

Two-Way Cost Automata and Cost Logics over Infinite Trees^{*}

Achim Blumensath[†]

TU Darmstadt
blumensath@mathematik.tu-
darmstadt.de

Thomas Colcombet

Université Paris Diderot
thomas.colcombet@liafa.univ-paris-
diderot.fr

Denis Kuperberg

University of Warsaw
denis.kuperberg@gmail.com

Paweł Parys[‡]

University of Warsaw
parys@mimuw.edu.pl

Michael Vanden Boom

University of Oxford
michael.vandenboom@cs.ox.ac.uk

Abstract

Regular cost functions provide a quantitative extension of regular languages that retains most of their important properties, such as expressive power and decidability, at least over finite and infinite words and over finite trees. Much less is known over infinite trees.

We consider cost functions over infinite trees defined by an extension of weak monadic second-order logic with a new fixed-point-like operator. We show this logic to be decidable, improving previously known decidability results for cost logics over infinite trees. The proof relies on an equivalence with a form of automata with counters called quasi-weak cost automata, as well as results about converting two-way alternating cost automata to one-way alternating cost automata.

Categories and Subject Descriptors Theory of computation [Automata over infinite objects]

1. Introduction

Boundedness is a central notion arising in both mathematics and computer science. Indeed, being able to determine the existence of bounds is an important form of quantitative reasoning in many contexts, ranging from verification to model theory.

Consider the following boundedness questions:

- Given a finite state automaton with some costly transitions, is there a bound n such that any finite word accepted by the automaton has an accepting run which takes these costly edges at most n times?

^{*}The research leading to these results has received funding from the European Union's Seventh Framework Programme (FP7/2007-2013) under grant agreement n°259454 and the project ANR 2010 BLAN 0202 02 FREC.

[†] Work partially supported by DFG grant BL 1127/2-1.

[‡] Work supported by the fellowship of the Foundation for Polish Science, during the author's post-doc stay at Université Paris Diderot.

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

CSL-LICS 2014, July 14–18, 2014, Vienna, Austria.
Copyright © 2014 ACM 978-1-4503-2886-9...\$15.00.
<http://dx.doi.org/10.1145/2603088.2603104>

- Given some regular language L of finite words, is there a bound n such that L^* (consisting of any finite concatenation of words from L) is equal to $L^0 \cup L^1 \cup \dots \cup L^n$?
- Given a monadic second-order formula $\varphi(X, x)$ positive in X , is there a bound n such that the least fixed point of φ over finite words is always reached within n iterations?

We could add to this list many other questions like the star height problem from language theory [11, 15–17], the boundedness problem for fixed points of monadic second-order formulae in model theory [3, 4], the bounded repair problem in database theory [2], and the resource bounded reachability problem in verification [20]. We could also consider these problems over different classes of structures (like finite trees or infinite words). The *theory of regular cost functions* introduced in [9] (and developed in subsequent papers) is a powerful extension of the theory of regular languages that enables all of these boundedness questions to be decided in a uniform way.

In the framework of cost functions, the notion of a language is extended to the richer notion of a function from structures (like words or trees) to $\mathbb{N} \cup \{\infty\}$. The central decidability question is whether a function is bounded by some $n \in \mathbb{N}$ over its domain (or a regular subset of its domain). By identifying a language with its characteristic function mapping structures in the language to 0 and everything else to ∞ , these boundedness questions subsume classical language decision problems like universality and emptiness.

Like regular languages, these regular cost functions can be described using a number of different formalisms. One way to define these regular cost functions is using *cost monadic logic*, an extension of monadic second-order logic with an operator $|X| \leq N$ which enables some limited reasoning about the cardinality of sets (enough to distinguish notions of “large” and “small” for the purposes of boundedness). These regular cost functions can also be defined in terms of *cost automata*, which are traditional automata enriched with a finite set of counters that can be incremented, reset, or left unchanged on each transition, and are used to assign a value from $\mathbb{N} \cup \{\infty\}$ to each input.

But the connections with regular languages do not end there. Since the works of Rabin, Scott, and Büchi [7, 23, 24], there have not been many attempts to extend the expressive power of the associated formalisms beyond regular languages, while at the same time keeping most of their good properties – closure, decidability, and the equivalent presentations in terms of regular expressions, algebra, automata, and logic. To the best of our knowledge, the theory of regular cost functions is the only extension that achieves faithful

extensions of classical results (in terms of closure, decidability, and equivalent presentations) over finite words [9, 10], infinite words [19], and finite trees [13]. This unique position makes it an attractive framework, since it subsumes many classical results while allowing further questions about boundedness to be answered.

Central open question

The central open question in the theory is whether it can also be faithfully extended to infinite trees. In particular the question “can we solve satisfiability of cost monadic logic over infinite trees?” – the counterpart of the theorem of Rabin for monadic second-order logic – is still open.

This paper can be seen as a step to close this gap, by showing that there is a robust subclass of decidable cost functions over infinite trees that have natural correspondences with cost logics and cost automata.

Showing decidability of the full logic would help answer some fundamental questions about language and automata theory that have been open for decades. One of the motivating open problems like this is the (nondeterministic) *Mostowski index problem*, or parity index problem, which asks, given a regular language of infinite trees and a set of priorities P , is there a nondeterministic parity automaton that accepts this language using only priorities P ? It turns out that this problem can be reduced to a boundedness question for regular cost functions over infinite trees [12]. Being able to solve this problem is useful, since the number of priorities in a parity automaton (even more so than the number of states) reflects how complicated the automaton is. Indeed, minimizing the number of priorities is especially important in verification, because model checking against a property described by a parity automaton is essentially exponential in the number of priorities.

Thus, our desire to understand and solve cost logics over infinite trees comes from a desire to close the gap in the theory of regular cost functions, as well as a desire to solve some long-standing and important problems in automata theory.

Weak and quasi-weak cost functions

In trying to solve the full logic, the weak fragment of cost monadic logic (where second-order quantifiers range over finite sets only) was a natural starting point for investigation. In [27], weak cost monadic logic was shown to be decidable, and equivalent to a form of weak alternating automata with counters.

However, weakness as defined for regular languages is not canonical in the context of cost functions. Indeed, a new, richer form of weakness emerged naturally in [18], in the form of *quasi-weak cost automata*, that enjoy many of the good properties of weak cost automata, but are strictly more expressive.

At the same time, there was a branch of work [3, 4] that sought to study the boundedness problem for fixed point formulas, and it turned out that the natural class of automata for problems like this was the quasi-weak class.

Recently in [14], it was also shown that the quasi-weak class of cost functions is strong enough to solve a special case of the *weak definability problem*. The weak definability problem asks, given a regular language of infinite trees, whether or not there is a weak automaton (equivalently, weak monadic second-order logic formula) that captures the same language. Like the parity index problem, the general version of this problem is open, and is of interest in verification because the weakly definable languages constitute an expressive but computationally feasible class of properties.

Thus, the quasi-weak class of cost functions emerged from a variety of sources as an interesting subclass of regular cost functions over infinite trees, with links to concrete decidable results and natural logical questions. However, it was not known whether this class corresponded to a natural extension of monadic second-order

logic. This paper addresses this question, and develops further tools for understanding and working with cost logics over infinite trees.

Contributions

In this paper, we demonstrate that the class of quasi-weak cost functions corresponds to natural cost logics.

- We characterize the class of quasi-weak cost functions in terms of an extension of monadic second-order logic with a new form of fixed points that allows only bounded unfolding.
- We describe an extension of the μ -calculus with this bounded unfolding operator, and explain that the alternation-free fragment of this logic also characterizes quasi-weak cost functions.
- We show how to translate a 2-way quasi-weak automaton (in fact, any 2-way cost parity automaton) into an equivalent 1-way automaton. This is the technical core of our contribution, which is central to showing the equivalence of quasi-weak cost automata and logic. Our translation differs from other similar constructions in the context of regular languages in the sense that the translation does not rely on a global positional determinacy result in the underlying game, and instead uses only local approximations of the two-way plays. Although the proof is technical, we believe the ideas behind this construction may be of independent interest, even outside of the theory of regular cost functions.

Related work

This work builds on the classical results due to Rabin, Scott and Büchi [7, 23, 24] about regularity for languages of words and trees, both finite and infinite. It is also indebted to the work of Hashiguchi, Leung, Simon and Kirsten [15, 17, 21, 25] on limitedness questions (mainly in the context of solving the star height problem), and work by Bojańczyk and Colcombet [5, 6] on $\text{MSO}+\cup$ (another logic related to boundedness).

Structure of this document

The article is organized as follows. In Section 2, we introduce cost monadic logic, and its weak and quasi-weak fragments. We then describe 2-way and 1-way alternating quasi-weak automata in Section 3. In Section 4, our main results concerning equivalence and decidability are given, and the essential ideas of the proof are sketched. Finally, in Section 5 we mention some links with a cost form of the μ -calculus that provide additional insight into the quasi-weak class of cost functions.

Notation and conventions

\mathbb{A} will denote a fixed finite alphabet. The set of finite and infinite words over \mathbb{A} is \mathbb{A}^* and \mathbb{A}^ω , respectively. The empty word is ϵ . For simplicity we work primarily with infinite binary trees. Let $\mathcal{T} = \{0, 1\}^*$ be the unlabeled infinite binary tree. The set $\mathcal{T}_{\mathbb{A}}$ of complete \mathbb{A} -labeled binary trees consists of mappings $t : \mathcal{T} \rightarrow \mathbb{A}$. A branch π is a word $\{0, 1\}^\omega$.

2. Quasi-weak cost logic

2.1 Cost functions

We write \mathbb{N} for the set of non-negative integers and \mathbb{N}_∞ for the set $\mathbb{N} \cup \{\infty\}$ with the obvious ordering. Non-decreasing functions $\mathbb{N} \rightarrow \mathbb{N}$ are called *correction functions*. We denote them by α, β, \dots . Any such function is implicitly extended to \mathbb{N}_∞ by $\alpha(\infty) = \infty$.

Let E be a set and \mathcal{F}_E the set of functions $E \rightarrow \mathbb{N}_\infty$. For $f, g \in \mathcal{F}_E$ and a correction function α , we write $f \preceq_\alpha g$ if $f \leq \alpha \circ g$ (or if we are comparing single values $n, m \in \mathbb{N}$, $n \preceq_\alpha m$

if $n \leq \alpha(m)$). We write $f \approx_\alpha g$ if $f \preceq_\alpha g$ and $g \preceq_\alpha f$. Finally, $f \approx g$ (respectively, $f \preceq g$) if $f \approx_\alpha g$ (respectively, $f \preceq_\alpha g$) for some α . Note that $f \preceq g$ holds if, and only if, f is bounded over every set $X \subseteq E$ over which g is bounded. The idea is that the *boundedness relation* \approx does not pay attention to exact values, but does preserve the existence of bounds over all subsets of the domain. A *cost function over E* is an equivalence class of \mathcal{F}_E/\approx . In practice, a cost function (denoted f, g, \dots) will be represented by one of its elements in \mathcal{F}_E . In this paper, E will usually be $\mathcal{T}_{\mathbb{A}}$ and functions defined by logics and automata will always be considered as cost functions, *i.e.*, up to \approx .

We can identify a language L with its characteristic function χ_L mapping structures in the language to 0 and everything else to ∞ . Note that for languages K and L , $K \subseteq L$ iff $\chi_L \preceq \chi_K$, so deciding cost function equivalence subsumes language inclusion testing in the classical setting.

Given two mathematical expressions $E(\bar{x}), F(\bar{x})$ denoting elements of \mathbb{N}_∞ and depending on variables \bar{x} (ranging over an implicit domain clear from the context), we write $E \preceq^{\bar{x}} F$ to specify that $\lambda \bar{x}.E \preceq \lambda \bar{x}.F$ where $\lambda \bar{x}.E$ denotes the function mapping \bar{x} to $E(\bar{x}) \in \mathbb{N}_\infty$. The relation $\approx^{\bar{x}}$ is defined accordingly.

2.2 Cost monadic logic

Cost monadic second-order logic (CMSO) was introduced in [8] as a quantitative extension of monadic second-order logic. As usual, the logic can be defined over any relational structure, but we restrict our attention to CMSO over trees. In addition to first-order variables ranging over nodes of the tree and set variables ranging over sets of nodes, CMSO uses a single additional variable N , called the *bound variable*, which ranges over \mathbb{N} .

The atomic formulae in CMSO are those from MSO (the membership relation $x \in X$ and relations $a(x, x_1, x_2)$ asserting that $a \in \mathbb{A}$ is the label at position x with children x_1, x_2 from left to right), as well as a new predicate $|X| \leq N$ where X is any set variable and N is the bound variable. Arbitrary CMSO formulae are built inductively by applying boolean connectives and by quantifying (existentially or universally) over first-order or set variables. We require that any predicates of the form $|X| \leq N$ appear positively in the formula (*i.e.*, within the scope of an even number of negations).

If we fix a value n for N , the semantics of $|X| \leq N$ is what one would expect: the predicate holds iff the value of X has cardinality at most n . If, however, no value for N is specified then a sentence φ in cost monadic logic defines a function $\llbracket \varphi \rrbracket : \mathcal{T}_{\mathbb{A}} \rightarrow \mathbb{N}_\infty$ by $\llbracket \varphi \rrbracket(t) := \inf \{n : t, n \models \varphi\}$, where we write $t, n \models \varphi$ if t satisfies φ when all occurrences of N take value n . Notice that in case φ is a pure MSO-sentence (not containing the predicates $|X| \leq N$), $\llbracket \varphi \rrbracket(t)$ is 0, if t satisfies the sentence φ , and ∞ otherwise.

The weak variant of CMSO (denoted WCMSO) restricts the second-order quantification to finite sets, and has been studied in [27].

2.3 Bounded expansion and quasi-weak cost monadic logic.

In this paper, we introduce a new logic called *quasi-weak cost monadic logic*, QWCMSO, which extends WCMSO by a *bounded expansion* operator μ^N . This operator takes a function F mapping sets of nodes to sets of nodes which is *monotonic* (*i.e.*, $X \subseteq Y$ implies $F(X) \subseteq F(Y)$), and it computes $F^{(N)}$, where $F^{(0)} = \emptyset$ and $F^{(\ell+1)} = F(F^{(\ell)})$. We denote this value as $\mu^N X.F(X)$ in order to make the name of the variable over which the construction is performed explicit.

To add this operator to WCMSO, we define the following new *bounded expansion* construct

$$x \in \mu^N Z. \{y : \varphi(y, Z)\},$$

where x, y are first order variables, Z is a set variable and $\varphi(y, Z)$ is a formula that *uses Z positively*, *i.e.*, every predicate of the form $z \in Z$ appears below an even number of negations. This operator binds the variables y and Z , while it leaves x free. We also require that any such bounded expansion construct appears positively in the formula. The semantics are as one would expect: $\{y : \varphi(y, Z)\}$ is a set which depends on Z , and is thus subject to the application of $\mu^N Z$. The variable N used here is the same bound variable used in the predicates $|X| \leq N$. As before, the semantics associate to a sentence the least n which can be substituted for all occurrences of N and make the sentence true.

Example 1. Let $\mathbb{A} = \{a, b\}$. We use the operator μ^N to define a formula counting the maximal number of consecutive a 's on a branch starting at the root, where $\text{root}(w)$ identifies the root of the tree:

$$\begin{aligned} \exists w [\text{root}(w) \wedge \\ x \in \mu^N X. \{x : \exists yz [b(x, y, z) \vee \\ (a(x, y, z) \wedge y \in X \wedge z \in X)]\}]. \end{aligned}$$

This is equivalent (in the sense of the \approx relation) to the CMSO formula

$$\forall X. [\text{downclose}(X) \wedge ((\forall x \in X) a(x)) \rightarrow |X| \leq N],$$

where $\text{downclose}(X)$ asserts that X is closed under the ancestor relation and $a(x)$ is shorthand for $\exists y, z. a(x, y, z)$.

3. Cost automata and games

In this section we introduce the model of automata used in this paper: the alternating 1-way/2-way B-quasi-weak automata. We consider classical parity automata over trees equipped with a finite set of counters Γ that can be *incremented* ic , *reset* r , or left *unchanged* ε (but whose values do not affect the flow of the automaton). Let $\mathbb{C} := \{\text{ic}, \text{r}, \varepsilon\}$ be the alphabet of counter actions. Each counter starts with value zero, and the value of a sequence of actions is the supremum of the values achieved during this sequence. For instance $(\text{ic})(\text{ic})\text{r}\varepsilon(\text{ic})\varepsilon$ has value 2, $((\text{ic})\text{r})^\omega$ has value 1, and $(\text{ic})\text{r}(\text{ic})^2\text{r}(\text{ic})^3\text{r}\dots$ has value ∞ . The set $\text{Act}_P^\Gamma := \mathbb{C}^\Gamma \times P$ collects the counter actions for a finite set Γ of counters and some finite set of priorities P . To an infinite sequence over Act_P^Γ , we assign the value ∞ if the maximum priority occurring infinitely often in it is odd (*i.e.*, if it does not satisfy the parity condition); otherwise, the value is the supremum of the values achieved by the counters (in case of several counters, we take the counter with the maximal value).

