RWTH Aachen
Mathematische Grundlagen der Informatik

Provenance Analysis for Temporal and Dynamic Logics

Bachelorarbeit

Lovro Mrkonjić

Erstgutachter: Prof. Dr. Erich Grädel
Zweitgutachter: Prof. Dr. Martin Grohe

25. September 2018

Contents

Chapter 1

Introduction

A logical formula φ is usually interpreted over a structure $\mathfrak{A}$, which yields a truth value indicating whether φ is true or false in $\mathfrak{A}$. The goal of *provenance analysis* is to systematically find the facts in $\mathfrak{A}$ that contribute to the truth value of φ and to establish how the truth value of φ depends on those facts. This can be done by generalizing the interpretation of logical formulas to semirings. Instead of assigning truth values from the semiring $\mathbb{B} = \{\bot, \top\}$ to the literals in $\mathfrak{A}$ and evaluating φ to another truth value, we assign values from a fixed semiring K to the literals and interpret φ to obtain a new value $[\![\varphi]\!]$ in K. In other words, we extend the standard interpretations of logical formulas to *semiring interpretations*. Depending on the choice of K, we can obtain additional information aside from the truth value of φ in $\mathfrak{A}$. In particular, provenance analysis can be performed by interpreting φ over $\mathfrak{A}$ in specific polynomial semirings K, we will call those semirings *provenance semirings*.

In this thesis, we will develop semiring interpretations for the well-known temporal logics LTL (Linear Temporal Logic) and CTL (Computation Tree Logic) as well as the positive fragment of the dynamic logic PDL (Propositional Dynamic Logic). These logics can be used, for instance, for verification. Given a program, we could build a transition system from its possible states and transitions. LTL, which we will evaluate on paths in this thesis, could be used to analyse a single run of the program. CTL is evaluated on transition systems and could therefore be used to analyse all possible runs of the program. PDL allows us to model multiple transitions and combine or iterate transitions, in particular, using the positive fragment of PDL, we could for example ask if some target state in a multi-transition system is reachable from a starting state by iterating one specific transition type. Provenance analysis opens up new possibilities. In this example, we could even identify the transitions that form paths between the starting state and the target state using the transition type we specified.

The idea to perform provenance analysis for logics originated from the field of relational databases. For a short introduction, we will summarize some results obtained by Green, Karvounarakis and Tannen in 2007 on semirings and provenance for databases [GKT07]. A *relational database* can be thought of as a finite set of finite relations where each relation has an arity n and contains a set of n-tuples. Given a database, we can interpret *queries* to obtain new relations, where queries are strings of a formal language that indicate how the new relation is calculated from the in-

put data. In practice, relations are usually depicted as tables, their tuples are the table entries and queries can be written, for instance, in SQL (Structured Query Language). For our simplified view on databases, we will also assume that all tuple elements from all tuples in the database are in the same, finite domain D.

Green, Karvounarakis and Tannen introduced *K-relations*, where all possible tuples in a relation are tagged with an element of a set K. Note that when we use $\mathbb{B} = \{\bot, \top\}$ for K, we obtain the usual notion of relations by tagging the tuples that are present in a relation with $\top$ and the remaining tuples with $\bot$. However, we could also use the natural numbers $\mathbb{N}$ for K and tag tuples in a relation with their multiplicity. We notice that, with suitable operations, both $\mathbb{B}$ and $\mathbb{N}$ form the commutative semirings $(\mathbb{B}, \vee, \wedge, \bot, \top)$ and $(\mathbb{N}, +, \cdot, 0, 1)$. Indeed, Green, Karvounarakis and Tannen generalized query interpretation to K-relations for any commutative semiring $(K, +, \cdot, 0, 1)$ and argued that only commutative semirings are suitable for this purpose.

Moreover, they presented some interesting applications of this theory. Suppose there are two relations R_1 and R_2 in our database and we want to interpret a simple query q, for example a join of R_1 and R_2 chained with a projection. Interpreting q yields a new relation R as the query result. Suppose there are two distinct tuples $t_1, t_2 \in R_1$ and a tuple $t_3 \in R_2$ such that q can produce the tuple t by combining t_1 with t_3 or t_2 with t_3 and then applying projection and assume that there is no other possibility to obtain t. Obviously, we will have $t \in R$. If we would now like to add multiplicity to the tuples, we would use $\mathbb{N}$-relations and tag tuples with their multiplicities, for example we could tag t_1 with 2, t_2 with 3 and t_3 with 2. Using the generalized query interpretation would yield an $\mathbb{N}$-relation R and the tuple t would be tagged with the multiplicity $2 \cdot 2 + 3 \cdot 2 = 10$, which is exactly the number of ways that t can be obtained from t_1, t_2 and t_3.

Provenance analysis can now be performed by tagging each tuple with its own unique token. Let X be the set of all those tokens and $\mathbb{N}[X]$ the set of all polynomials with coefficients in $\mathbb{N}$ and free variables in X. Crucially, $(\mathbb{N}[X], +, \cdot, 0, 1)$ is a commutative semiring with respect to standard polynomial addition and multiplication. Therefore, using the example from above, we could tag t_1 with x, t_2 with y and t_3 with z so that $\{x, y, z\} \subseteq X$. Now, interpreting q in $\mathbb{N}[X]$ would yield the tag $x \cdot z + y \cdot z \in \mathbb{N}[X]$ for the resulting tuple t. Clearly, this gives us an insight on the *provenance* of the tuple t, we can see how many different ways there are to produce t, which tuples from the input produce t and even how often those tuples are used.

The last result by Green, Karvounarakis and Tannen that we will mention for now is that they have shown their extended query interpretation to be compatible with semiring homomorphisms. They argued that any valuation of variables in X with elements from a commutative semiring K induces a unique semiring homomorphism from $\mathbb{N}[X]$ to K, so they concluded that $\mathbb{N}[X]$ is the most "general" semiring to interpret queries. Indeed, we can see that once we have computed the provenance polynomial $x \cdot z + y \cdot z$ for t in our example, all we have to do is to evaluate x, y and z with the multiplicities 2, 3 and 2 of their corresponding tuples and the induced homomorphism from $\mathbb{N}[X]$ to $\mathbb{N}$ would evaluate $x \cdot z + y \cdot z$ to $2 \cdot 2 + 3 \cdot 2 = 10$, which is exactly the multiplicity of t in our result.

Now, this approach can be adapted to first-order logic. In 2017, Grädel and Tannen

developed a provenance analysis for full first-order logic based on semiring interpretations [GT17]. To illustrate the connection between relational databases and first-order logic, recall that first-order sentences φ are interpreted over structures $\mathfrak{A} = (A, \tau)$, where A is a universe and τ is a set of function and relation symbols. We restrict A to be finite and τ to be finite and relational. Note that a function can be viewed as a relation by using the function's graph instead of the function itself. Now, interpreting φ over $\mathfrak{A}$ yields a value in $\mathbb{B}$, indicating whether φ is true in $\mathfrak{A}$ or not.

Let $R \in \tau$ be a relation of arity n and consider an arbitrary n-tuple $(a_1, ..., a_n)$ in A. $Ra_1...a_n$ and $\neg Ra_1...a_n$ are called *literals* and each of them is true or false in $\mathfrak{A}$. Notice the similarity between $\mathfrak{A}$ and a database. If R was a relation in a simplified database where all elements are from A, we could informally say that the literal $Ra_1...a_n$ was true if the tuple $(a_1, ..., a_n)$ was in R in our database and false otherwise. An important difference is that the negative literals $\neg Ra_1...a_n$ are not considered in databases, but they are important in logics, since formulas can usually be negated.

Similarly to extending relations to K-relations, Grädel and Tannen built a semiring interpretation for first-order logic by fixing an arbitrary, commutative semiring K and allowing each literal to be assigned a value in K rather than just true or false. Since for general semirings, it is not clear how to "negate" a value $a \in K$, they built their semiring interpretation on the *negation normal form* of first-order sentences, where only literals can appear in negated form and required both positive and negative literals to be tagged with values in K. In negation normal form, first-order sentences are just combinations of conjunctions and disjunctions aside from quantifiers and equalities. However, since the universe A is finite, quantifiers can be seen as finite conjunctions and disjunctions as well.

Grädel and Tannen have defined their semiring interpretation by interpreting disjunctions as additions in the semiring and conjunctions as a multiplications. For example, consider a first-order formula $\varphi = L_1 \vee (L_2 \wedge L_3)$ in negation normal form where L_1, L_2 and L_3 are literals. Interpreting φ in a semiring K would yield $[\![\varphi]\!] = [\![L_1]\!] + ([\![L_2]\!] \cdot [\![L_3]\!])$. Just as for the database example, we can choose a semiring K and interpret the literals accordingly to obtain useful results. For instance, as Grädel and Tannen pointed out, the Viterbi semiring $\mathbb{V} = ([0, 1], \max, \cdot, 0, 1)$ whose elements are confidence scores could be used. If we set the confidence values for our literals to $[\![L_1]\!] = 0.3$, $[\![L_2]\!] = 0.8$ and $[\![L_3]\!] = 0.5$, we obtain $[\![\varphi]\!] = \max\{0.3, 0.8 \cdot 0.5\} = 0.4$. This is the confidence score for φ being true.

Moreover, similarly to databases, the semirings to provide the most general results are polynomial semirings. In the above example, we can use $\mathbb{N}[X]$ with $\{x, y, z\} \subseteq X$ for provenance analysis. Assuming that L_1, L_2 and L_3 are true in $\mathfrak{A}$, we can track them by setting $[\![L_1]\!] = x$, $[\![L_2]\!] = y$ and $[\![L_3]\!] = z$. Interpreting φ yields $[\![\varphi]\!] = x + yz$. As we can see, each monomial corresponds to a proof of φ in $\mathfrak{A}$. The variables in a monomial indicate which literals are used in the proof, their exponents indicate how often they are used and coefficients indicate how many distinct proofs there are that use the same literals. Grädel and Tannen have also shown that this is generally true, performing provenance analysis in first-order logic is therefore done by interpreting formulas in polynomial semirings and the resulting polynomial describes all the proofs of the formula.

This semiring interpretation is also compatible with homomorphisms, so if we were only interested in counting the number of proofs for φ in the above example, we could interpret the formula in $\mathbb{N}$ and assign 1 to the true literals L_1, L_2 and L_3, or instead, we could just plug the values into the polynomial $x + yz$ for the corresponding variables. This would yield $1 + 1 \cdot 1 = 2$ which is the correct number of different proofs for φ.

Since, as we will see later on, LTL, CTL and PDL are logics that are interpreted over relational structures, we will use a similar approach as Grädel and Tannen to define semiring interpretations for these logics. In other words, we will tag literals with elements of a semiring K and then interpret our formulas in negation normal form. In particular, LTL, CTL and PDL have disjunctions and conjunctions as well, and we will also interpret those by addition and multiplication in K. However, these logics admit additional operators that are not always expressible in first-order logic, so especially for CTL and PDL, we will have to think of new ways to interpret them. Our interpretations will require additional conditions that are not met in all commutative semirings. Therefore, we will first of all consider smaller classes of semirings, introduce their important characteristics and look for new provenance semiring candidates in those smaller classes in the next chapter.

Chapter 2

Semirings

We will first provide a formal definition of semirings according to Grädel and Tannen [GT17]. Unless otherwise stated, the definitions and results in the beginning of this chapter and the following section about ω-continuous semirings are adapted from their unpublished work on provenance for logic and games [GT18].

(2.1) Definition (Semiring)**.** A *semiring* is an algebraic structure $(K, +, \cdot, 0, 1)$ where $0 \neq 1$, $(K, +, 0)$ is a commutative monoid, $(K, \cdot, 1)$ is a monoid, multiplication distributes over addition and multiplication by 0 annihilates elements, that is,

$$\begin{array}{lllll} (1) & a \cdot (b + c) = a \cdot b + a \cdot c & \text{and} & (b + c) \cdot a = b \cdot a + c \cdot a & \text{for } a, b, c \in K \text{ and} \\ (2) & a \cdot 0 = 0 & \text{and} & 0 \cdot a = 0 & \text{for } a \in K. \end{array}$$

The semiring is commutative if the monoid $(K, \cdot, 1)$ is commutative.

We will write K instead of $(K, +, \cdot, 0, 1)$ if $(+, \cdot, 0, 1)$ is clear from the context. Also, we only consider commutative semirings. This is justified, because, as stated in the introduction, we would like to interpret conjunctions as multiplications, so a formula $\varphi \wedge \psi$ would be interpreted as $[\![\varphi]\!] \cdot [\![\psi]\!]$ in K. Naturally, one would expect the formula $\psi \wedge \varphi$ that is interpreted as $[\![\psi]\!] \cdot [\![\varphi]\!]$ to yield the same semiring value. However, this is not generally true unless the semiring K is commutative, therefore we will implicitly assume all semirings in this thesis to be commutative.

To illustrate the above definition, we will show some examples of commutative semirings provided by Grädel and Tannen. In the introduction, we have already mentioned the Boolean semiring $(\mathbb{B}, \vee, \wedge, \bot, \top)$, the semiring of natural numbers $(\mathbb{N}, +, \cdot, 0, 1)$, polynomial semirings $(\mathbb{N}[X], +, \cdot, 0, 1)$ for any set X and the Viterbi semiring $\mathbb{V} = ([0, 1], \max, \cdot, 0, 1)$. Additionally, $\mathbb{T} = (\mathbb{R}_+^\infty, \min, +, \infty, 0)$ is called the *tropical* semiring, which can be used for the calculation of shortest paths, $\mathbb{F} = ([0, 1], \max, \min, 0, 1)$ is the *fuzzy* semiring, which is usable for fuzzy logic and $\mathbb{A} = (\{\mathrm{P} < \mathrm{C} < \mathrm{S} < \mathrm{T} < 0\}, \min, \max, 0, \mathrm{P})$ also forms a semiring, where the values are used for *access control* (P is public, C is confidential, S is secret, T is top secret and 0 is inaccessible).

While first-order sentences can be interpreted in any commutative semiring, the interpretations of CTL and PDL formulas will often be solutions of recursive equations. These solutions can be found by computing fixed points of functions $f : K \to K$. Of course, it is unreasonable to assume that any function $f : K \to K$ has a fixed

point in K, we just have to consider a function f with $f(0) = 1$ and $f(a) = 0$ for any $a \in K$ with $a \neq 0$. However, the functions that we are interested in will be expressible by variables, constants, additions and multiplications. Therefore, we say that a semiring K *admits* least (greatest) *fixed points* if any function $f : K \to K$ that is composed of constants, variables, addition and multiplication has a least (greatest) fixed point in K. Clearly, there are commutative semirings that do not admit any fixed points, for example, $f : \mathbb{N} \to \mathbb{N}$ with $f(n) = n + 1$ for $n \in \mathbb{N}$ does not have a fixed point in $\mathbb{N}$. Therefore, we will need semirings that satisfy additional conditions to interpret CTL and PDL formulas.

First of all, the notion of least and greatest fixed points requires an order on semirings. For any semiring K, we can define the relation $\leq$ on K by setting

$$a \leq b \quad \text{iff} \quad \text{there is a } d \in K \text{ such that } a + d = b \quad \text{for } a, b \in K.$$

Since $(K, +, 0)$ has a neutral element, $\leq$ is reflexive. The associativity and commutativity of $(K, +, 0)$ also guarantee that $\leq$ is transitive. Therefore, $\leq$ is a partial order on K if and only if it is antisymmetric. For example, on $\mathbb{N}$, $\leq$ is antisymmetric, but if K is a ring, $\leq$ is always symmetric because $a + d = b$ implies $b + (-d) = a$. We would like to consider semirings where $\leq$ is a partial order.

(2.2) Definition (Naturally Ordered Semiring)**.** A semiring K is *naturally ordered* if the relation $\leq$ is a partial order. In that case, we call $\leq$ the *natural order* on K.

In naturally ordered semirings, we always talk about least and greatest fixed points, suprema, infima, maxima or minima with respect to the natural order.

2.1 ω-Continuous Semirings

In this thesis, we admit the axiom of choice and use ordinal numbers, where ω refers to the ordinal number $\{0, 1, ...\}$. Since we have an appropriate order, we can describe a class of semirings that admits least fixed points.

(2.3) Definition (ω-Continuous Semiring)**.** A naturally ordered semiring K is *ω-continuous* if every ascending ω-chain $a_0 \leq a_1 \leq ...$ in K has a supremum $\sup_{i \in \omega} a_i$ in K and addition and multiplication are ω-continuous, that is, for $c \in K$, we have

$$\begin{aligned} c + \sup_{i \in \omega} a_i &= \sup_{i \in \omega}(c + a_i) \quad \text{and} \\ c \cdot \sup_{i \in \omega} a_i &= \sup_{i \in \omega}(c \cdot a_i). \end{aligned}$$

Unfortunately, $\mathbb{N}$ is not ω-continuous, since the ascending chain $0 \leq 1 \leq ...$ does not have a supremum in $\mathbb{N}$. However, we can adjoin infinity to $\mathbb{N}$ to obtain an ω-continuous semiring $\mathbb{N}^\infty = \mathbb{N} \cup \{\infty\}$ with

$$\begin{aligned} n + \infty &= \infty \quad \text{for } n \in \mathbb{N}^\infty \text{ and} \\ n \cdot \infty &= \infty \quad \text{for } n \in \mathbb{N}^\infty \setminus \{0\}. \end{aligned}$$

The properties of ω-continuous semirings allow us to define a summation for any countable sequence of elements in K.

(2.4) Definition (Countable Summation)**.** Let K be an ω-continuous semiring, the summation of a countable sequence $b_0, b_1, ...$ in K is defined as

$$\sum_{i\in\omega} b_i = \sup_{i\in\omega}(b_0 + ... + b_i).$$

Notice that the partial sums $b_0 + ... + b_i$ form an ascending chain and that this summation is compatible with the usual, finite summation. Distributivity, associativity and commutativity also extend to countable summation.

(2.5) Proposition (Countable Summation Laws)**.** In an ω-continuous semiring K, let $c, b_0, b_1, ... \in K$ and $(I_j)_{j\in J}$ be a partition of ω, then we have

$$(1) \quad c \cdot \sum_{i\in\omega} b_i = \sum_{i\in\omega}(c \cdot b_i) \quad \text{and}$$
$$(2) \quad \sum_{j\in J}\sum_{i\in I_j} b_i = \sum_{i\in\omega} b_i.$$

Aside from countable summation, ω-continuous semirings also admit least fixed points. We use Kleene's fixed-point theorem as stated by Baranga [Bar91]. In an ω-continuous semiring K, a function $f : K \to K$ is called ω-continuous if for ascending chains $a_0 \leq a_1 \leq ...$ in K, $f(\sup_{i\in\omega} a_i) = \sup_{i\in\omega} f(a_i)$. In particular, any function composed of addition and multiplication is ω-continuous, since chaining ω-continuous functions yields ω-continuous functions as well.

(2.6) Theorem. Let K be an ω-continuous semiring, then any ω-continuous function $f : K \to K$ has a least fixed point $\mathrm{lfp}(f)$ in K. Moreover, we have

$$\mathrm{lfp}(f) = \sup_{i\in\omega} f^i(0).$$

For example, the function $f(n) = n + 1$ in $\mathbb{N}^\infty$ has a least fixed point $\mathrm{lfp}(f) = \infty$. We can also verify that $\sup_{i\in\omega} f^i(0) = \sup\{0, 1, ...\} = \infty$.

For provenance analysis, we are also interested in more complex ω-continuous semirings. However, the polynomial semiring $\mathbb{N}[X]$ is not ω-continuous. The first problem is that the coefficients are in $\mathbb{N}$, so $\sup_{i\in\omega}(i \cdot x)$ does not exist for any variable $x \in X$ and the second problem is that polynomials are finite, therefore $\sum_{i\in\omega} x^i$ does not exist in $\mathbb{N}[X]$. These restrictions can be bypassed by choosing a different set of coefficients and allowing infinite polynomials.

(2.7) Definition (Formal Power Series)**.** Let K be a semiring and $X = \{x_1, ..., x_n\}$ a finite set of variables. A *formal power series* with coefficients in K and variables in X is a possibly infinite sum of monomials of the form $a \cdot x_1^{e_1} \cdot ... \cdot x_n^{e_n}$ where $a \in K$ and $e_1, ..., e_n \in \mathbb{N}$. We denote the set of formal power series with coefficients in K and variables in X as $K[\![X]\!]$.

$K[\![X]\!]$ forms a semiring with the operations defined as usual. If K is ω-continuous, then $K[\![X]\!]$ is ω-continuous as well. In particular, $\mathbb{N}^\infty[\![X]\!]$ is ω-continuous and we can

use it for provenance analysis. We observe that both problems that we mentioned above are solved in $\mathbb{N}^\infty[\![X]\!]$, since for $x \in X$, we have

$$\sup_{i\in\omega}(i \cdot x) = \infty \cdot x \qquad \in \mathbb{N}^\infty[\![X]\!] \text{ and}$$
$$\sum_{i\in\omega} x^i = 1 + x + x^2 + ... \in \mathbb{N}^\infty[\![X]\!].$$

To summarize, ω-continuous semirings admit least fixed points and countable summation. There is also an ω-continuous semiring $\mathbb{N}^\infty[\![X]\!]$ that is suitable for provenance analysis. However, ω-continuous semirings do not admit greatest fixed points, therefore we will introduce another class of semirings that admits both least and greatest fixed points.

2.2 Absorptive Lattice Semirings

The idea to use absorptive semirings whose natural order forms a complete lattice originates from Grädel and Tannen [GT18].

A semiring K is called *absorptive*, if for any $a, b \in K$,

$$a + a \cdot b = a.$$

Absorption in provenance analysis is justified by the observation that, as seen in the introduction, formulas of the form $\varphi \vee (\varphi \wedge \psi)$ would be interpreted as $[\![\varphi]\!] + [\![\varphi]\!] \cdot [\![\psi]\!]$. Clearly, in any logic, that formula is equivalent to φ. When we perform provenance analysis, we find all the proofs for $\varphi \vee (\varphi \wedge \psi)$. However, proving φ suffices to prove the entire formula, so one could argue that it is now justified to disregard the proofs for $\varphi \wedge \psi$, since they are "unnecessarily long". So, in absorptive semirings, since $[\![\varphi]\!] + [\![\varphi]\!] \cdot [\![\psi]\!] = [\![\varphi]\!]$, we do not find all the proofs for a formula, but we can still find the "shortest" proofs. In turn, absorptive, naturally ordered semirings have useful properties.

(2.8) Lemma. Let K be naturally ordered and absorptive, then, for $a, b \in K$,

(1)	$a + b = \sup\{a, b\}$	(addition yields the supremum),
(2)	$a \cdot b \leq a, b$	(multiplication decreases elements),
(3)	$0 \leq a$	(0 is the bottom element) and
(4)	$a \leq 1$	(1 is the top element).

Proof. For (1), $a + b$ is clearly an upper bound on $\{a, b\}$. Now, consider an upper bound $c \in K$ on $\{a, b\}$. Then, we have $d, e \in K$ with $a + d = b + e = c$. We conclude that $a + b \leq c$, since $(a + b) + (d + e) = (a + d) + (b + e) = c + c = c + c \cdot 1 = c$ by absorption. Therefore, $a + b$ is the least upper bound on $\{a, b\}$.
(2) is due to $a \cdot b + a = a + a \cdot b = a$ and $a \cdot b + b = b + b \cdot a = b$ by absorption.
(3) and (4) are shown by $0 + a = a$ and $a + 1 = 1 + 1 \cdot a = 1$. □

We will call a function $f : K \to K$ on a naturally ordered semiring K monotonic if $a \leq b$ implies $f(a) \leq f(b)$ for $a, b \in K$. Addition and multiplication are both

monotonic in each argument, as we can easily verify by observing that $a \leq b$ implies $a + d = b$ for some $d \in K$, so for any $c \in K$, we can infer

$$\begin{aligned}(a + c) + d = (a + d) + c = b + c \quad &\text{and}\\ (a \cdot c) + (d \cdot c) = (a + d) \cdot c \ &= b \cdot c,\end{aligned}$$

therefore $a+c \leq b+c$ and $a \cdot c \leq b \cdot c$. Since composition of monotonic functions yields monotonic functions as well, if the natural order on K was a complete lattice, the Knaster-Tarski fixed-point theorem would imply that K admits least and greatest fixed points.

(2.9) Definition (Complete Lattice)**.** A partial order $(K, \leq)$ is a *complete lattice* if every subset $S \subseteq K$ has an infimum and a supremum in K. Additionally, $(K, \leq)$ is called *completely distributive* [Ran52], if for any $(I_j)_{j \in J}$ where J is a set of indices and $I_j \subseteq K$ for $j \in J$, the equation

$$\inf_{j \in J} \sup I_j = \sup_{f \in F} \inf_{j \in J} f(j)$$

holds, where F is the set of choice functions $f : J \to K$ with $f(j) \in I_j$ for $j \in J$.

We can now define a class of semirings K that admit least and greatest fixed points. However, we would also like the addition and multiplication in K to be compatible with infima and suprema, therefore K should be absorptive, completely distributive and satisfy additional conditions as well.

(2.10) Definition (Absorptive Lattice Semiring)**.** An *absorptive lattice semiring* is an absorptive, naturally ordered semiring K such that the natural order $(K, \leq)$ is a completely distributive lattice. Moreover, for any $c \in K$, any subset $S \subseteq K$ and any descending ω-chain $a_0 \geq a_1 \geq ...$, the equations

$$\begin{aligned}&(1) \quad c \cdot \sup S = \sup(c \cdot S) \quad \text{and}\\ &(2) \quad c \cdot \inf_{i \in \omega} a_i = \inf_{i \in \omega}(c \cdot a_i)\end{aligned}$$

are satisfied, where $c \cdot S = \{c \cdot a \mid a \in S\}$.

The additional conditions ensure the compatibility of multiplication with suprema of arbitrary sets and infima of descending ω-chains. Using lemma (2.8), we can also prove that addition is compatible with suprema and infima.

(2.11) Lemma. Let K be an absorptive lattice semiring. For any $c \in K$ and any non-empty subset $\emptyset \subsetneq S \subseteq K$,

$$\begin{aligned}&(1) \quad c + \sup S = \sup(c + S) \quad \text{and}\\ &(2) \quad c + \inf S = \inf(c + S)\end{aligned}$$

hold with $c + S = \{c + a \mid a \in S\}$.

Proof. Since K is absorptive and naturally ordered, lemma (2.8) implies that addition yields the supremum of two elements. Therefore, (1) is equivalent to $\sup\{c, \sup S\} = \sup\{\sup\{c, a\} \mid a \in S\}$. We prove both directions:

"$\geq$": $\sup\{c, \sup S\}$ is greater than c and $\sup S$, therefore it is also greater than any $a \in S$, which makes it greater than $\sup\{c, a\}$ for any $a \in S$. We conclude that $\sup\{c, \sup S\} \geq \sup\{\sup\{c, a\} \mid a \in S\}$.

"$\leq$": $\sup\{\sup\{c, a\} \mid a \in S\}$ is greater than $\sup\{c, a\}$ for any $a \in S$, so, since S is not empty, it is greater than c and greater than a for any $a \in S$. We conclude that it is greater than $\sup S$ as well, so $\sup\{\sup\{c, a\} \mid a \in S\} \geq \sup\{c, \sup S\}$.

For (2), we first show the direction $c + \inf S \leq \inf(c + S)$ using the monotonicity of addition. Since $\inf S \leq a$ for any $a \in S$, we have $c + \inf S \leq c + a$ for any $a \in S$, and therefore $c + \inf S$ is a lower bound on $c + S$.

The other direction is equivalent to $\sup\{c, \inf S\} \geq \inf\{\sup\{c, a\} \mid a \in S\}$ and we can show it using the complete distributivity of $(K, \leq)$. We have

$$\inf\{\sup\{c, a\} \mid a \in S\} = \inf_{a \in S} \sup\{c, a\} = \sup_{f \in F} \inf_{a \in S} f(a)$$

where F is the set of functions $f : S \to K$ with $f(a) \in \{c, a\}$ for all $a \in S$. Consider such a function f. If $f(a) = c$ for some $a \in S$, then $\inf_{a \in S} f(a) \leq c$. Otherwise, $f(a) = a$ for all $a \in S$, so $f = \mathrm{id}_S$ and $\inf_{a \in S} f(a) = \inf S$, therefore, $\sup\{c, \inf S\}$ is an upper bound on $\{\inf_{a \in S} f(a) \mid f \in F\}$ and the lemma is proven. □

(2.12) Corollary. Any absorptive lattice semiring K is also ω-continuous.

For ω-continuous semirings, it is possible to define a summation for countable sequences. We will see that absorptive lattice semirings can extend this summation to work for arbitrary sets and that we can even define a multiplication for countable sequences.

(2.13) Definition (Infinite Summation)**.** Let K be an absorptive lattice semiring and $S \subseteq K$. We define

$$\sum S = \sup S.$$

Since by lemma (2.8), addition and suprema coincide, this definition clearly coincides with the regular summation for finite sets S and the summation in ω-continuous semirings for countable sets S. We will now perform some further sanity checks.

(2.14) Proposition (Infinite Summation Laws)**.** Let K be an absorptive lattice semiring, $c \in K$, $S \subseteq K$ and $(I_j)_{j \in J}$ be a partition of S, then the equations

$$\begin{aligned} &(1) \quad c \cdot \sum S = \sum (c \cdot S) \quad \text{and} \\ &(2) \quad \sum_{j \in J} \sum I_j = \sum S \end{aligned}$$

hold with $c \cdot S = \{c \cdot a \mid a \in S\}$.

Proof. Since summation and suprema are the same, (1) is true by definition (2.10) (1). Also, (2) is equivalent to $\sup_{j \in J} \sup I_j = \sup S$. We show both directions:

"$\leq$": Since $\sup S$ is an upper bound on S, it is also an upper bound on all the subsets I_j of S for $j \in J$. Therefore, $\sup S \geq \sup I_j$ for $j \in J$ and $\sup S$ is an upper bound on $\{\sup I_j \mid j \in J\}$.

"$\geq$": For any $a \in S$, there is a $j \in J$ with $a \in I_j$. This implies $\sup I_j \geq a$ and $\sup_{j \in J} \sup I_j \geq a$ for any $a \in S$. It follows that $\sup_{j \in J} \sup I_j$ is an upper bound on S, which ends the proof. □

(2.15) Definition (Countable Multiplication)**.** In an absorptive lattice semiring

K, we define the multiplication for a countable sequence $b_0, b_1, ...$ in K as

$$\prod_{i\in\omega} b_i = \inf_{i\in\omega}(b_0 \cdot ... \cdot b_i).$$

Since multiplication decreases elements, the partial products form a descending chain and this is compatible with the normal, finite multiplication. Rearranging the order of the elements of the sequence does not change the value of the product. This can be shown by considering an arbitrary bijection $\pi : \omega \to \omega$ and observing that

$$\prod_{i\in\omega} b_i = \inf_{i\in\omega}(b_0 \cdot ... \cdot b_i) = \inf_{i\in\omega}(b_{\pi(0)} \cdot ... \cdot b_{\pi(i)}) = \prod_{i\in\omega} b_{\pi(i)}.$$

We will just show the direction "$\leq$". For any partial product in $\prod_{i\in\omega} b_{\pi(i)}$ of the form $b_{\pi(0)}\cdot ...\cdot b_{\pi(i)}$, there is a partial product in $\prod_{i\in\omega} b_i$ of the form $b_0\cdot ...\cdot b_j$ that is less than $b_{\pi(0)} \cdot ... \cdot b_{\pi(i)}$, because since $\{0, ..., i\}$ is finite, we can set $j = \max \pi(\{0, ..., i\})$, and then $b_0 \cdot ... \cdot b_j$ contains all the factors of $b_{\pi(0)} \cdot ... \cdot b_{\pi(i)}$ and is therefore smaller, since multiplication decreases elements. The converse direction can be shown the same way by considering the inverse function π^{-1} of π. It is still left to show that this multiplication is invariant under partitioning.

(2.16) Proposition (Countable Multiplication Law)**.** Let K be an absorptive lattice semiring, $b_0, b_1, ...$ be a countable sequence in K and $(I_j)_{j\in J}$ a partition of ω, then

$$\prod_{j\in J}\prod_{i\in I_j} b_i = \prod_{i\in\omega} b_i.$$

Proof. Since $(I_j)_{j\in J}$ is a partition of ω, each I_j and J is countable. Since the order of the elements does not matter and countable products are compatible with finite products, it is justified to say w.l.o.g. that $J = \alpha$ for some possibly finite ordinal number $\alpha \leq \omega$. We will now restate the proposition using the infimum definitions of the products as

$$\prod_{j\in J}\prod_{i\in I_j} b_i = \inf_{k<\alpha}\left(\prod_{k'\leq k}\prod_{l\in I_{k'}} b_l\right) = \inf_{i\in\omega}(b_0 \cdot ... \cdot b_i) = \prod_{i\in\omega} b_i.$$

"$\leq$": It suffices to show that for every $b_0 \cdot ... \cdot b_i$, there is a $k < \alpha$ such that

$$\prod_{k'\leq k}\prod_{l\in I_{k'}} b_l \leq b_0 \cdot ... \cdot b_i.$$

Since $\{0, ..., i\}$ is finite, there is a finite subset $J' \subseteq J$ such that each element in $\{0, ..., i\}$ is contained in an I_j with $j \in J'$. We pick $k = \max J' < \alpha$. Then, for any $k' \in J'$, we have $k' \leq k$. Since multiplication decreases elements, we obtain

$$\prod_{k'\leq k}\prod_{l\in I_{k'}} b_l \leq \prod_{k'\in J'}\prod_{l\in I_{k'}} b_l,$$

so it remains to show $\prod_{k'\in J'}\prod_{l\in I_{k'}} b_l \leq b_0\cdot ...\cdot b_i$. Consider the product $\prod_{l\in I_{k'}} b_l$ for an arbitrary $k' \in J'$. Since $\{0, ..., i\}$ is finite, there is a partial product of $\prod_{l\in I_{k'}} b_l$ that

contains all the factors b_l with $l \in I_{k'} \cap \{0, ..., i\}$ and with multiplication decreasing elements, we have $\prod_{l \in I_{k'}} b_l \leq \prod_{l \in I_{k'} \cap \{0,...,i\}} b_l$ for each $k' \in J'$. The monotonicity of multiplication in each argument yields

$$\prod_{k' \in J'} \prod_{l \in I_{k'}} b_l \leq \prod_{k' \in J'} \prod_{l \in I_{k'} \cap \{0,...,i\}} b_l.$$

We notice that all the products are finite and each of the elements in $\{0, ..., i\}$ is contained in exactly one $I_{k'}$ with $k' \in J'$, so precisely the elements $b_0, ..., b_i$ constitute the product. Associativity and commutativity yield

$$\prod_{k' \in J'} \prod_{l \in I_{k'} \cap \{0,...,i\}} b_l = b_0 \cdot ... \cdot b_i,$$

which ends this direction of the proof.

"$\geq$": For this direction, we will show that for every $k < \alpha$,

$$\inf_{i \in \omega}(b_0 \cdot ... \cdot b_i) \leq \prod_{k' \leq k} \prod_{l \in I_{k'}} b_l =: P_k.$$

Since the inner products in P_k could be infinite, we will use the same argument as above and say w.l.o.g. that each I_j has a possibly finite cardinality $\alpha_j \leq \omega$ and we can fix an order on the indices in I_j by setting $I_j = \{i_{j,l} \mid l \in \alpha_j\}$. We will also write $b(i)$ instead of b_i for better readability and rewrite the inner products of P_k as

$$P_k = \prod_{k' \leq k} \inf_{l < \alpha_{k'}} \left(\prod_{l' \leq l} b(i_{k',l'}) \right).$$

We will prove by induction that for any $k < \alpha$, we have

$$P_k = \inf_{(l_0,...,l_k) < (\alpha_0,...,\alpha_k)} \left(\prod_{k' \leq k} \prod_{l'_{k'} \leq l_{k'}} b(i_{k',l'_{k'}}) \right),$$

where $<$ is defined component-wise on tuples.

For $k = 0$, we just rename indices and obtain

$$P_0 = \inf_{l < \alpha_0} \left(\prod_{l' \leq l} b(i_{0,l'}) \right) = \inf_{(l_0) < (\alpha_0)} \left(\prod_{l'_0 \leq l_0} b(i_{0,l'_0}) \right).$$

For $k+1$, we rewrite P_{k+1}

$$
\begin{aligned}
&= \prod_{k' \leq k+1} \inf_{l < \alpha_{k'}} \left(\prod_{l' \leq l} b(i_{k',l'}) \right) \\
&= \left(\prod_{k' \leq k} \inf_{l < \alpha_{k'}} \left(\prod_{l' \leq l} b(i_{k',l'}) \right) \right) \cdot \left(\inf_{l_{k+1} < \alpha_{k+1}} \left(\prod_{l'_{k+1} \leq l_{k+1}} b(i_{k+1,l'_{k+1}}) \right) \right) \\
&\overset{(1)}{=} \left(\inf_{(l_0,\ldots,l_k) < (\alpha_0,\ldots,\alpha_k)} \underbrace{\left(\prod_{k' \leq k} \prod_{l'_{k'} \leq l_{k'}} b(i_{k',l'_{k'}}) \right)}_{p(l_0,\ldots,l_k)} \right) \cdot \underbrace{\left(\inf_{l_{k+1} < \alpha_{k+1}} \left(\prod_{l'_{k+1} \leq l_{k+1}} b(i_{k+1,l'_{k+1}}) \right) \right)}_{c_1} \\
&\overset{(2)}{=} \inf_{(l_0,\ldots,l_k) < (\alpha_0,\ldots,\alpha_k)} \left(\underbrace{\left(\prod_{k' \leq k} \prod_{l'_{k'} \leq l_{k'}} b(i_{k',l'_{k'}}) \right)}_{c_2} \cdot \left(\inf_{l_{k+1} < \alpha_{k+1}} \underbrace{\left(\prod_{l'_{k+1} \leq l_{k+1}} b(i_{k+1,l'_{k+1}}) \right)}_{p'(l_{k+1})} \right) \right) \\
&\overset{(3)}{=} \inf_{(l_0,\ldots,l_k) < (\alpha_0,\ldots,\alpha_k)} \left(\inf_{l_{k+1} < \alpha_{k+1}} \left(\left(\prod_{k' \leq k} \prod_{l'_{k'} \leq l_{k'}} b(i_{k',l'_{k'}}) \right) \cdot \left(\prod_{l'_{k+1} \leq l_{k+1}} b(i_{k+1,l'_{k+1}}) \right) \right) \right) \\
&\overset{(4)}{=} \inf_{(l_0,\ldots,l_k,l_{k+1}) < (\alpha_0,\ldots,\alpha_k,\alpha_{k+1})} \left(\left(\prod_{k' \leq k} \prod_{l'_{k'} \leq l_{k'}} b(i_{k',l'_{k'}}) \right) \cdot \left(\prod_{l'_{k+1} \leq l_{k+1}} b(i_{k+1,l'_{k+1}}) \right) \right) \\
&= \inf_{(l_0,\ldots,l_k,l_{k+1}) < (\alpha_0,\ldots,\alpha_k,\alpha_{k+1})} \left(\prod_{k' \leq k+1} \prod_{l'_{k'} \leq l_{k'}} b(i_{k',l'_{k'}}) \right).
\end{aligned}
$$

We have to verify the transformations (1) to (4) before ending the induction. In transformation (1), we make use of the induction hypothesis. For (2), we can enumerate the tuples $(l_0, ..., l_k) < (a_0, ..., a_k)$ in a way that the partial products $p(l_0, ..., l_k)$ form a descending chain by never decreasing any $l_{k'}$. Then, we apply condition (2) of definition (2.10) to pull c_1 into the infimum. The same argument is used for (3), where we observe that the partial products $p'(l_{k+1})$ form a descending chain and pull c_2 into the infimum. Transformation (4) is verified by observing that for index sets A, B with $x_{a,b} \in K$ for $a \in A$ and $b \in B$, we have $\inf_{a \in A} \inf_{b \in B} x_{a,b} = \inf_{(a,b) \in A \times B}$, because infima are invariant under partition, therefore we can merge the two infimum computations. This ends the induction.

For an arbitrary P_k, we now know that P_k is the infimum of the partial products $P(l_0, ..., l_k) = \prod_{k' \leq k} \prod_{l'_{k'} \leq l_{k'}} b(i_{k',l'_{k'}})$ where $(l_0, ..., l_k) < (\alpha_0, ..., \alpha_k)$. Each of these products is finite and contains each factor $b_{i'}$ for $i' \in \omega$ at most once, because $(I_j)_{j \in J}$ was a partition of ω. Therefore, there is an $i \in \omega$ such that $b_0 \cdot ... \cdot b_i$ contains all the factors in $P(l_0, ..., l_k)$. So, the infimum $\inf_{i \in \omega}(b_0 \cdot ... \cdot b_i)$ is less than P_k and the proof is complete. □

We can now return to our goal of admitting least and greatest fixed points. Patrick

and Radhia Cousot have stated the Knaster-Tarski fixed-point theorem in their work [CC79] and they have also shown how to construct the fixed points that we are looking for. We can derive the following theorem.

(2.17) Theorem. Let K be an absorptive lattice semiring. Any monotonic function $f : K \to K$ has a least fixed point $\mathrm{lfp}(f)$ and a greatest fixed point $\mathrm{gfp}(f)$ in K. Moreover, we can define the transfinite sequences

$$\begin{aligned} x_0 &= 0 & y_0 &= 1 && \text{for the ordinal 0,} \\ x_{\alpha+1} &= f(x_\alpha) & y_{\alpha+1} &= f(y_\alpha) && \text{for successor ordinals } \alpha+1 \text{ and} \\ x_\lambda &= \sup_{\alpha<\lambda} x_\alpha & y_\lambda &= \inf_{\alpha<\lambda} y_\alpha && \text{for limit ordinals } \lambda. \end{aligned}$$

Then, there is a least ordinal β such that $f(x_\beta) = x_\beta$ and a least ordinal γ with $f(y_\gamma) = y_\gamma$ and we have $\mathrm{lfp}(f) = x_\beta$ and $\mathrm{gfp}(f) = y_\gamma$.

Absorptive lattice semirings admit infinite summation, countable multiplication and least and greatest fixed points. We will now consider an absorptive polynomial semiring that will be used for provenance analysis later on. The following definitions are due to Grädel and Tannen [GT18].

Let X be a finite set of variables $\{x_1, ..., x_n\}$. A *monomial* over X with exponents in $\mathbb{N}^\infty$ is a function $m : X \to \mathbb{N}^\infty$. Informally, we write $x_1^{m(x_1)} \cdot ... \cdot x_n^{m(x_n)}$. Multiplication of monomials is defined by adding the exponents. An order on the monomials can be defined like

$$m_1 \leq m_2 \quad \text{iff} \quad m_1(x) \geq m_2(x) \quad \text{for all } x \in X.$$

We say that m_2 absorbs m_1 iff $m_1 \leq m_2$. Intuitively speaking, "shorter" monomials, that is, monomials with smaller exponents, absorb "longer" monomials with greater exponents.

(2.18) Definition (Absorptive Polynomial)**.** An absorptive polynomial over a finite set of variables X is an antichain of monomials over X with exponents in $\mathbb{N}^\infty$ with respect to the monomial order. $\mathbb{S}^\infty[X]$ denotes the set of all absorptive polynomials over X.

Addition and multiplication on absorptive polynomials are defined the usual way, but we only keep the maximal monomials of the result. Also, there are no coefficients for the monomials, duplicate monomials are only kept once. This ensures that the result of an addition or multiplication is still an antichain of monomials. While absorptive polynomials are in fact sets of monomials, we will often write them as sums of monomials. For example, if we have $X = \{x, y, z\}$, then

$$(x + y^3 + z) + (x^\infty + y^2 + z) = x + y^2 + z,$$

because x absorbs x^∞, y^2 absorbs y^3 and z occurs in both polynomials.

Grädel and Tannen have shown that $(\mathbb{S}^\infty[X], +, \cdot, 0, 1)$ is an absorptive semiring with addition and multiplication defined as above, where 0 is the empty antichain of monomials and 1 is the antichain that consists of a single monomial 1, which is the monomial where all exponents are 0. They have also shown that $\mathbb{S}^\infty[X]$ is naturally ordered and the natural order on absorptive polynomials can be characterized in terms of the monomial order as follows. For $P_1, P_2 \in \mathbb{S}^\infty[X]$, we have

$$P_1 \leq P_2 \quad \text{iff} \quad \text{for every } m_1 \in P_1, \text{ there is an } m_2 \in P_2 \text{ such that } m_1 \leq m_2.$$

They have also shown that $(\mathbb{S}^\infty[X], \leq)$ is a complete lattice and for any $S \subseteq \mathbb{S}^\infty[X]$,

$$\sup S = \text{maximals}\left(\bigcup_{P \in S} P\right),$$

where maximals(M) for a set of monomials M denotes the antichain of the maximal monomials in M with respect to monomial order. In other words, the supremum of an arbitrary set of absorptive polynomials can be obtained by computing the union of all the monomials in the polynomials and then picking the maximal polynomials out of this union.

Additionally, every absorptive polynomial is finite, that is, every absorptive polynomial contains only finitely many monomials according to Grädel and Tannen.

(2.19) Proposition. $\mathbb{S}^\infty[X]$ is an absorptive lattice semiring for any finite X.

Proof. We have to show that $\mathbb{S}^\infty[X]$ satisfies the conditions from definition (2.10). We already know that $\mathbb{S}^\infty[X]$ is an absorptive, naturally ordered semiring and that the natural order $(\mathbb{S}^\infty[X], \leq)$ is a complete lattice. It is left to show that $(\mathbb{S}^\infty[X], \leq)$ is completely distributive and the conditions (1) and (2) from the definition are satisfied.

For the complete distributivity, we show that for any $(I_j)_{j \in J}$ where J is a set of indices and $I_j \subseteq \mathbb{S}^\infty[X]$ for $j \in J$,

$$\inf_{j \in J} \sup I_j = \sup_{f \in F} \inf_{j \in J} f(j)$$

is fulfilled and F is the set of choice functions with $f(j) \in I_j$ for $j \in J$. As stated above, this is equivalent to

$$P \coloneqq \inf_{j \in J} \text{maximals}\left(\bigcup_{R \in I_j} R\right) = \text{maximals}\left(\bigcup_{f \in F} \inf_{j \in J} f(j)\right) =: Q.$$

"$\leq$": Let $m_p \in P$ be a monomial in P. This implies $m_p \leq \text{maximals}\left(\bigcup_{R \in I_j} R\right)$ for all $j \in J$, so there is a monomial $m_j \in \text{maximals}\left(\bigcup_{R \in I_j} R\right)$ with $m_p \leq m_j$. Also, m_j is contained in a polynomial R_j in I_j. Now, let $f_{m_p} \in F$ be a choice function such that for each $j \in J$, $f(j) \in I_j$ is the polynomial R_j that contains m_j with $m_p \leq m_j$. We know that such a polynomial R in I_j must exist for each $j \in J$. So, we have $m_p \leq f_{m_p}(j)$ for each $j \in J$, therefore $m_p \leq \inf_{j \in J} f_{m_p}(j)$ and there is a monomial $m \in \inf_{j \in J} f_{m_p}(j)$ that absorbs m_p. Since $f_{m_p} \in F$, we have a monomial $m_q \in Q$ that absorbs m. Thus, $m_p \leq m \leq m_q$ and for each $m_p \in P$ there is an $m_q \in Q$ which absorbs m_p, so we have $P \leq Q$.

