

International Journal of Foundations of Computer Science  
 © World Scientific Publishing Company

## FPsolve: A Generic Solver for Fixpoint Equations over Semirings\*

Javier Esparza

Michael Luttenberger

Maximilian Schlund

*Department of Computer Science  
 Technische Universität München  
 Boltzmannstraße 3, 85748 Garching, Germany  
 {esparza,luttenbe,schlund}@in.tum.de*

Received (Day Month Year)

Accepted (Day Month Year)

Communicated by (xxxxxxxxxx)

We introduce FPSOLVE, an implementation of generic algorithms for solving fixpoint equations over semirings. We first illustrate the interest of generic solvers by means of a scenario. We then succinctly describe some of the algorithms implemented in the tool, and provide some implementation details.

*Keywords:* Keyword1; keyword2; keyword3.

### 1. Introduction

We present FPSOLVE<sup>a</sup>, a solver for *algebraic systems of equations* first introduced in [23]. These are systems of equations of the form

$$X_1 = f_1(X_1, X_2, \dots, X_n) \quad \cdots \quad X_n = f_n(X_1, X_2, \dots, X_n)$$

where  $f_1, \dots, f_n$  are polynomials in the variables  $X_1, X_2, \dots, X_n$ . The coefficients of the polynomials can be elements of any semiring satisfying some weak conditions, which ensure that there exists a unique smallest solution. FPSOLVE implements a number of *generic* algorithms, i.e. algorithms parametric in the semiring operations of addition and multiplication, plus possibly the Kleene star operation.

Algebraic systems naturally arise in various settings:

- The language of a context-free grammar like  $X \rightarrow aXX \mid b$  is the least solution of the equation  $X = aXX + b$  over the semiring whose elements are languages, with union and concatenation of languages as sum and product.

\*This work was funded by the DFG project “Polynomial Systems on Semirings: Foundations, Algorithms, Applications”

<sup>a</sup>Freely available from <https://github.com/mschlund/FPsolve>.

2 *Javier Esparza, Michael Luttenberger, and Maximilian Schlund*

- Shortest-paths problems on finite graphs and on some infinite graphs, like those generated by weighted pushdown automata, can be reduced to solving fixed-point equations over a semiring having the possible edge weights as elements [8, 22].
- Data-flow equations associated to many intra- and interprocedural dataflow analyses are fixed-point equations over complete lattices [18], which can often be recast as equations on semirings [22, 6].
- Authorization problems (like, for instance, the authorization problem for the SPKI/SDSI authorization system), can be recast as a reachability problem in weighted pushdown automata [17], and thus to algebraic systems [6].
- Computing the reputation of a principal in a reputation system (a system in which principals can recommend other principals, and rules are used to compute reputation out of a set of direct recommendations) reduces to solving an algebraic system [5].
- Evaluating a Datalog query can be reformulated as the problem of deciding whether a non-terminal of a context-free grammar is productive or not, and so it also amounts to solving a system of equations. Moreover, several problems concerning the computation of *provenance information*, an important research topic in database theory, reduces to solving an associated algebraic system over different semirings. [14]

The paper is structured as follows. Section 2 motivates by means of a scenario the interest of generic solvers for algebraic systems. Section 3 describes the basic algorithms and data structures used in FPSOLVE. Finally, Section 5 briefly describes the implementation.

## 2. Scenario: A Recommendation System

We succinctly describe SDSIREC, a recommendation system inspired by the SDSI authorization system [17], and very close to the reputation system described in [5].

SDSIREC distinguishes *customers* (denoted by  $x, y, z$ ) and *products* (denoted by  $p$ ). Given a collection of individual recommendations of products by customers, SDSIREC computes an aggregated customer rating for each product. Individual recommendations are described in SDSIREC by means of *rules* of the form:

$$x.\text{Rec} \xrightarrow{w} p \tag{1}$$

$$x.\text{Trust} \xrightarrow{w} y \tag{2}$$

The term  $x.\text{Rec}$  denotes the fuzzy set of all products recommended by customer  $x$ . The rule  $x.\text{Rec} \xrightarrow{w} p$  denotes that  $p$  belongs to  $x.\text{Rec}$  with weight  $w$ , i.e., that  $x$  recommends  $p$  with “rating”  $w$ . Analogously,  $x.\text{Trust}$  denotes the fuzzy set of all customers (whose recommendations are) trusted by  $x$ . The set of all weights, denoted by  $S$ , contains the special weight  $0$ , which explicitly states that  $p$  resp.  $y$  does not belong to  $x.\text{Rec}$  resp.  $x.\text{Trust}$ ; assigning a rule the weight  $0$  is equivalent to removing the rule from the input.

Besides direct recommendation and direct trust, SDSIREC also takes into account indirect recommendation of products via trust in other customers. For instance, consider the following scenario:

$$\begin{array}{l} \text{JESSE.Trust} \xrightarrow{w_1} \text{WALT} \\ \text{WALT.Rec} \xrightarrow{w_2} \text{FPSOLVE} \end{array}$$

Since WALT recommends FPSOLVE with weight  $w_2$ , and JESSE trusts the recommendations of WALT (his former high-school teacher) with weight  $w_1$ , SDSIREC infers that JESSE indirectly recommends FPSOLVE with some weight  $w_1 \odot w_2$ , where  $\odot$  abstracts from the concrete way how the weights should be combined into a new weight. The operator must satisfy  $0 \odot w = 0 = w \odot 0$ , so that the interpretation of  $0$  as a non-existing rule is preserved. The inference is modeled by the following (hard coded) rules:

$$x.\text{Trust} \xrightarrow{\lambda} x.\text{Trust.Trust} \tag{3}$$

$$x.\text{Rec} \xrightarrow{\mu} x.\text{Trust.Rec} \tag{4}$$

Rule (3) states that the set of customers trusted by  $x$  contains the set of customers trusted by customers trusted by  $x$ . Analogously, rule (4) states that the set of products recommended by  $x$  contains all products recommended by customers trusted by  $x$ . As these rules may lead to cycles, i.e.  $x$  might trust herself, thereby recommending to herself the products recommended by her, SDSIREC allows one to specify discount factors  $\lambda$  and  $\mu$  to dampen resp. penalize these effects. The special weight  $\mathbb{1}$  (required to satisfy  $w \odot \mathbb{1} = w = \mathbb{1} \odot w$ ) can be used to disable this discounting. SDSIREC then treats the rules (1) to (4) as rewrite rules in the sense of a pushdown system [22]. For instance, we get

$$\begin{array}{l} \text{JESSE.Rec} \xrightarrow{\mu} \text{JESSE.Trust.Rec} \xrightarrow{w_1} \text{WALT.Rec} \xrightarrow{w_2} \text{FPSOLVE} \\ \text{JESSE.Rec} \xrightarrow{\mu} \text{JESSE.Trust.Rec} \xrightarrow{\lambda} \text{JESSE.Trust.Trust.Rec} \xrightarrow{w_1} \text{WALT.Trust.Rec} \end{array}$$

The first “path” with weight  $\mu \odot w_1 \odot w_2$  captures that JESSE indirectly recommends FPSOLVE. The second path is an example of a path that cannot be extended to a recommendation of  $p$ : Since WALT trusts nobody (as specified by the input system), SDSIREC can never rewrite WALT.Trust to WALT.