Formally, an *(alternating) two-way B-parity automaton* over the alphabet \mathbb{A} is a tuple $\langle Q, \mathbb{A}, q_0, \Gamma, P, \delta \rangle$ consisting of a finite set of states Q , an initial state $q_0 \in Q$, a finite set Γ of counters, a finite set P of priorities, and a transition function

$$\delta : Q \times \mathbb{A} \rightarrow \mathcal{B}^+ (\{\uparrow, \swarrow, \searrow, \circ\} \times \text{Act}_P^\Gamma \times Q)$$

mapping a state and a letter to a positive boolean combination of triples of the form (d, c, q) . Such a triple encodes the instruction to send the automaton to state q in direction d while performing action c . The directions \swarrow and \searrow move to the left or right child, \uparrow moves to the parent, and \circ stays in place. We assume that $\delta(q, a)$ is written in disjunctive normal form for all q and a . Without loss of generality, we assume that the automaton never proceeds in direction \uparrow from the root of the tree.

Acceptance of an input tree t by a B-automaton \mathcal{A} is defined in terms of a game (\mathcal{A}, t) between two players: Eve is in charge of

the disjunctive choices. She tries to minimize counter values and to satisfy the parity condition. Adam, on the other hand, is in charge of the conjunctive choices and tries to maximize counter values or to sabotage the parity condition. As the transition function is given in disjunctive normal form, each turn of the game consists of Eve choosing a disjunct and Adam then selecting a single tuple (d, c, q) from it. We assume that each disjunction is nonempty, and each disjunct contains a tuple with direction other than \uparrow ; in other words, from every position there is some move (i.e., the automaton cannot get stuck at the root).

A *play* of \mathcal{A} on the tree t is a sequence

$$q_0, (d_1, c_1, q_1), (d_2, c_2, q_2), \dots$$

compatible with t and δ , i.e., q_0 is initial, and for all $i \in \mathbb{N}$, $(d_{i+1}, c_{i+1}, q_{i+1})$ appears in $\delta(q_i, t(x_i))$ where x_i is the node of t after following the directions $d_1 d_2 \dots d_i$ starting from the root. The value $\text{val}(\pi)$ of a play π is the value of the sequence $c_1 c_2 \dots$ as defined above. We will say that π is *n-winning* (for Eve) if $\text{val}(\pi) \leq n$.

A *strategy* for one of the players in the game (\mathcal{A}, t) is a function that returns the next choice given the history of the play. If this function depends only on the current position in the game (rather than the full history), then it is *positional*. Note that choosing a strategy for Eve and a strategy for Adam fixes a play in (\mathcal{A}, t) . We say that a play π is *compatible* with a strategy σ if there is some strategy σ' for the other player such that σ and σ' together yield the play π . A strategy for Eve is *n-winning* if every play compatible with it is *n-winning*. We say that Eve *n-wins the game* if there is some *n-winning* strategy for Eve. An automaton *n-accepts* a tree t if Eve *n-wins* the game (\mathcal{A}, t) . We denote by $\llbracket \mathcal{A} \rrbracket : \mathcal{T}_{\mathbb{A}} \rightarrow \mathbb{N}_{\infty}$ the function given by $\llbracket \mathcal{A} \rrbracket(t) := \inf \{n : \mathcal{A} \text{ n-accepts } t\}$.

We will sometimes use automata that start from some position other than the root. We will call such automata *localized*. For localized automata, we use the notation $\llbracket \mathcal{A} \rrbracket_v(t)$ to specify the node v the automaton starts at. We can also consider cost automata with other well-known acceptance conditions. For instance, a B-Büchi automaton is a B-parity automaton using priorities $\{1, 2\}$. In this case, we often assume that priorities label states rather than edges (always possible by adding intermediate states), and refer to *Büchi states* (priority 2) and *non-Büchi states* (priority 1).

If every $\delta(q, a)$ uses only directions \swarrow and \searrow , then we call \mathcal{A} *one-way*. Moreover, if every $\delta(q, a)$ is of the form

$$\bigvee_i (\swarrow, c_i, q_i) \wedge (\searrow, c'_i, q'_i),$$

\mathcal{A} is *nondeterministic*. For such a one-way nondeterministic automaton, we define a *run* to be the set of possible plays compatible with some fixed strategy of Eve. Since the only choices of Adam are in the branching, a run labels the entire binary tree with states, and choosing a branch yields a unique play of the automaton. A run is accepting if it is accepting on all branches, and the value assigned to a run of a B-automaton is the supremum of the values across all branches. For nondeterministic automata, the choices of Eve and Adam in the game described above can be viewed as Eve picking a transition $(\swarrow, c_i, q_i) \wedge (\searrow, c'_i, q'_i)$, and Adam choosing a direction (which uniquely determines which atom Adam picks from the conjunction). Unless otherwise indicated, we assume automata to be alternating.

Variants of weakness for cost automata

Weak alternating automata are Büchi automata with the restriction that no cycle of the automaton visits both Büchi and non-Büchi states (this is equivalent to the original definition in [22]). We consider two variants of this classical notion of weakness for cost automata. An alternating B-Büchi automaton is called

B-weak if in all cycles, either all states are Büchi, or no state is Büchi,

B-quasi-weak if in all cycles that contain both a Büchi and a non-Büchi state, there is a counter that is incremented and never reset in this cycle.

The weakness property implies that every play in the game associated with this automaton has to eventually stabilize, either in a strongly connected component where all states are Büchi (so the play is winning for Eve), or in a strongly connected component where no state is Büchi (so the play is winning for Adam). In fact, this stabilization occurs after at most $|Q|$ -many *changes of mode* between Büchi states and non-Büchi states (where Q is the set of states of the automaton). These automata were studied in [27], and have the same expressive power as WCMO.

Similarly, the quasi-weakness property implies that any play that does not stabilize after kn -changes of mode (for some constant k depending on the automaton) has a counter with value greater than n . Such plays cannot be *n-winning* for Eve, independent of any other consideration. Quasi-weak automata were introduced in [18]. They are strictly more expressive than weak automata (i.e., WCMO), but not as expressive as general cost automata (and in particular not as expressive as full CMSO).

Thus, the difference between these models is in the number of allowed mode changes: unrestricted for B-Büchi automata; bounded by some function of n for B-quasi-weak automata; and bounded by some constant for B-weak automata.

4. Equivalences

4.1 The main theorem and the general approach

Our goal is to show the equivalence of the various models of quasi-weak logics and automata we have introduced, as stated by the following main theorem.

Theorem 2. *Let f be a cost function over infinite trees. Then the following statements are equivalent:*

- f is recognizable by a 1-way B-quasi-weak automaton;
- f is recognizable by a 2-way B-quasi-weak automaton;
- f is definable in QWCMO.

Moreover, the translations between these formalisms are effective. We call such a cost function f a *quasi-weak cost function*.

Since $f \preceq g$ is decidable for functions f, g defined by 1-way B-quasi-weak automata [18], we obtain the following decidability result as a corollary.

Corollary 3. *Given quasi-weak cost functions f and g over infinite trees, it is decidable whether or not $f \preceq g$.*

We now describe our approach for proving Theorem 2. The translation from B-quasi-weak automata to the logic is standard, so we concentrate on the translation from QWCMO to B-quasi-weak automata, which goes via 2-way automata.

We use a well-known technique to enable automata to refer to free set variables. Given $t \in \mathcal{T}_{\mathbb{A}}$ and sets of nodes E_1, \dots, E_k , we write (t, E_1, \dots, E_k) for the tree over the alphabet $\mathbb{A} \times \{0, 1\}^k$ obtained from t by labeling each node v by k bits of extra information, the i -th bit being 1 if $v \in E_i$ and 0 otherwise. First-order variables are treated as singleton sets. Using this encoding, every statement relating automata to logic can be used in a context with free variables, so we can consider cost automata \mathcal{A} that correspond to formulae with free set variables \bar{X} . Given values \bar{E} for these variables, $\llbracket \mathcal{A} \rrbracket(t, \bar{E})$ denotes the evaluation of the automaton on the tree (t, \bar{E}) .

As usual, translating formulae to automata amounts to showing the closure of the automaton model under operations that simulate the constructions of the logic. For simulating disjunction and conjunction, we must show closure under taking the minimum and maximum (respectively) of the functions computed by B-quasi-weak automata. This is obvious for alternating automata. B-quasi-weak automata must also be provided for the atomic predicates. All of this is standard.

There are three non-trivial constructs: weak existential quantifiers, weak universal quantifiers, and bounded expansions. Essentially, by adapting the proof from [27], one can show that 1-way B-quasi-weak automata have the closure properties corresponding to the weak quantifiers (it corresponds to the closure under weak inf- and sup-projection of 1-way B-weak automata in [27]).

The natural way to deal with bounded expansion is to perform it on localized 2-way B-quasi-weak automata, so it remains to show how to (i) translate a 1-way B-quasi-weak automaton into a localized 2-way B-quasi-weak automaton (Section 4.3), (ii) construct from this a localized 2-way B-quasi-weak automaton for the bounded expansion (Section 4.2), and (iii) translate a (localized) 2-way B-quasi-weak automaton into a 1-way B-quasi-weak automaton (Section 4.4).

4.2 The bounded expansion operation on automata

The bounded expansion operator is applied to a formula $\varphi(x, Y)$ which is syntactically monotonic in Y . Testing if

$$z \in \mu^N Y. \{x : \varphi(x, Y)\}$$

can be viewed as a game which starts in node $x = z$ with a given value n for the number of iterations and that proceeds in turns as follows:

- Eve chooses some set Y such that $\varphi(x, Y)$ holds (if it is not possible, she loses), then
- Adam chooses some element $x \in Y$ (if it is not possible, he loses), and the game proceeds to the next turn.

If the game exceeds n turns, Adam is declared the winner. It is straightforward to check that Eve wins this game if and only if $z \in \mu^N Y. \{x : \varphi(x, Y)\}$ for $N \leq n$.

We implement this idea by taking an automaton \mathcal{A} equivalent to $\varphi(x, Y)$, and transforming it into a new automaton that simulates the game. For \mathcal{A} we take a localized 2-way B-quasi-weak automaton with one free variable Y . We assume priorities label states. This automaton n -accepts a tree (t, Y) starting from position x if $t, n \models \varphi(x, Y)$ (in fact, modulo \approx).

The construction outputs a new localized 2-way B-quasi-weak automaton \mathcal{B} with no free variables, which n -accepts a tree t from position x iff Eve wins the above game from initial position x . This new automaton \mathcal{B} has the same states, the same initial state, the same priority function, and the same counters as \mathcal{A} , along with one additional counter, say, γ . The transition function $\delta_{\mathcal{B}}$ is defined for all states p and all letters a as:

$$\delta_{\mathcal{B}}(p, a) = \delta_{\mathcal{A}}(p, (a, 0)) \vee [\delta_{\mathcal{A}}(p, (a, 1)) \wedge (\odot, \mathbb{I}_{\gamma}, q_0)],$$

where \mathbb{I}_{γ} resets all counters of \mathcal{A} and increments the counter γ . This transition function expresses that Eve is required to say whether she uses the assumption that the current a -labelled position, say y , belongs to the set Y . If she assumes it does not, she can take the transition allowed in \mathcal{A} for nodes not in Y (the first disjunct in the transition function above). If she assumes $y \in Y$ (the second disjunct), Adam and Eve can continue to play in the current round according to $\delta_{\mathcal{A}}(p, (a, 1))$ (intuitively, this corresponds to Adam checking that $t, n \models \varphi(x, Y)$ truly holds). Otherwise, Adam can ask to use y as his choice in the game by using the transition $(\odot, \mathbb{I}_{\gamma}, q_0)$. This advances one turn in the game, so the new

counter is incremented, all counters of \mathcal{A} are reset, and the automaton \mathcal{A} is restarted in state q_0 at y .

Theorem 4. *For every localized 2-way quasi-weak automaton \mathcal{A} with one free set variable, there exists a localized 2-way quasi-weak automaton \mathcal{B} with no free variables such that*

$$\llbracket \mathcal{B} \rrbracket_z(t) \approx^{t,z} \llbracket z \in \mu^N X. \{x : \llbracket \mathcal{A} \rrbracket_x(X) \leq N \} \rrbracket(t).$$

4.3 From 1-way to localized 2-way

Translating a 1-way automaton into an equivalent 2-way automaton is straightforward. The subtlety here is that we need to transform a 1-way automaton that has a first-order variable x as input, *i.e.*, reading the word $(t, \{x\})$ into a localized 2-way automaton that starts from node x , but reads only t .

Theorem 5. *For every 1-way B-quasi-weak automaton \mathcal{A} with one free first-order variable, there exists a localized 2-way B-quasi-weak automaton \mathcal{A}^{ℓ} with no free variables such that*

$$\llbracket \mathcal{A} \rrbracket(t, \{x\}) \approx^{t,x} \llbracket \mathcal{A}^{\ell} \rrbracket_x(t).$$

The difficulty here is that the automaton \mathcal{A} over the input $(t, \{x\})$ may cross the node x several times (on different plays), and make use of the fact that x is labeled. This is problematic since a localized automaton cannot remember this position, so the automaton would not know if/when it returns to x .

This would not be a problem if \mathcal{A} were nondeterministic. Unfortunately, \mathcal{A} is alternating, and making it nondeterministic would leave the class of quasi-weak automata (this situation occurs already for weakly definable languages). There is a known solution to this problem (due to Muller, Saoudi, and Schupp [22]), which is to make the automaton \mathcal{A} nondeterministic, but only on a finite portion of the tree containing both the root and x . We are able to accomplish this using machinery already present in the proof for the closure of 1-way cost automata under the weak existential quantifier.

4.4 From 2-way to 1-way

We now turn to the key technical contribution of the paper: transforming 2-way B-quasi-weak automata into their 1-way version. We remark that it is trivial to transform a localized 2-way B-quasi-weak automaton into an equivalent non-localized one, so the localized property is irrelevant here.

Before proceeding, let us briefly review the construction from Vardi [28] that transforms a 2-way parity automaton \mathcal{A} into an equivalent 1-way parity automaton \mathcal{B} (this result can also be deduced from other constructions like the unfolding or iteration). The behavior of \mathcal{A} on t can be represented as a parity game. By positional determinacy of parity games, Eve has a positional winning strategy if \mathcal{A} accepts t . Moreover, for any position $x \in \mathcal{T}$, loops (*i.e.*, finite paths from x to x) in this strategy can be summarized by their starting state, ending state, and maximum priority. The 1-way version guesses a labeling of t with a positional strategy together with these loop summaries, and then runs a 1-way deterministic parity automaton that checks that the labeling is valid and that every play consistent with the strategy satisfies the parity condition. The loop summaries are used to avoid backtracking.

The above approach fails in our case for two reasons. First, there is no known result of positional or finite memory determinacy for the games produced by 2-way B-quasi-weak automata (the results are known for acyclic arenas like those produced by 1-way B-quasi-weak automata [18], but the arenas produced by 2-way automata may be cyclic). Second, the construction described above naturally outputs a nondeterministic automaton, and hence the output automaton cannot be quasi-weak (since quasi-weak and weak nondeterministic automata are strictly less expressive than their al-

ternating versions). Thus, we have to use a less direct approach to prove the following theorem.

Theorem 6. *Given an alternating 2-way B-parity automaton \mathcal{A}_2 , there effectively exists an alternating 1-way B-parity automaton \mathcal{A}_1 such that*

$$\llbracket \mathcal{A}_2 \rrbracket(t) \approx^t \llbracket \mathcal{A}_1 \rrbracket(t).$$

Moreover, if \mathcal{A}_2 is quasi-weak, then \mathcal{A}_1 is also quasi-weak.

The proof consists of four steps.

(1) We first describe a way to summarize the history of a play of \mathcal{A}_2 . The idea is that the summary of a play collapses loops (finite paths that start and end at the same position in the tree) in this play to a single move described by a global action. This global action records the maximum priority, and (for each counter) whether the counter was reset at least once, incremented at least once but not reset, or left unchanged.

We then prove that Eve always has a *summary-dependent strategy*, i.e., a strategy where Eve’s choices depend only on the summary of the history of the play rather than the full history of the play.

Lemma 7. *Assume that \mathcal{A}_2 n -accepts a tree t for some $n \in \mathbb{N}$. Then Eve has an $\alpha(n)$ -winning summary-dependent strategy, for α independent of t .*

This summary-dependent strategy is constructed from an arbitrary n -winning strategy for Eve. In case there are several loops in the original strategy with the same summary, the summary-dependent strategy chooses the loop in which the counters had the greatest value, and then continues from there. This ensures that the counter values in the resulting summary-dependent strategy will not grow too much (and in fact are bounded by $\alpha(n)$ for some correction function α independent of the input tree). The parity condition will also be satisfied, since replacing loops by other loops having the same maximal priority, does not change whether or not the play is winning. We give the details of this proof in the next subsection.

