

Finite-Choice Logic Programming

ROBERT J. SIMMONS, Unaffiliated, USA

MICHAEL ARNTZENIUS, RelationalAI, United Kingdom

CHRIS MARTENS, Northeastern University, USA

Logic programming, as exemplified by datalog, defines the meaning of a program as its unique smallest model: the deductive closure of its inference rules. However, many problems call for an enumeration of models that vary along some set of choices while maintaining structural and logical constraints—there is no single canonical model. The notion of *stable models* for logic programs with negation has successfully captured programmer intuition about the set of valid solutions for such problems, giving rise to a family of programming languages and associated solvers known as answer set programming. Unfortunately, the definition of a stable model is frustratingly indirect, especially in the presence of rules containing free variables.

We propose a new formalism, *finite-choice logic programming*, that treats choice, not negation, as the fundamental primitive of nondeterminism. Finite-choice logic programming allows for straightforward expression of common idioms in answer set programming, gives meaning to a new and useful class of programs, and enjoys a least-fixed-point interpretation over a novel domain. We present a tractable algorithm for nondeterministically exploring the solution space and prove it correct with respect to our semantics. Our implementation, the DUSA logic programming language, has performance that compares favorably with state-of-the-art answer set solvers and exhibits more predictable scaling with problem size.

CCS Concepts: • **Computing methodologies** → **Logic programming and answer set programming**; • **Theory of computation** → **Constraint and logic programming**; **Denotational semantics**; • **Software and its engineering** → **Constraint and logic languages**.

Additional Key Words and Phrases: Datalog, answer set programming, semi-naive forward-chaining evaluation, possibility spaces

ACM Reference Format:

Robert J. Simmons, Michael Arntzenius, and Chris Martens. 2018. Finite-Choice Logic Programming. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (Draft)*. ACM, New York, NY, USA, 49 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 INTRODUCTION

Many important problem domains involve generating varied data according to structural and logical constraints. Examples include property-based random testing for typed functional programs [Claessen and Hughes 2011; Goldstein et al. 2023; Goldstein and Pierce 2022; Lampropoulos et al. 2017; Paraskevopoulou et al. 2022; Seidel et al. 2015], procedural content generation in games [Dormans and Bakkes 2011; Shaker et al. 2016; Short and Adams 2017], and software configuration [Czarnecki et al. 2004, 2002; Le et al. 2011]. To solve these problems, we write *generative*

Authors' Contact Information: Robert J. Simmons, robsimmons@gmail.com, Unaffiliated, Boston, Massachusetts, USA; Michael Arntzenius, RelationalAI, Somewhere in the, United Kingdom; Chris Martens, c.martens@northeastern.edu, Northeastern University, Boston, Massachusetts, USA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Draft, Draft, Draft

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-XXXX-X/18/06

<https://doi.org/XXXXXXX.XXXXXXX>

programs, characterized by *choice points* that create multiple possible outcomes and *constraints* that eliminate undesirable outcomes. Generative programs often make use of randomness to produce different solutions when the program is run multiple times. From a declarative perspective, these generative programs characterize *possibility spaces*, and the meaning of a program is the entire space, which may be considered independently of particular methods for sampling individual solutions.

Using logic to describe possibility spaces has a long history, and its expressive power has been harnessed by established tools such as satisfiability solvers and logic programming languages. However, different logic-based tools disagree in important ways about how logical statements should be interpreted and thus how possibility spaces should be expressed.

Consider the problem of procedurally generating maps for a virtual world of connected regions where every region has one of three terrain types—*mountain*, *ocean*, or *forest*—and where oceans never directly adjoin mountains. We can model the possibility space of terrain maps as a Boolean satisfiability (SAT) problem by saying that the truth of a proposition *mountain*(*r*) means that a region *r* has mountainous terrain (and so on):

$$(\text{mountain}(r) \vee \text{forest}(r) \vee \text{ocean}(r)) \leftrightarrow \text{region}(r) \quad (1)$$

$$(\neg \text{mountain}(r) \wedge \neg \text{forest}(r)) \vee (\neg \text{mountain}(r) \wedge \neg \text{ocean}(r)) \vee (\neg \text{forest}(r) \wedge \neg \text{ocean}(r)) \quad (2)$$

$$(\text{adjacent}(r_1, r_2) \wedge \text{ocean}(r_1)) \rightarrow (\text{ocean}(r_2) \vee \text{forest}(r_2)) \quad (3)$$

The conjunction of these three propositions neatly captures the desired possibility space: when augmented with a finite universe of terms (possible regions) and a collection of input literals characterizing the *adjacent* relation, a SAT solver would return valid terrain-to-region assignments.

Now suppose we additionally wish to enumerate all regions reachable from a given starting region by traversing only forests, perhaps with the goal of checking whether a player's initial position can reach both mountains and ocean. Reachability is naturally characterized by introducing a recursively-defined proposition *reach*(*r*) which indicates that the region *r* is reachable from the starting region:

$$\text{reach}(r) \leftarrow \text{start}(r) \quad (4)$$

$$\text{reach}(r_1) \leftarrow \text{adjacent}(r_1, r_2), \text{reach}(r_2), \text{forest}(r_2) \quad (5)$$

In this instance, logic programming gives the desired interpretation: from a rule $p \leftarrow q_1 \dots q_n$, and a database containing $q_1 \dots q_n$, we deduce *p*, where *p* and *q_i* are all positive assertions. Such a program has a single canonical model, the smallest set of assertions closed under the given implications; any other assertion is deemed false [Van Gelder et al. 1991].

Unfortunately, these two approaches don't play well together. On one hand, logic programming lacks a notion of *choice point*, because propositions are interpreted as rules: not merely justifications for what *may* live in the model, but assertions about what *must* appear in the model, so the first program has no direct analog in, say, datalog. On the other hand, defining transitive closure is notoriously difficult for SAT solvers: interpreting the reachability program as a logical formula using Clark's completion [1978], a SAT solver would validate solutions with spurious *reach* facts that have no justification imparted by the programmer.

The programming model we desire seems to alternate between two modes of operation: on the one hand, making *mutually exclusive choices* that multiply our possibility space and rule out certain combinations with other choices; and on the other hand, computing the *deductive consequences* of those choices. Unifying these perspectives is precisely what we achieve in this paper with finite-choice logic programming, a new approach to logic programming in which the generative

constraints in eqs. (1) to (3) can be expressed as follows:

$$\text{terrain}(r) \text{ is } \{\text{mountain}, \text{forest}, \text{ocean}\} \leftarrow \text{region}(r) \quad (6)$$

$$\text{terrain}(r_1) \text{ is } \{\text{forest}, \text{ocean}\} \leftarrow \text{adjacent}(r_1, r_2), \text{terrain}(r_2) \text{ is ocean} \quad (7)$$

The is syntax signals a *functional dependency* from the region r to its terrain type, and the rules represent mutually exclusive choices closed under deduction: a region must have one of three terrain types, and a region adjacent to the ocean must contain oceans or forests. In finite-choice logic programming, such rules combine naturally with rules like eqs. (4) and (5); we would rewrite the latter rule as $(\text{reach}(r_1) \leftarrow \text{adjacent}(r_1, r_2), \text{reach}(r_2), \text{terrain}(r_2) \text{ is forest})$.

1.1 Answer set programming and its discontents

Historically, the most successful approach for logic programming with mutually exclusive choices has been the *stable model semantics* introduced by Gelfond and Lifschitz [1988]. Stable models combine the Boolean-satisfiability intuitions behind Clark’s completion with a rejection of circular justification, and provide the foundation for *answer set programming*. Systems for computing the stable models of answer set programs have fruitfully co-evolved with the advancements in Boolean satisfiability solving, leading to sophisticated heuristics that make many problems fast in practice [Gebser et al. 2019]. Answer set programming has seen considerable and ongoing success in procedural content generation [Dabral and Martens 2020; Neufeld et al. 2015; Smith and Bryson 2014; Smith et al. 2013; Smith and Mateas 2011; Summerville et al. 2018], analysis and scenario generation for legal regulations, policies, and contracts [Dabral et al. 2023; Lim et al. 2022], spatial reasoning for built structures [Li et al. 2020], and distributed system reasoning [Alvaro et al. 2011].

In a modern answer set programming language like Clingo [Gebser et al. 2011], the map-generation criteria expressed in eqs. (1) to (3) can be expressed concisely as follows:

$$1 \{ \text{mountain}(r); \text{forest}(r); \text{ocean}(r) \} 1 \leftarrow \text{region}(r) \quad (8)$$

$$\leftarrow \text{adjacent}(r_1, r_2), \text{ocean}(r_1), \text{mountain}(r_2) \quad (9)$$

Equation (8) contains a “cardinality constraint” that insists on precisely one of the three terrain types holding for each region. Equation (9) is a “headless” rule, interpreted as a constraint forbidding the conjunction of all three premises from holding simultaneously.

Unfortunately, the interpretation of answer set programs involves multiple levels of indirection that make it challenging to reason about them. Firstly, the stable model semantics that justifies answer set programming is defined only for logic programs without free variables. Answer set programming presents clients with an interface of rules containing free variables, such as eq. (9) above; however, in a program with sixteen regions, the stable model semantics applies after expanding this rule into 256 variable-free rules, one for each assignment of regions to r_1 and r_2 . This is reflected in essentially all implementations of answer set programming, which involve the interaction of a *solver* that only understands variable-free rules and a *grounder* that generates variable-free rules, usually incorporating heuristics to minimize the number of rules the solver must deal with.

The second level of indirection is that even propositional answer set programs do not directly have a semantics: the definition of stable models involves a syntactic transformation of an answer set program into a logic program without negation (the *reduct*, see [Gelfond and Lifschitz 1988]). This transformation is defined with respect to a candidate model, and if the unique model of the reduct is the same as the candidate model, the candidate model is accepted as an actual model. This fixed-point-like definition is what the “stability” of stable models refers to.

The third level of indirection is that stable models are defined in terms of *negation* and not in terms of the higher-level constructs — choice, cardinality, and headless rules — that are foundational

for essentially all modern applications of answer set programming. Higher-level constructs are either justified by translation to “pure” answer set programming with only negation (as in [Sacca and Zaniolo 1990]) or by appeal to a nonstandard definition of stable models (as in [Simons et al. 2002]).

1.2 A constructive semantics based on mutually-exclusive choices

In this paper we develop finite-choice logic programming as a logic programming language with *choice*, not negation, as primitive. We present and connect two semantics for finite-choice logic programming: a set-based semantics that describes the incremental construction of solutions (Section 2) and an interpretation of programs as the least fixed point of a monotonic immediate consequence operator over the domain of possibility spaces (Section 5). We demonstrate that finite-choice logic programming subsumes both datalog (Section 2.2) and answer set programming (Section 2.3), and present new examples of idiomatic finite-choice logic programs (Section 3). We present an algorithm for nondeterministic enumeration of solutions (Section 6) and describe an implementation (Section 7) that outperforms state-of-the-art answer set programming engines on a variety of examples, in part because our implementation avoids the “grounding bottleneck” encountered by ASP solvers that take a ground-then-solve approach to program execution.

2 DEFINING FINITE-CHOICE LOGIC PROGRAMMING

At the core of finite-choice logic programming are facts, which take the form $p(\bar{t}) \text{ IS } v$, where $p(\bar{t})$ is the attribute and v is the unique value assigned to that attribute. A set of facts (which we’ll call a *database*) must map each attribute to at most one value. Concretely, a fact like *terrain(home) IS forest* indicates that the region named *home* contains forest terrain, and cannot contain any other terrain. This is called a *functional dependency*: the predicate *terrain* is a partial function from region names to terrain types.

Definition 2.1 (Terms). As common in logic programming settings, terms are Herbrand structures, either variables x, y, z, \dots or uninterpreted functions $f(t_1, \dots, t_n)$ where the arguments t_i are terms. Constants are uninterpreted functions with no arguments, and as usual we’ll leave the parentheses off and just write *b* or *c* instead of $b()$ or $c()$. We’ll often abbreviate sequences of terms t_1, \dots, t_n as \bar{t} when the indices aren’t important.

Definition 2.2 (Facts). A fact has the form $p(t_1, \dots, t_n) \text{ IS } v$, where p is a *predicate* and the t_i and v are variable-free (i.e. *ground*) terms. We call $p(t_1, \dots, t_n)$ the fact’s *attribute* and call v the fact’s *value*. We will sometimes use a to stand in for a variable-free attributes $p(\bar{t})$.

Finite-choice logic programming admits specification of possibility spaces through the interplay of two kinds of rules:

Definition 2.3 (Rules). Rules $H \leftarrow F$ have one of two forms, *open* and *closed*.

$$p(\bar{t}) \text{ IS? } v \leftarrow F \quad (\text{open form rule})$$

$$p(\bar{t}) \text{ IS } \{v_1, \dots, v_m\} \leftarrow F \quad (\text{closed form rule, } m \geq 1)$$

In both cases, F is a conjunction of *premises* of the form $p(\bar{t}) \text{ IS } v$. The rule’s *conclusion* (or *head*) H is the part to the left of the \leftarrow symbol. Both the head and the premises may contain variables, but every variable in the head must appear in a premise.

For intuition: a *closed* conclusion, such as *terrain(port) IS {ocean, forest}*, *requires* that the region named *port* be either ocean or forest. An *open* conclusion, such as *terrain(goal) IS? meadow*, *requires* that the attribute *terrain(goal)* takes some value in the solution, and *permits* that value to be meadow. We make the semantics of these rule forms precise in the next section.

Definition 2.4 (Programs). A program P is a finite set of rules.

Definition 2.5 (Substitutions). A substitution σ is a total function from variables to ground terms. Applying a substitution to a term (σt) or a formula (σF) replaces all variables x in the term or formula with the term $\sigma(x)$.

2.1 Fact-set semantics

In this section, we present the meaning of finite-choice logic programs by describing their construction according to a nondeterministic semantics.

Definition 2.6 (Databases). A database D is a set of variable-free facts $p(\bar{t}) \text{ IS } v$ that is *consistent*, meaning that each attribute $p(\bar{t})$ maps to at most one value v : if $p(\bar{t}) \text{ IS } v \in D$ and $p(\bar{t}) \text{ IS } v' \in D$, then $v = v'$.

Definition 2.7 (Satisfaction). We say that a substitution σ *satisfies* F in the database D when, for each $p(\bar{t}) \text{ IS } v$ in F , the fact $p(\sigma\bar{t}) \text{ IS } \sigma v$ is present in D .

Definition 2.8 (Fact-set evolution). The relation $D \Rightarrow_P S$ relates a database (a consistent set of facts) to a set of databases (a set of consistent sets of facts):

- If P contains the closed-form rule $p(\bar{t}) \text{ IS } \{v_1, \dots, v_m\} \leftarrow F$ and σ satisfies F in D , then $D \Rightarrow_P S$, where S is the set of all $D \cup \{p(\sigma\bar{t}) \text{ IS } \sigma v_i\}$ for $1 \leq i \leq m$ where the result of the union is consistent (and therefore a database).
- If P contains the open-form rule $p(t_1, \dots, t_n) \text{ IS? } v \leftarrow F$ and σ satisfies F in D , then $D \Rightarrow_P S$, where S contains one or two elements. S always contains D itself, and if $D \cup \{p(\sigma\bar{t}) \text{ IS } \sigma v\}$ is a consistent set of facts then S contains $D \cup \{p(\sigma\bar{t}) \text{ IS } \sigma v\}$ as well.

To avoid a corner case with empty programs, we'll also say that $D \Rightarrow_P \{D\}$ always.

Definition 2.9 (Steps). We say that the program P allows D to *step* to D' if $D \Rightarrow_P S$ and $D' \in S$. Steps gives rise to *step sequences*: $D_1 \dots D_k$ is a step sequence for P if $k \in \mathbb{N}$ (that is, if the sequence is finite) and if, for each $i > 1$, the program P allows D_{i-1} to step to D_i .

Definition 2.10 (Saturation). A database D is *saturated under a program* P if its *only possible evolution* under the \Rightarrow_P relation is the singleton set containing itself. In other words, D is saturated under P if, for all S such that $D \Rightarrow_P S$, it is the case that $S = \{D\}$.

Definition 2.11 (Solutions). A *solution* to the program P is a saturated database D where $\emptyset \dots D$ is a step sequence for P .

The implications of this definition of fact-set evolution are a bit subtle. In general, for a given program P and database D , there may be many S such that $D \Rightarrow_P S$, as many as there are pairs of substitutions σ and rules $H \leftarrow F$ such that σ satisfies F in P . This means that each step in a step sequence resolves two levels of nondeterminism: to take a step from D to D' , first it is necessary to pick one of the possibly many S such that $D \Rightarrow_P S$, and then it is necessary to pick a database D' from that set S .

2.2 Simulating datalog

A datalog program without negation is a set of rules (Horn clauses) where all variables in the conclusion appear somewhere in a premise.

$$p(\bar{t}) \leftarrow p_1(\bar{t}_1), \dots, p_n(\bar{t}_n) \quad (\text{datalog rule})$$

This is a generic use of “datalog,” as often people take “Datalog” to specifically refer to “function-free” logic programs where term constants have no arguments, a condition sufficient to ensure

$$p \text{ is? ff} \leftarrow \quad (10)$$

$$q \text{ is? ff} \leftarrow \quad (11)$$

$$p \text{ is \{tt\}} \leftarrow q \text{ is ff} \quad (12)$$

$$q \text{ is \{tt\}} \leftarrow p \text{ is ff} \quad (13)$$

Fig. 1. The finite-choice logic program corresponding to the two rules $p \leftarrow \neg q$ and $q \leftarrow \neg p$ in ASP.

that every program has a finite model. We follow many theoretical developments and practical implementations of datalog in ignoring the function-free requirement.

Finite-choice logic programming can simulate a datalog proposition $p(\bar{t})$ as a fact of the form $p(\bar{t}) \text{ is unit}$, where **unit** is a newly introduced constant. This is consistent with the traditional interpretation of datalog so long as the predicate p only appears in premises of the form $p(\bar{t}) \text{ is unit}$ or in conclusions of the form $p(\bar{t}) \text{ is \{unit\}}$. This observation justifies our the use of value-free predicates in finite-choice logic programs, which we already saw with the use of *region*(r) in rule 6 and *adjacent*(r_1, r_2) in rule 7.

2.3 Simulating answer-set programming

Finite-choice logic programming can use the interplay of open and closed rules to obtain all the expressiveness of stable models without any reference to logical negation. An answer set program containing the two rules $p \leftarrow \neg q$ and $q \leftarrow \neg p$ corresponds to the finite-choice logic program in Figure 1. The open rules 10 and 11 unconditionally *permit* p or q to have the “false” value **ff**, and the closed rules 12 and 13 ensure that the assignment of **ff** to either p or q will force the other attribute to take the “true” value **tt**. (This asymmetry between the handling of truth and falsehood reflects their asymmetric treatment in answer set programming.) The two solutions for this program are $\{p \text{ is tt}, q \text{ is ff}\}$ and $\{p \text{ is ff}, q \text{ is tt}\}$, which correspond to the two solutions that answer set programming assigns to the source program.

Answer set programming is usually defined in terms of variable-free rules that have both non-negated premises p_i and negated premises $\neg q_i$.

$$p \leftarrow p_1, \dots, p_n, \neg q_1, \dots, \neg q_m \quad (\text{ASP rule})$$

A stable model of an answer set program takes the form of a set X of variable-free propositions: any proposition in X is treated as true, and any proposition not in X is treated as false. (For the full definition of stable models, see Gelfond and Lifschitz [1988] or Appendix A.2.)

THEOREM 2.12. *Let P be a finite collection of ASP rules. There exists a translation of P to a finite-choice logic program $\langle P \rangle$ such that the following hold:*

- For all stable models X of P , there is a solution D to $\langle P \rangle$ such that $X = \{p \mid p \text{ is tt} \in D\}$.
- For all solutions D of $\langle P \rangle$, the set $\{p \mid p \text{ is tt} \in D\}$ is a stable model of P .

The proof of Theorem 2.12 is available in supplemental materials (Appendix A).