In order to compute to what extent  $p$  belongs to  $x.\text{Rec}$  SDSIREC finally aggregates the weight of the (possibly infinitely many) paths leading from  $x.\text{Rec}$  to  $p$ . We use  $\oplus$  to denote the operator that is used to aggregate the weights of different paths. It is well-known that if  $\langle S, \oplus, \odot, 0, \mathbb{1} \rangle$  forms an  $\omega$ -continuous semiring, then the problem of aggregating over all possible paths can be recast as computing the least solution of an algebraic system (see below) [10, 22, 6]. Recall that  $\langle S, \oplus, \odot, 0, \mathbb{1} \rangle$  is a semiring if  $\oplus$  and  $\odot$  are associative and have neutral elements  $0$  and  $\mathbb{1}$ , respectively,  $\oplus$  is commutative,  $\odot$  distributes over  $\oplus$ , and any product with  $0$  as factor evaluates to  $0$ . Given  $a, b \in S$ , we say  $a \sqsubseteq b$  if there is  $c \in S$  such that  $a + c = b$ .

4 *Javier Esparza, Michael Luttenberger, and Maximilian Schlund*

A semiring is *naturally ordered* if the relation  $\sqsubseteq$  is a partial order. An  $\omega$ -continuous semiring is a naturally ordered semiring such that every  $\sqsubseteq$ -non-decreasing sequence ( $\omega$ -chain)  $c_0 \sqsubseteq c_1 \sqsubseteq c_2 \sqsubseteq \dots$  has a supremum in  $S$ , and both addition and multiplication are  $\omega$ -continuous in both arguments.<sup>b</sup> In an  $\omega$ -continuous semiring, the infinite sum  $\sum_{i \in \mathbb{N}} a_i$  is then defined as  $\sup^{\sqsubseteq} \{ \sum_{0 \leq i \leq k} a_i \mid k \in \mathbb{N} \}$ . This sum essentially shares the same properties as absolutely convergent series over the reals, in particular, summation is order independent, and multiplication also distributes over sum [19].

Let  $R_{xp}$  and  $T_{xy}$  be variables standing for the total weights with which  $x$  recommends  $p$  or trusts  $y$ , and let  $r_{xp}$ ,  $t_{xy}$  denote the weights of the direct recommendation, or the direct trust of  $x$  in  $p$  and  $y$ , respectively (i.e.  $x.\text{Rec} \xrightarrow{r_{xp}} p$  resp.  $x.\text{Trust} \xrightarrow{t_{xy}} y$ ).

If  $\langle S, \oplus, \odot, 0, 1 \rangle$  is an  $\omega$ -continuous semiring, then the total weights are the unique smallest solution w.r.t.  $\sqsubseteq$  of the following algebraic system (cf. [10, 22, 6])

$$\begin{aligned} R_{xp} &= r_{xp} \oplus \bigoplus_y \mu \odot T_{xy} \odot R_{yp} && \text{for all consumers } x, \text{ products } p \\ T_{xy} &= t_{xy} \oplus \bigoplus_z \lambda \odot T_{xz} \odot T_{zy} && \text{for all consumers } x, y \end{aligned}$$

The key point of our argumentation is that in an application like the above we are interested in solving the same set of equations over many different semirings. Even further, users of the system may be interested in first defining their own semiring, and then solving the system. To illustrate this, let us examine several different interpretations of “weight”, all of them very natural:

**Weights as scores.** The most natural interpretation of weights is perhaps as scores. A consumer  $x$  gives a product  $p$  or another consumer  $Y$  a score, corresponding to its degree of satisfaction with  $p$ , or its degree of trust in the recommendations of  $Y$ . If we assume that scores are real numbers in the interval  $[0, 1]$ , and choose  $\oplus$  and  $\otimes$  as sum and product of real numbers, we obtain the probabilistic semiring. Then  $R_{xp}$  represents the total weight of all recommendation paths leading from  $x$  to  $p$ . If we choose the Viterbi semiring  $\langle [0, 1], \max, \cdot, 0, 1 \rangle$  instead, then  $R_{xp}$  returns the weight of the strongest recommendation path.

**Weights as expire times.** Direct recommendations and trust, represented by rules of types (1) and (2,) can (and should) have an expire time. If we choose  $\oplus$  to be the maximum and  $\odot$  the minimum over the reals, then  $R_{xp}$  returns the earliest time at which all recommendation paths from  $x$  to  $p$  will have expired.

**Weights as provenance information.** If a system user does not trust some consumers, she may wish to compute, for each recommendation path from  $x$  to  $p$ , the

<sup>b</sup>A function  $f: S \rightarrow S$  is  $\omega$ -continuous if  $f(\sup_{i \in \mathbb{N}}^{\sqsubseteq} c_i) = \sup_{i \in \mathbb{N}}^{\sqsubseteq} f(c_i)$  for every  $\omega$ -chain  $(c_i)_{i \in \mathbb{N}}$ .

set of consumers in the path. Or she may want to know the set of consumers visited along the recommendation path of maximal weight. Such provenance information can be computed within the semiring framework. For this it is convenient to treat all non-zero parameters  $r_{xp}, t_{xy}, \lambda, \mu$  as formal parameters (free variables).

- To compute for each path the set of consumers involved, one can use the Why-semiring, well-known in provenance theory. Semiring elements are sets of sets of consumers. We set  $r_{xp} = \{\{x\}\}$  and  $t_{xy} = \{\{x, y\}\}$  (and treat  $\lambda$  as  $\{\{\lambda\}\}$  and  $\mu$  as  $\{\{\mu\}\}$ ), and define:

$$\begin{aligned} & - \{X_1, \dots, X_n\} \odot \{Y_1, \dots, Y_m\} := \{X_1 \cup Y_1, x_1 \cup Y_2, \dots, X_n \cup Y_m\}, \text{ and} \\ & - \{X_1, \dots, X_n\} \oplus \{Y_1, \dots, Y_m\} := \{X_1, \dots, X_n, Y_1, \dots, Y_m\}. \end{aligned}$$

- If we wish to compute the provenance of the recommendation of maximal weight, we can use the following semiring: as semiring elements we choose the pairs  $(\alpha, X)$ , where  $\alpha \in [0, 1]$  and  $X$  is a set of consumers. We set  $t_{xy} = (w, \{x, y\})$  with  $w \in (0, 1]$ , and analogously for  $r_{xp}$ . The abstract operators are instantiated as follows:

$$\begin{aligned} & - (\alpha_1, X_1) \odot (\alpha_2, X_2) := (\alpha_1 \alpha_2, X_1 \cup X_2), \text{ and} \\ & - (\alpha_1, X_1) \oplus (\alpha_2, X_2) := (\max\{\alpha_1, \alpha_2\}, \text{if } \alpha_1 \geq \alpha_2 \text{ then } X_1 \text{ else } X_2) \end{aligned}$$

These examples show that, instead of creating new tools for each new semiring, it can be better to implement a generic tool, with generic algorithms applicable to any semiring, or at least to any semiring in a broad class.

### 3. Algorithms and Data Structures

The two main generic schemes implemented in FPSOLVE for the approximation (and sometimes exact computation) of the least solution of an algebraic system are classical fixpoint iteration and Newton’s method. Following [21, 11], we introduce them as procedures that “unfold” the algebraic system up to a certain depth which allows one both to unify and at the same time simplify their presentation.