(2) We construct a new alternating 2-way automaton $\mathcal{A}_{1.5}$ which never goes up (but may stay in the same node). It is not a B-parity automaton because it uses a more complicated acceptance condition, described below.

In general, the automaton $\mathcal{A}_{1.5}$ simulates the run of \mathcal{A}_2 on some t . However, when \mathcal{A}_2 wants to go down from some node x , Eve has to make some declarations about the two parts of the rest of the game: the part of the game before returning to x (i.e., the part in the subtree of t below x), and the part after returning to x (which may visit the children of x again). For the first part she gives a set D of *constraints*, that is of pairs (c, q) that assert that there is a play compatible with her strategy which returns to x , finishing in state q and performing global action c . To describe the second part, Eve declares another set $C_{c,q}$ of constraints for each pair $(c, q) \in D$. This set $C_{c,q}$ describes the different ways her plays could go up to the parent of x if she returns to x in state q after performing global action c . In other words, Eve is declaring relevant information about loops and upward moves, to help us simulate the operation of \mathcal{A}_2 without actually moving upwards in t .

Indeed, after Eve makes these declarations, Adam can decide to verify either the part after a loop (in which case he stays in the same node) or he can decide to verify the assumptions about a loop (in which case he moves down in the tree). Later, when \mathcal{A}_2 wants to go up, we just check whether the current summary and state are in the previously declared set of constraints, and immediately win or lose in $\mathcal{A}_{1.5}$ based on this.

What is the winning condition of $\mathcal{A}_{1.5}$? Whenever $\mathcal{A}_{1.5}$ directly simulates \mathcal{A}_2 , we output the same actions. When $\mathcal{A}_{1.5}$ follows some declared loop with global action c , we just output c . The

sequence of actions output in this way should be bounded by a number n . But this is not enough, since it does not bound the values on paths of \mathcal{A}_2 which go up in the tree. To deal with such paths, whenever $\mathcal{A}_{1.5}$ moves down, we also output Eve’s declarations after coming back from a loop (based on the sets $C_{c,q}$ for each $(c, q) \in D$ described above). We can view these declarations as a graph. The graph connects the pairs responsible for returning to a node x with the pairs describing moves up to the parent of x ; an edge in the graph is labeled by the global action of such a path going up. At each moment when we go down in $\mathcal{A}_{1.5}$ we output a slice of such a graph, and the winning condition requires that in the whole graph constructed from such slices, each path should have value bounded by n . As mentioned earlier, this ensures that the value coming from upward paths in the plays consistent with Eve’s strategy is also bounded.

Notice that in this construction, Eve must make the same declarations after coming up from two loops that have the same summary. Thus, to show that \mathcal{A}_2 and $\mathcal{A}_{1.5}$ are equivalent, it is important that Eve has a summary-dependent strategy in \mathcal{A}_2 , as obtained in point (1).

(3) To replace our winning condition by a B-parity condition, we form a product of $\mathcal{A}_{1.5}$ with an automaton recognizing the winning condition of $\mathcal{A}_{1.5}$. Usually, such a construction would use a deterministic automaton. Unfortunately, cost automata cannot always be determinized [8] so this is not possible. However, every cost automaton can be made “history-deterministic”, which is enough to ensure that it can be run on every branch of the game tree of $\mathcal{A}_{1.5}$ without causing conflicts. We refer the interested reader to [8].

(4) Finally, we eliminate from $\mathcal{A}_{1.5}$ moves using the \odot direction. For each state q , letter a , and goal set G of downward moves, we consider a *local game* $\mathcal{G}(q, a, G)$ describing the local (\odot direction) moves that are possible before moving downwards. Downward moves are terminal positions in the game: they are winning if they are in the goal set G , and losing otherwise. It can be shown that it is decidable whether Eve wins such a game. The idea is to transform this B-parity game into a game with a Streett winning condition, and solve the resulting Streett game using standard methods (see [26] for more details). The desired 1-way automaton \mathcal{A}_1 has the same input alphabet, set of states, set of counters, and initial state as $\mathcal{A}_{1.5}$. For a state q and input letter a , its transition function is defined as follows: Eve chooses a goal set G such that she wins in $\mathcal{G}(q, a, G)$; then Adam chooses any transition $(d, c, q') \in G$ and performs it.

Combining steps (1)–(4), we obtain an alternating 1-way B-parity automaton \mathcal{A}_1 that is equivalent to the original alternating 2-way B-parity automaton \mathcal{A}_2 (up to \approx). A closer examination of the construction shows that quasi-weakness is preserved.

4.5 Summary-dependent strategies (Proof of Lemma 7)

As mentioned earlier, it is an open problem whether in each B-parity game that is n -winning for Eve, she has a finite memory strategy that is $\alpha(n)$ -winning, for some correction function α that is independent of the size of the game. In this section, we seek to prove a weaker property, saying that there exists a strategy which depends only on a summary of the play (Lemma 7).

Let $\mathcal{A}_2 = \langle Q, \mathbb{A}, q_0, \Gamma, \delta \rangle$ be a 2-way B-parity automaton. The transition function δ has type

$$Q \times \mathbb{A} \rightarrow \mathcal{B}^+(\{\uparrow, \swarrow, \searrow, \odot\} \times Act \times Q).$$

The set of actions Act gathers counter actions and parity ranks: $Act = \mathbb{C} \times \{i, i+1, \dots, j\}$, where $\mathbb{C} = \{ic, \varepsilon, \tau\}^\Gamma$ and $i \leq j$ are natural numbers. For simplicity, we will assume that the automaton \mathcal{A}_2 satisfies the following conditions.

- We assume that the counter actions are *hierarchical*, which means that when some counter is incremented or reset, then simultaneously all counters with higher numbers are reset, *i.e.*, every action is of the form

$$(\varepsilon, \dots, \varepsilon, \mathbf{r}, \dots, \mathbf{r}) \quad \text{or} \quad (\varepsilon, \dots, \varepsilon, \mathbf{ic}, \mathbf{r}, \dots, \mathbf{r}).$$

It is known that every B-parity automaton can be converted into an equivalent one using only hierarchical actions [8].

- We assume that our alphabet is a product $\mathbb{A} = \mathbb{A}' \times \{0, 1\}$, and that each input tree t has the root marked by 1 in the second coordinate; additionally we assume that \mathcal{A}_2 never tries to go in the \uparrow direction from the root of the tree. This assumption does not decrease the generality of the result.
- We assume that the states in \mathcal{A}_2 can be partitioned as follows.
 - *d-states* $q \in Q_d$ for $d \in \{\uparrow, \swarrow, \searrow\}$ are such that for all $a \in \mathbb{A}$, $\delta(q, a)$ is a single transition performing action ε in direction d .
 - *Universal states* $q \in Q_\wedge$ are such that for all a , $\delta(q, a)$ is a conjunction of transitions staying in the same node (direction \circ).
 - *Existential states* $q \in Q_\vee$ are such that for all a , $\delta(q, a)$ is a disjunction of transitions staying in the same node (direction \circ).

Moreover, we assume that in the game (\mathcal{A}_2, t) (for each tree t), there is some move possible from each position, and between each pair of positions there is at most one move. It is always possible to achieve this normal form, by using additional intermediate states.

We are now ready to analyze the operation of \mathcal{A}_2 . The set of actions *Act* in the transition function is naturally equipped with a product operation, describing the *global action* of a sequence of actions. This product is defined component-wise: on the part from $\{i, i+1, \dots, j\}$, it corresponds to the maximum of the priorities, and on the part from $\{\mathbf{ic}, \varepsilon, \mathbf{r}\}$, it corresponds to the maximum according to the order $\varepsilon \leq \mathbf{ic} \leq \mathbf{r}$. Therefore, if c and c' are actions in *Act*, we can talk about the resulting action cc' . We use the symbol ε also to denote the neutral element of this product, that is $((\varepsilon, \dots, \varepsilon), i)$. We define the *value* of a finite sequence of counters as the maximum value of any counter in any moment while executing this sequence of actions. Recall that the value of an infinite sequence of actions takes into account also the parity condition. When we have a path whose edges are labeled by actions, we define the global action or the value of this path as the global action or the value of the sequence of actions on this path.

Notice that the global action contains all pertinent information about priorities, but loses some information about counter values, namely the value of a path, since for instance \mathbf{ic}^n is contracted to \mathbf{ic} . This means that we will have to be able to retrieve this information in some way when doing such a contraction.

We must now define a *summary* of a play. Formally, let

$$q_0, (d_1, c_1, q_1), (d_2, c_2, q_2), \dots, (d_m, c_m, q_m)$$

be an initial fragment of a play in the game (\mathcal{A}_2, t) . Let x_i be the node in the tree after following the directions $d_1 d_2 \dots d_i$ starting from the root. The summary starts from q_0 . Then for $i = 1, 2, \dots, m$ we proceed as follows. If x_{i-1} is such that all x_j for $j \geq i$ are proper descendants of x_{i-1} , then we simply append (d_i, c_i, q_i) to the summary. Otherwise, let $j \geq i$ be the smallest index for which $x_j = x_{i-1}$. Then to the summary we append $(\swarrow, c_i c_{i+1} \dots c_j, q_j)$ (on the second coordinate we have the product of the actions), and we continue generating the summary from $i := j+1$. For example if $d_1 \dots d_7 = \swarrow \swarrow \swarrow \uparrow \uparrow \circ \swarrow$, the summary

is

$$q_0, (\swarrow, c_1, q_1), (\swarrow, c_2 c_3 c_4 c_5, q_5), (\circ, c_6, q_6), (\swarrow, c_7, q_7).$$

Notice that in a summary we never use the \uparrow direction.

A strategy (of Eve) is called *summary-dependent* if the choices of Eve depend only on the summary of the history of the play. We will show that it is enough to consider summary-dependent strategies.

We are now ready to prove the following strengthening of Lemma 7:

Assume that \mathcal{A}_2 n -accepts a tree t for some $n \in \mathbb{N}$. Then Eve has an $(n+1)^k$ -winning summary-dependent strategy, where k is the number of hierarchical counters in \mathcal{A}_2 .

Proof. Let t be an \mathbb{A} -labeled binary tree, and σ an n -winning strategy in the game (\mathcal{A}_2, t) . Let t_σ be the strategy tree describing all plays compatible with σ . Each node of t_σ is labeled by a position in (\mathcal{A}_2, t) , which is a pair: a node of t and a state. Recall that, due to our normalization of \mathcal{A}_2 , two consecutive positions of the game determine the action of the move between these positions. To each node x of t_σ we assign a tuple $v_\sigma(x) = (n_1, \dots, n_{|\Gamma|})$ of values of the counters after performing the actions on the path from the root of t_σ to x . We will be comparing such tuples lexicographically. For a summary $\eta = q_0, (d_1, c_1, q_1), \dots, (d_m, c_m, q_m)$ we define *node*(η) to be the node of t reachable by following the directions $d_1 d_2 \dots d_i$ starting from the root (where \swarrow corresponds to staying in the same node).

First, some summaries $\eta = q_0, (d_1, c_1, q_1), \dots, (d_m, c_m, q_m)$ are assigned a node $g(\eta)$ of t_σ , labeled by $(\text{node}(\eta), q_m)$. This is done by induction on m . If $m = 0$, as $g(\eta)$ we just take the root of t_σ . Otherwise, let $\eta' := q_0, (d_1, c_1, q_1), \dots, (d_{m-1}, c_{m-1}, q_{m-1})$. If $g(\eta')$ is not defined, we also leave $g(\eta)$ undefined. Otherwise, assume first that $d_m \neq \swarrow$. If $g(\eta')$ has a child y labeled by $(\text{node}(\eta), q_m)$ (by our normalization assumption, there is at most one such child), we take $g(\eta) := y$. Otherwise, we leave $g(\eta)$ undefined. The other possibility is that $d_m = \swarrow$. Consider the set Y of all descendants y of $g(\eta')$ which are labeled by $(\text{node}(\eta), q_m)$, such that all nodes on the path from $g(\eta')$ to y , excluding $g(\eta')$ and y , are labeled by proper descendants of $\text{node}(\eta)$, and that the path from $g(\eta')$ to y has global action c_m . As $g(\eta)$ we take a node $y \in Y$ such that $v_\sigma(y) = \max_{z \in Y} v_\sigma(z)$, where the max-operator refers to the lexicographic order. If there are multiple nodes $y \in Y$ with maximal $v_\sigma(y)$, we can take any of them. If Y is empty, we leave $g(\eta)$ undefined. Observe that in both cases, $g(\eta)$ (if defined) is a descendant of $g(\eta')$, and the global action of the path from $g(\eta')$ to $g(\eta)$ is c_m .

Next, we construct a summary-dependent strategy ρ . When the history of the play has summary η , we simply look at the node $g(\eta)$, and we move in the same way as σ . Notice that $g(\eta)$ is labeled by our current position, so this move is legal. If $g(\eta)$ is undefined, we can move in any way; we will see below that when playing according to ρ we will never reach such situation.

It remains to see that ρ is winning. Take any play $\pi = q_0, (d_1, c_1, q_1), (d_2, c_2, q_2), \dots$ consistent with ρ . Let π_i be the prefix of π of length i , and let η_i be its summary. For each i , let y_i denote the child of $g(\eta_{i-1})$ such that the transition from $g(\eta_{i-1})$ to y_i is (d_i, c_i, q_i) . Such a node exists (assuming that $g(\eta_{i-1})$ is defined): recall that the position after π_{i-1} is the same as in $g(\eta_{i-1})$; either it is a position of Eve and (d_i, c_i, q_i) is the move to the only child of $g(\eta_{i-1})$, or it is a position of Adam and we have children corresponding to all moves from this position.

Now, by induction on $i > 0$ we will show that $g(\eta_i)$ is defined, and that $v_\sigma(g(\eta_i)) \geq v_\sigma(y_i)$. Let

$$q'_0, (d'_1, c'_1, q'_1), \dots, (d'_m, c'_m, q'_m) := \eta_i, \\ \eta' := q'_0, (d'_1, c'_1, q'_1), \dots, (d'_{m-1}, c'_{m-1}, q'_{m-1}).$$

We have two possibilities. First assume that $d'_m \neq \smile$. Then we simply have $\eta_{i-1} = \eta'$, and we see that $g(\eta_i) = y_i$. The other possibility is that $d'_m = \smile$. Then η' is a prefix of η_{i-1} , so $g(\eta_{i-1})$ (and y_i) is a descendant of $g(\eta')$, and the global action of path from $g(\eta')$ to y_i is c'_m . This means that y_i belongs to the set Y used to define $g(\eta_i)$. So $g(\eta_i)$ is defined, and $v_\sigma(g(\eta_i)) \geq v_\sigma(y_i)$.

Next, we observe the value of counter k . Consider k -tuples $v_k(x)$ which are defined as $v_\sigma(x)$, but contain only the values of the first k counters. Notice that $v_\sigma(x) \leq v_\sigma(y)$ implies $v_k(x) \leq v_k(y)$. If c_i increments counter k , then all counters with smaller numbers are not changed (recall that in \mathcal{A}_2 we only have hierarchical actions), so $v_k(g(\eta_{i-1})) < v_k(y_i) \leq v_k(g(\eta_i))$. And if the k -th counter is not changed by c_i , then all counters with smaller numbers are unchanged as well, so $v_k(g(\eta_{i-1})) = v_k(y_i) \leq v_k(g(\eta_i))$. We have only $(n+1)^k$ tuples with k values between 0 and n . This means that in π we cannot have a fragment where the k -th counter is incremented more than $(n+1)^k$ times and never reset. Thus the value of the counter is always at most $(n+1)^k$.

It remains to verify the parity condition. We can construct the summary also for the infinite play π ; it is an infinite sequence $\eta = q'_0, (d'_1, c'_1, q'_1), (d'_2, c'_2, q'_2), \dots$. Consider the (unique) sequence j_0, j_1, j_2, \dots such that η_{j_i} is the prefix of η of length i (in other words, j_i is the length of the prefix π that yields the summary η_{j_i} of length i). Notice that the global action of the fragment of π between the j_{i-1} -th and j_i -th step has global action c'_i , and also the path from $g(\eta_{j_{i-1}})$ to $g(\eta_{j_i})$ has global action c_i . It follows that the maximal priority appearing infinitely often in π is the same as in the branch of t_σ containing all $g(\eta_{j_i})$. All branches of t_σ are winning, so this priority is even. \square

5. The mu-calculus view of quasi-weakness

In this section, we look at quasi-weakness from the point of view of the μ -calculus. This section is independent of the rest of the paper. We assume the reader to have some familiarity with the μ -calculus (see, e.g., [1]).