"$\geq$": Consider a monomial $m_q \in Q$, then $m_q \in \bigcup_{f \in F} \inf_{j \in J} f(j)$. There is an $f_{m_q} \in F$ such that $m_q \in \inf_{j \in J} f_{m_q}(j)$, so $m_q \leq f_{m_q}(j)$ for each $j \in J$, which means that there is an $m_j \in f_{m_q}(j)$ with $m_q \leq m_j$. Since $f_{m_q}(j) \in I_j$, we have $m_j \in \bigcup_{R \in I_j} R$ and therefore, there is an $m'_j \in \text{maximals}\left(\bigcup_{R \in I_j} R\right)$ that absorbs m_j, so $m_q \leq m_j \leq m'_j$. Since this implies $m_q \leq \text{maximals}\left(\bigcup_{R \in I_j} R\right)$ for each $j \in J$,

we have $m_q \leq P$ and there is an $m_p \in P$ with $m_q \leq m_p$. This proves $Q \leq P$, because m_q was an arbitrary monomial in Q, and therefore $P = Q$.

Condition (1) of definition (2.10) states that for any $c \in \mathbb{S}^\infty[X]$ and $S \subseteq \mathbb{S}^\infty[X]$,

$$c \cdot \sup S = \sup(c \cdot S)$$

holds where $c \cdot S = \{c \cdot a \mid a \in S\}$.

"$\geq$": This direction can be proved using the monotonicity of multiplication. Since $\sup S \geq a$ for any $a \in S$, we have $c \cdot \sup S \geq c \cdot a$ for any $a \in S$ which immediately implies $c \cdot \sup S \geq \sup(c \cdot S)$.

"$\leq$": For this direction, we restate the equation as

$$P \coloneqq c \cdot \text{maximals}\left(\bigcup_{R \in S} R\right) = \text{maximals}\left(\bigcup_{R' \in c \cdot S} R'\right) =: Q.$$

Consider an arbitrary monomial $m_p \in P$. We know that $m_p = m_c \cdot m$ for two monomials $m_c \in c$ and $m \in \text{maximals}\left(\bigcup_{R \in S} R\right)$, so $m \in R$ for some $R \in S$. Now, $m_c \cdot m$ occurs in $c \cdot R$, but $c \cdot R$ is also a member of $c \cdot S$, so $m_c \cdot m$ occurs in $\left(\bigcup_{R' \in c \cdot S} R'\right)$ and therefore we have a monomial $m_q \in Q$ that absorbs $m_c \cdot m = m_p$. Since m_p was arbitrary, $P \leq Q$ is proven.

Finally, we prove condition (2) of definition (2.10). For $c \in \mathbb{S}^\infty[X]$ and any descending ω-chain $a_0 \geq a_1 \geq ...$ in $\mathbb{S}^\infty[X]$, we have to show

$$P \coloneqq c \cdot \inf_{i \in \omega} a_i = \inf_{i \in \omega}(c \cdot a_i) =: Q.$$

"$\leq$": Again, we can show this direction by using monotonicity of multiplication. We have $\inf_{i \in \omega} a_i \leq a_i$ for all $i \in \omega$, so $c \cdot \inf_{i \in \omega} a_i \leq c \cdot a_i$ for $i \in \omega$, which implies $c \cdot \inf_{i \in \omega} a_i \leq \inf_{i \in \omega}(c \cdot a_i)$.

"$\geq$": We will first prove the statement for the cases where c is empty or contains only one monomial and then use this to prove the general case. For empty c, $c = 0$ and both sides are 0. Now, assume that c only contains one monomial m_c. We will show that $Q = c \cdot Q'$ for some $Q' \in \mathbb{S}^\infty[X]$ and $Q' \leq \inf_{i \in \omega} a_i$. Clearly, because of monotonicity, this suffices to prove the statement.

We now have to find Q'. Intuitively, we would like to divide Q by c and we would obtain $Q' = \frac{Q}{c}$. However, as division is not defined, we will exploit the assumption that c consists of a single monomial m_c and the fact that Q, being an absorptive polynomial, has finitely many monomials, so w.l.o.g. we set $Q = m_1 + ... + m_k$ for some $k \in \mathbb{N}$. Our intuitive argument would now yield

$$Q' = \frac{Q}{c} = \frac{m_1 + ... + m_k}{m_c} = \frac{m_1}{m_c} + ... + \frac{m_k}{m_c}.$$

Of course, monomial division is undefined as well. However, monomial multiplication is defined by adding the exponents of the monomials, so we could define monomial division by subtracting the exponents. Recall that monomials m are defined as functions $m : X \to \mathbb{N}^\infty$ and monomials with smaller exponents are greater than

monomials with larger exponents with respect to monomial order. So, for $1 \leq i \leq k$, we would define $m_i' = \frac{m_i}{m_c}$ as

$$m_i'(x) = m_i(x) - m_c(x) \quad \text{for } x \in X.$$

However, subtraction is not defined yet in $\mathbb{N}^\infty$. First of all, we have to make sure that $m_i(x)$ is never less than $m_c(x)$. This is shown by observing that since $Q = \inf_{i\in\omega}(c \cdot a_i)$, we have $Q \leq c \cdot a_0 \leq c$. Therefore, any monomial m_i in Q is absorbed by some monomial in c. Since the only monomial in c is m_c, we have for each $1 \leq i \leq k$, that m_i is absorbed by m_c, so $m_i(x) \geq m_c(x)$ for each $x \in X$. Now, we can define the subtraction on $\mathbb{N}^\infty$. When subtracting two natural numbers, we use the normal subtraction. We set $\infty - n = \infty$ for each $n \in \mathbb{N}$ and $\infty - \infty = \infty$ as well. We do not have to define $n - \infty$ for $n \in \mathbb{N}$ since $n < \infty$, but we know that $m_i(x)$ is never less than $m_c(x)$. So, any $m_i'(x)$ is now well-defined and we can set $Q' = m_1' + ... + m_k'$.

First, we verify $Q = c \cdot Q'$. We have $c \cdot Q' = m_c \cdot (m_1' + ... + m_k') = m_c \cdot m_1' + ... + m_c \cdot m_k'$. We can show that for $1 \leq i \leq k$, $m_c \cdot m_i' = m_i$ by comparing exponents. For $x \in X$,

$$(m_c \cdot m_i')(x) = m_c(x) + m_i'(x) = m_c(x) + (m_i(x) - m_c(x)) \overset{(*)}{=} m_i(x).$$

The transformation $(*)$ can be shown by case distinction. If $m_i(x) \in \mathbb{N}$, then $m_c(x) \leq m_i(x)$ is also a natural number, therefore addition and subtraction are defined as normal and cancel each other out. Otherwise, $m_i(x) = \infty$, but in that case, regardless of $m_c(x)$, we have $(m_i(x) - m_c(x)) = \infty$ and therefore $m_c(x) + (m_i(x) - m_c(x)) = \infty = m_i(x)$ as well. This yields $c \cdot Q' = m_1 + ... + m_k = Q$.

It remains to show that $Q' \leq \inf_{i\in\omega} a_i$. For that, consider an arbitrary $i \in \omega$. By the definition of Q, we know that $Q \leq c \cdot a_i$. So, for each of the monomials m_j in Q with $1 \leq j \leq k$, there is a monomial $m \in c \cdot a_i$ with $m_j \leq m$. Since c only consists of a single monomial, every monomial in $c \cdot a_i$ is the product of m_c with some monomial in a_i, therefore we have $m = m' \cdot m_c$ for some $m' \in a_i$. So, we can express the exponents of m as $m(x) = m'(x) + m_c(x)$ for all $x \in X$. Since $m_j \leq m$, we also have

$$m_j(x) \geq m(x) = m'(x) + m_c(x)$$

for all $x \in X$. We claim that this implies

$$m_j'(x) = m_j(x) - m_c(x) \geq m'(x)$$

for all $x \in X$. This is again shown by case distinction. If $m_j(x) = \infty$, then $m_j'(x) = m_j(x) - m_c(x) = \infty$ regardless of $m_c(x)$, and therefore $m_j'(x) \geq m'(x)$ regardless of $m'(x)$. Otherwise, $m_j(x) \in \mathbb{N}$ and since $m_j(x) \geq m'(x) + m_c(x)$, both $m'(x)$ and $m_c(x)$ are natural numbers as well and the subtraction is defined as usual. Over the natural numbers, the inequality $m_j(x) \geq m'(x) + m_c(x)$ is equivalent to $m_j(x) - m_c(x) \geq m'(x)$ and the claim is proven. Since $m_j'(x) \geq m'(x)$ for all $x \in X$, we have $m_j' \leq m'$ and m' was in a_i, so for every m_j', there is a monomial in a_i that absorbs m_j'. This implies $Q' = m_1' + ... + m_j' \leq a_i$. Since i was arbitrary, we have shown $Q' \leq \inf_{i\in\omega} a_i$ which ends the case where c has one monomial.

Now, suppose c has more than one monomial. We know that since $c \in \mathbb{S}^\infty[X]$, c is the sum of finitely many monomials. We assume w.l.o.g. that $c = c_1 + ... + c_k$ for

some $k \in \mathbb{N}$ where $c_1, ..., c_k$ are single monomials. We can rewrite

$$\begin{aligned} Q &= \inf_{i\in\omega}(c \cdot a_i) \\ &= \inf_{i\in\omega}((c_1 + ... + c_k) \cdot a_i) \\ &= \inf_{i\in\omega}(c_1 \cdot a_i + ... + c_k \cdot a_i) \\ &= \inf_{i\in\omega} \sup_{1\leq j\leq k} c_j \cdot a_i. \end{aligned}$$

In the last step, we have replaced summation with a supremum by applying lemma (2.8). Since we have already shown that $\mathbb{S}^\infty[X]$ is completely distributive, we have

$$Q = \inf_{i\in\omega} \sup_{1\leq j\leq k} c_j \cdot a_i = \sup_{f\in F} \inf_{i\in\omega} f(i)$$

where F is the set of choice functions $f : \omega \to K$ such that $f(i) = c_l \cdot a_i$ for some $1 \leq l \leq k$. On the other side, we have

$$\begin{aligned} P &= c \cdot \inf_{i\in\omega} a_i \\ &= (c_1 + ... + c_k) \cdot \inf_{i\in\omega} a_i \\ &= c_1 \cdot \inf_{i\in\omega} a_i + ... + c_k \cdot \inf_{i\in\omega} a_i \\ &\overset{(*)}{=} \inf_{i\in\omega}(c_1 \cdot a_i) + ... + \inf_{i\in\omega}(c_k \cdot a_i) \\ &= \sup_{1\leq j\leq k} \inf_{i\in\omega} (c_j \cdot a_i). \end{aligned}$$

For $(*)$, we have already shown that this transformation works for single monomials $c_1, ..., c_k$. So, it is left to show that

$$P = \sup_{1\leq j\leq k} \inf_{i\in\omega} (c_j \cdot a_i) \geq \sup_{f\in F} \inf_{i\in\omega} f(i) = Q.$$

It suffices to prove that for every $f \in F$, there is a j with $1 \leq j \leq k$ such that $\inf_{i\in\omega} f(i) \leq \inf_{i\in\omega} c_j \cdot a_i$. We fix an arbitrary $f \in F$. Recall that for $i \in \omega$, $f(i) = c_l \cdot a_i$ for some $1 \leq l \leq k$. Intuitively, we could say that for each $i \in \omega$, f picks some $l \in \{1, ..., k\}$. Since $\{1, ..., k\}$ is finite, there is an element $j \in \{1, ..., k\}$ that is chosen infinitely many times, so for each $i \in \omega$, there is an $i' \geq i$ such that $f(i') = c_j \cdot a_{i'}$. To verify that $\inf_{i\in\omega} f(i) \leq \inf_{i\in\omega} c_j \cdot a_i$, consider an arbitrary $i \in \omega$. We know that there is an $i' \geq i$ with $f(i') = c_j \cdot a_{i'}$. Since $a_0 \geq a_1 \geq ...$ is a descending chain, we have $a_{i'} \leq a_i$ and monotonicity yields $f(i') \leq c_j \cdot a_i$, therefore, $\inf_{i\in\omega} f(i) \leq c_j \cdot a_i$. Since i was arbitrary, $\inf_{i\in\omega} f(i)$ is a lower bound on $\{c_j \cdot a_i \mid i \in \omega\}$, which ends the proof. □

There are other examples of absorptive lattice semirings. Let $(K, \leq)$ be a completely distributive lattice. The commutative semiring $(K, \vee, \wedge, \bot, \top)$ can be built by defining the constants and operations as follows for $a, b \in K$:

$$\begin{aligned} a \vee b &= \sup\{a, b\}, \\ a \wedge b &= \inf\{a, b\}, \\ \bot &= \sup \emptyset \quad \text{and} \\ \top &= \inf \emptyset. \end{aligned}$$

(2.20) Proposition. Let $(K, \leq)$ be a completely distributive lattice. The induced semiring $(K, \vee, \wedge, \bot, \top)$ is an absorptive lattice semiring.

Proof. Absorption is verified by $a \vee (a \wedge b) = \sup\{a, \inf\{a, b\}\} = a$.

K is also naturally ordered and the natural order is the same as the lattice order. To see this, we show for $a, b \in K$ that $a \leq b$ iff there is a $c \in K$ such that $a \vee c = b$.

"$\Rightarrow$": If $a \leq b$, then we have $a \vee b = \sup\{a, b\} = b$.

"$\Leftarrow$": If $a \vee c = b$ for some $c \in K$, then $\sup\{a, c\} = b$, which implies $a \leq b$.

Clearly, the natural order forms a completely distributive lattice. It is left to verify the conditions (1) and (2) from definition (2.10). For condition (1), we show for any $c \in K$ and $S \subseteq K$ that

$$c \wedge \sup S = \sup(c \wedge S)$$

with $c \wedge S = \{c \wedge a \mid a \in S\}$. Notice that $\wedge$ yields the infimum, so we set $J = \{0, 1\}$ and $I_0 = \{c\}$ and $I_1 = S$. Then, we have

$$c \wedge \sup S = \inf\{\sup\{c\}, \sup S\} = \inf_{j \in J} \sup I_j.$$

Complete distributivity implies

$$c \wedge \sup S = \inf_{j \in J} \sup I_j = \sup_{f \in F} \inf_{j \in J} f(j)$$

where F is the set of choice functions with $f(j) \in I_j$. In particular, we have $f(0) = c$ for any $f \in F$ and $f(1) \in S$. So, for every $c \wedge a \in c \wedge S$, we have a $f \in F$ with $f(1) = a$ and $\inf_{j \in J} f(j) = \inf\{c, a\} = c \wedge a$ and for every $f \in F$ we have $\inf_{j \in J} f(j) = \inf\{c, f(1)\} = c \wedge f(1) \in c \wedge S$. This yields

$$c \wedge \sup S = \sup_{f \in F} \inf_{j \in J} f(j) = \sup(c \wedge S).$$

Condition (2) dictates that for every $c \in K$ and any descending chain $a_0 \geq a_1 \geq ...$ in K, we have

$$c \wedge \inf_{i \in \omega} a_i = \inf_{i \in \omega}(c \wedge a_i).$$

Since $\wedge$ is the infimum, we rewrite to

$$\inf\left\{c, \inf_{i \in \omega} a_i\right\} = \inf_{i \in \omega}(\inf\{c, a_i\}).$$

"$\leq$": Clearly, $\inf\{c, \inf_{i \in \omega} a_i\} \leq c$ and $\inf\{c, \inf_{i \in \omega} a_i\} \leq \inf_{i \in \omega} a_i \leq a_i$ for each $i \in \omega$, so $\inf\{c, \inf_{i \in \omega} a_i\} \leq \inf\{c, a_i\}$ for each $i \in \omega$ and therefore, we have shown $\inf\{c, \inf_{i \in \omega} a_i\} \leq \inf_{i \in \omega}(\inf\{c, a_i\})$.

"$\geq$": First, we observe $\inf_{i \in \omega}(\inf\{c, a_i\}) \leq \inf\{c, a_0\} \leq c$. Moreover, we have $\inf_{i \in \omega}(\inf\{c, a_i\}) \leq \inf\{c, a_i\} \leq a_i$ for each $i \in \omega$, so $\inf_{i \in \omega}(\inf\{c, a_i\}) \leq \inf_{i \in \omega} a_i$, which yields $\inf_{i \in \omega}(\inf\{c, a_i\}) \leq \inf\{c, \inf_{i \in \omega} a_i\}$.

This verifies condition (2) and ends the proof. □

We observe that we can obtain an absorptive lattice semiring from a completely distributive lattice by defining the addition as the supremum in the lattice and

multiplication as the infimum in the lattice. Lemma (2.8) also guarantees that the addition in an absorptive lattice semiring is always the supremum in the underlying lattice. However, a similar proposition for the multiplication does not hold. The absorptive lattice semiring $\mathbb{S}^{\infty}[X]$ is a counterexample that shows that multiplication is not always the same as the infimum in the underlying lattice, since for $x \in X$, we have $x \cdot x = x^2 \neq x = \inf\{x, x\}$. We conclude that every completely distributive lattice induces an absorptive lattice semiring, but the multiplication is not uniquely defined by the underlying lattice and there are absorptive lattice semirings that are not constructed by means of proposition (2.20). In particular, the multiplication in absorptive lattice semirings is not always idempotent.

It is an open question whether any of the conditions for absorptive lattice semirings that we stated in definition (2.10) is redundant. In particular, it is not known whether every absorptive, naturally ordered semiring K is an absorptive lattice semiring.

2.3 Families over Semirings

In the previous sections, we have established classes of semirings where functions $f : K \to K$ have fixed points under specific conditions. However, we will need fixed points of functions that operate on multiple elements of K at the same time. In this section, we will generalize the results from above to these functions.

Let I be an arbitrary set of indices and K a semiring. K^I denotes the set of *families* $(a_i)_{i\in I}$ with $a_i \in K$ for $i \in I$. Assume that K is naturally ordered by $\leq$. We can define an order on K^I as follows. For $(a_i)_{i\in I}, (b_i)_{i\in I} \in K^I$,

$$(a_i)_{i\in I} \leq (b_i)_{i\in I} \quad \text{iff} \quad a_i \leq b_i \text{ for all } i \in I.$$

(2.21) Lemma. Let K be a naturally ordered semiring and I a set of indices. For an arbitrary subset $S \subseteq K^I$, let $S_i = \{a_i \mid (a_i)_{i\in I} \in S\} \subseteq K$ for $i \in I$. If each S_i has a supremum (an infimum) in K, then S has a supremum (an infimum) in K^I and

$$\sup S = (\sup S_i)_{i\in I} \quad \text{or} \quad \inf S = (\inf S_i)_{i\in I} \quad \text{respectively.}$$

Proof. We will just show the lemma for suprema, the proof for infima is dual. Since $(\sup S_i)_{i\in I}$ exists, we only have to verify that this is indeed the supremum of S. We first show that it is an upper bound. Consider a family $(a_i)_{i\in I} \in S$. For each $i \in I$, we have $a_i \in S_i$ and therefore $a_i \leq \sup S_i$, so $(a_i)_{i\in I} \leq (\sup S_i)_{i\in I}$.

Now, let $u = (u_i)_{i\in I} \in K^I$ be an upper bound of S. Let $i \in I$ be arbitrary and consider an arbitrary $a_i \in S_i$. Since a_i belongs to some family $(a_j)_{j\in I} \in S$ and $(a_j)_{j\in I} \leq u$, we have in particular $a_i \leq u_i$, so u_i is an upper bound on S_i, since a_i was arbitrary. This implies $\sup S_i \leq u_i$ and since i was arbitrary, we have $(\sup S_i)_{i\in I} \leq u$, so $(\sup S_i)_{i\in I}$ is the least upper bound on S. □

(2.22) Corollary. For a naturally ordered semiring K and a set of indices I, if every subset $S \subseteq K$ has a supremum (an infimum) in K, then every subset $S' \subseteq K^I$ has a supremum (an infimum) in K^I which can be obtained component-wise. Moreover, if every ascending ω-chain $a_0 \leq a_1 \leq ...$ in K has a supremum in K, then every

ascending ω-chain $(a_0)_{i\in I} \leq (a_1)_{i\in I} \leq ...$ in K^I has a supremum in K^I which can be obtained component-wise.

Proof sketch. The first part of the corollary is a direct consequence of lemma (2.21). We will just verify the second part of the corollary by observing that for an ascending ω-chain $(a_0)_{i\in I} \leq (a_1)_{i\in I} \leq ...$ in K^I, we can write the chain as a set $S = \{(a_j)_{i\in I} \mid j \in \omega\}$. But then, lemma (2.21) can be applied, because the sets $S_i = \{(a_j)_i \mid j \in \omega\}$ form ascending ω-chains, because $(a_0)_{i\in I} \leq (a_1)_{i\in I} \leq ...$ implies $(a_0)_i \leq (a_1)_i \leq ...$ for each $i \in I$. □

The above statements allow us to extend the fixed-point theorems for ω-continuous and absorptive lattice semirings K to families over K.

(2.23) Theorem. Let K be an ω-continuous semiring and I a set of indices. If a function $f : K^I \to K^I$ is component-wise ω-continuous in each argument, then f has a least fixed point $\mathrm{lfp}(f)$ in K^I and

$$\mathrm{lfp}(f) = \sup_{i\in\omega} f^i(0).$$

Proof. Due to corollary (2.22), we already know that K^I is ω-complete [Bar91], that is, ascending chains in K^I have suprema. We want to apply Kleene's fixed-point theorem to obtain the desired result, so we first have to show that f is ω-continuous, that is, for any ascending ω-chain $(a_0)_{i\in I} \leq (a_1)_{i\in I} \leq ...$ in K^I, we have

$$\sup_{j\in\omega} f((a_j)_{i\in I}) = f(\sup_{j\in\omega}(a_j)_{i\in I}).$$

We can split f into its components $f_k : K^I \to K$ for $k \in I$ and obtain

$$\sup_{j\in\omega} f((a_j)_{i\in I}) = \sup_{j\in\omega}(f_k((a_j)_{i\in I}))_{k\in I}.$$

Since suprema are calculated component-wise, we have

$$\sup_{j\in\omega}(f_k((a_j)_{i\in I}))_{k\in I} = \left(\sup_{j\in\omega} f_k((a_j)_{i\in I})\right)_{k\in I}.$$

As every component f_k of f is ω-continuous in each argument, we can write

$$\left(\sup_{j\in\omega} f_k((a_j)_{i\in I})\right)_{k\in I} = \left(f_k((\sup_{j\in\omega} a_{j,i})_{i\in I})\right)_{k\in I}.$$

Now, we reverse the component-wise supremum to

$$\left(f_k((\sup_{j\in\omega} a_{j,i})_{i\in I})\right)_{k\in I} = \left(f_k(\sup_{j\in\omega}(a_j)_{i\in I})\right)_{k\in I}$$

and put the components f_k of f together to finally obtain

$$\left(f_k(\sup_{j\in\omega}(a_j)_{i\in I})\right)_{k\in I} = f(\sup_{j\in\omega}(a_j)_{i\in I}).$$

This ends the proof as Kleene's fixed-point theorem can now be applied. □

In particular, a function $f : K^I \to K^I$ on families over an ω-continuous semiring K whose components are only composed of addition and multiplication with constants and variables is component-wise ω-continuous in each argument and therefore has a least fixed point in K^I. We can show a similar result for absorptive lattice semirings.

(2.24) Theorem. Let K be an absorptive lattice semiring and I a set of indices, then K^I is a complete lattice with respect to the order induced by the natural order of K and any function $f : K^I \to K^I$ that is component-wise monotonic in each argument with respect to the natural order on K is also monotonic in K^I.

Proof. Since the natural order on K forms a complete lattice, it admits arbitrary suprema and infima. Then, by corollary (2.22), the induced order on K^I also admits arbitrary suprema and infima, which makes it a complete lattice. Now, consider a function $f : K^I \to K^I$ that is component-wise monotonic in each argument. We claim that for $(a_i)_{i\in I}, (b_i)_{i\in I} \in K^I$,

$$(a_i)_{i\in I} \leq (b_i)_{i\in I} \quad \text{implies} \quad f((a_i)_{i\in I}) \leq f((b_i)_{i\in I}).$$

Assume $(a_i)_{i\in I} \leq (b_i)_{i\in I}$ and split f in components $f_k : K^I \to K$ for $k \in I$, then

$$(f_k((a_i)_{i\in I}))_{k\in I} \leq (f_k((b_i)_{i\in I}))_{k\in I}$$

has to be shown. We verify this component-wise for $k \in I$. Since $a_i \leq b_i$ for each $i \in I$, each argument of $f_k((a_i)_{i\in I})$ is bounded by the corresponding argument of $f_k((b_i)_{i\in I})$. Since f_k is monotonic in each argument, this implies the desired statement $f_k((a_i)_{i\in I}) \leq f_k((b_i)_{i\in I})$ and ends the proof. □

The Knaster-Tarski fixed-point theorem yields the following corollary.

(2.25) Corollary. For an absorptive lattice semiring K and an index set I, any monotonic function $f : K^I \to K^I$ has a least fixed point $\mathrm{lfp}(f)$ and a greatest fixed point $\mathrm{gfp}(f)$ in K^I and they can by obtained the same way as in theorem (2.17).

We have seen that the results from this chapter can be extended to families over the appropriate semirings and we will now use ω-continuous semirings whenever we need least fixed points and absorptive lattice semirings whenever we need both least and greatest fixed points. The infinitary summations and multiplications that were defined in this chapter will prove to be useful as well.

Chapter 3

Semiring Interpretations for Logics

It is our goal to define semiring interpretations for LTL, CTL and a fragment of PDL, which we will do separately for each logic. However, we will first explain the syntax and the usual two-value semantics of these logics before extending them to semirings. Finally, we will use the polynomial semirings defined in the previous chapter to provide examples for provenance analysis.

3.1 Linear Temporal Logic (LTL)

Linear Temporal Logic (LTL) is interpreted over path structures $\mathcal{W} = (W, <, (P_i)_{i \in I})$ where $<$ is a linear order, I is an index set and P_i are atomic propositions which can be true or false at any $w \in W$. We can see the P_i as unary relations. In provenance analysis, we are dealing with finite inputs, so we assume W to be finite, therefore we set $|W| = n \in \mathbb{N}$. We notice that $\mathcal{W}$ is isomorphic to a unique path structure $(\{0, ..., n-1\}, <, (P_i)_{i \in I})$, so it is justified to assume w.l.o.g. that $W = \{0, ..., n-1\}$. We call $\tau = \{<\} \cup \{P_i \mid i \in I\}$ the *signature* of $\mathcal{W}$ and we will often use $P, Q, ...$ as atomic proposition symbols.

LTL is called a temporal logic, because formulas are interpreted at some node $j \in \{0, ..., n-1\}$ and we can imagine j to be the current state and the successor nodes along the path $\mathcal{W}$ to be future states that LTL can reason about.

First of all, we will define the syntax and standard semantics for LTL. These definitions are based on Bauer, Leucker and Schallhart's work on LTL [BLS10]. Although they define LTL with a different set of operators than usual, they make sure that their set of operators admits a simple transformation to negation normal form, which is why it is very useful for provenance analysis.

(3.1) Definition (Syntax of LTL)**.** Given a signature τ, the formulas $\varphi \in \mathrm{LTL}(\tau)$ are generated by the grammar

$$\varphi = 0 \mid 1 \mid P_i \mid \neg\varphi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \mathrm{X}\varphi \mid \overline{\mathrm{X}}\varphi \mid \varphi \mathrm{U} \varphi \mid \varphi \mathrm{R} \varphi \quad \text{with } P_i \in \tau.$$

We also use the abbreviations $\mathrm{F}\psi = 1\mathrm{U}\psi$ and $\mathrm{G}\psi = 0\mathrm{R}\psi$ for $\psi \in \mathrm{LTL}(\tau)$.

As for the semantics, formulas $\varphi \in \mathrm{LTL}(\tau)$ are interpreted in matching path structures $\mathcal{W}$. Assuming $|W| = n$, we write $\mathcal{W}, j \models \varphi$ if φ is true at the node $j < n$ in

$\mathcal{W}$, otherwise, we write $\mathcal{W}, j \not\models \varphi$. We can now define the semantics inductively.

(3.2) Definition (Semantics of LTL)**.** Let $\mathcal{W} = (W, \tau)$ be a path structure with $|W| = n$. The formula 0 is always false and 1 is always true, so we have

$$\mathcal{W}, j \not\models 0 \quad \text{and} \quad \mathcal{W}, j \models 1 \quad \text{for all } j \in \{0, ..., n-1\}.$$

The semantics of P_i are given by the structure $\mathcal{W}$. We write $P_i^{\mathcal{W}} \subseteq W$ for the subset of all nodes where the atomic proposition P_i holds, then we have

$$\mathcal{W}, j \models P_i \quad \text{iff} \quad j \in P_i^{\mathcal{W}}.$$

For the compound formulas, let φ and ψ be in LTL(τ). The boolean connectives $\neg$, $\vee$ and $\wedge$ are interpreted as usual, that is,

$$\begin{aligned}
&\mathcal{W}, j \models \neg\varphi && \text{iff} \quad \mathcal{W}, j \not\models \varphi, \\
&\mathcal{W}, j \models \varphi \vee \psi && \text{iff} \quad \mathcal{W}, j \models \varphi \text{ or } \quad \mathcal{W}, j \models \psi \quad \text{and} \\
&\mathcal{W}, j \models \varphi \wedge \psi && \text{iff} \quad \mathcal{W}, j \models \varphi \text{ and } \mathcal{W}, j \models \psi.
\end{aligned}$$

The operator X is called the *next operator*. The formula $\mathrm{X}\varphi$ says that φ holds at the successor node of the current node. Notice that this implies the existence of a successor. Therefore, the dual operator $\overline{\mathrm{X}}$, which we will call the *weak next operator*, does not make this assertion. The semantics are

$$\begin{aligned}
&\mathcal{W}, j \models \mathrm{X}\varphi && \text{iff} \quad j+1 < n \text{ and } \mathcal{W}, j+1 \models \varphi \quad \text{and} \\
&\mathcal{W}, j \models \overline{\mathrm{X}}\varphi && \text{iff} \quad j+1 = n \text{ or } \quad \mathcal{W}, j+1 \models \varphi.
\end{aligned}$$

The remaining two operators U and R are called *until operator* and *release operator* respectively. Intuitively speaking, the formula $\varphi \mathrm{U} \psi$ asserts that φ holds true until eventually, ψ comes true at some point in the future. The release formula $\varphi \mathrm{R} \psi$ says that ψ must stay true unless it is released by φ, that is, ψ may only become false after φ was true. Figure (3.3) below illustrates some models for these formulas when evaluating at node j. The semantics are given by

$$\begin{aligned}
\mathcal{W}, j \models \varphi \mathrm{U} \psi \quad \text{iff} \quad & \mathcal{W}, k \models \psi \text{ for some } j \leq k \ < n \text{ such that} \\
& \mathcal{W}, l \models \varphi \text{ for all } \quad j \leq l \ < k \qquad \text{and} \\
\mathcal{W}, j \models \varphi \mathrm{R} \psi \quad \text{iff} \quad & \mathcal{W}, m \models \psi \text{ for all } \quad j \leq m < n \text{ or instead,} \\
& \mathcal{W}, k \models \varphi \text{ for some } j \leq k \ < n \text{ with} \\
& \mathcal{W}, l \models \psi \text{ for all } \quad j \leq l \ \leq k.
\end{aligned}$$

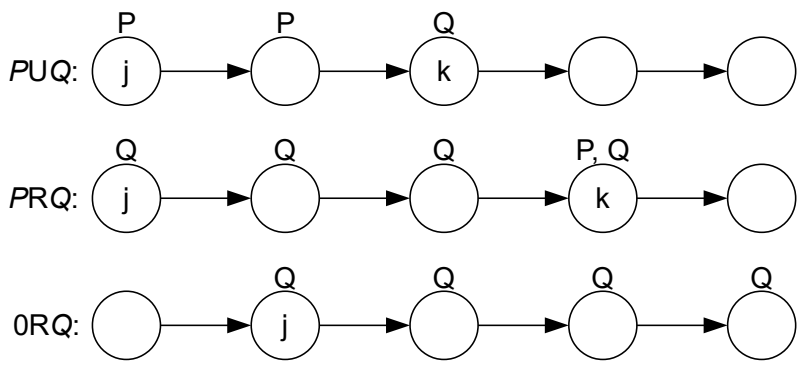

Figure (3.3): Models of until and release formulas in LTL.

Consequently, F is the *finally operator* and $\mathrm{F}\psi = 1\mathrm{U}\psi$ states that ψ will eventually come true. It is a special case of $\varphi\mathrm{U}\psi$ where there is no precondition φ. Similarly, G is the *globally operator* and $\mathrm{G}\psi = 0\mathrm{R}\psi$ states that ψ stays true forever, since it is a special case of $\varphi\mathrm{R}\psi$ without a release condition φ. For example, in figure (3.3), the first model is also a model of $\mathrm{F}Q$ and the third model also fulfills $\mathrm{G}Q$.

Before defining our semiring interpretation for LTL, we observe that this definition of LTL admits a negation normal form. For $\varphi, \psi \in \mathrm{LTL}(\tau)$, the equivalences

$$\begin{aligned} \neg 0 &\equiv 1, & \neg 1 &\equiv 0, \\ \neg(\varphi \vee \psi) &\equiv (\neg\varphi) \wedge (\neg\psi), & \neg(\varphi \wedge \psi) &\equiv (\neg\varphi) \vee (\neg\psi), \\ \neg(\mathrm{X}\varphi) &\equiv \overline{\mathrm{X}}(\neg\varphi), & \neg(\overline{\mathrm{X}}\varphi) &\equiv \mathrm{X}(\neg\varphi), \\ \neg(\varphi\mathrm{U}\psi) &\equiv (\neg\varphi)\mathrm{R}(\neg\psi) & \neg(\varphi\mathrm{R}\psi) &\equiv (\neg\varphi)\mathrm{U}(\neg\psi) \end{aligned}$$

hold. Now, as announced in the introduction, we will use the same approach as Grädel and Tannen to interpret LTL formulas in semirings. Let K be a commutative semiring. For LTL, this will suffice. Let τ be a signature and W the finite universe of a path structure. The set of literals $\mathrm{Lit}_W(\tau)$ is defined as

$$\mathrm{Lit}_W(\tau) = \{P_i w \mid P_i \in \tau, w \in W\} \cup \{\neg P_i w \mid P_i \in \tau, w \in W\}.$$

A *K-interpretation* is a function $\pi : \mathrm{Lit}_W(\tau) \to K$ that maps literals to values in K. Under π, any formula $\varphi \in \mathrm{LTL}(\tau)$ can be interpreted at any node $w \in W$.

(3.4) Definition (Semiring Interpretation for LTL)**.** Let K be a commutative semiring, $W = \{0, ..., n-1\}$ for $n \in \mathbb{N}$ a finite universe and τ a signature. Given a K-interpretation π, we can interpret a formula $\theta \in \mathrm{LTL}(\tau)$ at node $j \in W$ to obtain a semiring value $\pi[\![\theta]\!]_j \in K$, thereby extending π to a function $\mathrm{LTL}(\tau) \times W \to K$. We define this semiring interpretation inductively on the negation normal form of θ. Let $P_i \in \tau$ and φ and ψ be formulas in $\mathrm{LTL}(\tau)$, then we set

$$\begin{aligned} \pi[\![0]\!]_j &= 0, & \pi[\![1]\!]_j &= 1, \\ \pi[\![P_i]\!]_j &= \pi(P_i j), & \pi[\![\neg P_i]\!]_j &= \pi(\neg P_i j), \\ \pi[\![\varphi \vee \psi]\!]_j &= \pi[\![\varphi]\!]_j + \pi[\![\psi]\!]_j, & \pi[\![\varphi \wedge \psi]\!]_j &= \pi[\![\varphi]\!]_j \cdot \pi[\![\psi]\!]_j, \end{aligned}$$

$$\pi[\![\mathrm{X}\varphi]\!]_j = \begin{cases} \pi[\![\varphi]\!]_{j+1} & \text{if } j+1 < n \\ 0 & \text{otherwise,} \end{cases}$$

$$\pi[\![\overline{\mathrm{X}}\varphi]\!]_j = \begin{cases} \pi[\![\varphi]\!]_{j+1} & \text{if } j+1 < n \\ 1 & \text{otherwise,} \end{cases}$$

$$\pi[\![\varphi\mathrm{U}\psi]\!]_j = \sum_{j \le k < n} \left(\pi[\![\psi]\!]_k \cdot \prod_{j \le l < k} \pi[\![\varphi]\!]_l \right) \qquad \text{and}$$

$$\pi[\![\varphi\mathrm{R}\psi]\!]_j = \prod_{j \le m < n} \pi[\![\psi]\!]_m + \sum_{j \le k < n} \left(\pi[\![\varphi]\!]_k \cdot \prod_{j \le l \le k} \pi[\![\psi]\!]_l \right).$$

Since the semiring $(K, +, \cdot, 0, 1)$ has constants for 0 and 1, the first two definitions are not surprising. Also, as stated in the introduction, disjunctions are interpreted

as additions and conjunctions are interpreted as multiplications. The interpretations of the atomic propositions are given by π directly. As for the next operators $\mathrm{X}\varphi$ and $\overline{\mathrm{X}}\varphi$, we evaluate the formulas at the next node, if present. Otherwise, the normal next operator defaults to 0 whereas the weak next operator defaults to 1.

For the until operator, we look at its semantics and notice that it can be expressed as a disjunction. There can be any k with $j \leq k < n$ where ψ becomes true, which is a disjunction over $j \leq k < n$. Also, for this k, φ must be true for all $j \leq l < k$ to fulfill the formula, so we have an inner conjunction over $j \leq l < k$. Translating disjunctions to summations and conjunctions to multiplications yields the above definition. The same approach is used for the release operator.

As a further sanity check, for $j \in W$ and $\psi \in \mathrm{LTL}(\tau)$, we can use the above definitions to evaluate finally and globally operators as well, which yields exactly the expected results

$$\pi[\![\mathrm{F}\psi]\!]_j = \pi[\![1\mathrm{U}\psi]\!]_j = \sum_{j \leq k < n} \pi[\![\psi]\!]_k \quad \text{and} \quad \pi[\![\mathrm{G}\psi]\!]_j = \pi[\![0\mathrm{R}\psi]\!]_j = \prod_{j \leq m < n} \pi[\![\psi]\!]_m.$$

Assuming that addition and multiplication can be performed in constant time, the until operator and the release operator can be interpreted in $\mathcal{O}(n^2)$ whereas the remaining operators can be evaluated in constant time. We will close the section about LTL with an example.

(3.5) Example. Let P and Q be atomic propositions and consider the $\mathbb{N}[X]$-interpretation π with $X = \{p, q, r, s, x, y\}$ given in figure (3.6).

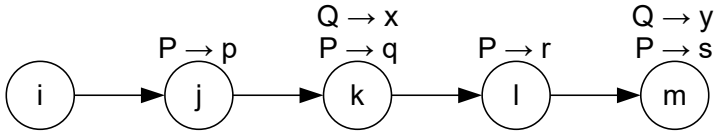

Figure (3.6): The $\mathbb{N}[X]$-interpretation π.

Normally, we would represent structures graphically by tagging the nodes with the atomic propositions they fulfill. For K-interpretations, we also need to know which values in K are assigned to the literals, so $P \to p$ at node j in the above example means that $\pi(Pj) = p$. Any literals whose values are not explicitly given in the figure, including all negative literals, are assumed to be tagged with 0.

We would like to interpret the formula $\varphi = P\mathrm{U}Q$ under π at node j. Since φ is in negation normal form, we can use definition (3.4) directly and obtain

$$\pi[\![P\mathrm{U}Q]\!]_j = px + pqry,$$

since k and m are the only nodes where Q is tagged with a nonzero value.

We can interpret another formula $\psi = \neg\mathrm{F}(\neg P)$ under π at j. First, we translate it into its negation normal form by

$$\neg\mathrm{F}(\neg P) = \neg(1\mathrm{U}(\neg P)) \equiv (\neg 1)\mathrm{R}(\neg(\neg P)) \equiv 0\mathrm{R}P = \mathrm{G}P.$$

Now, we can calculate the interpretation

$$\pi[\![\neg\mathrm{F}(\neg P)]\!]_j = \pi[\![\mathrm{G}P]\!]_j = pqrs.$$

3.2 Computation Tree Logic (CTL)

Since LTL was evaluated on paths, from every node on the path, there was only one "past" and one "future" sequence of nodes and there were no cycles. Computation Tree Logic (CTL) introduces branching. The nodes do not have to be arranged linearly, but instead, there may be multiple predecessors and successors for a single node and there could also be cycles.

Formally, CTL is interpreted on transition systems $\mathcal{G} = (V, E, (P_i)_{i \in I})$ where P_i are atomic propositions and I is an index set like for LTL. However, we now have a binary relation $E \subseteq V \times V$ that represents the edges between the nodes in V. We assume V to be finite, but paths over V can still be infinite because of possible cycles. We also assume $\mathcal{G}$ to be *non-terminating*, that is, for each $v \in V$, there is a successor $w \in V$ with $(v, w) \in E$. The signature of $\mathcal{G}$ is $\tau = \{E\} \cup \{P_i \mid i \in I\}$. Syntax and standard semantics of CTL are adapted from a book by Huth and Ryan. [HR00].

(3.7) Definition (Syntax of CTL)**.** Let τ be a signature, the formulas in CTL(τ) are generated by the symbol φ of the grammar

$$\begin{aligned} &\varphi = 0 \mid 1 \mid P_i \mid \neg\varphi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \mathrm{E}(\phi) \mid \mathrm{A}(\phi) \quad \text{with } P_i \in \tau, \\ &\phi = \mathrm{X}\varphi \mid \varphi \mathrm{U} \varphi \mid \varphi \mathrm{R} \varphi. \end{aligned}$$

The formulas generated by the symbol ϕ are called *path formulas* and the next, until and release operators from LTL have the same meaning here. However, the path formulas are no valid CTL formulas and they can only appear directly after a *path quantifier* in CTL, where E is the existential and A the universal path quantifier. So, we could eliminate the symbol ϕ and include all six possible formulas formed by $\mathrm{E}(\phi)$ and $\mathrm{A}(\phi)$ directly, that would be

$$\mathrm{E}(\mathrm{X}\varphi), \mathrm{A}(\mathrm{X}\varphi), \mathrm{E}(\varphi \mathrm{U} \varphi), \mathrm{A}(\varphi \mathrm{U} \varphi), \mathrm{E}(\varphi \mathrm{R} \varphi) \text{ and } \mathrm{A}(\varphi \mathrm{R} \varphi).$$

For $\psi \in \mathrm{CTL}(\tau)$, the abbreviations $\mathrm{F}\psi = 1\mathrm{U}\psi$ and $\mathrm{G}\psi = 0\mathrm{R}\psi$ are still valid for path formulas. With the above observations and the LTL semantics in mind, it is straightforward to define semantics for CTL inductively.

(3.8) Definition (Semantics of CTL)**.** Let $\mathcal{G} = (V, \tau)$ be a transition system. We always evaluate a formula $\theta \in \mathrm{CTL}(\tau)$ at a node $v \in V$. The formulas 0, 1, P_i for $P_i \in \tau$ and the boolean connectives $\neg$, $\vee$, $\wedge$ are interpreted as in LTL.

Let φ and ψ be formulas in CTL(τ). We first define the semantics of path formulas ϕ. Let $\pi = v_0 v_1 ...$ be an infinite path in $\mathcal{G}$, that is, $v_0, v_1, ... \in V$ and $(v_i, v_{i+1}) \in E$ for $i \in \omega$. We write $\pi \models \phi$ iff the path π fulfills the path formula ϕ.

The only difference between the LTL semantics and the evaluation of ϕ on π is the fact that π is infinite, whereas we assumed LTL path structures $\mathcal{W}$ to be finite. However, the next, until and release operators still work the same way on infinite

paths with only slight changes. For $\pi = v_0 v_1 ...$, we set

$$\begin{aligned}
\pi \models \mathrm{X}\varphi \quad &\text{iff} \quad \mathcal{G}, v_1 \models \varphi, \\
\pi \models \varphi \mathrm{U} \psi \quad &\text{iff} \quad \mathcal{G}, v_k \models \psi \text{ for some } k \in \omega \text{ such that} \\
&\qquad \mathcal{G}, v_l \models \varphi \text{ for all } \quad l < k \qquad\qquad \text{and} \\
\pi \models \varphi \mathrm{R} \psi \quad &\text{iff} \quad \mathcal{G}, v_j \models \psi \text{ for all } \quad j \in \omega \text{ or instead,} \\
&\qquad \mathcal{G}, v_k \models \varphi \text{ for some } k \in \omega \text{ with} \\
&\qquad \mathcal{G}, v_l \models \psi \text{ for all } \quad l \leq k.
\end{aligned}$$

Since the semantics of path formulas ϕ are defined, we can now define

$$\begin{aligned}
\mathcal{G}, v \models \mathrm{E}(\phi) \quad &\text{iff} \quad \pi \models \phi \text{ for some infinite path } \pi \text{ starting at } v \text{ and} \\
\mathcal{G}, v \models \mathrm{A}(\phi) \quad &\text{iff} \quad \pi \models \phi \text{ for all infinite paths } \pi \text{ starting at } v.
\end{aligned}$$

Notice that the restriction of the quantifiers to infinite paths does not cause counter-intuitive behaviour, because we assumed that $\mathcal{G}$ is non-terminating, therefore any path can be extended into an infinite path. Also, $\mathrm{E}(\mathrm{X}\varphi)$ corresponds to $\Diamond\varphi$ and $\mathrm{A}(\mathrm{X}\varphi)$ corresponds to $\Box\varphi$ in modal logic. Figure (3.9) illustrates some models for specific CTL formulas when evaluating at node v.

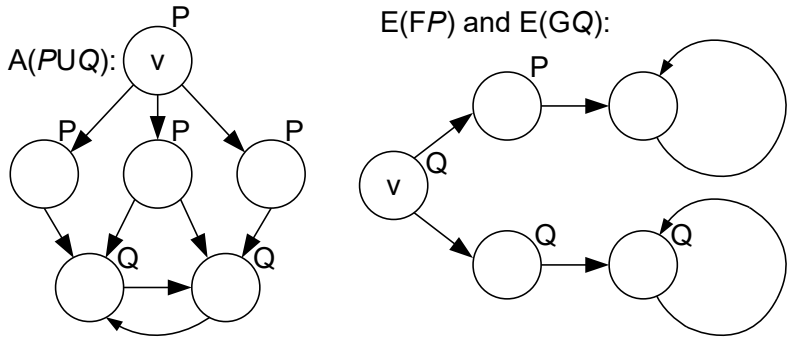

Figure (3.9): CTL formulas with models.

CTL admits a negation normal form, since for $\varphi, \psi \in \mathrm{CTL}(\tau)$, the equivalences

$$\begin{aligned}
\neg\mathrm{E}(\mathrm{X}\varphi) &\equiv \mathrm{A}(\mathrm{X}(\neg\varphi)), & \neg\mathrm{A}(\mathrm{X}\varphi) &\equiv \mathrm{E}(\mathrm{X}(\neg\varphi)), \\
\neg\mathrm{E}(\varphi\mathrm{U}\psi) &\equiv \mathrm{A}((\neg\varphi)\mathrm{R}(\neg\psi)), & \neg\mathrm{A}(\varphi\mathrm{U}\psi) &\equiv \mathrm{E}((\neg\varphi)\mathrm{R}(\neg\psi)), \\
\neg\mathrm{E}(\varphi\mathrm{R}\psi) &\equiv \mathrm{A}((\neg\varphi)\mathrm{U}(\neg\psi)), & \neg\mathrm{A}(\varphi\mathrm{R}\psi) &\equiv \mathrm{E}((\neg\varphi)\mathrm{U}(\neg\psi))
\end{aligned}$$

hold. Therefore, we would like to define a semiring interpretation for CTL. Given a set of nodes V and a signature τ, we will now have to consider the positive edge literals as well, because we might want to track them when performing provenance analysis. We set

$$\mathrm{Lit}_V(\tau) = \{P_i v \mid P_i \in \tau, v \in V\} \cup \{\neg P_i v \mid P_i \in \tau, v \in V\} \cup \{Evw \mid v, w \in V\}.$$

Now, a K-interpretation $\pi : \mathrm{Lit}_V(\tau) \to K$ assigns semiring values to all the literals and we would like to interpret CTL formulas under π. Clearly, this approach is very similar to our approach for LTL. We interpret formulas in their negation normal form and the formulas 0 and 1 as well as atomic propositions and the $\vee$ and $\wedge$ connectives can be interpreted the same way as in LTL.