**Classical fixed-point iteration.** Given an algebraic system of equations  $X = F(X)$  over an  $\omega$ -continuous semiring, Kleene’s theorem states that the system has a unique least solution  $\mu F$  with respect to the natural order  $\sqsubseteq$ , and that  $\mu F$  is the supremum of the sequence  $F(0), F^2(0), \dots, F^i(0)$  [19]. So  $\mu F$  can be approximated by computing successive elements of the sequence. If the semiring further satisfies the *ascending chain property* (for every  $\omega$ -chain  $a_1 \sqsubseteq a_2 \sqsubseteq \dots$  eventually  $a_k = a_{k+1} = a_{k+2} = \dots$ ) then  $\mu F = F^i(0)$  for some  $i \geq 0$ , and so  $\mu F$  can be effectively computed.

As explained in e.g. [21], an algebraic system can be associated a context-free grammar. For instance, for the system

$$X = a \odot X \odot X \oplus b \tag{5}$$

6 *Javier Esparza, Michael Luttenberger, and Maximilian Schlund*

we obtain the grammar

$$X \rightarrow aXX \mid b \tag{6}$$

Conversely, we assign to a derivation tree of the grammar a value in the semiring, given by the product of its leaves (the ordered product if the semiring is not commutative); further, we assign to a set of derivation trees the sum of the values of its elements. The following result can be proved by a simple induction on  $k$ :

$F^k(0)$  is the sum over the set of all derivation trees of the grammar of height less than  $k$ .

Building on this observation, for every  $k$  we “unfold” (5) into an acyclic system over variables  $X^{<h}$  and  $X^{=h}$  for every  $h \leq k$ , such that the solutions of  $X^{<h}$  and  $X^{=h}$  are the values of the derivation trees of height *less* than  $h$  and *equal* to  $h$ , respectively. For this, we obviously have to set

$$X^{<0} = 0 \quad X^{=0} = b \quad \text{and} \quad X^{<h+1} = X^{<h} \oplus X^{=h} \quad \text{for all } h \in \mathbb{N} \tag{7}$$

In order to obtain the defining equation of  $X^{=h}$  we observe that the trees of height  $h$  can be partitioned into those whose left subtree has height  $h-1$ , and those whose left subtree has height strictly smaller than  $h-1$  *and* whose right subtree has height  $h-1$ . This leads to

$$X^{=h} = a \odot X^{<h} \odot X^{=h-1} \oplus a \odot X^{<h-1} \odot X^{<h-1}. \tag{8}$$

We can see the unfolding up to depth  $k$  as a symbolic representation of  $F^k(0)$ , which implicitly uses subterm sharing (*arithmetic circuit*). When the coefficients of the algebraic system ( $a, b$  in our example) are formal parameters, we can efficiently compute  $F^k(0)$  for different values  $a, b$ , by just plugging them into the unfolded system.

**Newton’s method.** Newton’s method for arbitrary  $\omega$ -continuous semirings, as described in [10], can be much faster than Kleene iteration. It is shown in [11] that the method can also be presented as an unfolding of the algebraic system: This time, the system is unfolded w.r.t. the *Strahler number* or *dimension* of its associated derivation trees (see [11, 21]). The dimension of a rooted tree  $t$  is defined as the height of the largest perfect binary tree that is a minor of  $t$ .

Consider again equation (5). We split  $X$  into a family of variables  $X^{<d}$  and  $X^{=d}$  for  $d \in \mathbb{N}$ . The solutions of the unfolded system for  $X^{<d}$  and  $X^{=d}$  will now be the value of the derivation trees of dimension *less* than  $d$ , and *equal* to  $d$ , respectively. Just as before, we have

$$X^{<0} = 0 \quad X^{=0} = b \quad \text{and} \quad X^{<d+1} = X^{<d} \oplus X^{=d} \quad \text{for all } d \in \mathbb{N}. \tag{9}$$

In order to derive the defining equation of  $X^{=d}$ , observe that there are three possible cases for a tree of dimension  $d$ : either the left subtree has dimension  $d$ , and the right subtree has dimension at most  $d-1$ ; or vice versa; or both subtrees have dimension

exactly  $d-1$  (this is the case in which the root of the minor coincides with the root of the tree). So we get

$$X^{=d} = a \odot X^{=d} \odot X^{<d} \oplus a \odot X^{<d} \odot X^{=d} \oplus a \odot X^{=d-1} \odot X^{=d-1}. \quad (10)$$

However, this unfolding does not represent an arithmetic circuit as it is not yet acyclic:  $X^{=d}$  appears on both sides of the equation. But equation (10) is linear in  $X^{=d}$ , and so, if multiplication is commutative, we can replace it by (with  $\mathbb{1} \oplus \mathbb{1} = 2$ )

$$X^{=d} = 2 \odot a \odot X^{=d} \odot X^{<d} \oplus a \odot X^{=d-1} \odot X^{=d-1} \quad (11)$$

and use Kleene's theorem [19] to replace it by

$$X^{=d} = (2 \odot a X^{<d})^* \odot a \odot (X^{=d-1})^2 \quad (12)$$

where the Kleene star is defined, as usual, by  $x^* := \sum_{k \in \mathbb{N}} x^k$  (and is well defined for any  $\omega$ -continuous semiring).<sup>c</sup> The new system is acyclic, i.e. an arithmetic circuit w.r.t.  $\oplus$ ,  $\odot$ , and  $*$ , and as in the previous case, can be used as a compact symbolic representation of the  $d$ -th Newton approximation, useful when  $a, b$  are formal parameters (see Fig. 1 for an example). In particular, every Newton approximation can be represented by means of a rational expression.

To actually compute the solution for particular values of  $a$  and  $b$ , we can then use straight-forward constant propagation going from bottom ( $X^{=0}, X^{<0}$ ) to top ( $X^{<d}$ ). However, for this the Kleene star  $x^*$  must be effectively computable in the given semiring representation. This is indeed the case for several important semirings. The simplest example is the probability semiring, where for every rational number  $x \in [0, 1)$  we have  $x^* = 1/(1-x)$ . Tropical semirings are another example. For instance, over the integers extended by least ( $-\infty$ ) and greatest element ( $+\infty$ ), with addition given by  $\min$  and multiplication given by  $+$  (on  $\mathbb{Z}$ ), we have  $x^* = 0$  if  $x \geq 0$  and  $x = -\infty$  otherwise. A third example is the semiring of semilinear sets of vectors with components in  $\mathbb{N} \cup \{\infty\}$ , with  $X \oplus Y := X \cup Y$ ,  $X \odot Y := \{x+y \mid x \in X, y \in Y\}$ , and  $X^* := \sum_{i \geq 0} X^i$ . Note that these definitions are not effective. We define all operations on semilinear sets effectively in Section 4.

**Connection to Newton's method over the reals.** Applying Newton's method to  $g(X) := f(X) - X = aX^2 - X + b$  (interpreted over the reals) starting from the initial approximation  $X = 0$  we obtain the sequence:

$$X_0 := 0 \quad X_{d+1} = X_d - \frac{g(X_d)}{g'(X_d)} = X_d - \frac{aX_d^2 - X_d + b}{2aX_d - 1} = X_d + \frac{aX_d^2 - X_d + b}{1 - 2aX_d}.$$

Setting  $Y_d := X_{d+1} - X_d$  this can be written as

$$Y_d = 2aX_dY_d + (aX_d^2 - X_d + b).$$

<sup>c</sup>In the noncommutative case, one may resort to an instance of the semiring of contexts in order to obtain a rational tree expression.