It is well-known that μ -calculus is equivalent to alternating automata. The simplicity of the corresponding translations implies that many properties can be expressed equally well at the automaton level and at the logic level. For instance, the structure of the acceptance condition of the automaton is tightly related to the nesting structure of fixed points in μ -calculus formulae: least fixed points correspond to odd priorities, and greatest fixed points to even priorities. Moreover, the number of priorities in the automaton reflects the nesting of μ and ν operators in the formula. As a special case that we are particularly interested in, the *alternation free* fragment of μ -calculus (see below) corresponds exactly to weak alternating automata. We can extend some of these relationships to cost functions.

Let us recall the definition of the μ -calculus. We assume an infinite set of variables X, Y, Z, \dots ranging over sets of nodes. The semantics of a μ -calculus formula with free variables \bar{X} is given by a monotonic function from tuples of sets indexed by \bar{X} to sets. The following constructs are allowed. There are *modal operators* that consist here of constant modalities a , for all letters a , and successor modalities $\mathbf{X}(E, F)$. A modality of the first kind, evaluates to the set of nodes carrying the letter a , while a modality of the second kind returns, given sets of nodes E and F , the set of nodes such that the left child is in E and the right child in F . We also have boolean operators \vee, \wedge, \neg and fixed point operators. The *least fixed*

point $\mu X.F(X, \bar{V})$ is the least set E such that $E = F(E, \bar{V})$ and the *greatest fixed point* $\nu X.F(X, \bar{V})$ is the greatest such set. Every μ -calculus formula with free variables \bar{X} *evaluates* to a mapping from a tree t and a tuple of sets of nodes indexed by \bar{X} to a set of nodes.

The μ^N -calculus extends μ -calculus in a way similar to CMSO, meaning that the semantics of a formula is parametrized by some non-negative integer n . So a μ^N -calculus formula *n-evaluates* to a function. One adds the new construct $\mu^N X.\psi(X)$ which evaluates to $F_n^{(n)}$, assuming that ψ *n-evaluates* to F_n for all n (where $F^{(n)}$ is defined as in Section 2.3).

Given a μ^N -calculus formula ψ without free variables, it defines a function $\mathcal{T}_\mathbb{A} \rightarrow \mathbb{N}_\infty$ by

$$\llbracket \psi \rrbracket(t) = \inf \{n : \psi \text{ n-evaluates to } E \text{ with } \epsilon \in E\},$$

where ϵ is the root of t .

Example 8. Let $\mathbb{A} := \{a, b\}$. The formula

$$\mu^N X.(b \vee (a \wedge \mathbf{X}(X, X)))$$

computes the maximum number of consecutive a 's that occur on some branch starting in the root (the same cost function as Example 1). The idea is that we can start evaluating at the root. If the current node is a -labelled, then both successors are forced to “loop” back to X , and continue as before. Because the fixed point operator corresponding to X is a μ^N , we count these unfoldings. The least number of unfoldings that are needed is exactly the maximum number of consecutive a 's on some branch starting in the root.

Now consider the formula

$$\mu^N X.\nu Y.[(b \wedge \mathbf{X}(Y, Y)) \vee (a \wedge \mathbf{X}(X, X))]$$

which computes the maximum number of a 's (not necessarily consecutive) that occur on some branch. As above, an a -labelled node forces a loop back on X for both successors, and these unfoldings are counted. However, when a b -labelled node is encountered the successors loop back to Y . This fixed point variable Y corresponds to a ν operator, so these unfoldings are not counted, and indeed can be taken infinitely many times. Because the ν operator is nested inside of μ^N , looping back on Y does not “reset” the count of the number of unfoldings we have used for X , so this formula computes the desired cost function.

Finally, the formula

$$\mu^N X.\nu Y.[(b \wedge (\mathbf{X}(Y, \top) \vee \mathbf{X}(\top, Y))) \vee \\ (a \wedge (\mathbf{X}(X, \top) \vee \mathbf{X}(\top, X)))]$$

computes the minimal number of a 's that occur on some branch (\top is an abbreviation for $\nu Z.Z$). This is similar to the previous example, except that only a single branch is picked out during the evaluation of the formula (intuitively, the branch with the minimum number of a 's).

So far, we have described μ^N -calculus in its most general form, which is equivalent to B-parity automata. However, we are interested in the quasi-weak form of such automata. The nice thing about μ -calculus is that, again, quasi-weakness has a natural interpretation. Renaming bound variables, we may ensure that no variable is used in two distinct fixed points (μ, ν or μ^N) in a formula. Then the scope of the variable is the set of nodes of the formula (seen as a tree) that are below the fixed-point operator binding the variable, and above some use of the variable. If the variable is introduced with a μ (resp. ν or μ^N), we call it a μ -scope (resp. a ν -scope or a μ^N -scope). A μ^N -calculus formula is

weak if no ν -scope intersects a μ -scope, and no μ^N -scope simultaneously intersects a μ -scope and a ν -scope,

quasi-weak if no ν -scope intersects a μ -scope.

Note that for μ -calculus formulae without μ^N operators, the two definitions coincide. The classical equivalence between automata and μ -calculus formulae carries over without surprise to cost functions (the proofs are essentially the same).

Theorem 9. *In terms of recognizing cost functions,*

- *B-parity automata are effectively equivalent to μ^N -formulae,*
- *B-weak automata are effectively equivalent to weak μ^N -formulae,*
- *B-quasi-weak automata are effectively equivalent to quasi-weak μ^N -formulae.*

What is interesting about these notions is that the definition of quasi-weakness is significantly simpler than the one of weakness. Moreover, the definition for a quasi-weak μ^N -formula is the standard definition of the alternation-free fragment of the μ -calculus. This supports the fact that, in the cost setting, quasi-weakness is the correct counterpart to the notion of weakness from the classical setting.

6. Conclusion

We have described some natural logics that correspond to the class of quasi-weak cost functions over infinite trees, a robust and natural class of cost functions. Along the way, we have developed some technical tools, like summary-dependent strategies and the conversion of 2-way cost automata to 1-way cost automata.

Overall, the hope is that these results and tools may eventually prove useful for solving full cost monadic logic over infinite trees, and for solving challenging boundedness questions like the parity index problem and weak definability problem.

References

- [1] A. Arnold and D. Niwiński. *Rudiments of μ -Calculus*, volume 146 of *Studies in Logic and The Foundations of Computer Science*. North-Holland, 2001.
- [2] M. Benedikt, G. Puppis, and C. Riveros. Regular repair of specifications. In *LICS*, pages 335–344. IEEE Computer Society, 2011. ISBN 978-0-7695-4412-0.
- [3] A. Blumensath, M. Otto, and M. Weyer. Decidability Results for the Boundedness Problem. *Logical Methods in Computer Science*. to appear.
- [4] A. Blumensath, M. Otto, and M. Weyer. Boundedness of Monadic Second-Order Formulae over Finite Words. In *Proc. 36th Int. Colloquium on Automata, Languages and Programming, ICALP, Part II, LNCS 5556*, pages 67–78, 2009.
- [5] M. Bojańczyk. A bounding quantifier. In J. Marcinkowski and A. Tarlecki, editors, *CSL*, volume 3210 of *LNCS*, pages 41–55. Springer, 2004. ISBN 3-540-23024-6.
- [6] M. Bojańczyk and T. Colcombet. Bounds in ω -regularity. In *LICS*, pages 285–296. IEEE Computer Society, 2006.
- [7] J. R. Büchi. Weak second-order arithmetic and finite automata. *Z. Math. Logik Grundlagen Math.*, 6:66–92, 1960.
- [8] T. Colcombet. Regular cost functions over words, 2009. Manuscript at <http://www.liafa.jussieu.fr/~colcombe/>.
- [9] T. Colcombet. The theory of stabilisation monoids and regular cost functions. In *ICALP (2)*, volume 5556 of *LNCS*, pages 139–150. Springer, 2009. ISBN 978-3-642-02929-5.
- [10] T. Colcombet. Regular cost functions, part i: Logic and algebra over words. *Logical Methods in Computer Science*, 9(3), 2013.
- [11] T. Colcombet and C. Löding. The nesting-depth of disjunctive μ -calculus. In *CSL*, volume 5213 of *LNCS*, pages 416–430. Springer, 2008. ISBN 978-3-540-87530-7.
- [12] T. Colcombet and C. Löding. The non-deterministic Mostowski hierarchy and distance-parity automata. In L. Aceto, I. Damgard, L. A. Goldberg, M. M. Halldórsson, A. Ingólfssdóttir, and I. Walukiewicz, editors, *ICALP (2)*, volume 5126 of *LNCS*, pages 398–409. Springer, 2008. ISBN 978-3-540-70582-6.
- [13] T. Colcombet and C. Löding. Regular cost functions over finite trees. In *LICS*, pages 70–79. IEEE Computer Society, 2010. ISBN 978-0-7695-4114-3. Online at <http://www.liafa.jussieu.fr/~colcombe/>.
- [14] T. Colcombet, D. Kuperberg, C. Löding, and M. Vanden Boom. Deciding the weak definability of büchi definable tree languages. In S. R. D. Rocca, editor, *CSL*, volume 23 of *LIPICs*, pages 215–230. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2013. ISBN 978-3-939897-60-6.
- [15] K. Hashiguchi. Limitedness theorem on finite automata with distance functions. *J. Comput. Syst. Sci.*, 24(2):233–244, 1982.
- [16] K. Hashiguchi. Relative star height, star height and finite automata with distance functions. In J.-E. Pin, editor, *Formal Properties of Finite Automata and Applications*, volume 386 of *LNCS*, pages 74–88. Springer, 1988. ISBN 3-540-51631-X.
- [17] D. Kirsten. Distance desert automata and the star height problem. *ITA*, 39(3):455–509, 2005.
- [18] D. Kuperberg and M. Vanden Boom. Quasi-weak cost automata: a new variant of weakness. In S. Chakraborty and A. Kumar, editors, *FSTTCS*, volume 13 of *LIPICs*, pages 66–77. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2011. ISBN 978-3-939897-34-7. Online at <http://www.liafa.jussieu.fr/~dkuperbe/>.
- [19] D. Kuperberg and M. Vanden Boom. On the expressive power of cost logics over infinite words. In A. Czumaj, K. Mehlhorn, A. M. Pitts, and R. Wattenhofer, editors, *ICALP (2)*, volume 7392 of *Lecture Notes in Computer Science*, pages 287–298. Springer, 2012. ISBN 978-3-642-31584-8.
- [20] M. Lang. Resource-bounded reachability on pushdown systems. Master’s thesis, RWTH Aachen University, 2011.
- [21] H. Leung. On the topological structure of a finitely generated semi-group of matrices. *Semigroup Forum*, 37:273–287, 1988.
- [22] D. E. Muller, A. Saoudi, and P. E. Schupp. Alternating automata. The weak monadic theory of the tree, and its complexity. In L. Kott, editor, *ICALP*, volume 226 of *LNCS*, pages 275–283. Springer, 1986. ISBN 3-540-16761-7.
- [23] M. O. Rabin. Decidability of second-order theories and automata on infinite trees. *Trans. Amer. Math. Soc.*, 141:1–35, 1969. ISSN 0002-9947.
- [24] M. O. Rabin and D. Scott. Finite automata and their decision problems. *IBM J. Res. Dev.*, 3(2):114–125, Apr. 1959. ISSN 0018-8646. URL <http://dx.doi.org/10.1147/rd.32.0114>.
- [25] I. Simon. Limited subsets of a free monoid. In *FOCS*, pages 143–150. IEEE, 1978.
- [26] W. Thomas. Languages, automata, and logic. In G. Rozenberg and A. Salomaa, editors, *Handbook of Formal Languages*, volume 3, pages 389–455. Springer, 1997.
- [27] M. Vanden Boom. Weak cost monadic logic over infinite trees. In F. Murlak and P. Sankowski, editors, *MFCS*, volume 6907 of *LNCS*, pages 580–591. Springer, 2011. ISBN 978-3-642-22992-3.
- [28] M. Y. Vardi. Reasoning about the past with two-way automata. In K. G. Larsen, S. Skyum, and G. Winskel, editors, *ICALP*, volume 1443 of *Lecture Notes in Computer Science*, pages 628–641. Springer, 1998. ISBN 3-540-64781-3.

A. From 2-way automata to 1-way automata

In this section, we provide the details for the proof of Theorem 6.

We will write B-2PA for two-way B-parity alternating automata, and B-1PA for the one-way variant.

Given a B-2PA $\mathcal{A}_2 = \langle Q, \mathbb{A}, q_0, \Gamma, \delta \rangle$, we want to build an equivalent B-1PA. The transition function δ has type $Q \times \mathbb{A} \rightarrow \mathcal{B}^+(\{\uparrow, \swarrow, \searrow, \circ\} \times Act \times Q)$. The set of actions Act gathers counter actions and parity ranks: $Act = \mathbb{C} \times \{i, i+1, \dots, j\}$, where $\mathbb{C} = \{i, \varepsilon, \tau\}^\Gamma$ and $i \leq j \in \mathbb{N}$.

We assume \mathcal{A}_2 satisfies the assumptions described in Section 4.5.

A.1 Construction of new automaton

As an intermediate step, before building the B-1PA equivalent to \mathcal{A}_2 , we create an automaton $\mathcal{A}_{1.5}$. It will be allowed to go down or stay in the same node, but it will never go up. Moreover it will have a more complicated winning condition than the B-parity condition. Namely, a label of a transition will be not just a vector of counter actions, but a fragment of a graph, whose edges are labeled by counter actions. The winning condition will say that each path in the graph (created from parts output by each transition) should have bounded value.

We first explain the principle of the automaton $\mathcal{A}_{1.5}$. Whenever \mathcal{A}_2 wants to go down from some node x , the rest of the game consists of two parts: the part in the subtree, before going up to x , and the part after coming back to x (notice that in the second part, the automaton might want to go down again). In such a situation Eve will declare her expectations about the two parts. Namely, for the first part she gives a set D of constraints, that is of pairs (c, q) such that there is a play compatible with her strategy which returns to x , finishing in state q and performing global action c . To describe the second part, for each such pair $(c, q) \in D$, Eve also declares a set $C_{c,q}$ of constraints of the same form, now it describes how plays can go up to the parent of x . Later, whenever \mathcal{A}_2 wants to go up, we just check whether our action and state are in the previously declared set of constraints, and we immediately win or lose.

Now we construct the automaton $\mathcal{A}_{1.5} = \langle Q', \mathbb{A}, q'_0, \Gamma, \delta' \rangle$. Its state set is $Q' := Q \times \mathcal{P}(Const) \cup \{\text{win}, \text{lose}\}$, where $Const = Act \times Q$ is the set of constraints. The initial state is $q'_0 := (\varepsilon, q_0, \emptyset)$. In a pair (q, C) , q is the state of \mathcal{A}_2 , and C is the set of constraints declared by Eve for the current node. Transitions will be labeled by pairs from $Act \times Graph$, where $Graph := \mathcal{P}(Const \times Act \times Const)$. An element $G \in Graph$ can be seen as a bipartite graph with edges labeled by Act : its nodes are two copies of $Const$, and a triple $(\alpha, d, \beta) \in G$ means that there is a d -labeled edge to α in the first copy from β in the second copy. These graphs will be used to trace values of paths (of the original automaton) going up in a tree.

For an action $c \in Act$ and set of constraints $C \subseteq Const$ we define set $C|c := \{(c', q) : (cc', q) \in C\}$, which describes how C is modified after performing action c . We have a graph corresponding to such a shift, which is $G_c := \{((cc', q), \varepsilon, (c', q)) : (c', q) \in Const\}$.

Transitions from a pair $((q, C), a)$ are defined as follows.

- For $q \in Q_\vee$ with $\delta(q, a) = \bigvee_{i \in I} (\circ, c_i, q_i)$, let $\delta'((q, C), a) := \bigvee_{i \in I} (\circ, (c_i, G_{c_i}), (q_i, C|c_i))$.
- For $q \in Q_\wedge$ with $\delta(q, a) = \bigwedge_{i \in I} (\circ, c_i, q_i)$, let $\delta'((q, C), a) := \bigwedge_{i \in I} (\circ, (c_i, G_{c_i}), (q_i, C|c_i))$.
- For $q \in Q_\downarrow$ with $\delta(q, a) = (d, \varepsilon, q')$, Eve chooses a set of constraints $D \subseteq Const$, and for each $(c, p) \in D$ she also chooses a set of constraints $C_{c,p} \subseteq C|c$. Then one possibility of Adam is to move to $(d, (\varepsilon, G), (q', D))$, where $G :=$

$\{((cc', p'), c', (c, p)) : (c, p) \in D, (c', p') \in C_{c,p}\}$. Moreover, for each $(c, p) \in D$ Adam can move to $(\circ, (c, G_c), (p, C_{c,p}))$.