The two operators $\mathrm{E}(\mathrm{X}\varphi)$ and $\mathrm{A}(\mathrm{X}\varphi)$ are interpreted the same way that Grädel and Tannen interpreted their equivalents $\Diamond\varphi$ and $\Box\varphi$ in modal logic [GT18]. To illustrate their approach, consider the incomplete $\mathbb{N}[X]$-interpretation π in figure (3.10) with $\{p, q, r, s\} \subseteq X$. Relevant edges (v, w) are labelled with $\pi(Evw)$ and missing edges are implicitly tagged with 0.

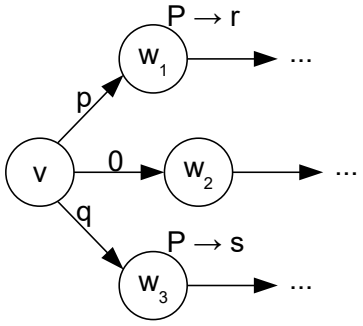

Figure (3.10): Incomplete $\mathbb{N}[X]$-interpretation π.

For $v \in V$, the set of successors is defined as

$$vE = \{w \in V \mid \pi(Evw) \neq 0\}.$$

We assume K-interpretations to be non-terminating, so $vE \neq \emptyset$ for all $v \in V$.

Generally, we interpret $\mathrm{E}(\mathrm{X}\varphi)$ and $\mathrm{A}(\mathrm{X}\varphi)$ at v for $\varphi \in \mathrm{CTL}(\tau)$ as

$$\pi[\![\mathrm{E}(\mathrm{X}\varphi)]\!]_v = \sum_{w \in vE} \pi(Evw) \cdot \pi[\![\varphi]\!]_w \quad \text{and}$$
$$\pi[\![\mathrm{A}(\mathrm{X}\varphi)]\!]_v = \prod_{w \in vE} \pi(Evw) \cdot \pi[\![\varphi]\!]_w.$$

So, we understand $\mathrm{E}(\mathrm{X}\varphi)$ as a disjunction of φ over all successors, but we also evaluate the edges that we used. $\mathrm{A}(\mathrm{X}\varphi)$ is interpreted as a conjunction over all successors. In the above example, we would obtain

$$\pi[\![\mathrm{E}(\mathrm{X}P)]\!]_v = pr + qs \quad \text{and}$$
$$\pi[\![\mathrm{A}(\mathrm{X}P)]\!]_v = pqrs.$$

Especially for the universal quantifier, it is crucial to disregard the edges tagged with 0, since they would make the entire product 0, which is counter-intuitive. As we can see, showing $\mathrm{A}(\mathrm{X}P)$ at v in the above example makes use of the fact that P holds at w_1 and w_3, hence the factors r and s, and also, the edges (v, w_1) and (v, w_3) are used, therefore, we have the factors p and q. The node w_2 is irrelevant, as (v, w_2) is tagged with zero, so $pqrs$ represents a full proof of $\mathrm{A}(\mathrm{X}P)$ at v.

Finally, we move on to interpreting the until and release formulas. Let φ and ψ be in $\mathrm{CTL}(\tau)$ and consider the formula $\mathrm{E}(\varphi \mathrm{U} \psi)$. Instead of evaluating the formula directly, we consider the equivalence

$$\mathrm{E}(\varphi \mathrm{U} \psi) \equiv \psi \vee (\varphi \wedge \mathrm{E}(\mathrm{X}(\mathrm{E}(\varphi \mathrm{U} \psi)))).$$

If we evaluate $\mathrm{E}(\varphi \mathrm{U} \psi)$ at a node v, clearly, if ψ is true at v, we know that $\mathrm{E}(\varphi \mathrm{U} \psi)$ holds at v. The other possibility for $\mathrm{E}(\varphi \mathrm{U} \psi)$ to be true at v would be that φ holds

at v and there is some successor of v where $\mathrm{E}(\varphi \mathrm{U} \psi)$ is true. This way, we can "push" the formula $\mathrm{E}(\varphi \mathrm{U} \psi)$ on to the successors of the current node.

Since we already know how to evaluate the next operator, we can translate this into an equation system. For all $v \in V$, we have the equations

$$\pi[\![\mathrm{E}(\varphi \mathrm{U} \psi)]\!]_v = \pi[\![\psi]\!]_v + \pi[\![\varphi]\!]_v \cdot \sum_{w \in vE} \pi(Evw) \cdot \pi[\![\mathrm{E}(\varphi \mathrm{U} \psi)]\!]_w.$$

If K is ω-continuous, we can solve this system. Consider the function $f : K^V \to K^V$ which is defined component-wise for $v \in V$ as

$$f_v(X) = \pi[\![\psi]\!]_v + \pi[\![\varphi]\!]_v \cdot \sum_{w \in vE} \pi(Evw) \cdot X_w \quad \text{where } X \in K^V.$$

By theorem (2.23), since f is component-wise ω-continuous in each argument, it has a least fixed point $\mathrm{lfp}(f) \in K^V$. Setting

$$\pi[\![\mathrm{E}(\varphi \mathrm{U} \psi)]\!]_v = \mathrm{lfp}(f)_v$$

for each $v \in V$ fulfills the above equations, since

$$\mathrm{lfp}(f)_v = f_v(\mathrm{lfp}(f)) = \pi[\![\psi]\!]_v + \pi[\![\varphi]\!]_v \cdot \sum_{w \in vE} \pi(Evw) \cdot \mathrm{lfp}(f)_w \quad \text{for } v \in V.$$

Obviously, the equation would be fulfilled by any other fixed point of f as well, but we use the least fixed point, since the until operator specifies a *reachability condition*, that is, when $\mathrm{E}(\varphi \mathrm{U} \psi)$ holds at v, a node where ψ holds must be reachable. Another justification for using the least fixed point can be found when we translate $\mathrm{E}(\varphi \mathrm{U} \psi)$ to the modal μ-calculus L_μ. Although we will not define L_μ here, we note that

$$\mathrm{E}(\varphi \mathrm{U} \psi) \text{ translates to } \mu X.(\psi \vee (\varphi \wedge \Diamond X)),$$

which calls for a least fixed point and exactly justifies the definition of f.

The remaining formulas $\mathrm{A}(\varphi \mathrm{U} \psi)$, $\mathrm{E}(\varphi \mathrm{R} \psi)$ and $\mathrm{A}(\varphi \mathrm{R} \psi)$ can be dealt with in a similar way. For $\mathrm{A}(\varphi \mathrm{U} \psi)$, we have the equivalence

$$\mathrm{A}(\varphi \mathrm{U} \psi) \equiv \psi \vee (\varphi \wedge \mathrm{A}(\mathrm{X}(\mathrm{A}(\varphi \mathrm{U} \psi)))),$$

so it suffices to just replace the summation in our operator f by a multiplication. We obtain $g : K^V \to K^V$ with

$$g_v(X) = \pi[\![\psi]\!]_v + \pi[\![\varphi]\!]_v \cdot \prod_{w \in vE} \pi(Evw) \cdot X_w \quad \text{for } v \in V.$$

The translation of

$$\mathrm{A}(\varphi \mathrm{U} \psi) \text{ to } \mu X.(\psi \vee (\varphi \wedge \Box X)) \text{ in } L_\mu$$

justifies the definition $\pi[\![\mathrm{A}(\varphi \mathrm{U} \psi)]\!]_v = \mathrm{lfp}(g)_v$.

For the release operators, we have to slightly change the functions f and g. Consider the formula $\mathrm{E}(\varphi \mathrm{R} \psi)$ at some node v. By the semantics of $\mathrm{E}(\varphi \mathrm{R} \psi)$, ψ must be true in any case at v as a necessary condition to fulfill $\mathrm{E}(\varphi \mathrm{R} \psi)$. Now, if φ is true at v

as well, this is sufficient to show $\mathrm{E}(\varphi \mathrm{R} \psi)$, or, the other possibility is that $\mathrm{E}(\varphi \mathrm{R} \psi)$ is true at some successor of v, so we have

$$\mathrm{E}(\varphi \mathrm{R} \psi) \equiv \psi \wedge (\varphi \vee \mathrm{E}(\mathrm{X}(\mathrm{E}(\varphi \mathrm{R} \psi)))).$$

This gives rise to the function $f' : K^V \to K^V$ with

$$f'_v(X) = \pi[\![\psi]\!]_v \cdot \left(\pi[\![\varphi]\!]_v + \sum_{w \in vE} \pi(Evw) \cdot X_w \right) \quad \text{for } v \in V.$$

Similarly, for $\mathrm{A}(\varphi \mathrm{U} \psi)$, we use $g' : K^V \to K^V$ defined as

$$g'_v(X) = \pi[\![\psi]\!]_v \cdot \left(\pi[\![\varphi]\!]_v + \prod_{w \in vE} \pi(Evw) \cdot X_w \right) \quad \text{for } v \in V.$$

However, release operators do not specify reachability conditions. For example, to fulfill $\mathrm{E}(\varphi \mathrm{R} \psi)$ at a node v, no node where φ holds has to be reachable, in fact, an infinite path of nodes where ψ holds is enough. So, the release operator specifies the *safety condition* that ψ stays true indefinitely or until it is released by φ. Therefore, we use greatest fixed points. To ensure their existence, we have to assume that K is an absorptive lattice semiring and use corollary (2.25) that implies the existence of greatest fixed points for f' and g', because they are component-wise monotonic in each argument. We set

$$\begin{aligned} \pi[\![\mathrm{E}(\varphi \mathrm{R} \psi)]\!]_v &= \mathrm{gfp}(f')_v \quad \text{and} \\ \pi[\![\mathrm{A}(\varphi \mathrm{R} \psi)]\!]_v &= \mathrm{gfp}(g')_v \quad \text{for } v \in V. \end{aligned}$$

The translations to L_μ of

$$\begin{aligned} \mathrm{E}(\varphi \mathrm{R} \psi) \quad &\text{to} \quad \nu X.(\psi \wedge (\varphi \vee \Diamond X)) \quad \text{and} \\ \mathrm{A}(\varphi \mathrm{R} \psi) \quad &\text{to} \quad \nu X.(\psi \wedge (\varphi \vee \Box X)) \end{aligned}$$

back our definition. We will now summarize the considerations above and inductively define the semiring interpretation for CTL.

(3.11) Definition (Semiring Interpretation for CTL)**.** Let K be an absorptive lattice semiring, V a finite set of nodes, τ a signature and $\pi : \mathrm{Lit}_V(\tau) \to K$ a K-interpretation. We define the semiring interpretation of any formula $\theta \in \mathrm{CTL}(\tau)$ at any node $v \in V$ inductively on the negation normal form of θ by setting

$$\begin{aligned} \pi[\![0]\!]_v &= 0, & \pi[\![1]\!]_v &= 1, \\ \pi[\![P_i]\!]_v &= \pi(P_i v), & \pi[\![\neg P_i]\!]_v &= \pi(\neg P_i v), \\ \pi[\![\varphi \vee \psi]\!]_v &= \pi[\![\varphi]\!]_v + \pi[\![\psi]\!]_v \quad \text{and} & \pi[\![\varphi \wedge \psi]\!]_v &= \pi[\![\varphi]\!]_v \cdot \pi[\![\psi]\!]_v, \end{aligned}$$

where $\varphi, \psi \in \mathrm{CTL}(\tau)$ as in LTL and further, we set

$$\begin{aligned} \pi[\![\mathrm{E}(\mathrm{X}\varphi)]\!]_v &= \sum_{w \in vE} \pi(Evw) \cdot \pi[\![\varphi]\!]_w, & \pi[\![\mathrm{A}(\mathrm{X}\varphi)]\!]_v &= \prod_{w \in vE} \pi(Evw) \cdot \pi[\![\varphi]\!]_w, \\ \pi[\![\mathrm{E}(\varphi \mathrm{U} \psi)]\!]_v &= \mathrm{lfp}(f^{\mathrm{E}(\varphi \mathrm{U} \psi)})_v, & \pi[\![\mathrm{A}(\varphi \mathrm{U} \psi)]\!]_v &= \mathrm{lfp}(f^{\mathrm{A}(\varphi \mathrm{U} \psi)})_v, \\ \pi[\![\mathrm{E}(\varphi \mathrm{R} \psi)]\!]_v &= \mathrm{gfp}(f^{\mathrm{E}(\varphi \mathrm{R} \psi)})_v \quad \text{and} & \pi[\![\mathrm{A}(\varphi \mathrm{R} \psi)]\!]_v &= \mathrm{gfp}(f^{\mathrm{A}(\varphi \mathrm{R} \psi)})_v \end{aligned}$$

with the corresponding functions $K^V \to K^V$ being defined as

$$f_v^{\mathrm{E}(\varphi \mathrm{U} \psi)}(X) = \pi[\![\psi]\!]_v + \pi[\![\varphi]\!]_v \cdot \sum_{w \in vE} \pi(Evw) \cdot X_w,$$

$$f_v^{\mathrm{A}(\varphi \mathrm{U} \psi)}(X) = \pi[\![\psi]\!]_v + \pi[\![\varphi]\!]_v \cdot \prod_{w \in vE} \pi(Evw) \cdot X_w,$$

$$f_v^{\mathrm{E}(\varphi \mathrm{R} \psi)}(X) = \pi[\![\psi]\!]_v \cdot \left(\pi[\![\varphi]\!]_v + \sum_{w \in vE} \pi(Evw) \cdot X_w \right) \quad \text{and}$$

$$f_v^{\mathrm{A}(\varphi \mathrm{R} \psi)}(X) = \pi[\![\psi]\!]_v \cdot \left(\pi[\![\varphi]\!]_v + \prod_{w \in vE} \pi(Evw) \cdot X_w \right).$$

Since absorptive lattice semirings have to fulfill very strong conditions and we only need them for the release operators, it is worth noting that if there were no release operators, ω-continuous semirings would be enough to interpret CTL.

(3.12) Definition (Positive CTL)**.** We define $\mathrm{posCTL}(\tau) \subseteq \mathrm{CTL}(\tau)$ for a signature τ as the subset of formulas $\varphi \in \mathrm{CTL}(\tau)$ where release operators do not occur in the negation normal form of φ.

Formulas in $\mathrm{posCTL}(\tau)$ can be interpreted by means of definition (3.11) in any ω-continuous semiring K. We will provide two examples for the semiring interpretation of CTL formulas.

(3.13) Example. Let $\tau = \{E, P\}$ and consider the $\mathbb{N}^\infty[\![X]\!]$-interpretation π from figure (3.14) below where $X = \{p, q, r\}$.

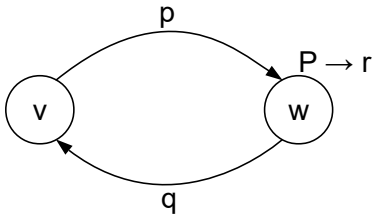

Figure (3.14): $\mathbb{N}^\infty[\![X]\!]$-interpretation π.

We will interpret the formula $\mathrm{E}(\mathrm{F}P) \in \mathrm{CTL}(\tau)$ in $\mathbb{N}^\infty[\![X]\!]$ under π at node v. Notice that $\mathrm{E}(\mathrm{F}P) = \mathrm{E}(1\mathrm{U}P) \in \mathrm{posCTL}(\tau)$, so that the ω-continuous semiring $\mathbb{N}^\infty[\![X]\!]$ is suited for that purpose. We have

$$\pi[\![\mathrm{E}(\mathrm{F}P)]\!]_v = \pi[\![\mathrm{E}(1\mathrm{U}P)]\!]_v = \mathrm{lfp}(f^{\mathrm{E}(1\mathrm{U}P)})_v.$$

Since we have two nodes v and w and $vE = \{w\}$ and $wE = \{v\}$, we can write the two components of $f^{\mathrm{E}(1\mathrm{U}P)} : K^{\{v,w\}} \to K^{\{v,w\}}$ as

$$f_v^{\mathrm{E}(1\mathrm{U}P)}(Y) = \pi[\![P]\!]_v + \pi[\![1]\!]_v \cdot \pi(Evw) \cdot Y_w \quad \text{and}$$
$$f_w^{\mathrm{E}(1\mathrm{U}P)}(Y) = \pi[\![P]\!]_w + \pi[\![1]\!]_w \cdot \pi(Ewv) \cdot Y_v.$$

By theorem (2.23), $\mathrm{lfp}(f^{\mathrm{E}(1\mathrm{U}P)}) = \sup_{i \in \omega}(f^{\mathrm{E}(1\mathrm{U}P)})^i(0)$, so we can iterate $f^{\mathrm{E}(1\mathrm{U}P)}$ to obtain the fixed point. We start with $Y_0 = (0, 0)$ and set $Y_{i+1} = f^{\mathrm{E}(1\mathrm{U}P)}(Y_i)$ for $i \in \omega$. This yields the iteration rules

$$(Y_{i+1})_v = f_v^{\mathrm{E}(1\mathrm{U}P)}(Y_i) = \pi[\![P]\!]_v + \pi[\![1]\!]_v \cdot \pi(Evw) \cdot (Y_i)_w = p \cdot (Y_i)_w \qquad \text{and}$$
$$(Y_{i+1})_w = f_w^{\mathrm{E}(1\mathrm{U}P)}(Y_i) = \pi[\![P]\!]_w + \pi[\![1]\!]_w \cdot \pi(Ewv) \cdot (Y_i)_v = r + q \cdot (Y_i)_v.$$

We can create the following table of $(Y_i)_v$ and $(Y_i)_w$ for $0 \leq i \leq 5$:

i	0	1	2	3	4	5	...
$(Y_i)_v$	0	0	pr	pr	$pr + p^2qr$	$pr + p^2qr$	...
$(Y_i)_w$	0	r	r	$r + pqr$	$r + pqr$	$r + pqr + p^2q^2r$	...

Of course, we cannot enumerate all Y_i for $i \in \omega$, but from the table, we can conclude that $(Y_i)_v$ contains all monomials of the form $p(pq)^j r$ with $j \leq k$ for some $k \in \omega$ and k increases with i, that is, more monomials are added as the iteration progresses. A similar observation can be made for $(Y_i)_w$, which contains all the monomials $(pq)^{j'} r$ with $j' \leq k'$ for some increasing k'. We claim

$$\sup_{i\in\omega}(Y_i)_v = \sum_{j\in\omega} p(pq)^j r = p(pq)^* r \quad \text{and}$$
$$\sup_{i\in\omega}(Y_i)_w = \sum_{j\in\omega} (pq)^j r \quad = (pq)^* r$$

where a^* is short for $\sum_{j\in\omega} a^j$ for $a \in K$ when K is ω-continuous. By lemma (2.21), these are the components of the supremum of $\{Y_i \mid i \in \omega\}$, so

$$\mathrm{lfp}(f^{\mathrm{E}(1\mathrm{U}P)}) = \sup_{i\in\omega}(f^{\mathrm{E}(1\mathrm{U}P)})^i(0) = \sup_{i\in\omega} Y_i = (p(pq)^* r, (pq)^* r).$$

Therefore, we have evaluated $\mathrm{E}(\mathrm{F}P)$ at v, and as a coproduct, we have evaluated it at w as well and obtain

$$\pi[\![\mathrm{E}(\mathrm{F}P)]\!]_v = \pi[\![\mathrm{E}(1\mathrm{U}P)]\!]_v = \mathrm{lfp}(f^{\mathrm{E}(1\mathrm{U}P)})_v = p(pq)^* r \quad \text{and}$$
$$\pi[\![\mathrm{E}(\mathrm{F}P)]\!]_w = \pi[\![\mathrm{E}(1\mathrm{U}P)]\!]_w = \mathrm{lfp}(f^{\mathrm{E}(1\mathrm{U}P)})_w = (pq)^* r.$$

As a sanity check for this result, we take a look at the literals that are used when proving $\mathrm{E}(\mathrm{F}P)$ at v. We can take the p-labelled edge (v, w) to w and then show P at w, which is labelled with r. However, before showing P at w, we can iterate through the cycle (w, v, w) arbitrarily often, and since the cycle's edges are labelled with p and q, the result $p(pq)^* r$ represents exactly all the ways to prove $\mathrm{E}(\mathrm{F}P)$ at v. For w, it is very similar, except that we do not have to take the p-transition in the beginning, since we are starting at w, so the result $(pq)^* r$ is justified. We will now provide another example with a greatest fixed point.

(3.15) Example. In this example, $\tau = \{E, Q\}$ and we use the $\mathbb{S}^\infty[X]$-interpretation π in figure (3.16) below with $X = \{p, q, r, s\}$.

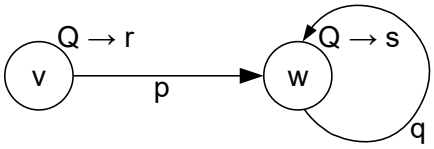

Figure (3.16): $\mathbb{S}^\infty[X]$-interpretation π.

The formula that we will interpret in $\mathbb{S}^\infty[X]$ under π at node v is $\mathrm{E}(\mathrm{G}Q) \in \mathrm{CTL}(\tau)$, which is equal to $\mathrm{E}(0\mathrm{R}Q)$, so

$$\pi[\![\mathrm{E}(\mathrm{G}Q)]\!]_v = \pi[\![\mathrm{E}(0\mathrm{R}Q)]\!]_v = \mathrm{gfp}(f^{\mathrm{E}(0\mathrm{R}Q)})_v.$$

We have $vE = wE = \{w\}$, so that the components of $f^{\mathrm{E(0RQ)}} : K^{\{v,w\}} \to K^{\{v,w\}}$ are

$$f_v^{\mathrm{E(0RQ)}}(Y) = \pi[\![Q]\!]_v \cdot (\pi[\![0]\!]_v + \pi(Evw) \cdot Y_w) \quad \text{and}$$
$$f_w^{\mathrm{E(0RQ)}}(Y) = \pi[\![Q]\!]_w \cdot (\pi[\![0]\!]_w + \pi(Eww) \cdot Y_w).$$

By corollary (2.25), $\mathrm{gfp}(f^{\mathrm{E(0RQ)}})$ exists and can be obtained by means of theorem (2.17), that is, we start iterating at $Y_0 = (1, 1)$, set $Y_{i+1} = f^{\mathrm{E(0RQ)}}(Y_i)$ for $i \in \omega$ and $Y_\omega = \inf_{i\in\omega} Y_i$. We will see in this example that Y_ω is already a fixed point, so there is no need to iterate further than Y_ω. We simplify

$$(Y_{i+1})_v = f_v^{\mathrm{E(0RQ)}}(Y_i) = \pi[\![Q]\!]_v \cdot (\pi[\![0]\!]_v + \pi(Evw) \cdot Y_w) \quad = pr \cdot (Y_i)_w \qquad \text{and}$$
$$(Y_{i+1})_w = f_w^{\mathrm{E(0RQ)}}(Y_i) = \pi[\![Q]\!]_w \cdot (\pi[\![0]\!]_w + \pi(Eww) \cdot Y_w) = qs \cdot (Y_i)_w \quad \text{for } i \in \omega.$$

This yields the following table for $0 \leq i \leq 5$:

i	0	1	2	3	4	5	...
$(Y_i)_v$	1	pr	$pqrs$	pq^2rs^2	pq^3rs^3	pq^4rs^4	...
$(Y_i)_w$	1	qs	q^2s^2	q^3s^3	q^4s^4	q^5s^5	...

We use the abbreviation $a^\infty = \prod_{i\in\omega} a$ for $a \in K$ where K is an absorptive lattice semiring. Notice that in $\mathbb{S}^\infty[X]$, greater exponents make monomials smaller, so we have $(qs)^\infty = q^\infty s^\infty \leq q^j s^j = (qs)^j$ for all $j \in \omega$. Now, using the above table and lemma (2.21), it is easy to see that

$$(Y_\omega)_v = \inf_{i\in\omega}(Y_i)_v = pr(qs)^\infty \quad \text{and}$$
$$(Y_\omega)_w = \inf_{i\in\omega}(Y_i)_w = (qs)^\infty.$$

This is already a fixed point, since $f^{\mathrm{E(0RQ)}}(Y_\omega) = Y_\omega$, because

$$f_v^{\mathrm{E(0RQ)}}(Y_\omega) = pr \cdot (Y_\omega)_w = pr(qs)^\infty \qquad = (Y_\omega)_v \qquad \text{and}$$
$$f_w^{\mathrm{E(0RQ)}}(Y_\omega) = qs \cdot (Y_\omega)_w = qs \cdot (qs)^\infty = (qs)^\infty = (Y_\omega)_w.$$

So, it is the greatest fixed point and we have

$$\pi[\![\mathrm{E(GQ)}]\!]_v = \pi[\![\mathrm{E(0RQ)}]\!]_v = \mathrm{gfp}(f^{\mathrm{E(0RQ)}})_v = (Y_\omega)_v = pr(qs)^\infty \quad \text{and}$$
$$\pi[\![\mathrm{E(GQ)}]\!]_w = \pi[\![\mathrm{E(0RQ)}]\!]_w = \mathrm{gfp}(f^{\mathrm{E(0RQ)}})_w = (Y_\omega)_w = (qs)^\infty.$$

Indeed, the only way to prove $\mathrm{E(GQ)}$ at v is by using the infinite path $(v, w, w, ...)$, so we use the fact that Q holds at v and the edge (v, w) once, hence the factor pr, and we use the edge (w, w) and the fact that Q holds at w infinitely often, hence $(qs)^\infty$. At the node w, we use the path $(w, w, ...)$, so we just have $(qs)^\infty$.

Unfortunately, as seen in the two examples, we can only interpret CTL formulas manually for now and it requires infinite iterations, so we have to intuitively see a pattern in the iterations and use induction to find the fixed points. However, we can see that cycles are causing the iterations to be infinite, which raises the question whether the results can be obtained systematically. Indeed, in the next chapter, we will present algorithms that are capable of interpreting CTL formulas without infinite iterations.

3.3 Propositional Dynamic Logic (PDL)

Propositional Dynamic Logic (PDL) is interpreted over Kripke structures $\mathcal{K} = (V, (E_a)_{a\in A}, (P_i)_{i\in I})$ where I is an index set and A is a set of actions. Intuitively, we can interpret the actions as atomic programs and if $(v, w) \in E_a$ for some $a \in A$, we can say that performing action a at state v takes the system to state w. Again, V is assumed to be finite and $\tau = \{E_a \mid a \in A\} \cup \{P_i \mid i \in I\}$ is called the signature of $\mathcal{K}$. The syntax and usual semantics of PDL are adapted from Berwanger's work from 2005 [Ber05].

(3.17) Definition (Syntax of PDL)**.** Given a signature τ, PDL(τ) is generated by the symbol φ of the grammar

$$\begin{aligned} \varphi &= 0 \mid 1 \mid P_i \mid \neg\varphi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \langle\rho\rangle\varphi \mid [\rho]\varphi && \text{with } P_i \in \tau, \\ \rho &= a \mid \rho \cup \rho \mid \rho;\rho \mid \rho^* \mid \varphi? && \text{with } E_a \in \tau. \end{aligned}$$

The symbol ρ generates the *programs*. The elementary programs are the actions a. The operator $\rho;\rho$ chains two programs, $\rho \cup \rho$ indeterministically chooses one of the given programs and ρ^* indeterministically iterates ρ zero or more times. Additionally, $\varphi?$ checks if φ is true at the current state or not. Knowing the meaning of the programs, we say that $\langle\rho\rangle\varphi$ is true iff there is a run of ρ that ends in a state where φ holds and $[\rho]\varphi$ holds iff all runs of ρ end in a state where φ is true.

Although we will usually interpret formulas φ at a single node $v \in V$, the easiest way to define the semantics of PDL is by setting $[\![\varphi]\!]^{\mathcal{K}} \subseteq V = \{v \in V \mid \mathcal{K}, v \models \varphi\}$ and interpreting programs ρ as binary relations $[\![\rho]\!]^{\mathcal{K}} \subseteq V \times V$.

(3.18) Definition (Semantics of PDL)**.** Given a Kripke structure $\mathcal{K} = (V, \tau)$, we define the semantics for formulas and programs inductively at the same time. Assume $\varphi, \psi \in \text{PDL}(\tau)$, we set

$$\begin{aligned} [\![0]\!]^{\mathcal{K}} &= \emptyset, & [\![1]\!]^{\mathcal{K}} &= V, \\ [\![P_i]\!]^{\mathcal{K}} &= P_i^{\mathcal{K}}, & [\![\neg\varphi]\!]^{\mathcal{K}} &= V \setminus [\![\varphi]\!]^{\mathcal{K}}, \\ [\![\varphi \vee \psi]\!]^{\mathcal{K}} &= [\![\varphi]\!]^{\mathcal{K}} \cup [\![\psi]\!]^{\mathcal{K}}, & [\![\varphi \wedge \psi]\!]^{\mathcal{K}} &= [\![\varphi]\!]^{\mathcal{K}} \cap [\![\psi]\!]^{\mathcal{K}}, \end{aligned}$$

$$[\![\langle\rho\rangle\varphi]\!]^{\mathcal{K}} = \{v \in V \mid \text{there is a } w \in V \text{ with } (v, w) \in [\![\rho]\!]^{\mathcal{K}} \text{ and } \mathcal{K}, w \models \varphi\} \quad \text{and}$$
$$[\![[\rho]\varphi]\!]^{\mathcal{K}} = \{v \in V \mid \text{for all } w \in V \text{ with } (v, w) \in [\![\rho]\!]^{\mathcal{K}}, \text{ we have } \mathcal{K}, w \models \varphi\}.$$

For the programs, assume ρ, ρ_1 and ρ_2 are subprograms and $\varphi \in \text{PDL}(\tau)$, then

$$\begin{aligned} [\![a]\!]^{\mathcal{K}} &= E_a^{\mathcal{K}}, \\ [\![\rho_1 \cup \rho_2]\!]^{\mathcal{K}} &= [\![\rho_1]\!]^{\mathcal{K}} \cup [\![\rho_2]\!]^{\mathcal{K}}, \\ [\![\rho_1;\rho_2]\!]^{\mathcal{K}} &= \{(v, w) \mid \text{there is a } u \in V \text{ with } (v, u) \in [\![\rho_1]\!]^{\mathcal{K}} \text{ and } (u, w) \in [\![\rho_2]\!]^{\mathcal{K}}\}, \\ [\![\rho^*]\!]^{\mathcal{K}} &= \{(v, w) \mid v = w \text{ or } w \text{ is reachable from } v \text{ via edges in } [\![\rho]\!]^{\mathcal{K}}\} \quad \text{and} \\ [\![\varphi?]\!]^{\mathcal{K}} &= \{(v, v) \mid \mathcal{K}, v \models \varphi\}. \end{aligned}$$

Figure (3.19) illustrates some models of PDL formulas when evaluating at node v. We label the edges with the action they belong to.

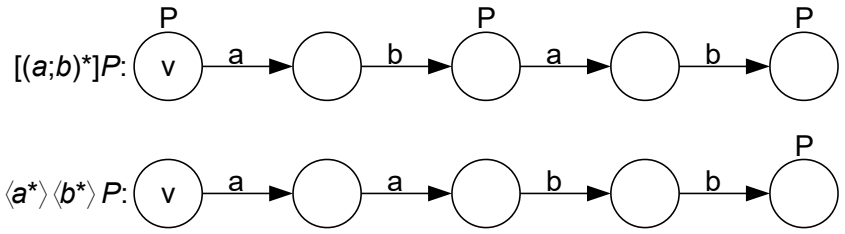

Figure (3.19): PDL formulas with models.

For programs ρ and formulas φ, we call formulas of the form $\langle\rho\rangle\varphi$ existential and formulas of the form $[\rho]\varphi$ universal. Clearly, PDL has a negation normal form when we consider

$$\neg\langle\rho\rangle\varphi \equiv [\rho](\neg\varphi) \quad \text{and} \quad \neg[\rho]\varphi \equiv \langle\rho\rangle(\neg\varphi).$$

(3.20) Definition (Positive PDL)**.** Let τ be an arbitrary signature, we define $\mathrm{posPDL}(\tau) \subseteq \mathrm{PDL}(\tau)$ as the subset of $\mathrm{PDL}(\tau)$ formulas whose negation normal form does not contain a universal subformula $[\rho]\varphi$ where ρ is a program and $\varphi \in \mathrm{PDL}(\tau)$.

In this thesis, we will only define a semiring interpretation for the positive fragment of PDL. For that, the semirings K that we use will have to be ω-continuous. For a set of nodes V and signature τ, we have the literals

$$\begin{aligned}\mathrm{Lit}_V(\tau) =& \{P_i v \mid P_i \in \tau, v \in V\} \cup \{\neg P_i v \mid P_i \in \tau, v \in V\} \\ &\cup \{E_a vw \mid E_a \in \tau, v, w \in V\}.\end{aligned}$$

A K-interpretation $\pi : \mathrm{Lit}_V(\tau) \to K$ is extended to all $\mathrm{posPDL}(\tau)$ formulas. We can handle all formulas except for $\langle\rho\rangle\varphi$ the same way as in CTL, where ρ is a program and $\varphi \in \mathrm{posPDL}(\tau)$. Formulas of the form $\langle\rho\rangle\varphi$ require us to assign a meaning to programs ρ first. Under the standard semantics, programs were interpreted as binary relations. Therefore, under π, we interpret programs as binary K-relations, that is for a program ρ and $(v, w) \in V \times V$, we will define

$$\pi[\![\rho]\!]_{(v,w)} \in K,$$

thereby extending π to a function $\pi : \mathrm{Prog}(\tau) \times (V \times V) \to K$ where $\mathrm{Prog}(\tau)$ is the set of programs over τ. Now, it is left to define $\pi[\![\rho]\!]_{(v,w)}$ inductively for all programs ρ and $(v, w) \in V \times V$.

The meaning of atomic programs a is given directly by the edge relation E_a, we set

$$\pi[\![a]\!]_{(v,w)} = \pi(E_a vw).$$

Let ρ_1, ρ_2 and ρ be programs and $\psi \in \mathrm{PDL}(\tau)$. For the compound programs, we will translate the standard semantics into semiring semantics. $\rho_1 \cup \rho_2$ can be seen as a disjunction, so we set

$$\pi[\![\rho_1 \cup \rho_2]\!]_{(v,w)} = \pi[\![\rho_1]\!]_{(v,w)} + \pi[\![\rho_2]\!]_{(v,w)}.$$

The program $\rho_1; \rho_2$ can also be seen as a disjunction. To evaluate $\rho_1; \rho_2$ at (v, w), we have to look for a middle node u and evaluate ρ_1 at (v, u) and ρ_2 at (u, w). This yields the definition

$$\pi[\![\rho_1; \rho_2]\!]_{(v,w)} = \sum_{u \in V} \pi[\![\rho_1]\!]_{(v,u)} \cdot \pi[\![\rho_2]\!]_{(u,w)}.$$

For $\psi?$, we notice that $\psi?$ is a program that does not change the state. Therefore, evaluating $\psi?$ at (v, w) with $v \neq w$ should automatically yield 0. When evaluating $\psi?$ at (v, v), we have to check whether ψ holds at v, so we have the definition

$$\pi[\![\psi?]\!]_{(v,w)} = \begin{cases} \pi[\![\psi]\!]_v & \text{if } v = w \\ 0 & \text{otherwise.} \end{cases}$$

The most complicated program is ρ^*. When we evaluate ρ^* at (v, w), we want to know if w is reachable from v by ρ-transitions. Obviously, this is automatically true if $v = w$. The other possibility is that v has a ρ-successor u from which w is reachable, so we could recursively evaluate ρ^* at (u, w). For $(v, w) \in V \times V$, we can translate this to the equation system

$$\pi[\![\rho^*]\!]_{(v,w)} = \pi[\![1?]\!]_{(v,w)} + \sum_{u \in V} \pi[\![\rho]\!]_{(v,u)} \cdot \pi[\![\rho^*]\!]_{(u,w)}.$$

The term $\pi[\![1?]\!]_{(v,w)}$ is 1 iff $v = w$, because in that case, w is reachable from v without using any transitions at all. For $v \neq w$, this does not apply and $\pi[\![1?]\!]_{(v,w)}$ is 0. Since K is ω-continuous, we can solve this equation system by iterating the function $f : K^{V \times V} \to K^{V \times V}$, which is defined component-wise as

$$f_{(v,w)}(X) = \pi[\![1?]\!]_{(v,w)} + \sum_{u \in V} \pi[\![\rho]\!]_{(v,u)} \cdot X_{(u,w)} \quad \text{where } X \in K^{V \times V}.$$

Since f is component-wise ω-continuous in each argument, by theorem (2.23), f has a least fixed point $\mathrm{lfp}(f) \in K^{V \times V}$ and we can verify that $\mathrm{lfp}(f)$ fulfills the equation

$$\mathrm{lfp}(f)_{(v,w)} = f_{(v,w)}(\mathrm{lfp}(f)) = \pi[\![1?]\!]_{(v,w)} + \sum_{u \in V} \pi[\![\rho]\!]_{(v,u)} \cdot \mathrm{lfp}(f)_{(u,w)}.$$

So, it is justified to set

$$\pi[\![\rho^*]\!]_{(v,w)} = \mathrm{lfp}(f)_{(v,w)}.$$

We notice that this is very similar to our approach for the existential until operator in CTL, except that we iterate over binary K-relations $K^{V \times V}$ instead of unary K-relations K^V. Taking the least fixed point is again justified by observing that $\langle \rho^* \rangle \varphi$ specifies a reachability condition.

Having defined the meanings of programs, it is easy to interpret $\langle \rho \rangle \varphi$ at a node v by noticing that this formula is true iff there is a ρ-successor of v where φ holds. This can be seen as a disjunction over all successors of V, so we define

$$\pi[\![\langle \rho \rangle \varphi]\!]_v = \sum_{w \in V} \pi[\![\rho]\!]_{(v,w)} \cdot \pi[\![\varphi]\!]_w.$$

Now, we can interpret positive PDL in ω-continuous semirings.

(3.21) Definition (Semiring Interpretation for Positive PDL)**.** Let K be an ω-continuous semiring, V a finite set of nodes, τ a signature and $\pi : \mathrm{Lit}_V(\tau) \to K$ a K-interpretation. We define the semiring interpretation of any program and any formula $\theta \in \mathrm{PDL}(\tau)$ by simultaneous induction on the negation normal form of

formulas and on the programs. Let $\varphi, \psi \in \mathrm{posPDL}(\tau)$, ρ_1, ρ_2 and ρ be programs and $v, w \in V$, we set

$$\begin{aligned} \pi[\![0]\!]_v &= 0, & \pi[\![1]\!]_v &= 1, \\ \pi[\![P_i]\!]_v &= \pi(P_i v), & \pi[\![\neg P_i]\!]_v &= \pi(\neg P_i v), \\ \pi[\![\varphi \vee \psi]\!]_v &= \pi[\![\varphi]\!]_v + \pi[\![\psi]\!]_v, & \pi[\![\varphi \wedge \psi]\!]_v &= \pi[\![\varphi]\!]_v \cdot \pi[\![\psi]\!]_v \quad \text{and} \end{aligned}$$

$$\pi[\![\langle\rho\rangle\varphi]\!]_v = \sum_{w \in V} \pi[\![\rho]\!]_{(v,w)} \cdot \pi[\![\varphi]\!]_w$$

for the formulas and for the programs, we set

$$\begin{aligned} \pi[\![a]\!]_{(v,w)} &= \pi(E_a v w), \\ \pi[\![\rho_1 \cup \rho_2]\!]_{(v,w)} &= \pi[\![\rho_1]\!]_{(v,w)} + \pi[\![\rho_2]\!]_{(v,w)}, \\ \pi[\![\rho_1; \rho_2]\!]_{(v,w)} &= \sum_{u \in V} \pi[\![\rho_1]\!]_{(v,u)} \cdot \pi[\![\rho_2]\!]_{(u,w)}, \\ \pi[\![\rho^*]\!]_{(v,w)} &= \mathrm{lfp}(f^{\rho^*})_{(v,w)} \quad \text{and} \\ \pi[\![\varphi?]\!]_{(v,w)} &= \begin{cases} \pi[\![\varphi]\!]_v & \text{if } v = w \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

The function $f^{\rho^*} : K^{V \times V} \to K^{V \times V}$ is defined as

$$f^{\rho^*}_{(v,w)}(X) = \pi[\![1?]\!]_{(v,w)} + \sum_{u \in V} \pi[\![\rho]\!]_{(v,u)} \cdot X_{(u,w)}.$$

(3.22) Example. We illustrate this with an example for $\tau = \{E_a, P\}$ and a $\mathbb{N}^\infty[\![X]\!]$-interpretation π given in figure (3.23) with $X = \{p, q, r, s\}$. Edges (v, w) are labelled with actions a and their appropriate values $\pi(E_a v w)$.

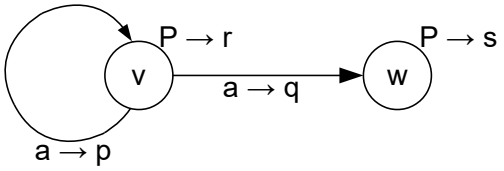

Figure (3.23): $\mathbb{N}^\infty[\![X]\!]$-interpretation π.

We would like to evaluate $\langle a^* \rangle P$ at node v in $\mathbb{N}^\infty[\![X]\!]$ under π. First of all, we will evaluate the program a^* at all pairs of nodes, therefore, we use lemma (2.23) and compute $\mathrm{lfp}(f^{a^*}) = \sup_{i \in \omega}(f^{a^*})^i(0)$. We have four components, so we start at

$Y_0 = (0,0,0,0)$ and set $Y_{i+1} = f^{a^*}(Y_i)$ for $i \in \omega$. This can be simplified to

$$\begin{aligned}
(Y_{i+1})_{(v,v)} = f^{a^*}_{(v,v)}(Y_i) &= \pi[\![1?]\!]_{(v,v)} + \sum_{u \in V} \pi[\![a]\!]_{(v,u)} \cdot (Y_i)_{(u,v)} \\
&= 1 + p \cdot (Y_i)_{(v,v)} + q \cdot (Y_i)_{(w,v)} \\
(Y_{i+1})_{(v,w)} = f^{a^*}_{(v,w)}(Y_i) &= \pi[\![1?]\!]_{(v,w)} + \sum_{u \in V} \pi[\![a]\!]_{(v,u)} \cdot (Y_i)_{(u,w)} \\
&= p \cdot (Y_i)_{(v,w)} + q \cdot (Y_i)_{(w,w)} \\
(Y_{i+1})_{(w,v)} = f^{a^*}_{(w,v)}(Y_i) &= \pi[\![1?]\!]_{(w,v)} + \sum_{u \in V} \pi[\![a]\!]_{(w,u)} \cdot (Y_i)_{(u,v)} \\
&= 0 \qquad \text{and} \\
(Y_{i+1})_{(w,w)} = f^{a^*}_{(w,w)}(Y_i) &= \pi[\![1?]\!]_{(w,w)} + \sum_{u \in V} \pi[\![a]\!]_{(w,u)} \cdot (Y_i)_{(u,w)} \\
&= 1.
\end{aligned}$$

We compute a table for $0 \leq i \leq 5$:

i	0	1	2	3	4	5	...
$(Y_i)_{(v,v)}$	0	1	$1+p$	$1+p+p^2$	$1+p+p^2+p^3$	$1+p+p^2+p^3+p^4$	...
$(Y_i)_{(v,w)}$	0	0	q	$q+pq$	$q+pq+p^2q$	$q+pq+p^2q+p^3q$	...
$(Y_i)_{(w,v)}$	0	0	0	0	0	0	...
$(Y_i)_{(w,w)}$	0	1	1	1	1	1	...

Clearly, we have $\sup_{i\in\omega}(Y_i)_{(v,v)} = p^*$ and $\sup_{i\in\omega}(Y_i)_{(v,w)} = p^*q$. Therefore,

$$\begin{aligned}
\pi[\![a^*]\!]_{(v,v)} = \mathrm{lfp}(f^{a^*})_{(v,v)} &= \sup_{i\in\omega}(Y_i)_{(v,v)} = p^* \quad \text{and} \\
\pi[\![a^*]\!]_{(v,w)} = \mathrm{lfp}(f^{a^*})_{(v,w)} &= \sup_{i\in\omega}(Y_i)_{(v,w)} = p^*q.
\end{aligned}$$

With these results, we can evaluate $\langle a^* \rangle P$ at v as

$$\begin{aligned}
\pi[\![\langle a^* \rangle P]\!]_v &= \sum_{u \in V} \pi[\![a^*]\!]_{(v,u)} \cdot \pi[\![P]\!]_u \\
&= \pi[\![a^*]\!]_{(v,v)} \cdot \pi[\![P]\!]_v + \pi[\![a^*]\!]_{(v,w)} \cdot \pi[\![P]\!]_w \\
&= p^*r + p^*qs.
\end{aligned}$$

We can check this result for sanity. To prove $\langle a^* \rangle P$ at node v, we can clearly take the p-labelled edge $(v,v) \in E_a$ arbitrarily often, hence the factor p^*. Then, we can use the fact P at v, which is labelled with r, or we can take the edge $(v,w) \in E_a$ and the fact P at w, which are labelled with q and s respectively to close the proof. So, the expected result is indeed $p^*(r+qs) = p^*r + p^*qs$.

(3.24) Remark. There is no obvious way to extend the above approach to universal formulas $[\rho]\varphi$ where ρ is a program and $\varphi \in \mathrm{PDL}(\tau)$.

The naive approach to handle formulas $[\rho]\varphi$ would be setting

$$\pi[\![[\rho]\varphi]\!]_v = \prod_{w \in V, \pi[\![\rho]\!]_{(v,w)} \neq 0} \pi[\![\rho]\!]_{(v,w)} \cdot \pi[\![\varphi]\!]_w.$$

The drawback of this is that programs change their meanings depending on whether they are used existentially or universally. Consider the formulas $\langle a \cup b \rangle P$ and $[a \cup b]P$ in PDL(τ) with $\{E_a, E_b, P\} \subseteq \tau$ and the $\mathbb{S}^\infty[X]$-interpretation π in figure (3.25) with $X = \{p, q, r, s, x, y, z\}$.

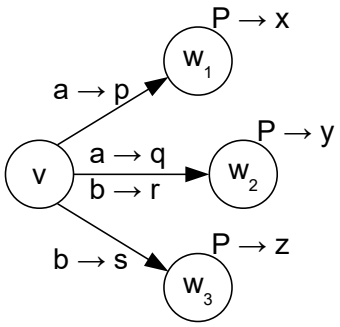

Figure (3.25): $\mathbb{S}^\infty[X]$-interpretation π.

Using the naive approach, we would obtain

$$\begin{aligned} \pi[\![\langle a \cup b \rangle P]\!]_v &= px + (q + r)y + sz = px + qy + ry + sz \quad \text{and} \\ \pi[\![[a \cup b]P]\!]_v &= px \cdot (q + r)y \cdot sz \quad = pxqysz + pxrysz. \end{aligned}$$

The existential formula is interpreted correctly, since any monomial represents a valid proof for $\langle a \cup b \rangle P$ at v by taking an edge labelled with a or b and then showing P at the successor. For the universal formula, the problem lies in

$$\pi[\![a \cup b]\!]_{(v,w_2)} = \pi[\![a]\!]_{(v,w_2)} + \pi[\![b]\!]_{(v,w_2)} = q + r.$$

This is counter-intuitive as it would be expected for the universal formula to use both q and r instead of choosing between them, since (v, w_2) consists of two edges, one labelled with a and one with b. We could try to fix this by interpreting programs differently depending on whether they are used existentially or universally, for example, we would set

$$\begin{aligned} \pi[\![\langle a \cup b \rangle]\!]_{(v,w_2)} &= \pi[\![\langle a \rangle]\!]_{(v,w_2)} + \pi[\![\langle b \rangle]\!]_{(v,w_2)} = q + r \quad \text{as before and} \\ \pi[\![[a \cup b]]\!]_{(v,w_2)} &= \pi[\![[a]]\!]_{(v,w_2)} \cdot \pi[\![[b]]\!]_{(v,w_2)} \quad = qr. \end{aligned}$$

This would yield $\pi[\![[a \cup b]P]\!]_v = px \cdot qry \cdot sz = pxqrysz$. However, this is still unsatisfactory, as we would need the fact that P holds at w_2 twice, since w_2 is reached from v in two separate ways, so the expected result is $pxqry^2sz$.