2.4 Example program execution

Let P be the four rule program from Figure 1. We demonstrate the fact-set semantics step-by-step, starting from the database $D_0 = \emptyset$. Two rules have satisfied premises and thus generate evolutions:

Rules: $p \text{ IS? ff} \leftarrow$	Evolutions: $\emptyset \Rightarrow_P \{ \emptyset, \{p \text{ IS ff}\} \} = S_1$
$q \text{ IS? ff} \leftarrow$	$\emptyset \Rightarrow_P \{ \emptyset, \{q \text{ IS ff}\} \} = S_2$
$p \text{ IS \{tt\}} \leftarrow q \text{ IS ff}$	<i>rule does not apply</i>
$q \text{ IS \{tt\}} \leftarrow p \text{ IS ff}$	<i>rule does not apply</i>

There are three sets S such that $\emptyset \Rightarrow_P S$: by rule 10 we have $\emptyset \Rightarrow_P S_1$, by rule 11 we have $\emptyset \Rightarrow_P S_2$, and by the trivial evolution $\emptyset \Rightarrow_P \{\emptyset\}$.

It requires only a slight misreading of the definitions to incorrectly conclude that \emptyset is saturated: after all, in all cases that $\emptyset \Rightarrow_P S$, it is the case that $\emptyset \in S$. However, \emptyset is *not* saturated by Definition 2.7, because \emptyset can step to other databases as well.

If we choose to step from \emptyset to $\{p \text{ IS ff}\}$, three rules apply:

$p \text{ IS? ff} \leftarrow$	$\{p \text{ IS ff}\} \Rightarrow_P \{ \{p \text{ IS ff}\} \} = S_3$
$q \text{ IS? ff} \leftarrow$	$\{p \text{ IS ff}\} \Rightarrow_P \{ \{p \text{ IS ff}\}, \{p \text{ IS ff}, q \text{ IS ff}\} \} = S_4$
$p \text{ IS \{tt\}} \leftarrow q \text{ IS ff}$	<i>rule does not apply</i>
$q \text{ IS \{tt\}} \leftarrow p \text{ IS ff}$	$\{p \text{ IS ff}\} \Rightarrow_P \{ \{p \text{ IS ff}, q \text{ IS tt}\} \} = S_5$

As before, there are three possible evolutions and one of them is trivial. If we step to $\{p \text{ IS ff}, q \text{ IS ff}\}$ via S_4 , we will find ourselves in trouble:

$p \text{ IS? ff} \leftarrow$	$\{p \text{ IS ff}, q \text{ IS ff}\} \Rightarrow_P \{ \{p \text{ IS ff}, q \text{ IS ff}\} \}$
$q \text{ IS? ff} \leftarrow$	$\{p \text{ IS ff}, q \text{ IS ff}\} \Rightarrow_P \{ \{p \text{ IS ff}, q \text{ IS ff}\} \}$
$p \text{ IS \{tt\}} \leftarrow q \text{ IS ff}$	$\{p \text{ IS ff}, q \text{ IS ff}\} \Rightarrow_P \emptyset$
$q \text{ IS \{tt\}} \leftarrow p \text{ IS ff}$	$\{p \text{ IS ff}, q \text{ IS ff}\} \Rightarrow_P \emptyset$

This is yet another case where a slight misreading of the definitions could lead one to incorrectly conclude that $\{p \text{ IS ff}, q \text{ IS ff}\}$ is saturated: after all, we've just demonstrated that it is a database that can only step to itself. However, $\{p \text{ IS ff}, q \text{ IS ff}\}$ is not saturated by Definition 2.10: if we look at the last rule, it would seem to require that the program derive $q \text{ IS tt}$, conflicting with the existing fact $q \text{ IS ff}$. This conflict means that $\{p \text{ IS ff}, q \text{ IS ff}\} \Rightarrow_P \emptyset$, so the database is not saturated by Definition 2.10. The database also cannot take a step to any other database aside from itself, and so this represents a failed search: no series of steps will lead to a solution. Note that this case demonstrates why we cannot define a saturated database as "a database that can only step to itself." Such a definition would—undesirably—admit databases that evolve to the empty set.

If we back up and instead let $\{p \text{ IS ff}\}$ step to $\{p \text{ IS ff}, q \text{ IS tt}\}$ via S_5 , we can observe that the only way the database can evolve is to the singleton set containing itself, which is what Definition 2.10 required:

$p \text{ IS? ff} \leftarrow$	$\{p \text{ IS ff}, q \text{ IS tt}\} \Rightarrow_P \{ \{p \text{ IS ff}, q \text{ IS tt}\} \}$
$q \text{ IS? ff} \leftarrow$	$\{p \text{ IS ff}, q \text{ IS tt}\} \Rightarrow_P \{ \{p \text{ IS ff}, q \text{ IS tt}\} \}$
$p \text{ IS \{tt\}} \leftarrow q \text{ IS ff}$	<i>rule does not apply</i>
$q \text{ IS \{tt\}} \leftarrow p \text{ IS ff}$	$\{p \text{ IS ff}, q \text{ IS tt}\} \Rightarrow_P \{ \{p \text{ IS ff}, q \text{ IS tt}\} \}$

$$edge(x, y) \leftarrow edge(y, x) \quad (14)$$

$$root \text{ IS? } x \leftarrow edge(x, y) \quad (15)$$

$$parent(x) \text{ IS } \{x\} \leftarrow root \text{ IS } x \quad (16)$$

$$parent(y) \text{ IS? } x \leftarrow edge(x, y), parent(x) \text{ IS } z \quad (17)$$

Fig. 2. Calculating a spanning tree over an undirected graph.

$$edge(x, y) \leftarrow edge(y, x) \quad (18)$$

$$representative(x) \text{ IS? } x \leftarrow node(x) \quad (19)$$

$$representative(y) \text{ IS } \{z\} \leftarrow edge(x, y), representative(x) \text{ IS } z \quad (20)$$

Fig. 3. Appointing a canonical representative for each connected component in an undirected graph.

Because $\{p \text{ IS ff}, q \text{ IS tt}\}$ is saturated, it is a solution. Symmetric reasoning applies to see that $\{p \text{ IS tt}, q \text{ IS ff}\}$ is a solution. By inspection, there are no solutions where both p and q are assigned the same value.

3 FINITE-CHOICE LOGIC PROGRAMMING BY EXAMPLE

This section presents examples to demonstrate common idioms that arise naturally in writing and reasoning about finite-choice logic programs.

3.1 Spanning tree creation

Seeded with an *edge* relation, the finite-choice logic program in Figure 2 will pick an arbitrary node and construct a spanning tree rooted at that node. The structure of this program is such that it's not possible to make forward progress that indirectly leads to conflicts: rule 15 can only apply once in a series of deductions, and rules 16 and 17 cannot fire at all until some root is chosen. A node can only be added to the tree once, with a parent that already exists in the tree, so this is effectively a declarative description of Prim's algorithm without weights.

The creation of an arbitrary spanning tree for an undirected graph is a common first benchmark for datalog extensions that admit multiple solutions. Most previous work makes a selection greedily, either discards any future contradictory selections [Giannotti et al. 2001; Greco and Zaniolo 2001; Hu et al. 2021; Krishnamurthy and Naqvi 1988], or else avoiding contradictory deductions by consuming linear resources [Simmons and Pfenning 2008].

3.2 Appointing canonical representatives

When we want to check whether two nodes in an undirected graph are in the same connected component, one option is to compute the transitive closure of the *edge* relation. However, in a sparse graph, that can require computing $O(n^2)$ facts for a graph with n edges.

An alternative is to appoint an arbitrary member of each connected component as the canonical representative of that connected component: then, two nodes are in the same connected component if and only if they have the same canonical representative. This is the purpose of the program in Figure 3. In principle, it's quite possible for this program to get stuck in dead ends: if nodes *a* and *b* are connected by an edge, then rule 19 could appoint both nodes as a canonical representative, and

$visit(z) \leftarrow$	$visit(z) \text{ is \{tt\} \leftarrow}$	(21)
$visit(s(n)) \leftarrow more(n)$	$visit(s(n)) \text{ is \{tt\} \leftarrow more(n) \text{ is tt}$	(22)
	$more(n) \text{ is? ff \leftarrow visit(n) \text{ is tt}$	(23)
$stop(n) \leftarrow visit(n), \neg more(n)$	$stop(n) \text{ is \{tt\} \leftarrow visit(n) \text{ is tt, more(n) \text{ is ff}$	(24)
	$stop(n) \text{ is? ff \leftarrow visit(n) \text{ is tt}$	(25)
$more(n) \leftarrow visit(n), \neg stop(n)$	$more(n) \text{ is \{tt\} \leftarrow visit(n) \text{ is tt, stop(n) \text{ is ff}$	(26)

Fig. 4. At left, an answer set program with no finite grounding. At right, the translation of this program to a finite-choice logic program.

rule 16 would then prevent any extension of that database from being a solution. In the greedy-choice languages mentioned in the previous section, this would be a problem for correctness: incorrectly firing an analogue of rule 19 would mean that a final database might contain two canonical representatives in a connected component. In finite-choice logic programming, because closed rules can lead to the outright rejection of a database, this is merely a problem of efficiency: we would like to avoid going down these dead ends.

In finite-choice logic programs like this one, we can reason about avoiding certain dead ends by assuming a mode of execution that we call **deduce, then choose**. The deduce-then-choose strategy dictates that, when picking the next step in a step sequence, we will always choose a non-trivial evolution to a singleton ($D \Rightarrow_P \{D'\}$ with $D \neq D'$) over any evolution to a set containing two or more databases.

Endowed with the deduce-then-choose execution strategy, the program in Figure 3 will never make deductions that indirectly lead to conflicts. First, rule 18 will ensure that the edge relation is symmetric. Once that is complete, execution will be forced to choose some canonical representative using rule 19 in order to make forward progress. Once a representative is chosen, rule 20 will exhaustively assign that newly-appointed representative to every other node in the connected component. Only when it is done may rule 19 fire again for a node in another connected component.

3.3 Lazy answer set programming

Theorem 2.12 only describes a correspondence between variable-free finite-choice logic programs and variable-free answer set programs: given that this is how answer set programming is usually formally defined, it was difficult to do otherwise. However, the program transformation underlying Theorem 2.12 applies to the non-ground answer set programs that are usually written down in practice, and all the translations we have attempted are faithful translations of the source answer set program. We confidently conjecture that the correctness of the translation extends to non-ground programs, but we leave the formal details for future work.

Translating non-ground answer set programs is interesting in part because of answer set programs like the one in Figure 4. That program has a well-defined set of solutions, but none of these solutions can be enumerated by mainstream answer set programming implementations because the solver invokes an initial grounding step that is forced to generate an infinite set of ground rules. The translation of this program as a finite-choice logic program, on the other hand, has a simple operational interpretation under the deduce-then-choose strategy: starting from zero, each simple round of deduction will *visit* a successively larger natural number, at which point a choice will be made to either *stop* at that number or to continue to visit *more* numbers.

$$p \text{ IS } \{\text{tt}, \text{ff}\} \leftarrow \quad (27)$$

$$q \text{ IS } \{\text{tt}, \text{ff}\} \leftarrow \quad (28)$$

$$r \text{ IS } \{\text{tt}, \text{ff}\} \leftarrow \quad (29)$$

$$ok \text{ IS } \{\text{yes}\} \leftarrow \quad (30)$$

$$ok \text{ IS } \{\text{no}\} \leftarrow p \text{ IS } \text{ff}, q \text{ IS } \text{tt} \quad (31)$$

$$ok \text{ IS } \{\text{no}\} \leftarrow p \text{ IS } \text{tt}, q \text{ IS } \text{ff}, r \text{ IS } \text{ff} \quad (32)$$

Fig. 5. A finite-choice logic program representing the SAT instance $(p \vee \neg q) \wedge (\neg p \vee q \vee r)$.

This example suggests a connection between finite-choice logic programming and the strategy of using *lazy grounding* in answer set programming [Dal Palù et al. 2009], a strategy which similarly enables the evaluation of answer set programs with no finite grounding [Comploi-Taupe et al. 2023]. We will return to this point in Section 7.2.

3.4 Satisfiability

The previous examples all show how the evaluation of finite-choice logic programming can entirely avoid reaching databases that are not solutions. However, the full expressive power of finite-choice logic programming comes from the ability to represent problems where that avoidance is not possible. Since answer set programming generalizes boolean satisfiability, it is no surprise that boolean satisfiability problems can be represented straightforwardly in finite-choice logic programming.

A Boolean satisfiability problem can be written as a conjunction of clauses where each clause is a disjunction of propositions p and negated propositions $\neg p$. We represent such problems by explicitly assigning each proposition to **tt** or **ff** by a closed rule, and adding a rule for each clause that causes a value conflict for the *ok* predicate if the clause's negation holds; see Figure 5 for an example. The deduce-then-choose execution strategy gives no advantages here: the only deduction is observing inconsistencies that result from already-selected choices.

4 NONDETERMINISTIC IMMEDIATE CONSEQUENCES

The semantics given in Section 2 allow us to interpret the meaning of a finite-choice logic program as the set of the program's solutions (Definition 2.11). In some ways, though, this semantics leaves a lot to be desired. Consider the following program P :

$$p \text{ IS } \{\text{a}, \text{b}\} \leftarrow \quad (33)$$

$$p \text{ IS } \{\text{b}, \text{c}\} \leftarrow \quad (34)$$

$$q \text{ IS? } \text{ff} \leftarrow \quad (35)$$

$$q \text{ IS } \{\text{tt}\} \leftarrow p \text{ IS } x \quad (36)$$

Under the semantics in Section 2, there is a two-step sequence from \emptyset to $\{p \text{ IS } \text{c}, q \text{ IS } \text{ff}\}$:

$$\emptyset \Rightarrow_P \{ \{p \text{ IS } \text{b}\}, \{p \text{ IS } \text{c}\} \} \quad (\text{by rule 34})$$

$$\{p \text{ IS } \text{c}\} \Rightarrow_P \{ \{p \text{ IS } \text{c}\}, \{q \text{ IS } \text{ff}, p \text{ IS } \text{c}\} \} \quad (\text{by rule 35})$$

The result is not a solution, cannot be extended to a solution, and also cannot step to any other database. The first step led us, irrevocably, to a dead end: the step sequence $\emptyset, \{p \text{ IS } \text{c}\}, \dots$ can never be extended to reach a solution.

Some dead ends are unavoidable when conflicting assignments only occur down significant chains of deduction. That possibility is part of what gives finite-choice logic programming its expressive power! In the program above, though, the conflict is in some sense immediate: at each step, we have enough information to know that rule 33 and rule 34 both apply, and the overlap of these closed rules means that p can only be given the value b in any solution.

In this section, we will revisit the semantics of finite-choice logic programming to present a *immediate consequence* operator that captures global information about what is “immediately derivable” from a given program and database. This development will provide the basis for a least-fixed-point semantics of finite-choice logic programming (Section 5), as well as the foundation for our implementation of finite-choice logic programming (Section 6).

4.1 Bounded-complete posets

This development introduces several new concepts: *constraints*, *constraint databases*, and *choice sets*. All these are instances of the same semilattice-like structure: bounded-complete posets.

In Section 2, we define databases as sets with an auxiliary definition of consistency. Now we will generalize consistency to a notion of *compatibility*:

Definition 4.1 (Compatibility). If \mathcal{D} is a set equipped with a partial order \leq , then a subset $X \subseteq \mathcal{D}$ is *compatible* when it has an upper bound, i.e. $\exists y \in \mathcal{D}. \forall x \in X. x \leq y$. We write $\|X$ to assert that X is compatible and $\nparallel X$ for its negation. As a binary operator $x \parallel y = \|\{x, y\}$ and $x \nparallel y = \nparallel\{x, y\}$.

We will frequently use this basic fact about (in)compatibility:

LEMMA 4.2. *Compatibility is anti-monotone and incompatibility is monotone: if $D \leq D'$ and $E \leq E'$, then $D' \parallel E' \implies D \parallel E$, and contrapositively $D \nparallel E \implies D' \nparallel E'$.*

PROOF. $D' \parallel E'$ means D', E' have some upper bound D^* ; if $D \leq D'$ and $E \leq E'$, then D^* is also an upper bound for D, E . \square

Definition 4.3. A *bounded-complete poset* $(\mathcal{D}, \leq_{\mathcal{D}}, \perp_{\mathcal{D}}, \bigvee_{\mathcal{D}})$ is a poset with a least element where all compatible subsets have least upper bounds. In detail:

- (1) $\leq_{\mathcal{D}} \subseteq \mathcal{D} \times \mathcal{D}$ is a partial order (a reflexive, transitive, antisymmetric relation).
- (2) $\perp_{\mathcal{D}} \in \mathcal{D}$ is the least element: $\forall x \in \mathcal{D}. \perp \leq x$.
- (3) $\bigvee_{\mathcal{D}} : \{X \subseteq \mathcal{D} : \|X\} \rightarrow \mathcal{D}$ finds the least upper bound of a compatible set of elements: $(\forall x \in X. x \leq z) \implies \bigvee_{\mathcal{D}} X \leq z$. As a binary operator, $x \vee_{\mathcal{D}} y = \bigvee_{\mathcal{D}} \{x, y\}$.

Example 4.4. Databases as presented in Section 2.1 are bounded-complete posets, where $D_1 \leq D_2$ is the subset relation and \perp is the empty set \emptyset . The least upper bound operation is set union: if a set of databases are each individually subsets of some consistent set D of facts, then their union is a subset of D and so must also be a consistent set of facts.

4.2 Constraints and constraint databases

Definition 4.5. A *constraint*, $c \in \text{Constraint}$, is either $(\text{just } t)$ for some ground term t or $(\text{noneOf } X)$ for some set X of ground terms. Constraints form a bounded-complete poset as follows:

- (1) $\leq_{\text{Constraint}}$ is defined by cases:

$$\text{noneOf } X \leq \text{noneOf } Y \iff X \subseteq Y$$

$$\text{noneOf } X \leq \text{just } t \iff t \notin X$$

$$\text{just } t \leq \text{just } t' \iff t = t'$$

$$\text{just } t \not\leq \text{noneOf } X$$

(2) $\perp_{\text{Constraint}} = \text{noneOf } \emptyset$.

(3) $\bigvee_{\text{Constraint}} C$ is defined by cases. Because the least upper bound is only defined for compatible sets, if we know $\text{just } t \in C$ then $\text{just } t$ is the least upper bound: any other upper bound must have the form $\text{just } t'$ with $t = t'$. Otherwise, every $c_i \in C$ is of the form $\text{noneOf } X_i$, and their least upper bound is $\text{noneOf } (\bigcup_i X_i)$. By way of illustration, in the binary case:

$$\begin{aligned} \text{noneOf } X \vee \text{noneOf } Y &= \text{noneOf } (X \cup Y) \\ \text{noneOf } X \vee \text{just } t &= \text{just } t && \text{if } t \notin X \\ \text{just } t \vee \text{noneOf } X &= \text{just } t && \text{if } t \notin X \\ \text{just } t \vee \text{just } t' &= \text{just } t && \text{if } t = t' \end{aligned}$$

If none of these cases apply, $c_1 \vee c_2$ is undefined because $c_1 \not\ll c_2$.

Definition 4.6. A constraint database, $D, E \in \text{DB}$, is a map from ground attributes a to constraints $D[a]$. Constraint databases form a bounded-complete poset with structure inherited pointwise from *Constraint*:

(1) $D \leq_{\text{DB}} E \iff \forall a. D[a] \leq E[a]$.

(2) \perp_{DB} is mapping that takes every attribute a to $\perp_{\text{Constraint}} = \text{noneOf } \emptyset$.

(3) $\bigvee_{\text{DB}} S = a \mapsto \bigvee_{\text{Constraint}} \{D[a] : D \in S\}$, and so is defined whenever $\forall a. \parallel \{D[a] : D \in S\}$.

Definition 4.7. A constraint database D is *positive* if $D[a] = \text{noneOf } X$ implies $X = \emptyset$.

Definition 4.8. A constraint database D is *finite* if $\{a : D[a] \neq \text{noneOf } \emptyset\}$ is finite and all X such that $D[a] = \text{noneOf } X$ are finite.

We will introduce a new notation to describe finite constraint databases: $(p \mapsto \text{just ff}, q \mapsto \text{noneOf } \{\text{ff}\})$ represents the constraint database that takes p to just ff , takes q to $\text{noneOf } \{\text{ff}\}$, and takes every other attribute to $\perp_{\text{Constraint}} = \text{noneOf } \emptyset$.

There is an obvious isomorphism between positive constraint databases and the databases-as-consistent-sets-of-facts as introduced in Definition 2.6. The consistent sets of facts \emptyset , $\{p \text{ is tt}, q \text{ is tt}\}$, and $\{\text{edge}(a, b) \text{ is unit}\}$ correspond, respectively, to the constraint databases \perp_{DB} , $(p \mapsto \text{just tt}, q \mapsto \text{just ff})$, and $(\text{edge}(a, b) \mapsto \text{just unit})$.

