi}$, $\langle 2 \rangle \varphi$, and $\langle 2 \rangle \bar{\varphi}$. If we now consider $\text{split}(\text{split}(\varphi))$, we once again add only four formulas to the lean: $\langle 1 \rangle \text{split}(\varphi)$, $\langle 1 \rangle \bar{\text{split}(\varphi)}$, $\langle 2 \rangle \text{split}(\varphi)$, and $\langle 2 \rangle \bar{\text{split}(\varphi)}$. This linear growth, even though the formula's size increases exponentially, is due to the fact that modalities are considered atomically and are not split up in their components (e.g., $\langle 1 \rangle (\varphi \wedge \psi)$ is not split up into $\langle 1 \rangle \varphi$ and $\langle 1 \rangle \psi$).

3.3 Satisfiability-Testing Algorithms based on the Lean

A typical approach to decide the satisfiability of a formula is to first build the lean, as described above, then to use a tableau-based algorithm implemented with BDDs [14,12,5,13]. The time complexity of this approach is shown to be exponential in the size of the formula. More precisely, it is exponential in the size of the lean, which is in turn linear in the size of the formula.

The essence of this paper is to realize that the lean may grow much more slowly than the formula when sub-formulas are duplicated. This opens the way for solving a wide range of problems in exponential-time even though their direct translation into the modal logic is exponential, as illustrated previously and in the following section.

4 Applications

In this section, we present several instances where advanced properties on the underlying data structure (here a tree) can be formulated using combinators, and for which our result applies. In particular, after an introductory example, we consider numerical constraints on the global number of occurrences, and properties on the sequence of leaves in a tree. This means, for instance, that if one extends a query language (such as XPath) with such kinds of features, then our result applies: problems such as query satisfiability and query containment

would not be harder to solve in terms of computational complexity than they already are for the language without the extensions.

An online implementation of a μ -calculus satisfiability solver is available since the work of [5]. Each example given in this section is provided with a boxed version which can directly be copied and pasted in order to be tested with this online implementation.

4.1 A Very Simple Example: Split

The combinator `split(X)` introduced in Section 3.2 may generate arbitrary large formulas: for instance, let $\varphi = a \wedge \langle 1 \rangle b \wedge \langle 2 \rangle \mu y. c \vee \langle 2 \rangle y$, the expanded formula $\psi = \text{split}(\text{split}(\text{split}(\varphi)))$ uses 8 occurrences of φ . To give this formula to the implementation of [7], we write it as follows.

```
phi() = a & <1>b & <2>let $y = c | <2>$y in $y;
split($x) = <1>$x & <2>$x;
split(split(split(phi())))
```

We then observe that ψ contains 24 atomic propositions, 38 modalities, 23 conjunctions, and 8 disjunctions (including duplicates). The size of the lean is only 18 (14 modalities and 4 atomic propositions.)⁵ Each new `split(-)` around the formula then only adds two elements to the lean.

4.2 Document-Order Relation and Global Counting




Fig. 4. Tree Navigation

We illustrate tree navigation in Figure 4. A very simple example of a combinator is the descendant relation that checks that a node satisfying some formula

⁵ Notice that in the implementation of [5], the lean is not closed under negation. Closing the lean under negation adds a modal formula $\langle a \rangle \bar{\varphi}$ for every modal formula $\langle a \rangle \varphi$ where φ is not \top . In this particular case, the lean would contain 10 other formulas. Each new `split(-)` would then add 4 formulas to the lean.

X is accessible in the subtree by any sequence of forward modalities. It is encoded as follows:

$$\text{descendant}(X) = \langle 1 \rangle \mu z. X \vee \langle 1 \rangle z \vee \langle 2 \rangle z$$

A whole range of combinators to navigate in a tree can be defined in a similar manner. In particular we can encode:

$$\text{following}(X) = \text{ancestor_or_self}(\psi)$$

where

$$\psi = \text{following_sibling}(\text{descendant_or_self}(X))$$

These combinators, whose intuition is illustrated in Figure 4, are predefined in [7]. They can be used as such to encode the so-called *document-order* relation \ll . This relation corresponds to the ordering of nodes given by a depth-first tree traversal: $x \ll y$ iff node y is visited after node x in a depth-first tree traversal. We define the combinator $\text{next}(X) = \text{descendant}(X) \vee \text{following}(X)$ with which we can mimic the document-order relation (we write $X \wedge \text{next}(Y)$ for $x \ll y$). Notice that this combinator duplicates formulas, since the placeholder X appears twice in its definition.