Resolving this problem would require a new definition of the semiring interpretation that breaks the separation between programs and formulas. When evaluating the program $[a \cup b]$ in the above example, we would have to keep in mind that w_2 is reached twice and that the formula P will have to be evaluated twice at w_2. In order to avoid these problems, we restrict our semiring interpretation to positive PDL formulas.

Chapter 4

Algorithms for Semiring Interpretation

The next goal is to show how the semiring interpretations from the previous chapter can be performed algorithmically. Given a formula θ in LTL, CTL or posPDL in negation normal form and a suitable K-interpretation π over a finite set of nodes V, we would like to compute the semiring interpretation $\pi[\![\theta]\!]_v$ for some $v \in V$. Assuming that addition and multiplication in K can be performed in constant time, most types of formulas can be interpreted trivially. For example, if $\theta = \varphi \vee \psi$ is a disjunction and the interpretations $\pi[\![\varphi]\!]_v$ and $\pi[\![\psi]\!]_v$ are already known, then we have $\pi[\![\theta]\!]_v = \pi[\![\varphi]\!]_v + \pi[\![\psi]\!]_v$, which is easy to compute under the above assumption.

However, the following types of formulas or programs are interpreted in terms of fixed points:

- until formulas $\mathrm{E}(\varphi \mathrm{U} \psi)$ and $\mathrm{A}(\varphi \mathrm{U} \psi)$ in CTL,
- release formulas $\mathrm{E}(\varphi \mathrm{R} \psi)$ and $\mathrm{A}(\varphi \mathrm{R} \psi)$ in CTL and
- program iterations ρ^* in PDL.

It is not immediately clear how to interpret the above formulas and programs algorithmically, since we have to find the corresponding fixed points, but so far, we have only established theorem (2.23) and corollary (2.25) to do that, and both make use of infinite or transfinite sequences. In this chapter, we will introduce alternative approaches to find the desired fixed points without infinite iterations and show that it can be done algorithmically.

4.1 Paths and Complete Trees

Looking back at the examples (3.13), (3.15) and (3.22), where we evaluated existential until and release formulas in CTL as well as a program iteration in PDL, we notice that all the monomials in the results represent paths over the given transition system. In the following sections, we will see that this is not a coincidence and we will capture this concept formally. Consider a formula $\mathrm{E}(\varphi \mathrm{U} \psi)$ in $\mathrm{CTL}(\tau)$ that is

interpreted at a node $v \in V$, where $\mathcal{G} = (V, \tau)$ is a transition system. Proving that $\mathrm{E}(\varphi \mathrm{U} \psi)$ holds at v requires a path in $\mathcal{G}$ that starts at v and ends at some node $w \in V$ where ψ holds, while φ is true at the non-terminal nodes of the path. Figure (4.1) illustrates a transition system $\mathcal{G}$ and a path over $\mathcal{G}$ that witnesses that $\mathrm{E}(P\mathrm{U}Q)$ holds at v.

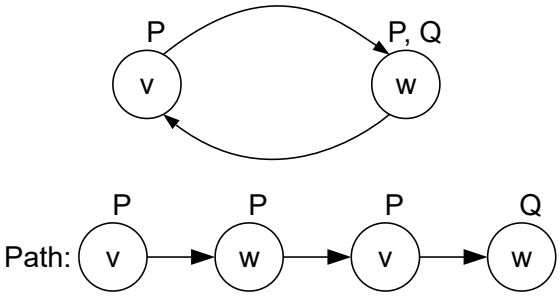

Figure (4.1): Transition system with path that witnesses $\mathrm{E}(P\mathrm{U}Q)$ at v.

We have labelled the nodes of the path with the facts that we used for the proof of $\mathrm{E}(P\mathrm{U}Q)$ at v. So, on the first traversal of w, we have used the fact P at w and continued the proof, but on the second traversal, we have used the fact Q and ended the proof, since this is enough to witness that $P\mathrm{U}Q$ holds on the path, and as the path starts at v, we know that $\mathrm{E}(P\mathrm{U}Q)$ is true at v. Of course, we could have obtained a shorter proof by using the fact Q at the first traversal of w, or we could build arbitrarily long proofs by repeatedly using P at the first $n \in \mathbb{N}$ traversals of w and then using Q. The example gives rise to the following formal definition.

(4.2) Definition (Path)**.** A *path over* V p is a finite or infinite path whose nodes are labelled with elements of V. We denote

the set of nodes of p	as $n(p)$,
the set of edges of p	as $e(p)$ and
the set of non-terminal, or internal nodes of p	as $i(p) \subseteq n(p)$.

If p is finite, then there is also a terminal node $t(p) \in n(p)$. The path has a label function $L_p : n(p) \to V$, which can be extended to edges by setting $L_p((x, y)) = (L_p(x), L_p(y)) \in V \times V$ for $(x, y) \in e(p)$.

$P(V)$ denotes the set of all paths over V. We can add subscripts or superscripts to refer to specific subsets of $P(V)$. Subscripts denote the starting node of the paths, for example, if $v \in V$, then $P_v(V) \subseteq P(V)$ denotes the set of paths over V that start at a node that is labelled with v. Superscripts will be used to indicate the length of the path, which we define as $|e(p)|$, the number of edges on the path. Notice that the length of a path is a cardinal number κ and that the length of infinite paths is ω. With that in mind, for example, $P^{\geq\omega}(V)$ refers to the set of infinite paths over V and $P^{<52}(V)$ refers to the set of paths over V with less than 52 edges. We also introduce $P^{\mathrm{fin}}(V)$ and $P^{\mathrm{inf}}(V)$ to refer to the set of finite and infinite paths over V respectively. Of course, we will often combine subscripts and superscripts.

Also, we will usually identify the nodes of a path with their labels. For example, consider the path from figure (4.1). We denote it as $p = (v, w, v, w)$ and we would say that it starts at v and ends at w instead of saying that it starts at a node labelled with v and ends at a node labelled with w. We would also say that v occurs twice in p, meaning that there are two separate nodes in p that are labelled with v. The

label function will only be used for formal definitions.

Paths are useful for semiring interpretations of existential formulas. In figure (4.3) below, we extend the transition system from figure (4.1) to a $\mathbb{S}^\infty[X]$-interpretation π with $X = \{p, q, r, s, t\}$.

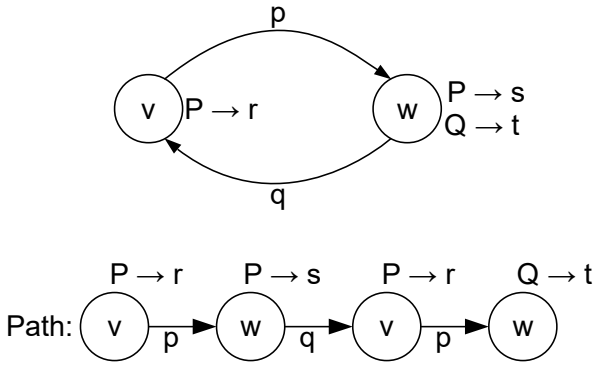

Figure (4.3): $\mathbb{S}^\infty[X]$-interpretation π with path that witnesses $\mathrm{E}(P\mathrm{U}Q)$ at v.

The same path $p_0 = (v, w, v, w)$ as in figure (4.1) is also depicted here in figure (4.3), but the facts are now labelled with values from $\mathbb{S}^\infty[X]$. We can now define the $P\mathrm{U}Q$-*cost* of p_0 under π as

$$\pi[\![P\mathrm{U}Q]\!]_{p_0} = r \cdot p \cdot s \cdot q \cdot r \cdot p \cdot t = p^2qr^2st.$$

The $P\mathrm{U}Q$-cost of p_0 indicates which facts are used to prove that $P\mathrm{U}Q$ holds on p. First, we have to make sure that the edges exist, hence the factors p^2q. Then, P must be true at the internal nodes of p_0, therefore we have the factors r^2s. Finally, Q must be true at the terminal node, so we have the last factor t.

This concept can be generalized to arbitrary until formulas $\varphi\mathrm{U}\psi$ and finite paths p. The $\varphi\mathrm{U}\psi$-cost of p is defined by evaluating φ at the internal nodes, taking the edges into account and then evaluating ψ at the terminal node. The formal definition is given below.

(4.4) Definition (Until-Costs for Paths)**.** Let K be ω-continuous. For a K-interpretation π and a finite path p over V, we define the $\varphi\mathrm{U}\psi$-cost of p as

$$\pi[\![\varphi\mathrm{U}\psi]\!]_p = \left(\prod_{x \in i(p)} \pi[\![\varphi]\!]_{L_p(x)}\right) \cdot \left(\prod_{(x,y) \in e(p)} \pi(EL_p(x)L_p(y))\right) \cdot \pi[\![\psi]\!]_{L_p(t(p))}$$

where φ, ψ are CTL formulas.

Clearly, $\varphi\mathrm{U}\psi$-costs for infinite paths are not defined, since $\varphi\mathrm{U}\psi$ requires ψ to come true at some point on a path, so infinite paths are not suited to witness $\mathrm{E}(\varphi\mathrm{U}\psi)$.

A similar concept can be introduced for release formulas. Consider a path p and a release formula $\varphi\mathrm{R}\psi$. If p is finite, proving $\varphi\mathrm{R}\psi$ on p requires us to use the edges, show that ψ holds on all nodes and show that φ is true at the terminal node. Also, unlike until formulas, release formulas can be witnessed by infinite paths. If p is infinite, it witnesses $\varphi\mathrm{R}\psi$ if we can show that the edges exist and that ψ is true on all nodes. So, the $\varphi\mathrm{R}\psi$-cost of an infinite path is a countable product, but this is no problem, since release formulas are interpreted in absorptive lattice semirings, which

admit countable products according to definition (2.15). We obtain the following definition.

(4.5) Definition (Release-Costs for Paths)**.** Let K be an absorptive lattice semiring. If π is a K-interpretation and p is a path over V, then for CTL formulas φ and ψ, the $\varphi \mathrm{R} \psi$-cost of p is defined as

$$\pi[\![\varphi \mathrm{R} \psi]\!]_p = \left(\prod_{x \in n(p)} \pi[\![\psi]\!]_{L_p(x)} \right) \cdot \left(\prod_{(x,y) \in e(p)} \pi(EL_p(x)L_p(y)) \right) \cdot \pi[\![\varphi]\!]_{L_p(t(p))}$$

if p is finite, otherwise, it is defined as

$$\pi[\![\varphi \mathrm{R} \psi]\!]_p = \left(\prod_{x \in n(p)} \pi[\![\psi]\!]_{L_p(x)} \right) \cdot \left(\prod_{(x,y) \in e(p)} \pi(EL_p(x)L_p(y)) \right).$$

The third and final type of costs that we will define for paths is the ρ-cost of a path p, where ρ is a program in PDL. If p is a finite path over V starting at $v \in V$ and ending at $w \in V$, it witnesses that w is reachable by ρ-transitions from v if we evaluate ρ on every edge in p. We obtain a straightforward definition of the ρ-costs of paths.

(4.6) Definition (Program-Costs for Paths)**.** Let K be an ω-continuous semiring. For a K-interpretation π and a finite path p over V, the ρ-cost of p, where ρ is a PDL program, is defined as

$$\pi[\![\rho]\!]_p = \prod_{e \in e(p)} \pi[\![\rho]\!]_{L_p(e)}.$$

As for until formulas, infinite paths are disregarded, because infinite paths are not suited to witness that a node w is reachable from a node v. Figure (4.7) informally illustrates the three types of path costs that we have defined above.

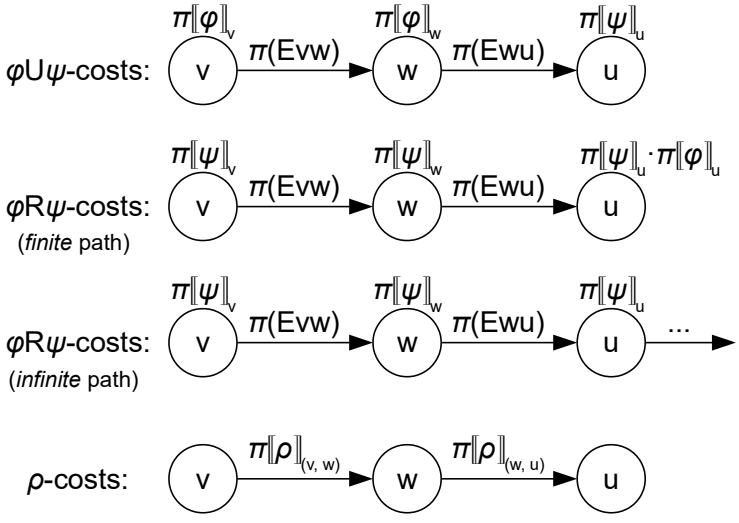

Figure (4.7): Informal illustration of path costs.

So far, we have only considered existential formulas and paths. Obviously, to show that an existential formula like $\mathrm{E}(\varphi \mathrm{U} \psi)$ in CTL is true at $v \in V$, we just need one

path p that starts at v and fulfills $\varphi \mathrm{U} \psi$. However, when we consider a universal formula $\mathrm{A}(\varphi \mathrm{U} \psi)$ in CTL, a single path will generally not be suited to witness the truth of this formula at a node v. We will use trees instead of paths for universal formulas. To justify this, consider the formula $\mathrm{A}(P \mathrm{U} Q)$ interpreted in the transition system $\mathcal{G}$ from figure (4.8) at node v as an example.

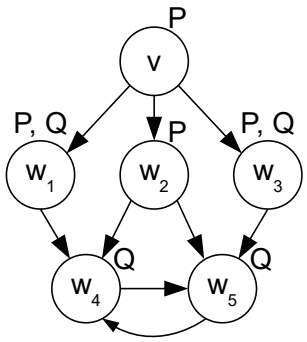

Figure (4.8): Transition system $\mathcal{G}$.

We claim that the following two trees t_1 and t_2 from figure (4.9) each witness the truth of $\mathrm{A}(P \mathrm{U} Q)$ at node v in $\mathcal{G}$. The nodes of the trees are labelled with the atomic propositions that are used to prove $\mathrm{A}(P \mathrm{U} Q)$ with the corresponding tree.

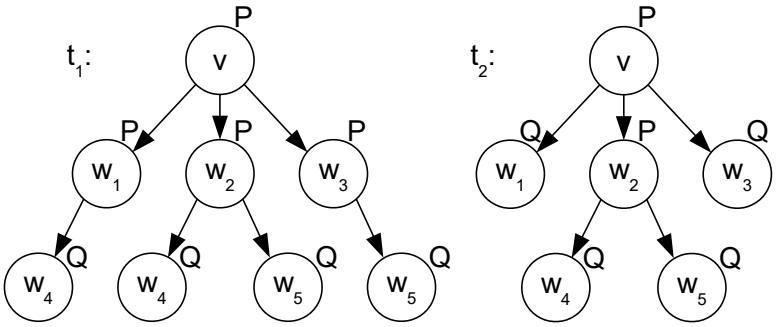

Figure (4.9): Two trees over $\mathcal{G}$.

The trees are constructed by starting at v and trying to prove $\mathrm{A}(P \mathrm{U} Q)$. Since Q is not true at v, in both cases, we have to use P at v and show $\mathrm{A}(P \mathrm{U} Q)$ at all three successors of v. Since Q also does not hold at w_2, we have no choice but using P and continuing to the successors of w_2. As for w_1 and w_3, since both P and Q are true there, we have the option to use Q directly, which is done in t_2, or we can use P and continue to the successors of w_1 and w_3, like in t_1. At w_4 or w_5, we always use Q and end the branch of the tree.

We call both t_1 and t_2 *complete* trees with respect to $\mathcal{G}$. In a complete tree, any non-leaf node has exactly the same successors as the corresponding node in the transition system $\mathcal{G}$. This property is crucial for proving universal formulas, because it ensures that all infinite paths starting from v are covered by the tree and eventually cross one of the leaves. Since we have P at all internal nodes and Q at all the leaf nodes, we know that any infinite path starting at v has to cross a leaf and therefore fulfill $P \mathrm{U} Q$. Like for the paths, we will capture this concept formally.

(4.10) Definition (Complete Tree)**.** A *complete tree over* V t with respect to a transition system $\mathcal{G} = (V, E, (P_i)_{i \in I})$ is a finite or infinite tree whose nodes are

labelled with elements of V. We denote

the set of nodes of t	as $n(t)$,
the set of edges of t	as $e(t)$,
the set of terminal, or leaf nodes of t	as $l(t) \subseteq n(p)$ and
the set of non-terminal, or internal nodes of p	as $i(t) \subseteq n(p)$.

The label function is denoted as $L_t : n(t) \to V$ and the tree fulfills the completeness property, that is, for every node $x \in n(t)$, if $L_t(x) = v \in V$, then for every outgoing edge of v $(v, w) \in E$, there is exactly one successor y of x with $(x, y) \in e(t)$ and $L(y) = w$ and for every successor z of x with $(x, z) \in e(t)$ and $L_t(z) = u$, there is an outgoing edge of v $(v, u) \in E$. As for paths, the label function can be extended to edges by setting $L_t((x, y)) = (L_t(x), L_t(y)) \in E$ for $(x, y) \in e(t)$.

Notice that this definition can be extended to K-interpretations π instead of transition systems $\mathcal{G}$. If K is a semiring, V is a set of nodes and $\pi : \mathrm{Lit}_V(\tau) \to K$ is a K-interpretation where $E \in \tau$, then π induces an edge relation

$$E_\pi = \{(v, w) \in V \times V \mid \pi(Evw) \neq 0\}.$$

Therefore, we can also define complete trees with respect to a K-interpretation π by applying the completeness property to E_π. This is justified, because edges that are tagged with 0 by a K-interpretation are "absent" edges, and edges tagged with nonzero values are "present" edges when performing provenance analysis.

If the edge relation, given by a transition system or a K-interpretation, is clear from the context, then $T(V)$ denotes the set of all complete trees over V. The same conventions that we established for paths apply here as well, for example, we will often disregard the label functions. Also, since all trees that we will consider are complete, we will often write "trees" instead of "complete trees". Subscripts denote the root of the trees, so $T_v(V)$ denotes the set of all trees over V rooted at v. Superscripts indicate the height of the trees, which is defined as the length of a path with maximal length in the tree that starts at the root. Notice that this is a cardinal number, which is finite for finite trees and ω for infinite trees. For example, $T^{\leq 2}(V)$ denotes the set of trees over V up to height 2 and $T^{\mathrm{fin}}(V)$ and $T^{\mathrm{inf}}(V)$ denote the finite and infinite trees over V respectively.

Just like we have defined $\varphi \mathrm{U} \psi$-costs for paths, we can define $\varphi \mathrm{U} \psi$-costs for trees. Suppose that t is a complete tree rooted at v and we want to use t to witness that $\mathrm{A}(\varphi \mathrm{U} \psi)$ holds at v. First, we would have to show that the edges exist, then we would have to make sure that ψ holds at all the leaf nodes and also, we would have to prove that φ is true at the internal nodes. Notice that t has to be finite, because any path from the root has to end in a leaf where ψ holds. An infinite φ-path is not enough to satisfy an until formula. This yields the following definition.

(4.11) Definition (Until-Costs for Trees)**.** Let K be ω-continuous. For a K-interpretation π and a finite, complete tree t over V, the $\varphi \mathrm{U} \psi$-cost of t for CTL formulas φ and ψ is

$$\pi[\![\varphi \mathrm{U} \psi]\!]_t = \left(\prod_{x \in i(t)} \pi[\![\varphi]\!]_{L_t(x)} \right) \cdot \left(\prod_{(x,y) \in e(t)} \pi(EL_t(x)L_t(y)) \right) \cdot \left(\prod_{x \in l(t)} \pi[\![\psi]\!]_{L_t(x)} \right).$$

For the $\varphi \mathrm{R} \psi$-costs of trees, we use a similar approach. Assume that t is a complete tree rooted at v and we want to use it to show that $\mathrm{A}(\varphi \mathrm{R} \psi)$ holds at v. We would have to prove that ψ is true at all nodes of t, φ is true at the leaf nodes and the edges all exist. However, t is not necessarily finite. For example, some branches could end in leaves where φ is true, whereas other branches could be infinite ψ-paths. In fact, t does not need to have any leaves at all. Therefore, we need countable products to define the $\varphi \mathrm{R} \psi$-costs of trees.

(4.12) Definition (Release-Costs for Trees)**.** Let K be an absorptive lattice semiring. If π is a K-interpretation and t is a complete tree over V, we define the $\varphi \mathrm{R} \psi$-cost of t as

$$\pi[\![\varphi \mathrm{R} \psi]\!]_t = \left(\prod_{x \in n(t)} \pi[\![\psi]\!]_{L_t(x)} \right) \cdot \left(\prod_{(x,y) \in e(t)} \pi(E L_t(x) L_t(y)) \right) \cdot \left(\prod_{x \in l(t)} \pi[\![\varphi]\!]_{L_t(x)} \right),$$

where φ and ψ are CTL formulas.

For now, we have used intuitive arguments to define paths, trees and their costs. In the following sections, we will prove that the fixed points that are defined in the previous chapter can be expressed in terms of path or tree costs and from that, we will derive algorithms that compute the desired fixed points.

4.2 Until Operators in CTL

Recall example (3.13) where we had the $\mathbb{N}^\infty[\![X]\!]$-interpretation π given below in figure (4.13) and obtained $\pi[\![\mathrm{E}(\mathrm{F} P)]\!]_v = \pi[\![\mathrm{E}(1 \mathrm{U} P)]\!]_v = p(pq)^* r$.

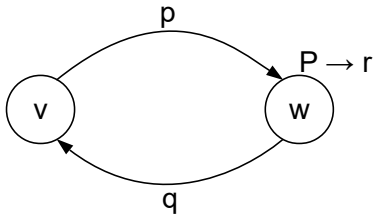

Figure (4.13): $\mathbb{N}^\infty[\![X]\!]$-interpretation π.

Now, consider the set $P_v^{\mathrm{fin}}(V)$ of all finite paths over the above transition system that start at v. When we consider the $1 \mathrm{U} P$-costs of those paths, many of them will have the costs 0, for example, if their terminal node is not w or if they use transitions other than (v, w) and (w, v), because those are labelled with 0. So, all the paths $p_0 \in P_v^{\mathrm{fin}}(V)$ with $\pi[\![1 \mathrm{U} P]\!]_{p_0} \neq 0$ have the form $(v, (w, v)^i, w)$ for some $i \in \mathbb{N}$, that is, they start at v and end at w and the cycle (w, v, w) is repeated i times in between. With the observation that

$$\pi[\![1 \mathrm{U} P]\!]_{(v,(w,v)^i,w)} = p(pq)^i r$$

for all $i \in \mathbb{N}$, we can draw the interesting conclusion that

$$\pi[\![\mathrm{E}(1 \mathrm{U} P)]\!]_v = p(pq)^* r = \sum_{i \in \mathbb{N}} p(pq)^i r = \sum_{p_0 \in P_v^{\mathrm{fin}}(V)} \pi[\![1 \mathrm{U} P]\!]_{p_0}.$$

Putting this in words, the interpretation of $\mathrm{E}(1\mathrm{U}P)$ at v is the same as the sum of the $1\mathrm{U}P$-costs of all finite paths starting at v. This is not surprising, since the $1\mathrm{U}P$-costs of finite paths starting at v exactly represent proofs for $\mathrm{E}(1\mathrm{U}P)$ at v.

We will generalize this to arbitrary formulas $\mathrm{E}(\varphi\mathrm{U}\psi)$. Moreover, we can obtain a similar result for universal formulas by replacing paths with complete trees. We claim that the interpretation of $\mathrm{A}(\varphi\mathrm{U}\psi)$ at v yields the sum of the $\varphi\mathrm{U}\psi$-costs of all finite, complete trees rooted at v. This yields the following theorem.

(4.14) Theorem. Let K be an ω-continuous semiring, V a finite set of nodes and π a K-interpretation over V. For arbitrary formulas φ and ψ in CTL and any $v \in V$, we have

$$(1) \quad \pi[\![\mathrm{E}(\varphi\mathrm{U}\psi)]\!]_v = \sum_{p \in P_v^{\mathrm{fin}}(V)} \pi[\![\varphi\mathrm{U}\psi]\!]_p \quad \text{and}$$

$$(2) \quad \pi[\![\mathrm{A}(\varphi\mathrm{U}\psi)]\!]_v = \sum_{t \in T_v^{\mathrm{fin}}(V)} \pi[\![\varphi\mathrm{U}\psi]\!]_t.$$

Proof. First of all, we notice that the above summations are well-defined in ω-continuous semirings, since $P_v^{\mathrm{fin}}(V)$ and $T_v^{\mathrm{fin}}(V)$ are both countable, as we can enumerate the paths by their length and the trees by their height. The finiteness of V ensures that for any given length or height i, there are only finitely many distinct paths or trees over V respectively.

Now, we prove part (1) of the theorem by calculating $\mathrm{lfp}(f^{\mathrm{E}(\varphi\mathrm{U}\psi)})$ according to definition (3.11). Using lemma (2.23), this is done by setting $X_i = (f^{\mathrm{E}(\varphi\mathrm{U}\psi)})^i(0)$ for $i \in \omega$ and then calculating $\sup_{i\in\omega} X_i$. We will show by induction on i that

$$(X_i)_v = \sum_{p \in P_v^{<i}(V)} \pi[\![\varphi\mathrm{U}\psi]\!]_p \quad \text{for all } v \in V.$$

For $i = 0$, this is clearly true as $(X_0)_v = 0$ for $v \in V$, and since there are no paths of length less than 0, $P_v^{<0}(V)$ is empty and the empty sum is always 0.

Now, assume the hypothesis holds for i and consider $i+1$. We have

$$\begin{aligned}(X_{i+1})_v &= \pi[\![\psi]\!]_v + \pi[\![\varphi]\!]_v \cdot \sum_{w \in vE} \pi(Evw) \cdot (X_i)_w \\ &= \pi[\![\psi]\!]_v + \sum_{w \in vE} \pi[\![\varphi]\!]_v \cdot \pi(Evw) \cdot \left(\sum_{p \in P_w^{<i}(V)} \pi[\![\varphi\mathrm{U}\psi]\!]_p \right)\end{aligned}$$

Notice that $\pi[\![\psi]\!]_v$ is exactly the cost of the path (v). Also, multiplying $\pi[\![\varphi]\!]_v$ and $\pi(Evw)$ to the cost of a path $p \in P_w^{<i}(V)$ yields the costs of the path (v,p), which is obtained by appending v to the start of p. Notice that (v,p) starts at v and has a length between 1 and i, so we have

$$\begin{aligned}(X_{i+1})_v &= \pi[\![\varphi\mathrm{U}\psi]\!]_{(v)} + \sum_{w \in vE} \sum_{p \in P_w^{<i}(V)} \pi[\![\varphi\mathrm{U}\psi]\!]_{(v,p)} \\ &= \sum_{p \in P_v^{=0}(V)} \pi[\![\varphi\mathrm{U}\psi]\!]_p + \sum_{p \in \bigcup_{1 \le j \le i} P_v^{=j}(V)} \pi[\![\varphi\mathrm{U}\psi]\!]_p \\ &= \sum_{p \in P_v^{<i+1}} \pi[\![\varphi\mathrm{U}\psi]\!]_p,\end{aligned}$$

since (v) is the only path of length 0 starting at v. This ends the induction.

Now, we can use proposition (2.5) that states that countable summation in ω-continuous semirings is invariant under partition. Therefore, we partition the finite paths over V starting at v by their length and obtain

$$\begin{aligned}
\sum_{p \in P_v^{\mathrm{fin}}(V)} \pi[\![\varphi \mathrm{U} \psi]\!]_p &= \sum_{i \in \omega} \sum_{p \in P_v^{=i}(V)} \pi[\![\varphi \mathrm{U} \psi]\!]_p \\
&= \sup_{i \in \omega} \left(\sum_{p \in P_v^{=0}(V)} \pi[\![\varphi \mathrm{U} \psi]\!]_p + \ldots + \sum_{p \in P_v^{=i}(V)} \pi[\![\varphi \mathrm{U} \psi]\!]_p \right) \\
&= \sup_{i \in \omega} \sum_{p \in P_v^{<i+1}(V)} \pi[\![\varphi \mathrm{U} \psi]\!]_p \\
&= \sup_{i \in \omega} (X_{i+1})_v \\
&= \sup_{i \in \omega \setminus \{0\}} (X_i)_v.
\end{aligned}$$

Since $(X_0)_v$ is 0, we can add it into the supremum. Using lemma (2.21), we have

$$\begin{aligned}
\sum_{p \in P_v^{\mathrm{fin}}(V)} \pi[\![\varphi \mathrm{U} \psi]\!]_p &= \sup_{i \in \omega \setminus \{0\}} (X_i)_v \\
&= \sup_{i \in \omega} (X_i)_v \\
&= (\sup_{i \in \omega} X_i)_v \\
&= \mathrm{lfp}(f^{\mathrm{E}(\varphi \mathrm{U} \psi)})_v \\
&= \pi[\![\mathrm{E}(\varphi \mathrm{U} \psi)]\!]_v.
\end{aligned}$$

This ends the proof for (1). The proof of part (2) is very similar. Here, we need to calculate $\mathrm{lfp}(f^{\mathrm{A}(\varphi \mathrm{U} \psi)})$, which is done by setting $Y_i = (f^{\mathrm{A}(\varphi \mathrm{U} \psi)})^i(0)$ for $i \in \omega$ and then calculating the supremum. It will suffice to show by induction that for $i \in \omega$, we have

$$(Y_i)_v = \sum_{t \in T_v^{<i}(V)} \pi[\![\varphi \mathrm{U} \psi]\!]_t.$$

The rest of the proof is the same as for part (1), except that we replace X with Y and paths with trees. For $i = 0$, there is nothing to be shown, since there are no trees of height less than 0 and $(Y_0)_v = 0$.

If the induction hypothesis holds for i, then for $i + 1$, we have

$$\begin{aligned}
(Y_{i+1})_v &= \pi[\![\psi]\!]_v + \pi[\![\varphi]\!]_v \cdot \prod_{w \in vE} \pi(Evw) \cdot (Y_i)_w \\
&= \pi[\![\psi]\!]_v + \pi[\![\varphi]\!]_v \cdot \prod_{w \in vE} \pi(Evw) \cdot \left(\sum_{t \in T_w^{<i}(V)} \pi[\![\varphi \mathrm{U} \psi]\!]_t \right).
\end{aligned}$$

Since π is non-terminating and V is finite, vE is non-empty and finite, so w.l.o.g., we can write $vE = \{w_1, \ldots, w_l\} \subseteq V$ for some $l \in \mathbb{N}$. Every tuple $(t_1, \ldots, t_l)$ with $t_j \in T_{w_j}^{<i}(V)$ along with the edges (v, w_j) for $1 \leq j \leq l$ and the root node v forms

a complete tree over V with a height h between 1 and i. Conversely, any complete tree over V rooted at v with height $1 \leq h \leq i$ can be split into v, the edges (v, w_j) for $1 \leq j \leq l$ and a tuple of trees $(t_1, ..., t_l)$ such that $t_j \in T^{<i}_{w_j}(V)$. Therefore, we can simplify the product to

$$\begin{aligned}(Y_{i+1})_v &= \pi[\![\psi]\!]_v + \sum_{(t_1,...,t_l)\in T^{<i}_{w_1}(V)\times\ldots\times T^{<i}_{w_l}(V)} \pi[\![\varphi]\!]_v \cdot \left(\prod_{j=1}^{l} \pi(Evw_l)\right) \cdot \left(\prod_{j=1}^{l} \pi[\![\varphi \mathrm{U} \psi]\!]_{t_j}\right) \\ &= \pi[\![\psi]\!]_v + \sum_{t\in\bigcup_{1\leq h\leq i} T^{=i}_v(V)} \pi[\![\varphi \mathrm{U} \psi]\!]_t \\ &= \sum_{t\in T^{<i+1}_v(V)} \pi[\![\varphi \mathrm{U} \psi]\!]_t.\end{aligned}$$

The last step is justified by noticing that (v) is the only complete tree of height 0 rooted at v and its cost is $\pi[\![\psi]\!]_v$. Since this closes the induction, we can now use the same argument as above to obtain

$$\begin{aligned}\sum_{t\in T^{\mathrm{fin}}_v(V)} \pi[\![\varphi \mathrm{U} \psi]\!]_t &= (\sup_{i\in\omega} Y_i)_v \\ &= \mathrm{lfp}(f^{\mathrm{A}(\varphi \mathrm{U} \psi)})_v \\ &= \pi[\![\mathrm{A}(\varphi \mathrm{U} \psi)]\!]_v,\end{aligned}$$

which ends the proof. □

Of course, even with theorem (4.14) above, there is still no obvious way to compute $\pi[\![\mathrm{E}(\varphi \mathrm{U} \psi)]\!]_v$ and $\pi[\![\mathrm{A}(\varphi \mathrm{U} \psi)]\!]_v$ algorithmically, since we have just transformed the fixed points that required infinite iterations into infinite sums. However, since our transition systems only have a finite set of nodes V, long paths and large trees are bound to contain repetitions and recurring patterns. For example, if we have n nodes and a path is longer than n, the path must traverse a cycle. Obviously, the presence of cycles implies that there are infinitely many paths and we cannot hope to enumerate all the paths. However, if there is a single cycle with cost c, then the expression c^* covers the costs of arbitrarily many repetitions of the cycle. We will exploit this to compute representations of $\pi[\![\mathrm{E}(\varphi \mathrm{U} \psi)]\!]_v$ and $\pi[\![\mathrm{A}(\varphi \mathrm{U} \psi)]\!]_v$.

4.2.1 Existential Until Operators

First, we will look at the existential until formulas $\mathrm{E}(\varphi \mathrm{U} \psi)$. Let K be ω-continuous, V a finite set of nodes and π an arbitrary K-interpretation. In order to calculate $\pi[\![\mathrm{E}(\varphi \mathrm{U} \psi)]\!]_v$ for some $v \in V$, we will use results from automata theory. First of all, we will transform the K-interpretation π into a K-automaton $\mathcal{A}_v(\pi)$. Intuitively, a K-automaton $\mathcal{A}$ is a transition system with a starting state s and a final state t and edges are labelled with elements in K by a cost function. Our goal is to build $\mathcal{A}_v(\pi)$ for $\mathrm{E}(\varphi \mathrm{U} \psi)$ in such a way that $\mathcal{A}_v(\pi)$ "recognizes" proofs for $\mathrm{E}(\varphi \mathrm{U} \psi)$ at v, that is, the paths p from s to t in $\mathcal{A}_v(\pi)$ should have the same costs as paths $p' \in P^{\mathrm{fin}}_v(V)$ over V that witness $\mathrm{E}(\varphi \mathrm{U} \psi)$ at v. We will first define K-automata formally and then describe how $\mathcal{A}_v(\pi)$ can be built.

(4.15) Definition (K-Automaton)**.** A K-automaton for an ω-continuous semiring K is a structure $\mathcal{A} = (Q, C, s, t)$ where Q is a finite set of states, $C : Q \times Q \to K$ is

a cost function and s and t are starting and final states respectively, so $C(q, s) = 0$ and $C(t, q) = 0$ for $q \in Q$. We extend C to finite paths over Q by setting $C(p)$ for $p \in P^{\text{fin}}(Q)$ as the product of the costs of the edges in p, that is

$$C(p) = \prod_{(x,y)\in e(p)} C(L(x), L(y)).$$

Moreover, $P_{s\to t}(Q)$ is the set of all paths over Q going from s to t and we set

$$C(\mathcal{A}) = \sum_{p\in P_{s\to t}(Q)} C(p).$$

When transforming π into $\mathcal{A}_v(\pi) = (Q, C, s, t)$, we would like to accomplish

$$C(\mathcal{A}_v(\pi)) = \pi[\![\mathrm{E}(\varphi \mathrm{U} \psi)]\!]_v = \sum_{p\in P_v^{\text{fin}}(V)} \pi[\![\varphi \mathrm{U} \psi]\!]_p,$$

therefore, the s-t-paths over Q should exactly represent the proofs of $\mathrm{E}(\varphi \mathrm{U} \psi)$ at v. We set $Q = V \cup \{s, t\}$ and w.l.o.g., $s, t \notin V$. Since the proof of $\mathrm{E}(\varphi \mathrm{U} \psi)$ should start at v, we connect the starting state s with v, so we set

$$C(s, v) = 1 \quad \text{and} \quad C(s, w) = 0 \quad \text{for all } w \in Q \setminus \{v\}.$$

Notice that the edge (s, v) effectively has no cost, since 1 is the neutral element of multiplication. Once we arrive from the starting state to v, we have two options. We can show that ψ holds at v and move to the final state t, since in that case, $\mathrm{E}(\varphi \mathrm{U} \psi)$ is proven at v. The second possibility is to show φ at v and take an edge (v, w) to another node w, and then prove $\mathrm{E}(\varphi \mathrm{U} \psi)$ at w. From w, we have the same two options, that is, showing ψ and moving to the final state or proving φ and taking yet another edge (w, u) to a node u. Therefore, we set

$$\begin{aligned} C(w, t) &= \pi[\![\psi]\!]_w && \text{for all } w \in V \text{ and} \\ C(w, u) &= \pi[\![\varphi]\!]_w \cdot \pi(Ewu) && \text{for all } w, u \in V. \end{aligned}$$

It is left to prove that these intuitive definitions of $\mathcal{A}_v(\pi)$ yield the desired result.

(4.16) Proposition. Let K be an ω-continuous semiring, V a finite set of nodes, π a K-interpretation over V and $v \in V$. The K-automaton $\mathcal{A}_v^{\mathrm{E}(\varphi \mathrm{U} \psi)}(\pi) = (Q, C, s, t)$ with $s, t \notin V$ is defined by setting $Q = V \cup \{s, t\}$ and defining C by

$$\begin{aligned} C(q, s) &= 0 && \text{for all } q \in Q, \\ C(t, q) &= 0 && \text{for all } q \in Q, \\ C(s, v) &= 1, \\ C(s, q) &= 0 && \text{for all } q \in Q \setminus \{v\}, \\ C(w, t) &= \pi[\![\psi]\!]_w && \text{for all } w \in V \quad \text{and} \\ C(w, u) &= \pi[\![\varphi]\!]_w \cdot \pi(Ewu) && \text{for all } w, u \in V. \end{aligned}$$

Then, $\mathcal{A}_v^{\mathrm{E}(\varphi \mathrm{U} \psi)}(\pi)$ has the cost

$$C(\mathcal{A}_v^{\mathrm{E}(\varphi \mathrm{U} \psi)}(\pi)) = \sum_{p\in P_v^{\text{fin}}(V)} \pi[\![\varphi \mathrm{U} \psi]\!]_p.$$

Proof. By the definition of the cost of $\mathcal{A}_v^{\mathrm{E}(\varphi \mathrm{U} \psi)}(\pi)$, we have to show that

$$(1) \qquad \sum_{p' \in P_{s \to t}(Q)} C(p') = \sum_{p \in P_v^{\mathrm{fin}}(V)} \pi[\![\varphi \mathrm{U} \psi]\!]_p.$$

It suffices to consider $N_{s \to t}(Q)$, the set of s-t-paths over Q with nonzero cost and the set of finite paths starting at v with nonzero $\varphi \mathrm{U} \psi$-costs $N_v^{\mathrm{fin}}(V)$. We first observe that there is a bijection $f : N_{s \to t}(Q) \to N_v^{\mathrm{fin}}(V)$ which is defined as follows. Consider a path $p' \in N_{s \to t}(Q)$. The states s and t occur exactly at the start and the end of the path, since p' is an s-t-path with nonzero cost. Therefore, there is some $n \in \mathbb{N}$ with $p' = (s, v_0, ..., v_n, t)$ and $v_0, ..., v_n \in V$. Also, $v_0 = v$, since p' would otherwise have cost 0. We set

$$f(p') = f((s, v_0, ..., v_n, t)) = (v_0, ..., v_n) = p.$$

Clearly, we have $p \in P_v^{\mathrm{fin}}(V)$. The function f also preserves costs, that is, we have $C(p') = \pi[\![\varphi \mathrm{U} \psi]\!]_p$. We verify this by looking at the definition of C and finding that (s, v_0) has cost 1, so that $C((s, v_0, ..., v_n, t)) = C((v_0, ..., v_n, t))$. Moreover, (v_n, t) has the cost $\pi[\![\psi]\!]_{v_n}$ and the edges in $(v_0, ..., v_n)$ cover $\pi[\![\varphi]\!]_{v_i}$ and $\pi(E v_i v_{i+1})$ for $0 \leq i < n$, so the cost of p' is exactly $\pi[\![\varphi \mathrm{U} \psi]\!]_{(v_0, ..., v_n)} = \pi[\![\varphi \mathrm{U} \psi]\!]_p$.

This proves that f is well-defined, since we now know that p is in $N_v^{\mathrm{fin}}(V)$ for $p' \in N_{s \to t}(Q)$. Also, f can be easily inverted. Let $p'' \in N_v^{\mathrm{fin}}(V)$, then we obtain (s, p'', t) by appending s to the first node of p'' and t to the last node. This is also cost-preserving, so (s, p'', t) is in $N_{s \to t}(Q)$ and f is indeed bijective. Also, this ends the proof, since we can see that the two sums in (1) have exactly the same nonzero summands and are therefore equal. □

To provide an example for this transformation, we will transform the $\mathbb{N}^\infty[\![X]\!]$-interpretation π from example (3.13) to a K-automaton. We want to evaluate $\mathrm{E}(\mathrm{F}P) = \mathrm{E}(1 \mathrm{U} P)$ at v. Figure (4.17) shows how $\mathcal{A}_v^{\mathrm{E}(1 \mathrm{U} P)}(\pi)$ is built from π. Transitions with cost 0 are omitted.

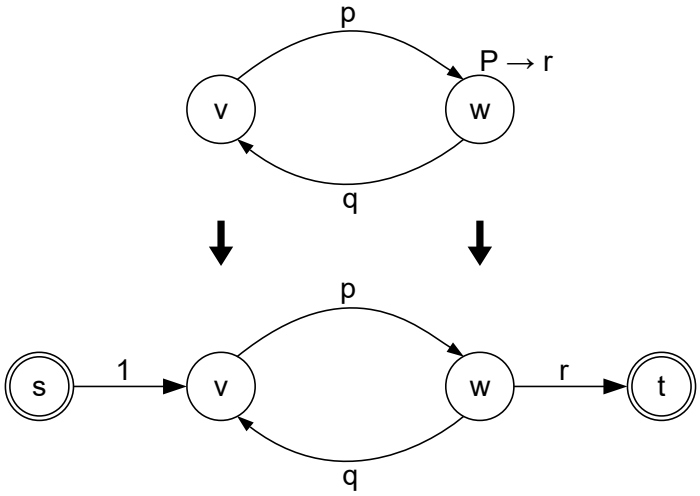

Figure (4.17): $\mathbb{N}^\infty[\![X]\!]$-interpretation π (above) and $\mathcal{A}_v^{\mathrm{E}(1 \mathrm{U} P)}(\pi)$ (below).

The next goal is to calculate $C(\mathcal{A}_v^{\mathrm{E}(1 \mathrm{U} P)}(\pi))$, since we have shown that this is equal to $\pi[\![\mathrm{E}(1 \mathrm{U} P)]\!]_v$. We will now show how to obtain a representation of $C(\mathcal{A})$ for general K-automata $\mathcal{A}$ and then apply this method to the above example.

Our approach is inspired by the well-known *state removal* method from automata theory. The state removal method is used to build a regular expression for a language that is recognized by a given automaton. A description of the state removal method was given, for instance, by Neumann in 2005 [Neu05]. Of course, we will have to adapt this method to work with K-automata, but the idea behind state removal remains the same.

Given a K-automaton $\mathcal{A} = (Q, C, s, t)$, we pick a state $q \in Q \setminus \{s, t\}$ and remove it, so that $Q' = Q \setminus \{q\}$. Then, we adapt the cost function to compensate for the removed state. For $q_1, q_2 \in Q \setminus \{q\}$, this is done by setting

$$C'(q_1, q_2) = C(q_1, q_2) + C(q_1, q) \cdot (C(q, q))^* \cdot C(q, q_2).$$

The intuitive understanding is that before q was removed, q_2 was reachable from q_1 via q, that is, by taking the edge (q_1, q) and then (q, q_2). Of course, (q, q) could be iterated arbitrarily often before taking (q, q_2) and ending up at q_2. However, since q is missing in the new automaton $\mathcal{A}' = (Q', C', s, t)$, the new cost function C' compensates for this by adding $C(q_1, q) \cdot (C(q, q))^* \cdot C(q, q_2)$ to $C(q_1, q_2)$.

We will show that removing a state as above does not change the cost of the automaton, that is

$$C(\mathcal{A}) = C(\mathcal{A}').$$

With this result, we can compute $C(\mathcal{A})$ by repeatedly removing states until only s and t are left. When only s and t are left, it is easy to see the cost of the automaton. We will now formally state the algorithm and and prove its correctness.

(4.18) Algorithm (State Removal)**.** The input for the state removal algorithm is a K-automaton $\mathcal{A} = (Q, C, s, t)$. The algorithm then computes a representation of $C(\mathcal{A})$ using addition, multiplication and the star operator $(*)$ in K.

1. Start with $\mathcal{A}_0 = (Q_0, C_0, s, t) = \mathcal{A}$.
2. Repeat for $i = 0, 1, ...$ as long as $Q_i \neq \{s, t\}$:
   (a) Pick an arbitrary state $q_i \in Q_i \setminus \{s, t\}$.
   (b) Remove q_i by setting $Q_{i+1} = Q_i \setminus \{q_i\}$.
   (c) For all $q_{\text{src}}, q_{\text{dest}} \in Q_{i+1}$, set
   $$C_{i+1}(q_{\text{src}}, q_{\text{dest}}) = C_i(q_{\text{src}}, q_{\text{dest}}) + C_i(q_{\text{src}}, q_i) \cdot (C_i(q_i, q_i))^* \cdot C_i(q_i, q_{\text{dest}}).$$
   (d) This yields $\mathcal{A}_{i+1} = (Q_{i+1}, C_{i+1}, s, t)$.
3. The loop terminates at $\mathcal{A}_n = (Q_n, C_n, s, t)$ with $Q_n = \{s, t\}$ for some $n \in \mathbb{N}$.
4. The output is $C_n(s, t)$.

Proof. We first verify that the algorithm actually terminates and that all the operations are well-defined. Let n be $|Q \setminus \{s, t\}|$, the number of "normal" states in $\mathcal{A}$. Clearly, after n iterations, the loop (2) terminates and only s and t are left as states in $\mathcal{A}_n$. Also, since the runtime of step (c) in the inner loop is in $\mathcal{O}(n^2)$, the

runtime of the entire algorithm is in $\mathcal{O}(n^3)$, assuming that the operations $+$, $\cdot$ and $*$ in the semiring are performed in constant time.