4.3 Why constraint databases aren't enough

The move from “sets of facts” to “functions from attributes to constraints” is an instance of a common pattern encountered in extensions of logic programming to non-Boolean values. In the bilattice-annotated logic programming setting, Fitting [1993] calls the analogue of a constraint database a *valuation*, and Komendantskaya and Seda [2009] call the analogue of a constraint database an *annotation Herbrand model*. In the weighted logic programming setting, Eisner [2023] maintains a similar map ω from items to weights.

To our knowledge, all of this related work proceeds to define an immediate consequence operator as a function from database-analogues to database-analogues, and the meaning of a program is given as the unique least fixed point of this operator. But in finite-choice logic programming, as in answer set programming, programs can have multiple incompatible solutions, and so the meaning of a program *cannot* be a singular least fixed point of a function from constraint databases to constraint databases. One way forward is to present a function where solutions correspond to *any* fixed point, or any fixed point satisfying some additional condition: this is essentially the approach used by Fitting [1993] to bound the set of stable models for an answer set program using bilattices. However, this formulation doesn't directly characterize the stable models: it is still necessary to carry out the Gelfond-Lifschitz program transformation as a post-hoc check to see if one of

Fitting's bounded fixed points is a stable model. Our attempts to precisely characterize solutions for finite-choice logic programs as an arbitrary fixed point of some function from databases to databases encountered similar difficulties.

In this work, we take a different approach, which we believe is novel: in Section 5, the least-fixed-point interpretation of finite-choice logic programming is defined to be a *set of pairwise incompatible constraint databases*, which we call a *choice set*. Choice sets are therefore the next step towards defining our immediate consequence operator.

4.4 Choice sets

Choice sets form not merely a bounded-complete poset, but a complete lattice that has all least upper bounds. We will establish this in two steps, first defining choice sets as a pointed partial order, and then defining least upper bounds.

Definition 4.9. A choice set $C \in \text{Choice}$, is a *pairwise-incompatible* set of constraint databases, meaning that $(\forall D, E \in C. D \parallel E \implies D = E)$. *Choice* is a pointed partial order:

- (1) $C_1 \leq_{\text{Choice}} C_2 \iff (\forall D_2 \in C_2. \exists D_1 \in C_1. D_1 \leq D_2) \iff (\forall D_2 \in C_2. \exists! D_1 \in C_1. D_1 \leq D_2)$.
Existence $\exists D_1$ implies *unique* existence $\exists! D_1$ by pairwise incompatibility: if $D_1, D'_1 \in C_1$ are both $\leq D_2$ they are compatible and thus equal.
- (2) $\perp_{\text{Choice}} = \{\perp_{DB}\}$, that is, the set containing one item, the constraint database that maps every attribute to **noneOf** \emptyset .
- (3) $\top_{\text{Choice}} = \emptyset$.

In a constraint database D , each rule and rule-satisfying substitution in a finite-choice logic program will induce a choice set C_i , and deriving the immediate consequences of D involves a “parallel composition” of all these choices. The least upper bound $\bigvee_i C_i$ calculates this parallel composition. We will start with a finite example to build intuition:

Example 4.10. If we have $C_1 = \{D_a, D_b, D_c\}$ and $C_2 = \{D_x, D_y\}$, then $C_1 \vee C_2$ contains between zero and six elements: all of the least upper bounds out of $(D_a \vee D_x)$, $(D_b \vee D_x)$, $(D_c \vee D_x)$, $(D_a \vee D_y)$, $(D_b \vee D_y)$, and $(D_c \vee D_y)$ that are actually defined.

For the least upper bound of n choice sets, a database is in the least upper bound exactly when it is the least upper bound of n compatible databases, one drawn from each of the choice sets. Formally, we define the least upper bound of a collection $C_{i \in I}$ of choice sets indexed by some set I . Let $f : I \rightarrow DB$ be a function choosing one database from each choice set, so that $f(i) \in C_i$. If the set of databases thus chosen, $\text{Im}(f) = \{f(i) : i \in I\}$, is compatible, we include its least upper bound $\bigvee_i f(i)$ in the resulting choice set $\bigvee_i C_i$. The set of chosen databases, $\text{Im}(f) = \{f(i) : i \in I\}$, can be seen as a *candidate set* for the least upper bound. If $\text{Im}(f)$ is compatible, the least upper bound of $\text{Im}(f)$ exists and we include its least upper bound in the resulting choice set $\bigvee_i C_i$.

Definition 4.11 (Least upper bounds for Choice). Take any $\{C_i : i \in I\} \subseteq \text{Choice}$, and let $\prod_{i \in I} C_i$ be the set of all functions $f : I \rightarrow DB$ such that $f(i) \in C_i$. Then:

$$\bigvee_{i \in I} C_i = \left\{ \bigvee \text{Im}(f) : f \in \prod_{i \in I} C_i, \parallel \text{Im}(f) \right\}$$

It is not entirely trivial to show that $\bigvee_i C_i$ is a least upper bound in *Choice*: the proof is available in the supplemental material (Appendix B). The main subtleties are ensuring that we avoid combining incompatible databases and ensuring that the resulting set remains pairwise incompatible.

The partial order \leq_{Choice} , unlike the partial orders on *Constraint* and *DB*, has a greatest element (\emptyset), and so any collection of choice sets has an upper bound. Therefore, *Choice* is a complete lattice, not merely a bounded-complete poset.

4.5 Immediate consequence

Almost everything is now in place define an immediate consequence operator $\tau_P : DB \rightarrow Choice$. We only need to replay the definition of satisfaction from Section 2.1 in terms of constraint databases and choice sets, and then we can define the immediate consequence as the least upper bound of the consequence of every rule that can fire.

Definition 4.12 (Satisfaction). We say that a substitution σ satisfies F in the constraint database D when, for each premise $p(\bar{t}) \text{ is } v$ in F , we have $(\text{just } \sigma v) \leq D[p(\sigma \bar{t})]$.

Definition 4.13. A ground rule conclusion H defines a element of *Choice*, which we write as $\langle H \rangle$, in the following way:

- $\langle p(\bar{t}) \text{ is? } v \rangle = \{ (p(\bar{t}) \mapsto \text{just } v), (p(\bar{t}) \mapsto \text{noneOf } \{v\}) \}$
- $\langle p(\bar{t}) \text{ is } \{v_1, \dots, v_n\} \rangle = \{ (p(\bar{t}) \mapsto \text{just } v_1), \dots, (p(\bar{t}) \mapsto \text{just } v_n) \}$

Definition 4.14 (Immediate consequence). The immediate consequence operator $\tau_P : DB \rightarrow Choice$ is the least upper bound of every head of a rule with satisfied premises:

$$\tau_P(D) = \{D\} \vee \left(\bigvee \{ \langle \sigma H \rangle : (H \leftarrow F) \in P, \sigma F \leq D \} \right)$$

We forcibly ensure that $\{D\} \leq \tau_P(D)$ by making the result the least upper bound of $\{D\}$ itself and the least upper bound of all the satisfied rule heads. (We conjecture this is redundant for our subsequent developments, but it is also harmless.)

Example 4.15. Returning to the four rule program P at the beginning of this section (rules 33-36), these are all examples of how the immediate consequence operator for that program behaves on different inputs:

$$\begin{aligned} \tau_P(\perp_{DB}) &= \{ (p \mapsto \text{just } b, q \mapsto \text{just } ff), \\ &\quad (p \mapsto \text{just } b, q \mapsto \text{noneOf } \{ff\}) \} \\ \tau_P(p \mapsto \text{just } b, q \mapsto \text{just } ff) &= \emptyset \\ \tau_P(p \mapsto \text{just } b, q \mapsto \text{noneOf } \{ff\}) &= \{ (p \mapsto \text{just } b, q \mapsto \text{just } tt) \} \\ \tau_P(q \mapsto \text{just } ff) &= \{ (p \mapsto \text{just } b, q \mapsto \text{just } ff) \} \\ \tau_P(p \mapsto \text{just } b) &= \{ (p \mapsto \text{just } b, q \mapsto \text{just } tt) \} \\ \tau_P(p \mapsto \text{just } b, q \mapsto \text{just } tt) &= \{ (p \mapsto \text{just } b, q \mapsto \text{just } tt) \} \\ \tau_P(q \mapsto \text{just } \text{someOtherTerm}) &= \{ (p \mapsto \text{just } b, q \mapsto \text{just } \text{someOtherTerm}) \} \end{aligned}$$

The immediate consequence operator on \perp_{DB} captures the fact that that, due to rules 33 and 34 in combination, p can only be assigned the value b . Any other step, like a step to $p \text{ is } a$ by rule 33, is a dead end. The immediate consequence operator does not, however, preclude the dead-end step from an empty database to a database where $q \text{ is } ff$, since that information is not immediately available: rule 36 only applies when a database where p has a definite value.

5 FIXED-POINT SEMANTICS

We have one notion of what a finite-choice logic program means: the meaning of a program P is the set of databases that are solutions to P according to Definition 2.11. So why do we need to provide another explanation for what a program means?

In our view, there are three satisfying ways of explaining the foundations of a logic programming language. In no particular order:

- A logic program defines a set of derivable propositions in a constructive proof theory. The methodology of *uniform proofs* brings with it a built-in operational semantics in terms of proof search in a focused sequent calculus [Miller et al. 1991].
- A logic program defines a monotonic *immediate consequence* function, and the meaning of a logic program is the least fixed point of this function. This methodology brings with it a built-in operational semantics, because one can repeatedly apply immediate consequence operation to a least element in hopes of reaching a fixed point [Van Gelder et al. 1991].
- A logic program is a proposition in classical logic obtained from a program through the program completion of Clark [1978]. This interpretation admits spurious circular justifications: the program containing one rule $p \leftarrow q$ has the Clark completion $p \leftrightarrow q$, which admits both the expected interpretation (both propositions are false) as well as an undesirable interpretation (both propositions are true). Some programs, however, have a unique *least* model — the smallest set of true propositions — which can be interpreted as the canonical solution.

A very nice property of datalog is that it admits all three justifications, and they all agree. Answer set programming provides a greatly desirable property — the ability to give a program multiple, mutually-exclusive solutions — but it does so at the cost of being able to give programs any of these foundationally satisfying semantics.

In this section, we will present the interpretation of a finite-choice logic program as a member of *Choice*, the least fixed point of a “lifted” immediate consequence operator that has choice sets as both its domain and range. Theorem 5.12 establishes that solutions (Definition 2.11) correspond to databases in this least fixed point that are positive (Definition 4.7) and finite (Definition 4.8).

Example 5.1. First, let’s develop some intuition about what it means for a choice set to provide the interpretation of a finite-choice logic program. In this example only, allow yourself to squint and reinterpret consistent sets of facts as constraint databases: in this hazy light, the set of solutions for a given program looks a lot like a choice set.

The program with no rules corresponds to the choice set $\perp_{\text{Choice}} = \{\emptyset\}$.

The program with one rule ($p \text{ IS } \{tt, ff\} \leftarrow$) has two pairwise-incompatible solutions that form the choice set $\{ \{p \text{ IS } tt\}, \{p \text{ IS } ff\} \}$.

$$\perp_{\text{Choice}} \leq_{\text{Choice}} \{ \{p \text{ IS } tt\}, \{p \text{ IS } ff\} \}$$

One way additional rules can create greater choice sets is by adding facts. A second rule ($q \text{ IS } \{tt\} \leftarrow$) results in a program that still has two solutions.

$$\{ \{p \text{ IS } tt\}, \{p \text{ IS } ff\} \} \leq_{\text{Choice}} \{ \{p \text{ IS } tt, q \text{ IS } tt\}, \{p \text{ IS } ff, q \text{ IS } tt\} \}$$

If we instead added a more constrained second rule ($q \text{ IS } \{tt\} \leftarrow p \text{ IS } tt$), we would instead have these solutions:

$$\{ \{p \text{ IS } tt\}, \{p \text{ IS } ff\} \} \leq_{\text{Choice}} \{ \{p \text{ IS } tt, q \text{ IS } tt\}, \{p \text{ IS } ff\} \}$$

Another way additional rules can create greater choice sets is by *removing solutions*. If we return to the program with just the rule ($p \text{ IS } \{tt, ff\} \leftarrow$) and add a second rule ($p \text{ IS } \{tt\} \leftarrow p \text{ IS } x$), the database where $p \text{ IS } ff$ is no longer a solution:

$$\{ \{p \text{ IS } tt\}, \{p \text{ IS } ff\} \} \leq_{\text{Choice}} \{ \{p \text{ IS } tt\} \}$$

A choice set that is greater according to \leq_{Choice} may alternatively contain *more* constraint databases, which would occur if the second rule was instead ($q \text{ IS } \{\text{tt}, \text{ff}\} \leftarrow p \text{ IS } \text{ff}$):

$$\{ \{p \text{ IS } \text{tt}\}, \{p \text{ IS } \text{ff}\} \} \leq_{\text{Choice}} \{ \{p \text{ IS } \text{tt}\}, \{p \text{ IS } \text{ff}, q \text{ IS } \text{tt}\}, \{p \text{ IS } \text{ff}, q \text{ IS } \text{ff}\} \}$$

Finally, if we have a fully contradictory program, such as one containing two rules ($p \text{ IS } \{\text{tt}, \text{ff}\} \leftarrow$) and ($p \text{ IS } \{\text{meadow}\} \leftarrow p \text{ IS } x$), its interpretation is the set containing zero solutions. This is a valid choice set, and is in fact the greatest element of *Choice*.

$$\{ \{p \text{ IS } \text{tt}\}, \{p \text{ IS } \text{ff}\} \} \leq_{\text{Choice}} \emptyset = \top_{\text{Choice}}$$

Running through this example, one can observe a kind of monotonicity at work: adding rules to a program always results in the meaning of that larger program being a greater choice set according to the partial order \leq_{Choice} from Definition 4.9.

5.1 Lifted immediate consequence

The typical move when defining the meaning of a forward-chaining logic programs is to interpret the program as the least fixed point of an immediate consequence function. Immediate consequence presented in Section 4.5 is a function $\tau_P : DB \rightarrow \text{Choice}$, and because the domain and range are different we can't take the fixed point. Recalling the definition of saturation in Definition 2.10 gives us a limited (but important!) notion of fixed points that we call *models*:

Definition 5.2. A constraint database D is a *model* of the program P in either of these equivalent conditions:

$$D \in \tau_P(D) \iff \tau_P(D) = \{D\}$$

We have no hope of giving a least-fixed-point interpretation of finite-choice logic programs this way, as most interesting finite-choice logic programs do not have unique minimal models. Instead, we will “lift” τ_P to a function $\mathcal{T}_P : \text{Choice} \rightarrow \text{Choice}$:

$$\text{Definition 5.3. } \mathcal{T}_P(C) = \bigcup_{D \in C} \tau_P(D).$$

For Definition 5.3 to be a candidate for iterating to a fixed point, we must show that $\bigcup_{D \in C} \tau_P(D)$ is in *Choice*, i.e. that all the databases it contains are pairwise incompatible:

LEMMA 5.4. *If C is pairwise incompatible, then so is $\bigcup_{D \in C} \tau_P(D)$. That is, if $E_1, E_2 \in \bigcup_{D \in C} \tau_P(D)$ are compatible, then they are equal.*

PROOF. Consider some $D_1, D_2 \in C$ and $E_1 \in \tau_P(D_1)$ and $E_2 \in \tau_P(D_2)$. Suppose $E_1 \parallel E_2$. Since $D_1 \leq E_1$ and $D_2 \leq E_2$, by Lemma 4.2 we know $D_1 \parallel D_2$, thus $D_1 = D_2$ (since both are in C). And since $\tau_P(D_1) = \tau_P(D_2)$ is pairwise incompatible, $E_1 = E_2$. \square

Now we would like to define the meaning of a finite-choice logic program as the unique least fixed point of \mathcal{T}_P . Because choice sets form a complete lattice, we can apply Knaster-Tarski [Tarski 1955] as long as \mathcal{T}_P is monotone, which it is:

$$\text{LEMMA 5.5 (}\tau_P \text{ IS MONOTONE). } \text{If } D_1 \leq_{DB} D_2, \text{ then } \tau_P(D_1) \leq_{\text{Choice}} \tau_P(D_2).$$

PROOF. Because $\tau_P(D_1)$ and $\tau_P(D_2)$ are the least upper bound of the rule heads satisfied in D_1 and D_2 respectively, it suffices to show that if σ satisfies F in D_1 , then σ satisfies F in D_2 . This follows from Definition 4.12: for each premise $p(\bar{t}) \text{ IS } v$ in F , we have $\text{just } \sigma v \leq D_1[p(\sigma\bar{t})] \leq D_2[p(\sigma\bar{t})]$. \square

$$\text{LEMMA 5.6 (}\mathcal{T}_P \text{ IS MONOTONE). } \text{If } C_1 \leq C_2, \text{ then } \mathcal{T}_P(C_1) \leq \mathcal{T}_P(C_2).$$

PROOF. We wish to show $\bigcup_{D \in C} \tau_P(D) \leq_{\text{Choice}} \bigcup_{D' \in C'} \tau_P(D')$. So, fixing $D' \in C'$ and $E' \in \tau_P(D')$, we wish to find a $D \in C$ and $E \in \tau_P(D)$ with $E \leq E'$. Since $C \leq C'$, for $D' \in C'$ there exists a unique $D \in C$ with $D \leq D'$. By monotonicity of τ_P we have $\tau_P(D) \leq_{\text{Choice}} \tau_P(D')$. Thus for $E' \in \tau_P(D')$ we have a unique $E \in \tau_P(D)$ with $E \leq E'$. \square

THEOREM 5.7 (LEAST FIXED POINTS). \mathcal{T}_P has a least fixed point, written as $\text{lfp } \mathcal{T}_P$.

PROOF. Because \mathcal{T}_P is monotone and *Choice* is a complete lattice, by Tarski [1955], the set of fixed points of \mathcal{T}_P forms a complete lattice. The least fixed point is the least element of this lattice, i.e. $\text{lfp } \mathcal{T}_P = \bigwedge \{C : \mathcal{T}_P(C) \leq C\}$ (where $\bigwedge X = \bigvee \{C : \forall x \in X, C \leq x\}$). \square

COROLLARY 5.8. Every model E has a lower bound $D \in \text{lfp } \mathcal{T}_P$ with $D \leq E$.

PROOF. Since $\{E\}$ is a fixed point of \mathcal{T}_P , we know $\text{lfp } \mathcal{T}_P \leq_{\text{Choice}} \{E\}$, i.e. $\exists D \in \text{lfp } \mathcal{T}_P. D \leq E$. \square

THEOREM 5.9 (MINIMAL MODELS). $\text{lfp } \mathcal{T}_P$ contains exactly the minimal models of P , meaning models D such that any model $D' \leq D$ is equal to D .

PROOF. No model outside $\text{lfp } \mathcal{T}_P$ is minimal, because by Corollary 5.8 it has a lower bound in $\text{lfp } \mathcal{T}_P$. And every model $D \in \text{lfp } \mathcal{T}_P$ is minimal: given a model $E \leq D$, by Corollary 5.8, there is some $D' \in \text{lfp } \mathcal{T}_P$ with $D' \leq E \leq D$. But then $D' \parallel D$ and by pairwise incompatibility $D' = D = E$. \square

5.2 Agreement of step semantics and least-fixed-point semantics

Theorem 5.7 establishes that it is reasonable to say that a finite-choice logic program has a canonical model, $\text{lfp } \mathcal{T}_P$. Defining the meaning of a program this way, instead of the step-by-step definition given in Section 2.1, has a number of advantages. However, we need to take care with stating the correspondence between our two ways of assigning meaning to a finite-choice logic program.

Example 5.10. The one-rule program ($p \text{ is? } b \leftarrow$) has one solution according to Definition 2.11, the set $\{p \text{ is } b\}$. The least-fixed-point interpretation of this program contains two models. The first, ($p \mapsto \text{just } b$), obviously corresponds to the unique solution. The second, ($p \mapsto \text{noneOf } \{b\}$), does not correspond to any solution.

This is a relatively straightforward issue to resolve: we will only expect *positive* models (Definition 4.7) to correspond to solutions.

Another challenge for connecting the step semantics and the least-fixed-point interpretation has to do with infinite choice sets and constraint databases.