The document-order relation can be used to express global counting properties in trees. For instance, if we want to encode the so-called concept of a *nominal* – or more generally the fact that some formula ψ is satisfied by one and only one node in the tree – we can write:

```
psi() = a & <1>b & <2>let $y = c | <2>$y in $y;
next($x) = descendant($x) | following($x);
root() = ~<-1>T & ~<-2>T;
nominal($x) = ancestor_or_self(root()
    & next($x & ~next($x)));
nominal(psi())
```

If we now want to force the existence of at least 4 different tree nodes that satisfy ψ , we can write:

```
psi() & next(psi() & next(psi() & next(psi())))
```

The full expansion of the above formula is notably large (if we count duplicates, the formula contains 468 atomic propositions, 933 modalities, 343 conjunctions, and 590 disjunctions). However, the size of its lean is 48 (44 modalities and 4 atomic propositions).

4.3 The Tree Frontier

We borrow from [1] another advanced example: the description of properties on a tree frontier. A tree frontier is the set of leaves (nodes without an outgoing “1” edge) ordered from left to right. A frontier node y is the successor of a frontier node x iff $x \ll y$ and there is no leaf node in between x and y in the




Fig. 5. Tree Frontier Case Analysis

```

leaf() = ~<1>T;
down_to_first_leaf($z) = let $x = (leaf() & <0>$z) | <1>$x in $x;
up_until_rsibl($x) = (<2>T & $x)
    | let $w = <-1>((<2>T & $x) | $w) | <-2>$w in $w;
next_frontier_node($y) = (leaf() & <2>($y & leaf()))
    | (leaf() & up_until_rsibl(<2>down_to_first_leaf($y)));
down_to_first_leaf(a & next_frontier_node(a & next_frontier_node(a)))

```

Fig. 6. Tree Frontier Example (including all the necessary definitions for use with [7]).

document order. A simple case analysis shows that node y is the successor of a frontier node x in one of three cases. These are depicted in Figure 5, where the current leaf is black, the next leaf is white, and grayed nodes are not selected. Dotted arrows correspond to sequences of navigation.

1. Either x is a leaf with an immediate next sibling which is also a leaf (y);
2. or x is a leaf with an immediate next sibling which is not a leaf, in which case, by navigating downward in its subtree we reach the leftmost leaf (y);
3. or x is a leaf with no next sibling, in which case, by going up to the parent node recursively until we reach a parent node which has a next sibling, then going to this next sibling, and then, from this node, navigating downward in its subtree we reach the leftmost leaf (y).

This yields the following definition of a combinator in which the first member of the disjunction captures the first case and the second member of the disjunction captures the two remaining cases:

$$\text{next_frontier_node}(Y) = (\text{leaf} \wedge \langle 2 \rangle (Y \wedge \text{leaf})) \vee (\text{leaf} \wedge \text{up_until_rsibl}(\psi))$$

In this definition, the placeholder Y is to be replaced by a formula that holds at the successor node, and:

$$\begin{aligned}
\text{leaf} &= \neg \langle 1 \rangle \top \\
\psi &= \langle 2 \rangle \text{down_to_first_leaf}(Y) \\
\text{up_until_rsibl}(X) &= (\langle 2 \rangle \top \wedge X) \\
&\quad \vee \mu x. \langle \bar{1} \rangle ((\langle 2 \rangle \top \wedge X) \vee x) \vee \langle \bar{2} \rangle x \\
\text{down_to_first_leaf}(Z) &= \mu x. (\text{leaf} \wedge Z) \vee \langle 1 \rangle x
\end{aligned}$$

Using these combinators, we can now express properties on the tree frontier. For instance, the formula shown on Figure 6 states that the leftmost leaf is labeled “a”, and, by further navigation on the tree frontier, we encounter two other leaves labeled “a”. If we count duplicates, the corresponding formula contains 13 atomic propositions, 58 modalities, 30 variables, 13 fixpoint binders, 21 negations, 33 conjunctions, and 25 disjunctions. However, the size of the corresponding lean is 23. The lean is only composed of 21 distinct modalities and 2 distinct atomic propositions. Each additional nested call to the combinator $_ \wedge \text{next_frontier_node}(Y)$ extends the lean by 7 modalities. However, the corresponding global formula goes from 58 modalities to 184 for the first addition, then it goes to 562 for the second addition.