Also, all incoming edges to s and all outgoing edges from t are labelled with 0 in each iteration. This is easily verified by observing that step (c) of the algorithm preserves this property. Therefore, the only edge labelled with a nonzero value in $\mathcal{A}_n$ is (s,t) and the only s-t-path with nonzero cost in $\mathcal{A}_n$ is (s,t). We have

$$C_n(s,t) = \sum_{p \in P_{s\to t}(Q_n)} C_n(p) = C_n(\mathcal{A}_n).$$

The final step of the proof is to verify that this is indeed the cost of $\mathcal{A}$, so it is left to show

$$C(\mathcal{A}) = C_0(\mathcal{A}_0) = C_n(\mathcal{A}_n).$$

This can be shown inductively by proving for $i \in \mathbb{N}$ that

$$C_i(\mathcal{A}_i) = C_{i+1}(\mathcal{A}_{i+1}).$$

This is equivalent to

$$\sum_{p_{i+1} \in P_{s\to t}(Q_{i+1})} C_{i+1}(p_{i+1}) = \sum_{p_i \in P_{s\to t}(Q_i)} C_i(p_i).$$

We define the reduction function $\mathrm{R} : P_{s\to t}(Q_i) \to P_{s\to t}(Q_{i+1})$. The reduction function removes all occurrences of q_i from paths in $P_{s\to t}(Q_i)$, so for $p_i \in P_{s\to t}(Q_i)$, $\mathrm{R}(p_i) = p_{i+1}$ is built from p_i by removing all occurrences of q_i. Clearly, we have $p_{i+1} \in P_{s\to t}(Q_{i+1})$. Since by proposition (2.5), summation is invariant under partitioning and the fibers of R partition $P_{s\to t}(Q_i)$, we have

$$\sum_{p_i \in P_{s\to t}(Q_i)} C_i(p_i) = \sum_{p_{i+1} \in P_{s\to t}(Q_{i+1})} \left(\sum_{p_i \in \mathrm{R}^{-1}(\{p_{i+1}\})} C_i(p_i). \right)$$

It remains to show for each $p_{i+1} \in P_{s\to t}(Q_{i+1})$ that

$$C_{i+1}(p_{i+1}) = \sum_{p_i \in \mathrm{R}^{-1}(\{p_{i+1}\})} C_i(p_i).$$

Intuitively speaking, this means that the cost of p_{i+1} under C_{i+1} exactly covers the C_i-costs of all paths mapped to p_{i+1} by R. In order to show this, we will first prove the following statement by induction. Let $(a_0, ..., a_k)$ be a path in $P_{s\to t}(Q_{i+1})$ of length $k \geq 1$. Then, we have

$$C_{i+1}((a_0, ..., a_k)) = \sum_{(j_1,...,j_k) \in \omega^k} C_i((a_0, q_i^{j_1}, a_1, ..., a_{k-1}, q_i^{j_k}, a_k)).$$

In the above equation, q_i^j for $j \in \omega$ means that the state q_i is traversed j times at the corresponding location in the path. To start the induction, consider $k = 1$ and

the path (a_0, a_1). The hypothesis is true because of

$$\begin{aligned} C_{i+1}((a_0,a_1)) &= C_{i+1}(a_0,a_1) \\ &= C_i(a_0,a_1) + C_i(a_0,q_i)\cdot(C_i(q_i,q_i))^*\cdot C_i(q_i,a_1) \\ &= C_i(a_0,a_1) + \sum_{l\in\omega} C_i(a_0,q_i)\cdot(C_i(q_i,q_i))^l\cdot C_i(q_i,a_1) \\ &= C_i((a_0,a_1)) + \sum_{l\in\omega\setminus\{0\}} C_i((a_0,q_i^l,a_1)) \\ &= \sum_{(j_1)\in\omega^1} C_i((a_0,q_i^{j_1},a_1)). \end{aligned}$$

Now, we assume the hypothesis to be true for k. For $k+1$, we obtain

$$\begin{aligned} & C_{i+1}((a_0,...,a_k,a_{k+1})) \\ = \; & C_{i+1}((a_0,...,a_k))\cdot C_{i+1}((a_k,a_{k+1})) \\ \overset{(1)}{=} \; & \left(\sum_{(j_1,...,j_k)\in\omega^k} C_i((a_0,q_i^{j_1},...,q_i^{j_k},a_k))\right)\cdot\left(\sum_{j_{k+1}\in\omega} C_i((a_k,q_i^{j_{k+1}},a_{k+1}))\right) \\ \overset{(2)}{=} \; & \sum_{(j_1,...,j_k)\in\omega^k}\left(\sum_{j_{k+1}\in\omega} C_i((a_0,q_i^{j_1},...,q_i^{j_k},a_k))\cdot C_i((a_k,q_i^{j_{k+1}},a_{k+1}))\right) \\ \overset{(3)}{=} \; & \sum_{(j_1,...,j_{k+1})\in\omega^{k+1}} C_i((a_0,q_i^{j_1},...,q_i^{j_k},a_k,q_i^{j_{k+1}},a_{k+1})). \end{aligned}$$

For (1), we used the induction hypothesis for the first part, and for $C_{i+1}((a_k, a_{k+1}))$, the same argument as for the case $k = 1$ applies. (2) is shown by applying distributivity of multiplication over countable sums twice, as stated in proposition (2.5). The transformation (3) uses the invariance of summation under partitioning from the same proposition and ends the induction.

Now, we return to $C_{i+1}(p_{i+1})$ for $p_{i+1} \in P_{s\to t}(Q_{i+1})$. We can write

$$p_{i+1} = (s, a_1, ..., a_{k-1}, t)$$

for some $k \geq 1$. Clearly, the statement that we have proved above yields

$$C_{i+1}(p_{i+1}) = C_{i+1}((s,a_1,...,a_{k-1},t)) = \sum_{(j_1,...,j_k)\in\omega^k} C_i((s,q_i^{j_1},a_1,...,a_{k-1},q_i^{j_k},t)).$$

Now, consider the set $\mathrm{R}^{-1}(\{p_{i+1}\})$. Remember that R removes all occurrences of q_i from paths $p_i \in P_{s\to t}(Q_i)$, so the paths that R maps to p_{i+1} are exactly the paths in the set

$$\mathrm{R}^{-1}(\{p_{i+1}\}) = \{(s,q_i^{j_1},a_1,...,a_{k-1},q_i^{j_k},t) \mid (j_1,...,j_k)\in\omega^k\} \subseteq P_{s\to t}(Q_i).$$

Therefore, for any $p_{i+1} \in P_{s\to t}(Q_{i+1})$, we have some k with

$$C_{i+1}(p_{i+1}) = \sum_{(j_1,...,j_k)\in\omega^k} C_i((s,q_i^{j_1},a_1,...,a_{k-1},q_i^{j_k},t)) = \sum_{p_i\in\mathrm{R}^{-1}(\{p_{i+1}\})} C_i(p_i),$$

which ends the proof. □

We can now use state removal to evaluate $\mathrm{E}(\varphi \mathrm{U} \psi)$ under a K-interpretation π over V at some node $v \in V$. We just have to put the above results together.

1. Build $\mathcal{A}_v^{\mathrm{E}(\varphi \mathrm{U} \psi)}(\pi)$ as in proposition (4.16).
   - Together with theorem (4.14), we have
   $$C(\mathcal{A}_v^{\mathrm{E}(\varphi \mathrm{U} \psi)}(\pi)) = \sum_{p \in P_v^{\mathrm{fin}}(V)} \pi[\![\varphi \mathrm{U} \psi]\!]_p = \pi[\![\mathrm{E}(\varphi \mathrm{U} \psi)]\!]_v.$$
2. Use state removal to obtain $C(\mathcal{A}_v^{\mathrm{E}(\varphi \mathrm{U} \psi)}(\pi))$.
   - This is exactly the interpretation $\pi[\![\mathrm{E}(\varphi \mathrm{U} \psi)]\!]_v$.

Let $n = |V|$, then $\mathcal{A}_v^{\mathrm{E}(\varphi \mathrm{U} \psi)}(\pi)$ has $n+2$ states, so the runtime of step (1) is in $\mathcal{O}(n^2)$ and the runtime of state removal is in $\mathcal{O}(n^3)$, which yields a total runtime in $\mathcal{O}(n^3)$.

We close this subsection by applying state removal to $\mathcal{A}_v^{\mathrm{E}(1\mathrm{U}P)}(\pi)$, which we built earlier in figure (4.17). The two steps of state removal are shown in figure (4.19) below.

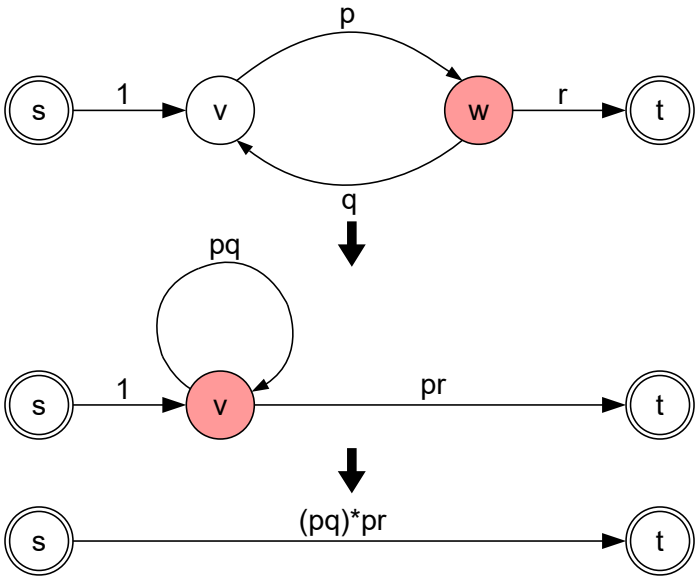

Figure (4.19): State removal performed on $\mathcal{A}_v^{\mathrm{E}(1\mathrm{U}P)}(\pi)$.

The result of the state removal algorithm is $(pq)^*pr = p(pq)^*r$, which is the label of (s,t) in the last step. So, we have

$$C(\mathcal{A}_v^{\mathrm{E}(1\mathrm{U}P)}(\pi)) = \pi[\![\mathrm{E}(1\mathrm{U}P)]\!]_v = \pi[\![\mathrm{E}(\mathrm{F}P)]\!]_v = p(pq)^*r.$$

This is the same result that we obtained in example (3.13) by manually performing the fixed-point iteration. We conclude that the state removal algorithm is a systematic approach to interpret existential until operators in CTL.

4.2.2 Universal Until Operators

Consider an ω-continuous semiring K, a formula $\mathrm{A}(\varphi \mathrm{U} \psi)$ in CTL and a matching K-interpretation π over a finite set of nodes V. As seen in theorem (4.14), we have

$$\pi[\![\mathrm{A}(\varphi \mathrm{U} \psi)]\!]_v = \sum_{t \in T_v^{\mathrm{fin}}(V)} \pi[\![\varphi \mathrm{U} \psi]\!]_t \quad \text{for all } v \in V.$$

Same as for the existential formulas, we will use an approach from automata theory to calculate this. However, since universal formulas are witnessed by trees instead of paths, we will use methods for tree automata. In a book by Comon, Dauchet, Gilleron, Jacquemard, Lugiez, Löding, Tison and Tommasi [CDG+07], they describe a method that converts tree automata to regular tree expressions. This subsection presents an adaptation of their method to obtain an expression for $\pi[\![\mathrm{A}(\varphi \mathrm{U} \psi)]\!]_v$.

First of all, recall the $\varphi \mathrm{U} \psi$-costs for trees that we introduced in definition (4.11). We would like to be able to split trees as shown in figure (4.20) below. We have marked the node in t_1 where t_2 should be appended with a star ($*$). Edge costs are disregarded and nodes are labelled with their contributions to the $\varphi \mathrm{U} \psi$-cost of their respective tree.

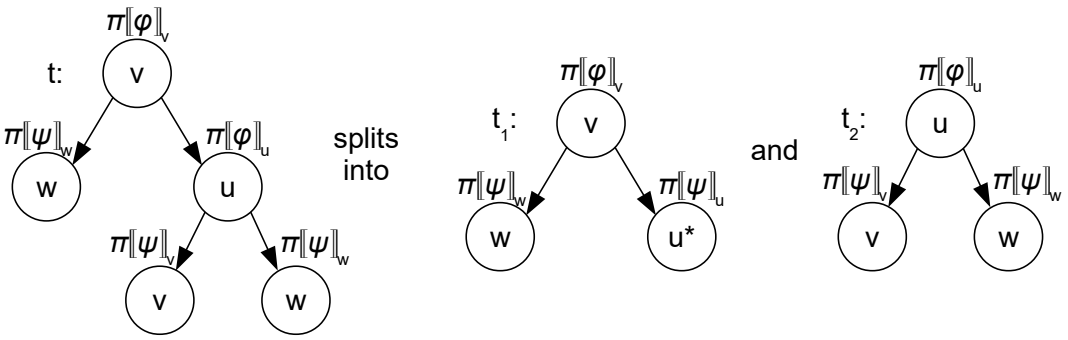

Figure (4.20): A tree t being split into two trees t_1 and t_2.

Obviously, splitting trees like that does not play along with our definition of $\varphi \mathrm{U} \psi$-costs. The problem is the marked u-leaf in t_1. For $\pi[\![\varphi \mathrm{U} \psi]\!]_t$, we would have obtained $\pi[\![\varphi]\!]_u$ as the cost of the u-node. However, in t_1 and t_2, the u-node appears twice. In t_1, it is a leaf, therefore we get the cost $\pi[\![\psi]\!]_u$ and in t_2, it is the root and we get the correct cost $\pi[\![\varphi]\!]_u$. Because of the factor $\pi[\![\psi]\!]_u$ in t_1, we generally have

$$\pi[\![\varphi \mathrm{U} \psi]\!]_t \neq \pi[\![\varphi \mathrm{U} \psi]\!]_{t_1} \cdot \pi[\![\varphi \mathrm{U} \psi]\!]_{t_2},$$

but we would naturally expect both sides to be equal, since t_1 and t_2 form t when t_2 is inserted into t_1 at the marked node.

Overcoming this issue requires a new definition of tree costs where marked leaves are treated differently. Instead of evaluating u in t_1 to $\pi[\![\psi]\!]_u$, we evaluate it to a variable $x_u \in X$ where $X = \{x_v \mid v \in V\}$ is a set of variables. When appending t_2 to t_1, we simply insert the cost of t_2 into the cost of t_1, which is a polynomial. To see that this works, we disregard the edge costs in the above example and calculate

$$\pi[\![\varphi \mathrm{U} \psi]\!]_t = \pi[\![\varphi]\!]_v \cdot \pi[\![\psi]\!]_w \cdot \pi[\![\varphi]\!]_u \cdot \pi[\![\psi]\!]_v \cdot \pi[\![\psi]\!]_w.$$

Then, we calculate the costs of t_1 and t_2 with the new evaluation strategy and obtain

$$\pi[\![\varphi \mathrm{U} \psi]\!]_{t_1} = \pi[\![\varphi]\!]_v \cdot \pi[\![\psi]\!]_w \cdot x_u \quad \text{and}$$
$$\pi[\![\varphi \mathrm{U} \psi]\!]_{t_2} = \pi[\![\varphi]\!]_u \cdot \pi[\![\psi]\!]_v \cdot \pi[\![\psi]\!]_w.$$

Clearly, inserting the cost of t_2 into x_u in the costs of t_1 yields exactly the cost of t. Now, we will provide a formal definition for the polynomial costs of trees with marked leaves.

(4.21) Definition (Complete Tree with Marked Leaves)**.** A complete tree with marked leaves over V is a tree $t \in T(V)$ with a subset $m(t) \subseteq l(t)$ of marked leaves.

We will sometimes restrict the nodes that may be marked to a subset $M \subseteq V$. That means that only leaves which are labelled with an element from M may be marked. We write $T(V, M)$ for the set of all complete trees over V with marked leaves in M. In the above example, only u-leaves are marked in t_1, so $t_1 \in T(V, \{u\})$. Since t and t_2 do not contain marked leaves, we have $t, t_2 \in T(V, \emptyset)$. Notice that for trees in $T(V, M)$, some or all leaves in M may be unmarked, so we have $T(V, M) \subseteq T(V, V)$ for all $M \subseteq V$. Now, we will specify how to handle marked leaves when calculating the $\varphi \mathrm{U} \psi$-costs of trees.

(4.22) Definition (Until-Costs for Trees with Marked Leaves)**.** Let K be ω-continuous and V a finite set of nodes. Let π be a K-interpretation and t a finite, complete tree over V with marked leaves. We set $X = \{x_v \mid v \in V\}$ and for φ and ψ in CTL, the $\varphi \mathrm{U} \psi$-cost of t is

$$\pi[\![\varphi \mathrm{U} \psi]\!]_t = \left(\prod_{x \in i(t)} \pi[\![\varphi]\!]_{L_t(x)} \right) \cdot \left(\prod_{(x,y) \in e(t)} \pi(EL_t(x)L_t(y)) \right) \cdot \left(\prod_{x \in l(t) \setminus m(t)} \pi[\![\psi]\!]_{L_t(x)} \right)$$
$$\cdot \left(\prod_{y \in m(t)} x_{L_t(y)} \right) \in K[\![X]\!].$$

Intuitively speaking, the marked leaves are ignored for the normal evaluation and a marked leaf v gets evaluated as x_v. The result is formally an element of $K[\![X]\!]$, because there are variables in X and coefficients in K. We use formal power series instead of normal polynomials, because generally, unlike $K[X]$, $K[\![X]\!]$ is again ω-continuous.

Next, we will define *insertion* for formal power series $K[\![X]\!]$ with $X = \{x_v \mid v \in V\}$. Let p and q be formal power series $p \in K[\![X_1]\!]$ and $q \in K[\![X_2]\!]$ with $X_1, X_2 \subseteq X$. For $v \in V$, we define the insertion $p \cdot_v q$. Informally, we insert q into p for each occurrence of the variable x_v.

The formal definition is based on an observation by Green, Karvounarakis and Tannen [GKT07]. Notice that q induces a unique ω-continuous homomorphism $h_q^v : K[\![X_1]\!] \to K[\![(X_1 \setminus \{x_v\}) \cup X_2]\!]$ with $h_q^v(x_v) = q$ and $h_q^v(r) = r$ if x_v does not occur in r. So, we can set

$$p \cdot_v q = h_q^v(p).$$

The fact that h_q^v is uniquely defined can be shown by observing that p is an element of $K[\![X_1]\!]$, but we can see it as an element of $K[\![X_1 \setminus \{x_v\}]\!][\![x_v]\!]$. So, we can write p

as

$$p = \sum_{i \in \omega} p(i) \cdot x_v^i$$

where $p(i) \in K[\![X_1 \setminus \{x_v\}]\!]$ is the coefficient of x_v^i. Clearly, as h_q^v is an ω-continuous homomorphism and $h_q^v(p(i)) = p(i)$, we have

$$h_q^v(p) = h_q^v\left(\sum_{i \in \omega} p(i) \cdot x_v^i\right) = \sum_{i \in \omega} p(i) \cdot h_q^v(x_v)^i = \sum_{i \in \omega} p(i) \cdot q^i,$$

which is a uniquely defined element of $K[\![(X_1 \setminus \{x_v\}) \cup X_2]\!]$.

Now that we have defined polynomial insertion, we can prove some properties.

(4.23) Lemma. Let K be an ω-continuous semiring, V a finite set of nodes and $X = \{x_v \mid v \in V\}$. For any $v \in V$, formal power series $p, q, r \in K[\![X]\!]$ and ascending chains $s_0 \leq s_1 \leq ...$ in $K[\![X]\!]$, we have

$$\begin{aligned}
&(1) \quad (p \cdot_v q) \cdot_v r = p \cdot_v (q \cdot_v r) && (\cdot_v \text{ is associative}),\\
&(2) \quad (p + q) \cdot_v r = p \cdot_v r + q \cdot_v r,\\
&(3) \quad (p \cdot q) \cdot_v r = (p \cdot_v r) \cdot (q \cdot_v r),\\
&(4) \quad (\sup_{i \in \omega} s_i) \cdot_v p = \sup_{i \in \omega}(s_i \cdot_v p) && \text{and}\\
&(5) \quad p \cdot_v (\sup_{i \in \omega} s_i) = \sup_{i \in \omega}(p \cdot_v s_i).
\end{aligned}$$

Proof. For (1), we rewrite the equation using the definition of $\cdot_v$ to

$$(h_r^v \circ h_q^v)(p) = h^v_{h_r^v(q)}(p).$$

Since h_r^v and h_q^v are both ω-continuous homomorphisms, $h_r^v \circ h_q^v$ is an ω-continuous homomorphism as well. We have $(h_r^v \circ h_q^v)(t) = t$ for any polynomial that does not contain x_v and

$$(h_r^v \circ h_q^v)(x_v) = h_r^v(h_q^v(x_v)) = h_r^v(q).$$

As $h^v_{h_r^v(q)}$ is the unique ω-continuous homomorphism that fulfills these properties, we conclude that $h_r^v \circ h_q^v = h_{h_r^v(q)}$ and (1) is proven.

For (2) and (3), we simply use the definition and obtain

$$\begin{aligned}
(p + q) \cdot_v r = h_r^v(p + q) &= h_r^v(p) + h_r^v(q) = p \cdot_v r + q \cdot_v r \quad \text{and}\\
(p \cdot q) \cdot_v r = h_r^v(p \cdot q) \ &= h_r^v(p) \cdot h_r^v(q) \ = (p \cdot_v r) \cdot (q \cdot_v r),
\end{aligned}$$

since h_r^v is a homomorphism.

(4) uses the fact that h_p^v is ω-continuous, so

$$(\sup_{i \in \omega} s_i) \cdot_v p = h_p^v((\sup_{i \in \omega} s_i)) = \sup_{i \in \omega} h_p^v(s_i) = \sup_{i \in \omega}(s_i \cdot_v p).$$

The final statement (5) states that the polynomial function f_p^v induced by p is ω-continuous. The polynomial function f_p^v is defined by $f_p^v(q) = h_q^v(p)$ for $q \in K[\![X]\!]$. Grädel and Tannen have stated that this function is ω-continuous [GT18]. Therefore,

we will only provide a proof sketch for (5) by induction on p. If p does not contain x_v, then $p \cdot_v (\sup_{i\in\omega} s_i) = p$ and $p \cdot_v s_i = p$ for each $i \in \omega$, so there is nothing to prove. For $p = x_v$, we have $p \cdot_v t = t$ for any power series t, so clearly, both sides are equal. Knowing that addition and multiplication are ω-continuous and using the statements (2), (3) and (4), we can inductively conclude that (5) is true for any formal power series $p \in K\langle\langle X\rangle\rangle$. □

Before getting back to the trees, we will define another operation on formal power series. The idea is that a tree t that is rooted at w and has a marked leaf w at the same time could be appended to itself arbitrarily often to build arbitrarily large trees. Therefore, we will need *iteration* for formal power series. Let $p \in K\langle\langle X\rangle\rangle$ be a formal power series and $v \in V$. For $i \in \omega$, we inductively define

$$\begin{aligned} p^{0,v} &= x_v \quad \text{and} \\ p^{i+1,v} &= x_v + p \cdot_v p^{i,v}. \end{aligned}$$

Informally, $p^{i,v}$ represents the sum of all the polynomials built by inserting p into itself for the variable x_v up to a depth of i. For example, if $p = ax_v$ for some $a \in K$, then we have

$$\begin{aligned} p^{0,v} &= x_v, \\ p^{1,v} &= x_v + ax_v, \\ p^{2,v} &= x_v + ax_v + a^2x_v, \dots . \end{aligned}$$

Since $K\langle\langle X\rangle\rangle$ is ω-continuous, we can define

$$p^{*_v} = \sup_{i\in\omega} p^{i,v}.$$

Notice that $p^{0,v}, p^{1,v}, \dots$ is an ascending chain, which is verified by showing that $p^{i,v} \leq p^{i+1,v}$ inductively. For $i = 0$, we have $p^{0,v} = x_v$ and $p^{1,v} = x_v + p \cdot_v p^{0,v}$, so $p^{0,v} \leq p^{1,v}$ is true. For $i > 0$, we have $p^{i+1,v} = x_v + p \cdot_v p^{i,v}$. Since addition and $\cdot_v$ are ω-continuous, they are also monotonic and therefore, with the induction hypothesis, we have

$$p^{i+1,v} = x_v + p \cdot_v p^{i,v} \geq x_v + p \cdot_v p^{i-1,v} = p^{i,v}.$$

This shows that p^{*_v} is well-defined. In the above example, we would obtain

$$(ax_v)^{*_v} = x_v + ax_v + a^2x_v + \dots = a^*x_v.$$

As we can see, the insertion depth is unlimited here. Note that the objects described by the iteration operator $*_v$ do not always have a simple representation as in the example above. Consider a non-linear power series ax_v^2, then we have

$$(ax_v^2)^{*_v} = x_v + ax_v^2 + 2a^2x_v^3 + 5a^3x_v^4 + \dots .$$

Looking at the monomial $2a^2x_v^3$ explains the issue. We can obtain $a^2x_v^3$ by inserting ax_v^2 into itself for *one* of the variables x_v. However, since $x_v^2 = x_v \cdot x_v$, there are two options to do this since x_v technically occurs twice in ax_v^2. Hence, the monomial $2a^2x_v^3$ has the coefficient 2.

Even though expressions with $*_v$ are harder to read than the usual star expressions from the previous subsection, we will see that they are very useful for describing the $\varphi \mathrm{U} \psi$-costs of infinite sets of trees.

Fix an ω-continuous semiring K, a K-interpretation π over a finite set of nodes V and a formula $\mathrm{A}(\varphi \mathrm{U} \psi)$ in CTL. We can assign an arbitrary order to the nodes in V, so w.l.o.g., we assume that $V = \{1, ..., n\}$ for some $n \in \mathbb{N}$. By theorem (4.14), we know that for $v \in V$, we have

$$\pi[\![\mathrm{A}(\varphi \mathrm{U} \psi)]\!]_v = \sum_{t \in T_v^{\mathrm{fin}}(V)} \pi[\![\varphi \mathrm{U} \psi]\!]_t.$$

We can compute a representation of $\pi[\![\mathrm{A}(\varphi \mathrm{U} \psi)]\!]_v$ using $+$, $\cdot$, $\cdot_w$ and $*_w$ for $w \in V$ inductively. We modify the approach that is used in [CDG+07] to convert tree automata to regular tree expressions. For any $1 \leq i \leq n$, $0 \leq j \leq n$ and $L \subseteq V$, we define $T(i, j, L)$ as the set of all finite, complete trees t over V with the following properties:

- t is rooted at i,
- t may contain marked leaves, but only leaves that are labelled with an element of L can be marked,
- t is not trivial, meaning that the root of t must not be a marked leaf and
- all unmarked nodes in t except for the root node must be labelled with an element of $\{1, ..., j\}$.

Parallelly, we define the costs

$$C(i, j, L) = \sum_{t \in T(i,j,L)} \pi[\![\varphi \mathrm{U} \psi]\!]_t.$$

Clearly, we have $T(v, n, \emptyset) = T_v^{\mathrm{fin}}(V)$, so we are looking for $C(v, n, \emptyset)$. In the following, we will show how all $C(i, j, L)$ can be computed inductively over j.

For $j = 0$, $C(i, 0, L)$ is easy to compute, since $T(i, 0, L)$ only contains trees without any unmarked nodes except for the root. There are no more than two such trees, one of them is the tree that only consists of the root i and the second one is the tree that consists of i with all of its successors and all of them are marked.

Moving on to $j > 0$, we assume that we have already computed $C(i, j', L)$ for any $j' < j$. We claim that

$$C(i, j, L) = \begin{cases} C(i, j-1, L \cup \{j\}) \cdot_j C(j, j-1, L \cup \{j\})^{*_j} & \text{if } j \in L \\ C(i, j-1, L \cup \{j\}) \cdot_j C(j, j-1, L \cup \{j\})^{*_j} \cdot_j 0 & \text{otherwise.} \end{cases}$$

Figure (4.24) provides an intuitive justification for this equation.

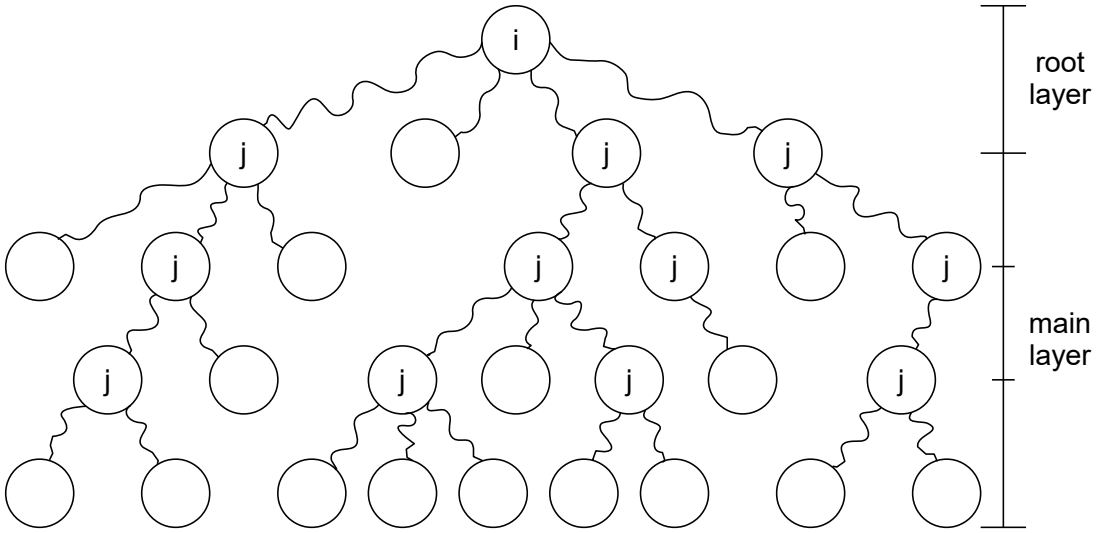

Figure (4.24): A t tree in $T(i, j, L)$ split in two layers.

The figure shows an informal representation of a tree t in $T(i, j, L)$. As seen in the picture, we can split the tree along any occurrence of the node j. Clearly, the subtrees will not contain any instances of j except for root nodes and marked leaves. The root layer, which is the subtree that contains the root, will be a tree rooted at i with unmarked nodes in $\{1, ..., j-1\}$ and marked leaves in $L \cup \{j\}$, since we have split it at j. Therefore, it is an element of $T(i, j-1, L \cup \{j\})$. The main layer consists of trees that are rooted at j and contain marked leaves in $L \cup \{j\}$, so they are elements of $T(j, j-1, L \cup \{j\})$. Clearly, these trees can be chained arbitrarily often, therefore, we use the iteration operator $*_j$. Notice that if $j \notin L$, there must not be any marked j-leaves at the bottom of t, therefore we use $\cdot_j 0$ to eliminate any trees with marked j-leaves that were generated by the iteration operator.

Now, it is left to prove the above equation formally. We will first introduce the concept of the j-height of a tree. For $j \in V$, the j-height of a tree t over V is defined as the maximum number of occurrences of j along a path from the root, not including marked leaves. For example, the j-height of the tree from figure (4.24) is 3, assuming that $i \neq j$. Additionally, we will need the following lemma.

(4.25) Lemma. Let $t \in T(V, V)$ be a finite, complete tree over V with marked leaves and $T \subseteq T_j(V, V)$ a countable set of trees rooted at j. Then, with π, φ and ψ given as above, we have

$$\pi[\![\varphi \mathrm{U} \psi]\!]_t \cdot_j \sum_{t' \in T} \pi[\![\varphi \mathrm{U} \psi]\!]_{t'} = \sum_{t'' \in T'} \pi[\![\varphi \mathrm{U} \psi]\!]_{t''}$$

where T' is the set of trees that are obtained by inserting a combination of trees from T into all marked j-leaves of t.

Proof. Since t is finite, there is a $k \in \mathbb{N}$ such that t contains exactly k marked

j-leaves. According to definition (4.22), we have

$$\pi[\![\varphi \mathrm{U} \psi]\!]_t = \left(\prod_{x \in i(t)} \pi[\![\varphi]\!]_{L_t(x)} \right) \cdot \left(\prod_{(x,y) \in e(t)} \pi(EL_t(x)L_t(y)) \right) \cdot \left(\prod_{x \in l(t) \setminus m(t)} \pi[\![\psi]\!]_{L_t(x)} \right) \cdot \left(\prod_{y \in m(t), L_t(y) \neq j} x_{L_t(y)} \right) \cdot x_j^k.$$

So, we can write $\pi[\![\varphi \mathrm{U} \psi]\!]_t = p \cdot x_j^k$ for some $p \in K[\![X \setminus \{x_j\}]\!]$. Since p does not contain x_j, it follows that

$$\begin{aligned}
\pi[\![\varphi \mathrm{U} \psi]\!]_t \cdot_j \sum_{t' \in T} \pi[\![\varphi \mathrm{U} \psi]\!]_{t'} &= (p \cdot x_j^k) \cdot_j \sum_{t' \in T} \pi[\![\varphi \mathrm{U} \psi]\!]_{t'} \\
&= p \cdot \left(\sum_{t' \in T} \pi[\![\varphi \mathrm{U} \psi]\!]_{t'} \right)^k \\
&= p \cdot \left(\sum_{t_1 \in T} \pi[\![\varphi \mathrm{U} \psi]\!]_{t_1} \right) \cdot \ldots \cdot \left(\sum_{t_k \in T} \pi[\![\varphi \mathrm{U} \psi]\!]_{t_k} \right) \\
&= p \cdot \left(\sum_{t_k \in T} \ldots \sum_{t_1 \in T} \pi[\![\varphi \mathrm{U} \psi]\!]_{t_1} \cdot \ldots \cdot \pi[\![\varphi \mathrm{U} \psi]\!]_{t_k} \right) \\
&= p \cdot \sum_{(t_1, \ldots, t_k) \in T^k} \pi[\![\varphi \mathrm{U} \psi]\!]_{t_1} \cdot \ldots \cdot \pi[\![\varphi \mathrm{U} \psi]\!]_{t_k} \\
&= \sum_{(t_1, \ldots, t_k) \in T^k} p \cdot \pi[\![\varphi \mathrm{U} \psi]\!]_{t_1} \cdot \ldots \cdot \pi[\![\varphi \mathrm{U} \psi]\!]_{t_k}
\end{aligned}$$

We used the partition-invariance of possibly infinite sums and the distributivity of multiplication over sums from proposition (2.5). Now, to close the proof, it remains to show that

$$\sum_{(t_1, \ldots, t_k) \in T^k} p \cdot \pi[\![\varphi \mathrm{U} \psi]\!]_{t_1} \cdot \ldots \cdot \pi[\![\varphi \mathrm{U} \psi]\!]_{t_k} = \sum_{t'' \in T'} \pi[\![\varphi \mathrm{U} \psi]\!]_{t''}.$$

We observe that there is a bijection between T^k and T'. Fix an arbitrary enumeration of the k marked j-leaves in t. Every tuple $(t_1, \ldots, t_k) \in T^k$ can be mapped to the element t'' of T' that is obtained by inserting t_i for the i-th marked j-leaf in t with $1 \leq i \leq k$. Clearly, we have

$$\pi[\![\varphi \mathrm{U} \psi]\!]_{t''} = p \cdot \pi[\![\varphi \mathrm{U} \psi]\!]_{t_1} \cdot \ldots \cdot \pi[\![\varphi \mathrm{U} \psi]\!]_{t_k}$$

when t'' is constructed from $(t_1, \ldots, t_k)$ like this. To see that this is a one-to-one correspondence, observe that any $t'' \in T'$ is built by taking some $t_1, \ldots, t_k \in T^k$ and inserting them for the marked j-leaves in t. This ends the proof of the lemma, since the two sums above contain exactly the same elements. □

Now, we return to proving the original claim that

$$C(i, j, L) = \begin{cases} C(i, j-1, L \cup \{j\}) \cdot_j C(j, j-1, L \cup \{j\})^{*_j} & \text{if } j \in L \\ C(i, j-1, L \cup \{j\}) \cdot_j C(j, j-1, L \cup \{j\})^{*_j} \cdot_j 0 & \text{otherwise.} \end{cases}$$

We start by evaluating the iteration $C(j, j-1, L \cup \{j\})^{*_j}$. Let $T'(j, j, L \cup \{j\})$ be the set $T(j, j, L \cup \{j\})$ adjoined with the trivial tree $(j*)$ that consists of only one marked leaf j. We claim that

$$C(j, j-1, L \cup \{j\})^{*_j} = \sum_{t \in T'(j,j,L \cup \{j\})} \pi[\![\varphi \mathrm{U} \psi]\!]_t.$$

In order to prove that, we partition $T'(j, j, L \cup \{j\})$ by the j-height of the trees that it contains. Define $T'^{=h}(j, j, L \cup \{j\}) = \{t \in T'(j, j, L \cup \{j\}) \mid t \text{ has the } j\text{-height } h\}$ for $h \in \omega$. $T'^{<h}(j, j, L \cup \{j\})$ and $T'^{\leq h}(j, j, L \cup \{j\})$ for $h \in \omega$ have the obvious meanings. Then, we obtain

$$\begin{aligned}
\sum_{t \in T'(j,j,L \cup \{j\})} \pi[\![\varphi \mathrm{U} \psi]\!]_t &= \sum_{h \in \omega} \left(\sum_{t \in T'^{=h}(j,j,L \cup \{j\})} \pi[\![\varphi \mathrm{U} \psi]\!]_t \right) \\
&= \sup_{h \in \omega} \left(\sum_{k=0}^{h} \sum_{t \in T'^{=k}(j,j,L \cup \{j\})} \pi[\![\varphi \mathrm{U} \psi]\!]_t \right) \\
&= \sup_{h \in \omega} \sum_{t \in T'^{\leq h}(j,j,L \cup \{j\})} \pi[\![\varphi \mathrm{U} \psi]\!]_t.
\end{aligned}$$

On the other side, we have

$$C(j, j-1, L \cup \{j\})^{*_j} = \sup_{h \in \omega} C(j, j-1, L \cup \{j\})^{h,j}.$$

It remains to show by induction on h that

$$C(j, j-1, L \cup \{j\})^{h,j} = \sum_{t \in T'^{\leq h}(j,j,L \cup \{j\})} \pi[\![\varphi \mathrm{U} \psi]\!]_t.$$

In the base case $h = 0$, we have $C(j, j-1, L \cup \{j\})^{0,j} = x_j$ and indeed, the only tree rooted at j with a j-height of 0 or less is the tree $(j*)$ that consists of a single marked leaf j, and therefore has cost x_j.

For $h+1$, we look at the inductive definition of $C(j, j-1, L \cup \{j\})^{h+1,j}$. Using the induction hypothesis for (1), the properties of $\cdot_j$ from lemma (4.23) for (2) and lemma (4.25) for (3), we obtain

$$\begin{aligned}
& C(j, j-1, L \cup \{j\})^{h+1,j} \\
&= x_j + C(j, j-1, L \cup \{j\}) \cdot_j C(j, j-1, L \cup \{j\})^{h,j} \\
&\overset{(1)}{=} x_j + \left(\sum_{t \in T(j,j-1,L \cup \{j\})} \pi[\![\varphi \mathrm{U} \psi]\!]_t \right) \cdot_j \left(\sum_{t' \in T'^{\leq h}(j,j,L \cup \{j\})} \pi[\![\varphi \mathrm{U} \psi]\!]_{t'} \right) \\
&\overset{(2)}{=} x_j + \sum_{t \in T(j,j-1,L \cup \{j\})} \pi[\![\varphi \mathrm{U} \psi]\!]_t \cdot_j \left(\sum_{t' \in T'^{\leq h}(j,j,L \cup \{j\})} \pi[\![\varphi \mathrm{U} \psi]\!]_{t'} \right) \\
&\overset{(3)}{=} x_j + \sum_{t \in T(j,j-1,L \cup \{j\})} \left(\sum_{t' \in T'^{\leq h}_t(j,j,L \cup \{j\})} \pi[\![\varphi \mathrm{U} \psi]\!]_{t'} \right),
\end{aligned}$$

where $T_t'^{\leq h}(j, j, L \cup \{j\})$ is the set of trees that are built by inserting trees from $T'^{\leq h}(j, j, L \cup \{j\})$ into all the marked j-leaves of t. It remains to show that

$$\sum_{t \in T'^{\leq h+1}(j,j,L\cup\{j\})} \pi[\![\varphi \mathrm{U} \psi]\!]_t = x_j + \sum_{t \in T(j,j-1,L\cup\{j\})} \left(\sum_{t' \in T_t'^{\leq h}(j,j,L\cup\{j\})} \pi[\![\varphi \mathrm{U} \psi]\!]_{t'} \right).$$

Let $\mathrm{R}_j : T'^{\leq h+1}(j, j, L \cup \{j\}) \setminus \{(j*)\} \to T(j, j-1, L \cup \{j\})$ be the top-down cutting function. For a given $t \in T'^{\leq h+1}(j, j, L \cup \{j\}) \setminus \{(j*)\}$, $\mathrm{R}_j(t)$ is the tree that we obtain by cutting off t at all the first occurrences of j viewed from the top-down direction, excluding the root itself. The subtrees that are cut off are replaced with marked j-leaves. For example, in figure (4.24), the top-down cutting operation would return the root layer of the depicted tree.

Notice that we have excluded the trivial tree $(j*)$ from the definition of R_j, since it would be mapped to itself, but $T(j, j-1, L \cup \{j\})$ does not contain $(j*)$ by definition. The fibers $\mathrm{R}_j^{-1}(\{t\})$ for $t \in T(j, j-1, L \cup \{j\})$ along with $(j*)$ form a partition of $T'^{\leq h+1}(j, j, L \cup \{j\})$, therefore, we have

$$\sum_{t \in T'^{\leq h+1}(j,j,L\cup\{j\})} \pi[\![\varphi \mathrm{U} \psi]\!]_t = x_j + \sum_{t \in T(j,j-1,L\cup\{j\})} \left(\sum_{t' \in \mathrm{R}_j^{-1}(\{t\})} \pi[\![\varphi \mathrm{U} \psi]\!]_{t'} \right).$$

Since by definition, $\mathrm{R}_j^{-1}(\{t\})$ for $t \in T(j, j-1, L \cup \{j\})$ is the set of all trees in $T'^{\leq h+1}(j, j, L \cup \{j\}) \setminus \{(j*)\}$ that R_j maps to t, we have

$$\mathrm{R}_j^{-1}(\{t\}) = T_t'^{\leq h}(j, j, L \cup \{j\}),$$

or in words, the trees that R_j maps to t are exactly the trees obtained by inserting arbitrary trees from $T'^{\leq h}(j, j, L \cup \{j\})$ into all marked j-leaves of t. We will prove that both sets contain each other.

"$\subseteq$": Let t' be an element of $\mathrm{R}_j^{-1}(\{t\}) \subseteq T'^{\leq h+1}(j, j, L \cup \{j\}) \setminus \{(j*)\}$. Then, $\mathrm{R}_j(t') = t$ and t can be obtained from t' by cutting off some subtrees that are rooted at j. Since t' had a j-height that did not exceed $h+1$ and t still has j as an unmarked root node, the subtrees that were cut off have a j-height that does not exceed h, therefore they are elements of $T'^{\leq h}(j, j, L \cup \{j\})$. This implies that t' can be obtained from t by inserting trees from $T'^{\leq h}(j, j, L \cup \{j\})$ into the j-leaves of t, therefore we have $t' \in T_t'^{\leq h}(j, j, L \cup \{j\})$.

"$\supseteq$": Now, let t' be an element of $T_t'^{\leq h}(j, j, L \cup \{j\})$. Since t' is obtained by inserting trees from $T_t'^{\leq h}(j, j, L \cup \{j\})$ into the marked j-leaves of t, applying R_j to t' will reverse this and yield t, as t itself does not contain any unmarked j-nodes outside of the root. Thereby, we conclude $\mathrm{R}_j(t') = t$ and $t' \in \mathrm{R}_j^{-1}(\{t\})$.

This immediately implies

$$C(j, j-1, L \cup \{j\})^{h+1,j} = \sum_{t \in T'^{\leq h+1}(j,j,L\cup\{j\})} \pi[\![\varphi \mathrm{U} \psi]\!]_t,$$

which ends the induction. It also proves that

$$C(j, j-1, L \cup \{j\})^{*_j} = \sum_{t \in T'(j,j,L\cup\{j\})} \pi[\![\varphi \mathrm{U} \psi]\!]_t.$$

Next, we want to use this to show

$$C(i, j-1, L \cup \{j\}) \cdot_j C(j, j-1, L \cup \{j\})^{*_j} = \sum_{t \in T(i,j,L \cup \{j\})} \pi[\![\varphi \mathrm{U} \psi]\!]_t.$$

The left side can be transformed using lemma (4.23) for (1) and lemma (4.25) for (2), which yields

$$\begin{aligned}
& C(i, j-1, L \cup \{j\}) \cdot_j C(j, j-1, L \cup \{j\})^{*_j} \\
= & \left(\sum_{t \in T(i,j-1,L \cup \{j\})} \pi[\![\varphi \mathrm{U} \psi]\!]_t \right) \cdot_j \left(\sum_{t' \in T'(j,j,L \cup \{j\})} \pi[\![\varphi \mathrm{U} \psi]\!]_{t'} \right) \\
\overset{(1)}{=} & \sum_{t \in T(i,j-1,L \cup \{j\})} \pi[\![\varphi \mathrm{U} \psi]\!]_t \cdot_j \left(\sum_{t' \in T'(j,j,L \cup \{j\})} \pi[\![\varphi \mathrm{U} \psi]\!]_{t'} \right) \\
\overset{(2)}{=} & \sum_{t \in T(i,j-1,L \cup \{j\})} \left(\sum_{t' \in T'_t(j,j,L \cup \{j\})} \pi[\![\varphi \mathrm{U} \psi]\!]_{t'} \right),
\end{aligned}$$

where $T'_t(j, j, L \cup \{j\})$ for $t \in T(i, j-1, L \cup \{j\})$ is the set of trees that can be obtained by inserting an arbitrary combination of trees from $T'(j, j, L \cup \{j\})$ into all the marked j-leaves of t. Now, to show that

$$\sum_{t \in T(i,j,L \cup \{j\})} \pi[\![\varphi \mathrm{U} \psi]\!]_t = \sum_{t \in T(i,j-1,L \cup \{j\})} \left(\sum_{t' \in T'_t(j,j,L \cup \{j\})} \pi[\![\varphi \mathrm{U} \psi]\!]_{t'} \right),$$

we can use the same approach as above. Let $\mathrm{R}'_j : T(i, j, L \cup \{j\}) \to T(i, j-1, L \cup \{j\})$ be the same top-down cutting function as above, defined for a different domain $T(i, j, L \cup \{j\})$. Notice that even if $i = j$, the sets $T(i, j, L \cup \{j\})$ and $T(i, j-1, L \cup \{j\})$ by definition do not contain the trivial tree $(j*)$ and the function R'_j never cuts trees off at the root. Clearly, applying R'_j to a tree in $T(i, j, L \cup \{j\})$ yields a tree in $T(i, j-1, L \cup \{j\})$, even if $i = j$, since the resulting tree can only contain j as a root node or a marked leaf. Therefore, the fibers $\mathrm{R}'^{-1}_j(\{t\})$ for $t \in T(i, j-1, L \cup \{j\})$ are a partition of $T(i, j, L \cup \{j\})$. This yields

$$\sum_{t \in T(i,j,L \cup \{j\})} \pi[\![\varphi \mathrm{U} \psi]\!]_t = \sum_{t \in T(i,j-1,L \cup \{j\})} \left(\sum_{t' \in \mathrm{R}'^{-1}_j(\{t\})} \pi[\![\varphi \mathrm{U} \psi]\!]_{t'} \right).$$

We observe that $\mathrm{R}'^{-1}_j(\{t\}) = T'_t(j, j, L \cup \{j\})$ for $t \in T(i, j-1, L \cup \{j\})$.