8 *Javier Esparza, Michael Luttenberger, and Maximilian Schlund*

Straight-forward induction now shows that over the nonnegative reals the values of  $X_d$  and  $X^{<d}$  resp.  $Y_d$  and  $X^{=d}$  coincide [11]. In particular, the defining equation of  $X^{=d}$  can be seen as the generalization of the derivative of  $aX^2$  in the noncommutative case.

**Multivariate case.** Both unfoldings immediately generalize to the setting of multiple variables  $X, Y, \dots$ . As mentioned above, in the univariate case we use the fact that the solution of an equation  $X = aX + b$  is given by  $a^*b$ . In the multivariate case, when the semiring is commutative, we have to deal with systems of linear equations  $X = AX + B$  for a matrix  $A$  and a vector  $B$  over the semiring. It is well known that the solution is given by  $A^*B$ , where  $A^* = \sum_{i=0}^{\infty} A^i$ , and matrix multiplication is defined as for the natural or the real numbers, but replacing sum and product by the operations of the semiring being considered. In the next section, we describe the two algorithms for computing  $A^*$  implemented in FPSOLVE.

### 3.1. Solving Linear Equations

As mentioned above, solving a linear equation amounts to computing  $A^*$  for a given square matrix  $A$  over a semiring. Given a matrix  $A$ , the algorithm returns a matrix whose elements are semiring expressions over the semiring operations and the Kleene star. So, intuitively, the algorithms reduce the problem of computing the star of a matrix to computing the star of semiring elements.

FPSOLVE implements both the well-known Floyd-Warshall algorithm, and the recursive divide-and-conquer approach.

#### 3.1.1. Generalized Floyd-Warshall

The Floyd-Warshall algorithm (see Algorithm 1) for solving the all-pairs-shortest-path problem in weighted (finite) graphs carries directly over to the setting of generic  $\omega$ -continuous semirings, even if addition  $\oplus$  is not idempotent (cf. [8]). (In fact, it suffices that the semiring  $\langle S, \oplus, \odot, \mathbb{0}, \mathbb{1} \rangle$  is closed but not necessarily  $\omega$ -continuous.) From the description of the algorithm it is easy to count that the total number of semiring operations (i.e.  $+, \cdot, *$ ) needed is  $T(n) = 2n^3 - 2n^2 + n \in \Theta(n^3)$ .

#### 3.1.2. Divide-and-conquer

This algorithm recursively applies the formula for computing the Kleene star of a  $2 \times 2$ -matrix:

$$\mathbf{M} = \begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{bmatrix} \quad \mathbf{M}^* = \begin{bmatrix} \mathbf{F} & \alpha \mathbf{G}^* \\ \mathbf{G}^* \beta & \mathbf{G}^* \end{bmatrix} \quad \text{with} \quad \begin{array}{l} \alpha = \mathbf{A}^* \mathbf{B} \\ \beta = \mathbf{C} \mathbf{A}^* \\ \mathbf{G} = \mathbf{D} + \mathbf{C} \alpha \\ \mathbf{F} = \alpha \mathbf{G}^* \beta + \mathbf{A}^* \end{array} .$$

Given a  $n \times n$ -matrix  $\mathbf{M}$  ( $n > 2$ ), the entries  $\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}$  become submatrices of  $\mathbf{M}$  to which the algorithm is then applied recursively; the recursion stops when either



---

**Algorithm 1:** Generalized Floyd-Warshall algorithm over semirings.

---

**input** : Matrix  $\mathbf{A} \in S^{n \times n}$  over a semiring  $S$ .  
**output**: Reflexive-transitive closure  $\mathbf{A}^*$ .

```

B := A
for k = 1 .. n do
  Bk,k := Bk,k*
  for i = 1 .. n, i ≠ k do
    Bi,k = Bi,k ⊙ Bk,k
    for j = 1 .. n, j ≠ k do
      Bi,j := Bi,j ⊕ Bi,k ⊙ Bk,j
    end
  end
  for j = 1 .. n, j ≠ k do
    Bk,j = Bk,k ⊙ Bk,j
  end
end
return B

```

---

$n = 2$  or  $n = 1$ . A formal proof of correctness (for any *Conway* semiring) goes back to Ésik and Kuich (cf. [8]).

Altogether we need two recursive calls, six matrix multiplications (the term  $\alpha\mathbf{G}^*$  appears twice and thus needs to be evaluated only once), and two matrix additions. Hence, the number of operations needed by this algorithm can be expressed by the recurrence relation <sup>d</sup>

$$\begin{aligned}
 T(n) &= 2T\left(\frac{n}{2}\right) + 6 \left[ 2 \left(\frac{n}{2}\right)^3 - \left(\frac{n}{2}\right)^2 \right] + 2 \left(\frac{n}{2}\right)^2 \\
 &= 2T\left(\frac{n}{2}\right) + \frac{3}{2}n^3 - n^2
 \end{aligned}$$

Which can be solved exactly (setting  $T(n) = 1$  and  $n = 2^l$ ) resulting in  $T(n) = 2n^3 - 2n^2 + n \in \Theta(n^3)$ . Hence this algorithm uses the same number of operations as Floyd-Warshall. Both algorithms need  $n^3 - 2n^2 + n$  additions,  $n^3 - n$  multiplications, and  $n$  Kleene stars.

### 3.1.3. Symbolic solving

Recall our initial example

$$X = a \odot X \odot X \oplus b$$

<sup>d</sup>Note that multiplying two  $n \times n$  matrices requires  $n^3 - n^2$  operations (via the schoolbook method – we cannot use e.g. Strassen’s algorithm as we lack a difference operator!).

10 *Javier Esparza, Michael Luttenberger, and Maximilian Schlund*

and its unfolding w.r.t. dimension for an arbitrary  $d \in \mathbb{N}$  (assuming  $\odot$  is commutative)

$$X^{=d} = 2 \odot a \odot X^{=d} \odot X^{<d} \oplus a \odot X^{=d-1} \odot X^{=d-1}$$

As  $X^{<d} = X^{<d-1} \oplus X^{=d-1}$ , every iteration of Newton’s method essentially consists of solving this linear equation after substituting for the variables  $X^{<d}$  and  $X^{=d-1}$  the already computed solution. Analogously, in the multivariate setting essentially the same linear equation system has to be solved over and over again. FPSOLVE thus can first compute a symbolic solution of the linear system by treating  $X^{<d}$  and  $X^{=d-1}$  as formal parameters which allows us to obtain a succinct symbolic representation of a Newton approximation by sharing common subexpressions. This can be used to efficiently evaluate a Newton approximation of an algebraic system for several different semiring interpretations.