Example 5.11. The translated answer set program with no finite grounding that we showed in Figure 4 makes for an interesting example. All but one of the models in the least-fixed-point interpretation of this program are finite. Here is one finite model, corresponding to the visiting only 0 and then stopping:

$$(\text{visit}(z) \mapsto \text{just tt}, \text{more}(z) \mapsto \text{just ff}, \text{stop}(z) \mapsto \text{just tt})$$

That constraint database corresponds to the fact set $\{\text{visit}(z) \text{ is tt}, \text{more}(z) \text{ is ff}, \text{stop}(z) \text{ is tt}\}$, which is a solution according to Definition 2.11.

In addition to countably infinite constraint databases that correspond to solutions, the least-fixed-point interpretation of the program in Figure 4 also includes a single infinite constraint database that does not correspond to any solution:

$$\bigvee_{i \in \mathbb{N}} (\text{visit}(s^i(z)) \mapsto \text{just tt}, \text{stop}(s^i(z)) \mapsto \text{just ff}, \text{more}(s^i(z)) \mapsto \text{just tt})$$

Example 5.11 demonstrates an advantage the least-fixed-point interpretation: it lets us reason mathematically about the meaning of programs that cannot be reached in finitely many steps. The least-fixed-point interpretation of datalog confers analogous advantages: a simple forward-chaining interpreter cannot fully evaluate the datalog program with two rules $(p(z) \leftarrow)$ and $(p(s(x)) \leftarrow p(x))$, but the least-fixed-point interpretation assigns the program a canonical interpretation as the infinite set $\{p(s^n(z)) \mid n \in \mathbb{N}\}$.

Because we defined solutions to be finite sets, we have to account for the fact that the least fixed point interpretation of a finite-choice logic program may contain infinite constraint databases that do not correspond to solutions.

The following theorem establishes that there are no other caveats beyond the two highlighted by Examples 5.10 and 5.11:

THEOREM 5.12. *For $D \in DB$, the following are equivalent:*

- (1) D is a solution to P by Definition 2.11.
- (2) $D \in \text{lfp } \mathcal{T}_P$ and D is positive and finite.

PROOF. Details are in the supplemental materials (Appendix C.1), but we will provide an outline here:

Going from (1) to (2) we first establish that any step which is precluded by the immediate consequence operator (as discussed in the introduction to Section 4) is actually a dead end, which means that productive steps taken towards a solution are always suitably “within” the set of immediate consequences. We then observe that all solutions are models, fixed points as defined by Definition 5.2, and are therefore bounded below by some minimal model in $\text{lfp } \mathcal{T}_P$. But any step sequence deriving a solution must be bounded above by a constraint database in $\text{lfp } \mathcal{T}_P$, and this squeezes solutions into membership in $\text{lfp } \mathcal{T}_P$.

Going from (2) to (1) we establish that any finite constraint database in $\text{lfp } \mathcal{T}_P$ belongs to a finite iteration of \mathcal{T}_P , and the action of each of those finitely-many iterations can be simulated with finite step sequences. \square

6 ABSTRACT ALGORITHM

Choice sets are critical both in defining the immediate consequence operators τ_P and \mathcal{T}_P and in defining the meaning of a finite-choice logic program as $\text{lfp } \mathcal{T}_P$. However, the infrastructure of choice sets is more heavyweight than was strictly required to define τ_P . For any program P and constraint database D , we can express $\tau_P(D)$ in a simplified form, as the least upper bound of *singular* choice sets:

Definition 6.1. A choice set C is *singular* to an attribute a when, for all $D \in C$, we have that $D[a'] = \text{noneOf } \emptyset$ whenever $a' \neq a$.

The least upper bound of multiple choice sets singular to a is also a choice set singular to a , so we can view the output of immediate consequence as a *map from attributes to singular choice sets*. This critical shift in perspective is illustrated in Figure 6 and is specified formally in supplemental materials (Appendix C.1).

Note that while $\tau_P(D)$ can always be written as the least upper bound of singular choice sets, not all choice sets can be described this way. A simple example is the program from Figure 1, whose interpretation is given by this choice set:

$$\{ (p \mapsto \text{just ff}, q \mapsto \text{just tt}), (p \mapsto \text{just tt}, q \mapsto \text{just ff}) \}$$

6.1 An algorithm for exploring the interpretation of a finite-choice logic program

This nondeterministic algorithm attempts to return a single element from $\text{lfp } \mathcal{T}_P$.

$$\begin{aligned}
& \{ (p_1 \mapsto \text{just } b) \} \vee & \{ (p_1 \mapsto \text{just } b, p_2 \mapsto \text{just } b, p_3 \mapsto \text{just } b), \\
& \{ (p_2 \mapsto \text{just } b), (p_2 \mapsto \text{just } c) \} \vee & (p_1 \mapsto \text{just } b, p_2 \mapsto \text{just } b, p_3 \mapsto \text{noneOf } \{b\}), \\
& \{ (p_3 \mapsto \text{just } b), (p_3 \mapsto \text{noneOf } \{b\}) \} & (p_1 \mapsto \text{just } b, p_2 \mapsto \text{just } c, p_3 \mapsto \text{just } b), \\
& & (p_1 \mapsto \text{just } b, p_2 \mapsto \text{just } c, p_3 \mapsto \text{noneOf } \{b\}) \}
\end{aligned}$$

Fig. 6. Two views of the same choice set. On the right, the choice set is represented straightforwardly as a set of pairwise-incompatible constraint databases. On the left, the choice set is represented as the least upper bound of singular choice sets.

- Initially, let D be \perp_{DB} .
- While $\tau_P(D) \neq \{D\}$ and $\tau_P(D) \neq \emptyset$:
 - (1) Interpret $\tau_P(D)$ as a map from attributes a to choice sets singular in a .
 - (2) Nondeterministically pick some attribute a where $\tau_P(D)[a] \neq \{(a \mapsto D[a])\}$.
(The loop guard $\tau_P(D) \neq \{D\}$ ensures some such a exists.)
 - (3) Nondeterministically pick some $(a \mapsto c)$ in $\tau_P(D)[a]$.
(The loop guard $\tau_P(D) \neq \emptyset$ ensures that $\tau_P(D)[a] \neq \emptyset$.)
 - (4) Modify D by setting $D[a]$ to be c .
- Successfully return D if $\tau_P(D) = \{D\}$, and return failure if $\tau_P(D) = \emptyset$.

This algorithm captures the intuition discussed at the start of Section 4: we can use the immediate consequences of a database to guide the process of picking the next step in the step-by-step computation of solutions to a finite-choice logic program.

THEOREM 6.2. *For positive constraint databases D , the following are equivalent:*

- (1) *The algorithm described above may successfully return D .*
- (2) *$D \in \text{lfp } \mathcal{T}_P$ and D is finite.*

This is a corollary of the lemmas used to establish Theorem 5.12, which are presented in the supplemental materials (Appendix C.1).

The key to efficient evaluation of this algorithm is that we can incrementally update and efficiently query the portion of $\tau_P(D)$ where $\tau_P(D)[a] \neq \{(a \mapsto D[a])\}$. This is not original to finite-choice logic programming or our implementation: maintaining a data structure capturing all immediate consequences is perhaps *the* fundamental move for efficient implementation of semi-naïve, tuple-at-a-time forward-chaining evaluation of logic programs. To give two examples, this is precisely the purpose of the queue Q in [McAllester 2002] and of the *agenda* A in [Eisner 2023]. Our setting is novel because the immediate consequence of a partial solution is not another deterministically-defined partial solution — as it is for McAllester, Eisner, and all other work we are aware of — but a pairwise-incompatible set of partial solutions, which we explore nondeterministically.

6.2 Resolving nondeterminism

There are two nondeterministic choice points in the abstract algorithm described in Section 6.1: the choice of an attribute a in step 2 of the loop, and the choice from $\tau_P(D)[a]$ in step 3 of the loop. The nondeterministic choice in step 2 is common to all similar semi-naïve, tuple-at-a-time forward-chaining evaluation algorithms, but the nondeterministic choice in step 3 is not.

We claim that the selection of an attribute (in step 2 of the abstract algorithm's loop) isn't "lossy:" when the algorithm might pick one of two attributes, the choice of one won't preclude finding any answer that the program might have returned if it had instead picked the other attribute. This

means that a backtracking version of our algorithm that searches for all solutions doesn't need to reconsider the choices made in step 2.

While each *individual* attribute choice cannot cut off the path to any particular solution, a *systematic* bias in the way attributes are selected can interfere with the algorithm's nondeterministic completeness. Nevertheless, we choose for our implementation to be systematically biased by the *deduce-then-choose* strategy introduced in Section 3.2. In the context of our abstract algorithm, the deduce-then-choose strategy always picks an attribute a where $\tau_P(D)[a]$ is a singleton when such a choice is possible. This strategy does force non-termination on the abstract algorithm where it might otherwise be able to terminate, as the following example shows:

Example 6.3. Consider the following five-rule finite-choice logic program:

$$p \text{ IS } \{\text{red}\} \leftarrow \quad (37)$$

$$q \text{ IS } \{\text{blue, yellow}\} \leftarrow \text{num}(s(s(z))) \quad (38)$$

$$p \text{ IS } \{x\} \leftarrow q \text{ IS } x \quad (39)$$

$$\text{num}(z) \leftarrow \quad (40)$$

$$\text{num}(s(x)) \leftarrow \text{num}(x) \quad (41)$$

This program has no solutions — its least-fixed-point interpretation is the empty set — and it is possible for the abstract algorithm in Section 6.1 to finitely return failure or to loop endlessly, depending on how nondeterministic choices are resolved. However, the additional constraint imposed by the deduce-then-choose strategy means that our algorithm cannot finitely return failure, and can only fail to terminate.

Unlike the choice of which attribute to consider in step 2 of the abstract algorithm, the choice in step 3 to pick a specific value necessarily cuts off possible solutions. However, we can turn our abstract algorithm into a partially correct algorithm for enumerating the elements of a least-fixed-point interpretation by backtracking over the choices made in step 3.

Alongside backtracking, we add in one final bias. While our abstract algorithm enumerates arbitrary members of $\text{lfp } \mathcal{T}_P$, we are interested only in *solutions*, which by Theorem 5.12 are the *positive* constraint databases in $\text{lfp } \mathcal{T}_P$. Accordingly, we restrict our backtracking strategy to prefer constraints of the form (*just t*) prior to backtracking. This means that if the algorithm returns without backtracking, the result will be a positive model.

This combination of strategies—the deduce-then-choose strategy, no backtracking over choices of attribute, and some form of backtracking over the choice of constraint while preferring positive constraints—means that we can think of an execution of the abstract algorithm as a process of expanding a (potentially infinite) tree of decisions.

Example 6.4. Consider the following five-rule finite-choice logic program:

$$p \text{ IS } \{\text{tt, ff}\} \leftarrow \quad (42)$$

$$q \text{ IS } \{\text{tt, ff}\} \leftarrow \quad (43)$$

$$r \text{ IS? } a \leftarrow \quad (44)$$

$$r \text{ IS } \{b, c\} \leftarrow p \text{ IS } \text{ff} \quad (45)$$

$$r \text{ IS } \{x\} \leftarrow p \text{ IS } x, q \text{ IS } x \quad (46)$$

Two of the many possible induced execution trees are shown in Figure 7. The upper tree Figure 7 represents an execution where p is considered first in step 2 of the abstract algorithm's loop, and the lower tree represents an execution where r is considered first.




Fig. 7. Two trees of nondeterministic execution induced by different ways of selecting attributes in course of executing the finite-choice logic program from Example 6.4. Locations representing failure-return are crossed out, and locations representing successful terms are highlighted.

The algorithm potentially returns the same five models regardless of the order in which attributes are selected. The related models are outlined in the same color in the two trees. For example, the paths that corresponds to returning the model $(p \mapsto \text{just tt}, q \mapsto \text{just tt}, r \mapsto \text{noneOf}\{a\})$ is highlighted in brown and can be found on the left side of the upper tree and near the center of the lower tree.

In the first example where p is considered first, when the algorithm decides to give p the value **just tt**, the next attribute considered is q , but when p is given the value **just ff**, the next attribute considered is r . This highlights that we are not insisting on a single global ordering of attributes.

6.3 Completeness of the constrained algorithm

Example 6.3 demonstrated that the restrictions we put on the nondeterministic algorithm in this section force the non-backtracking version of our algorithm into non-termination where it might have otherwise terminated signaling failure. It is an open question whether the restrictions in this section interfere with the nondeterministic completeness of the algorithm: in other words, we don't know whether there is a program P and a database $D \in \text{lfp } \mathcal{T}_P$ such that the more restricted algorithm is unable return D . We conjecture that, even with the refinements we have discussed, any finite $D \in \text{lfp } \mathcal{T}_P$ can be reached with a suitable choice of attributes in step 2 of the algorithm's loop.

Even if these restrictions do sacrifice completeness, we believe they are worthwhile. The deduce-then-choose strategy makes the performance of finite-choice logic programs vastly more predictable, as demonstrated by the examples in Section 3.

7 IMPLEMENTATION

Our implementation of the constrained abstract algorithm for finite-choice logic programming is called DUSA. The implementation supports three modes of interaction: a TypeScript API, a command-line program, and a browser-based editor (pictured in Figure 8). All modes allow the




Fig. 8. The online editor for the DUSA implementation of finite-choice logic programming.

client to request incremental enumeration of all solutions to a finite-choice logic program, and the first two modes allow the client to sample individual solutions.

Despite implementing the algorithm in Section 6.1, which enumerates finite models in $\text{lfp } \mathcal{T}_P$, our implementation discards all non-positive models and only exposes solutions—the positive models in the program’s interpretation—to the client. That’s why, while the interpretation of the program in Example 6.4 includes five different models, the web interface in Figure 8 only indicates the presence of four solutions.

Our implementation of finite-choice logic programming was initially designed with procedural generation and possibility-space exploration in mind. This makes our design considerations very similar to those that led Horswill to design CatSAT as a way to facilitate procedural generation in Unity games [Horswill 2018]. Horswill identifies four barriers to the use of answer set programming in video games, all of which are relevant to finite-choice logic programming and the DUSA implementation:

Designer transparency. By this, Horswill refers to tools that allow “non-programmers” to understand and manipulate programs. CatSAT does not directly address this barrier, and while DUSA has not been systematically tested with designers, our web interface does make it possible for non-experts to inspect, manipulate, and share programs without downloading or installing software (meeting many of the criteria for *casual creators* [Compton and Mateas 2015]).

Performance. CatSAT uses a simple implementation of the WalkSAT algorithm to get reasonable time performance and excellent memory performance on satisfiability problems that are not excessively constrained. DUSA approaches performance from a different angle: our implementation is relatively memory-intensive, but we aim to provide *predictable* performance, which we discuss more in Section 7.1.

Run-time integration. Whereas answer set programming tools are primarily written as standalone programs, CatSAT is implemented as a DSL in C# to naturally integrate with programs in the Unity ecosystem. DUSA is implemented in TypeScript to facilitate integration with the web ecosystem.

Determinism. Horswill observes that most SAT solvers and answer set programming languages make it difficult or impossible to access suitably random behavior. Like CatSAT, the DUSA implementation defaults to randomness. After constraining the abstract algorithm as described in Section 6.2, our implementation resolves all remaining nondeterministic choices uniformly at random. Furthermore, rather than using a stack-based backtracking algorithm, our implementation creates in-memory representations of the trees described in Figure 7. After a solution is discovered, the algorithm prunes fully-explored parts of the tree and then returns to the root of the tree to begin random exploration again. The goal of this strategy is to avoid returning a second solution




Fig. 9. Performance of the spanning tree program from Section 3.1 and the canonical representative algorithm from Section 3.2 on graphs with different characteristics. All data points are the median of three runs requesting a single solution.

that is very similar to the first, a behavior that is commonly observed when exploring a state space with a stack-based (or depth-first) backtracking search.

While our strategy for randomized exploration does not truly sample the interpretation at random, in practice it has proven suitable for procedural generation. This strategy also does an empirically good job of locating all the “small” solutions for programs with an infinite interpretation, like the one in Figure 4, and will eventually return any specific finite solution for that particular program with probability 1. The Alpha implementation of ASP with lazy grounding [Weinzierl 2017] exhibits a different behavior that is wholly unsuitable for procedural generation: when we ask Alpha to enumerate solutions to the answer set program in Figure 4, we observe that it skips most small models and successively returns larger and larger models with large gaps in between.

7.1 Predictable performance with prefix firings

The implementation of the abstract evaluation algorithm in DUSA is closely modeled after the semi-naive, tuple-at-a-time forward-chaining interpreter for datalog presented by McAllester [2002]. McAllester’s algorithm was designed to facilitate a *cost semantics*, a way of reasoning about the run-time cost of evaluation without reasoning about the details of how the implementation works.

McAllester’s cost semantics is based on counting the number of *prefix firings* in a solution. If a rule $H \leftarrow F$ has the premises $F = p_1(\bar{t}_1) \text{ is } v_1, \dots, p_n(\bar{t}_n) \text{ is } v_n$, then a prefix firing for the database D is a unique variable-free instantiation of the first $i \leq n$ premises $p_1(\sigma\bar{t}_1) \text{ is } v_1, \dots, p_i(\sigma\bar{t}_i) \text{ is } v_i$ that is satisfied by the D . The central result of McAllester’s work is that, assuming constant-time hashtable operations, a datalog program that returns the solution D_{final} can be evaluated in time proportional to the prefix firings for D_{final} . This cost semantics allows programmers to make decisions about how to write declarative programs: the McAllester cost semantics suggests defining the recursive rule for a transitive closure as $path(x, z) \leftarrow edge(x, y), path(y, z)$ rather than $path(x, z) \leftarrow path(x, y), path(y, z)$, because in a graph where the number of edges is proportional to n , the number of vertices, the former rule has $O(n^2)$ prefix firings and the latter has $O(n^3)$.

We do not present a theorem about cost semantics in this development, in part because it requires careful introduction of a number of extra concepts and details, and in part because our implementation is based on functional maps instead of hashtables in order to support the very general backtracking methodology described above; this means that a variant of McAllester’s result will not apply directly to our implementation. Nevertheless, the cost semantics has *predictive* power. The programs for spanning tree generation in Figure 2 and for canonical representative appointment in Figure 3 are expected, in any deduce-then-choose execution, to find some solution without backtracking, and solutions for both programs will have prefix firings proportional to the




Fig. 10. Performance of Alpha, Clingo, and Dusa at finding canonical representatives. All data points are the median of three runs requesting a single solution, discarding runs exceeding 100 seconds.

number of edges. Figure 9 demonstrates that the actual behavior of DUSA generally conforms to the running time predicted by the cost semantics.

7.2 Comparing Dusa to answer set programming

There’s not an obvious way to ask an answer set programming engine to run the spanning tree program from Section 3.1 or the canonical representative program from Section 3.2, because we have not defined a translation from finite-choice logic programs to answer set programs. However, solutions to the rooted spanning tree problem and the canonical representative problem can be expressed in answer set programming, and these answer set programs can be translated into finite-choice logic programs as discussed in Sections 2.3 and 3.3.

In Figure 10 we compare the canonical representative problem from Figure 3 against an idiomatic Clingo program that uses features like cardinality constraints, as well as comparing the performance of a single “pure” answer set program executed in Clingo, in the Alpha implementation of ASP with lazy grounding [Weinzierl 2017], and in DUSA by direct translation of ASP to finite-choice logic programming. (The yellow lines in Figure 10 repeat the red and blue lines from the right-hand side of Figure 9.) In contrast with the predictable behavior shown in Figure 9, the performance of the other four solutions varies significantly depending on characteristics of the graph. In the head-to-head comparison, DUSA outperforms the state of the art on dense graphs but does worse as graphs get sparser.

The poor performance of Clingo in the head-to-head comparison can be attributed to the “grounding bottleneck” encountered by traditional approaches to ASP. Both Alpha’s lazy grounding and DUSA’s deduce-then-choose execution model represent ways of avoiding this grounding bottleneck, and this translates to better asymptotic performance for this example. DUSA’s relatively poor performance on sparse graphs in the head-to-head comparison is almost certainly attributable to our implementation’s very naive backtracking strategy, a level of sophistication comparable to 1990’s-era DPLL SAT solvers that lacked non-chronological backtracking.

An extended set of benchmarking results, including the programs used in this comparison, is available in supplemental materials (Appendix D).

8 RELATED WORK

Our least-fixed-point interpretation of finite-choice logic programming describes a domain-like structure for nondeterminism that draws inspiration from domain theory writ large [Scott 1982], particularly powerdomains for nondeterministic lambda calculi and imperative programs [Kenaway and Hoare 1980; Plotkin 1976; Smyth 1976]. Our partial order on choice sets matches the one used for Smyth powerdomains in particular [Smyth 1976]. Our construction requires certain

completeness criteria on posets that differ from these and other domain definitions we have found in the literature.