"$\subseteq$": If t' is an element of $\mathrm{R}'^{-1}_j(\{t\})$, then we have $t' \in T(i, j, L \cup \{j\})$ and $\mathrm{R}'_j(t') = t$. We know that t is obtained by cutting off some subtrees from t', and those subtrees are rooted at j, so they are elements of $T'(j, j, L \cup \{j\})$. Therefore, t' can be built by inserting trees from $T'(j, j, L \cup \{j\})$ into the marked j-leaves of t, so we have $t' \in T'_t(j, j, L \cup \{j\})$.

"$\supseteq$": For any $t' \in T'_t(j, j, L \cup \{j\})$, we clearly have $\mathrm{R}'_j(t') = t$ and $t' \in R'^{-1}_j(\{t\})$.

Thereby, we have shown

$$C(i, j-1, L \cup \{j\}) \cdot_j C(j, j-1, L \cup \{j\})^{*_j} = \sum_{t \in T(i,j,L \cup \{j\})} \pi[\![\varphi \mathrm{U} \psi]\!]_t = C(i, j, L \cup \{j\}).$$

In case that $j \in L$, we have

$$C(i, j, L) = C(i, j, L \cup \{j\}) = C(i, j-1, L \cup \{j\}) \cdot_j C(j, j-1, L \cup \{j\})^{*_j}.$$

Otherwise, we have $j \notin L$ and claim that

$$C(i, j, L) = C(i, j, L \cup \{j\}) \cdot_j 0 = C(i, j-1, L \cup \{j\}) \cdot_j C(j, j-1, L \cup \{j\})^{*_j} \cdot_j 0.$$

The first equality is proven by

$$C(i, j, L \cup \{j\}) \cdot_j 0 = \left(\sum_{t \in T(i,j,L \cup \{j\})} \pi[\![\varphi \mathrm{U} \psi]\!]_t \right) \cdot_j 0 = \sum_{t \in T(i,j,L \cup \{j\})} \pi[\![\varphi \mathrm{U} \psi]\!]_t \cdot_j 0.$$

Now, if $t \in T(i, j, L)$, then t does not contain a marked j-leaf, since $j \notin L$, so $\pi[\![\varphi \mathrm{U} \psi]\!]_t \cdot_j 0 = \pi[\![\varphi \mathrm{U} \psi]\!]_t$. Otherwise, t must contain a marked j-leaf, but then we have $\pi[\![\varphi \mathrm{U} \psi]\!]_t \cdot_j 0 = 0$. Thus, we conclude

$$C(i, j, L \cup \{j\}) \cdot_j 0 = \sum_{t \in T(i,j,L \cup \{j\})} \pi[\![\varphi \mathrm{U} \psi]\!]_t \cdot_j 0 = \sum_{t \in T(i,j,L)} \pi[\![\varphi \mathrm{U} \psi]\!]_t = C(i, j, L),$$

which ends the proof of the claim

$$C(i, j, L) = \begin{cases} C(i, j-1, L \cup \{j\}) \cdot_j C(j, j-1, L \cup \{j\})^{*_j} & \text{if } j \in L \\ C(i, j-1, L \cup \{j\}) \cdot_j C(j, j-1, L \cup \{j\})^{*_j} \cdot_j 0 & \text{otherwise.} \end{cases}$$

We can close this subsection by summarizing the algorithm.

(4.26) Algorithm. Let K be an ω-continuous semiring. The input of the algorithm is a formula $\mathrm{A}(\varphi \mathrm{U} \psi)$ in CTL and a matching K-interpretation π over a finite transition system V with $|V| = n$ and a node $v \in V$. The output is a representation of $\pi[\![\mathrm{A}(\varphi \mathrm{U} \psi)]\!]_v$ using addition, multiplication and the operators $\cdot_j$ and $*_j$ for $1 \leq j \leq n$.

1. Rename the elements of V arbitrarily to $V = \{1, ..., n\}$.
2. Compute $C(i, 0, L)$ for $1 \leq i \leq n$ and $L \subseteq V$ directly.
3. Repeat for $j = 1, ..., n$:

(a) Compute $C(i,j,L)$ for $1 \leq i \leq n$ and $L \subseteq V$ as

$$C(i,j,L) = \begin{cases} C(i,j-1,L \cup \{j\}) \cdot_j C(j,j-1,L \cup \{j\})^{*_j} & \text{if } j \in L \\ C(i,j-1,L \cup \{j\}) \cdot_j C(j,j-1,L \cup \{j\})^{*_j} \cdot_j 0 & \text{otherwise.} \end{cases}$$

4. The output is $C(v,n,\emptyset)$.

Proof. We first observe that the algorithm terminates. In step 2, $n \cdot 2^n$ values have to be calculated, each $C(i,0,L)$ can be found by evaluating at most two trees with up to $n+1$ nodes, which yields a runtime in $\mathcal{O}(n^2 \cdot 2^n)$. In step 3, a total of $n^2 \cdot 2^n$ values is calculated, we assume that each of them can be calculated in constant time, so the total runtime of the algorithm is in $\mathcal{O}(n^2 \cdot 2^n)$.

The correctness of step 3 (a) has been shown above. The result is $C(v,n,\emptyset)$, which is equal to

$$C(v,n,\emptyset) = \sum_{t \in T(v,n,\emptyset)} \pi[\![\varphi \mathrm{U} \psi]\!]_t.$$

Since $T(v,n,\emptyset) = T_v^{\mathrm{fin}}(V)$, we can use theorem (4.14) to obtain

$$C(v,n,\emptyset) = \sum_{t \in T_v^{\mathrm{fin}}(V)} \pi[\![\varphi \mathrm{U} \psi]\!]_t = \pi[\![\mathrm{A}(\varphi \mathrm{U} \psi)]\!]_v.$$

This ends the proof of the correctness of the algorithm. □

We conclude that the universal until operator can be evaluated in exponential time.

4.3 Release Operators in CTL

Formulas with a release operator $\mathrm{E}(\varphi \mathrm{R} \psi)$ or $\mathrm{A}(\varphi \mathrm{R} \psi)$ in CTL have to be interpreted in absorptive lattice semirings K. Looking at theorem (4.14) for until formulas, we expect that a similar theorem can be shown for release formulas. Let π be a K-interpretation over V and $v \in V$. The truth of $\mathrm{E}(\varphi \mathrm{R} \psi)$ at node v can be witnessed by any path $p \in P_v(V)$ and the truth of $\mathrm{A}(\varphi \mathrm{R} \psi)$ can be witnessed by any tree $t \in T_v(V)$. We have already introduced $\varphi \mathrm{R} \psi$-costs for paths and trees to evaluate them under π, so we can state the following theorem for release formulas.

(4.27) Theorem. Let K be an absorptive lattice semiring and V a finite set of nodes. If π is a K-interpretation over V, v is an element of V and φ and ψ are formulas in CTL, then

$$(1) \quad \pi[\![\mathrm{E}(\varphi \mathrm{R} \psi)]\!]_v = \sum_{p \in P_v(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_p \quad \text{and}$$

$$(2) \quad \pi[\![\mathrm{A}(\varphi \mathrm{R} \psi)]\!]_v = \sum_{t \in T_v(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_t.$$

We will provide a proof for this theorem in the following subsections. First of all, notice the difference to theorem (4.14), where only finite paths were included. For release formulas, infinite paths have to be included as well. The sums in the above

theorem are therefore uncountable in the general case. Also, we are working with absorptive lattice semirings K. Later, we will see that absorption will simplify the above theorem and that we will not have to compute the sum of the $\varphi \mathrm{R} \psi$-costs of all the paths and trees over V, but instead, many of them will be absorbed by other paths and trees.

4.3.1 Existential Release Operators

The goal of this subsection is to prove part (1) of theorem (4.27). Let K be an absorptive lattice semiring and $\mathrm{E}(\varphi \mathrm{U} \psi)$ a CTL formula with a matching K-interpretation π over a finite set of nodes V. In order to illustrate how we can use absorption, consider the paths p, p' and p'' from figure (4.28). The nodes and edges are labelled with their contributions to the $\varphi \mathrm{R} \psi$-cost of the respective path.

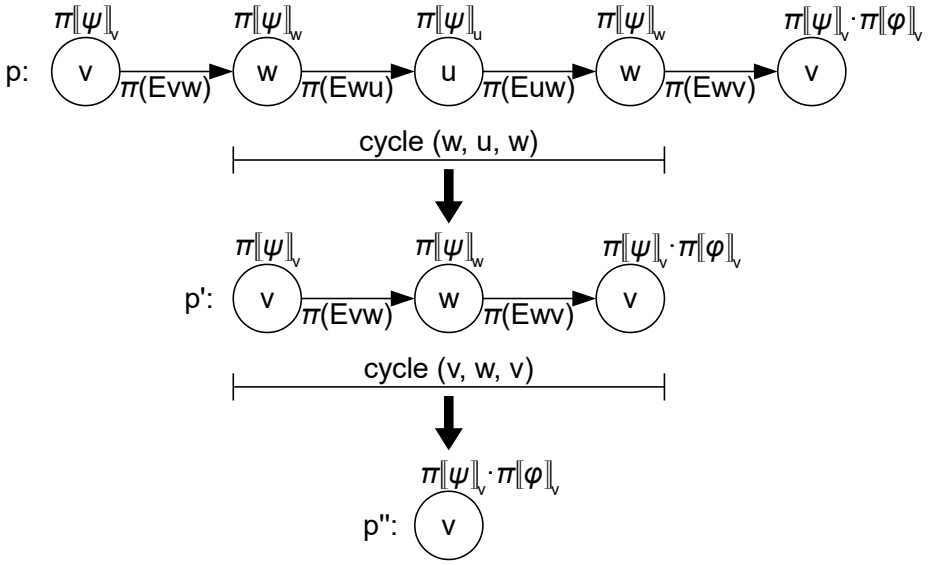

Figure (4.28): Three paths p, p' and p''.

Observe that p is the longest path, p' is obtained from p by removing the w-cycle (w, u, w) in the middle of p and replacing it with w and p'' is obtained from p' by removing yet another cycle, that is, the v-cycle (v, w, v), from p'. Recall that multiplication decreases elements in absorptive lattice semirings, and since we have

$$\begin{aligned} \pi[\![\varphi \mathrm{R} \psi]\!]_p &= \pi[\![\varphi \mathrm{R} \psi]\!]_{p'} \cdot \pi[\![\psi]\!]_w \cdot \pi(Ewu) \cdot \pi[\![\psi]\!]_u \cdot \pi(Euw) \quad \text{and} \\ \pi[\![\varphi \mathrm{R} \psi]\!]_{p'} &= \pi[\![\varphi \mathrm{R} \psi]\!]_{p''} \cdot \pi[\![\psi]\!]_v \cdot \pi(Evw) \cdot \pi[\![\psi]\!]_w \cdot \pi(Ewv), \end{aligned}$$

it follows that $\pi[\![\varphi \mathrm{R} \psi]\!]_p \leq \pi[\![\varphi \mathrm{R} \psi]\!]_{p'} \leq \pi[\![\varphi \mathrm{R} \psi]\!]_{p''}$. Absorption implies that

$$\pi[\![\varphi \mathrm{R} \psi]\!]_p + \pi[\![\varphi \mathrm{R} \psi]\!]_{p'} + \pi[\![\varphi \mathrm{R} \psi]\!]_{p''} = \pi[\![\varphi \mathrm{R} \psi]\!]_{p''},$$

meaning that the costs of p'' absorb the costs of p and p'. We will often simplify this and say that p'' absorbs p and p', which refers to the $\varphi \mathrm{R} \psi$-costs of the respective paths. This observation shows us that removing cycles from paths yields paths with greater costs, therefore, the "shortest" paths, or more precisely, the cycle-free paths have the greatest $\varphi \mathrm{R} \psi$-costs and absorb "longer" paths.

Also, the above example motivates us to introduce ψ-costs for partial, incomplete paths or cycles. If we set

$$\pi[\![\psi]\!]_{(w,u,w)} = \pi[\![\psi]\!]_w \cdot \pi(Ewu) \cdot \pi[\![\psi]\!]_u \cdot \pi(Euw) \quad \text{and}$$
$$\pi[\![\psi]\!]_{(v,w,v)} = \pi[\![\psi]\!]_v \cdot \pi(Evw) \cdot \pi[\![\psi]\!]_w \cdot \pi(Ewv),$$

then we obtain

$$\pi[\![\varphi \mathrm{R} \psi]\!]_p = \pi[\![\varphi \mathrm{R} \psi]\!]_{p'} \cdot \pi[\![\psi]\!]_{(w,u,w)} \quad \text{and}$$
$$\pi[\![\varphi \mathrm{R} \psi]\!]_{p'} = \pi[\![\varphi \mathrm{R} \psi]\!]_{p''} \cdot \pi[\![\psi]\!]_{(v,w,v)}.$$

Notice that for the ψ-costs, we evaluate the edges and we evaluate ψ at the internal nodes normally, we would do the same when computing the $\varphi \mathrm{R} \psi$-costs for the respective paths, but when computing ψ-costs, we do not evaluate the terminal node at all. This is very useful if we define the *appending* operation for paths $q_1 \circ q_2$, where q_1 and q_2 are paths such that q_1 ends at the starting node of q_2. Then, $q_1 \circ q_2$ is defined by replacing the last node of q_1 with q_2. If we now append the cycle (w, u, w) to itself, we obtain $(w, u, w) \circ (w, u, w) = (w, u, w, u, w)$ and conveniently, because we have defined ψ-costs to ignore the terminal node, we have

$$\pi[\![\psi]\!]_{(w,u,w)\circ(w,u,w)} = \pi[\![\psi]\!]_{(w,u,w)} \cdot \pi[\![\psi]\!]_{(w,u,w)}.$$

We will now define ψ-costs formally and show some properties.

(4.29) Definition (ψ-Costs for Paths)**.** Let K be an absorptive lattice semiring. If π is a K-interpretation and p is a finite or infinite path over V and φ and ψ are CTL formulas, the ψ-cost of p is defined as

$$\pi[\![\psi]\!]_p = \left(\prod_{x \in i(p)} \pi[\![\psi]\!]_{L_p(x)} \right) \cdot \left(\prod_{(x,y) \in e(p)} \pi(E L_p(x) L_p(y)) \right).$$

As argued above, for finite paths p and q where p ends at the starting node of q, we generally have

$$\pi[\![\psi]\!]_{(p \circ q)} = \pi[\![\psi]\!]_p \cdot \pi[\![\psi]\!]_q.$$

Also, if c is a v-cycle for some $v \in V$, we can chain c with itself arbitrarily often. We recursively define

$$c^0 = (v) \quad \text{and}$$
$$c^{k+1} = c^k \circ c \quad \text{for } k \in \omega.$$

Then, it follows that

$$\pi[\![\psi]\!]_{c^k} = \pi[\![\psi]\!]_c^k \quad \text{for } k \in \omega.$$

Let c^∞ be the infinite path obtained by appending c to itself infinitely often, then

$$\pi[\![\psi]\!]_{c^\infty} = \pi[\![\psi]\!]_c^\infty.$$

It is also worth noticing the relationship of $\varphi \mathrm{R} \psi$-costs and ψ-costs. For finite paths p, we have

$$\pi[\![\varphi \mathrm{R} \psi]\!]_p = \pi[\![\psi]\!]_p \cdot \pi[\![\psi]\!]_{L_p(t(p))} \cdot \pi[\![\varphi]\!]_{L_p(t(p))},$$

or in words, we can obtain the $\varphi \mathrm{R} \psi$-costs of p by computing the ψ-costs of p and evaluating the terminal node. If p is an infinite path, then we even have

$$\pi[\![\varphi \mathrm{R} \psi]\!]_p = \pi[\![\psi]\!]_p,$$

since infinite paths only have internal nodes, and therefore $n(p) = i(p)$.

Another interesting observation can be made by defining $p|_k$ for any path p over V and any $k \in \omega$ as the path that consists of the first k nodes in p. If p is finite and has less than k nodes, then $p|_k = p$. For infinite paths p, we notice that

$$\pi[\![\varphi \mathrm{R} \psi]\!]_p = \pi[\![\psi]\!]_p = \inf_{k \in \omega} \pi[\![\psi]\!]_{p|_k}.$$

This can be shown by using the partition-invariance of countable multiplication from proposition (2.16) and rearranging the factors of $\pi[\![\psi]\!]_p$.

Crucially, we conclude that a finite or infinite path p that contains a v-cycle c for some $v \in V$ as a subpath can be transformed into a path p' by removing c from p and replacing it with v. The definitions and observations above allow us to infer that

$$\pi[\![\varphi \mathrm{R} \psi]\!]_p = \pi[\![\varphi \mathrm{R} \psi]\!]_{p'} \cdot \pi[\![\psi]\!]_c,$$

so that p' absorbs p. The observation that we can remove cycles while increasing the paths' $\varphi \mathrm{R} \psi$-costs will be very useful when proving theorem (4.27).

Recall that we want to prove part (1) of theorem (4.27), which we can now restate to

$$\pi[\![\mathrm{E}(\varphi \mathrm{R} \psi)]\!]_v = \sum_{p \in P_v^{\mathrm{fin}}(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_p + \sum_{p \in P_v^{\mathrm{inf}}(V)} \pi[\![\psi]\!]_p,$$

since we have shown that $\varphi \mathrm{R} \psi$-costs of infinite paths correspond to their ψ-costs.

According to definition (3.11), we have

$$\pi[\![\mathrm{E}(\varphi \mathrm{R} \psi)]\!]_v = \mathrm{gfp}(f^{\mathrm{E}(\varphi \mathrm{R} \psi)})_v,$$

which can be computed by means of theorem (2.17), that is, we set $X_0 = 1 \in K^V$ and start a transfinite iteration of $f^{\mathrm{E}(\varphi \mathrm{R} \psi)}$. We define

$$\begin{aligned} X_{i+1} &= f^{\mathrm{E}(\varphi \mathrm{R} \psi)}(X_i) \quad \text{for } i \in \omega \text{ and} \\ X_\omega &= \inf_{i \in \omega} X_i. \end{aligned}$$

As we will see, X_ω is already a fixed point so there will be no need to iterate any further.

(4.30) Lemma. With X_i for $i \in \omega$ defined as above, we have

$$(X_i)_v = \sum_{p \in P_v^{<i}(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_p + \sum_{p \in P_v^{=i}(V)} \pi[\![\psi]\!]_p \quad \text{for } v \in V.$$

Proof. We will prove the lemma by induction on i. For $i = 0$, $P_v^{<0}(V)$ is empty and $P_v^{=0}(V)$ contains only the path (v). Since (v) does not have any edges or internal nodes, we obtain $\pi[\![\psi]\!]_{(v)} = 1 = (X_0)_v$.

For $i+1$, we have $(X_{i+1})_v$

$$
\begin{aligned}
&= f_v^{\mathrm{E}(\varphi \mathrm{R} \psi)}(X_i) \\
&= \pi[\![\psi]\!]_v \cdot \left(\pi[\![\varphi]\!]_v + \sum_{w \in vE} \pi(Evw) \cdot (X_i)_w \right) \\
&= \pi[\![\psi]\!]_v \cdot \pi[\![\varphi]\!]_v + \sum_{w \in vE} \pi[\![\psi]\!]_v \cdot \pi(Evw) \cdot (X_i)_w \\
&= \pi[\![\psi]\!]_v \cdot \pi[\![\varphi]\!]_v + \sum_{w \in vE} \pi[\![\psi]\!]_v \cdot \pi(Evw) \cdot \left(\sum_{p \in P_w^{<i}(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_p + \sum_{p \in P_w^{=i}(V)} \pi[\![\psi]\!]_p \right)
\end{aligned}
$$

by induction hypothesis, which can be further simplified to $(X_{i+1})_v$

$$
\begin{aligned}
&= \pi[\![\psi]\!]_v \cdot \pi[\![\varphi]\!]_v + \left(\sum_{w \in vE} \sum_{p \in P_w^{<i}(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_{(v,p)} \right) + \left(\sum_{w \in vE} \sum_{p \in P_w^{=i}(V)} \pi[\![\psi]\!]_{(v,p)} \right) \\
&= \sum_{p \in P_v^{=0}(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_p + \sum_{p \in \bigcup_{1 \le j \le i} P_v^{=j}(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_p + \sum_{p \in P_v^{=i+1}(V)} \pi[\![\psi]\!]_p \\
&= \sum_{p \in P_v^{<i+1}(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_p + \sum_{p \in P_v^{=i+1}(V)} \pi[\![\psi]\!]_p,
\end{aligned}
$$

where (v,p) is the path obtained by appending v to the beginning of p. Clearly, (v,p) where $p \in P_w^{<i}(V)$ for $w \in vE$ yields exactly all the paths starting at v with a length between 1 and i, disregarding paths with cost zero, and (v,p) with $p \in P_w^{=i}(V)$ for $w \in vE$ yields exactly the paths starting at v with the length $i+1$, again omitting zero-cost paths. This ends the induction and proves the lemma. □

The next step is to prove that

$$
(X_\omega)_v = \sum_{p \in P_v^{\mathrm{fin}}(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_p + \sum_{p \in P_v^{\mathrm{inf}}(V)} \pi[\![\psi]\!]_p \quad \text{for } v \in V.
$$

Lemma (2.21) implies

$$
(X_\omega)_v = \left(\inf_{i \in \omega} X_i \right)_v = \inf_{i \in \omega} (X_i)_v.
$$

Using lemma (4.30), it remains to show

$$
\sum_{p \in P_v^{\mathrm{fin}}(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_p + \sum_{p \in P_v^{\mathrm{inf}}(V)} \pi[\![\psi]\!]_p = \inf_{i \in \omega} \underbrace{\left(\sum_{p \in P_v^{<i}(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_p + \sum_{p \in P_v^{=i}(V)} \pi[\![\psi]\!]_p \right)}_{(X_i)_v}.
$$

"$\le$": In absorptive lattice semirings, summations are the same as suprema, so we have to show that each summand s on the left side is a lower bound of

$$
\{(X_i)_v \mid i \in \omega\} = \left\{ \sum_{p \in P_v^{<i}(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_p + \sum_{p \in P_v^{=i}(V)} \pi[\![\psi]\!]_p \mid i \in \omega \right\},
$$

thereby showing that the infimum is greater than s. We distinguish two cases.

If $s = \pi[\![\varphi \mathrm{R} \psi]\!]_p$ for some $p \in P_v^{\mathrm{fin}}(V)$, let k be the length of the path p. For $i \leq k$, consider $p|_i \in P_v^{=i}(V)$. Clearly the ψ-cost of $p|_i$ is greater than the $\varphi \mathrm{R} \psi$-cost of p, since it is a subpath of p. So, we have $s \leq \pi[\![\psi]\!]_{p|_i} \leq (X_i)_v$. For $i > k$, we have $p \in P_v^{<i}(V)$, which immediately implies $s \leq (X_i)_v$. Therefore, s is a lower bound of $\{(X_i)_v \mid i \in \omega\}$.

The other possibility is that $s = \pi[\![\psi]\!]_p$ for some $p \in P_v^{\mathrm{inf}}(V)$. In that case, for each $i \in \omega$, we have $p|_i \in P_v^{=i}(V)$ with $\pi[\![\psi]\!]_{p|_i} \geq \pi[\![\psi]\!]_p = s$, so $s \leq (X_i)_v$, which makes s a lower bound of $\{(X_i)_v \mid i \in \omega\}$. This ends the direction "$\leq$", but before proceeding to the other direction, notice that from this case, we get the corollary

$$(*) \qquad \sum_{p \in P_v^{\mathrm{inf}}(V)} \pi[\![\psi]\!]_p \leq \inf_{i \in \omega} \left(\sum_{p \in P_v^{=i}(V)} \pi[\![\psi]\!]_p \right).$$

"$\geq$": Instead of proving this direction directly, we will show

$$\begin{aligned} \sum_{p \in P_v^{\mathrm{fin}}(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_p + \sum_{p \in P_v^{\mathrm{inf}}(V)} \pi[\![\psi]\!]_p &\geq \inf_{i \in \omega} \left(\sum_{p \in P_v^{\mathrm{fin}}(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_p + \sum_{p \in P_v^{=i}(V)} \pi[\![\psi]\!]_p \right) \\ &\geq \inf_{i \in \omega} \left(\sum_{p \in P_v^{<i}(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_p + \sum_{p \in P_v^{=i}(V)} \pi[\![\psi]\!]_p \right). \end{aligned}$$

Clearly, the second inequality is fulfilled, so we have to show that the first inequality

$$\sum_{p \in P_v^{\mathrm{fin}}(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_p + \sum_{p \in P_v^{\mathrm{inf}}(V)} \pi[\![\psi]\!]_p \geq \inf_{i \in \omega} \left(\underbrace{\sum_{p \in P_v^{\mathrm{fin}}(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_p}_{:= c} + \sum_{p \in P_v^{=i}(V)} \pi[\![\psi]\!]_p \right)$$

is true as well. Since the term that we labelled with c does not depend on i, we can use part (2) of lemma (2.11) to obtain

$$\begin{aligned} \inf_{i \in \omega} \left(c + \sum_{p \in P_v^{=i}(V)} \pi[\![\psi]\!]_p \right) &= c + \inf_{i \in \omega} \left(\sum_{p \in P_v^{=i}(V)} \pi[\![\psi]\!]_p \right) \\ &= \sum_{p \in P_v^{\mathrm{fin}}(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_p + \inf_{i \in \omega} \left(\sum_{p \in P_v^{=i}(V)} \pi[\![\psi]\!]_p \right) \end{aligned}$$

Thanks to the monotonicity of addition, we only have to show

$$\sum_{p \in P_v^{\mathrm{inf}}(V)} \pi[\![\psi]\!]_p \geq \inf_{i \in \omega} \left(\sum_{p \in P_v^{=i}(V)} \pi[\![\psi]\!]_p \right)$$

to complete the direction "$\geq$". We state this as a lemma.

(4.31) Lemma. In an absorptive lattice semiring K, for any CTL formula ψ and

any matching K-interpretation π over a finite set of nodes V , we have

$$\sum_{p \in P_v^{\inf}(V)} \pi[\![\psi]\!]_p = \inf_{i \in \omega} \left(\sum_{p \in P_v^{=i}(V)} \pi[\![\psi]\!]_p \right)$$

Proof. The direction "$\leq$" is already shown above as a corollary $(*)$. We will now prove the converse direction "$\geq$". First, we rewrite the sum on the left side as a supremum and invoke the complete distributivity of K to obtain

$$\inf_{i \in \omega} \left(\sum_{p \in P_v^{=i}(V)} \pi[\![\psi]\!]_p \right) = \inf_{i \in \omega} \left(\sup_{p \in P_v^{=i}(V)} \pi[\![\psi]\!]_p \right) = \sup_{f \in F} \inf_{i \in \omega} \pi[\![\psi]\!]_{f(i)},$$

where F is the set of choice functions $f : \omega \to P_v^{\text{fin}}(V)$ such that $f(i) \in P_v^{=i}(V)$ for each $i \in \omega$. Our goal is to show that

$$\sup_{f \in F} \inf_{i \in \omega} \pi[\![\psi]\!]_{f(i)} \leq \sum_{p \in P_v^{\inf}(V)} \pi[\![\psi]\!]_p = \sup_{p \in P_v^{\inf}(V)} \pi[\![\psi]\!]_p.$$

For each $f \in F$, we will find a $p \in P_v^{\inf}(V)$ such that

$$\inf_{i \in \omega} \pi[\![\psi]\!]_{f(i)} \leq \pi[\![\psi]\!]_p.$$

This will be sufficient to prove the claim. Therefore, we fix an arbitrary $f \in F$. We can picture f as an infinite sequence of paths with increasing length. An example is shown in figure (4.32).

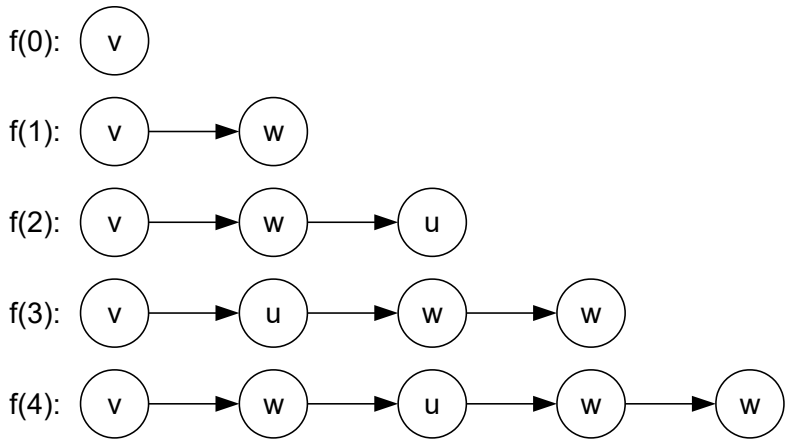

Figure (4.32): An example for the first 5 values of a function $f \in F$.

The difficulty of this proof is that the paths chosen by f are generally unrelated. In the example above, we see that $f(1)$ is a continuation of $f(0)$ and $f(2)$ is a continuation of $f(1)$, however, $f(3)$ is a completely unrelated path, whereas $f(4)$ is a continuation of $f(2)$. Nevertheless, we will show that there must be repeating patterns in f. In the above example, imagine that $V = \{v, w, u\}$. Since $f(3)$ and $f(4)$ have a length greater than 2, there must be at least one node repetition in those paths.

Returning from the example to the general case, we claim that there must be a node $w \in V$ that *occurs arbitrarily often* in f. Formally, this means that for any $j \in \omega$,

there is an $i \in \omega$ such that $f(i)$ contains at least j occurrences of w. If the opposite were true, then for any $w \in V$, there would be an $M_w \in \omega$ such that no $f(i)$ contains M_w or more occurrences of w. This would be a contradiction, since

$$M = \sum_{w \in V} M_w$$

would be an upper bound on the length of $f(i)$ for all $i \in \omega$, but this is impossible, since $f(M+1)$ by definition has the length $M+1$. Therefore, we can fix a $w \in V$ that occurs arbitrarily often in f.

Now, we define the w-reduction function $\mathrm{R}_w : P_v^{\mathrm{fin}}(V) \to P_v^{\mathrm{fin}}(V)$. Suppose that p is a path in $P_v^{\mathrm{fin}}(V)$. Then, we can divide p into sections at each occurrence of w. More precisely, the first section starts at the beginning of the path at the node v and at each occurrence of w, a new section starts. So, each section is a path of the form $(w, ...)$ with the exception of the first section $(v, ...)$ and each section contains at most one occurrence of w. We now define $\mathrm{R}_w(p)$ as the path that is obtained by removing all the cycles of p in each of its sections. Formally, removing a cycle $(u, ..., u)$ from p is done by replacing the subpath $(u, ..., u)$ in p by a single occurrence of u. As argued above, this increases the costs of paths, so we have

$$\pi[\![\psi]\!]_{\mathrm{R}_w(p)} \geq \pi[\![\psi]\!]_p$$

for each $p \in P_v^{\mathrm{fin}}(V)$, since R_w only removes cycles from p. Also, we notice that R_w never removes any occurrences of w, since any section contains at most one occurrence of w, so there cannot be any w-cycles in the sections of p. Figure (4.33) illustrates how R_w is performed on a path $p \in P_v^{\mathrm{fin}}(V)$.

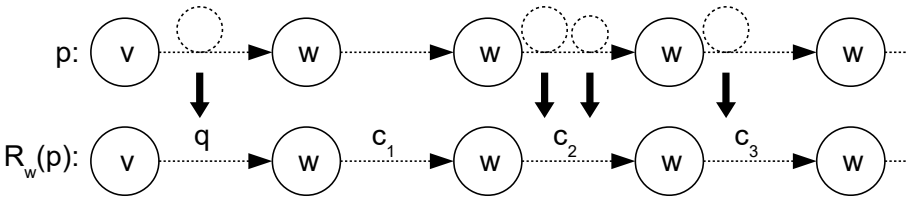

Figure (4.33): Informal illustration of a path p and $\mathrm{R}_w(p)$.

We observe that $\mathrm{R}_w(p)$ starts with an initial segment q, which a path that starts at v and ends at w, and multiple w-cycles $c_1, c_2, ...$ are appended to q. Thanks to the reduction R_w, neither q nor $c_1, c_2, ...$ contain any node repetitions, aside from the fact that the cycles start and end at w. Therefore, we consider the set $P'_{v \to w}(V)$ of all paths over V without node repetitions that start at v and end at w and the set $C'_w(V)$ of all w-cycles over V without node repetitions aside from the two endpoints. Clearly, both of those sets are finite.

Now, with the same argument as above, we claim that there is a $c \in C'_w(V)$ that occurs arbitrarily often in $\mathrm{R}_w \circ f$, that is, for each $j \in \omega$, there is an $i \in \omega$ such that $\mathrm{R}_w(f(i))$ contains at least j occurrences of c. If the number of occurrences of each cycle $c \in C'_w(V)$ was bounded, then there would be a bound C on the number of w-cycles that can occur in $\mathrm{R}_w(f(i))$ for each $i \in \omega$. However, such a bound C cannot exist, since w occurs arbitrarily often in f, therefore, there would be an i such that $f(i)$ would contain at least $C+2$ occurrences of w. Since R_w does not remove any occurrences of w, $\mathrm{R}_w(f(i))$ would then contain at least $C+1$ w-cycles in $C'_w(V)$

between the $C+2$ occurrences of w, which would be a contradiction. Consequently, we fix a $c \in C'_w(V)$ that occurs arbitrarily often in $\mathrm{R}_w \circ f$.

Before concluding the proof, we also need a $q \in P'_{v\to w}(V)$ such that c occurs arbitrarily often in $\mathrm{R}_w \circ f$ with the initial segment q. Formally, this means that for all $j \in \omega$, there is an $i \in \omega$ such that $\mathrm{R}_w(f(i))$ starts with q and contains at least j occurrences of c. Since $P'_{v\to w}(V)$ is finite, such a $q \in P'_{v\to w}(V)$ must exist, because otherwise, the number of occurrences of c in $\mathrm{R}_w \circ f$ with the initial segment q would be bounded by m_q for each $q \in P'_{v\to w}(V)$, which is impossible, because

$$\max\{m_q \mid q \in P'_{v\to w}(V)\} \in \omega$$

would exist and be an upper bound on the occurrences of c in $\mathrm{R}_w \circ f$, contradicting the choice of c. So, we also fix a $q \in P'_{v\to w}(V)$ such that c occurs arbitrarily often in $\mathrm{R}_w \circ f$ with the initial segment q.

Now, consider the infinite path $p = q \circ c^\infty \in P_v^{\inf}(V)$, which is built by starting with q and appending c infinitely many times. We know that

$$\pi[\![\psi]\!]_p = \inf_{k\in\omega} \pi[\![\psi]\!]_{p|_k}.$$

Let $k \in \omega$ be arbitrary, then we can extend $p|_k$ to the next occurrence of w in p and obtain a path p' of the form $p' = q \circ c^l$ for some $l \in \omega$. Clearly, since p' is an extension of $p|_k$, we have

$$\pi[\![\psi]\!]_{p'} \leq \pi[\![\psi]\!]_{p|_k}.$$

By choice of q and c, we know that there is an $i \in \omega$, such that $\mathrm{R}_w(f(i))$ starts with q and contains c at least l times, so we have

$$\pi[\![\psi]\!]_{f(i)} \leq \pi[\![\psi]\!]_{\mathrm{R}_w(f(i))} \leq \pi[\![\psi]\!]_{q\circ c^l} = \pi[\![\psi]\!]_{p'} \leq \pi[\![\psi]\!]_{p|_k}.$$

Since for each $k \in \omega$, there is an $i \in \omega$ with

$$\pi[\![\psi]\!]_{f(i)} \leq \pi[\![\psi]\!]_{p|_k},$$

we conclude that

$$\inf_{i\in\omega} \pi[\![\psi]\!]_{f(i)} \leq \inf_{k\in\omega} \pi[\![\psi]\!]_{p|_k} = \pi[\![\psi]\!]_p.$$

This ends the proof. □

With lemma (4.31), we can conclude

$$(X_\omega)_v = \sum_{p\in P_v^{\mathrm{fin}}(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_p + \sum_{p\in P_v^{\inf}(V)} \pi[\![\psi]\!]_p \quad \text{for } v \in V.$$

Before showing that this is a fixed point, we observe that the proof of the lemma yields an interesting corollary for the infinite paths over V.

(4.34) Corollary. Let K be an absorptive lattice semiring, V a finite set of nodes, ψ a formula in CTL and π matching K-interpretation over V. If we define $P'_{v\to w}(V)$ and $C'_w(V)$ for $w \in V$ as above, we obtain

$$\sum_{p\in P_v^{\inf}(V)} \pi[\![\psi]\!]_p = \sum_{w\in V} \sum_{(q,c)\in P'_{v\to w}(V)\times C'_w(V)} \pi[\![\psi]\!]_q \cdot \pi[\![\psi]\!]_c^\infty.$$

Proof sketch. In the direction "$\geq$" of the proof of lemma (4.31) we have only used infinite paths of the form $q \circ c^\infty$ for $q \in P'_{v\to w}(V)$ and $c \in C'_w(V)$ for some $w \in V$ instead of all infinite paths over V, we have actually shown that

$$\sum_{w\in V} \sum_{(q,c)\in P'_{v\to w}(V)\times C'_w(V)} \pi[\![\psi]\!]_{q\circ c^\infty} \geq \inf_{i\in\omega} \left(\sum_{p\in P_v^{=i}(V)} \pi[\![\psi]\!]_p \right)$$

Clearly, the paths $q \circ c^\infty$ are a subset of all infinite paths starting at v, so we have

$$\sum_{w\in V} \sum_{(q,c)\in P'_{v\to w}(V)\times C'_w(V)} \pi[\![\psi]\!]_{q\circ c^\infty} \leq \sum_{p\in P_v^{\inf}(V)} \pi[\![\psi]\!]_p.$$

Additionally, the direction "$\leq$" of lemma (4.31) implies that

$$\sum_{p\in P_v^{\inf}(V)} \pi[\![\psi]\!]_p \leq \inf_{i\in\omega} \left(\sum_{p\in P_v^{=i}(V)} \pi[\![\psi]\!]_p \right) \leq \sum_{w\in V} \sum_{(q,c)\in P'_{v\to w}(V)\times C'_w(V)} \pi[\![\psi]\!]_{q\circ c^\infty}.$$

The circular inequality between these three values implies that they are all equal, that is

$$\sum_{p\in P_v^{\inf}(V)} \pi[\![\psi]\!]_p = \inf_{i\in\omega} \left(\sum_{p\in P_v^{=i}(V)} \pi[\![\psi]\!]_p \right) = \sum_{w\in V} \sum_{(q,c)\in P'_{v\to w}(V)\times C'_w(V)} \pi[\![\psi]\!]_{q\circ c^\infty}.$$

Finally, the observation that

$$\pi[\![\psi]\!]_{q\circ c^\infty} = \pi[\![\psi]\!]_q \cdot \pi[\![\psi]\!]_c^\infty$$

for any $(q, c) \in P'_{v\to w}(V) \times C'_w(V)$ proves the corollary. □

We will use this corollary later. Now, it is still left to prove that (X_ω) is a fixed point of $f^{\mathrm{E}(\varphi \mathrm{R} \psi)}$. For any $v \in V$, we have

$$\begin{aligned} f_v^{\mathrm{E}(\varphi \mathrm{R} \psi)}(X_\omega) &= \pi[\![\psi]\!]_v \cdot \left(\pi[\![\varphi]\!]_v \cdot \sum_{w\in vE} \pi(Evw) \cdot (X_\omega)_w \right) \\ &= \pi[\![\psi]\!]_v \cdot \pi[\![\varphi]\!]_v \cdot \sum_{w\in vE} \pi[\![\psi]\!]_v \cdot \pi(Evw) \cdot (X_\omega)_w \end{aligned}$$

We have calculated $(X_\omega)_w$ for $w \in V$ above, so we obtain $f_v^{\mathrm{E}(\varphi \mathrm{R} \psi)}(X_\omega)$

$$\begin{aligned} &= \pi[\![\psi]\!]_v \cdot \pi[\![\varphi]\!]_v \cdot \sum_{w\in vE} \pi[\![\psi]\!]_v \cdot \pi(Evw) \cdot \left(\sum_{p\in P_w^{\mathrm{fin}}(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_p + \sum_{p\in P_w^{\inf}(V)} \pi[\![\psi]\!]_p \right) \\ &= \pi[\![\psi]\!]_v \cdot \pi[\![\varphi]\!]_v \cdot \left(\sum_{w\in vE} \sum_{p\in P_w^{\mathrm{fin}}(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_{(v,p)} \right) + \left(\sum_{w\in vE} \sum_{p\in P_w^{\inf}(V)} \pi[\![\psi]\!]_{(v,p)} \right) \\ &= \sum_{p\in P_v^{=0}(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_p + \sum_{p\in \bigcup_{1\leq j<\omega} P_v^{=j}(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_p + \sum_{p\in P_v^{\inf}(V)} \pi[\![\psi]\!]_p \\ &= \sum_{p\in P_v^{\mathrm{fin}}(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_p + \sum_{p\in P_v^{\inf}(V)} \pi[\![\psi]\!]_p \\ &= (X_\omega)_v, \end{aligned}$$

where (v, p) for a path $p \in P_v(V)$ is the path obtained by appending v to the start of the path p. Clearly, the paths (v, p) with $p \in P_w^{\text{fin}}(V)$ for some $w \in vE$ are exactly the finite paths starting at v of length at least 1 and the paths (v, p) with $p \in P_w^{\text{inf}}(V)$ for a $w \in vE$ are exactly the infinite paths starting at v, if we disregard paths with cost 0. This is very similar to the inductive proof of lemma (4.30). Since v was arbitrary, we conclude that $f^{\mathrm{E}(\varphi \mathrm{R} \psi)}(X_\omega) = X_\omega$, therefore we have $\mathrm{gfp}(f^{\mathrm{E}(\varphi \mathrm{R} \psi)}) = X_\omega$ and

$$\pi[\![\mathrm{E}(\varphi \mathrm{R} \psi)]\!]_v = \mathrm{gfp}(f^{\mathrm{E}(\varphi \mathrm{R} \psi)})_v = (X_\omega)_v = \sum_{p \in P_v^{\text{fin}}(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_v + \sum_{p \in P_v^{\text{inf}}(V)} \pi[\![\psi]\!]_p$$

for any $v \in V$, which proves part (1) of theorem (4.27).

The question of how to compute $\pi[\![\mathrm{E}(\varphi \mathrm{R} \psi)]\!]_v$ is immediately answered by corollary (4.34) and the following lemma.

(4.35) Lemma. Let K be an absorptive lattice semiring, $\mathrm{E}(\varphi \mathrm{R} \psi)$ a formula in CTL and π a matching K-interpretation over a finite set of nodes V, then for any $v \in V$, we have

$$\sum_{p \in P_v^{\text{fin}}(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_p = \sum_{p \in P'_v(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_p,$$

where $P'_v(V)$ denotes the set of all paths over V without node repetition that start at V.

Proof. For the direction "$\leq$", recall that removing cycles from paths increases their costs. Therefore, any path $p \in P_v^{\text{fin}}(V)$ is absorbed by some path $p' \in P'_v(V)$, which is obtained by removing all cycles from p, so we have $\pi[\![\varphi \mathrm{R} \psi]\!]_p \leq \pi[\![\varphi \mathrm{R} \psi]\!]_{p'}$. Since sums and suprema are the same in absorptive lattice semirings, the direction "$\leq$" follows.

The converse direction "$\geq$" is immediately clear, since $P'_v(V) \subseteq P_v^{\text{fin}}(V)$. □

Putting part (1) of theorem (4.27) together with lemma (4.35) and corollary (4.34) yields

$$\pi[\![\mathrm{E}(\varphi \mathrm{R} \psi)]\!]_v = \sum_{p \in P'_v(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_p + \sum_{w \in V} \sum_{(q,c) \in P'_{v \to w}(V) \times C'_w(V)} \pi[\![\psi]\!]_q \cdot \pi[\![\psi]\!]_c^\infty.$$

Since all the sets V, $P'_v(V)$, $P'_{v \to w}(V)$ and $C'_w(V)$ are finite and p, q and c are also finite, we can directly compute this value. If $|V| = n$, we observe that $P'_v(V)$, $P'_{v \to w}(V)$ and $C'_w(V)$ each have a cardinality in $\mathcal{O}(n!)$, which is also the time it takes to enumerate their elements. Additionally, computing the costs of p, q and c respectively can be done in $\mathcal{O}(n)$. So, the first sum can be calculated in $\mathcal{O}(n \cdot n!)$ and the second sum can be calculated in $\mathcal{O}(n^2 \cdot (n!)^2)$ operations. We have thereby defined a naive evaluation algorithm for $\pi[\![\mathrm{E}(\varphi \mathrm{R} \psi)]\!]_v$ with a runtime in $\mathcal{O}(n^2 \cdot (n!)^2)$.

The naive algorithm's exponential runtime may be unsatisfactory, therefore it is worth mentioning that it might be possible to adapt the state removal algorithm (4.18) from subsection (4.2.1) to evaluate $\pi[\![\mathrm{E}(\varphi \mathrm{R} \psi)]\!]_v$ in polynomial time. However, since K-automata only keep track of finite paths, whereas release formulas can be witnessed by infinite paths, we would have to introduce a new variation of K-automata where infinite paths are not ignored. Additionally, when performing state

removal, we would have to keep track of infinite paths. For example, when removing the state w, we would have to modify the costs of the other edges to make up for the loss of the infinite path $(w, w, ...)$. Although we will not prove it here, we state the following conjecture.

(4.36) Conjecture. The runtime for the evaluation of $\pi[\![\mathrm{E}(\varphi \mathrm{R} \psi)]\!]_v$ can be reduced to $\mathcal{O}(n^3)$ by using a modified version of the state removal algorithm.

We close this subsection by recalling the $\mathbb{S}^\infty[X]$-interpretation π with $X = \{p, q, r, s\}$ for $\tau = \{E, Q\}$ which was given in example (3.15) and is shown again in figure (4.37).

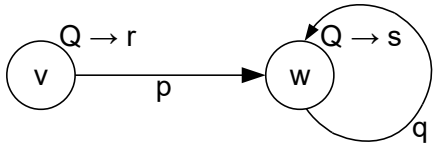

Figure (4.37): $\mathbb{S}^\infty[X]$-interpretation π.

We computed $\pi[\![\mathrm{E}(\mathrm{G}Q)]\!]_v = \pi[\![\mathrm{E}(0\mathrm{R}Q)]\!]_v = pr(qs)^\infty$ with a fixed-point iteration. Using the alternative approach yields

$$\pi[\![\mathrm{E}(0\mathrm{R}Q)]\!]_v = \sum_{p \in P'_v(V)} \pi[\![\mathrm{E}(0\mathrm{R}Q)]\!]_p + \sum_{u \in V} \sum_{(q,c) \in P'_{v \to u}(V) \times C'_u(V)} \pi[\![Q]\!]_q \cdot \pi[\![Q]\!]_c^\infty.$$

Clearly, we can disregard $P'_v(V)$, since $\pi[\![\mathrm{E}(0\mathrm{R}Q)]\!]_p = 0$ for any finite path p. We also disregard paths and cycles with a Q-cost of 0, so we have $C'_v(V) = \emptyset$ and the expression is simplified to

$$\pi[\![\mathrm{E}(0\mathrm{R}Q)]\!]_v = \sum_{(q,c) \in P'_{v \to w}(V) \times C'_w(V)} \pi[\![Q]\!]_q \cdot \pi[\![Q]\!]_c^\infty.$$

Again disregarding paths and cycles with a Q-cost of 0 yields $C'_w(V) = \{(w, w)\}$ and $P'_{v \to w}(V) = \{(v, w)\}$, so we have

$$\pi[\![\mathrm{E}(0\mathrm{R}Q)]\!]_v = \pi[\![Q]\!]_{(v,w)} \cdot \pi[\![Q]\!]_{(w,w)}^\infty = pr(qs)^\infty,$$

which is the expected result.