Consider the generic linear equation system

$$\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \cdot \begin{pmatrix} x \\ y \end{pmatrix} + \begin{pmatrix} e \\ f \end{pmatrix}$$

Treating  $a, \dots, f$  as formal parameters over some semiring, the (symbolic) solution of this system is given by

$$\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} a^*b(ca^*b \oplus d)^*ca^*e \oplus a^*b(ca^*b \oplus d)^*f \oplus a^*e \\ (ca^*b \oplus d)^*ca^*e \oplus (ca^*b \oplus d)^*f \end{pmatrix}$$

where we have omitted the  $\odot$  for readability. The internal representation of these terms is shown in Fig. 1: FPSOLVE stores the expressions as part of an “abstract syntax DAG” (reversing the direction of the edges we obtain an arithmetic circuit with gates for addition, multiplication, and Kleene star) similar to BDD libraries like CUDD [24], where we have colored the  $x$ - resp.  $y$ -component in light resp. dark grey; this representation allows one to reduce both the memory consumption, and the re-evaluation of identical subterms.

In this simple  $2 \times 2$  case the concrete recursive approach (as stated above) computes 10 semiring operations, the same if the symbolic solution is computed (using the same recursive algorithm). This holds in general if all elements of the input matrix are different. However, in general input matrices can have the same element in many different positions, then even the recursive algorithm will compute some identical subexpressions multiple times (that occur in different execution branches) since it cannot guess them a priori. In this case, symbolic solving allows for a *global* subexpression detection after the whole matrix-star has been computed.

Although the symbolic approach significantly reduces the number of semiring operations needed, the overhead from computing and storing the symbolic solution is not always negligible. This is particularly true for numeric semirings (like the semiring of positive reals) that are implemented using machine precision floating point numbers – for these the semiring operations are so fast that the overhead outweighs the benefits of symbolic solving.




Fig. 1. Succinctly representing all terms of the product  $A^* \cdot (e, f)^T$  via a BDD-like sharing of subexpressions. By reversing the direction of all edges this can be read as an arithmetic circuit with output gates colored in grey.

We therefore give the user the freedom of choice whether to use the concrete (i.e. in every iteration) or symbolic (i.e. solve once then plug in in every iteration) method of solving linear equations.

### 3.2. Decomposition into strongly-connected components

To efficiently process large algebraic systems, FPSOLVE supports a decomposition of the system into strongly connected components (SCCs). To make this precise, recall the definition of *dependency graph*: Its nodes are the variables occurring in the algebraic system and its edges are induced by the defining equations: we have an edge from variable  $X$  to variable  $Y$  if  $Y$  occurs in the defining equation of  $X$ . We say that  $X$  *depends* on  $Y$  if there is a path from  $X$  to  $Y$  in the dependency graph. To determine the value of variable  $X$  it then suffices to determine the values of all variables on which  $X$  depends. Using Tarjan's algorithm we therefore partition

12 *Javier Esparza, Michael Luttenberger, and Maximilian Schlund*

the dependency graph into SCCs, and process these SCCs in reverse topological ordering (“bottom up”). In particular when using Newton’s method this can lead to a noticeable speed-up in the computation of the Kleene star.

#### 4. Representing Semirings – Case Study: Semilinear sets

So far, we have described the part of FPSOLVE independent of the actual semiring. FPSOLVE comes with support for several standard semirings e.g. the extended reals  $([0, \infty], +, \cdot, 0, 1)$ , the tropical semiring  $(\mathbb{N}_\infty, \min, +, \infty, 0)$ , or the why-semiring used in database provenance (see Section 2). In this section we will focus on the implementation of the *counting semiring*, more precisely that of its subsemiring of *semilinear* (also: rational) sets [13].

##### 4.1. The counting semiring

Let  $X, Y \subseteq \mathbb{N}^k$  for some  $k \in \mathbb{N}$ , and define

$$X + Y := X \cup Y \quad \text{and} \quad X \cdot Y := \{x + y \mid x \in X, y \in Y\}.$$

Further, set  $\mathbf{0} := \emptyset$ , and  $\mathbf{1} := \{\vec{0}\}$  for  $\vec{0} \in \mathbb{N}^k$  the null vector. Then  $\mathbf{C}_k = \langle 2^{\mathbb{N}^k}, +, \cdot, \mathbf{0}, \mathbf{1} \rangle$  is called the *counting semiring* (of dimension  $k$ ). This semiring is commutative and idempotent. It is also  $\omega$ -continuous, and the Kleene star is given by  $X^* = \bigcup_{m \in \mathbb{N}} X^m$ . Because of commutativity and idempotence, the Kleene star satisfies the following two additional equalities:

$$(XY^*)^* = \mathbf{1} + XX^*Y^* \tag{13}$$

$$(X + Y)^* = X^*Y^* \tag{14}$$

The counting semiring  $\mathbf{C}_k$  can be characterized in several ways: it is (isomorphic to) the semiring  $\mathbb{N}_1 \llbracket y_1, \dots, y_k \rrbracket = \mathbb{N}_1 \langle \langle \{y_1, \dots, y_k\}^\oplus \rangle \rangle$  of formal power series in commuting variables  $\{y_1, \dots, y_k\}$  with coefficients in  $\mathbb{N}_1 = \{0, 1\}$ ; it is also (isomorphic to) the commutative and idempotent semiring freely generated by  $k$  elements.

Rational expressions on  $\mathbf{C}_k$  are constructed as usual: the atomic terms are the singleton sets of  $\mathbb{N}^k$  (resp. monomials w.r.t.  $\mathbb{N}_1 \llbracket y_1, \dots, y_k \rrbracket$ ); any further terms are obtained using the operators  $+$ ,  $\cdot$ , and  $^*$ . A rational expression of the form  $x_0 x_1^* \dots x_l^*$  is called *linear* if each  $x_i$  denotes a singleton set.  $S \subseteq \mathbb{N}^k$  is called *semilinear* (*linear*) if it is definable by a (linear) rational expression [13].

By virtue of above equalities Equations 13 and 14, any rational expression can be rewritten to a finite sum of linear expressions by “pushing the Kleene star downward in the term structure”. That is, any semilinear set is a finite sum (union) of linear sets; in particular, the semilinear sets form a non- $\omega$ -continuous subsemiring of  $\mathbf{C}_k$  which is closed under the Kleene star.

When solving algebraic systems over the counting semiring (or any other commutative and idempotent semiring) using Newton’s method, all Newton approximations are by definition always rational. Furthermore, Newton’s method computes

the solution of the system in a finite number of steps. More precisely,  $n + 1$  steps of the method suffice to compute the solution of a system of  $n$  equations [9]. Hence, for the implementation of Newton's method we can concentrate on the subsemiring of semilinear sets. In the rest of the section we discuss two representations of the semilinear sets, and their advantages and disadvantages.

#### 4.2. Constants-periods representation

Given a finite subset  $P = \{p_1, \dots, p_n\}$  of  $\mathbb{N}_k$  and  $c \in \mathbb{N}^k$ , we use  $(c, P)$  as a representation for  $cp_1^* \dots p_n^*$ , i.e., for the set  $\{c + \sum_{p \in P} n_p \cdot p \mid \forall p \in P: n_p \in \mathbb{N}\}$ . Accordingly,  $P$  is called the set of *periods*, and  $c$  the *constant* of the set.