Both the least-fixed-point interpretation and the step-sequence interpretation of finite-choice logic programming represent ways of reasoning about the incremental development of solutions. Because we provide a translation of answer set programming into finite-choice logic programming, these ways of reasoning about incremental development of solutions are applicable to ASP as well. Previous work in this area includes the *stable backtracking fixedpoint* algorithm described by Sacca and Zaniolo [1990]. (We interpret their algorithm as giving the first direct account for answer set programming without grounding, though it seems to us that this significant fact was not noticed or exploited by the authors or anyone else.) There has also been a disjunctive extension to datalog [Eiter et al. 1997] and characterizations of its stable models [Przymusiński 1991]; Leone et al. [1997] present an algorithm for incrementally deriving stable models for disjunctive datalog. Their approach is based on initially creating a single canonical model by letting some propositions be “partially true.”

At the core of finite-choice logic programming is the propositional form $p(\bar{t}) \text{ is } v$, which establishes a *functional dependency* from an attribute to a value. This mirrors the proposition form $R[x_1, \dots, x_{n-1}] = x_n$ used to enforce integrity constraints in LogicQL [Aref et al. 2015]. Soufflé’s functional dependencies allow programs to express multiple mutually contradictory conclusions [Hu et al. 2021], but as discussed in Section 3.1, functional dependencies in the tradition of Krishnamurthy and Naqvi [1988] cannot express that the violation of a functional dependency should cause a potential solution to be rejected. Roughly speaking, LogicQL works like a finite-choice logic programming language that only has closed rules without multiple choices, and Soufflé works like a finite-choice logic programming language that has only open rules. These systems are less expressive than ours, but they can completely avoid expensive backtracking.

The partially-ordered set *Constraint* introduced in Section 4.2 suggests a relationship between finite-choice logic programming and other languages that assign non-Boolean values to propositions: values drawn from a semiring in weighted logic programming [Eisner et al. 2005], various interpretations of degree-of-truth in annotated logic programming [Kifer and Subrahmanian 1992] and bilattice-based logic programming [Fitting 1991], and probabilities in probabilistic logic programming [Sterling 1995]. As discussed in Section 4.3, we’re unaware of work along these lines that can give an account for stable models without falling back on a variant of Gelfond and Lifschitz’s syntactic transformation. However, the related mathematical structures shared across these approaches points towards generalizing *Constraint* to accept alternate partial orders, similar to the lattice-based functional dependencies seen in datalog extensions like Bloom [Conway et al. 2012], Flix [Madsen and Lhoták 2020], and Egglog [Zhang et al. 2023], which could expand the expressiveness of finite-choice logic programming to account for monotonic data aggregation [Ross and Sagiv 1992].

Fandinno et al. [2023] extend a translation due to Niemlā [2008] to formalize the “idea from folklore” that ASP can be based on an ASP choice rules along with integrity constraints. This is similar to our observation that choice, rather than negation, is a suitable foundation for answer set programming, though a limited (stratified) negation is still necessary in their translation. They present a conceptual operational semantics based on making *all* possible choices up-front, followed by a phase of deduction; this operational interpretation is intended as a pedagogical tool rather than the basis of a practical implementation.

9 CONCLUSION

We have introduced the theory and implementation of finite-choice logic programming, an approach to logic programming that has connections to answer set programming but that treats *choice*, not

negation, as the fundamental primitive of nondeterminism. We establish that choice is a suitably expressive foundation by showing that answer set programming can be defined in terms of finite-choice logic programming.

We give a first definition of finite-choice logic programming in terms of nondeterministically augmenting a set of facts, a second definition as the unique least fixed point in a novel domain of mutually-exclusive models, and a proof that these two definitions agree on solutions.

Our DUSA implementation can enumerate solutions to finite-choice logic programs, and the runtime behavior of our implementation can often reliably (if approximately) be predicted by McAllester’s cost semantics based on prefix firings. Because our implementation is not subject to the “grounding bottleneck” that affects mainstream answer set programming solvers, our implementation outperforms the state-of-the-art Clingo ASP implementation in many cases despite its extremely naive backtracking strategy.

ACKNOWLEDGMENTS

The genesis of this work was performed at the Recurse Center, and we are grateful for the Recurse Center for nurturing a space where people can think about ideas in programming.

We want to thank Richard Comploi-Taupe, who helped us understand the state of research on answer set programming with lazy grounding. Many folks gave helpful feedback that helped us improve this draft, particularly Mitch Wand, Emma Tosch, Jason Reed, Philip Zucker, David Renshaw, and three anonymous reviewers.

REFERENCES

- Peter Alvaro, William R. Marczak, Neil Conway, Joseph M. Hellerstein, David Maier, and Russell Sears. 2011. Dedalus: Datalog in time and space. In *Datalog Reloaded*, Oege de Moor, Georg Gottlob, Tim Furche, and Andrew Sellers (Eds.). Springer, Oxford, UK, 262–281. https://doi.org/10.1007/978-3-642-24206-9_16
- Molham Aref, Balder ten Cate, Todd J. Green, Benny Kimelfeld, Dan Olteanu, Emir Pasalic, Todd L. Veldhuizen, and Geoffrey Washburn. 2015. Design and Implementation of the LogicBlox System. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data* (Melbourne, Victoria, Australia) (SIGMOD ’15). Association for Computing Machinery, New York, NY, USA, 1371–1382. <https://doi.org/10.1145/2723372.2742796>
- Koen Claessen and John Hughes. 2011. QuickCheck: a lightweight tool for random testing of Haskell programs. *ACM SIGPLAN Notices* 46, 4 (2011), 53–64. <https://doi.org/10.1145/1988042.1988046>
- Keith L. Clark. 1978. Negation as Failure. In *Logic and Data Bases*, Hervé Gallaire and Jack Minker (Eds.). Springer US, Boston, MA, 293–322. https://doi.org/10.1007/978-1-4684-3384-5_11
- Richard Comploi-Taupe, Gerhard Friedrich, Konstantin Schekotihin, and Antonius Weinzierl. 2023. Domain-Specific Heuristics in Answer Set Programming: A Declarative Non-Monotonic Approach. *J. Artif. Int. Res.* 76 (may 2023), 56 pages. <https://doi.org/10.1613/jair.1.14091>
- Kate Compton and Michael Mateas. 2015. Casual Creators. In *International Conference on Computational Creativity*. 228–235.
- Neil Conway, William R Marczak, Peter Alvaro, Joseph M Hellerstein, and David Maier. 2012. Logic and lattices for distributed programming. In *Proceedings of the Third ACM Symposium on Cloud Computing*. 1–14.
- Krzysztof Czarnecki, Simon Helsen, and Ulrich Eisenecker. 2004. Staged Configuration Using Feature Models. In *Software Product Lines*, Robert L. Nord (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 266–283.
- Krzysztof Czarnecki, Kasper Østerbye, and Markus Völter. 2002. Generative programming. In *European Conference on Object-Oriented Programming*. Springer, 15–29.
- Chinmaya Dabral and Chris Martens. 2020. Generating explorable narrative spaces with answer set programming. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, Vol. 16. AAAI Press, Washington, USA, 45–51. <https://doi.org/10.1609/aiide.v16i1.7406>
- Chinmaya Dabral, Emma Tosch, and Chris Martens. 2023. Exploring Consequences of Privacy Policies with Narrative Generation via Answer Set Programming. (2023). arXiv:2212.06719 Presented at Workshop on Programming Languages and the Law (ProLaLa@POPL).
- Alessandro Dal Palù, Agostino Dovier, Enrico Pontelli, and Gianfranco Rossi. 2009. GASP: Answer Set Programming with Lazy Grounding. *Fundam. Inf.* 96, 3 (2009), 297–322.
- Joris Dormans and Sander Bakkes. 2011. Generating missions and spaces for adaptable play experiences. *IEEE Transactions on Computational Intelligence and AI in Games* 3, 3 (2011), 216–228. <https://doi.org/10.1109/TCIAIG.2011.2149523>

- Jason Eisner. 2023. Time-and-Space-Efficient Weighted Deduction. *Transactions of the Association for Computational Linguistics* 11 (08 2023), 960–973. https://doi.org/10.1162/tacl_a_00588
- Jason Eisner, Eric Goldlust, and Noah A. Smith. 2005. Compiling Comp Ling: Weighted Dynamic Programming and the Dyna Language. In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, Raymond Mooney, Chris Brew, Lee-Feng Chien, and Katrin Kirchhoff (Eds.). Association for Computational Linguistics, Vancouver, British Columbia, Canada, 281–290. <https://aclanthology.org/H05-1036>
- Thomas Eiter, Georg Gottlob, and Heikki Mannila. 1997. Disjunctive datalog. *ACM Transactions on Database Systems (TODS)* 22, 3 (1997), 364–418. <https://doi.org/10.1145/261124.261126>
- Jorge Fandinno, Seemran Mishra, Javier Romero, and Torsten Schaub. 2023. *Answer Set Programming Made Easy*. Springer, Cham, 133–150. https://doi.org/10.1007/978-3-031-31476-6_7
- Melvin Fitting. 1991. Bilattices and the semantics of logic programming. *The Journal of Logic Programming* 11, 2 (1991), 91–116. [https://doi.org/10.1016/0743-1066\(91\)90014-G](https://doi.org/10.1016/0743-1066(91)90014-G)
- Melvin Fitting. 1993. The family of stable models. *The Journal of Logic Programming* 17, 2 (1993), 197–225. [https://doi.org/10.1016/0743-1066\(93\)90031-B](https://doi.org/10.1016/0743-1066(93)90031-B) Special Issue: Non-Monotonic Reasoning and Logic Programming.
- Martin Gebser, Roland Kaminski, Benjamin Kaufmann, and Torsten Schaub. 2019. Multi-shot ASP solving with clingo. *Theory and Practice of Logic Programming* 19, 1 (2019), 27–82. <https://doi.org/10.1017/S1471068418000054>
- Martin Gebser, Benjamin Kaufmann, Roland Kaminski, Max Ostrowski, Torsten Schaub, and Marius Schneider. 2011. Potassco: The Potsdam answer set solving collection. *AI Communications* 24, 2 (2011), 107–124.
- Michael Gelfond and Vladimir Lifschitz. 1988. The Stable Model Semantics for Logic Programming. In *Proceedings of International Logic Programming Conference and Symposium*, Robert Kowalski and Kenneth A. Bowen (Eds.). MIT Press, 1070–1080.
- Fosca Giannotti, Dino Pedreschi, and Carlo Zaniolo. 2001. Semantics and Expressive Power of Nondeterministic Constructs in Deductive Databases. *J. Comput. System Sci.* 62, 1 (2001), 15–42. <https://doi.org/10.1006/jcss.1999.1699>
- Harrison Goldstein, Samantha Frohlich, Meng Wang, and Benjamin C Pierce. 2023. Reflecting on Random Generation. *Proceedings of the ACM on Programming Languages* 7, ICFP (2023), 322–355. <https://doi.org/10.1145/3607842>
- Harrison Goldstein and Benjamin C Pierce. 2022. Parsing randomness. *Proceedings of the ACM on Programming Languages* 6, OOPSLA2 (2022), 89–113. <https://doi.org/10.1145/3563291>
- Sergio Greco and Carlo Zaniolo. 2001. Greedy algorithms in Datalog. *Theory and Practice of Logic Programming* 1, 4 (2001), 381–407. <https://doi.org/10.1017/S1471068401001090>
- Ian Horswill. 2018. CatSAT: A Practical, Embedded, SAT Language for Runtime PCG. *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment* 14, 1 (Sep. 2018), 38–44. <https://doi.org/10.1609/aiide.v14i1.13026>
- Xiaowen Hu, Joshua Karp, David Zhao, Abdul Zreika, Xi Wu, and Bernhard Scholz. 2021. The Choice Construct in the Soufflé Language. In *Programming Languages and Systems*, Hakjoo Oh (Ed.). Springer International Publishing, Cham, 163–181.
- J.R. Kennaway and C.A.R. Hoare. 1980. A theory of nondeterminism. In *Automata, Languages and Programming*, Jaco de Bakker and Jan van Leeuwen (Eds.). Springer, Berlin, Heidelberg, 338–350. https://doi.org/10.1007/3-540-10003-2_82
- Michael Kifer and V.S. Subrahmanian. 1992. Theory of generalized annotated logic programming and its applications. *The Journal of Logic Programming* 12, 4 (1992), 335–367. [https://doi.org/10.1016/0743-1066\(92\)90007-P](https://doi.org/10.1016/0743-1066(92)90007-P)
- Ekaterina Komendantskaya and Anthony Karel Seda. 2009. Sound and Complete SLD-Resolution for Bilattice-Based Annotated Logic Programs. *Electronic Notes in Theoretical Computer Science* 225 (2009), 141–159. <https://doi.org/10.1016/j.entcs.2008.12.071> Proceedings of the Irish Conference on the Mathematical Foundations of Computer Science and Information Technology (MFCSIT 2006).
- Ravi Krishnamurthy and Shamim Naqvi. 1988. Non-Deterministic Choice in Datalog. In *Proceedings of the Third International Conference on Data and Knowledge Bases*, C. BEERI, J.W. SCHMIDT, and U. DAYAL (Eds.). Morgan Kaufmann, 416–424. <https://doi.org/10.1016/B978-1-4832-1313-2.50038-X>
- Leonidas Lampropoulos, Zoe Paraskevopoulou, and Benjamin C Pierce. 2017. Generating good generators for inductive relations. *Proceedings of the ACM on Programming Languages* 2, POPL (2017), 1–30. <https://doi.org/10.1145/3158133>
- Duc Le, Eric Walkingshaw, and Martin Erwig. 2011. #ifdef confirmed harmful: Promoting understandable software variation. In *2011 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. IEEE, 143–150.
- Nicola Leone, Pasquale Rullo, and Francesco Scarcello. 1997. Disjunctive stable models: Unfounded sets, fixpoint semantics, and computation. *Information and computation* 135, 2 (1997), 69–112. <https://doi.org/10.1006/inco.1997.2630>
- Beidi Li, Jochen Teizer, and Carl Schultz. 2020. Non-monotonic Spatial Reasoning for Safety Analysis in Construction. In *Proceedings of the 22nd International Symposium on Principles and Practice of Declarative Programming (Bologna, Italy) (PPDP '20)*. Association for Computing Machinery, New York, NY, USA, Article 16, 12 pages. <https://doi.org/10.1145/3414080.3414096>

- How Khang Lim, Avishkar Mahajar, Martin Strecker, and Meng Weng Wong. 2022. Automating defeasible reasoning in law with answer set programming. *CEUR* (2022).
- Magnus Madsen and Ondřej Lhoták. 2020. Fixpoints for the masses: programming with first-class datalog constraints. *Proceedings of the ACM on Programming Languages* 4, OOPSLA (2020), 1–28.
- David McAllester. 2002. On the complexity analysis of static analyses. *J. ACM* 49, 4 (jul 2002), 512–537. <https://doi.org/10.1145/581771.581774>
- Dale Miller, Gopalan Nadathur, Frank Pfenning, and Andre Scedrov. 1991. Uniform proofs as a foundation for logic programming. *Annals of Pure and Applied Logic* 51, 1 (1991), 125–157. [https://doi.org/10.1016/0168-0072\(91\)90068-W](https://doi.org/10.1016/0168-0072(91)90068-W)
- Xenija Neufeld, Sanaz Mostaghim, and Diego Perez-Liebana. 2015. Procedural level generation with answer set programming for general video game playing. In *2015 7th Computer Science and Electronic Engineering Conference (CEEC)*. IEEE, 207–212. <https://doi.org/10.1109/CEEC.2015.7332726>
- Ilkka Niemelä. 2008. Answer Set Programming without Unstratified Negation. In *Logic Programming*, Maria Garcia de la Banda and Enrico Pontelli (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 88–92. https://doi.org/10.1007/978-3-540-89982-2_15
- Zoe Paraskevopoulou, Aaron Eline, and Leonidas Lampropoulos. 2022. Computing correctly with inductive relations. In *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation*. ACM, NY, USA, 966–980. <https://doi.org/10.1145/3519939.3523707>
- Gordon D. Plotkin. 1976. A powerdomain construction. *SIAM J. Comput.* 5, 3 (1976), 452–487. <https://doi.org/10.1137/0205035>
- Teodor C Przymusiński. 1991. Stable semantics for disjunctive programs. *New generation computing* 9, 3-4 (1991), 401–424. <https://doi.org/10.1007/BF03037171>
- Kenneth A. Ross and Yehoshua Sagiv. 1992. Monotonic aggregation in deductive databases. In *Proceedings of the Eleventh ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems* (San Diego, California, USA) (PODS '92). Association for Computing Machinery, New York, NY, USA, 114–126. <https://doi.org/10.1145/137097.137852>
- Domenico Sacca and Carlo Zaniolo. 1990. Stable models and non-determinism in logic programs with negation. In *Proceedings of the Ninth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems* (Nashville, USA) (PODS '90). Association for Computing Machinery, New York, NY, USA, 205–217. <https://doi.org/10.1145/298514.298572>
- Dana S Scott. 1982. Domains for denotational semantics. In *Automata, Languages and Programming: Ninth Colloquium Aarhus, Denmark, July 12–16, 1982* 9. Springer, Berlin, Heidelberg, 577–610. <https://doi.org/10.1007/BFb0012801>
- Eric L Seidel, Niki Vazou, and Ranjit Jhala. 2015. Type targeted testing. In *Programming Languages and Systems: 24th European Symposium on Programming, ESOP 2015, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2015, London, UK, April 11–18, 2015, Proceedings* 24. Springer, Berlin, Heidelberg, 812–836. https://doi.org/10.1007/978-3-662-46669-8_33
- Noor Shaker, Julian Togelius, and Mark J Nelson. 2016. *Procedural content generation in games*. Springer. <https://doi.org/10.1007/978-3-319-42716-4>
- Tanya Short and Tarn Adams. 2017. *Procedural generation in game design*. CRC Press.
- Robert J. Simmons and Frank Pfenning. 2008. Linear Logical Algorithms. In *Automata, Languages and Programming*, Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfssdóttir, and Igor Walukiewicz (Eds.). Springer, Berlin, Heidelberg, 336–347. https://doi.org/10.1007/978-3-540-70583-3_28
- Patrik Simons, Ilkka Niemelä, and Timo Soininen. 2002. Extending and implementing the stable model semantics. *Artificial Intelligence* 138, 1 (2002), 181–234. [https://doi.org/10.1016/S0004-3702\(02\)00187-X](https://doi.org/10.1016/S0004-3702(02)00187-X) Knowledge Representation and Logic Programming.
- Anthony J Smith and Joanna J Bryson. 2014. A logical approach to building dungeons: Answer set programming for hierarchical procedural content generation in roguelike games. In *Proceedings of the 50th Anniversary Convention of the AISB*.
- Adam M Smith, Eric Butler, and Zoran Popovic. 2013. Quantifying over play: Constraining undesirable solutions in puzzle design. In *Foundations of Digital Games*. 221–228.
- Adam M Smith and Michael Mateas. 2011. Answer set programming for procedural content generation: A design space approach. *IEEE Transactions on Computational Intelligence and AI in Games* 3, 3 (2011), 187–200. <https://doi.org/10.1109/TCIAIG.2011.2158545>
- Michael B Smyth. 1976. Powerdomains. In *International Symposium on Mathematical Foundations of Computer Science*. Springer, Berlin, Heidelberg, 537–543. https://doi.org/10.1007/3-540-07854-1_226
- Leon S. Sterling. 1995. *A Statistical Learning Method for Logic Programs with Distribution Semantics*. MIT Press, 715–729. <https://doi.org/10.7551/mitpress/4298.003.0069>
- Adam Summerville, Chris Martens, Ben Samuel, Joseph Osborn, Noah Wardrip-Fruin, and Michael Mateas. 2018. Gemini: Bidirectional generation and analysis of games via ASP. *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment* 14, 1 (2018), 123–129. <https://doi.org/10.1609/aiide.v14i1.13013>
- Alfred Tarski. 1955. A lattice-theoretical fixpoint theorem and its applications. *Pacific J. Math* 5, 2 (1955), 285–309.