5 Conclusion

We have presented the concept of logical combinators that avoid exponential increases in combined complexity due to sub-formula duplication. Our main result, of theoretical nature, has very practical consequences and applies for a large class of logical solvers such as the ones found in [14,11,5,13].

We have further illustrated this result in the context of one of these satisfiability solvers [5], for which we have presented an in-depth analysis. This analysis focuses on the time complexity of lean-based algorithms to decide the satisfiability of a tree logic equipped with inverse programs, nominals, and counting introduced via combinators. The analysis highlights our result by showing that the lean automatically factorizes duplicated sub-formulas even for such advanced features, thus the complexity of the algorithm should not be stated in terms of the size of the initial formula but in terms of the size of the lean. A direct consequence of this observation is that the addition of nominals and a more general form of counting to the initial tree logic has no impact on decidability nor on its precise complexity bound. We have also reported on the experimental tests of this analysis using an implementation available online.

As a direction for future work, it would be interesting to use this approach to investigate or revisit problems which have been avoided as they were leading to formula duplication.

References

1. Afanasiev, L., Blackburn, P., Dimitriou, I., Gaiffe, B., Goris, E., Marx, M., de Rijke, M.: PDL for ordered trees. *Journal of Applied Non-Classical Logics* 15(2), 115–135 (2005)
2. Chekol, M.W., Euzenat, J., Genevès, P., Layaïda, N.: SPARQL query containment under RDFS entailment regime. In: *IJCAR: Proceedings of the 6th International Joint Conference on Automated Reasoning*. pp. 134–148 (2012)
3. Comon, H., Dauchet, M., Gilleron, R., Löding, C., Jacquemard, F., Lugiez, D., Tison, S., Tommasi, M.: Tree automata techniques and applications. Available on: <http://www.grappa.univ-lille3.fr/tata> (2007), release October, 12th 2007
4. Dutertre, B., de Moura, L.: The Yices SMT solver. Tool paper at <http://yices.csl.sri.com/tool-paper.pdf> (Aug 2006)
5. Genevès, P., Layaïda, N., Schmitt, A.: Efficient static analysis of XML paths and types. In: *PLDI* (2007)
6. Genevès, P., Layaïda, N., Quint, V.: On the analysis of cascading style sheets. In: *WWW’12: Proceedings of the 21st World Wide Web Conference*. pp. 809–818 (April 2012)
7. Genevès, P., Layaïda, N., Schmitt, A.: XML reasoning solver project, <http://wam.inrialpes.fr/websolver/>
8. Kolovski, V., Hendler, J., Parsia, B.: Analyzing web access control policies. In: *Proceedings of the 16th international conference on World Wide Web*. pp. 677–686. *WWW ’07*, ACM, New York, NY, USA (2007), <http://doi.acm.org/10.1145/1242572.1242664>
9. Libkin, L., Sirangelo, C.: Reasoning about xml with temporal logics and automata. *J. Applied Logic* 8(2), 210–232 (2010)
10. de Moura, L.M., Bjørner, N.: Z3: An efficient smt solver. In: Ramakrishnan, C.R., Rehof, J. (eds.) *TACAS. Lecture Notes in Computer Science*, vol. 4963, pp. 337–340. Springer (2008)
11. Pan, G., Sattler, U., Vardi, M.Y.: Bdd-based decision procedures for the modal logic k . *Journal of Applied Non-Classical Logics* 16(1-2), 169–208 (2006)
12. Pan, G., Sattler, U., Vardi, M.Y.: BDD-based decision procedures for the modal logic K . *Journal of Applied Non-classical Logics* 16(1-2), 169–208 (2006)
13. Tanabe, Y., Takahashi, K., Hagiya, M.: A decision procedure for alternation-free modal μ -calculi. In: Areces, C., Goldblatt, R. (eds.) *Advances in Modal Logic*. pp. 341–362. College Publications (2008)
14. Tanabe, Y., Takahashi, K., Yamamoto, M., Tozawa, A., Hagiya, M.: A decision procedure for the alternation-free two-way modal μ -calculus. In: Beckert, B. (ed.) *TABLEAUX. Lecture Notes in Computer Science*, vol. 3702, pp. 277–291. Springer (2005)

A Proofs

We define the *number of recursive expansions* of F or φ , written $\mathcal{E}_\emptyset(F)$, in Figure 7.