This representation allows for a straightforward computation of both multiplication and the Kleene star:

$$\begin{aligned} (c_1, P_1) \cdot (c_2, P_2) &= (c_1 + c_2, P_1 \cup P_2) \\ (c, P)^* &= \mathbb{1} + (c, P \cup \{c\}) \end{aligned}$$

where the latter equation follows easily from (13). Analogously, we represent a semilinear set  $S$  as a set of linear sets. Given two representations  $S_1 = \{(c_1, P_1), \dots, (c_l, P_l)\}$  and  $S_2 = \{(c'_1, P'_1), \dots, (c'_m, P'_m)\}$  of semilinear sets we then have, using (14):

$$\begin{aligned} S_1 + S_2 &:= S_1 \cup S_2 \\ S_1 \cdot S_2 &:= \{(c_i + c'_j, P_i \cup P'_j) : i \in [l], j \in [m]\} \\ S_1^* &:= (c_1, P_1)^* \cdot \dots \cdot (c_l, P_l)^*. \end{aligned}$$

While this representation is simple and appealing, it also has several disadvantages:

- In the course of computing the Newton approximations, linear sets often share common period vectors and semilinear sets share linear sets. However, the representation does not profit from this to reduce the space necessary to store a set.
- The multiplication and Kleene star operations have quadratic and exponential complexity, respectively, in the number of linear sets of the operands.
- Checking inclusion or equivalence of representations is hard ( $\Pi_p^2$ -complete [16]).

In [23] we mitigate the first two problems by using extensive sharing of common (sets of) vectors and by eliminating redundant linear sets or periods after every semiring operation. These optimizations enable us to solve small systems of equations in a few milliseconds (less than 20 equations and 3-4 constant symbols). As an illustration we have built a simple inequivalence-checker for context-free grammars that computes their commutative image and checks it for equivalence. The equivalence-check for the computed semilinear sets is performed by converting them to their (unique) NDD representation (see the next section). The main bottleneck for this application is the size of the set of terminals of the grammar.

TODO: Table of runtimes here for the cfg-analyzer-benchmarks

### 4.3. NDD representation

In order to address the third drawback (high complexity of checking inclusion and equivalence), FPSOLVE also allows one to represent semilinear sets by *number decision diagrams (NDDs)* [7, 26]. NDDs are minimal deterministic finite automata that recognize sets of vectors of natural numbers. A vector of  $\mathbb{N}^n$  is converted into a string by first encoding each of its components in binary—using either a least significant or a most significant bit first encoding—and then converting the resulting  $n$  binary strings into one single string, for which there are also several possibilities. NDDs have numerous applications, e.g. in decision procedures for Presburger arithmetic and MSO, or analysis of infinite state systems, and there exist several libraries available for their manipulation, like GENEPI, MONA, and LASH [15, 12].

In the encodings used by these libraries, semilinear sets of vectors is converted into regular languages, and so a semilinear set can be canonically represented by its unique NDD. Canonicity follows immediately from the uniqueness of the minimal deterministic automaton for a regular language.

Although the NDD representation can be just as prohibitively space consuming as the constants-periods representation, NDDs often have smaller size; further, all operations and queries such as equivalence of two semilinear sets, membership of a vector in a set, intersection of sets, etc. can be performed in polynomial time in the size of the NDD.

While there are algorithms to extract an explicit representation of a semilinear set from an NDD (under the promise that the NDD indeed represents a semilinear set, since NDDs can represent a larger class of sets), they are rather complicated and mostly of theoretical interest [20]. Conversions between representations should therefore be avoided, and so we define the semiring operations on NDDs directly. For this, let  $\mathcal{A}_1, \mathcal{A}_2$  be two NDDs representing the semilinear sets  $S_1, S_2 \subseteq \mathbb{N}^k$ .

Computing  $S_1 + S_2$  simply translates to applying the standard product construction to  $\mathcal{A}_1$  and  $\mathcal{A}_2$ , followed by minimization of the result.

The algorithm for computing  $S_1 \cdot S_2$  is as follows:

- build an intermediate NDD  $I$  recognizing the semilinear set  $\{(u, v, u + v) \in \mathbb{N}^{3k} \mid u, v \in \mathbb{N}^k\}$ ;  
(actually,  $I$  does not have to be explicitly constructed, it can be kept implicit)
- build NDDs  $\mathcal{A}'_1, \mathcal{A}'_2$  for the sets  $S_1 \times \mathbb{N}^{2k}$  and  $\mathbb{N}^k \times S_2 \times \mathbb{N}^k$ , respectively;
- build an NDD  $\mathcal{A}$  for the intersection of  $I, \mathcal{A}'_1, \mathcal{A}'_2$ ;
- build an NDD for the projection of the vectors recognized by  $\mathcal{A}$  onto their last  $k$  components, followed by determinization and minimization.

Computing the Kleene star of a semilinear set represented as an NDD is more challenging, since it is not obvious how to define it by a Presburger formula. We

propose a novel algorithm. To get the idea of our construction it is useful to look at the representation of semilinear sets by rational expressions. For example, the semilinear set

$$S = \left\{ \left( (2, 1), \{(2, 0), (1, 1)\} \right), \left( (1, 3), \{(0, 3), (2, 3)\} \right) \right\}$$

can be equivalently represented as

$$r = (y_1^2 y_2)(y_1^2)^*(y_1 y_2)^* + (y_1 y_2^3)(y_2^3)^*(y_1^2 y_2^3)^*.$$

Suppose a semilinear set  $S$  is described by an expression  $r = c_1 p_1^* + c_2 p_2^*$ . Then  $S^*$  can be expressed (modulo commutativity and idempotence) as

$$\begin{aligned} r^* &= (c_1 p_1^* + c_2 p_2^*)^* \\ &= (c_1 p_1^*)^* \cdot (c_2 p_2^*)^* \\ &= (1 + c_1 c_1^* p_1^*)(1 + c_2 c_2^* p_2^*) \\ &= 1 + c_1 p_1^* c_1^* + c_2 p_2^* c_2^* + c_1 c_2 p_2^* p_1^* (c_1^* c_2^*) \\ &= 1 + (r + r^2)(c_1^* c_2^*) \end{aligned}$$

where the last equality holds because of idempotence and since  $c_1 \in S$  implies  $c_1^* \in S^*$ . Note that the constants of the semilinear set are simply added to the set of periods. For a semilinear set  $S$  consisting of  $n$  linear sets with constants  $C = \{c_1, \dots, c_n\}$  we obtain by induction

$$S^* = \mathbb{1} + \sum_{i=1}^n S^i \cdot C^*$$

which, by idempotence, can be rewritten to

$$S^* = \mathbb{1} + S \cdot (\mathbb{1} + S)^{n-1} \cdot C^* \tag{15}$$

so that only  $\lceil \log_2 n \rceil$  many multiplications are required to compute  $S^*$ .

Equation 15 immediately translate into an NDD-based algorithm, *as long as we know the cset  $C$  of constants  $C$* . However, for the implementation of Newton's method this is not a problem: Since all NDDs required to compute the Newton approximations are built from smaller ones, we keep track of the set of constants  $C$  for every semilinear set. More precisely, we represent a semilinear set by two NDDs: a first NDD recognizing the set itself, and a second NDD recognizing its set of constants. At first this seems very inefficient, but since any set  $B$  of constants satisfying  $B^* = C^*$  suffices to compute the Kleene star, in fact it suffices to store an NDD recognizing a *minimal* set  $B$ .

Despite these improvements, the practical performance of the NDD representation is still rather poor. Currently, the main bottleneck is the multiplication of two NDDs, and future work will focus on it. But even this implementation is already very useful for checking membership or inclusion. We implemented a translation from the constant-period representation to the NDD representation, which we use whenever equivalence or inclusion must be checked.