4.3.2 Universal Release Operators

In this subsection, we will prove part (2) of theorem (4.27). In order to do that, the ideas from the previous section have to be adapted to trees. First of all, recall the concept of trees with marked leaves from definition (4.21). Trees with marked leaves are useful when we want to split trees such that their costs are preserved. In definition (4.22), we defined $\varphi \mathrm{U} \psi$-costs for trees with marked leaves using formal power series. Marked leaves v were evaluated as a variable x_v. In this section, we will define $\varphi \mathrm{R} \psi$-costs for trees with marked leaves, but instead of evaluating marked leaves with variables, we will simply ignore them. This yields the following definition.

(4.38) Definition (Release-Costs for Trees with Marked Leaves)**.** For an absorptive lattice semiring K, a finite set of nodes V and a K-interpretation π, if φ and ψ are

formulas in CTL and t is a finite or infinite complete tree over V with marked leaves, we define the $\varphi \mathrm{R} \psi$-cost of t as

$$\pi[\![\varphi \mathrm{R} \psi]\!]_t = \left(\prod_{x \in n(t) \setminus m(t)} \pi[\![\psi]\!]_{L_t(x)} \right) \cdot \left(\prod_{(x,y) \in e(t)} \pi(EL_t(x)L_t(y)) \right) \cdot \left(\prod_{x \in l(t) \setminus m(t)} \pi[\![\varphi]\!]_{L_t(x)} \right).$$

This definition is justified by figure (4.39), which shows a tree t split into t_1 and t_2. The nodes are labelled with their contributions to the $\varphi \mathrm{R} \psi$-costs of the corresponding tree and the insertion point for t_2 in t_1 is marked with $(*)$. Edge costs are ignored.

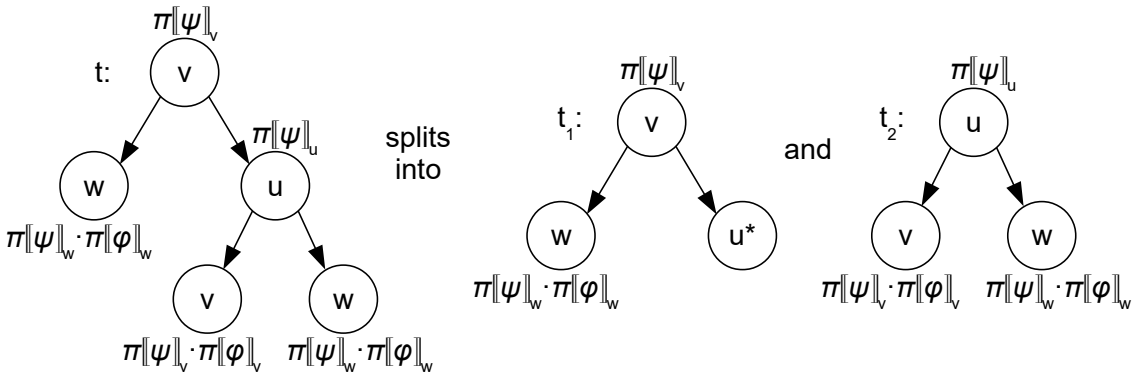

Figure (4.39): A tree t split into t_1 and t_2.

Since we have decided to ignore marked leaves, we obtain the convenient equality

$$\pi[\![\varphi \mathrm{R} \psi]\!]_t = \pi[\![\varphi \mathrm{R} \psi]\!]_{t_1} \cdot \pi[\![\varphi \mathrm{R} \psi]\!]_{t_2},$$

which is generally true when splitting a tree t into t_1 and t_2 along at any node x. The splitting operation can be formally defined. Recall that $T_v(V, M)$ refers to the set of trees with marked leaves where only nodes in M are allowed to be marked.

(4.40) Definition (Tree Split)**.** Let $t \in T_v(V, M)$ be a tree rooted at a node $v \in V$ for some finite set V and $M \subseteq V$. If x is a node in $n(t)$ that is labelled with $L_t(x) = w$, then t can be split at x into t_1 and t_2, where $t_2 \in T_w(V, M)$ is the subtree of t rooted at x and $t_1 \in T_v(V, M \cup \{w\})$ is the tree obtained by replacing t_2 with a single marked w-leaf in t.

Aside from splitting trees, we can also cut trees off at a specific height $h \in \omega$. A node x in a tree t is said to be at height h if the distance of x to the root is exactly h. This yields the following definition.

(4.41) Definition (Tree Cut)**.** Let $t \in T_v(V)$ for a node $v \in V$ and a finite set V. For $h \in \omega$, the *h-cut* of t, denoted as $t|_h$, is defined as the tree that is obtained by cutting t off at any node at height h, that is, for any $x \in n(t)$ at height h, we replace the subtree rooted by x with a marked leaf that has the same label as x. $t|_h$ is a finite tree in $T_v(V, V)$ and its height does not exceed h. Notice that if the height of t is less than h, $t|_h$ is still defined and $t|_h = t$.

This definition can also be reversed. A tree $t \in T(V, M)$ is said to be a *h-cut tree* for $h \in \omega$ if any node $x \in n(t)$ at height h is a marked leaf and there are no other marked leaves. Notice that any tree with a height that is less than h without marked leaves is also a h-cut tree.

Figure (4.42) illustrates this definition by showing the 1-cut $t|_1$ of the tree t above.

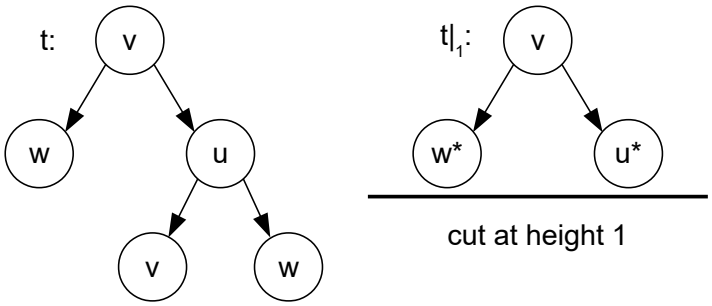

Figure (4.42): A tree t cut at height 1.

(4.43) Lemma. Let K be an absorptive lattice semiring, V a finite set of nodes, φ and ψ CTL formulas and π a matching K-interpretation over V, then for any $t \in T_v(V)$, we have

$$\pi[\![\varphi \mathrm{R} \psi]\!]_t = \inf_{h \in \omega} \pi[\![\varphi \mathrm{R} \psi]\!]_{t|_h}.$$

Proof sketch. For finite trees, this is obvious because the sequence $t|_h$ for $h \in \omega$ converges to t. For infinite trees t, we can partition the nodes and edges in t by their distance to the root, which is finite for any given node and edge, and applying partition-invariance of multiplication from proposition (2.16) and the definition of countable multiplication via infima yields the above result. □

The observation from lemma (4.43) allows us to represent the costs of infinite trees as the infimum of a sequence of costs of finite trees. Later, we will prove that the converse is also true, but for now, we return to evaluating $\mathrm{A}(\varphi \mathrm{R} \psi)$.

Fix an absorptive lattice semiring K, a finite set of nodes V, a CTL formula $\mathrm{A}(\varphi \mathrm{U} \psi)$ and a matching K-interpretation π. We would like to calculate $\pi[\![\mathrm{A}(\varphi \mathrm{R} \psi)]\!]_v$ for some $v \in V$. According to definition (3.11), this is done by computing $\mathrm{gfp}(f^{\mathrm{A}(\varphi \mathrm{R} \psi)})$. We use theorem (2.17) and start a fixed-point iteration at $X_0 = 1 \in K^V$ and set

$$\begin{aligned} X_{i+1} &= f^{\mathrm{A}(\varphi \mathrm{R} \psi)}(X_i) \quad \text{for } i \in \omega \text{ and} \\ X_\omega &= \inf_{i \in \omega} X_i. \end{aligned}$$

Just as we did for existential release formulas $\mathrm{E}(\varphi \mathrm{R} \psi)$, we will first compute X_i for any $i \in \omega$ and then compute X_ω, which is already a fixed point of $f^{\mathrm{A}(\varphi \mathrm{R} \psi)}$ as we will see later, so the iteration will end at X_ω.

(4.44) Lemma. With X_i for $i \in \omega$ defined as above, we have

$$(X_i)_v = \sum_{t \in T_v^{|i}(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_t \quad \text{for } v \in V,$$

where $T_v^{|h}$ for $h \in \omega$ refers to the set of all h-cut trees over V rooted at v.

Proof. We show this by induction on i. For $i = 0$, the only 0-cut tree over V rooted at v is $(v*)$, since the root itself has to be a marked leaf. Since $\pi[\![\varphi \mathrm{R} \psi]\!]_{(v*)} = 1 = (X_0)_v$, the hypothesis is true in this case.

For $i + 1$, we use the definition and obtain

$$\begin{aligned}(X_{i+1})_v &= f_v^{\mathrm{A}(\varphi \mathrm{R} \psi)}(X_i) \\ &= \pi[\![\psi]\!]_v \cdot \left(\pi[\![\varphi]\!]_v + \prod_{w \in vE} \pi(Evw) \cdot (X_i)_w \right) \\ &= \pi[\![\psi]\!]_v \cdot \pi[\![\varphi]\!]_v + \pi[\![\psi]\!]_v \cdot \prod_{w \in vE} \pi(Evw) \cdot (X_i)_w.\end{aligned}$$

Using the induction hypothesis yields

$$\begin{aligned}(X_{i+1})_v &= \pi[\![\psi]\!]_v \cdot \pi[\![\varphi]\!]_v + \pi[\![\psi]\!]_v \cdot \prod_{w \in vE} \pi(Evw) \cdot \left(\sum_{t \in T_w^{|i}(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_t \right) \\ &= \pi[\![\psi]\!]_v \cdot \pi[\![\varphi]\!]_v + \pi[\![\psi]\!]_v \cdot \prod_{w \in vE} \left(\sum_{t \in T_w^{|i}(V)} \pi(Evw) \cdot \pi[\![\varphi \mathrm{R} \psi]\!]_t \right).\end{aligned}$$

Let $vE = \{w_1, ..., w_l\}$ for some $1 \leq l < \omega$. This is justified by the observation that $vE \subseteq V$ is finite and non-empty, because π describes a non-terminating transition system. Then, the product of sums above can be expressed as a sum over all possible combinations of summands, which yields

$$\begin{aligned}(X_{i+1})_v &= \pi[\![\psi]\!]_v \cdot \pi[\![\varphi]\!]_v + \pi[\![\psi]\!]_v \cdot \sum_{(t_1,...,t_l) \in T_{w_1}^{|i}(V) \times ... \times T_{w_l}^{|i}(V)} \prod_{j=1}^{l} \pi(Evw_j) \cdot \pi[\![\varphi \mathrm{R} \psi]\!]_{t_j} \\ &= \pi[\![\psi]\!]_v \cdot \pi[\![\varphi]\!]_v + \sum_{(t_1,...,t_l) \in T_{w_1}^{|i}(V) \times ... \times T_{w_l}^{|i}(V)} \pi[\![\psi]\!]_v \cdot \prod_{j=1}^{l} \pi(Evw_j) \cdot \pi[\![\varphi \mathrm{R} \psi]\!]_{t_j} \\ &= \sum_{t \in T_v^{|i+1}(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_t.\end{aligned}$$

The last transformation is verified by the observation that there is a bijection between $T_{w_1}^{|i}(V) \times ... \times T_{w_l}^{|i}(V)$ and $T_v^{|i+1}(V) \setminus \{(v)\}$, where (v) is the tree that consists of only one unmarked node v. Consider a tuple $(t_1, ..., t_l) \in T_{w_1}^{|i}(V) \times ... \times T_{w_l}^{|i}(V)$, connecting v to the root nodes of t_1, ..., t_l yields a tree $t \in T_v^{|i+1}(V) \setminus \{(v)\}$ with the cost

$$\pi[\![\varphi \mathrm{R} \psi]\!]_t = \pi[\![\psi]\!]_v \cdot \prod_{j=1}^{l} \pi(Evw_j) \cdot \pi[\![\varphi \mathrm{R} \psi]\!]_{t_j}.$$

This is due to the fact that the nodes that were at height i in $t_1, ..., t_l$ are exactly the nodes at height $i + 1$ in t. Conversely, an arbitrary tree $t \in T_v^{|i+1}(V) \setminus \{(v)\}$ can be mapped back to a tuple $(t_1, ..., t_l) \in T_{w_1}^{|i}(V) \times ... \times T_{w_l}^{|i}(V)$, since $t \neq (v)$ and t is a compete tree, the root v of t has to be connected to the nodes $(w_1, ..., w_l)$ and

$(t_1, ..., t_l)$ are chosen as the subtrees rooted at $(w_1, ..., w_l)$ of t. Clearly, the nodes at height $i+1$ of t are exactly the nodes at height i in $t_1, ..., t_l$ and we have

$$\pi[\![\psi]\!]_v \cdot \prod_{j=1}^{l} \pi(Evw_j) \cdot \pi[\![\varphi \mathrm{R} \psi]\!]_{t_j} = \pi[\![\varphi \mathrm{R} \psi]\!]_t.$$

Finally, we observe that $\pi[\![\varphi \mathrm{R} \psi]\!]_{(v)} = \pi[\![\psi]\!]_v \cdot \pi[\![\varphi]\!]_v$, which completes the proof. □

To prove part (2) of theorem (4.27), it is left to show that

$$(*) \quad (X_\omega)_v = \sum_{t \in T_v(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_t \quad \text{for } v \in V$$

and that this is already a fixed point of $f^{\mathrm{A}(\varphi \mathrm{R} \psi)}$. For now, we will assume that the equation $(*)$ is true and prove it later. In that case, we can verify $f^{\mathrm{A}(\varphi \mathrm{R} \psi)}(X_\omega) = X_\omega$ component-wise. We have $f_v^{\mathrm{A}(\varphi \mathrm{R} \psi)}(X_\omega)$

$$\begin{aligned}
&= \pi[\![\psi]\!]_v \cdot \left(\pi[\![\varphi]\!]_v \cdot \prod_{w \in vE} \pi(Evw) \cdot (X_\omega)_w \right) \\
&\overset{(*)}{=} \pi[\![\psi]\!]_v \cdot \pi[\![\varphi]\!]_v + \pi[\![\psi]\!]_v \cdot \prod_{w \in vE} \pi(Evw) \cdot \left(\sum_{t \in T_w(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_t \right) \\
&= \pi[\![\psi]\!]_v \cdot \pi[\![\varphi]\!]_v + \pi[\![\psi]\!]_v \cdot \prod_{w \in vE} \left(\sum_{t \in T_w(V)} \pi(Evw) \cdot \pi[\![\varphi \mathrm{R} \psi]\!]_t \right) \\
&= \pi[\![\psi]\!]_v \cdot \pi[\![\varphi]\!]_v + \pi[\![\psi]\!]_v \cdot \sum_{(t_1, ..., t_l) \in T_{w_1}(V) \times ... \times T_{w_l}(V)} \prod_{j=1}^{l} \pi(Evw_j) \cdot \pi[\![\varphi \mathrm{R} \psi]\!]_{t_j} \\
&= \pi[\![\psi]\!]_v \cdot \pi[\![\varphi]\!]_v + \sum_{(t_1, ..., t_l) \in T_{w_1}(V) \times ... \times T_{w_l}(V)} \pi[\![\psi]\!]_v \cdot \prod_{j=1}^{l} \pi(Evw_j) \cdot \pi[\![\varphi \mathrm{R} \psi]\!]_{t_j} \\
&\overset{(1)}{=} \sum_{t \in T_v(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_t \\
&\overset{(*)}{=} (X_\omega)_v
\end{aligned}$$

for all $v \in V$. As in the proof of lemma (4.44), we assumed that $vE = \{w_1, ..., w_l\}$ for some $l \in \omega$ and the transformation (1) is verified by a straightforward bijection between $T_v(V) \setminus \{(v)\}$ and $T_{w_1}(V) \times ... \times T_{w_l}(V)$.

Assuming that $(*)$ is true, we have shown that X_ω is a fixed point of $f^{\mathrm{A}(\varphi \mathrm{R} \psi)}$, which also implies that it is the greatest fixed point, therefore, we obtain

$$\pi[\![\mathrm{A}(\varphi \mathrm{R} \psi)]\!]_v = \mathrm{gfp}(f^{\mathrm{A}(\varphi \mathrm{R} \psi)})_v = (X_\omega)_v = \sum_{t \in T_v(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_t,$$

which ends the proof for part (2) of theorem (4.27).

However, we still have to show $(*)$. First, lemma (2.21) yields

$$(X_\omega)_v = \left(\inf_{i \in \omega} X_i \right)_v = \inf_{i \in \omega} (X_i)_v.$$

Together with the results from lemma (4.44), we can transform $(*)$ to

$$\sum_{t \in T_v(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_t = \inf_{i \in \omega} \left(\sum_{t \in T_v^{|i}(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_t \right) \quad \text{for } v \in V.$$

Replacing the summation on the right side with a supremum and using complete distributivity of the order on K, we conclude that this is equivalent to

$$\sum_{t \in T_v(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_t = \sup_{f \in F} \inf_{i \in \omega} \pi[\![\varphi \mathrm{R} \psi]\!]_{f(i)},$$

where F refers to the set of choice functions $f : \omega \to T_v^{\mathrm{fin}}(V, V)$ with $f(i) \in T_v^{|i}(V)$ for all $i \in \omega$. In other words, for any i, $f(i)$ is an i-cut tree rooted at v.

The direction "$\leq$" follows immediately from lemma (4.43). For any $t \in T_v(V)$, the function f with $f(i) = t|_i \in T_v^{|i}(V)$ is in F, and we have

$$\pi[\![\varphi \mathrm{R} \psi]\!]_t = \inf_{h \in \omega} \pi[\![\varphi \mathrm{R} \psi]\!]_{t|_h} = \inf_{i \in \omega} \pi[\![\varphi \mathrm{R} \psi]\!]_{f(i)}.$$

To prove the direction "$\geq$", we will show that for any $f \in F$, we can find a $t \in T_v(V)$ with

$$\pi[\![\varphi \mathrm{R} \psi]\!]_t \geq \inf_{i \in \omega} \pi[\![\varphi \mathrm{R} \psi]\!]_{f(i)}.$$

Before we can do that, we will introduce some new definitions and concepts.

(4.45) Definition (Sequence of Trees)**.** Let K be an absorptive lattice semiring, V a finite set of nodes and π a K-interpretation over V. A *sequence* of v-trees is a function $f : \omega \to T_v^{\mathrm{fin}}(V, V)$ for some $v \in V$. f is called a *cut sequence* of v-trees if for any $i \in \omega$, $f(i)$ is a h-cut tree for some $h \geq i$. The $\varphi \mathrm{R} \psi$-cost of f for CTL formulas φ and ψ is defined as

$$\pi[\![\varphi \mathrm{R} \psi]\!]_f = \inf_{i \in \omega} \pi[\![\varphi \mathrm{R} \psi]\!]_{f(i)}.$$

We observe that F consists only of cut sequences of v-trees, since for any $f \in F$, $f(i)$ is an i-cut tree. Next, we will define costs for multisets of trees.

(4.46) Definition (Multisets of Trees)**.** Let m be a finite multiset of trees in $T^{\mathrm{fin}}(V, V)$ for some finite set of nodes V and a K-interpretation π, where K is an absorptive lattice semiring. For CTL formulas φ and ψ, we define the $\varphi \mathrm{R} \psi$-cost of m as

$$\pi[\![\varphi \mathrm{R} \psi]\!]_m = \prod_{m(t) \neq 0} \pi[\![\varphi \mathrm{R} \psi]\!]_t^{m(t)},$$

where $m(t)$ for $t \in T^{\mathrm{fin}}(V, V)$ denotes the multiplicity of t in m.

The set of multisets over $T(V, V)$ is denoted as $S(V, V)$. We will also identify trees $t \in T(V, V)$ with the multiset $\{t\}$, which is justified, because $\pi[\![\varphi \mathrm{R} \psi]\!]_t = \pi[\![\varphi \mathrm{R} \psi]\!]_{\{t\}}$.

Now, we can extend the splitting operation for trees that we have defined above.

(4.47) Definition (v-Split for Trees)**.** For a finite tree $t \in T(V, V)$ and a node $v \in V$, the *v-split* of t, denoted as $\mathrm{Sp}_v(t)$ is the multiset of trees obtained by splitting t at any occurrence of v except for roots and marked leaves.

Figure (4.48) illustrates the v-split of a tree t.

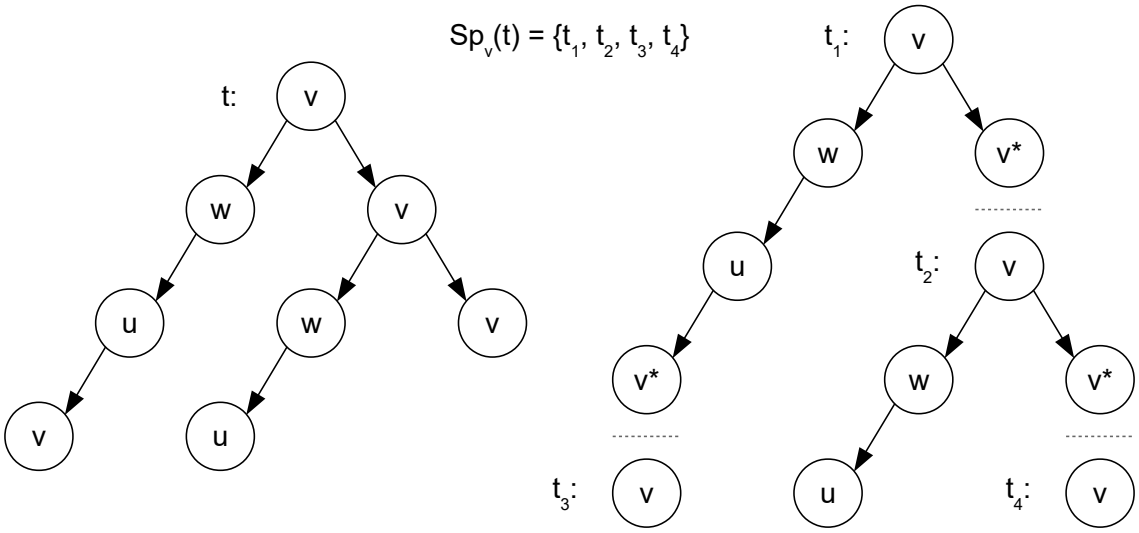

Figure (4.48): A tree t and its v-split $\mathrm{Sp}_v(t)$.

Notice that v-splitting trees preserves their costs, since single splits do not change the costs of trees as well. We have

$$\pi[\![\varphi \mathrm{R} \psi]\!]_{\mathrm{Sp}_v(t)} = \pi[\![\varphi \mathrm{R} \psi]\!]_t.$$

Also, v-splitting can be performed algorithmically, since we only split finite trees and there can only be finitely many occurrences of v. Each split removes one non-root and non-marked occurrence of v. We observe that splitting trees at their root or at marked leaves is pointless, because one of the resulting trees would be a trivial tree that only consists of a marked leaf and has cost 1. Therefore, the multiset $\mathrm{Sp}_v(t)$ depicted in figure (4.48) cannot be split any further along v-nodes.

We can also chain different split operations. Let t be a tree and $v, w \in V$. Since $\mathrm{Sp}_v(t)$ is a multiset, it would be useful to define splits on multisets. For a finite multiset of finite trees m, we define $\mathrm{Sp}_w(m)$ as the union of the w-splits of the elements in m, which is again a finite multiset of finite trees. So, the expression $\mathrm{Sp}_w \circ \mathrm{Sp}_v(t)$ is well-defined. A finite multiset multiset of finite trees m is called a *tree split* if $m = \mathrm{Sp}_{v_k} \circ ... \circ \mathrm{Sp}_{v_1}(t)$ for some $k \in \omega$, $t \in T(V)$ and $v_1, ..., v_k \in V$. Notice that k may be zero, so $\{t\}$ is also a tree split. We also call t the *original tree* of m. For any $t' \in m$, we know that t' is a subtree of t. We say that t' is rooted at height h if the root of t' in t has a distance of h to the root of t. For example, in figure (4.48), t_1 is rooted at height 0, t_2 is rooted at height 1, t_3 is rooted at height 3 and t_4 is rooted at height 2. Note that there is always exactly one $t' \in m$ that is rooted at height 0, which contains the original root of t.

Splitting operations can be applied to sequences of trees f. For $v \in V$, $\mathrm{Sp}_v(f)$ is the function $\mathrm{Sp}_v \circ f$. Note that $(\mathrm{Sp}_v(f))(i) = \mathrm{Sp}_v(f(i))$ is a multiset of trees for each $i \in \omega$. We call $g : \omega \to S(V, V)$ a sequence of v-tree splits if $g = \mathrm{Sp}_{v_k} \circ ... \circ \mathrm{Sp}_{v_1}(f)$ for some $k \in \omega$, $v_1, ..., v_k \in V$ and a sequence of v-trees f. If f was a cut sequence of v-trees, then we call g a *cut sequence of v-tree splits*. Notice that for each $i \in \omega$, $g(i) = \mathrm{Sp}_{v_k} \circ ... \circ \mathrm{Sp}_{v_1}(f(i))$ is a tree split whose original tree is $f(i)$.

Note that since we identify t with $\{t\}$ and $\{t\}$ is a tree split, in the following, any

definition that refers to tree splits or cut sequences of v-tree splits also applies to normal trees or sequences of v-trees respectively.

We can now return to proving the claim

$$\sum_{t \in T_v(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_t = \sup_{f \in F} \inf_{i \in \omega} \pi[\![\varphi \mathrm{R} \psi]\!]_{f(i)}$$

from above. Recall that F was the set of choice functions $f : \omega \to T_v^{\mathrm{fin}}(V, V)$ with $f(i) \in T_v^{\downarrow i}(V)$ for $i \in \omega$ and the direction "$\geq$" was left to show. For each $f \in F$, we have to find a $t \in T_v(V)$ such that

$$\pi[\![\varphi \mathrm{R} \psi]\!]_t \geq \pi[\![\varphi \mathrm{R} \psi]\!]_f = \inf_{i \in \omega} \pi[\![\varphi \mathrm{R} \psi]\!]_{f(i)}.$$

We will state this as a proposition.

(4.49) Proposition. Let K be an absorptive lattice semiring, V a finite set of nodes, φ and ψ CTL formulas, π a matching K-interpretation over V and $v \in V$. For every cut sequence of v-trees f, there is a $t \in T_v(V)$ such that

$$\pi[\![\varphi \mathrm{R} \psi]\!]_t \geq \pi[\![\varphi \mathrm{R} \psi]\!]_f.$$

The following pages will be dedicated to the proof of proposition (4.49). Since any $f \in F$ is a cut sequence of v-trees, this will suffice to prove the claim above. First, we define some properties of cut sequences of v-tree splits.

(4.50) Definition. Let f be a cut sequence of v-tree splits. We say that w *occurs arbitrarily often along a path* in f if for every $j \in \omega$, there is an $i \in \omega$ so that $f(i)$ contains a tree $t \in f(i)$ such that t contains at least j occurrences of the node w along a path from the root.

Figure (4.52) below provides an example of what the elements $g(0)$ to $g(5)$ of a cut sequence of v-trees g could look like.

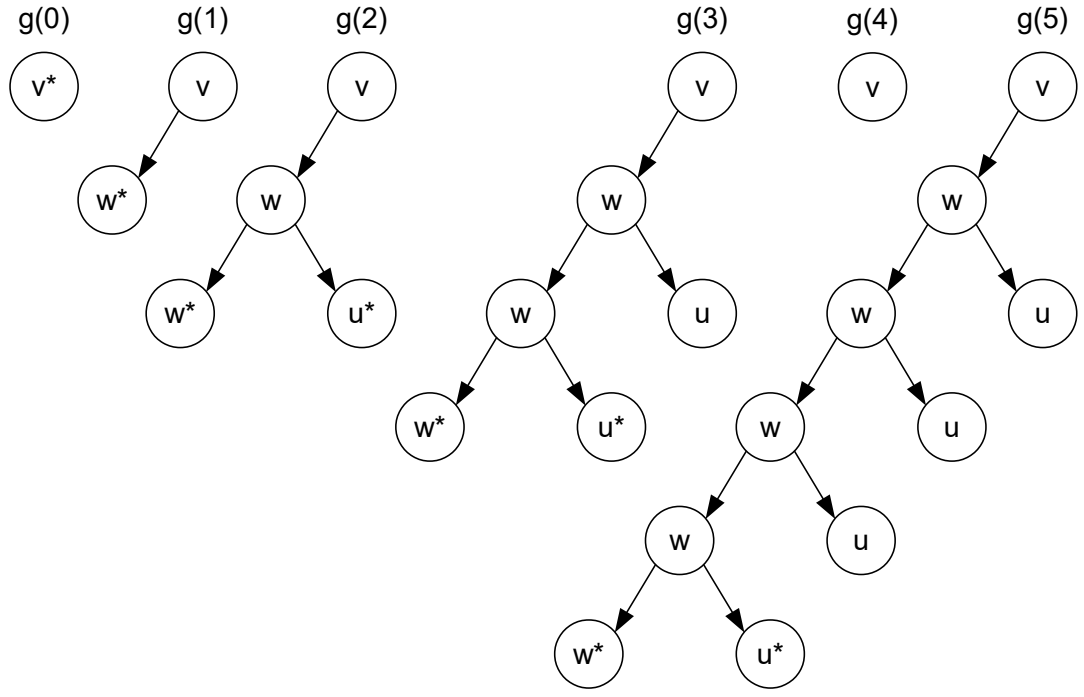

Figure (4.52): 6 elements of a cut sequence of v-trees g.

Clearly, all the trees except for $g(4)$ follow the pattern of appending the subtree $w(w*, u*)$ to the w-leaf. Imagine that the remaining elements $g(6), g(7), ...$ follow the same pattern. In that case, w occurs arbitrarily often along a path in g, since for every $j \in \omega \setminus \{4\}$, $g(j)$ contains a path starting at the root that contains j occurrences of w. The exception of $g(4)$ is not a problem, since for $j = 4$, we have the element $g(5)$ that contains at least 4 occurrences of w along a path from the root. However, notice that even though the node u occurs arbitrarily often in the trees $g(i)$, u does *not* occur arbitrarily often along a path in g. In fact, any path starting at the root in any tree $g(i)$ contains at most one occurrence of u.

Returning to the general case, if w occurs arbitrarily often in a cut sequence of v-tree splits f, we can "clean up" the sequence by only retaining those elements that actually contain the desired occurrences of w. This yields the following definition.

(4.51) Definition. Let f be a cut sequence of v-tree splits and w a node that occurs arbitrarily often along a path in f. We define the sequence $\mathrm{Cl}_w(f)$ for $j \in \omega$ by setting

$$(\mathrm{Cl}_w(f))(j) = f(i)$$

for the smallest $i \geq j$ such that there is a $t \in f(i)$ that contains at least j unmarked occurrences of w along a path from the root.

Obviously, if w did not occur arbitrarily often in f, then $\mathrm{Cl}_w(f)$ would not be well-defined. However, if w does indeed occur arbitrarily often in f, then we know that for each $(j + 1) \in \omega$, we have an $i' \in \omega$ such that $f(i')$ contains a tree t with at least $(j + 1)$ occurrences of w along a path from the root. Only the last of those occurrences can be marked, so t contains at least j unmarked occurrences of w along a path from the root. Additionally, we know that we can find an $i \geq j$ with this property, because otherwise, only the elements $f(0), ...f(j - 1)$ of f would contain at least j occurrences of w along a path from the root, but since $f(0), ..., f(j - 1)$ are finitely many finite multisets of finite trees, this would be a contradiction, since it would imply that the number of occurrences of w along a path from the root is limited to

$$\max\{m_l \mid 0 \leq l < j\},$$

where m_l is the maximal number of occurrences of w along a path from the root in any tree in $f(l)$, because we would know that none of the elements $f(j), f(j+1), ...$ could contain any more than j occurrences of w along a path from the root. This would contradict the assertion that w occurs arbitrarily often in f.

Since any element $(\mathrm{Cl}_w(f))(j)$ for $j \in \omega$ of $\mathrm{Cl}_w(f)$ is equal to $f(i)$ for some $i \geq j$, we can derive

$$\pi[\![\varphi \mathrm{R} \psi]\!]_{\mathrm{Cl}_w(f)} = \inf_{j \in \omega} \pi[\![\varphi \mathrm{R} \psi]\!]_{(\mathrm{Cl}_w(f))(j)} \geq \inf_{i \in \omega} \pi[\![\varphi \mathrm{R} \psi]\!]_{f(i)} = \pi[\![\varphi \mathrm{R} \psi]\!]_f.$$

We conclude that cleaning up a sequence increases its costs. Also, $\mathrm{Cl}_w(f)$ is still a cut sequence of v-tree splits, because for each $j \in \omega$, $(\mathrm{Cl}_w(f))(j) = g(i)$ is a tree split of a t tree that is h-cut for some $h \geq i \geq j$.

Finally, we provide an example for cleaning up sequences. Consider $\mathrm{Cl}_w(g)$ for the

example from figure (4.52) above. We have

$$\begin{aligned}(\mathrm{Cl}_w(g))(0) &= g(0),\\ (\mathrm{Cl}_w(g))(1) &= g(2),\\ (\mathrm{Cl}_w(g))(2) &= g(3),\\ (\mathrm{Cl}_w(g))(3) &= g(5) \quad \text{and}\\ (\mathrm{Cl}_w(g))(4) &= g(5).\end{aligned}$$

Continuing the pattern, we would obtain $(\mathrm{Cl}_w(g))(j) = g(j+1)$ for $j \geq 5$, the index shift of 1 is due to the assertion that the occurrences of w in the cleaned sequence must be unmarked and the irregularity at $(\mathrm{Cl}_w(g))(3)$ is caused by the irregularity of $g(4)$, which has to be skipped since it does not contain any w-node.

Having defined Cl_w, recall definition (4.47), where we have defined the w-split Sp_w. Chaining these two operations together defines the w-reduction $\mathrm{R}_w(f)$ of a sequence. If w occurs arbitrarily often along a path in g, we set

$$\mathrm{R}_w(f) = \mathrm{Sp}_w \circ \mathrm{Cl}_w(f).$$

Figure (4.53) below illustrates the first 3 values of $\mathrm{R}_w(g)$ for the example of g provided above in figure (4.52).

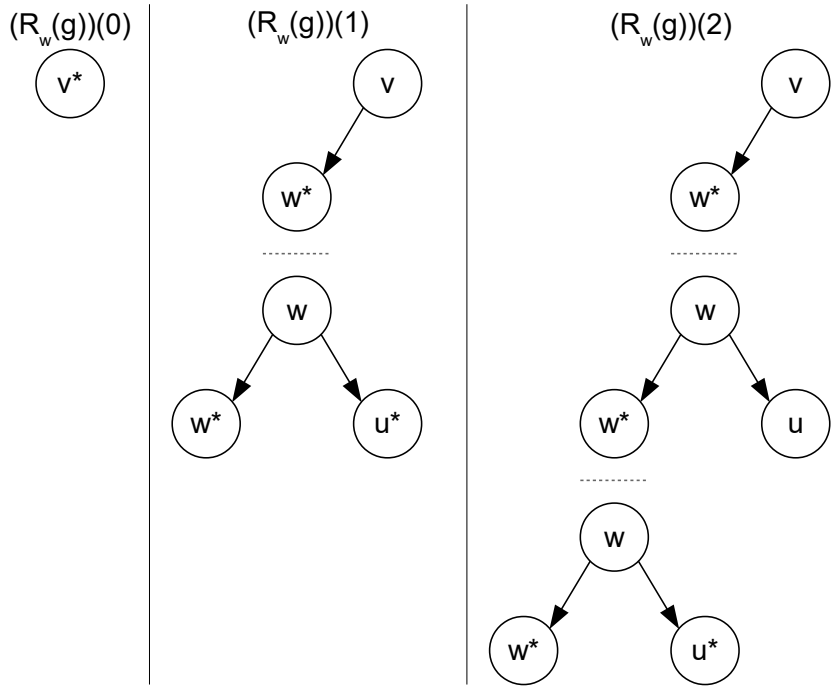

Figure (4.53): Values of $\mathrm{R}_w(g)$ for the sequence g in the above example.

Notice that for each $j \in \omega$, $(\mathrm{R}_w(g))(j)$ contains at least j non-trivial trees with root w, each of them rooted at a different height. This is easily explained by the fact that $(\mathrm{Cl}_w(g))(j)$ contains a tree where a path from the root contains at least j unmarked occurrences of w, when we w-split this tree, each of those occurrences of w forms the root of a tree in $(\mathrm{R}_w(g))(j) = (\mathrm{Sp}_w \circ \mathrm{Cl}_w(g))(j)$ and all of these trees are non-trivial and rooted at different heights.

The reason why $\mathrm{R}_w(g)$ is called the w-reduction of g is that we have effectively removed w from g, since in $\mathrm{R}_w(g)$, the node w may only occur as a root node or a

marked leaf, but due to the w-splitting, other unmarked occurrences of w outside of root nodes are not permitted. All of these observations are generally valid for any cut sequence of v-tree splits f.

Additionally, we can see in the above example that no node occurs arbitrarily often along a path in $\mathrm{R}_w(g)$ anymore. Unfortunately, this is not generally true. In figure (4.54), a new cut sequence of v-trees h is given to illustrate this issue. Again, we have only depicted the values $h(1)$ to $h(3)$ of h.

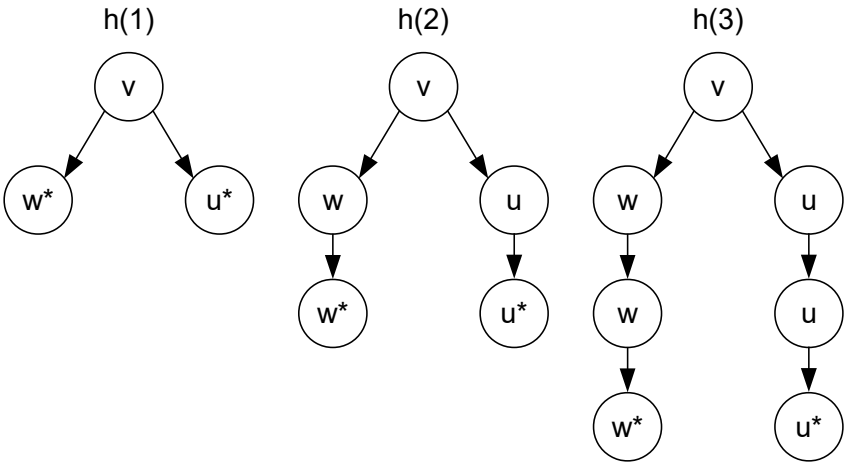

Figure (4.54): 3 elements of a cut sequence of v-trees h.

Assuming that the same pattern continues for all $h(i)$ with $i \in \omega$, we can clearly see that u occurs arbitrarily often along a path in $\mathrm{R}_w(h)$, because even after w-splitting, the path $(u, u, ..., u*)$ stays in the same subtree. However, we can overcome this issue by applying the u-reduction to $\mathrm{R}_w(h)$ to obtain $\mathrm{R}_u \circ \mathrm{R}_w(h)$. Figure (4.55) below shows $(\mathrm{R}_u \circ \mathrm{R}_w(h))(1)$ and $(\mathrm{R}_u \circ \mathrm{R}_w(h))(2)$ for the sequence h from figure (4.54).

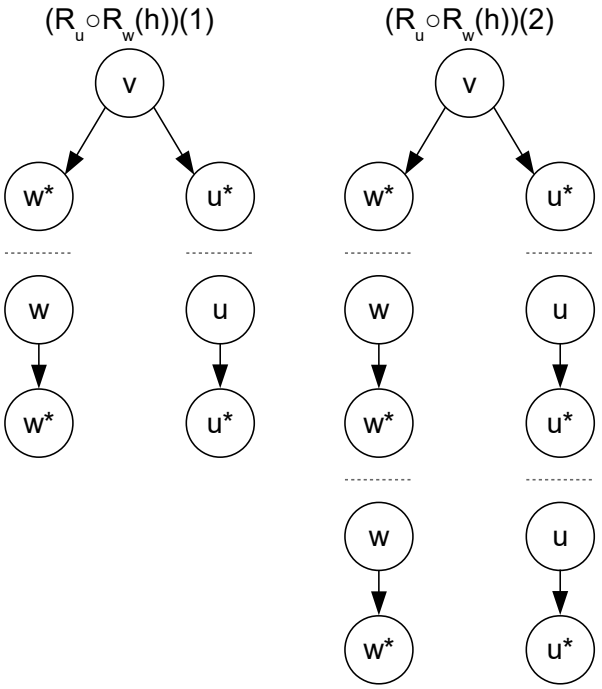

Figure (4.55): Values of $\mathrm{R}_u \circ \mathrm{R}_w(h)$ for h from figure (4.54).

This resolves the issue for the example from figure (4.54), but in the general case, if f is a cut sequence of v-tree splits, there may still be different nodes from w and u that occur arbitrarily often along a path in $\mathrm{R}_u \circ \mathrm{R}_w(f)$. Therefore, we will inductively apply reductions to f until there is no node left that occurs arbitrarily often in the resulting sequence. First, we fix an arbitrary linear order $<$ on V. This enables us to chain reductions.

(4.56) Definition. Let $W \subseteq V$. The W-reduction $\mathrm{R}_W(f)$ for cut sequences of v-tree splits f is defined inductively over the cardinality of W.

For $W = \emptyset$, $\mathrm{R}_\emptyset(f) = f$, so the $\emptyset$-reduction does nothing.

If $|W| = k + 1$ for some $k \in \omega$, let $w = \max_<(W)$, then $W = W' \cup \{w\}$ with $|W'| = k$. If $\mathrm{R}_{W'}(f)$ exists and w occurs arbitrarily often along a path in $\mathrm{R}_{W'}(f)$, then define $\mathrm{R}_W(f)$ as $\mathrm{R}_W(f) = \mathrm{R}_w \circ \mathrm{R}_{W'}(f)$. Otherwise, $\mathrm{R}_W(f)$ does not exist.

As an example, if $w_1 < w_2 < w_3$ are elements of V and $W = \{w_1, w_2, w_3\}$, we have

$$\mathrm{R}_W(g) = \mathrm{R}_{w_3} \circ \mathrm{R}_{w_2} \circ \mathrm{R}_{w_1}(g),$$

provided that $\mathrm{R}_W(g)$ exists at all. Now, we will prove a powerful lemma for $\mathrm{R}_W(f)$.

(4.57) Lemma. Let f be a cut sequence of v-tree splits. With $\mathrm{R}_W(f)$ defined as above, we claim that if $\mathrm{R}_W(f)$ exists, then the following conditions are met for all $j \in \omega$:

1. The trees in $(\mathrm{R}_W(f))(j)$ may only contain marked leaves in W except for leaves at the maximum cut-off height of the original tree,
2. there is at least one tree in $(\mathrm{R}_W(f))(j)$ rooted at v at height 0 and
3. for every $w \in W$, there are at least j non-trivial trees in $(\mathrm{R}_W(f))(j)$ with the root w, each of them is rooted at a different height in the original tree.

Proof. The conditions (1.) and (2.) are met, because R_W is a chain of w-splitting and cleaning operations for $w \in W$. Therefore, each $(\mathrm{R}_W(f))(j)$ is a tree split and the original tree is a h-cut tree $t \in T_v(V, V)$ for some $h \geq j$. Since a h-cut tree does not contain any marked leaves other than the leaves at height h, the trees in $(\mathrm{R}_W(f))(j)$ also do not have any other marked leaves than those at height h and the marked leaves induced by w-splitting for $w \in W$, which proves condition (1.). Also, since t is rooted at v, after splitting t, at least one of the resulting trees in $(\mathrm{R}_W(f))(j)$ retains the root v at height 0, so condition (2.) is true as well.

We prove condition (3.) by induction over the cardinality of W. For $W = \emptyset$, there is nothing to show, so we assume $|W| = k+1$ for a $k \in \omega$. Again, we pick $u = \max_<(W)$ and obtain $W = W' \cup \{u\}$ with $u \notin W'$. Since we assume that $\mathrm{R}_W(f)$ exists, by definition of R_W, $g \coloneqq \mathrm{R}_{W'}(f)$ must exist as well and we have $\mathrm{R}_W(f) = \mathrm{R}_u(g)$. We have to show condition (3.) for all $w \in W'$ and for u.

If $w \in W'$, we have $(\mathrm{R}_W(f))(j) = \mathrm{Sp}_u(g(i))$ for some $i \geq j$. By induction, we know that $g(i)$ contained at least i non-trivial trees with the root w and each of them was rooted at a different height. Clearly, after u-splitting $g(i)$, there are still at least $i \geq j$ non-trivial trees with the root w at different heights in $(\mathrm{R}_W(f))(j)$.

For u, we know that $(\mathrm{R}_W(f))(j) = \mathrm{Sp}_u(\mathrm{Cl}_u(g(j))$ and $\mathrm{Cl}_u(g(j))$ contains a tree t with at least j unmarked u-nodes along a path from the root. Each of these nodes is at a different height and after u-splitting, each of them becomes a root of its own non-trivial subtree, which ends the proof. □

Notice that in order to use lemma (4.57) for some cut sequence of v-tree splits f and $W \subseteq V$, we first need to make sure that $\mathrm{R}_W(f)$ exists. We claim that $\mathrm{R}_W(f)$ exists if W *occurs arbitrarily often along a path* in g. This notion is defined inductively over the cardinality of W.

For $W = \emptyset$, $\emptyset$ occurs arbitrarily often along a path in f for any cut sequence of v-tree splits f. Notice that $\mathrm{R}_\emptyset(f)$ always exists.

For $|W| = k + 1$ and $k \in \omega$, pick $w = \max_<(W)$, so that $W = W' \cup \{w\}$ with $w \notin W'$. We say that W occurs arbitrarily often along a path in f if W' occurs arbitrarily often along a path in f and w is the *minimal* element of V such that w occurs arbitrarily often along a path in $\mathrm{R}_{W'}(f)$. Notice that $\mathrm{R}_{W'}(f)$ exists by induction, since $|W'| = k$. Also, this implies that $\mathrm{R}_W(f)$ exists.

The minimality condition for w allows us to prove the statement that for any cut sequence of v-tree splits f, if W occurs arbitrarily often along a path in f and u occurs arbitrarily often along a path in $\mathrm{R}_W(f)$, then $u > w$ for all $w \in W$. We prove this by induction.

For $W = \emptyset$, there is nothing to show. Suppose now that $W = k + 1$ for a $k \in \omega$, set $m = \max_<(W)$ and $W = W' \cup \{m\}$ with $m \notin W'$. Additionally, suppose that u occurs arbitrarily often along a path in $\mathrm{R}_W(f)$. We now have to show that $u > m$, thereby showing that $u > w$ for all $w \in W$. Since W' occurs arbitrarily often along a path in f and u occurring arbitrarily often along a path in $\mathrm{R}_W(f)$ implies that it also occurs arbitrarily often along a path in $\mathrm{R}_{W'}(f)$, we conclude by induction that $u > w$ for all $w \in W$. Now, if $u < m$, then m would not be the minimal element that occurs arbitrarily often along a path in $\mathrm{R}_{W'}(f)$, which is a contradiction. Also, $u = m$ is impossible, because m does not occur arbitrarily often along a path in $\mathrm{R}_W(f)$. Therefore, we have $u > m$ and the claim is proven.

Now, we can prove proposition (4.49). Let f be an arbitrary cut sequence of v-trees. The goal is to find a tree $t \in T_v(V)$ such that $\pi[\![\varphi \mathrm{R} \psi]\!]_t \geq \pi[\![\varphi \mathrm{R} \psi]\!]_f$. First, we pick the maximal set $W \subseteq V$ such that W occurs arbitrarily often along a path in f. Such a set always exists, even if no node occurs arbitrarily often along a path in f, then we would have $W = \emptyset$.