- For $q \in Q_\uparrow$ with $\delta(q, a) = (\uparrow, \varepsilon, q')$, we win immediately if $(\varepsilon, q') \in C$, otherwise we lose immediately. Formally, this is realized by a transition to state win or lose, in which we have a loop labeled by $((\varepsilon, \dots, \varepsilon), i, G_\varepsilon)$, where i is, respectively, even or odd.

It remains to define when a play is n -winning. A play gives us an infinite sequence of pairs $(c_1, G_1), (c_2, G_2), \dots$. First, we require that c_1, c_2, \dots has value at most n (which includes that the maximal priority appearing infinitely often is even). Moreover we define a graph whose nodes are $\mathbb{N} \times Const$. There is an edge to $(i-1, \alpha)$ from (i, β) labeled by c if $(\alpha, c, \beta) \in G_i$. We require that for any finite path in this graph, its value does not exceed n .

A.2 From strategy in \mathcal{A}_2 to strategy in $\mathcal{A}_{1.5}$

In this section we prove the equivalence of \mathcal{A}_2 and $\mathcal{A}_{1.5}$ in one direction, as described by the following lemma.

Lemma 10. *Assume that Eve has an n -winning summary-dependent strategy in (\mathcal{A}_2, t) for some tree t . Then $\mathcal{A}_{1.5}$ n -accepts t .*

Fix a tree t , and an n -winning summary-dependent strategy σ in the game (\mathcal{A}_2, t) . Let t_σ be the strategy tree describing all plays compatible with σ . To prove the lemma we have to construct an n -winning strategy ρ in $(\mathcal{A}_{1.5}, t)$.

We are particularly interested in fragments of plays called returns. We say that a fragment (infix) π of a play of (\mathcal{A}_2, t) is a return when the node in which π ends is the parent of the node in which π starts, and a proper ancestor of all nodes visited by π in between. Moreover, when the global action of such π is c , and it ends in state p , we say that it is an (c, p) -return. For a node x of t_σ , and for constraint $\alpha \in Const$ we define a set $ret(x, \alpha)$ as containing all descendants y of x such that the path from x to y describes an α -return.

Recall that a strategy assigns a choice of Eve to every initial fragment of a play, ending in a position of Eve. Simultaneously with defining the strategy ρ , to every initial fragment π of a play consistent with the already defined part of ρ , we will assign a node x_π in t_σ . While doing this, we will ensure that when π ends in position $(v, (q, C))$, then x_π is labeled by (v, q) , and for each β such that $ret(x_\pi, \beta)$ is nonempty, it holds $\beta \in C$ (invariant 1). Moreover, for every extension π' of π by one transition labeled by (c_o, G_o) , consistent with the strategy ρ being defined, the following additional invariants will hold:

2. $x_{\pi'}$ is a proper descendant of x_π , and
3. c_o is the global action of the path from x_π to $x_{\pi'}$, and
4. if $(\alpha, c, \beta) \in G_o$, then for each $y' \in ret(x_{\pi'}, \beta)$ there exists $y \in ret(x_\pi, \alpha)$ which is a descendant of y' and such that the path from y' to y has global action c (allowing $y' = y$ when $c = \varepsilon$); moreover, if $c \neq \varepsilon$, the set $ret(x_{\pi'}, \beta)$ is nonempty.

We proceed by induction on the length of π . We take the root of t_σ as x_π for π of length 0. Now, take π such that x_π is already defined, and ρ is already defined for all plays shorter than π . Let (q, C) be the state of the position at the end of π . We have several possibilities:

- Assume that $q \in Q_\vee$. Then there is exactly one child of x_π , reached by some transition (\circ, c, q') . In this situation in ρ after π we use the transition $(\circ, (c, G_c), (q', C|c))$. Let π' be the extension of π by this transition. We define $x_{\pi'}$ to be the only child of x_π . Because the global action between x_π and $x_{\pi'}$ is c , and because they are labeled by the same node of t , we have $ret(x_{\pi'}, (c', p)) \subseteq ret(x_\pi, (cc', p))$ for each c'

and p . So, if the first set is nonempty, also the second one is nonempty. By the induction assumption (invariant 1 for x_π), $(cc', p) \in C$, so $(c', p) \in C \upharpoonright c$, which gives invariant 1 for $x_{\pi'}$. The next two invariants are trivial. In G_c we only have edges of the form $((cc', p), \varepsilon, (c', p))$. Again because $ret(x_{\pi'}, (c', p)) \subseteq ret(x_\pi, (cc', p))$, we can simply take $y' = y$ in the last invariant. The last part of this invariant is irrelevant, since in G_c we have only edges labeled by ε .

- Assume that $q \in Q_\vee$. There is Adam's turn. Consider any extension π' of π by one transition, which is of the form $(\odot, (c, G_c), (q', C \upharpoonright c))$. Then also x_π has a child, to which the transition (\odot, c, q') is performed. We take this child as $x_{\pi'}$. The invariants are satisfied, as in the previous case.
- Assume that $q \in Q_\downarrow$. In this situation Eve chooses D to be the set of all constraints α such that $ret(x', \alpha)$ is nonempty, where x' is the only child of x_π . Then for each $\alpha \in D$ we fix any $y_\alpha \in ret(x', \alpha)$. Eve chooses C_α to be the set of all constraints β for which $ret(y_\alpha, \beta)$ is nonempty.

Let first π' be the extension of π by the first possible choice of Adam, that is by the transition $(d, (\varepsilon, G), (q', D))$. We define $x_{\pi'} := x'$. By definition, whenever $ret(x_{\pi'}, \alpha)$ is nonempty then $\alpha \in D$, which gives invariant 1. The next two invariants are trivial. Now take any edge of G , it is of the form $((cc', p'), c', (c, p))$ where $(c, p) \in D$ and $(c', p') \in C_{c,p}$. First we see that $ret(x_{\pi'}, (c, p))$ is nonempty by definition of D . Then, take any $y' \in ret(x_{\pi'}, (c, p))$. Notice that the summary of the play described by the path to y' and to $y_{c,p}$ is the same; namely this is the summary of the path to x_π , with (\cup, c, p) appended. Thus, because σ is summary-dependent, the subtrees of σ rooted in y' and in $y_{c,p}$ are isomorphic. In particular, because $ret(y_{c,p}, (c', p'))$ is nonempty, also $ret(y', (c', p'))$ is nonempty. Because the global action between x_π and y' is c , and because they are labeled by the same node of t , we have $ret(y', (c', p')) \subseteq ret(x_\pi, (cc', p'))$. Thus any element $y \in ret(y', (c', p')) \subseteq ret(x_\pi, (cc', p'))$ witnesses the last invariant.

Next, let π' be the extension of π by some of the other possible choices of Adam, namely by a transition $(\odot, (c, G_c), (p, C_{c,p}))$ for some $(c, p) \in D$. We define $x_{\pi'} := y_{c,p}$. Invariants 2-3 are easily satisfied, since $y_{c,p} \in ret(x', (c, p))$. Invariants 1 and 4 are true for the same reasons as in the case $q \in Q_\vee$: again we have graph G_c , and the global action between x_π and $x_{\pi'}$ is c , and they are labeled by the same node of t .

- Assume that $q \in Q_\uparrow$. Let q' be the state of the only child x' of x_π . Notice that $x' \in ret(x_\pi, (\varepsilon, q'))$, and by invariant 1 for x_π we have $(\varepsilon, q') \in C$. Thus the game continues to state win.

This finishes the definition of the strategy ρ . Consider now any (infinite) play π consistent with ρ . We have to show that π is n -winning. Let $(c_1, G_1)(c_2, G_2), \dots$ be the sequence of labels on the transitions of π , and let π_i be the prefix of π of length i . As already observed, π never enters the lose state. When π does not enter the win state, in t_σ we have an infinite path containing all x_{π_i} (invariant 2), and each c_i is the global action of the path between $x_{\pi_{i-1}}$ and x_{π_i} . By contracting counter actions, we cannot increase the maximal value of a counter; additionally, the maximal priority appearing infinitely often stays the same. Since each path in t_σ is n -winning, we obtain that c_1, c_2, \dots has value at most n . When π at some moment enters the win state, we have the same for the prefix before this moment, so the counter values do not grow above n ; then we loop in the win state, using transitions with even priority, which does not change the counters.

Next, consider the graph defined by the sequence G_1, G_2, \dots . Take any (finite) path in this graph; we have to show that its value

does not exceed n . We need not consider paths which start by an ε -labeled edge (adding such edges does not change the value). Denote the nodes on our path by $(m, \beta_m), (m-1, \beta_{m-1}), \dots, (k, \beta_k)$. By the last part of invariant 4 for x_{π_m} , we can find some $y_m \in ret(x_{\pi_m}, \beta_m)$. Then by using invariant 4 for $x_{\pi_{i+1}}$ we show for each $i = m-1, \dots, k$ that there is some descendant $y_i \in ret(x_{\pi_i}, \beta_i)$ of y_{i+1} , such that the global action of the path from y_{i+1} to y_i is the label of the edge from (i, β_i) to $(i+1, \beta_{i+1})$ in our graph. Since the value of the path from y_m to y_k does not exceed n , the same holds for our path in the graph. This finishes the proof of Lemma 10.

A.3 From strategy in $\mathcal{A}_{1.5}$ to strategy in \mathcal{A}_2

In this section, we construct a strategy in \mathcal{A}_2 based on a winning strategy in $\mathcal{A}_{1.5}$. Notice that the winning condition in $\mathcal{A}_{1.5}$ bounds the value of summaries of prefixes of a play (the first coordinate of labels), and the value of fragments of a play going only up (the graph coordinate of labels). In the next section we see that these conditions imply that the value of the whole play is bounded. The result of this section is summarized in the following lemma.

Lemma 11. *Assume that $\mathcal{A}_{1.5}$ n -accepts t . Then there exists a strategy ρ in the game (\mathcal{A}_2, t) such that for each play π consistent with ρ :*

- for each prefix π' of π , the sequence of actions in the summary of π' has value at most n , and
- for each infix of π which is of the form $\pi_1\pi_2 \dots \pi_m$, where each π_i is a return with global action c_i , the sequence c_1, c_2, \dots, c_m has value at most $n+1$, and
- the maximal priority appearing infinitely often in the labels of transitions in π is even.

Fix a tree t , and an n -winning strategy σ in the game $(\mathcal{A}_{1.5}, t)$. Let t_σ be the strategy tree describing all plays compatible with σ . We will construct a strategy ρ in the game (\mathcal{A}_2, t) , which will be summary-dependent.

First, to each node x of t_σ , except nodes labeled by state win, we will assign a summary $s(x)$. The definition is by induction on the level of x (i.e., the length of the path from the root to x). To the root of t_σ we assign the empty summary q_0 . Then take a node x for which $s(x)$ is defined, and its child y . Let (q, \cdot) be the state labeling x , and let $(d, (c, \cdot), (q', \cdot))$ be the transition used to reach y from x . When $q \in Q_\downarrow$, and $d = \odot$, as $s(y)$ we take $s(x), (\cup, c, q')$; otherwise as $s(y)$ we take $s(x), (d, c, q')$. Notice that when $q \in Q_\uparrow$, the state of y is win, so we need not define $s(y)$. It is easy to see that $s(x) \neq s(x')$ for $x \neq x'$.

Next we observe (property $(*)$) that when x is labeled by state (q, C) , and $\pi\pi'$ is a partial play such that π has summary $s(x)$ and π' is a (c', p') -return for $(c', p') \in C$, then the summary of $\pi\pi'$ is equal to $s(z)$ for some node z . Again, we make induction on the level of x . For x being the root this is trivial, since C is empty. Take now a node x for which this property holds, and its child y . Let (q, C) be the state labeling x , and let $(d, (c, \cdot), (q', C'))$ be the transition used to reach y from x . Let also $\pi\pi'$ be a partial play such that π has summary $s(y)$, and π' is a (c', p') -return for $(c', p') \in C'$. The first case is that $q \in Q_\downarrow$ and $d = \odot$, or that $q \in Q_\vee \cup Q_\wedge$. Then $\pi = \pi_1\pi_2$, where π_1 has summary $s(x)$ and $\pi_2\pi'$ is a (cc', p') -return. Because $(c', p') \in C' \subseteq C \upharpoonright c$, it holds $(cc', p') \in C$. By the induction assumption the summary of $\pi_1\pi_2\pi'$ is equal to $s(z)$ for some node z , which is what we want. The other case is that $q \in Q_\downarrow$, and $d \in \{\swarrow, \searrow\}$. Then the summary of $\pi\pi'$ is $s(x), (\cup, c', p')$. Notice that C' is equal to the set D guessed by Eve in the node x . By definition of $\mathcal{A}_{1.5}$, the assumption $(c', p') \in C'$ implies that there exists a child z of x to which we go using transition $(\odot, (c', \cdot), (p', \cdot))$. To this z we have assigned summary $s(x), (\cup, c', p')$.

We define the strategy ρ as follows. Let η be the summary of the history of the play after which Eve has to make a decision. We look for the node x such that $s(x) = \eta$. If it does not exist, we move in any way. If such x exists, we see that it is labeled by $(v, (q, \cdot))$ when we are in position (v, q) ; it has exactly one child, reached by some transition $(\circlearrowleft, (c, \cdot), (q', \cdot))$. Then in ρ we use the transition $(\circlearrowleft, c, q')$.

We can see that for each partial play consistent with ρ , having summary η , there exists a node x such that $s(x) = \eta$. Indeed, it is true for the play of length 0. By induction, take some partial play π ending in position (v, q) , for which such x exists. Whenever $q \in Q_\wedge \cup Q_\downarrow$, and from (v, q) we have a transition (d, c, q') , then x has a child reached by transition $(d, (c, \cdot), (q', \cdot))$; the $s(\cdot)$ of this child gives the summary of π extended by the transition (d, c, q') . Similarly for $q \in Q_\vee$, but now x has only a child corresponding to the transition chosen by Eve according to ρ . The last case is that $q \in Q_\uparrow$; from (v, q) we have a transition $(\uparrow, \varepsilon, q')$. The one-step fragment performing this transition is a (ε, q') -return. Let $(v, (q, C))$ be the label of x ; notice that $(\varepsilon, q') \in C$, since the successor of $(v, (q, C))$ is win (we do not reach lose by the winning strategy σ). Thus, by the property (*) above, we can find y whose $s(y)$ is equal to the summary of π extended by the transition $(\uparrow, \varepsilon, q')$.

We observe that when the transitions used to reach some node x of t_σ are labeled by $(c_1, G_1), \dots, (c_m, G_m)$, then the sequence of actions in the summary $s(x)$ is also c_1, \dots, c_m ; because σ is n -winning, this sequence has value at most n . Since the summary of each partial play consistent with ρ is of form $s(x)$, this proves the first statement of the lemma. Similarly we prove the last statement: Consider the infinite summary η of some infinite play π consistent with ρ . For each of its finite prefixes η_i we have a node x_i such that $s(x_i) = \eta_i$. These nodes form an infinite, winning path in t_σ , so the maximal priority appearing infinitely often in the actions of η (hence also of π) is even.

It remains to check the second statement. For a node x , we can define the graph $G(x)$ in a similar fashion as the winning condition of $\mathcal{A}_{1.5}$: Let $(\cdot, G_1), \dots, (\cdot, G_{k(x)})$ be the labels of the transitions on the path of t_σ leading to x . Graph $G(x)$ has as nodes $\{0, 1, \dots, k(x)\} \times \text{Const}$, and there is a c -labeled edge to $(i-1, \alpha)$ from (i, β) iff $(\alpha, c, \beta) \in G_{k(x)}$. We will prove the following property: if a partial run $\pi_x \pi_1 \pi_2 \dots \pi_m$ is consistent with ρ , and each π_i is a (c_i, p_i) -return, and $s(x)$ is the summary of π_x , then in $G(x)$ there exists a path starting from $(k(x), (c_1, p_1))$, which is labeled by a sequence obtained from c_2, c_3, \dots, c_m by inserting some ε actions. Notice that this property already gives the second statement: because σ is n -winning, the value of any path in the graph does not exceed n ; inserting additional ε actions does not change this value.