Allen Van Gelder, Kenneth A. Ross, and John S. Schlipf. 1991. The Well-Founded Semantics for General Logic Programs. *J. ACM* 38, 3 (jul 1991), 619–649. <https://doi.org/10.1145/116825.116838>

Antonius Weinzierl. 2017. Blending Lazy-Grounding and CDNL Search for Answer-Set Solving. In *Logic Programming and Nonmonotonic Reasoning*, Marcello Balduccini and Tomi Janhunen (Eds.). Springer International Publishing, Cham, 191–204. https://doi.org/10.1007/978-3-319-61660-5_17

Yihong Zhang, Yisu Remy Wang, Oliver Flatt, David Cao, Philip Zucker, Eli Rosenthal, Zachary Tatlock, and Max Willsey. 2023. Better together: Unifying datalog and equality saturation. *Proceedings of the ACM on Programming Languages* 7, PLDI (2023), 468–492.

A DATALOG AND ANSWER SET PROGRAMMING CONNECTIONS

This section includes additional details from Sections 2.2 and 2.3 in the main body of the paper.

A.1 Simulating datalog

A datalog program is a finite collection of rules of this form:

$$p(\bar{t}) \leftarrow p_1(\bar{t}_1), \dots, p_n(\bar{t}_n) \quad (\text{datalog rule})$$

The model of a datalog program is traditionally understood as the least fixed point of the (datalog) immediate consequence operator, which takes a set X of datalog facts and returns the set of rule conclusions $p(\sigma\bar{t})$ such that all that rule's premises $p_i(\sigma\bar{t}_i) \in X$. These facts can be given a total order:

LEMMA A.1. *If the model of a datalog program is finite, it is possible to order the elements of that model in a sequence $p_1(\bar{t}_1) \dots p_n(\bar{t}_n)$, such that $p_i(\bar{t}_i)$ is always an immediate consequence of $p_j(\bar{t}_j)$ for $j < i$.*

PROOF. If a datalog model is finite, it can be derived by a finite number of applications of the (datalog) immediate consequence operator, where each application adds a finite number of additional facts. We construct the sequence by placing all the facts derived with the first application of the immediate consequence operator first, followed by any additional facts derived with the second application of the immediate consequence operator, and so on. \square

We translate datalog rules of the form shown above to the finite-choice logic programming rules that look like this:

$$p(\bar{t}) \text{ IS \{unit\} } \leftarrow p_1(\bar{t}_1) \text{ IS unit}, \dots, p_n(\bar{t}_n) \text{ IS unit}$$

The paper glosses over the statement of correctness for this translation, but for completeness we make it explicit here:

THEOREM A.2. *Let P be a datalog program with a finite model, and let $\langle P \rangle$ be the interpretation of P as a finite-choice logic program as above. Then there is a unique solution D to $\langle P \rangle$, and the model of P is $\{p(\bar{t}) \mid p(\bar{t}) \text{ IS unit} \in D\}$.*

PROOF. By Lemma A.1 we can get some sequence $X_{\text{seq}} = p_1(\bar{t}_1) \dots p_n(\bar{t}_n)$ containing exactly the elements in the model of P , where $p_i(\bar{t}_i)$ is always an immediate consequence of the set of facts that precede it in the sequence. By induction on k we can derive a step sequence from \emptyset to $\{p_i(\bar{t}_i) \text{ IS unit} \mid 1 \leq i \leq k \wedge p_i(\bar{t}_i) \in X_{\text{seq}}\}$ for any $1 \leq k \leq n$. We will let $D = \{p_i(\bar{t}_i) \text{ IS unit} \mid p_i(\bar{t}_i) \in X_{\text{seq}}\}$. There is a step sequence $\emptyset \dots D$, so if D is saturated, it is a solution. Assume some additional fact $p(\bar{t}) \text{ IS unit}$ can be derived in the translated program, a corresponding additional fact $p(\bar{t})$ can be derived in the model of P that isn't present in X_{seq} , a contradiction.

It remains to be shown that D is the unique solution. Consider an arbitrary solution D' . First, by induction on the step sequence that reaches D' , if $p(\bar{t}) \text{ IS } v \in D'$ it must be the case that $v = \text{unit}$. If D' was missing any elements in D then it wouldn't be a solution, and if D was missing any elements in D' we could use that to show that X_{seq} was missing some element in P 's unique model, so $D = D'$. \square

A.2 Simulating answer set programming

Answer set programs are defined in terms of the *stable model semantics* [Gelfond and Lifschitz 1988]. A rule in answer set programming has both non-negated (positive) premises, which we'll

write as p_i , and negated (negative) premises, which we'll write as q_i .

$$p \leftarrow p_1, \dots, p_n, \neg q_1, \dots, \neg q_m \quad (\text{ASP rule})$$

Let X denote a finite set of ground predicates. When modeling ASP, any predicate in the set is treated as true, and any predicate not in the set is false. To explain what a *stable* model is, the ASP literature [Gelfond and Lifschitz 1988] first defines P^X , the *reduct* of a program P over X obtained by:

- (1) removing any rules where a negated premise appears in X
- (2) removing all the negated premises from rules that weren't removed in step 1

The reduct P^X is always a regular datalog program with a unique finite model. If the model of P^X is exactly X , then X is a stable model for the ASP program P .

The translation of a single ASP rule with m negated premises of the form above produces $m + 1$ rules in the resulting finite-choice logic program. These introduced rules use two new program-wide introduced constants, **tt** and **ff**, which represent the answer set program's assignment of truth or falsehood, respectively, to the relevant attribute:

$$\begin{aligned} q_1 \text{ is? ff} &\leftarrow p_1 \text{ is tt}, \dots, p_n \text{ is tt} \\ q_2 \text{ is? ff} &\leftarrow p_1 \text{ is tt}, \dots, p_n \text{ is tt}, q_1 \text{ is ff} \\ &\dots \\ q_m \text{ is? ff} &\leftarrow p_1 \text{ is tt}, \dots, p_n \text{ is tt}, q_1 \text{ is ff}, \dots, q_{m-1} \text{ is ff} \\ p \text{ is \{tt\}} &\leftarrow p_1 \text{ is tt}, \dots, p_n \text{ is tt}, q_1 \text{ is ff}, \dots, q_{m-1} \text{ is ff}, q_m \text{ is ff} \end{aligned}$$

It is this translation that we will use to prove Theorem 2.12, but it is worth noting that nothing in the proof *requires* the introduced open rules to have all the premises shown above. The proof works equally well if arbitrary premises are removed from the open rules, so this translation should be seen as one end of a spectrum of possible translations, the version that is the most "hesitant" to derive facts of the form $q \text{ is? ff}$. At the other end of the spectrum is the translation that just introduces a rule $(q_i \text{ is? ff} \leftarrow)$ for every negated premise in the source program.

We will prove the two directions of Theorem 2.12 separately as Proposition A.3 and Proposition A.4

PROPOSITION A.3. *Let P be an ASP program, and let $\langle P \rangle$ be the interpretation of P as a finite-choice logic program as defined above. For all stable models X of P , there is a solution D to $\langle P \rangle$ such that $X = \{p \mid p \text{ is tt} \in D\}$.*

PROOF. Let X be a stable model of P . The reduct P^X is just a datalog program, so by Lemma A.1 we can get some sequence $X_{\text{seq}} = p_1 \dots p_k$ containing exactly the elements of X such that each p_i is the immediate consequence of the preceding facts according to some rule in the reduct P^X .

We will construct our solution D in two steps. In the first step, we establish by induction on j that, for any prefix $p_1 \dots p_j$ of X_{seq} , there is a database D_j and step sequence $\emptyset \dots D_j$ under such that $\langle P \rangle$ such that:

- $p' \text{ is tt} \in D_j$ if and only if $p' \in p_1, \dots, p_j$.
- $p' \text{ is ff} \in D_j$ only if $p' \notin X$.

Applying this inductive construction to the entire sequence X_{seq} , we get a D_k with $p' \text{ is tt} \in D$ if and only if $p' \in X$. In the second step, we add additional facts of the form $q' \text{ is ff}$ to D_k until we have a solution D to $\langle P \rangle$ where $p' \text{ is tt} \in D$ if and only if $p' \in X$.

First step. The base case of our induction is immediate for $D_0 = \emptyset$. In the inductive case, we're given a step sequence $\emptyset \dots D_{j-1}$ under $\langle P \rangle$ where D_{j-1} contains $p' \text{ is tt}$ for exactly the $p' \in p_1 \dots p_{j-1}$. We need to construct D_j by adding one fact $p_j \text{ is tt}$ to D_{j-1} . We can also add any facts of the form $p' \text{ is ff}$ so long as $p' \notin X$.

When we derive p_j in the reduct, we know that this was due to a rule in the reduct whose positive premises all appear in X_{seq} prior to p . This rule in the reduct P^X was itself derived from a rule $p_j \leftarrow p_1, \dots, p_n, \neg q_1, \dots, \neg q_m$ in the original program P . Having found this rule, we let $D_j = D_{j-1} \cup \{p \text{ is tt}, q_1 \text{ is ff}, \dots, q_m \text{ is ff}\}$. This immediately satisfies the first condition that $p' \text{ is tt} \in D$ if and only if $p' \in p_1, \dots, p$. Because the reduct survived into P^X , we know none of the q_i are in X , so the third condition that $p' \text{ is ff} \in D$ only if $p' \notin X$ is also satisfied.

We need only to construct a step sequence from D_{j-1} to D_j . We will first take m steps, one for every negative premise, using the m open rules in the translation to add $q_i \text{ is ff}$ for i in order from 1 to m . Having added all the relevant negative premises, we have satisfied all the premises of the one translated closed rule:

$$p_j \text{ is \{tt\}} \leftarrow p_1 \text{ is tt}, \dots, p_n \text{ is tt}, q_1 \text{ is ff}, \dots, q_m \text{ is ff}$$

Therefore, we can take a last step with this rule to add we can add $p_j \text{ is tt}$.

Second step. The database D_k we obtained in the first step might not be a solution, because it might not be saturated. In particular, there may be open rules that allow us to derive additional non-contradictory facts $p' \text{ is ff}$. This would be the case, for example, if our original program looked something like this:

$$\begin{aligned} c &\leftarrow \\ b &\leftarrow c \\ a &\leftarrow b, \neg d, \neg c \end{aligned}$$

In this case, D_k would only contain $c \text{ is tt}$ and $b \text{ is tt}$. In order to get a solution database, $d \text{ is ff}$ would need to be added.

First, we'll establish that no additional facts of the form $p \text{ is tt}$ can be added to D_k or any consistent extension of D_k that only adds facts of the form $q \text{ is ff}$. We consider every rule

$$p \text{ is \{tt\}} \leftarrow p_1 \text{ is tt}, \dots, p_n \text{ is tt}, q_1 \text{ is ff}, \dots, q_m \text{ is ff}$$

in the translated program. There are three cases to consider:

- (1) If $p \text{ is tt} \in D_k$, then the rule can't add any new facts.
- (2) If $p \text{ is tt} \notin D_k$ and the original ASP rule was *not* eliminated from the reduct P^X , then $p \leftarrow p_1, \dots, p_n$ appears in the reduct. Because p is not in the model of the reduct, there must be some premise $p_i \notin X$, meaning $p_i \text{ is tt} \notin D_k$, so the rule doesn't apply and, furthermore, adding additional consistent facts of the form $q \text{ is tt}$ will not change that.
- (3) If $p \text{ is tt} \notin D_k$ and the original ASP rule *was* eliminated from the reduct P^X , then that rule has a premise $q_i \text{ is ff}$ that cannot possibly be satisfied in a consistent extension of D_k , because $q_i \text{ is tt} \in D_k$. Therefore, the rule doesn't apply and, furthermore, adding additional consistent facts of the form $q \text{ is ff}$ will not change that.

Because no new facts of the form $p \text{ is tt}$ can be derived, and every rule that might introduce a $p \text{ is ff}$ fact is an open rule, we can iteratively apply rules that derive new consistent $p \text{ is ff}$ facts to D_k until no additional such facts can be consistently added. The result is a consistent, saturated database: a solution D where $p \text{ is tt} \in D$ if and only if $p \in X$. \square

PROPOSITION A.4. Let P be an ASP program, and let $\langle P \rangle$ be the interpretation of P as a finite-choice logic program as defined above. For all solutions D of $\langle P \rangle$, the set $\{p \mid p \text{ is tt} \in D\}$ is a stable model of P .

PROOF. Let D be a solution to $\langle P \rangle$, and let X be $\{p \mid p \text{ is tt} \in D\}$. We can compute the reduct P^X , and that reduct P^X has a unique (and finite) model M . If $X = M$, then X is a stable model of P .

First step: $M \subseteq X$. M is the model of P^X , and so by Lemma A.1 we can get some M_{seq} containing exactly the elements of M , and where each element in the sequence is the immediate consequence of the previous elements. We'll prove by induction that for every prefix of M_{seq} , the elements of that sequence are a subset of $X = \{p \mid p \text{ is tt} \in D\}$. The base case is immediate, because $\emptyset \subseteq X$.

When we want to extend our prefix with a new fact p , we have that p was derived by a rule of the form $p \leftarrow p_1, \dots, p_n$ in the reduct P^X , which itself comes from a rule $p \leftarrow p_1, \dots, p_n, \neg q_1, \dots, \neg q_m$ in the program P . By the induction hypothesis, $p_i \text{ is tt} \in D$ for $1 \leq i \leq n$.

By virtue of the fact that this rule was in the reduct P^X , it must *not* be the case that $q_i \text{ is tt} \in D$ for $1 \leq i \leq m$. There are m open rules of the form

$$q_j \text{ is ff} \leftarrow p_1 \text{ is tt}, \dots, p_n \text{ is tt}, q_1 \text{ is ff}, \dots, q_{j-1} \text{ is ff}$$

for $1 \leq j \leq m$, and we can show by induction on j that the presence of this rule, plus the fact that $q_j \text{ is tt} \notin D$, that it must be the case that $q_j \text{ is ff} \in D$, because otherwise D would not be saturated.

This means that all the premises of the closed rule deriving $p \text{ is tt}$ are satisfied, so by the saturation of D we must have $p \text{ is tt} \in D$ and therefore $p \in X$.

Second step: $X \subseteq M$. For every $p \text{ is tt} \in D$, we must show that p is in the model of P^X . We'll prove by induction on j that, if we have a step sequence $\emptyset \dots D_j$ and $D_j \subseteq D$, then $\{p \mid p \text{ is tt} \in D_j\}$ is a subset of the model of P^X .

In the base case, we have null step sequence from \emptyset to itself. It's necessarily the case that \emptyset is a subset of the model of P^X .

When we extend the derivation by deriving a new fact of the form $p \text{ is ff}$, we can directly use the induction hypothesis, because the set $\{p \mid p \text{ is tt} \in D_j\}$ is the same as the set $\{p \mid p \text{ is tt} \in D_{j-1}\}$.

When we extend the derivation with the rule

$$p \text{ is tt} \leftarrow p_1 \text{ is tt}, \dots, p_n \text{ is tt}, q_1 \text{ is ff}, \dots, p_m \text{ is ff}$$

then we need to show that p is in the model of P^X . We know by the definition of the translation that a corresponding rule $p \leftarrow p_1, \dots, p_n, \neg q_1, \dots, \neg q_m$ must appear in P .

For $1 \leq i \leq m$, we have $q_i \text{ is ff} \in D_{j-1}$ by the fact that the rule was applied, so its premises must have been satisfied. Because $D_{j-1} \subseteq D$ we must have $q_i \text{ is ff} \in D$, and because D is consistent, $q_i \text{ is tt} \notin D$, so $q_i \notin X$ for all $1 \leq i \leq m$. That means the rule $p \leftarrow p_1, \dots, p_n$ appears in the reduct P^X .

For $1 \leq i \leq n$, we have $p_i \text{ is tt} \in D_{j-1}$ by the fact that the rule was applied, so its premises must have been satisfied. By the induction hypothesis these p_i must all be in the model of P^X .

If the premises of a rule are in the model then the conclusion must be in the model, so p is in the model of P^X . \square

B LEAST UPPER BOUNDS FOR CHOICE SETS

The following theorems establish that the pointed partial order *Choice* is a complete lattice with the least upper bound defined in Definition 4.11 as follows: Take any $\{C_i : i \in I\} \subseteq \text{Choice}$, and let \mathcal{F} be the set of functions $f : I \rightarrow DB$ such that $f(i) \in C_i$. (Equivalently, \mathcal{F} is the direct product $\prod_{i \in I} C_i$.) We define $\bigvee \{C_i : i \in I\}$ as $\{\bigvee \mathcal{D} : f \in \mathcal{F}, \mathcal{D} = \{f(i) : i \in I\}, \|\mathcal{D}\}\}$.

We must prove that this function that takes sets of choice sets always produces a choice set, and that the result is actually a least upper bound.

LEMMA B.1. $\bigvee_i C_i$ is an element of Choice, i.e. is pairwise incompatible.

PROOF. Consider databases $D_1, D_2 \in \bigvee_i C_i$. These are the least upper bound of the sets of databases picked out by f_1 and f_2 , respectively. If $D_1 \neq D_2$, then there is some index i for which $f_1(i) \neq f_2(i)$, and $f_1(i)$ and $f_2(i)$ come from the same choice set C_i , so they must be incompatible. By monotonicity of incompatibility (Lemma 4.2), D_1 and D_2 are also incompatible. \square

LEMMA B.2. \bigvee_{Choice} is an upper bound in Choice. That is, if $C \in \mathcal{S}$ then $C \leq \bigvee \mathcal{S}$.

PROOF. Let $C_i \in \mathcal{S}$. Every element of $\bigvee \mathcal{S}$ is the join over some selection of compatible databases, exactly one of which is an element of C_i . That is, if $X_j \in \bigvee \mathcal{S}$, then $X_j = \bigvee_i f_j(i)$ where $f_j(i) = D_i \in C_i$. Thus $D_i \leq X_j$, as required. \square

THEOREM B.3. $\bigvee_i C_i$ is the least upper bound of $\{C_i : i \in I\} \subseteq \text{Choice}$.

PROOF. Let C' be an upper bound of $\{C_i : i \in I\}$. To show that $\bigvee_i C_i \leq C'$, it suffices to show that for every $D' \in C'$, there exists some $D \in \bigvee_i C_i$ such that $D \leq D'$.

Therefore, let D' be an arbitrary element of C' . For every $i \in I$, there is (by definition of \leq_{Choice}) a unique $D_i \in C_i$ such that $D_i \leq D'$, and we can define $f : I \rightarrow DB$ to pick out that D_i . The set of databases $\{f(i) : i \in I\}$ is compatible because D' is an upper bound. That means that $\bigvee_i C_i$ contains $\bigvee_i f(i)$, and because \bigvee_{DB} is already established as a least upper bound, $\bigvee_i f(i) \leq_{DB} D'$. Letting $D = \bigvee_i f(i)$, establishes that \bigvee_{Choice} is also a least upper bound. \square

C EQUIVALENCE OF SEMANTICS

This section contains the proof that three semantics for finite-choice logic programs—the step sequences in Section 2, the least-fixed-point interpretation in Section 5, and the algorithm in Section 6—all agree.

These proofs relate sequences of fact sets and the least fixed-point of an immediate consequence operator on choice sets, sets of mutually exclusive constraint databases. In the main body of the paper, we frequently conflated sets of facts and constraint databases. In this appendix, we will avoid this conflation, and will consistently use D to refer to databases that are comprised of sets of facts as defined in Section 2 but use Δ to refer to constraint databases as defined in Section 5. Fact sets and constraint databases are connected by two adjoint functors, erasure and promotion.

Definition C.1 (Erasure of a constraint database to a set of facts). For all constraint databases $\Delta \in DB$, define $(\Delta)^- = \{a \text{ is } v \mid \Delta(a) = \text{just } v\}$. This set is a consistent database of facts as defined in Definition 2.6.

Definition C.2 (Promotion of a set of facts to a constraint database). If the set of facts D is consistent by Definition 2.6, then for any a there is at most one v such that $a \text{ is } v \in D$. Therefore, we can define $(D)^+$ as a function from consistent sets of facts to constraint databases as follows:

- $(D)^+(a) = \text{just } v$ if $a \text{ is } v \in D$
- $(D)^+(a) = \text{noneOf } \emptyset$ if there does not exist a v such that $a \text{ is } v \in D$

Here are two examples of these definitions in action:

- $(\{p \text{ is } tt\})^+ = (p \mapsto \text{just } tt)$
- $(p \mapsto \text{just } tt, q \mapsto \text{noneOf } \{a, b\})^- = \{p \text{ is } tt\}$

LEMMA C.3 (PROPERTIES OF ERASURE AND PROMOTION). *The following all hold:*

- For all consistent sets of facts D , it is the case that $(D)^+$ is positive and $((D)^+)^- = D$.
- For all $\Delta \in DB$, it is the case that $((\Delta)^-)^+ \leq_{DB} \Delta$.
- For all $\Delta \in DB$, it is the case that $((\Delta)^-)^+ = \Delta$ if and only if Δ is positive.