Definition A1 Given φ , and Γ , we define $(\Gamma)_X^\varphi$ as follows.

$$\{G' \mid X \prec G' \wedge (G' \{\varphi/X\} \in \Gamma \vee \overline{G'} \{\varphi/X\} \in \Gamma)\}$$

$$\begin{aligned}
\mathcal{E}_\Gamma(\top) &= \mathcal{E}_\Gamma(\perp) = \mathcal{E}_\Gamma(x) = \mathcal{E}_\Gamma(X) = \mathcal{E}_\Gamma(\neg X) = 0 \\
\mathcal{E}_\Gamma(\sigma) &= \mathcal{E}_\Gamma(\neg\sigma) = \mathcal{E}_\Gamma(\neg\langle a \rangle \top) = 0 \\
\mathcal{E}_\Gamma(F \vee G) &= \mathcal{E}_\Gamma(F \wedge G) = \mathcal{E}_\Gamma(F) + \mathcal{E}_\Gamma(G) \\
\mathcal{E}_\Gamma(\langle a \rangle F) &= \mathcal{E}_\Gamma(F) \\
\mathcal{E}_\Gamma(\mu x.F) &\stackrel{\text{def}}{=} \mathcal{E}_\Gamma(F) \quad \text{if } x \not\prec F \\
\mathcal{E}_\Gamma(\mu x.F) &\stackrel{\text{def}}{=} 0 \quad \text{if } \mu x.F \in \Gamma \text{ or } \mu x.\bar{F} \in \Gamma \\
\mathcal{E}_\Gamma(\mu x.F) &\stackrel{\text{def}}{=} \mathcal{E}_{\Gamma \cup \{\mu x.F\}}(F\{\mu x.F/x\}) + 1 \quad \text{otherwise}
\end{aligned}$$

Fig. 7. Number of Recursive Expansion

Lemma A2 We have $\overline{F\{G/x\}} = \overline{F}\{\bar{G}/x\}$.

Proof. By induction on F , relying on the fact that $\bar{x} = x$.

Lemma A3 We have $\overline{F\{G/X\}} = \overline{F}\{G/X\}$.

Proof. By induction on F , relying on the fact that $\bar{X} = \neg X$.

Lemma A4 If $\Gamma \subseteq \Gamma'$, then $\bar{\mathcal{L}}_{\Gamma'}(F) \subseteq \bar{\mathcal{L}}_\Gamma(F)$.

Proof. By induction on the lexical order of $\mathcal{E}_\Gamma(F)$ and the size of F . Base cases are immediate. For conjunction and disjunction, we may apply the induction hypothesis because $\mathcal{E}_\Gamma(F_1)$ and $\mathcal{E}_\Gamma(F_2)$ do not increase and the formula size decreases. This is also the case for the modality case.

For the recursion case where $x \prec F$, we distinguish three cases (we do not mention the negation $\mu x.\bar{F}$ in these cases):

- if $\mu x.F \in \Gamma$, then necessarily $\mu x.F \in \Gamma'$ and we immediately conclude;
- if $\mu x.F \in \Gamma'$, then we conclude by $\emptyset \subseteq \mathcal{S}$ for any set \mathcal{S} ;
- otherwise, we have $\mu x.F$ in neither set, and we apply the induction hypothesis, as $\mathcal{E}_{\Gamma \cup \{\mu x.F\}}(F\{\mu x.F/x\})$ is strictly smaller than $\mathcal{E}_\Gamma(\mu x.F)$.

Lemma A5 We have $\bar{\mathcal{L}}_{\Gamma \cup \{\mu x.F\}}(G) = \bar{\mathcal{L}}_{\Gamma \cup \{\mu x.\bar{F}\}}(G)$.

Proof. By an immediate induction on the lexical order of $\mathcal{E}_\Gamma(G)$ and the size of G .

Lemma A6 For any F , we have $\bar{\mathcal{L}}_\Gamma(F) = \bar{\mathcal{L}}_\Gamma(\bar{F})$.

Proof. By induction on the lexical order of $\mathcal{E}_\Gamma(F)$ and the size of F .