The above property is proved by induction on the level of x . It is trivially true in the root of t_σ , because when we are in the root of t we cannot perform any return. Assume that the property is true for some x , and take its child y . Let $\pi_y \pi_1 \pi_2 \dots \pi_m$ be a partial run consistent with ρ , such that each π_i is a (c_i, p_i) -return, and $s(y)$ is the summary of π_y . Let (d, c, q) be the last element of $s(y)$. We can represent $\pi_y = \pi_x \pi'$, where π_x has summary $s(x)$. One possibility is that $d \in \{\circlearrowleft, \cup\}$. Then the transition between x and y is of the form $(\circlearrowleft, (c, G_c), (q, \cdot))$, and $\pi' \pi_1$ is a (cc_1, p_1) -return. By the induction assumption in $G(x)$ (hence also in $G(y)$) there is a path from $(k(x), (cc_1, p_1))$ which is labeled by a sequence obtained from c_2, c_3, \dots, c_m by inserting some ε actions. We can extend it using the edge $((cc_1, p_1), \varepsilon, (c_1, p_1)) \in G_c$, to obtain such path from $(k(y), (c_1, p_1))$. The other possibility is that $d \in \{\swarrow, \searrow\}$. Let z be the node for which $s(z)$ is equal to the summary of $\pi_y \pi_1$. Notice that y and z are labeled by states (q, D) and (p_1, C_{c_1, p_1}) , respectively, where D and C_{c_1, p_1} are the

sets of constraints guessed by Eve in x . By property (*) we have $(c_1, p_1) \in D$ and $(c_2, p_2) \in C_{c_1, p_1}$. Notice that $c = \varepsilon$ and $\pi' \pi_1 \pi_2$ is a $(c_1 c_2, p_2)$ -return. By the induction assumption in $G(x)$ (hence also in $G(y)$) there is a path from $(k(x), (c_1 c_2, p_2))$ which is labeled by a sequence obtained from c_3, c_4, \dots, c_m by inserting some ε actions. In the graph labeling the transition from x to y we have the edge $((c_1 c_2, p_2), c_2, (c_1, p_1))$, thus there is a path from $(k(y), (c_1, p_1))$ as required.

A.4 Expanding loops

Lemma 11 gives us a strategy for $\mathcal{A}_{1.5}$ for which each play satisfy three conditions. Now we want to show that such play is $\alpha(n)$ -winning for some $\alpha(n)$, as described by the following lemma.

Lemma 12. *Let π be a play in the game (\mathcal{A}_2, t) such that for some $n \in \mathbb{N}$:*

- for each prefix π' of π , the sequence of actions in the summary of π' has value at most n , and
- for each infix of π which is of the form $\pi_1 \pi_2 \dots \pi_m$, where each π_i is a return with global action c_i , the sequence c_1, c_2, \dots, c_m has value at most $n + 1$, and
- the maximal priority appearing infinitely often in the labels of transitions in π is even.

Then π is $\alpha(n)$ -winning, for some correction function α not depending on π .

We can analyze the behavior of each counter independently; moreover, we do not have to care about the parity condition. Thus below, when talking about actions, we think just about actions of a single counter, that is ic or r or ε . We are first interested in values of returns.

Lemma 13. *Let π be a return with global action ic or ε such that for some $n, n' \in \mathbb{N}$:*

- for each proper prefix π' of π , the sequence of actions in the summary of π' has value at most n , and
- for each suffix of π which is of the form $\pi_1 \pi_2 \dots \pi_m$, where each π_i is a return with global action c_i , the sequence c_1, c_2, \dots, c_m has value at most n' .

Then π has value at most $2^{n+n'}$.

Proof. We proceed by induction on the length of π . We have three cases depending on the direction d used in the first transition of π . When $d = \uparrow$, the return is of length 1, so its value is at most 1.

Next, assume that $d = \circlearrowleft$. Let c be the action of the first transition, and let π_C be the suffix of π without the first transition. Notice that π_C is also a return, and that the summary of a prefix of π_C , preceded by the first transition, gives the summary of a prefix of π . If $c = \varepsilon$, by the induction assumption π_C has value at most $2^{n+n'}$, hence π as well. If $c = \text{ic}$, we can use the induction assumption even for $n - 1$ and n' ; it says that π_C has value at most $2^{n-1+n'}$, hence π has value at most $2^{n-1+n'} + 1 \leq 2^{n+n'}$.

The last possibility is that $d = \downarrow$. Then $\pi = \pi_A \pi_B \pi_C$, where π_A performs one transition, and both π_B and π_C are returns. Recall that π_A has action ε . Let c_B and c_C be the global actions of π_B and π_C , respectively. Notice that the summary of a prefix of π_C , preceded by (\cup, c_B, \cdot) , gives the summary of a prefix of π (we have this also for the summary of a prefix of π_B preceded by the first transition). On the other hand, each suffix of π_B which is of the form $\pi_1 \pi_2 \dots \pi_m$, where each π_i is a return, can be extended by the return π_C to a suffix of π . Thus, if $c_B = \varepsilon$, by the induction assumption π_C has value at most $2^{n+n'}$, hence π as well, and if $c_C = \varepsilon$, by the induction assumption π_B has value at most $2^{n+n'}$, hence π as well. If $c_B = c_C = \text{ic}$, we can use the induction

assumption for n and $n' - 1$ for π_B , and for $n - 1$ and n' for π_C ; then the value of π is at most $2^{n+n'-1} + 2^{n-1+n'} = 2^{n+n'}$. \square

To deal with returns having global action \mathbf{r} we need the following definition. For a sequence of actions, its *head value* is the number of increments before the first reset, and its *tail value* is the number of increments after the last reset (in both cases, we count all increments when there are no resets).

Lemma 14. *Let π be a return such that for some $n, n_H, n_T \in \mathbb{N}$, where $n_H, n_T \leq n$:*

- for each proper prefix π' of π , the sequence of actions in the summary of π' has value at most n , and head value at most n_H , and
- for each suffix of π which is of the form $\pi_1\pi_2 \dots \pi_m$, where each π_i is a return with global action c_i , the sequence c_1, c_2, \dots, c_m has value at most n , and tail value at most n_T .

Then π has value at most 2^{2n+2} , head value at most 2^{n+n_H+1} , and tail value at most 2^{n+n_T+1} .

Proof. For returns with global action \mathbf{ic} or ε the previous lemma gives the result. Thus assume that π has global action \mathbf{r} . As previously, we proceed by induction on the length of π . Again we have three cases depending on the direction d used in the first transition of π , and again the case $d = \uparrow$ is trivial.

Assume that $d = \circlearrowleft$. Let c be the action of the first transition, and let π_C be the suffix of π without the first transition. Then, if $c = \varepsilon$, we can use the induction assumption for π_C for n, n_H, n_T ; if $c = \mathbf{ic}$, for $n, n_H - 1, n_T$; and if $c = \mathbf{r}$, for n, n, n_T . For $c = \varepsilon$ the induction assumption gives the result. In all cases the tail value of π and π_C is the same. For $c = \mathbf{ic}$, the head value of π is at most $2^{n+n_H} + 1 \leq 2^{n+n_H+1}$; the value of π is at most $\max(2^{n+n_H} + 1, 2^{2n+2}) = 2^{2n+2}$. For $c = \mathbf{r}$, the head value of π is 0, and the value of π and π_C is the same.

Next, assume that $d = \downarrow$. Then $\pi = \pi_A\pi_B\pi_C$, where π_A performs one transition, and both π_B and π_C are returns. Let c_B and c_C be the global actions of π_B and π_C , respectively. If $c_B = \varepsilon$, we can use the induction assumption for π_C for n, n_H, n_T ; if $c_B = \mathbf{ic}$, for $n, n_H - 1, n_T$; and if $c_B = \mathbf{r}$, for n, n, n_T . Similarly, if $c_C = \varepsilon$, we can use the induction assumption for π_B for n, n_H, n_T ; if $c_C = \mathbf{ic}$, for $n, n_H, n_T - 1$; and if $c_C = \mathbf{r}$, for n, n_H, n . When $c_B = \varepsilon$ or $c_C = \varepsilon$, the induction assumption gives the result. When $c_B = \mathbf{ic}$, the head value of π is at most $2^{n_H+n} + 2^{n+n_H} = 2^{n+n_H+1}$ (we use the previous lemma for π_B with n_H as n and n as n'); the tail value of π and π_C is the same (since $c_C = \mathbf{r}$); the value of π is at most $\max(2^{n+n_H+1}, 2^{2n+2}) = 2^{2n+2}$. The case $c_C = \mathbf{ic}$ is symmetric. Finally, when $c_B = c_C$, the head value of π and π_B is the same; the tail value of π and π_C is the same; the value of π is at most $\max(2^{2n+2}, 2^{n+n+1} + 2^{n+n+1}) = 2^{2n+2}$. \square

Finally, we have a lemma saying what happens when we replace single actions by sequences of actions.

Lemma 15. *Let π_{short} be a sequence of actions on one counter of value at most k . Let π be obtained from π_{short} by expanding each action c into a sequence of global action c and value at most m . Then the value of π is at most $(k + 2)m$.*

Proof. We add special symbols $\$$ in π to mark the frontiers between each of these sequences. Let τ be an infix of π not containing resets, and let n be the number of increments in τ . Before the first $\$$ in τ , and after the last $\$$ in τ (or when there are $\$$ in τ) we have at most m increments. Also between two consecutive $\$$ in τ we have at most m increments. But when there are some increments between two consecutive $\$$, this fragment was obtained by expanding an

increment in π_{short} . Thus we have at most k such fragments in τ . We get that n (hence also the value of π) is at most $(k + 2)m$. \square

Now it is easy to conclude with the proof of Lemma 12. Take any play π as in its assumptions, and concentrate on actions of one counter. Notice that each infix of π which is a return, satisfies the assumptions of Lemma 14 for $n + 1$ taken as n, n_H, n_T . Thus the value of each return is at most 2^{2n+4} . By assumption, the sequence of actions in the summary of π has value at most n . To obtain π from its summary, we have to replace elements of the form (\circlearrowleft, c, q) by corresponding fragments of π . Each such corresponding fragment consists of a transition going down (with action ε), and a return with global action c . Thus, by Lemma 15, the value of π is at most $(n + 2)2^{2n+4}$.

A.5 Simplifying the winning condition

In this section, we want to construct a B-parity automaton $\mathcal{A}'_{1.5}$ based on the automaton $\mathcal{A}_{1.5}$ (which uses a more complex acceptance condition). For this we observe that the acceptance condition can be recognized by a history-deterministic B-automaton. Informally, a non-deterministic automaton is history-deterministic if it is possible to choose deterministically the next transition of the automaton, basing on the already read prefix of the word (but not necessarily only on the current state, as in deterministic automata). It is known that B-automata cannot be determinized, but can be made history-deterministic (see [8] for details).

Consider the cost function over $Act \times Graph$, which maps a word $(c_1, G_1), (c_2, G_2), \dots$ into the smallest n such that c_1, c_2, \dots has value at most n , and each path in the graph defined by G_1, G_2, \dots (as in the winning condition of $\mathcal{A}_{1.5}$) has value at most n . We want to construct a history-deterministic B-automaton computing this function. The first part presents no difficulty: the automaton just deterministically simulates the actions c_1, c_2, \dots ; it has the same set of priorities as $\mathcal{A}_{1.5}$. In the second part, we want to verify a safety condition: the value of any (finite) path does not exceed n . Therefore it is sufficient to design an automaton recognizing this function on finite words from $Graph^*$. It is easy to do with a universal B-automaton: simply guess a backward path in the graph, guessing at each step an edge (notice that it does not matter whether we execute a sequence of actions from the end). This shows that the cost function we want to define is regular. So by [8], there is a history-deterministic B-automaton computing this cost function, up to some correction function. Moreover, this automaton has to accept every prefix of the input word, so it has a safety acceptance condition. By doing the product of the two automata for these two parts, we obtain a history-deterministic B-automaton \mathcal{A}_{acc} computing our function.

Since \mathcal{A}_{acc} is history-deterministic, we can take as $\mathcal{A}'_{1.5}$ the product of $\mathcal{A}_{1.5}$ and \mathcal{A}_{acc} (where \mathcal{A}_{acc} just reads the labels of transitions of $\mathcal{A}_{1.5}$). Due to history-determinism, the run of \mathcal{A}_{acc} on different branches of the game tree of $\mathcal{A}_{1.5}$ are the same on the common prefix of the branches. Thus the product computes the same cost function over infinite trees (up to a correction function) as $\mathcal{A}_{1.5}$.

A.6 Eliminating local moves

Now, starting from the automaton $\mathcal{A}'_{1.5} = \langle Q'', \mathbb{A}, q_0'', \Gamma', \delta \rangle$, which does not use the \uparrow direction, we will construct an equivalent automaton \mathcal{A}_1 which is one-way, i.e. it does not use also the \circlearrowleft direction.

As in the case of \mathcal{A}_2 , we assume that states of $\mathcal{A}'_{1.5}$ are divided into existential, universal, and d -states for $d \in \{\swarrow, \searrow\}$. Let $Goal = \mathcal{P}(\{\swarrow, \searrow\} \times Act \times Q'')$. For each $q_s \in Q''$, each $a \in \mathbb{A}$ and each $G \in Goal$ we define a local game $\mathcal{G}(q_s, a, G)$, which is a B-parity game. Its positions are $Act \times Q'' \cup \{\text{win, lose}\}$. Positions

of (c, q) with q universal belong to Adam, the others to Eve. The initial position is (ε, q_s) . Moves from a position (c, q) are defined as follows:

- When q is existential or universal with $\delta(q, a) = \bigvee_{i \in I} (\circlearrowleft, c_i, q_i)$ or $\delta(q, a) = \bigwedge_{i \in I} (\circlearrowleft, c_i, q_i)$, for each $i \in I$ we have a move to (cc_i, q_i) , labeled by c_i .
- For q is a d -state with $\delta(q, a) = (d, \varepsilon, q')$, we have an ε -labeled move to win when $(d, c, q') \in G$; otherwise an ε -labeled move to lose.

Position win is made winning, and lose losing, by appropriately labeled loop.

In order to make our translation effective, we now observe that it is decidable, whether Eve wins in such a game $\mathcal{G}(q_s, a, G)$. We describe briefly how we can turn the winning condition of $\mathcal{G}(q_s, a, G)$ into a Streett condition (see [26] for more details). Let P_0 be the parity condition from the game $\mathcal{G}(q_s, a, G)$. For every counter $\gamma \in \Gamma$, let P_γ be the following condition: if there are infinitely many increments for γ , then there must be infinitely many resets for γ . Taking the conjunction of P_0 and all the P_γ yields a Streett condition. Let \mathcal{G}_S be the corresponding Streett game, obtained from $\mathcal{G}(q_s, a, G)$, and equipped with this Streett condition. A Streett game can be solved, and if it is winning for Eve, she has a finite-memory strategy. Of course if Eve wins in $\mathcal{G}(q_s, a, G)$, she can also win in \mathcal{G}_S . On the other hand, if some finite-memory strategy is winning in \mathcal{G}_S , it is also winning in $\mathcal{G}(q_s, a, G)$: whenever the position and the content of the memory repeats, then each counter γ either was reset in between, or it was not changed (otherwise, by repeating this fragment we obtain a play consistent with the strategy with infinitely many increments and no resets); thus the number of increments in a row is bounded by the size of the memory times the size of the game.

Now we define the automaton \mathcal{A}_1 . It has the same input alphabet, set of states, set of counters, and initial state as $\mathcal{A}_{1.5}$. For a state q and input letter a , its transition function is defined as follows: Eve chooses a set $G \in \text{Goal}$ such that she wins in $\mathcal{G}(q, a, G)$; then Adam chooses any transition $(d, c, q') \in G$ and performs it.

It remains to prove that \mathcal{A}_1 and $\mathcal{A}_{1.5}$ are equivalent. Consider some input tree t , and assume that Eve has an n -winning strategy σ in $(\mathcal{A}_{1.5}, t)$. Let t_σ be its strategy tree. We will construct an n -winning strategy ρ in (\mathcal{A}_1, t) . While defining it, we keep track of a node of t_σ (labeled by the position at the end of our partial play). To the empty partial play we assign the root of t_σ . We proceed by induction on the length of the partial play. Consider some partial play, for which we want to define the behavior of ρ . Let x be the corresponding node of t_σ ; let (v_x, q_x) be the label of x , and a_x the label of v_x . We look for all descendants y of x such that the path from x to y ends by a \swarrow or \searrow direction, and before uses only \circlearrowleft direction. In ρ Eve takes as G the set of triples (d_y, c_y, q_y) for all such y , where d_y is the direction used by the last transition before y , c_y is the global action of the path from x to y , and q_y is the state labeling y . Notice that the strategy σ , starting from node x , is n -winning in the game $\mathcal{G}(q_x, a_x, G)$. Indeed, the counters never exceed the value n , and we either always use the \circlearrowleft direction and we continue along an infinite branch of t_σ , or we go down, and then we win by definition of G . Thus, choosing such G is a legal move of Eve. When Adam chooses a transition $(d, c, q) \in G$, to such a longer partial run we assign any node y for which $(d_y, c_y, q_y) = (d, c, q)$. We see that the actions on any play consistent with ρ are obtained as contractions of actions on some play consistent with σ , so ρ is n -winning.