PROOF. Straightforward unfolding of definitions. \square

With these definitions, we can precisely restate Theorem 5.12 without the imprecise punning of fact sets and constraint databases:

PROPOSITION C.4 (RESTATEMENT OF THEOREM 5.12). *For all P and D , the following are equivalent:*

- (1) *There is a fact-set step sequence $\emptyset \dots D$ and whenever $D \Rightarrow_P S$ then $S = \{D\}$.*
- (2) *$(D)^+ \in \text{lfp } \mathcal{T}_P$ and D is finite.*

C.1 Correspondence of the sequential, algorithmic, and least-fixed-point semantics

The algorithm described in Section 6.1 defines another kind of sequence, which we can call an *algorithmic step sequence* to distinguish it from a fact-set step sequence.

The definition of algorithmic steps takes advantage of the fact that all rule consequences are singular in some attribute a , which lets us define *attribute-specific* immediate consequence $\tau_{P[a]}(\Delta)$, mapping an attribute a to a choice set, the elements of which each map a to a constraint (by the definition of choice sets, those constraints are pairwise-incompatible):

Definition C.5 (*Attribute-specific immediate consequence*).

$$\tau_{P[a]}(\Delta) = \bigvee \{ \langle \sigma H \rangle : (H \leftarrow F) \in P, \sigma F \leq \Delta, \sigma H @ a \}$$

Where $\sigma H @ a$ means $\sigma H = a \text{ IS } S$ or $\sigma H = a \text{ IS? } v$.

In Section 6.1 we wrote $\tau_P(D)[a]$ for the concept we are precisely defining as $\tau_{P[a]}(D)$. The action of the algorithm described in Section 6.1 can be characterized in terms of $\tau_{P[a]}$ as follows:

Definition C.6 (*Algorithmic steps*). Let Δ be a constraint database. We say that the program P allows the constraint database Δ to make an *algorithmic step* to $\Delta' \in \{\Delta\} \vee \tau_{P[a]}(\Delta)$ whenever $\tau_P(\Delta)$ is nonempty.

Algorithmic steps give rise to algorithmic step sequences: $\Delta_0 \dots \Delta_k$ is an algorithmic step sequence for P if, for each $i > 0$, the program P allows Δ_{i-1} to take an algorithmic step to Δ_i . (As with fact-set step sequences, algorithmic step sequences are assumed to be finite.)

These definitions allow us to break our central theorem into several lemmas:

LEMMA C.7 (COMPLETENESS OF THE ALGORITHM). *If there is a fact-set step sequence $\emptyset \dots D$ under P and D is saturated by Definition 2.10, then there is an algorithmic step sequence $\perp_{DB} \dots (D)^+$ under P and $(D)^+$ is a model by Definition 5.2.*

PROOF. See Appendix C.2. \square

LEMMA C.8 (SOUNDNESS OF THE ALGORITHM). *If there is an algorithmic step sequence $\perp_{DB} \dots \Delta$ and Δ is a model of P by Definition 5.2, then $\Delta \in \text{lfp } \mathcal{T}_P$.*

PROOF. See Appendix C.3. \square

LEMMA C.9 (REACHABILITY OF FINITE FIXED POINTS). *If $\Delta \in \text{lfp } \mathcal{T}_P$ and Δ is finite, then there is a fact-set step sequence $\emptyset \dots (\Delta)^-$.*

PROOF. See Appendix C.4. \square

LEMMA C.10. *If Δ is a positive model, $(\Delta)^-$ is saturated (Definition 2.10).*

PROOF. Inspection of definitions: given a positive constraint database Δ , we can show that any non-trivial evolution of $(\Delta)^-$ would contradict the assumption that Δ is in $\text{lfp } \mathcal{T}_P$. (This would not be true if Δ were not positive.) \square

These four lemmas suffice to prove Theorem 5.12/Proposition C.4 and, in a different configuration, the correctness of the algorithm, Theorem 6.2.

The algorithm in Section 6.1 is really an algorithm for enumerating all finite minimal models, not just for finding the positive models that correspond to solutions. However, Theorem 6.2 only talks about positive models. We conjecture the following, which is just Theorem 6.2 with the word “positive” removed, but the proof requires defining a constraint-database analogue to fact-step sequences which we have not included.

CONJECTURE C.11. *For constraint databases Δ , the following are equivalent:*

- (1) *The algorithm described in Section 6.1 may successfully return Δ .*
- (2) *$\Delta \in \text{lfp } \mathcal{T}_P$ and Δ is finite.*

C.2 Completeness of the algorithm

Here we prove Lemma C.7: if there is a fact-set step sequence $\emptyset \dots D$ under P and D is saturated by Definition 2.10, then there is an algorithmic step sequence $\perp_{DB} \dots (D)^+$ under P and $(D)^+$ is a model of P by Definition 5.2.

First, a warm-up lemma:

LEMMA C.12. *If $\tau_P((D_0)^+) = \emptyset$ and $D_0 \dots D_n$ is a fact-step sequence, then D_n is not saturated.*

PROOF. By induction on n , in each D_i one of these two conditions apply:

- There is an attribute a and two rule groundings with satisfied premises in D_i and conclusions $a \text{ IS } \{v_1, \dots, v_n\}$ and $a \text{ IS } \{w_1, \dots, w_m\}$ where the v_i and w_i do not overlap.
- $\{a \text{ IS } v\} \in D_i$ and there is a rule grounding with satisfied premises in D_i and conclusion $a \text{ IS } \{w_1, \dots, w_m\}$, where v is not one of the w_i .

Either of these cases mean that D_i cannot be saturated. \square

By induction, we will prove that if there is a fact-set step sequence $D_1 \dots D_k$ under P , then either D_k cannot be extended to a saturated set of facts — there is no solution D' such that $D_k \dots D'$ is a saturated — or $(D_1)^+ \dots (D_k)^+$ is an algorithmic step sequence under P .

The base case is immediate: given a fact-set step sequence containing no steps, provide an algorithmic step sequence containing no steps.

In the inductive case, we have a fact-set step sequence $D_1 \dots, D_k, D_{k+1}$ and we know one of two things is true: either D_k cannot be extended to a solution, or else $(D_1)^+ \dots (D_k)^+$ is an algorithmic step sequence under P . We need to show either that D_{k+1} cannot be extended to a solution or else that $(D_1)^+ \dots (D_{k+1})^+$ is an algorithmic step sequence under P .

If D_k can't be extended to a solution neither can D_{k+1} , satisfying the first disjunctive option in the induction hypothesis.

If $D_k = D_{k+1}$ then $(D_1)^+ \dots (D_k)^+$ is the same thing as $(D_1)^+ \dots (D_{k+1})^+$, and we're done. Otherwise, by inspecting the definitions we see $D_{k+1} = \{a \text{ IS } v\} \cup D_k$ for some a and v . We identify three cases:

- If $\tau_P((D)^+) = \emptyset$, the first disjunctive option in the induction hypothesis is satisfied by Lemma C.12.
- If $\tau_P((D)^+) \neq \emptyset$ and $(a \mapsto \text{just } v) \in \tau_{P[a]}((D)^+)$, then $(D_{k+1})^+ \in \{(D_k)^+\} \vee \tau_{P[a]}((D_k)^+)$. This means $(D_k)^k$ can take an algorithmic step to $(D_{k+1})^+$, satisfying the second disjunctive option in the induction hypothesis.

- If $\tau_P((D)^+) \neq \emptyset$ and $(a \mapsto \text{just } v) \notin \tau_{P[a]}((D)^+)$, then we must be able to identify some closed rule $H \leftarrow F$ in P such that some σ satisfies F in $(D_k)^+$, where $\sigma H = a \text{is } \{v_1, \dots, v_n\}$, and where v is not one of the v_i . This implies that in D_{k+1} , or in any extension of it, evolves to the empty set (Definition 2.8) and so therefore cannot be a solution or extended to a solution, satisfying the first disjunctive option in the induction hypothesis.

Applying this proof by induction to the given fact-set step sequence, we obtain from the fact-step sequence $\emptyset \dots D$ an algorithmic step sequence $\perp_{DB} \dots (D)^+$, because it would be an immediate contradiction to say that the saturated database D could not be extended to a saturated database. By inspection we see that if D is saturated by Definition 2.10 then $(D)^+$ is a model, which completes the proof.

C.3 Soundness of the algorithm

Here we prove Lemma C.8: if there is an algorithmic step sequence on P from \perp_{DB} to Δ and Δ is a model of P , then $\Delta \in \text{lfp } \mathcal{T}_P$. Since $\text{lfp } \mathcal{T}_P$ is the set of minimal models of P (Theorem 5.9), to prove Lemma C.8 it suffices to show Δ is minimal, which we show in Lemma C.13:

LEMMA C.13. *Let $\perp_{DB} \dots \Delta$ be an algorithmic step sequence for a program P , with Δ a model of P . Then Δ is minimal: if M is a model and $M \leq \Delta$ then $M = \Delta$.*

PROOF. We will show inductively that each step Δ_i in the sequence $\perp_{DB} \dots \Delta$ is below M , and so in particular $\Delta \leq M$.

Base case: $\Delta_0 = \perp_{DB}$. We have $\perp_{DB} \leq M$ by definition.

Inductive case: Given $\Delta_i \leq M$, we need to show $\Delta_{i+1} \leq M$. By definition C.6, we know $\Delta_{i+1} \in \{\Delta_i\} \vee \tau_{P[a]}(\Delta_i)$. As a first step, we show $\tau_{P[a]}(\Delta_i) \leq \{M\}$:

$$\begin{aligned}
 \tau_{P[a]}(\Delta_i) &= \bigvee \{ \langle \sigma H \rangle : (H \leftarrow F) \in P, \sigma F \leq \Delta_i, \sigma F @ a \} \\
 &\leq \bigvee \{ \langle \sigma H \rangle : (H \leftarrow F) \in P, \sigma F \leq \Delta_i \} \\
 &\leq \bigvee \{ \langle \sigma H \rangle : (H \leftarrow F) \in P, \sigma F \leq M \} && \text{since } \Delta_i \leq M \\
 &\leq \tau_P(M) && \text{definition of } \tau_P \\
 &= \{M\} && \text{since } M \text{ is a model}
 \end{aligned}$$

This shows $\{M\}$ upper bounds $\tau_{P[a]}(\Delta_i)$; it is also an upper bound of $\{\Delta_i\}$, since $\Delta_i \leq M$. Therefore it is above their least upper bound: $\{\Delta_i\} \vee \tau_{P[a]}(\Delta_i) \leq \{M\}$. By definition of \leq_{Choice} this means there is some $\Delta' \in \{\Delta_i\} \vee \tau_{P[a]}(\Delta_i)$ with $\Delta' \leq M$. To show $\Delta_{i+1} \leq M$ it suffices to show that this Δ' is in fact Δ_{i+1} . Since they are elements of the choice set $\{\Delta_i\} \vee \tau_{P[a]}(\Delta_i)$, this follows if they are compatible. And indeed, our final destination Δ upper bounds both; the first because $\Delta' \leq M \leq \Delta$; the second because the step sequence is monotonically increasing, $\perp_{DB} \leq \dots \leq \Delta_{i+1} \leq \dots \leq \Delta$. \square

C.4 Reachability of finite fixed points

Here we prove Lemma C.9: If $\Delta \in \text{lfp } \mathcal{T}_P$ and Δ is finite, then there is a fact-set step sequence $\emptyset \dots (\Delta)^-$.

LEMMA C.14. *If $\Delta \in \text{lfp } \mathcal{T}_P$ and Δ is finite, then there is a n such that $\Delta \in \mathcal{T}_P^n(\perp)$*

PROOF. For all i , it is the case that $\mathcal{T}_P^i(\perp) \leq \text{lfp } \mathcal{T}_P$, so we can define an infinite sequence of Δ_i such that each $\Delta_i \leq \Delta$ and $\Delta_i \in \mathcal{T}_P^i(\perp)$.

If we ever have a i such that $\Delta_i = \Delta_{i+1}$, then $\Delta_i \in \text{lfp } \mathcal{T}_P$, meaning that i is the n we are looking for. (Proof: Δ_i is a model, so $\Delta \leq \Delta_i$, and because $\mathcal{T}_P^i(\perp) \leq \text{lfp } \mathcal{T}_P$, $\Delta_i \leq \Delta$.)

Therefore, to show that there exists an n such that $\Delta \in \mathcal{T}_p^n(\perp)$, it suffices to show that the sequence of Δ_i cannot be strictly ascending without limit.

If Δ_{i-1} and Δ_i agree on all attributes where $\Delta_i[a] = \text{just } t$, then $\Delta_i = \Delta_{i+1}$, so an infinite ascending sequence implies a that each Δ_i would have to contain at least i attributes that map to values **just** t . Because Δ is finite and $\Delta_i \leq \Delta$, this is impossible, so there must be the required finite n . \square

LEMMA C.15. *If Δ_i is finite and takes an algorithmic step to Δ_{i+1} , then $(\Delta_i)^- \dots (\Delta_{i+1})^-$ is a fact-set step sequence.*

PROOF. This amounts to showing two things. First, we observe that the same substitutions satisfy the same premises according to Definitions 2.7 and 4.12 (the extra information removed in the erasure transformation isn't relevant to which rules apply). Second, we observe that there are a finite number of relevant rule heads σH in any $\tau_P(\Delta)$, a straightforward consequence of the finite size of the database and the restriction in Definition 2.3 that all variables in the head of a rule appear in a premise. \square

The proof of Lemma C.9 follows from these two lemmas: we string together the required fact-set step sequence by concatenating finite step sequences between each of the finitely-many relevant iterations of the immediate consequence operator.

D EXAMPLES AND BENCHMARKS FOR DUSA

We compare DUSA against two existing systems: Clingo, part of the Postdam Answer Set Solving Collection (Potassco) and generally recognized as the industry standard [Gebser et al. 2011], and Alpha, a implementation of answer set programming with lazy grounding [Weinzierl 2017]. As discussed in Sections 3.3 and 7.2, lazy-grounding systems are, like DUSA, able to handle problems that pre-grounding systems like Clingo cannot, and are able to more efficiently compute solutions in cases where the fully-ground program is quite large (this is called the “grounding bottleneck”).

There are three other often-cited lazy-grounding ASP solvers: Gasp, ASPeRiX, and Omiga. We were unable to find source code for Gasp, and we were unable to successfully compile ASPeRiX.

All the results in this section were generated on AWS m7g.xlarge instances (64-bit 4-core virtual machines with 16 GB of memory running Ubuntu Linux) with processes limited to 10 GB of memory and 100 seconds of wall-clock time.¹ We discuss three types of benchmarks:

- In Appendix D.1, we cover spanning tree generation and canonical representative identification, examples from Section 3 well-suited for finite-choice logic programming and the DUSA implementation.
- In Appendix D.2, we cover N-queens and Adam Smith's “Map Generation Speedrun,” pure search problems that are expressible in finite-choice logic programming but that demonstrate the kinds of problems that the current DUSA implementation is ill-suited for.
- In Appendix D.3, we replicate a set of benchmarks designed for the Alpha implementation of answer set programming with lazy grounding.

Concrete DUSA syntax. The DUSA examples in this section use the concrete syntax from the implementation, which follows the finite-choice logic programming notation from the paper closely. As is common in logic programming languages, variables are distinguished by being capitalized and the arrow \leftarrow is replaced by “:-”. Less commonly, DUSA uses an un-curried logic programming notation, so we write “color X Y is red” to express a premise we would have written in this paper as $\text{color}(x, y) \text{ is red}$. In addition:

¹See <https://github.com/robsimmons/dusa-benchmarking/> for all code and analysis.

- The implementation directly supports attributes with no values: writing “edge $X\ Y$ ” is equivalent to `edge(x, y) is unit` as motivated in Section 2.2.
- When the conclusion of a closed rule only has one value, we are allowed to optionally omit the curly-braces: we can write “parent $X\ X$ is tt” for the head `parent(x, x) is {tt}`, though `parent(x, x) is {tt, ff}` requires curly braces and will be written “parent $X\ X$ is { tt, ff }”.
- The implementation supports #forbid rules, a list of premises which cannot be simultaneously satisfied in a valid solution. The rule “#forbid $p\ X, q\ X$ ” in DUSA’s concrete syntax can be desugared into the rule (`ok is {no} $\leftarrow p(x), q(x)$`) with the addition a new predicate `ok` and rule (`ok is {yes} \leftarrow`).

Translating answer set programs to finite-choice logic programs. As discussed in Section 7.2, we can perform fair head-to-head comparisons between ASP solvers and DUSA by using the translation from answer set programming to finite-choice logic programming presented in Appendix A.2. We will extend that translation in two ways. First, ASP features *constraints*, in the form of headless rules ($\leftarrow F$), which translate directly to the #forbid rules introduced above. The second feature is *choice rules*, which allow the ASP engine to treat a predicate as either true or false. A choice rule written in ASP as “{ p } :- q , not r ” can be represented in finite-choice logic programming with the rule (`p is {tt, ff} $\leftarrow q$ is tt, r is ff`). The general correctness of this extended translation is left for future work (as is the general correctness of the translation on open programs, as discussed in Section 3.3).

D.1 Examples from the paper

The spanning tree generation and canonical representative identification discussed in Section 3 are both cases where finite-choice logic programming allows the problem to be described quite concisely. ASP implementations of the same algorithm are less consistent and less efficient, as we see in this section, which expands on the discussion in Sections 7.1 and 7.2.

D.1.1 Spanning tree generation. A finite-choice logic program for spanning tree generation is presented in Section 3.1, and a pure ASP program for solving the same problem, as well as that program’s translation into DUSA’s concrete syntax, is given in Figure 11. The best solution we know of in idiomatic Clingo ASP replaces the four rules in Figure 11 labeled “Exactly one root” with the single cardinality constraint `1 {root(X) : edge(X, Y)} 1`.

These programs were tested on a graphs with different properties, with results shown in Figure 12. In most cases, DUSA is outperformed by the idiomatic Clingo ASP solution on smaller graphs, but Clingo’s performance seems to grow faster than linearly. In a fair fight, DUSA and Alpha outperform Clingo on some sparse graphs, which is what we would expect if Clingo’s performance issues were connected to the ground-then-solve methodology.

D.1.2 Appointing canonical representatives. The idiomatic finite-choice logic programming approach to appointing canonical representatives for connected components of an undirected graph is presented in Section 3.2, and a pure ASP program for solving the same program, as well as that program’s translation into DUSA’s concrete syntax, is given in Figure 13. As in the previous section, it’s possible to get a shorter and more performant Clingo program by using Clingo-supported cardinality constraints. In this case, the last three ASP rules in Figure 13 can be replaced by the single Clingo rule:

```
:- node( $X$ ), not 1{representative( $X, R$ ) : node( $R$ )} 1.
```

For this problem, idiomatic DUSA is the clear winner in all cases. Both DUSA and Alpha outperform Clingo in almost all cases when running the same ASP program, as discussed in Section 7.2. This is not terribly surprising: in a graph with v nodes, Clingo’s grounder is almost unavoidably going

```

1912 edge(X,Y) :- edge(Y,X).
1913
1914 % Exactly one root
1915 {root(X)} :- node(X).
1916 :- root(X), root(Y), X != Y.
1917 someRoot :- root(X).
1918
1919 :- not someRoot.
1920
1921 % The root has itself as a parent
1922 parent(X,X) :- root(X).
1923
1924 % Any tree node can be a parent
1925 inTree(P) :- parent(P,_).
1926 {parent(X,P)} :-
1927     edge(X,P), inTree(P).
1928
1929 % Only 1 parent
1930 :- parent(X,P1),
1931     parent(X,P2), P1 != P2.
1932
1933 % Tree covers connected component
1934
1935 :- edge(X,Y), inTree(X),
1936     not inTree(Y).
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960

```

```

edge X Y :- edge Y X.

# Exactly one root
root X is { tt, ff } :- node X.
#forbid root X is tt, root Y is tt, X != Y.
someRoot is tt :- root X is tt.
someRoot is? ff.
#forbid someRoot is ff.

# The root has itself as a parent
parent X X is tt :- root X.

# Any tree node can be a parent
inTree P is tt :- parent P _ is tt.
parent X P is { tt, ff } :-
    edge X P, inTree P.

# Only 1 parent
#forbid parent X P1 is tt,
    parent X P2 is tt, P1 != P2.

# Tree covers connected component
inTree Y is? ff :-
    edge X Y, inTree X is tt.
#forbid edge X Y, inTree X is tt,
    inTree Y is ff.