The base cases \top , \perp , x , X , $\neg X$, σ , $\neg\sigma$, $\neg\langle a \rangle \top$, and $\mu x.F$ where $x \not\prec F$ are immediate.

For $F \wedge G$, we compute as follows.

$$\begin{aligned}
\overline{\mathcal{L}}_I(F \wedge G) &= \overline{\mathcal{L}}_I(F) \cup \overline{\mathcal{L}}_I(G) \\
&= \overline{\mathcal{L}}_I(\overline{F}) \cup \overline{\mathcal{L}}_I(\overline{G}) && \text{by induction} \\
&= \overline{\mathcal{L}}_I(\overline{F} \vee \overline{G}) \\
&= \overline{\mathcal{L}}_I(\overline{F \wedge G})
\end{aligned}$$

The disjunction case is similar.

For the recursion case, if $\mu x.F$ or $\mu x.\overline{F}$ is in I , then $\mu x.\overline{F}$ or $\mu x.\overline{\overline{F}} = \mu x.F$ is also in I and the result follows.

Finally, if neither is in I , we compute as follows.

$$\begin{aligned}
\overline{\mathcal{L}}_I(\mu x.F) &= \overline{\mathcal{L}}_{I \cup \{\mu x.F\}}(F\{\mu x.F/x\}) \\
&= \overline{\mathcal{L}}_{I \cup \{\mu x.F\}}(\overline{F\{\mu x.F/x\}}) && \text{by induction} \\
&= \overline{\mathcal{L}}_{I \cup \{\mu x.F\}}(\overline{F}\{\mu x.\overline{F}/x\}) && \text{by Lemma A2} \\
&= \overline{\mathcal{L}}_{I \cup \{\mu x.\overline{F}\}}(\overline{F}\{\mu x.\overline{F}/x\}) && \text{by Lemma A5} \\
&= \overline{\mathcal{L}}_I(\mu x.\overline{F})
\end{aligned}$$

Lemma A7 For all I and F , if $X \not\prec F$, then $\overline{\mathcal{L}}_{(I)_X^\varphi}(F) = \overline{\mathcal{L}}_\emptyset(F)$.

Proof. We prove the more general result: for all I' , $\overline{\mathcal{L}}_{(I)_X^\varphi \cup I'}(F) = \overline{\mathcal{L}}_{I'}(F)$ by induction on the lexical order of $\mathcal{E}_{I'}(F)$ and the size of F .

The result is immediate for the base cases, and by induction for the conjunction, disjunction, and modality cases. For the recursion case, if $\mu x.F$ (or its negation) is in $(I)_X^\varphi \cup I'$, then it must be in I' as $X \not\prec F$ and members of $(I)_X^\varphi$ contain X by definition. Thus both sides are equal to \emptyset . If neither $\mu x.F$ nor its negation are in $(I)_X^\varphi \cup I'$, we have $\overline{\mathcal{L}}_{(I)_X^\varphi \cup I'}(\mu x.F) = \overline{\mathcal{L}}_{(I)_X^\varphi \cup I' \cup \{\mu x.F\}}(F\{\mu x.F/x\})$

We next apply the induction hypothesis with $I' \cup \{\mu x.F\}$ and $F\{\mu x.F/x\}$, thus we have $\overline{\mathcal{L}}_{(I)_X^\varphi \cup I' \cup \{\mu x.F\}}(F\{\mu x.F/x\}) = \overline{\mathcal{L}}_{I' \cup \{\mu x.F\}}(F\{\mu x.F/x\}) = \overline{\mathcal{L}}_{I'}(\mu x.F)$.

We conclude by taking $I' = \emptyset$.

Lemma A8 Let F be a formula mentioning X , and φ a closed formula. We have $\overline{\mathcal{L}}_\emptyset(F\{\varphi/X\}) \subseteq \overline{\mathcal{L}}_\emptyset(F)\{\varphi/X\} \cup \overline{\mathcal{L}}_\emptyset(\varphi)$.