We can infer that $\mathrm{R}_W(f)$ exists and we have

$$\pi[\![\varphi \mathrm{R} \psi]\!]_{\mathrm{R}_W(f)} \geq \pi[\![\varphi \mathrm{R} \psi]\!]_f,$$

because R_W is a chain of splitting and cleaning operations, and we have shown that splitting operations preserve $\varphi \mathrm{R} \psi$-costs and cleaning operations never decrease $\varphi \mathrm{R} \psi$-costs of sequences.

Additionally, we know that no node $u \in V$ occurs arbitrarily often along a path in $\mathrm{R}_W(f)$. Otherwise, we would have $u > w$ for all $w \in W$, but this would contradict the maximality of W, since in that case, $W \cup \{u'\}$ would occur arbitrarily often along a path in f for some $u' \leq u$ with $u' > w$ for all $w \in W$.

Therefore, for every $u \in V$, there is a $h_u \in \omega$ such that no tree in $\mathrm{R}_W(f)$ contains

more than h_u occurrences of u along a single path from the root. We can derive that

$$H = \sum_{u \in V} h_u$$

is an upper bound on the height of the trees in $\mathrm{R}_W(f)$. Now, consider the sequence g with $g(i) = (\mathrm{R}_W(f))(i + H + 1)$. By condition (3.) of lemma (4.57), we conclude that for each $w \in W$, $g(i)$ contains at least $i + H + 1$ non-trivial trees rooted at w at different heights. Also, since the height of those trees is bounded by H, at least k of those trees are rooted so far away from the cut-off height of the original tree that none of their leaves are cut off. This means by condition (1.) of lemma (4.57) that they only contain marked leaves in W. Now, let $T_w^{\leq H}(V, W)$ denote the set of all non-trivial trees over V rooted at w with marked leaves in W. Since the height is bounded, this set is finite.

Assume w.l.o.g. that $W = \{w_1, ..., w_k\}$ for some $k \in \omega$. The above argument yields that $g(i)$ contains at least i trees in $T_{w_j}^{\leq H}(V, W)$ for each w_j. If we group those trees to tuples, we can rephrase this and claim that $g(i)$ contains at least i tuples

$$(t_1, ..., t_k) \in T_{w_1}^{\leq H}(V, W) \times ... \times T_{w_k}^{\leq H}(V, W).$$

Since there are only finitely many such tuples, there must be one tuple $(t_1, ..., t_k) \in T_{w_1}^{\leq H}(V, W) \times ... \times T_{w_k}^{\leq H}(V, W)$ that occurs arbitrarily often in g, that is, for each $j \in \omega$, there is an $i \in \omega$ such that $g(i)$ contains $(t_1, ..., t_k)$ at least j times. We fix such a tuple $(t_1, ..., t_k)$ and define a new sequence g' with $g'(j) = g(i)$ for the smallest i such that $g(i)$ contains $(t_1, ..., t_k)$ at least j times.

Finally, we use condition (2.) of lemma (4.57), which states that each $g'(j)$ contains a tree t that is rooted at v at the height 0. Since $g'(j) = g(i) = (\mathrm{R}_W(f))(i + H + 1)$ for some $i \in \omega$, none of the leaves of t can be cut off, since t is rooted at 0 and its height is bounded by H. Therefore, the marked leaves of t can only be in W by condition (1.) of lemma (4.57), which implies $t \in T_v^{\leq H}(V, W)$. Again, $T_v^{\leq H}(V, W)$ being finite allows us to conclude that there is a $t_0 \in T_v^{\leq H}(V, W)$ that is contained in infinitely many members $g'(j)$ of g'. We fix such a t_0 and define the sequence g'' by $g''(l) = g'(j)$ for the smallest j such that $g'(j)$ contains at least l occurrences of $(t_1, ..., t_k)$ and t_0. Notice that our choice of t_0 guarantees the existence of such a j for each $l \in \omega$.

Recall that all elements of g'' are contained in g', all elements of g' are contained in g and all elements of g are contained in $\mathrm{R}_W(f)$, so we have

$$\pi[\![\varphi \mathrm{R} \psi]\!]_{g''} \geq \pi[\![\varphi \mathrm{R} \psi]\!]_{g'} \geq \pi[\![\varphi \mathrm{R} \psi]\!]_{g} \geq \pi[\![\varphi \mathrm{R} \psi]\!]_{\mathrm{R}_W(f)} \geq \pi[\![\varphi \mathrm{R} \psi]\!]_{f}.$$

We have also fixed a tuple

$$(t_0, t_1, ..., t_k) \in T_v^{\leq H}(V, W) \times T_{w_1}^{\leq H}(V, W) \times ... \times T_{w_k}^{\leq H}(V, W).$$

Now, consider the tree

$$t = t_0 \circ t_1^\infty \circ ... \circ t_k^\infty.$$

This tree is constructed by starting at t_0 and appending the trees $t_1, ..., t_k$ to any marked w_j-leaves whenever they are needed. In other words, whenever there is a marked w_j-leaf for $1 \leq j \leq k$, we append t_j, which is rooted at w_j, to this leaf.

We have $t \in T_v(V)$, since t_0 is rooted at v. Also, t has no marked leaves, because the trees $t_0, t_1, ..., t_k$ only have marked leaves in W, but each marked leaf $w_j \in W$ is eliminated by appending t_j. Of course, t is generally an infinite tree, because we have to append $t_1, ..., t_k$ infinitely many times, but since $t_1, ..., t_k$ are non-trivial trees, this is not a problem.

We claim that

$$\pi[\![\varphi \mathrm{R} \psi]\!]_t = \inf_{h \in \omega} \pi[\![\varphi \mathrm{R} \psi]\!]_{t|_h} \geq \inf_{l \in \omega} \pi[\![\varphi \mathrm{R} \psi]\!]_{g''(l)} = \pi[\![\varphi \mathrm{R} \psi]\!]_{g''}.$$

For an arbitrary $h \in \omega$, $t|_h$ contains t_0 and at most l full or partial occurrences of $t_1, ..., t_k$ for some $l \in \omega$, since $t|_h$ is finite. Therefore, we have

$$\pi[\![\varphi \mathrm{R} \psi]\!]_{t|_h} \geq \pi[\![\varphi \mathrm{R} \psi]\!]_{t_0} \cdot \pi[\![\varphi \mathrm{R} \psi]\!]_{t_1}^l \cdot ... \cdot \pi[\![\varphi \mathrm{R} \psi]\!]_{t_k}^l.$$

Also, since $g''(l)$ contains t_0 and at least l occurrences of $(t_1, ..., t_k)$, we have

$$\begin{aligned} \pi[\![\varphi \mathrm{R} \psi]\!]_{t|_h} &\geq \pi[\![\varphi \mathrm{R} \psi]\!]_{t_0} \cdot \pi[\![\varphi \mathrm{R} \psi]\!]_{t_1}^l \cdot ... \cdot \pi[\![\varphi \mathrm{R} \psi]\!]_{t_k}^l \\ &\geq \pi[\![\varphi \mathrm{R} \psi]\!]_{g''(l)} \\ &\geq \inf_{l \in \omega} \pi[\![\varphi \mathrm{R} \psi]\!]_{g''(l)}, \end{aligned}$$

which proves the claim, because $h \in \omega$ was arbitrary.

This implies that

$$\pi[\![\varphi \mathrm{R} \psi]\!]_t \geq \pi[\![\varphi \mathrm{R} \psi]\!]_f$$

and ends the proof for proposition (4.49), and with that, theorem (4.27) is also proven and we have

$$\pi[\![\mathrm{A}(\varphi \mathrm{R} \psi)]\!]_v = \sum_{t \in T_v(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_t.$$

This raises the question of how to compute $\pi[\![\mathrm{A}(\varphi \mathrm{R} \psi)]\!]_v$. For that, we can obtain a useful corollary from the proof above. Rather than just showing that any cut sequence of v-tree splits is absorbed by a tree $t \in T_v(V)$, we have shown that for every cut sequence of v-tree splits f, there is a $W = \{w_1, ..., w_k\} \subseteq V$, and a tuple

$$(t_0, t_1, ..., t_k) \in T_v^{\leq H}(V, W) \times T_{w_1}^{\leq H}(V, W) \times ... \times T_{w_k}^{\leq H}(V, W)$$

for some $H \in \omega$ such that $t_0 \circ t_1^\infty \circ ... \circ t_k^\infty$ absorbs f. It is important to notice that $t_0, t_1, ..., t_k$ are thereby all finite.

For any $w \in V$, let $T'_w(V, W)$ be the set of all trees over V without any node repetitions along a path from the root that are rooted at w and may only have marked leaves in W. Further, let $T_w^{\mathrm{CYC}}(V, W)$ for $w \in W$ be the set of all non-trivial trees over V without any node repetitions along a path from the root, with the exception that the node w appears in the root and may also appear as a marked leaf and, like above, marked leaves must be in W.

We observe that node repetitions along paths from the root can be eliminated as shown in figure (4.58) below, which shows a tree t with a node w repeating along a path from the root and the reduced tree t' obtained by eliminating the repetition.

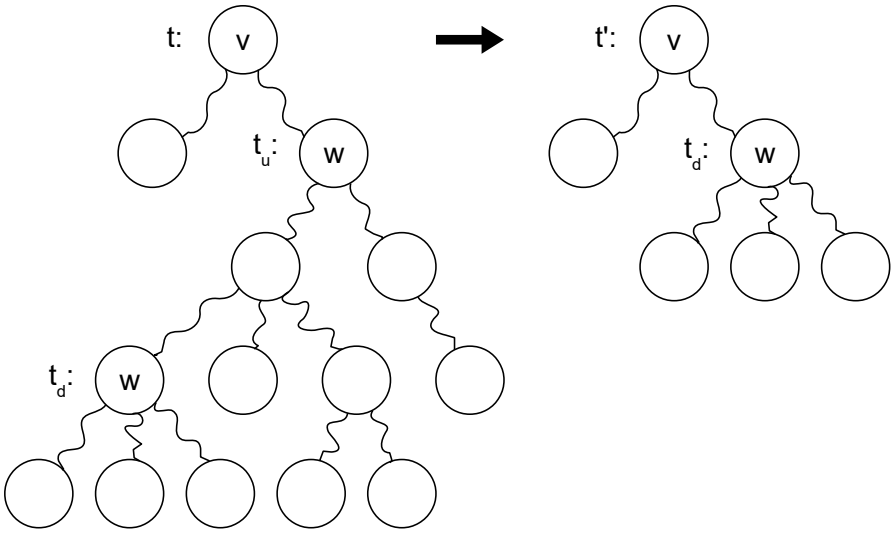

Figure (4.58): Elimination of a w-repetition in a tree t.

As shown in the picture, the subtree t_u of t rooted at the occurrence of w that is closer to the root simply has to be replaced with the subtree t_d that is rooted at the other occurrence of w. Clearly, we have $n(t') \subseteq n(t)$, $e(t') \subseteq e(t)$ and $l(t') \subseteq l(t)$, which implies that

$$\pi[\![\varphi \mathrm{R} \psi]\!]_{t'} \geq \pi[\![\varphi \mathrm{R} \psi]\!]_t.$$

This is a general result. Whenever there are node repetitions along a path from the root of a tree, we can simply eliminate them and obtain a tree with a higher $\varphi \mathrm{R} \psi$-cost. This procedure can be applied to the tuple $(t_0, t_1, ..., t_k)$ from above. For t_0, we remove all repetitions and since t_0 is finite, we obtain a tree $t_0' \in T_v'(V, W)$ with $\pi[\![\varphi \mathrm{R} \psi]\!]_{t_0'} \geq \pi[\![\varphi \mathrm{R} \psi]\!]_{t_0}$. For $t_1, ..., t_k$, we also remove all repetitions, except for repetitions where w is the root node and appears as a marked leaf at the same time, so we obtain

$$(t_1', ..., t_k') \in T_{w_1}^{\mathrm{CYC}}(V, W) \times ... \times T_{w_k}^{\mathrm{CYC}}(V, W)$$

with $\pi[\![\varphi \mathrm{R} \psi]\!]_{t_j'} \geq \pi[\![\varphi \mathrm{R} \psi]\!]_{t_j}$ for $1 \leq j \leq k$. It is important to note that the trees $t_1', ..., t_k'$ are not trivial. Therefore, we can use $t_0', t_1', ..., t_k'$ to build the new tree $t_0' \circ (t_1')^\infty \circ ... \circ (t_k')^\infty$, which is a well-defined tree in $T_v(V)$, using the same argument as for $t_0 \circ t_1^\infty \circ ... \circ t_k^\infty$. Clearly, we have

$$\pi[\![\varphi \mathrm{R} \psi]\!]_{t_0' \circ (t_1')^\infty \circ ... \circ (t_k')^\infty} \geq \pi[\![\varphi \mathrm{R} \psi]\!]_{t_0 \circ t_1^\infty \circ ... \circ t_k^\infty} \geq \pi[\![\varphi \mathrm{R} \psi]\!]_f.$$

For any $v \in V$ and $W = \{w_1, ..., w_k\} \subseteq V$, we define

$$\mathrm{Tup}_v(W) = T_v'(V, W) \times T_{w_1}^{\mathrm{CYC}}(V, W) \times ... \times T_{w_k}^{\mathrm{CYC}}(V, W).$$

Also, let F be the set of all cut sequences of v-trees f. As a corollary from the proof above, we obtain

$$\sup_{f \in F} \pi[\![\varphi \mathrm{R} \psi]\!]_f \leq \sum_{W \subseteq V} \sum_{(t_0, t_1, ..., t_k) \in \mathrm{Tup}_v(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_{t_0 \circ t_1^\infty \circ ... \circ t_k^\infty} \leq \sum_{t \in T_v(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_t.$$

Since we already know that

$$\sum_{t \in T_v(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_t \leq \sup_{f \in F} \pi[\![\varphi \mathrm{R} \psi]\!]_f$$

because of the sequence $f \in F$ with $f(i) = t|_i$ which can be constructed for any t, we conclude that all the values above are equal and obtain the following corollary.

(4.59) Corollary. With the definitions from above, we have

$$\pi[\![\mathrm{A}(\varphi \mathrm{R} \psi)]\!]_v = \sum_{t \in T_v(V)} \pi[\![\varphi \mathrm{R} \psi]\!]_t = \sum_{W \subseteq V} \sum_{(t_0, t_1, ..., t_k) \in \mathrm{Tup}_v(W)} \pi[\![\varphi \mathrm{R} \psi]\!]_{t_0 \circ t_1^\infty \circ ... \circ t_k^\infty}.$$

However, this is still not sufficient to compute $\pi[\![\mathrm{A}(\varphi \mathrm{R} \psi)]\!]_v$, even though $\mathrm{Tup}_v(W)$ is a finite set for each $W \subseteq V$. This is due to the fact that in general, we have

$$\pi[\![\varphi \mathrm{R} \psi]\!]_{t_0 \circ t_1^\infty \circ ... \circ t_k^\infty} \neq \pi[\![\varphi \mathrm{R} \psi]\!]_{t_0} \cdot \pi[\![\varphi \mathrm{R} \psi]\!]_{t_1}^\infty \cdot ... \cdot \pi[\![\varphi \mathrm{R} \psi]\!]_{t_k}^\infty$$

for $(t_0, t_1, ..., t_k) \in \mathrm{Tup}_v(W)$. Consider a case where $k > 0$. Since we have $t_0 \in T'_v(V, W)$, we know that t_0 *might* contain marked leaves in W. However, this is not necessarily the case. For example, we might have $t_0 = (v)$, so that t_0 is a tree that consists of a single unmarked node. In the definition of $t_0 \circ t_1^\infty \circ ... \circ t_k^\infty$, we have stated that the trees $t_1, ..., t_k$ are only used if they are needed, that is, t_j is only used to fill a marked leaf w_j. So in this case, we would obtain

$$t_0 \circ t_1^\infty \circ ... \circ t_k^\infty = t_0,$$

which implies

$$\pi[\![\varphi \mathrm{R} \psi]\!]_{t_0 \circ t_1^\infty \circ ... \circ t_k^\infty} = \pi[\![\varphi \mathrm{R} \psi]\!]_{t_0}.$$

Obviously, for $k > 0$, this witnesses the inequality from above. Therefore, we will have to find a way to ensure that the trees $t_1, ..., t_k$ are actually used infinitely many times in $t_0 \circ t_1^\infty \circ ... \circ t_k^\infty$.

(4.60) Definition (Dependency Graph)**.** Let $t = (t_0, t_1, ..., t_k) \in \mathrm{Tup}_v(W)$ for some $v \in V$ and $W \subseteq V$. We define the *dependency graph* $D(t)$ as follows. The nodes of $D(t)$ are $t_0, t_1, ..., t_k$. There is a directed edge from t_i to t_j for $0 \leq i \leq k$ and $1 \leq j \leq k$ if, and only if t_i contains at least one marked leaf w_j. The node t_0 has no incoming edges.

The dependency graph can be used to determine whether all elements of a tuple $t \in \mathrm{Tup}_v(W)$ are actually necessary to construct a tree without marked leaves. The first observation that we can make is the following.

(4.61) Lemma. For any $t = (t_0, t_1, ..., t_k) \in \mathrm{Tup}_v(W)$, if there is at least one t_i in $D(t)$ that is not reachable from t_0, then there is a tuple $t' = (t_0, t'_1, ..., t'_{k'}) \in \mathrm{Tup}_v(W')$ for some $W' \subsetneq W$ such that every node in $D(t')$ is reachable from t_0 and

$$\pi[\![\varphi \mathrm{R} \psi]\!]_{t_0 \circ t_1^\infty \circ ... \circ t_k^\infty} = \pi[\![\varphi \mathrm{R} \psi]\!]_{t_0 \circ (t'_1)^\infty \circ ... \circ (t'_{k'})^\infty}.$$

Proof sketch. This is very easy to see, since trees that are not reachable from t_0 in $D(t)$ are never actually used in the construction of $t_0 \circ t_1^\infty \circ ... \circ t_k^\infty$, so we obtain $t' = (t_0, t'_1, ..., t'_{k'})$ by leaving out all the trees that are unreachable from t_0 in $D(t)$. Clearly, $t' \in \mathrm{Tup}_v(W')$ for a smaller set W' and we know that the trees in t' do not contain any marked leaves outside of W', because otherwise, some of the trees that we left out would have been reachable from t_0 in $D(t)$. □

The lemma allows us to filter out some of the "bad" tuples, but even if all trees from t are reachable in $D(t)$, this does not guarantee that they have to be used

infinitely often. Figure (4.62) illustrates the dependency graphs $D(p)$ and $D(q)$ for two different tuples $p = (p_0, p_1, p_2, p_3)$ and $q = (q_0, q_1, q_2, q_3)$ in $\mathrm{Tup}_v(W)$.

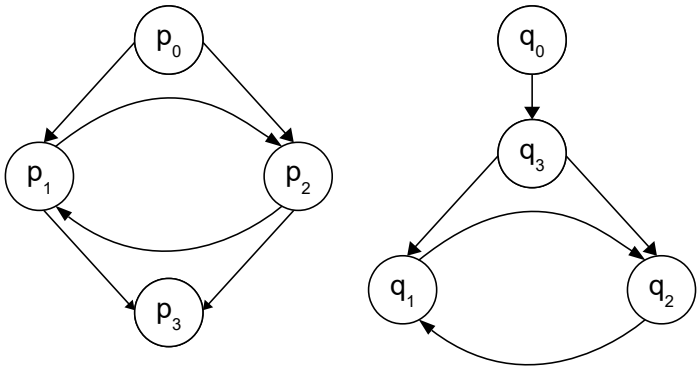

Figure (4.62): Dependency graphs of $p, q \in \mathrm{Tup}_v(W)$.

We can derive from the dependency graphs that p is a "good" tuple, because p_1, p_2 and p_3 are used infinitely often in $t_p = p_0 \circ p_1^\infty \circ p_2^\infty \circ p_3^\infty$. This is due to the cycle (p_1, p_2). Whenever p_1 is appended, p_2 has to be appended to its marked w_2-leaves and vice versa, leading to the conclusion that both p_1 and p_2 occur infinitely often as subtrees of t_p. Since any occurrence of p_1 or p_2 also requires an occurrence of p_3, because p_1 and p_2 have marked w_3-leaves, p_3 also occurs infinitely often in t_p.

However, the situation is different in $t_q = q_0 \circ q_1^\infty \circ q_2^\infty \circ q_3^\infty$. With the same argument as for t_p, q_1 and q_2 occur infinitely often in t_q. However, neither q_1 nor q_2 have any marked w_3-leaves, so q_3 is only needed to fill the marked w_3-leaves of q_0, but there are only finitely many of them. Since none of the other trees contains any w_3-leaves, q_3 occurs only finitely often in t_q. We conclude that q is a bad tuple, but we can also see that it is possible to turn it into a good tuple $r = (r_0, q_1, q_2) \in \mathrm{Tup}_v(W')$ with $W' = W \setminus \{w_3\}$. r_0 is obtained by merging q_0 and q_3 into a single tree, that is, appending q_3 to the w_3-leaves of q_0 and then removing repetitions along paths that may possibly arise. Since none of the trees r_0, q_1 and q_2 have any w_3-leaves, r is indeed an element of $\mathrm{Tup}_v(W')$ and $r_0 \circ q_1^\infty \circ q_2^\infty$ uses q_1 and q_2 infinitely many times. Now, we will generalize this approach.

Assume that $t = (t_0, t_1, ..., t_k) \in \mathrm{Tup}_v(W)$ is a tuple and all nodes of $D(t)$ are reachable from t_0. We say that t_i for $1 \leq i \leq k$ is *active* if, and only if t_i lies on a cycle in $D(t)$ or is reachable from a node that lies on a cycle. The *inactive graph* $I(t)$ is obtained by removing all active nodes from $D(t)$. We call t *valid* if, and only if $I(t)$ only contains t_0. Let $\mathrm{Tup}_v^*(W)$ be the set of all valid tuples in $\mathrm{Tup}_v(W)$.

(4.63) Lemma. For any $t = (t_0, t_1, ..., t_k) \in \mathrm{Tup}_v(W)$, t is valid or there is a valid $t' = (t'_0, t'_1, ..., t'_{k'}) \in \mathrm{Tup}_v^*(W')$ for some $W' \subsetneq W$ such that

$$\pi[\![\varphi \mathrm{R} \psi]\!]_{t_0 \circ t_1^\infty \circ ... \circ t_k^\infty} \leq \pi[\![\varphi \mathrm{R} \psi]\!]_{t'_0 \circ (t'_1)^\infty \circ ... \circ (t'_{k'})^\infty}.$$

Proof sketch. Thanks to lemma (4.61), we can assume that all nodes in $D(t)$ are reachable from t_0. Now, suppose that t is invalid, that is, $I(t)$ contains other inactive trees than t_0. By definition of $I(t)$, it is a directed acyclic graph, which implies we can merge the trees in $I(t)$ into a finite tree t_I, which can be turned into t'_0 by removing node repetitions along paths. Suppose w.l.o.g. that $t_1, ..., t_{k'}$ are the active nodes and $W' = \{1, ..., k'\} \subsetneq W$. The tuple $t' = (t'_0, t_1, ..., t_{k'})$ is then an element of $\mathrm{Tup}_v(W')$, since neither t'_0 nor any of the trees $t_1, ..., t_{k'}$ contain

any marked leaves in $W \setminus W'$, otherwise the inactive trees would be reachable from $t_1, ..., t_{k'}$ in $D(t)$, which is a contradiction. Also, since $t_1, ..., t_{k'}$ were all reachable from t_0, by construction of t'_0, they are also reachable from t'_0 and they are also active in $D(t')$, which implies that $I(t')$ only contains t'_0 and t' is valid, so we have $t' \in \mathrm{Tup}^*_v(W')$. The absorption

$$\pi[\![\varphi \mathrm{R} \psi]\!]_{t_0 \circ t_1^\infty \circ ... \circ t_k^\infty} \leq \pi[\![\varphi \mathrm{R} \psi]\!]_{t'_0 \circ t_1^\infty \circ ... \circ t_{k'}^\infty}$$

is verified by observing that $t'_0 \circ t_1^\infty \circ ... \circ t_{k'}^\infty$ can be obtained from $t_0 \circ t_1^\infty \circ ... \circ t_k^\infty$ by removing node repetitions along paths, which increases $\varphi \mathrm{R} \psi$-costs. □

The lemma implies that for calculating $\pi[\![\mathrm{A}(\varphi \mathrm{R} \psi)]\!]_v$, we can leave out invalid tuples, because they are absorbed by valid tuples anyway, so we have

$$\begin{aligned} \pi[\![\mathrm{A}(\varphi \mathrm{R} \psi)]\!]_v &= \sum_{W \subseteq V} \sum_{(t_0, t_1, ..., t_k) \in \mathrm{Tup}_v(W)} \pi[\![\varphi \mathrm{R} \psi]\!]_{t_0 \circ t_1^\infty \circ ... \circ t_k^\infty} \\ &= \sum_{W \subseteq V} \sum_{(t_0, t_1, ..., t_k) \in \mathrm{Tup}^*_v(W)} \pi[\![\varphi \mathrm{R} \psi]\!]_{t_0 \circ t_1^\infty \circ ... \circ t_k^\infty}. \end{aligned}$$

But for valid tuples $t = (t_0, t_1, ..., t_k) \in \mathrm{Tup}^*_v(W)$, we know that all nodes in $D(t)$ are reachable from t_0 and, except for t_0, all of them are active, which means that they are on a cycle or reachable from a node on a cycle. This implies that $t_1, ..., t_k$ are used infinitely often in the construction of $t_0 \circ t_1^\infty \circ ... \circ t_k^\infty$, so we have

$$\pi[\![\varphi \mathrm{R} \psi]\!]_{t_0 \circ t_1^\infty \circ ... \circ t_k^\infty} = \pi[\![\varphi \mathrm{R} \psi]\!]_{t_0} \cdot \pi[\![\varphi \mathrm{R} \psi]\!]_{t_1}^\infty \cdot ... \cdot \pi[\![\varphi \mathrm{R} \psi]\!]_{t_k}^\infty.$$

This yields the following proposition.

(4.64) Proposition. For an absorptive lattice semiring K, a finite set of nodes V, a formula $\mathrm{A}(\varphi \mathrm{R} \psi)$ in CTL and a matching K-interpretation π over V, we have

$$\pi[\![\mathrm{A}(\varphi \mathrm{R} \psi)]\!]_v = \sum_{W \subseteq V} \sum_{(t_0, t_1, ..., t_k) \in \mathrm{Tup}^*_v(W)} \pi[\![\varphi \mathrm{R} \psi]\!]_{t_0} \cdot \pi[\![\varphi \mathrm{R} \psi]\!]_{t_1}^\infty \cdot ... \cdot \pi[\![\varphi \mathrm{R} \psi]\!]_{t_k}^\infty$$

for any $v \in V$, where $\mathrm{Tup}^*_v(W)$ for $W \subseteq V$ is defined as above.

To close this subsection, we can derive an algorithm that computes $\pi[\![\mathrm{A}(\varphi \mathrm{R} \psi)]\!]_v$ from this proposition.

(4.65) Algorithm. The input for the algorithm is a formula $\mathrm{A}(\varphi \mathrm{R} \psi)$ with a matching K-interpretation π over a finite set of nodes V, where K is an absorptive lattice semiring and a node $v \in V$. The output $\pi[\![\mathrm{A}(\varphi \mathrm{R} \psi)]\!]_v$ is computed as follows.

1. Start with $S := 0$.
2. For any $W = \{w_1, ..., w_k\} \subseteq V$, repeat:
   (a) Find all trees in $T'_v(V, W)$ and $T^{\mathrm{CYC}}_{w_1}(V, W)$, ..., $T^{\mathrm{CYC}}_{w_k}(V, W)$.
   (b) For every combination of those trees $t = (t_0, t_1, ..., t_k) \in \mathrm{Tup}_v(W)$, repeat:
      i. Construct the dependency graph $D(t)$.
      ii. Check if every node in $D(t)$ is reachable from t_0.
         - If the check fails, discard t.

iii. Find all active nodes in $D(t)$.
- If a node other than t_0 is inactive, discard t.

iv. If t was not discarded, then $t \in \mathrm{Tup}_v^*(W)$.

v. Compute $\pi[\![\varphi \mathrm{R} \psi]\!]_{t_0}$, $\pi[\![\varphi \mathrm{R} \psi]\!]_{t_1}$, ... $\pi[\![\varphi \mathrm{R} \psi]\!]_{t_k}$.

vi. Update $S \coloneqq S + \pi[\![\varphi \mathrm{R} \psi]\!]_{t_0} \cdot \pi[\![\varphi \mathrm{R} \psi]\!]_{t_1}^{\infty} \cdot \ldots \cdot \pi[\![\varphi \mathrm{R} \psi]\!]_{t_k}^{\infty}$.

3. The output is

$$S = \sum_{W \subseteq V} \sum_{(t_0, t_1, \ldots, t_k) \in \mathrm{Tup}_v^*(W)} \pi[\![\varphi \mathrm{R} \psi]\!]_{t_0} \cdot \pi[\![\varphi \mathrm{R} \psi]\!]_{t_1}^{\infty} \cdot \ldots \cdot \pi[\![\varphi \mathrm{R} \psi]\!]_{t_k}^{\infty} = \pi[\![\mathrm{A}(\varphi \mathrm{R} \psi)]\!]_v.$$

Proof. The correctness is implied by proposition (4.64). For the runtime, let $|V| = n > 0$. Clearly, step (1) is repeated 2^n times.

Let $N(n)$ be an upper bound for the cardinality of $T'_v(V, V)$ for any given $n \in \omega$. We claim that $N(n) \in \mathcal{O}(2^{n!})$ and prove this by induction on n. For $n = 1$, there are only two possible trees in $T'_v(V, V)$, a tree with a single root (v) or a marked leaf $(v*)$, so $N(1) = 2$. For $n > 1$, we have those same two trees (v) and $(v*)$ in $T'_v(V, V)$, but additionally, if v is not marked, it could have successors. In that case, the successors of v are uniquely defined, since our trees must be complete, and there are at most $(n-1)$ successors, since v may not repeat, and therefore v cannot be a successor of v itself. Each of the successors w forms a subtree in $T'_w(V \setminus \{v\}, V \setminus \{v\})$, so there are at most $N(n-1)$ possibilities for each successor, which yields

$$N(n) \leq 2 + N(n-1)^{(n-1)} \leq N(n-1)^n \leq (2^{(n-1)!})^n = 2^{n!},$$

since $N(n-1) \geq 2$. Therefore, we conclude that $N(n) \in \mathcal{O}(2^{n!})$.

Moreover, the cardinality of $T_w^{\mathrm{CYC}}(V, V)$ for any $w \in W$ is in $\mathcal{O}(2^{n \cdot n!})$, because the number of trees in $T_w^{\mathrm{CYC}}(V, V)$ is bounded by $1 + N(n)^n$, since all trees in $T_w^{\mathrm{CYC}}(V, V)$ consist of a root node w with at most n successors, and each successor u induces a subtree in $T'_u(V, V)$.

We conclude that step (a) of the algorithm requires at most $\mathcal{O}(n \cdot 2^{n \cdot n!})$ operations, since all elements of $T'_v(V, W) \subseteq T'_v(V, V)$ and $T_{w_1}^{\mathrm{CYC}}(V, W) \subseteq T_{w_1}^{\mathrm{CYC}}(V, V)$, ..., $T_{w_k}^{\mathrm{CYC}}(V, W) \subseteq T_{w_k}^{\mathrm{CYC}}(V, V)$ have to be enumerated.

Step (b) is repeated at most $\mathcal{O}(2^{n!} \cdot (2^{n \cdot n!})^n)$ times, which is the maximum number of tuples in $\mathrm{Tup}_v(W)$.

Finally, the inner loop (i) to (vi) can be performed in $\mathcal{O}(n^3)$, the most expensive operation is finding the active nodes in $D(t)$, which can be done by running a breadth-first search from all $k \leq n+1$ nodes in $D(t)$, which takes $\mathcal{O}(k^2)$ operations each time.

Therefore, we obtain a total runtime in $\mathcal{O}(2^n \cdot 2^{n!} \cdot (2^{n \cdot n!})^n \cdot n^3)$, which we can simplify to

$$\begin{aligned} 2^n \cdot 2^{n!} \cdot (2^{n \cdot n!})^n \cdot n^3 &= 2^n \cdot 2^{n!} \cdot 2^{n^2 \cdot n!} \cdot 2^{3 \cdot \log(n)} \\ &= 2^{(n^2+1) \cdot n! + n + 3 \cdot \log(n)}. \end{aligned}$$

This shows that our algorithm terminates and ends the proof. With a runtime in $\mathcal{O}(2^{(n^2+1) \cdot n! + n + 3 \cdot \log(n)})$, which is a double exponential runtime, the algorithm is only relevant for theoretical purposes. □

We have now developed algorithms that are able to to evaluate all formulas in CTL in their respective semirings.

4.4 Program Iterations in PDL

In this section, we will see that interpreting program iterations ρ^* where ρ is a program in PDL is similar to the interpretation of existential until formulas in CTL from subsection (4.2.1). First of all, let K be an ω-continuous semiring and π a K-interpretation over a finite set of nodes V. For all $v, w \in V$, let $P_{v\to w}(V)$ denote the set of all paths over V from v to w. With the ρ-costs for paths from definition (4.6), we can prove a theorem for PDL programs that is similar to theorem (4.14).

(4.66) Theorem. For any ω-continuous semiring K, PDL program ρ, finite set V, K-interpretation π over V and $v, w \in V$, we have

$$\pi[\![\rho^*]\!]_{(v,w)} = \sum_{p\in P_{v\to w}(V)} \pi[\![\rho]\!]_p.$$

Proof. By definition (3.21), we have $\pi[\![\rho^*]\!]_{(v,w)} = \mathrm{lfp}(f^{\rho^*})_{(v,w)}$. Theorem (2.23) states that

$$\mathrm{lfp}(f^{\rho^*}) = \sup_{i\in\omega} (f^{\rho^*})^i(0).$$

Set $X_i = (f^{\rho^*})^i(0)$. We will show by induction on i that

$$(X_i)_{(v,w)} = \sum_{p\in P^{<i}_{v\to w}(V)} \pi[\![\rho]\!]_p$$

for all $v, w \in V$ where $P^{<i}_{v\to w}(V)$ refers to the paths from v to w that are shorter than i.

In the base case $i = 0$, there is nothing to show, since there are no paths that are shorter than 0 and $(X_0)_{(v,w)} = 0$.

If we assume the hypothesis to be true for i, then for $i + 1$, we have

$$\begin{aligned}(X_{i+1})_{(v,w)} &= f^{\rho^*}(X_i) \\ &= \pi[\![1?]\!]_{(v,w)} + \sum_{u\in V} \pi[\![\rho]\!]_{(v,u)} \cdot (X_i)_{(u,w)}.\end{aligned}$$

Using the induction hypothesis for $(X_i)_{(u,w)}$ yields

$$\begin{aligned}
(X_{i+1})_{(v,w)} &= \pi[\![1?]\!]_{(v,w)} + \sum_{u\in V} \pi[\![\rho]\!]_{(v,u)} \cdot \left(\sum_{p\in P^{<i}_{u\to w}(V)} \pi[\![\rho]\!]_p \right) \\
&= \pi[\![1?]\!]_{(v,w)} + \sum_{u\in V} \sum_{p\in P^{<i}_{u\to w}(V)} \pi[\![\rho]\!]_{(v,u)} \cdot \pi[\![\rho]\!]_p \\
&= \pi[\![1?]\!]_{(v,w)} + \sum_{u\in V} \sum_{p\in P^{<i}_{u\to w}(V)} \pi[\![\rho]\!]_{(v,p)} \\
&\overset{(*)}{=} \sum_{p\in P^{=0}_{v\to w}(V)} \pi[\![\rho]\!]_p + \sum_{p\in \bigcup_{1\le j\le i} P^{=j}_{v\to w}(V)} \pi[\![\rho]\!]_p \\
&= \sum_{p\in P^{<i+1}_{v\to w}(V)} \pi[\![\rho]\!]_p.
\end{aligned}$$

The transformation $(*)$ is verified by two observations. First, there is only a path from v to w of length 0 if $v = w$. In that case, the ρ-cost of this path is $\pi[\![1?]\!]_{(v,w)} = 1$, since it does not have any edges. Otherwise, if $v \neq w$, then $\pi[\![1?]\!]_{(v,w)} = 0$. The second observation is that for $p \in P^{<i}_{u\to w}(V)$ and $u \in V$, the paths (v, p) that are formed by adding v to the start of p are exactly all the paths from v to w of length 1 to i, and the ρ-cost of (v, p) is clearly $\pi[\![\rho]\!]_{(v,u)} \cdot \pi[\![\rho]\!]_p$. This ends the induction.

Next, we will show

$$\left(\sup_{i\in\omega} X_i\right)_{(v,w)} = \sum_{p\in P_{v\to w}(V)} \pi[\![\rho]\!]_p$$

for all $v, w \in V$. We apply lemma (2.21) to obtain

$$\begin{aligned}
\left(\sup_{i\in\omega} X_i\right)_{(v,w)} &= \sup_{i\in\omega} (X_i)_{(v,w)} \\
&= \sup_{i\in\omega} \sum_{p\in P^{<i}_{v\to w}(V)} \pi[\![\rho]\!]_p \\
&= \sup_{i\in\omega} \sum_{p\in P^{\le i}_{v\to w}(V)} \pi[\![\rho]\!]_p \\
&= \sum_{i\in\omega} \sum_{p\in P^{=i}_{v\to w}(V)} \pi[\![\rho]\!]_p \\
&= \sum_{p\in P_{v\to w}(V)} \pi[\![\rho]\!]_p.
\end{aligned}$$

Note that $P^{=i}_{v\to w}(V)$ for $i \in \omega$ partitions the set $P_{v\to w}(V)$, since all paths from v to w have some finite length $i \in \omega$. We conclude

$$\pi[\![\rho^*]\!]_{(v,w)} = \left(\sup_{i\in\omega} (f^{\rho^*})^i(0)\right)_{(v,w)} = \left(\sup_{i\in\omega} X_i\right)_{(v,w)} = \sum_{p\in P_{v\to w}(V)} \pi[\![\rho]\!]_p,$$

which ends the proof. □

Recall the definition (4.15) where we have defined K-automata. Existential until formulas in CTL were evaluated under a K-interpretation π by transforming π into a K-automaton and then applying state removal. We will do the same for PDL programs. To calculate $\pi[\![\rho^*]\!]_{(v,w)}$, we will first build an appropriate K-automaton $\mathcal{A}^{\rho^*}_{(v,w)}(\pi)$ and then calculate $C(\mathcal{A}^{\rho^*}_{(v,w)}(\pi))$.

(4.67) Proposition. For an ω-continuous semiring K, a finite set of nodes V, a K-interpretation π over V and a program ρ, we define the K-automaton $\mathcal{A}^{\rho^*}_{(v,w)}(\pi) = (Q, C, s, t)$ for all $v, w \in V$ by setting $Q = V \cup \{s, t\}$ and connecting s to v and w to t. We assume $s, t \notin V$. Formally, we set

$$\begin{aligned}
C(q, s) &= 0 && \text{for all } q \in Q,\\
C(t, q) &= 0 && \text{for all } q \in Q,\\
C(s, v) &= 1 &&\\
C(s, q) &= 0 && \text{for all } q \in Q \setminus \{v\},\\
C(w, t) &= 1 &&\\
C(q, t) &= 0 && \text{for all } q \in Q \setminus \{w\} \quad \text{and}\\
C(u_1, u_2) &= \pi[\![\rho]\!]_{(u_1,u_2)} && \text{for all } u_1, u_2 \in V.
\end{aligned}$$

With this construction, $\mathcal{A}^{\rho^*}_{(v,w)}(\pi)$ has the cost

$$C(\mathcal{A}^{\rho^*}_{(v,w)}(\pi)) = \sum_{p \in P_{v \to w}(V)} \pi[\![\rho]\!]_p.$$

Proof sketch. Recall the definition of $C(\mathcal{A}^{\rho^*}_{(v,w)}(\pi))$ as

$$C(\mathcal{A}^{\rho^*}_{(v,w)}(\pi)) = \sum_{p \in P_{s \to t}(Q)} C(p).$$

We have to prove

$$\sum_{p \in P_{s \to t}(Q)} C(p) = \sum_{p' \in P_{v \to w}(V)} \pi[\![\rho]\!]_{p'}.$$

This is easy to see, because any path $p \in P_{s \to t}(Q)$ with nonzero costs must be a path of the form

$$p = (s, v, ..., w, t).$$

Clearly, we can map p to $p' = (v, ..., w) \in P_{v \to w}(V)$ by cutting off s and t. Since the edges (u_1, u_2) in $\mathcal{A}^{\rho^*}_{(v,w)}(\pi)$ for $u_1, u_2 \in V$ are defined to have the cost $\pi[\![\rho]\!]_{(u_1,u_2)}$, we conclude that

$$C(p) = \pi[\![\rho]\!]_{p'}.$$

Therefore, we have a cost-preserving one-to-one correspondence between $P_{s \to t}(Q)$ and $P_{v \to w}(V)$ if we disregard paths with cost 0, which proves the claim. □

Since we have already introduced the state removal algorithm (4.18) that computes $C(\mathcal{A})$ for a given K-automaton $\mathcal{A}$, there is a straightforward way to interpret PDL program iterations $\pi[\![\rho^*]\!]_{(v,w)}$.

1. Construct $\mathcal{A}^{\rho^*}_{(v,w)}(\pi)$.

- Proposition (4.67) and theorem (4.66) yield

$$C(\mathcal{A}^{\rho^*}_{(v,w)}(\pi)) = \sum_{p \in P_{v \to w}(V)} \pi[\![\rho]\!]_p = \pi[\![\rho^*]\!]_{(v,w)}.$$

2. Use the state removal algorithm (4.18) to compute $C(\mathcal{A}^{\rho^*}_{(v,w)}(\pi))$.

   - The output is the interpretation $\pi[\![\rho^*]\!]_{(v,w)}$.

Assuming that $|V| = n$, the automaton $\mathcal{A}^{\rho^*}_{(v,w)}(\pi)$ can be constructed in $\mathcal{O}(n^2)$ and state removal has a runtime on $\mathcal{O}(n^3)$, which yields a total runtime in $\mathcal{O}(n^3)$ to interpret PDL program iterations.

We will close this section by calculating $\pi[\![a^*]\!]_{(v,w)}$ for the $\mathbb{N}^\infty[\![X]\!]$-interpretation π from example (3.22). Figure (4.68) below shows the $\mathbb{N}^\infty[\![X]\!]$-interpretation π and the $\mathbb{N}^\infty[\![X]\!]$-automaton $\mathcal{A}^{a^*}_{(v,w)}(\pi)$ that we will use to compute $\pi[\![a^*]\!]_{(v,w)}$.

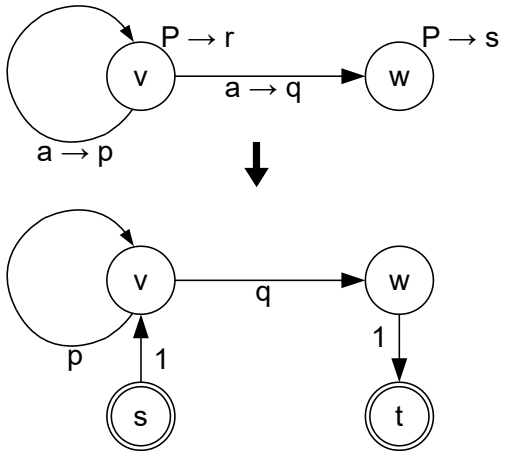

Figure (4.68): $\mathbb{N}^\infty[\![X]\!]$-interpretation π (above) transformed to $\mathcal{A}^{a^*}_{(v,w)}(\pi)$ (below).

Now, we can apply state removal, as shown in figure (4.69) below.

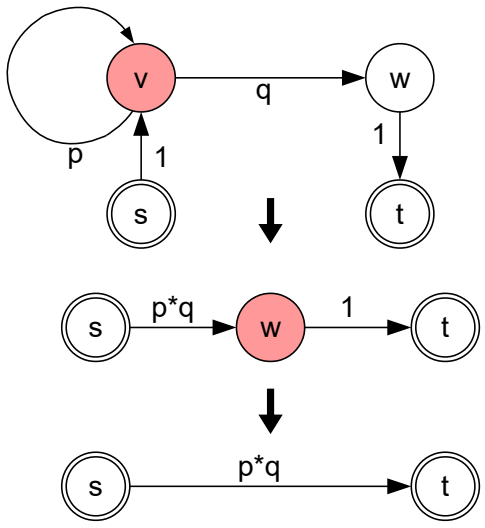

Figure (4.69): State removal performed on $\mathcal{A}^{a^*}_{(v,w)}(\pi)$.

The result is $\pi[\![a^*]\!]_{(v,w)} = p^*q$ which is also what we obtained in example (3.22).

Chapter 5

Conclusion

We have defined semiring interpretations for LTL, CTL and positive PDL, and with the algorithms from the previous section, we have also developed a way to compute finite representations of the resulting semiring elements. However, there are still open questions. For example, while we guarantee that the representations of the semiring values that we compute in the previous section are finite, we do not provide any guarantees for their readability. Also, it remains unknown whether the runtime of the algorithms that we provided is optimal. As conjecture (4.36) suggests, there might be more efficient algorithms for semiring interpretations.

Moreover, our entire work is based on ω-continuous semirings and absorptive lattice semirings. In particular, we use absorptive lattice semirings to evaluate release formulas in CTL. It is an open question whether there are any larger classes of semirings that are suited to interpret release formulas in CTL. While we know that the conditions that we asserted in definition (2.10) for absorptive lattice semirings are sufficient to enable the interpretation of release formulas, it is not known whether they are too strong. This is a topic that will be covered in future work.

Finally, we will mention some related topics. As pointed out in remark (3.24), there are still issues left to overcome in order to extend our semiring interpretations from positive PDL to full PDL. Moreover, the semiring interpretations for some CTL formulas were inspired by their translation into the modal μ-calculus, which is a stronger logic than LTL, CTL and PDL. This raises the question whether the approaches that we used to define the semiring interpretations for CTL could be generalized to the full modal μ-calculus.

Bibliography

[Bar91] Andrei Baranga. The contraction principle as a particular case of Kleene's fixed point theorem. *Discrete Mathematics*, 98(1):75–79, 1991.

[Ber05] Dietmar Berwanger. *Games and Logical Expressiveness*. PhD thesis, RWTH Aachen, 2005.

[BLS10] Andreas Bauer, Martin Leucker, and Christian Schallhart. Comparing LTL Semantics for Runtime Verification. *Journal of Logic and Computation*, 20(3):651–674, 2010.

[CC79] Patrick Cousot and Radhia Cousot. Constructive versions of Tarski's fixed point theorems. *Pacific Journal of Mathematics*, 82(1):43–57, 1979.

[CDG+07] H. Comon, M. Dauchet, R. Gilleron, F. Jacquemard, D. Lugiez, C. Löding, S. Tison, and M. Tommasi. Tree Automata Techniques and Applications. Available on: `http://www.grappa.univ-lille3.fr/tata`, 2007. release October, 12th 2007.

[GKT07] Todd J. Green, Grigoris Karvounarakis, and Val Tannen. Provenance Semirings. In *Proceedings of the twenty-sixth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 31–40. ACM, 2007.

[GT17] Erich Grädel and Val Tannen. Semiring Provenance for First-Order Model Checking. arXiv:1712.01980 [cs.LO], 2017.

[GT18] Erich Grädel and Val Tannen. Provenance in Logic and Games. unpublished, 2018.

[HR00] Michael Huth and Mark Ryan. *Logic in Computer Science: Modelling and Reasoning about Systems*. Cambridge University Press, 2000.

[Neu05] Christoph Neumann. Converting Deterministic Finite Automata to Regular Expressions. 2005.

[Ran52] George N. Raney. Completely Distributive Complete Lattices. *Proceedings of the American Mathematical Society*, 3(5):677–680, 1952.