```

Fig. 11. Spanning-tree description in ASP, and direct translation to Dusa.




Fig. 12. Performance of Dusa, Clingo, and Alpha on finding spanning trees of graphs with different features. The x-axis is the number of edges in the input graph. All data points are the median of three runs requesting a single solution, discarding timeouts.

<pre> 1961 edge(Y,X) :- edge(X,Y). 1962 1963 % Any node can be a rep 1964 {representative(X,X)} :- 1965 node(X). 1966 1967 % Connected nodes have same rep 1968 representative(Y,Rep) :- 1969 edge(X,Y), 1970 representative(X,Rep). 1971 1972 % Representatives must be unique 1973 :- representative(X,R1), 1974 representative(X,R2), R1 != R2. 1975 1976 % Every node has a rep 1977 hasRep(X) :- 1978 representative(X,_). 1979 1980 :- node(X), not hasRep(X). </pre>	<pre> edge Y X :- edge X Y. # Any node can be a rep representative X X is { tt, ff } :- node X. # Connected nodes have same rep representative X Rep is tt :- edge X Y, representative Y Rep is tt. # Representatives must be unique #forbid representative X R1 is tt, representative X R2 is tt, R1 != R2. # Every node has a rep hasRep X is tt :- representative X _ is tt. hasRep X is? ff :- node X. #forbid node X, hasRep X is ff. </pre>
---	---

Fig. 13. Canonical representative identification in ASP, and direct translation to DUSA.




Fig. 14. Performance of DUSA, Clingo, and Alpha on finding canonical representatives for the connected components of graphs with different features. The x-axis is the number of edges in the input graph. All data points are the median of three runs requesting a single solution, discarding timeouts.

to require $O(v^3)$ groundings of the key rule that requires connected nodes to have the same representative. It was more unexpected that, on all but the sparsest graphs, DUSA outperformed Alpha in a head-to-head matchup despite the DUSA implementation exerting almost no control over backtracking. Extremely sparse random graphs were a notable exception that caused DUSA to time out on all but the very smallest examples.

```

{ queen(1..n,1..n) }.
:- not { queen(I,J) } == n.
:- queen(I,J), queen(I,JJ), J != JJ.
:- queen(I,J), queen(II,J), I != II.
:- queen(I,J), queen(II,JJ), (I,J) != (II,JJ), I-J == II-JJ.
:- queen(I,J), queen(II,JJ), (I,J) != (II,JJ), I+J == II+JJ.

```

Fig. 15. Basic N-queens program in Clingo’s dialect of answer-set programming.

```

#builtin INT_MINUS minus
#builtin INT_PLUS plus
dim N :- size is N.
dim (minus N 1) :- dim N, N != 1.
# Expands to: dim Nminus1 :- dim N, N != 1, minus N 1 is Nminus1.

```

Fig. 16. DUSA declarations for using arithmetic and making a dim relation containing the numbers from 1 to declared size, inclusive.

D.2 Examples ill-suited to the current Dusa implementation

In this section we compare DUSA to Clingo on two problems where we expect DUSA to perform poorly, because they require brute-force search. These represent low-hanging fruit for optimization of DUSA.

D.2.1 N-queens. The N-queens problem is a well-known chess puzzle. The problem is to place N queens on an N-by-N chessboard such that no two queens can immediately attack each other. We used the 2014 notes “Basic modeling and ASP and more via the n-Queens puzzle” by Torsten Schaub as a guide to the ASP implementations in this sections.² Schaub’s solutions are not pure ASP, so we can’t just turn a crank to translate them to DUSA. Instead, we’ll consider a spectrum of two different Clingo solutions, and three different finite-choice logic programming solutions.

Basic Clingo implementation. The most basic Clingo implementation of N-Queens is given in Figure 15. Reading the rules in order, they express:

- Each spot in an N-by-N grid can have a queen or not.
- There must be exactly N queens.
- No two queens can share the same column.
- No two queens can share the same row.
- No two queens can share the same diagonal (in either direction).

This is an example of the “sculptural” approach to answer set programming [Smith and Mateas 2011]. In the sculptural approach, a possibility space is described in the most general terms, and then constraints are used to carve out parts of the problem space that are not wanted.

An aside on externally-computed predicates in Dusa. Our implementation doesn’t support the range notation $1..n$ shown in Figure 15, but this is simple to emulate as a relation `dim` that contains all the numbers from 1 to N. This is achieved by the code in Figure 16, which is implicitly included in subsequent finite-choice logic programs in this section. Figure 16 demonstrates the use of integers and integer operations in DUSA, which is supported by our implementation and our theory. The theory of finite-choice logic programming supports treating externally-computed predicates as infinite relations. The most fundamental such relations are equality and inequality on terms and comparison of numbers, and these have built-in syntax in the DUSA implementation. Premises like

²<https://www.cs.uni-potsdam.de/~torsten/Potassco/Videos/BasicModeling/modeling.pdf>

```

2059     location N is? (tup X Y) :-          col N is { X? } :-
2060         dim N, dim X, dim Y.              dim N, dim X.
2061                                         row N is { Y? } :-
2062                                             dim N, dim Y.
2063
2063     #forbid                               #forbid                               rowFor X is { Y? } :-
2064         location N is (tup X _),          col N is X, col M is X,              dim X, dim Y.
2065         location M is (tup X _),          N != M.
2066
2066     #forbid                               #forbid                               colFor Y is X :-
2067         location N is (tup _ Y),          row N is Y, row M is Y,              rowFor X is Y.
2068         location M is (tup _ Y),          N != M.
2069
2069     #forbid                               #forbid                               posDiag (plus X Y)
2070         location N is (tup X1 Y1),        row N is X1, col N is Y1,              is (tuple X Y) :-
2071         location M is (tup X2 Y2),        row M is X2, col M is Y2,              rowFor X is Y.
2072         N != M,                          N != M,
2073         minus X1 Y1 == minus X2 Y2.        minus X1 Y1 == minus X2 Y2.
2074
2074     #forbid                               #forbid                               negDiag (minus X Y)
2075         location N is (tup X1 Y1),        row N is X1, col N is Y1,              is (tuple X Y) :-
2076         location M is (tup X2 Y2),        row M is X2, col M is Y2,              rowFor X is Y.
2077         N != M,                          N != M,
2078         plus X1 Y1 == plus X2 Y2.          plus X1 Y1 == plus X2 Y2.

```

Fig. 17. Three different implementations of N-Queens in DUSA.

“ $4 > X$ ”, “ $N \neq 1$ ”, or “ $P == \text{pair } X \ Y$ ” can be thought of as referring to infinite relations: in the language of the paper, we would write $gt(4, x)$, $neg(n, 1)$, or $eq(p, \text{pair}(x, y))$, respectively.

Our theory and implementation are able to handle operating on infinite relations under two conditions:

- (1) Infinite relations are never present in the conclusion of a rule, and
- (2) It is tractable to find all the substitutions that satisfy a premise.

In the implementation, conservative syntactic checks prevent us from writing rules like $p(x) \leftarrow gt(4, x)$ or $p(x, y) \leftarrow eq(x, y)$ that might have infinitely many satisfying substitutions.

Additional built-in operations like addition and subtraction are included with `#builtin` directives. These are also treated like infinite relations. Figure 16 also demonstrates an additional language convenience that we picked up from LogiQL [Aref et al. 2015]: the language supports treating built-in relations with functional dependencies like functions, so we can write the conclusion “dim (minus N 1)” instead of needing to write “dim Nminus1” and including “minus N 1 is Nminus1” as an additional premise.

N-queens in DUSA. The notation $\{ \text{queen}(I, J) \} == n$ in Figure 15 is an aggregate operation: it counts the number of occurrences of the queen predicate. Aggregates in logic programming are, in many ways, similar to negation, because they allow one part of a program to reflect on the presence *and absence* of information in another part of the program. DUSA currently supports neither negation nor aggregation, but a future implementation of finite-choice logic programming could potentially support a stratified form of both. In order to adapt the N-queens problem to finite-choice logic programming, we need a way to ensure that there will be exactly N queens without aggregation.

Of the three programs in Figure 17, the first is the most sculptural: we number each of the N queens and assign a unique coordinate pair to each queen. Our prefix-firing cost semantics suggests that the first rule in this program has a complexity of $O(n^3)$ for n queens, and this observation

```

{ queen(I,1..n) } == 1 :- I = 1..n.
{ queen(1..n,J) } == 1 :- J = 1..n.

:- { queen(D-J,J) } >= 2, D = 2..2*n.
:- { queen(D+J,J) } >= 2, D = 1-n..n-1.

```

Fig. 18. Optimized N-queens implementation in Clingo.




Fig. 19. Performance of the previous five implementations of the N-queens problem in Dusa and Clingo. All data points are the median of three runs requesting 10 different solutions, discarding timeouts.

also suggests a slight improvement. The middle program in Figure 17 is a rewrite of the first one where we switch to sixth normal form (as the literature on LogiQL suggests we should prefer [Aref et al. 2015]). Our cost semantics predicts more manageable behavior, quadratic instead of cubic, for this middle program. Despite the fact that any polynomial effect is dwarfed by exponential backtracking, the intuition it provides seems good: the middle program performs slightly better than the left program when we test performance.

If we think a bit about the problem, we can avoid the use of `#forbid` constraints entirely, as shown in the third column of Figure 17. Each column is labeled by an x coordinate, and each row is labeled by a y coordinate, and there is a functional dependency in both directions: each row has a queen in one column, and each column has a queen in one row. This is expressed by the first two rules in the rightmost program. Furthermore, we can assign every cell's coordinates to a unique numbered diagonal by adding the coordinates (for the diagonals that go up and to the right) or subtracting the coordinates (for the diagonals that go down and to the right). The last two rules in the third program use functional dependencies to invalidate any assignment that places two queens on the same diagonal.

Better Clingo implementation. The program we worked our way towards in the previous section uses the same ideas as one of Torsten Schaub's improved Clingo programs, shown in Figure 18. Each of the four rules in Schaub's more efficient Clingo program corresponds loosely with one of the four rules in our final program in Figure 17, but Schaub's version uses Clingo's counting constraints to enforce the at-most-one-queen-per-diagonal property that is enforced by functional dependencies in finite-choice logic programming.

Performance. The performance of the two Clingo programs and three Dusa programs from this section is shown in Figure 19 with the problem size scaling from 1 to 100 queens. The first Clingo program scales quite badly and starts hitting the 100-second timeout in at least some runs starting at 12 queens, whereas the optimized Clingo program in Figure 18 is far and away the best. The three successive Dusa programs fall in the middle: the third program from Figure 17 is the best but starts to encounter timeouts at 28 queens.

2157		#builtin INT_PLUS plus
2158		#builtin INT_MINUS minus
2159	dim(1..W) :- width(W).	dim Width :- width is Width.
2160		dim (minus N 1) :- dim N, N != 1.
2161		
2162	{ solid(X,Y) :dim(X) :dim(Y) }.	solid X Y is { tt, ff } :-
2163		dim X, dim Y.
2164	start(1,1).	start 1 1.
2165	finish(W,W) :- width(W).	finish W W :- width is W.
2166	step(0,-1).	step 0 -1.
2167	step(0,1).	step 0 1.
2168	step(1,0).	step 1 0.
2169	step(-1,0).	step -1 0.
2170	reachable(X,Y) :-	reachable X Y :-
2171	start(X,Y),	start X Y,
2172	solid(X,Y).	solid X Y is tt.
2173	reachable(NX,NY) :-	reachable NX NY :-
2174	reachable(X,Y),	reachable X Y,
2175	step(DX,DY),	step DX DY,
2176	NX = X + DX,	NX == plus X DX,
2177	NY = Y + DY,	NY == plus Y DY,
2178	solid(NX,NY).	solid NX NY is tt.
2179	complete :- finish(X,Y),	#demand finish X Y,
2180	reachable(X,Y).	reachable X Y.
2181	:- not complete.	
2182	at(X,Y, 0) :-	at X Y 0 :-
2183	start(X,Y),	start X Y,
2184	solid(X,Y).	solid X Y is tt.
2185	at(NX,NY, T+1) :-	at NX NY (plus T 1) :-
2186	at(X,Y,T),	at X Y T,
2187	length(Len),	length is Len,
2188	T < Len,	T < Len,
2189	step(DX,DY),	step DX DY,
2190	NX = X + DX,	NX == plus X DX,
2191	NY = Y + DY,	NY == plus Y DY,
2192	solid(NX,NY).	solid NX NY is tt.
2193	speedrun :- finish(X,Y), at(X,Y,T).	speedrun :- finish X Y, at X Y _.
2194	:- speedrun.	#forbid speedrun.

Fig. 20. Adam Smith's Map Generation Speedrun in Clingo's dialect of ASP, and an analogous (indirect) interpretation of the speedrun in Dusa.

D.2.2 Map-generation speedrun. Adam Smith's blog post "A Map Generation Speedrun with Answer Set Programming" is a beautiful example of the sculptural approach to ASP.³ On a tile grid, we stipulate that every tile can be solid or not, and we request maps where only a circuitous route through solid regions can reach the bottom-right corner of the map from the upper-left corner.

³<https://eis-blog.soe.ucsc.edu/2011/10/map-generation-speedrun/>




Fig. 21. Performance of DUSA, Clingo, and Alpha on the map generation programs in Figure 20. All data points are the median of three runs requesting five different maps where the shortest path from upper-left to lower-right on an $n \times n$ grid was more than $2n + 1$ steps long, discarding timeouts.

```

p(X1,X2,X3,X4,X5,X6) :-
    select(X1), select(X2), select(X3),
    select(X4), select(X5), select(X6).

select(X) :- dom(X), not nselect(X).
nselect(X) :- dom(X), not select(X).
:- not nselect(Y), select(X), dom(Y), X != Y.

```

Fig. 22. Ground explosion benchmark.

Figure 20 shows that ASP’s expressiveness translates well into finite-choice logic programming: note that this is an indirect interpretation, not a direct or mechanical translation. The fundamental choice — is a tile solid or not? — is naturally expressed as a closed choice rule.

The map generation speedrun is an excellent demonstration of DUSA’s current lack of suitability for sculptural programming: it times out for all grids larger than 10×10 , and times out at least once at every size, including 5×5 grids. We did not expect Alpha’s relatively poor showing: it was unable to handle any grids larger than 6×6 .

D.3 Alpha benchmarks

Using a benchmark someone else wrote is a good way to keep yourself honest. We are therefore grateful to Weinzirel [2017], who published an easily accessible set of benchmarks for the Alpha implementation of lazy grounding ASP,⁴ and all the answer set programs in this benchmark are amenable to direct, mechanical translation to finite-choice logic programs.

Our hypothesis was that this benchmark would be somewhat kind to DUSA, since lazy grounding ASP and the finite-choice logic program translation of ASP share many of the same advantages over the ground-then-solve approach. By in large, the results were similar to our other benchmarks: DUSA performs impressively well except on problems that obviously rely heavily on exponential backtracking search.

Weinzirel’s test suite has four sections, which we will discuss in turn. Unless otherwise indicated, we evaluate DUSA on the direct translation of the ASP programs that are shown.

D.3.1 Ground explosion. The ground explosion benchmarks use the program in Figure 22 to trigger pathological behavior in ground-then-solve ASP solvers. For a domain with n elements, a ground-then-solve solver will need to ground 6^n variants of the first rule, despite the fact that there are only $n + 1$ solutions.

⁴<http://www.kr.tuwien.ac.at/research/systems/alpha/benchmarks.html>




Fig. 23. Performance of the Dusa, Clingo, and Alpha solvers on the ground explosion benchmark in Figure 22. The x-axis is the problem size. All data points are the median of three runs, discarding timeouts.

```

chosenColour(N,C) :- node(N), colour(C),
    not notChosenColour(N,C).
notChosenColour(N,C) :- node(N), colour(C),
    not chosenColour(N,C).

:- node(X), not colored(X).
colored(X) :- chosenColour(X,Fv1).

:- node(N), chosenColour(N,C1),
    chosenColour(N,C2), C1!=C2.

colour(red0).
colour(green0).
colour(blue0).
colour(yellow0).
colour(cyan0).

:- node(X), not chosenColour(X,red0),
    not chosenColour(X,blue0),
    not chosenColour(X,yellow0),
    not chosenColour(X,green0),
    not chosenColour(X,cyan0).

:- link(X,Y), node(X), node(Y), X<Y,      #forbid link X Y,  X < Y,
    chosenColour(X,C), chosenColour(Y,C).    color X is C, color Y is C.

color N is {
    red,
    green,
    blue,
    yellow,
    cyan
} :- node N.

```

Fig. 24. Graph 5-colorability expressed in ASP, and an analogous (indirect) interpretation in Dusa.

The results from this set of Alpha benchmarks are shown in Figure 23. This benchmark shows both the strength of the Dusa implementation and the degree to which backtracking is done naively and inefficiently in Dusa. The first solution is found immediately, matching the asymptotic behavior of the Alpha implementation and improving on constant factors, but enumerating additional solutions is done in a way that results in exponential backtracking, so Dusa is unable to return 10 distinct solutions for most examples. We expect it to be relatively straightforward to dramatically improve Dusa's behavior here.

D.3.2 Graph 5-colorability. A graph is 5-colorable if each node can be assigned one of five different colors such that no edge connects nodes of the same color. Perhaps there is a dramatically clearer way to express this concept in ASP than the ASP code from the benchmark shown in Figure 24, but it is remarkable how much simpler this is to express in finite-choice logic programming, as shown on the right side of Figure 24.

This is essentially a pure search problem, so it's not surprising that the Dusa implementation performs quite badly. We didn't bother testing the performance of the finite-choice logic program




Fig. 25. Performance of the Dusa, Clingo, and Alpha solvers on 5-colorability benchmark from Figure 24. All data points are the **average** of ten runs on different graphs requesting 10 different solutions each time, discarding timeouts. (The same data points with 50 nodes and 200 edges are present in both graphs.)

```

delete(X,Y) :- edge(X,Y), not keep(X,Y).
keep(X,Y) :- edge(X,Y), delete(X1,Y1), X1 != X.
keep(X,Y) :- edge(X,Y), delete(X1,Y1), Y1 != Y.
reachable(X,Y) :- keep(X,Y).
reachable(X,Y) :- special(Y),reachable(X,Z),reachable(Z,Y).

```

Fig. 26. “Cutedge” benchmark.




Fig. 27. Performance of the Dusa and Alpha solvers on the “cutedge” benchmark from Figure 26 on random graphs with different edge probabilities. All data points are the **average** of ten runs on different graphs requesting 10 different solutions, discarding timeouts.

```

reachable(X, Y) :- edge(X, Y).
reachable(S, Y) :- start(S), reachable(S, X), reachable(X, Y).

```

Fig. 28. Graph reachability.

directly and mechanically translated from the ASP program in Figure 24, since the idiomatic program already exhibited extremely poor performance.

D.3.3 “Cutedge” benchmark. The “cutedge” benchmark shown in Figure 26 removes an arbitrary edge from a graph and then calculates reachability from one endpoint of the removed edge and another specially designated node. The benchmark seems to be intended to force bad performance on ground-then-solve approaches without being as obviously unfair as the ground explosion benchmark. Figure 27 shows the very similar performance of Dusa and Alpha. Clingo only able to handle the smallest examples in under 100 seconds.

D.3.4 Graph reachability benchmark. The graph reachability benchmark (Figure 28) is included for completeness, but the test cases and results weren’t particularly interesting. It would be more interesting to rerun this test on the graphs used to test the spanning tree and canonical representative programs in order to tease apart different asymptotic behaviors.

Instance size	Our implementation	Clingo	Alpha
1000/4	0.4	0.05	0.8
1000/8	0.5	0.1	0.9
10000/2	1.9	1.1	1.2
10000/4	3.6	2.3	1.6
10000/8	2.3	5.1	2.1

Fig. 29. Performance of the Dusa, Clingo, and Alpha solvers on the graph reachability program from Figure 28. Instance size is number of nodes / number of edges per node, running time in **seconds**, and all data points are the **average** of ten runs on different graphs requesting 10 different solutions each time (though there is in fact only one unique solution in each case).