Proof. We prove the following more general property for any I by induction on the lexical order of $\mathcal{E}_{(I)_X^\varphi}(F)$ and the size of F .

$$\overline{\mathcal{L}}_I(F\{\varphi/X\}) \subseteq \overline{\mathcal{L}}_{(I)_X^\varphi}(F)\{\varphi/X\} \cup \overline{\mathcal{L}}_\emptyset(\varphi)$$

We first deal with every case where $X \not\vdash F$. In this case, X also does not occur in $\overline{\mathcal{L}}_F(F)$.

$$\begin{aligned}
\overline{\mathcal{L}}_F(F\{\varphi/X\}) &= \overline{\mathcal{L}}_F(F) \\
&\subseteq \overline{\mathcal{L}}_\emptyset(F) && \text{by Lemma A4} \\
&= \overline{\mathcal{L}}_{(F)_X^\varphi}(F) && \text{by Lemma A7} \\
&= \overline{\mathcal{L}}_{(F)_X^\varphi}(F)\{\varphi/X\} \\
&\subseteq \overline{\mathcal{L}}_{(F)_X^\varphi}(F)\{\varphi/X\} \cup \overline{\mathcal{L}}_\emptyset(\varphi)
\end{aligned}$$

Case X . We compute as follows, using Lemma A4 for the last inclusion.

$$\overline{\mathcal{L}}_F(X\{\varphi/X\}) = \overline{\mathcal{L}}_F(\varphi) \subseteq \overline{\mathcal{L}}_\emptyset(\varphi)$$

Case $\neg X$. We compute as follows, using Lemma A4 for the set inclusion, and Lemma A6 to conclude.

$$\overline{\mathcal{L}}_F((\neg X)\{\varphi/X\}) = \overline{\mathcal{L}}_F(\overline{\varphi}) \subseteq \overline{\mathcal{L}}_\emptyset(\overline{\varphi}) = \overline{\mathcal{L}}_\emptyset(\varphi)$$

Case $F \wedge G$. We compute as follows.

$$\begin{aligned}
&\overline{\mathcal{L}}_F((F \wedge G)\{\varphi/X\}) \\
&= \overline{\mathcal{L}}_F(F\{\varphi/X\}) \cup \overline{\mathcal{L}}_F(G\{\varphi/X\}) \\
&\subseteq \overline{\mathcal{L}}_{(F)_X^\varphi}(F)\{\varphi/X\} \cup \overline{\mathcal{L}}_{(G)_X^\varphi}(G)\{\varphi/X\} \cup \overline{\mathcal{L}}_\emptyset(\varphi) && \text{by induction} \\
&= \overline{\mathcal{L}}_{(F \wedge G)_X^\varphi}(F \wedge G)\{\varphi/X\} \cup \overline{\mathcal{L}}_\emptyset(\varphi)
\end{aligned}$$

Case $F \vee G$. Identical to the previous case.

Case $\langle a \rangle F$. We compute as follows, using Lemma A3 and the induction hypothesis.

$$\begin{aligned}
&\overline{\mathcal{L}}_F((\langle a \rangle F)\{\varphi/X\}) \\
&= \overline{\mathcal{L}}_F(\langle a \rangle (F\{\varphi/X\})) \\
&= \{\langle a \rangle F\{\varphi/X\}; \langle a \rangle \overline{F\{\varphi/X\}}; \langle a \rangle \top\} \cup \overline{\mathcal{L}}_F(F\{\varphi/X\}) \\
&= \{\langle a \rangle F; \langle a \rangle \overline{F}; \langle a \rangle \top\}\{\varphi/X\} \cup \overline{\mathcal{L}}_F(F\{\varphi/X\}) \\
&\subseteq \{\langle a \rangle F; \langle a \rangle \overline{F}; \langle a \rangle \top\}\{\varphi/X\} \cup \overline{\mathcal{L}}_{(F)_X^\varphi}(F)\{\varphi/X\} \cup \overline{\mathcal{L}}_\emptyset(\varphi) \\
&= \overline{\mathcal{L}}_{(\langle a \rangle F)_X^\varphi}(\langle a \rangle F)\{\varphi/X\} \cup \overline{\mathcal{L}}_\emptyset(\varphi)
\end{aligned}$$

Case $\mu x.F$ with $x \not\vdash F$. We compute as follows.

$$\begin{aligned}
&\overline{\mathcal{L}}_F((\mu x.F)\{\varphi/X\}) \\
&= \overline{\mathcal{L}}_F(\mu x.F\{\varphi/X\}) && \varphi \text{ closed} \\
&= \overline{\mathcal{L}}_F(F\{\varphi/X\}) && x \not\vdash F\{\varphi/X\} \\
&\subseteq \overline{\mathcal{L}}_{(F)_X^\varphi}(F)\{\varphi/X\} \cup \overline{\mathcal{L}}_\emptyset(\varphi) && \text{by induction} \\
&= \overline{\mathcal{L}}_{(\mu x.F)_X^\varphi}(\mu x.F)\{\varphi/X\} \cup \overline{\mathcal{L}}_\emptyset(\varphi) && x \not\vdash F
\end{aligned}$$

Case $\mu x.F$ with $x \prec F$.