For the opposite direction, take an n -winning strategy σ in the game (\mathcal{A}_1, t) . Fix also N such that Eve N -wins in each local game $\mathcal{G}(q, a, G)$ in which she can win for some N' . The strategy ρ in the game $(\mathcal{A}_{1.5}, t)$ is defined as follows. When we enter

some node v of t (and at the very beginning), we choose graph G according to σ . Let a be the label of v , and let q be state at this moment. Then we proceed according to the N -winning strategy in the game $\mathcal{G}(q, a, G)$. When the play in this game enters the win position, due to some transition $(d, c, q') \in G$, we simulate in σ the move of Adam using (d, c, q') , and we continue in the child of v . Consider now some play π consistent with ρ ; we will show that it is $(n+2)N$ -winning. One possibility is that this play visits infinitely many nodes of t . Then we have a corresponding play η of (\mathcal{A}_1, t) consistent with σ . Notice that the sequence of actions in π is obtained from the sequence of actions in η , by expanding each action c into a sequence of global action c which is a partial play in some local game, consistent with the N -winning strategy. Since η is n -winning, and the partial plays have value at most N , by Lemma 15, π is $(n+2)N$ -winning. The other possibility is that after finitely many steps we stay forever in some local game; this situation is analogous.

A.7 Putting the pieces together

By all the previous results, we obtain an B-IPA \mathcal{A}_1 which is equivalent to \mathcal{A}_2 .

Theorem 16. *This construction preserves the quasi-weak property.*

Proof. We want to show that if we start with a quasi-weak automaton \mathcal{A}_2 , this construction yields a quasi-weak automaton \mathcal{A}_1 .

Consider a loop π_1 in \mathcal{A}_1 , which is a path from $q'' \in Q''$ to itself. In particular the Q component of Q'' , keeping track of the state of the original automaton \mathcal{A}_2 , has to be the same. This means that π_1 is a contraction of a loop π_2 in \mathcal{A}_2 , from $q \in Q$ to itself. If both priorities are seen in π_1 , then it is also the case in π_2 . Therefore, there must be a counter γ which is incremented but not reset in π_2 . This property is preserved by contraction of paths, so γ is incremented but not reset in π_1 . We obtain that \mathcal{A}_1 is quasi-weak. \square

B. From QWCMSO to B-quasi-weak automata

We now prove the remaining properties and translations which were used in Section 4 to convert QWCMSO formulae to 1-way B-quasi-weak automata.

B.1 Closure under weak inf-projection and weak sup-projection

We first show that B-quasi-weak automata are closed under weak inf- and weak sup-projection, the operations corresponding to weak existential quantification and weak universal quantification in QWCMSO.

Let \mathcal{A} be a 1-way B-quasi-weak automaton over \mathbb{A}' . Fix $h : \mathbb{A}' \rightarrow \mathbb{A}$ such that $\mathbb{A}' \supseteq \mathbb{A}$ and $h(a) = a$ for $a \in \mathbb{A}$. For a tree t' , we write $h(t')$ for the relabeling of t' according to h , and $h_{\text{fin}}(t') = t$ if only finitely many positions of t' have labels in $\mathbb{A}' \setminus \mathbb{A}$. For all $g : \mathbb{A}' \rightarrow \mathbb{N}_\infty$ and $\text{op} \in \{\text{inf}, \text{sup}\}$, we define $g_{\text{op}, h_{\text{fin}}} : \mathbb{A} \rightarrow \mathbb{N}_\infty$ such that $g_{\text{op}, h_{\text{fin}}}(t) := \text{op} \{g(t') : h_{\text{fin}}(t') = t\}$. This is called the *weak op-projection of f by h* .

We seek to construct $\mathcal{A}_{\text{inf}, h_{\text{fin}}}$ and $\mathcal{A}_{\text{sup}, h_{\text{fin}}}$ such that $[\mathcal{A}_{\text{op}, h_{\text{fin}}}] \approx [\mathcal{A}]_{\text{op}, h_{\text{fin}}}$. Because the relabeling only happens on a finite prefix of the tree, the idea is to take an automaton which performs this op-projection on finite trees, and run it on larger and larger prefixes of the infinite tree, while also checking that there is an infinite continuation of the original automaton \mathcal{A} outside of the prefix.

We start with a lemma which allows us to take a regular cost function over finite trees and construct cost automata over infinite trees that compute the minimum or maximum value of this function over all prefixes. This relies on non-trivial results about regular cost functions over finite trees from [13].

Lemma 17. *For every regular cost function f over finite trees, there exists a nondeterministic B-reachability automaton \mathcal{F}_{inf} and an alternating B-safety automaton \mathcal{F}_{sup} such that $\llbracket \mathcal{F}_{\text{op}} \rrbracket(t) \approx_{\alpha_{\text{nd}}} \text{op} \{f(t|_P) : P \text{ is a prefix of } t\}$ over infinite trees t .*

Proof. Let \mathcal{F} be the nondeterministic B-automaton recognizing f over finite trees [13]. The nondeterministic B-reachability automaton recognizing the first function allows Eve to select a finite prefix P of t and simulate \mathcal{F} on this prefix. During the simulation, the automaton outputs only priority 1. At a position on the frontier P , if the simulation ends in an accepting (respectively, rejecting) state then the output is priority 2 (respectively, priority 1) for the remainder of the run. We call this a *B-reachability automaton* since any play in a corresponding cost game starts in priority 1, and is trying to reach a state where it can change to priority 2, and then output priority 2 forever.

By [13, Theorem 12], every regular cost function over finite trees can be recognized by both a nondeterministic B-automaton and a nondeterministic S-automaton. Nondeterministic S-automata are a dual form of cost automata that, roughly speaking, compute the maximum value over all runs of the automaton (see [13] for more information). Thus, for \mathcal{F}_{sup} it makes sense to start from the nondeterministic S-automaton for f and, using similar reasoning as above, get a nondeterministic S-reachability automaton recognizing the second function. Using an adaptation of a standard dualization procedure which switches conjunctions and disjunctions and dualizes the priority function, this nondeterministic S-reachability automaton can be converted into an alternating B-safety automaton (again, by [13, Theorem 12]). We call this a *B-safety automaton* since any play in a corresponding cost game starts in priority 2, and is trying to avoid reaching a state where it would change to priority 1, and output priority 1 forever. \square

We will use this lemma to help construct the desired automata.

Let \mathcal{A}_{ext} be an automaton on finite trees over the alphabet $\mathbb{A}' \times \mathcal{P}(Q)$ which simulates \mathcal{A} and checks that the state of \mathcal{A} is in the label at that position.

Let $f_{\text{op}} = \llbracket \mathcal{A}_{\text{ext}} \rrbracket_{\text{op},h}$ (over finite trees) obtained using [13]. Then apply Lemma 17 to get automata \mathcal{F}_{op} on infinite trees over the alphabet $\mathbb{A}' \times \mathcal{P}(Q)$.

Finally, let $\mathcal{A}_{\text{op},h_{\text{fin}}}$ be the B-Büchi automaton over the alphabet \mathbb{A} such that on input t , Adam and Eve start by playing the game induced by \mathcal{F}_{op} , with Eve additionally responsible for choosing annotations from $\mathcal{P}(Q)$. At any position corresponding to some $x \in \mathcal{T}$, Adam can challenge some $q \in Q$ in the annotation selected by Eve, and then the players switch to the game (\mathcal{A}_q, t_x) , where \mathcal{A}_q is \mathcal{A} starting from state q and t_x is the subtree of t starting at x . The output comes directly from the output from \mathcal{F}_{op} and \mathcal{A}_q (we assume that \mathcal{F}_{op} and \mathcal{A}_q do not share counters).

Notice that $\mathcal{A}_{\text{op},h_{\text{fin}}}$ is still quasi-weak, since running \mathcal{F}_{op} before launching copies of \mathcal{A} does not introduce any cycles with both priorities. It remains to show that it recognizes the weak op-projection as desired.

Theorem 18. $\llbracket \mathcal{A}_{\text{op},h_{\text{fin}}} \rrbracket \approx \llbracket \mathcal{A} \rrbracket_{\text{op},h_{\text{fin}}}$

Proof. We give the proof for $\text{op} = \text{sup}$ (the harder case, since \mathcal{F}_{sup} is alternating rather than nondeterministic. Let $\mathcal{B} := \mathcal{A}_{\text{sup},h_{\text{fin}}}$. Let $\alpha = \alpha_{\text{nd}}$ from Lemma 17.

$\llbracket \mathcal{B} \rrbracket \preceq_{\alpha} \llbracket \mathcal{A} \rrbracket_{\text{sup},h_{\text{fin}}}$. Assume $\llbracket \mathcal{A} \rrbracket_{\text{sup},h_{\text{fin}}}(t) \leq n < \infty$. Then for all prefixes P , and for all relabellings on this prefix resulting in t' with $h_{\text{fin}}(t') = t$, $\llbracket \mathcal{A} \rrbracket(t') \leq n$.

For each $x \in \mathcal{T}$, let $Q_x^t := \{q \in Q : \llbracket \mathcal{A}_q \rrbracket(t_x) \leq n\}$. Let $\sigma_{q,x}$ denote the strategy for Eve which witnesses $\llbracket \mathcal{A}_q \rrbracket(t_x) \leq n$ for $q \in Q_x$.

Fix some prefix P . For any relabeling on this prefix resulting in t' such that $h_{\text{fin}}(t') = t$, there is a strategy σ for Eve in (\mathcal{A}, t') such that $\text{val}(\sigma) \leq n$, and any play which ends in a position x on the frontier of P must be in a state $q \in Q_x$ (otherwise, $\text{val}(\sigma) > n$).

Let t_{ext} be the tree over $\mathbb{A}' \times \mathcal{P}(Q)$ such that $t_{\text{ext}}(x) = (t'(x), Q_x)$. This means that $\llbracket \mathcal{A}_{\text{ext}} \rrbracket(t_{\text{ext}}|_P) \leq n$ (this is witnessed by $\sigma|_P$). Since this is true for all such t' , this means that $\llbracket \mathcal{A}_{\text{ext}} \rrbracket_{\text{sup},h}(t_{\text{ext}}|_P) \leq n$.

Since this is true for all prefixes P , $\llbracket \mathcal{F}_{\text{sup}} \rrbracket(t_{\text{ext}}) \leq \alpha(n)$ by Lemma 17, witnessed by some strategy $\sigma_{\mathcal{F}}$.

We use this to construct a strategy for Eve in (\mathcal{B}, t) witnessing value at most $\alpha(n)$: initially, Eve selects the labels from $\mathcal{P}(Q)$ at position x according to Q_x , and uses her strategy from \mathcal{F}_{sup} ; if Adam challenges some q at position x , then Eve switches to strategy $\sigma_{q,x}$. Overall, the value is at most $\alpha(n)$.

$\llbracket \mathcal{A} \rrbracket_{\text{sup},h_{\text{fin}}} \preceq_{\alpha} \llbracket \mathcal{B} \rrbracket$. Assume $\llbracket \mathcal{B} \rrbracket(t) \leq n < \infty$. Then there is a strategy $\sigma_{\mathcal{B}}$ for Eve in (\mathcal{B}, t) such that $\text{val}(\sigma_{\mathcal{B}}) \leq n$. Let Q_x be the set of all states $q \in Q$ such that there is $\pi \in \sigma_{\mathcal{B}}$ which reaches a position corresponding to $x \in \mathcal{T}$ with q in the label selected by Eve. This implies that $\llbracket \mathcal{A}_q \rrbracket(t_x) \leq n$ for all $q \in Q_x$, witnessed by some strategy $\sigma_{q,x}$ in (\mathcal{A}_q, t_x) .

Because \mathcal{F}_{sup} is alternating, Eve may not make the choices of a label from $\mathcal{P}(Q)$ consistently in $\sigma_{\mathcal{B}}$ (i.e., Eve may select different labels for two plays which reach position x). However, it is straightforward to construct a new strategy $\sigma'_{\mathcal{B}}$ for Eve in (\mathcal{B}, t) with $\text{val}(\sigma'_{\mathcal{B}}) \leq n$ where Eve plays as before, except now she chooses the same label Q_x at all game positions corresponding to position $x \in \mathcal{T}$. This label Q_x is just the union of all of the labels from $\mathcal{P}(Q)$ that Eve selected in some partial play in $\sigma_{\mathcal{B}}$ ending in x . This means that if Adam challenges some $q \in Q_x$, then $\sigma'_{\mathcal{B}}$ can duplicate Eve's substrategy starting from any play in $\sigma_{\mathcal{B}}$ where she included q in her label at x .

Let $t_{\text{ext}}(x) = (t(x), Q_x)$. Since $\text{val}(\sigma'_{\mathcal{B}}) \leq n$, $\llbracket \mathcal{F}_{\text{sup}} \rrbracket(t_{\text{ext}}) \leq n$ (witnessed by the restriction of $\sigma'_{\mathcal{B}}$ to the \mathcal{F}_{sup} part of the plays). By Lemma 17, this means that for all prefixes P , $\llbracket \mathcal{A}_{\text{ext}} \rrbracket_{\text{sup},h}(t_{\text{ext}}|_P) \leq \alpha(n)$, witnessed by strategies σ_P .

Consider any tree t' with $h_{\text{fin}}(t') = t$. Let P be the minimal prefix such that t' is identical to t outside of P . We construct a strategy σ witnessing $\llbracket \mathcal{A} \rrbracket(t') \leq \alpha(n)$. Within the prefix P , the strategy σ plays annotations from $\mathcal{P}(Q)$ according to Q_x , and moves in \mathcal{A} according to σ_P . Starting from the frontier of P , the strategy uses $\sigma_{q,x}$ for the infinite continuation. Overall, the value of this strategy is at most $\alpha(n)$ from the part using σ_P , and at most n from the part using $\sigma_{q,x}$. Hence, overall, the value is bounded by $\alpha(n)$. Because this is true for all t' with $h_{\text{fin}}(t') = t$, $\llbracket \mathcal{A} \rrbracket_{\text{sup},h_{\text{fin}}}(t) \leq \alpha(n)$ as desired. \square

B.2 The bounded expansion operator on automata

We now prove Theorem 4.

We start with a 2-way localized B-quasi-weak automaton \mathcal{A} that takes input (t, X) where t is a tree over the alphabet \mathbb{A} , X is a distinguished set of positions.

We assume that $\llbracket \mathcal{A} \rrbracket$ is monotonic in X in the sense that for all $q \in Q$ and $n \in \mathbb{N}$, if $X \subseteq X'$ then $\{x : \llbracket \mathcal{A}_q \rrbracket_x(t, X) \leq n\} \subseteq \{x : \llbracket \mathcal{A}_q \rrbracket_x(t, X') \leq n\}$ where \mathcal{A}_q denotes \mathcal{A} starting from state q .

To help with the proof, we define the *m-bounded expansions* of \mathcal{A} on $t \in \mathcal{T}_{\mathbb{A}}$ such that $X_{\mathcal{A},t,m}^0 = \emptyset$ and $X_{\mathcal{A},t,m}^i = \{x : \llbracket \mathcal{A} \rrbracket_x(t, X_{\mathcal{A},t,m}^{i-1}) \leq m\}$. Note that monotonicity of $\llbracket \mathcal{A} \rrbracket$ implies that $X_{\mathcal{A},t,m}^i \subseteq X_{\mathcal{A},t,m}^j$ for $i \leq j$.

Rewritten in terms of these bounded expansions, Theorem 4 says that the localized 2-way automaton \mathcal{B} over the alphabet \mathbb{A} described in Section 4.2 satisfies $\llbracket \mathcal{B} \rrbracket_z(t) \approx^{t,z} \inf \{m : z \in X_{\mathcal{A},t,m}^m\}$.

We recall the construction of \mathcal{B} . The set of states is the same as \mathcal{A} , and the initial state of \mathcal{B} is the initial state of \mathcal{A} . In addition to the counters from \mathcal{A} we have one additional counter γ (which will only be incremented and checked, and never reset). We say that we start a *fresh copy* of \mathcal{A} from some position x , if we

- move to the initial state of \mathcal{A} ,
- reset all counters in \mathcal{A} , and
- perform `ic` on γ .

Given some input t and some starting position x , \mathcal{B} starts a fresh copy of \mathcal{A} at x , with Eve selecting labels over the extended alphabet which are consistent with t : if $t(y) = a$ then Eve can choose a label of $(a, 0)$ or $(a, 1)$. If Eve chooses $(a, 1)$ at y , then Adam is allowed to challenge this choice by starting a fresh copy of \mathcal{A} from y . The output is taken from the simulations of \mathcal{A} , along with the new counter actions from γ .

Lemma 19. *For all positions x and for all $m \in \mathbb{N}$, $\llbracket \mathcal{A} \rrbracket_x(t, X_{\mathcal{A}, t, m}^{n-1}) \leq m$ iff $\llbracket \mathcal{B} \rrbracket_x(t)$ has value at most n in counter γ and value at most m in other counters.*

Proof. Fix some position x and $m \in \mathbb{N}$. We proceed by induction on n . We write X_m^{n-1} for $X_{\mathcal{A}, t, m}^{n-1}$. The base case for $n = 1$ is straightforward (the base case is $n = 1$ rather than $n = 0$, since the automaton starts with a fresh copy of \mathcal{A} which immediately increments γ).