## 5. Implementation

Currently, FPSOLVE comprises roughly 8,000 lines of C++. The code can be obtained freely from <https://github.com/mschlund/FPSolve>. We use several existing frameworks and libraries:

- CPPUNIT for writing unit-tests.
- BOOST for IO-tasks (parsing, command-line arguments).
- GENEPI, MONA, and LASH for representing semilinear sets via NDDs.
- LIBFA for representing elements of “lossy” semirings (i.e. semirings satisfying  $1 \sqsubseteq a$  for any semiring element  $a \neq 0$  – this generalizes the downward closure of languages) as finite automata.

FPSOLVE features data structures for commutative as well as non-commutative polynomials, different solvers (semi-naive fixpoint iteration, Newton’s method), and several predefined semirings (semilinear sets, real numbers, tropical and boolean semiring) as well as some generic constructions (via C++ templates) to build new semirings from existing ones, like the direct product of two semirings, or the semiring of matrices over some semiring.

The focus of our library is to provide generic algorithms and to be easily extensible. One of our goals was to make it easy for users to write their own semiring-constructions or tailor the generic solving algorithms to their needs.

The library consists of three main parts:

- Data structures (polynomials, matrices, BDD-like DAG-structure to support subterm sharing)
- Semirings (semilinear sets, positive real numbers, why semiring, generic product semiring, ...)
- Solvers (Kleene solver, Newton solver)

Figure 2 shows a simplified view of the main structure of our library. First, observe that many classes are templated (e.g. `Commutative_Polynomial` depends on an abstract `Semiring` given as template parameter). This produces efficient code due to compile-time polymorphism. The heart of the library is the generic `NewtonSolver` which depends on a `Semiring`, a `Polynomial` (either commutative, or non-commutative), and an algorithm to solve linear equations. All these are again given as template parameters to the `NewtonSolver`.

### 5.1. Invocation of the Standalone Solver

FPSOLVE also includes a callable solver and a parser for equation systems that demonstrates the use of the library.

To apply the standalone solver, one has to describe the algebraic system as a BNF-style context-free grammar. Variables of the system are enclosed in angle brackets, multiplication is not explicitly written, the addition  $x + y$  is written as



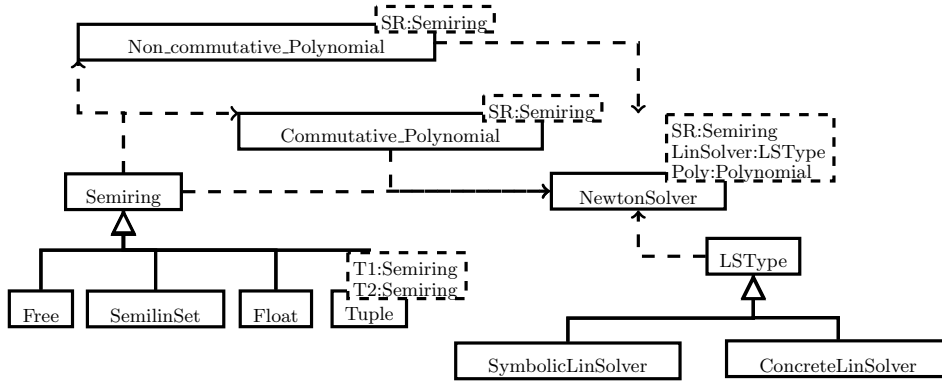


Fig. 2. A (simplified) part of our architecture.

$x \mid y$ . To solve the following system over the reals

$$X = 0.5XY + 0.5 \quad Y = 0.3Y + 0.7X$$

we would create a text file `test.g` containing:

```
<X> ::= 0.5 <X> <Y> | 0.5;
<Y> ::= 0.3 <Y> | 0.7 <X>;
```

The simplest invocation of the tool is then

```
$ ./fpsolve -f test.g --float
```

This minimal set of parameters specifies

- (1) the input file containing the algebraic system (`-f test.g`).
- (2) the semiring over which the system and its constants (like 0.3) are to be interpreted (here `--float`).

The tool outputs:

```
$ ./fpsolve -f test.g --float
Newton Concrete
Iterations: 3
Solving time: 0 ms (196 us)
X == 0.875
Y == 0.875
```

By default, the number of Newton iterations for a system of  $n$  equations is  $n + 1$  – for commutative, idempotent semirings this suffices to compute the exact solution [9].

A more sophisticated use of the tool’s options would be the following:

18 *Javier Esparza, Michael Luttenberger, and Maximilian Schlund*

```
$ ./fpsolve -f test.g --float -i 10 -s newtonSymb
Newton Symbolic
Iterations: 10
Solving time: 0 ms (536 us)
X == 0.999023
Y == 0.999023
```

Here, we select

- (1) the number of iterations (`-i 10`)
- (2) the solving algorithm to use (switch `-s`), possible choices are `newtonSymb`, `newtonConc`, `kleene`.

For larger equation systems there is the possibility to decompose the system into SCCs and solve them bottom-up (switch `--scc`).

## 5.2. Custom Semirings and Extensions

It is very easy and straightforward to extend our library with new semirings, it merely requires three steps:

- Implement all semiring operations (addition  $\oplus$ , multiplication  $\odot$ , star  $*$ )
- Define a constructor that takes a string-argument (effectively a small parser)
- Add a new command-line switch to the main-method together with a call to the solving function.

The second point delegates the IO/parsing task for semiring-elements to the implementer. This enables us to parse equation systems into the most general intermediate format (non-commutative polynomials over the free semiring) and then to map these to the user-defined semiring. Since our input-parser takes quite some time to compile (due to `boost::spirit` and templates), by this approach we avoid to touch the parser and the need for recompilation.

The semiring operations  $\oplus$ ,  $\odot$  (+ and  $*$ ) are implemented in the abstract base-class `Semiring` using `+=` and `*=`. Any new semiring should be derived from the abstract class `StarableSemiring` and has to implement the three operations `*=`, `+=`, `star()`. Take for instance the “MaxProvenance” semiring from the end of Section 2 consisting of pairs  $(\alpha, X)$  of real numbers and sets of variables with

$$\begin{aligned} (\alpha_1, X_1) \odot (\alpha_2, X_2) &:= (\alpha_1 \alpha_2, X_1 \cup X_2) \\ (\alpha_1, X_1) \oplus (\alpha_2, X_2) &:= (\max\{\alpha_1, \alpha_2\}, \text{if } \alpha_1 \geq \alpha_2 \text{ then } X_1 \text{ else } X_2) \\ (\alpha, X)^* &:= (1, \emptyset) \end{aligned}$$

To implement this simple semiring, we derive from `StarableSemiring` the new class `MaxProvSR` with members `weight` and `prov` storing  $\alpha$  (e.g. as a `float`) and  $X$  (e.g. as a `set<>`), respectively. What remains is then to implement the three operators `*=`, `+=`, `star()`. For instance, the `+=` operator could be implemented as

```

MaxProvSR MaxProvSR::operator+=(const MaxProvSR& elem)
{
    if(this->weight < elem.weight) {
        this->weight = elem.weight;
        this->prov = elem.prov;
    }
    return *this;
}

```

Inheritance then takes care of the implementation of the addition operator. Implementing the remaining two operators is just as straight-forward. To make the semiring available in the command line tool, a corresponding command line switch has to be defined. Also, since we parse the input into the most general intermediate representation (non-commutative polynomials over the free semiring) it remains to map this representation to the concrete semiring (which requires only one line of code essentially).