If we have $\mu x.F\{\varphi/X\} \in \Gamma$ or $\mu x.\overline{F}\{\varphi/X\} \in \Gamma$ then $\overline{\mathcal{L}}_\Gamma(\mu x.F\{\varphi/X\}) = \emptyset$ and the result is immediate.

Otherwise we compute as follows.

$$\begin{aligned}
& \overline{\mathcal{L}}_\Gamma((\mu x.F)\{\varphi/X\}) \\
&= \overline{\mathcal{L}}_\Gamma(\mu x.F\{\varphi/X\}) \quad \varphi \text{ closed} \\
&= \overline{\mathcal{L}}_{\Gamma \cup \{\mu x.F\{\varphi/X\}\}}(F\{\varphi/X\}\{\mu x.F\{\varphi/X\}/x\}) \\
&= \overline{\mathcal{L}}_{\Gamma \cup \{\mu x.F\{\varphi/X\}\}}(F\{\mu x.F/x\}\{\varphi/X\}) \quad \varphi \text{ closed}
\end{aligned}$$

To apply the induction hypothesis, we show that

$$\mathcal{E}_{(\Gamma)_X^\varphi}(\mu x.F) = \mathcal{E}_{(\Gamma \cup \{\mu x.F\{\varphi/X\}\})_X^\varphi}(F\{\mu x.F/x\}) + 1.$$

First, we have $\mu x.F \notin (\Gamma)_X^\varphi$ and $\mu x.\overline{F} \notin (\Gamma)_X^\varphi$, since otherwise, we would have $\mu x.F\{\varphi/X\} \in \Gamma$ or $\mu x.\overline{F}\{\varphi/X\} = \mu x.F\{\varphi/X\} \in \Gamma$, which we assumed to be false. Thus $\mathcal{E}_{(\Gamma)_X^\varphi}(\mu x.F) = \mathcal{E}_{(\Gamma)_X^\varphi \cup \{\mu x.F\}}(F\{\mu x.F/x\}) + 1$. Next, we have $(\Gamma \cup \{\mu x.F\{\varphi/X\}\})_X^\varphi = (\Gamma)_X^\varphi \cup \{\mu x.F\}$ by definition. Thus we have $\mathcal{E}_{(\Gamma)_X^\varphi}(\mu x.F) = \mathcal{E}_{(\Gamma \cup \{\mu x.F\{\varphi/X\}\})_X^\varphi}(F\{\mu x.F/x\}) + 1$.

We may thus apply the induction hypothesis and continue to compute.

$$\begin{aligned}
& \subseteq \overline{\mathcal{L}}_{(\Gamma \cup \{\mu x.F\{\varphi/X\}\})_X^\varphi}(F\{\mu x.F/x\})\{\varphi/X\} \cup \overline{\mathcal{L}}_\emptyset(\varphi) \\
&= \overline{\mathcal{L}}_{(\Gamma)_X^\varphi \cup \{\mu x.F\}}(F\{\mu x.F/x\})\{\varphi/X\} \cup \overline{\mathcal{L}}_\emptyset(\varphi)
\end{aligned}$$

As neither $\mu x.F$ nor $\mu x.\overline{F}$ are in $(\Gamma)_X^\varphi$, we have $\overline{\mathcal{L}}_{(\Gamma)_X^\varphi}(\mu x.F) = \overline{\mathcal{L}}_{(\Gamma)_X^\varphi \cup \{\mu x.F\}}(F\{\mu x.F/x\})$. We may conclude the computation.

$$= \overline{\mathcal{L}}_{(\Gamma)_X^\varphi}(\mu x.F)\{\varphi/X\} \cup \overline{\mathcal{L}}_\emptyset(\varphi)$$

We complete the proof by taking Γ to be \emptyset and remarking that $(\emptyset)_X^\varphi = \emptyset$.