To check our claim of “easy extendability”, we made a rather naive implementation of the Why-semiring for this paper which took about two hours (until all bugs were eliminated<sup>e</sup>). Once a new semiring is defined and the main-method is adapted, all solvers just work out-of-the-box to solve algebraic systems like the following (file `test/grammars/bintrees.g`):

```

<X> ::= a<X><X> | c;

$ ./fpsolve --why -f ../test/grammars/bintrees.g
Newton Concrete
Iterations: 2
Solving time: 0 ms (214 us)
X == {{a,c},{c}}

$ ./fpsolve --why -f ../test/grammars/bintrees.g -s kleene -i 2
Kleene solver
Iterations: 2
Solving time: 0 ms (281 us)
X == {{a,c},{c}}

```

## 6. Conclusions and Related Tools

We have introduced FPSOLVE, an implementation of generic algorithms for solving fixpoint equations on semirings. The algorithms are parametric on the semiring. New semirings can be easily added by defining implementations of the sum, product and (possibly) Kleene star operations.

<sup>e</sup>We developed a small collection of unit-tests (also generic tests that can be instantiated with any semiring) and encourage any user who implements new semirings to use and adapt them during development.

As mentioned in the introduction, many program analysis problems can be reduced to solving fixpoint equations on semirings. This has led to a number of implementations and tools. An early effort is the Fixpoint-Analysis Machine for solving systems of boolean fixpoint equations [25]. The tool can deal with hierarchical and alternating fixpoints, but is not parametric on the equation domain. The Weighted Pushdown Systems Library and Weighted Automata Library (see [22, 3, 2]), and GOBLINT (see [4, 1]) implement many sophisticated algorithms for semirings satisfying the ascending chain condition.

While FPSOLVE is currently an academic tool, we have illustrated its potential interest outside theoretical computer science by means of an application scenario, namely a recommendation system. Genericity allows the users of the system to aggregate the information given by individual recommendations in different, personalized ways, by defining their own semiring.

**Acknowledgments** We thank Michael Kerscher and Michał Terepeta for their help in developing and extending FPSOLVE.

## References

- [1] Goblint <http://goblint.in.tum.de/>.
- [2] WALi: the Weighted Automata Library <https://research.cs.wisc.edu/wpis/wpds/>.
- [3] Weighted Pushdown Systems Library <http://www2.informatik.uni-stuttgart.de/fmi/szs/tools/wpds/>.
- [4] K. Apinis, H. Seidl and V. Vojdani, How to combine widening and narrowing for non-monotonic systems of equations, *PLDI*, eds. H.-J. Boehm and C. Flanagan (ACM, 2013), pp. 377–386.
- [5] A. Bouajjani, J. Esparza, S. Schwoon and D. Suwimonteerabuth, SDSIrep: A Reputation System Based on SDSI, *TACAS*, (2008), pp. 501–516.
- [6] A. Bouajjani, J. Esparza and T. Touili, A generic approach to the static analysis of concurrent programs with procedures, *Int. J. Found. Comput. Sci.* **14**(4) (2003) 551–.
- [7] A. Boudet and H. Comon, Diophantine Equations, Presburger Arithmetic and Finite Automata, *Trees in Algebra and Programming CAAP '96*, ed. H. Kirchner, *Lecture Notes in Computer Science* **1059** (Springer Berlin Heidelberg, 1996), pp. 30–43.
- [8] M. Droste, W. Kuich and H. Vogler, *Handbook of Weighted Automata* (Springer, 2009).
- [9] J. Esparza, S. Kiefer and M. Luttenberger, On Fixed Point Equations over Commutative Semirings, *STACS*, (2007), pp. 296–307.
- [10] J. Esparza, S. Kiefer and M. Luttenberger, Newtonian Program Analysis, *J. ACM* **57**(6) (2010) p. 33.
- [11] J. Esparza and M. Luttenberger, Solving fixed-point equations by derivation tree analysis, *CALCO*, (2011), pp. 19–35.
- [12] A. Finkel and J. Leroux, How to Compose Presburger-Accelerations: Applications to Broadcast Protocols, *FST TCS 2002: Foundations of Software Technology and Theoretical Computer Science*, *LNCS* **2556** (Springer Berlin Heidelberg, 2002), pp. 145–156.
- [13] S. Ginsburg and E. H. Spanier, Bounded ALGOL-like languages, *Transactions of the American Mathematical Society* (1964) 333–368.

- [14] T. J. Green, G. Karvounarakis and V. Tannen, Provenance semirings, *PODS*, (2007), pp. 31–40.
- [15] J. Henriksen, J. Jensen, M. Jørgensen, N. Klarlund, B. Paige, T. Rauhe and A. Sandholm, Mona: Monadic Second-Order Logic in Practice, *Tools and Algorithms for the Construction and Analysis of Systems, First International Workshop, TACAS '95, LNCS 1019*, (1995).
- [16] T.-D. Huynh, The Complexity of Semilinear Sets, *Automata, Languages and Programming*, eds. J. de Bakker and J. van Leeuwen, *Lecture Notes in Computer Science* **85** (Springer Berlin Heidelberg, 1980), pp. 324–337.
- [17] S. Jha and T. W. Reps, Model checking SPKI/SDSI, *Journal of Computer Security* **12**(3-4) (2004) 317–353.
- [18] J. Knoop and B. Steffen, The interprocedural coincidence theorem, *CC*, (1992), pp. 125–140.
- [19] W. Kuich, *Handbook of Formal Languages* (Springer, 1997), ch. 9: Semirings and Formal Power Series: Their Relevance to Formal Languages and Automata, pp. 609 – 677.
- [20] J. Leroux, A Polynomial Time Presburger Criterion and Synthesis for Number Decision Diagrams, *20th IEEE Symposium on Logic in Computer Science (LICS 2005), 26-29 June 2005, Chicago, IL, USA, Proceedings*, (2005), pp. 147–156.
- [21] M. Luttenberger and M. Schlund, Convergence of Newton’s Method over Commutative Semirings, *LATA, LNCS 7810* (2013), pp. 407–418.
- [22] T. Reps, S. Schwoon, S. Jha and D. Melski, Weighted pushdown systems and their application to interprocedural dataflow analysis, *Science of Computer Programming* **58** (October 2005) 206–263, Special Issue on the Static Analysis Symposium 2003.
- [23] M. Schlund, M. Terepeta and M. Luttenberger, Putting Newton into Practice: A Solver for Polynomial Equations over Semirings, *LPAR*, (2013), pp. 727–734.
- [24] F. Somenzi, Colorado University Decision Diagram Package <ftp://vlsi.colorado.edu/pub>, (1998).
- [25] B. Steffen, A. Claßen, M. Klein, J. Knoop and T. Margaria, The Fixpoint-Analysis Machine, *CONCUR*, eds. I. Lee and S. A. Smolka *Lecture Notes in Computer Science* **962**, (Springer, 1995), pp. 72–87.
- [26] P. Wolper and B. Boigelot, An Automata-Theoretic Approach to Presburger Arithmetic Constraints (Extended Abstract), *Proceedings of the Second International Symposium on Static Analysis, SAS '95*, (Springer-Verlag, London, UK, UK, 1995), pp. 21–32.