Let $n > 1$ and assume that $\llbracket \mathcal{A} \rrbracket_x(t, X_m^{n-1}) \leq m$. Since $y \in X_m^{n-1}$ implies $\llbracket \mathcal{A} \rrbracket_y(t, X_m^{n-2}) \leq m$, the inductive hypothesis implies that $\llbracket \mathcal{B} \rrbracket_y(t)$ has value at most $n - 1$ from counter γ and value at most m from the other counters.

Consider the strategy in $(\mathcal{B}, t)_x$ where Eve guesses the labeling X_m^{n-1} and then plays according to her strategy in $(\mathcal{A}, t, X_m^{n-1})_x$, unless Adam challenges at some position $y \in X_m^{n-1}$ and then Eve uses her strategy from $(\mathcal{B}, t)_y$. It is not hard to check that this is indeed a strategy in $(\mathcal{B}, t)_x$. Before a challenge from Adam, the value from counter γ is at most 1 and the value from the other counters is at most m . Overall, this means that the value of counter γ is at most n and the value from the other counters is at most m (recall that the other counters are reset when Adam challenges).

Next, assume that $\llbracket \mathcal{B} \rrbracket_x(t)$ has value at most n from counter γ and value at most m from other counters. Let σ be the strategy witnessing these values, and let σ_1 be the restriction of σ to plays where γ has value at most 1.

Recall that Eve is allowed to select labels over the extended alphabet that are consistent with t . Let Y be the set of positions y such that there is a play π in σ_1 where Eve guesses a label for y in $\mathbb{A} \times \{1\}$. Because \mathcal{A} is alternating, Eve may not guess consistently that the label for y is in $\mathbb{A} \times \{1\}$. However, by monotonicity of \mathcal{A} , marking extra positions with 1 cannot increase the value. Hence, it must be the case that $\llbracket \mathcal{A} \rrbracket_x(t, Y) \leq m$.

The strategy σ also induces strategies $(\mathcal{B}, t)_y$ for $y \in Y$ such that the value of counter γ is at most $n - 1$ and the value from the other counters is at most m . Hence, by the inductive hypothesis $\llbracket \mathcal{A} \rrbracket_y(t, X_m^{n-2}) \leq m$ for all $y \in Y$. This means that $Y \subseteq X_m^{n-1}$, so by monotonicity, $\llbracket \mathcal{A} \rrbracket_x(t, X_m^{n-1}) \leq m$ as desired. \square

In particular, $\llbracket \mathcal{A} \rrbracket_x(t, X_{\mathcal{A}, t, m}^{m-1}) \leq m$ iff $\llbracket \mathcal{B} \rrbracket_x(t) \leq m$, which is enough to conclude that $\llbracket \mathcal{B} \rrbracket_x(t) \approx^{t, x} \inf \{m : x \in X_{\mathcal{A}, t, m}^{m-1}\}$ as desired. We remark that the correction function at this step is just the identity function (it computes exactly the same function).

B.3 From 1-way to localized 2-way

Next, we prove Theorem 5.

Let \mathcal{A} be a 1-way B-quasi-weak automaton over some alphabet $\mathbb{A}' \times \{0, 1\}$, where there is a single distinguished position x marked

using the extended alphabet. As usual, \mathcal{A} starts its computation from the root. We want a variant of this automaton over the alphabet \mathbb{A}' (i.e., on a tree with no position marked) that when started from position x recognizes the same cost function.

We do this by constructing a 2-way B-quasi-weak automaton \mathcal{A}^ℓ over the alphabet \mathbb{A}' such that for all positions x , the value of \mathcal{A}^ℓ on $t \in \mathcal{T}_{\mathbb{A}'}$ when starting at position x is equivalent to the value of \mathcal{A} on the tree t marked at position x and starting at the root. We will write this as $\llbracket \mathcal{A}^\ell \rrbracket_x(t) \approx^{t, x} \llbracket \mathcal{A} \rrbracket(t, x)$ where $\llbracket \mathcal{A}^\ell \rrbracket_x$ denotes the value of \mathcal{A}^ℓ when the computation starts at position x . We write (\mathcal{A}, t, x) for the game where \mathcal{A} acts on the tree t marked with position x but starting at the root, and $(\mathcal{A}^\ell, t)_x$ for the game where \mathcal{A}^ℓ acts on t starting at position x .

Let \mathcal{A}_{ext} be the automaton over finite trees over the alphabet $\mathbb{A}' \times \{0, 1\} \times \mathcal{P}(Q)$ that simulates \mathcal{A} , but also ensures that the state of \mathcal{A} at position y is contained in the third component of the label at y . There is a nondeterministic version $\mathcal{A}_{\text{ext}, \text{nd}}$ on finite trees which is equivalent up to some α_{nd} by [13].

We now describe the operation of \mathcal{A}^ℓ on input t starting at position x . Let τ be the path from the root to x . The idea is that Eve can move upwards from x to the root by guessing labels from $\mathcal{P}(Q)$ and simulating an accepting run of $\mathcal{A}_{\text{ext}, \text{nd}}$ in reverse on τ (a special case of a finite tree). At any point, Adam is allowed to launch a copy of the original automaton \mathcal{A} downwards, starting in a state from the label in $\mathcal{P}(Q)$ selected by Eve, on any path different from τ .

More formally, consider the game at a position corresponding to $y \in \mathcal{T}$ on the path upwards from x to the root (in particular, when the game starts at position x). For $y \neq x$, let yd be the successor of y which is on this path. At such a position y , Eve selects a state q , a transition $(q, (t(y), 0), (c_0, q_0), (c_1, q_1)) \in \Delta_{\text{ext}, \text{nd}}$ of $\mathcal{A}_{\text{ext}, \text{nd}}$, and a label $P \in \mathcal{P}(Q)$ such that the following conditions hold:

- if $y = \epsilon$, then q is initial for $\mathcal{A}_{\text{ext}, \text{nd}}$;
- if $y = x$, then q is accepting for $\mathcal{A}_{\text{ext}, \text{nd}}$;
- if $y \neq x$, then Eve selected q_d in her previous move.

Adam then selects the direction. If he moves upwards, the output is c_d and the play continues as above from the parent of y . Otherwise, Adam selects some $q' \in P$, and Eve and Adam switch to playing the game (\mathcal{A}, t, x) starting from position (q', y) (for the first move after switching to this game, there is an additional restriction that Adam must choose a direction d such that yd is not on τ). The output comes from the (partial) run of $\mathcal{A}_{\text{ext}, \text{nd}}$ on τ and/or a play from \mathcal{A} . We assume counters are not shared between $\mathcal{A}_{\text{ext}, \text{nd}}$ and \mathcal{A} .

Lemma 20. $\llbracket \mathcal{A} \rrbracket(t, x) \approx^{t, x} \llbracket \mathcal{A}^\ell \rrbracket_x(t)$.

Proof. We must show that this construction preserves the value.

$\llbracket \mathcal{A}^\ell \rrbracket_x(t) \preceq_{\alpha_{\text{nd}}} \llbracket \mathcal{A} \rrbracket(t, x)$ for α_{nd} coming from $\mathcal{A}_{\text{ext}, \text{nd}}$ described above. Assume $\llbracket \mathcal{A} \rrbracket(t, x) \leq n < \infty$, witnessed by strategy σ . Let $Q_y = \{q \in Q : \llbracket \mathcal{A}_q \rrbracket(t_y) \leq n\}$. Let t_{ext} be the tree such that $t_{\text{ext}}(y) = (t(y), b, Q_y)$ where $b = 1$ if $y = x$ and 0 otherwise. Consider the finite branch τ from the root to x . Then $\llbracket \mathcal{A}_{\text{ext}, \text{nd}} \rrbracket(t_{\text{ext}}|_\tau) \leq \alpha_{\text{nd}}(n)$. We can use this to construct a strategy in $(\mathcal{A}^\ell, t)_x$ in which Eve plays labels Q_y and the accepting run of $\mathcal{A}_{\text{ext}, \text{nd}}$ on $t_{\text{ext}}|_\tau$ while on τ , and switches to strategy σ outside of τ . Overall, the value would be at most $\alpha_{\text{nd}}(n)$ as desired.

$\llbracket \mathcal{A} \rrbracket(t, x) \preceq_\alpha \llbracket \mathcal{A}^\ell \rrbracket_x(t)$ for $\alpha(n) = \alpha_{\text{nd}}(n) + n$. Assume $\llbracket \mathcal{A}^\ell \rrbracket_x(t) \leq n < \infty$, witnessed by strategy σ^ℓ .

Consider the play which stays on path τ from x to the root. This play induces labels Q_y for each y on τ , and an accepting run of $\mathcal{A}_{\text{ext}, \text{nd}}$ on the finite tree τ_{ext} with $\tau_{\text{ext}}(y) = (t(y), b, Q_y)$ such that b is 1 if $y = x$ and 0 otherwise (and τ_{ext} is undefined outside

of τ). This play has value at most n , so this means there is a strategy σ_{ext} in $(\mathcal{A}_{\text{ext}}, \mathcal{T}_{\text{ext}})$ of value at most $\alpha_{\text{nd}}(n)$.

The strategy σ^ℓ also induces strategies $\sigma_{q,y}$ in the game (\mathcal{A}, t) starting from (q, y) for all y on τ and all $q \in Q_y$.

By combining the strategy σ_{ext} (ignoring the labels from $\mathcal{P}(Q)$) with the strategies $\sigma_{q,y}$, we can construct a strategy in (\mathcal{A}, t) starting from (q_0, ϵ) (i.e., starting from the root). Overall, this strategy must have value at most $\alpha_{\text{nd}}(n) + n$, so $\llbracket \mathcal{A} \rrbracket(t) \leq \alpha(n) + n$. \square

C. From B-quasi-weak automata to QWCMSO

We prove the remaining part of Theorem 2, writing a QWCMSO formula that describes the operation of a 1-way B-quasi-weak automaton.

For the restriction to B-weak automata, we can write a formula in WCMO.

Theorem 21 ([27]). *For every B-weak automaton \mathcal{A} over the alphabet \mathbb{A} , there exists effectively a formula $\psi(x)$ in WCMO and a correction function α_w such that for all $t \in \mathcal{T}_{\mathbb{A}}$ and all positions $v \in \mathcal{T}$, $\llbracket \mathcal{A} \rrbracket(t_v) \approx_{\alpha_w} \llbracket \psi(x) \rrbracket(t, v)$ where t_v denotes the subtree of t rooted at position v .*

In general, the formula is in QWCMSO.

Theorem 22. *For every B-quasi-weak automaton \mathcal{A} over alphabet \mathbb{A} , there exists effectively a formula $\psi(x)$ in QWCMSO and a correction function α such that for all $t \in \mathcal{T}_{\mathbb{A}}$ and all positions $v \in \mathcal{T}$, $\llbracket \mathcal{A} \rrbracket(t_v) \approx_{\alpha} \llbracket \psi(x) \rrbracket(t, v)$ where t_v denotes the subtree of t rooted at position v .*

Proof. We assume that \mathcal{A} has a counter γ_{alt} dedicated to counting alternations in priorities (in a quasi-weak automaton, a counter like this can always be added without changing the cost function computed). We also assume that for each state $q \in Q$, there exists some action c such that for all $a \in \mathbb{A}$, in $\delta(q, a)$ has action c . This means we can think of priorities and counter actions labeling states, rather than edges, in the automaton.

We proceed by induction on the number j of states that have some target according to δ that is not a sink state (q is a sink state if there exists c such that for all $a \in \mathbb{A}$, $\delta(q, a) = (0, c, q) \wedge (1, c, q)$).

We claim that $\llbracket \mathcal{A} \rrbracket(t_v) \approx_{\alpha_j} \llbracket \psi(x) \rrbracket(t, v)$ for $\alpha_0(n) = \alpha_w(n)$ and $\alpha_j(n) = (n+1)(\alpha_w(n) + \alpha_{j-1}(n)) + 1$.

Assume $j = 0$. Then from the initial state q_0 , \mathcal{A} must send copies only to sink states, so \mathcal{A} is weak (there is at most one alternation in priority). Then we can apply Theorem 21 to write a formula in WCMO as desired.

Now assume $j > 0$. Let q_j be the initial state. If q_j has only sink state targets then we proceed as in the previous case. Otherwise, let q_1, \dots, q_j be the states that do not have sink state targets.

The goal is to decompose the operation of \mathcal{A} into an initial weak part \mathcal{A}_j that is allowed to use any state but stops when it would alternate priorities, and quasi-weak parts $\mathcal{A}_1, \dots, \mathcal{A}_{j-1}$ where it can use any state but stops if it would enter q_j . We can use Theorem 21 to write a formula for the first part, and the inductive hypothesis to write a formula for the second part. We then write a formula combining these, looping back to the weak part if the automaton wants to enter q_j in the quasi-weak part. We use the bounded expansion operator to handle this looping back because γ_{alt} is incremented in each loop.

Let \mathcal{A}_j denote the automaton that takes input $(t, x_j, X_j, X_{j-1}, \dots, X_1)$ where $t \in \mathcal{T}_{\mathbb{A}}$ (so this is an automaton over the alphabet $\mathbb{A} \times \{0, 1\}^{j+1}$), and simulates \mathcal{A} starting from the position x_j . Moreover, if counter γ_{alt} is incremented in state q_i in position y , then the transition function sends the automaton from q_i to an accepting sink state (respectively, rejecting sink state) if $y \in X_i$ (respec-

tively, $y \notin X_i$). Notice that \mathcal{A}_j is weak, so we can get a WCMO formula $\psi_j(x, x_j, X_j, \dots, X_1)$ for it using Theorem 21.

Likewise, for $1 \leq i \leq j-1$, let \mathcal{A}_i denote the automaton that takes input (t, z, X_j) . This automaton simulates \mathcal{A} starting from state q_i and the position z . Moreover, if the automaton is in state q_j in position y , then the transition function sends the automaton from q_j to an accepting sink state (respectively, rejecting sink state) if $y \in X_j$ (respectively, $y \notin X_j$). Because the targets of q_j are now sink states, the inductive hypothesis can be applied to \mathcal{A}_i to get a formula $\psi_i(z, X_j)$ for it.

Let $\psi'(x_j, X_j) := \psi_j(x_j, X_j, X_{j-1}, \dots, X_1)[\psi_i(y, X_j)/y \in X_i : 1 \leq i \leq j-1]$, so ψ' is the result of substituting $\psi_i(y, X_j)$ for any statement of the form $y \in X_i$ in ψ_j . The desired formula is

$$\psi(x) := x \in \mu^N X_j. \{x_j : \psi'(x_j, X_j)\}.$$

To help with the proof, we define the *m-bounded expansions* of ψ' on $t \in \mathcal{T}_{\mathbb{A}}$ such that $X_{\psi', t, m}^0 = \emptyset$ and $X_{\psi', t, m}^i = \{x : \llbracket \psi' \rrbracket(t, X_{\psi', t, m}^{i-1}, x) \leq m\}$. Note that monotonicity implies that $X_{\psi', t, m}^i \subseteq X_{\psi', t, m}^{i+1}$ for $i \leq j$.

Let $\alpha_j^m(n) = (m+1)(\alpha_w(n) + \alpha_{j-1}(n)) + 1$, and we proceed with the proof.

$\llbracket \psi(x) \rrbracket(t, v) \preccurlyeq_{\alpha_j} \llbracket \mathcal{A} \rrbracket(t_v)$. We prove a slightly stronger result: if $\llbracket \mathcal{A} \rrbracket(t_v) \leq n < \infty$ and has γ_{alt} values of at most $m \leq n$ witnessed by some strategy σ , then $v \in X_{\psi', t, \alpha_j^m(n)}^{m+1}$. The result follows since $\llbracket \mathcal{A} \rrbracket(t_v) \leq n < \infty$ with γ_{alt} at most n implies that $v \in X_{\psi', t, \alpha_j^m(n)}^{m+1} \subseteq X_{\psi', t, \alpha_j^n(n)}^m$, so $\llbracket \psi(x) \rrbracket(t, v) \leq \alpha_j^n(n) = \alpha_j(n)$.

If $m = 0$, then γ_{alt} is never incremented, so σ can be viewed as a strategy in \mathcal{A}_j of value at most n on input (t, v, X_j, \dots, X_1) for any choice of X_i , and in particular for $X_i = \emptyset$ for all $0 \leq i \leq j$. By Theorem 21, this means that $\llbracket \psi_j(x_j, X_j, \dots, X_1) \rrbracket(t, v, \emptyset, \dots, \emptyset) \leq \alpha_w(n) \leq \alpha_j^n(n)$. Since ψ_j is monotonic in X_1, \dots, X_j , this means that $v \in X_{\psi', t, \alpha_j^n(n)}^1$.

Now let $m > 0$. For each $\pi \in \sigma$, let π' denote the restriction of π to a prefix $\tau_0 \tau$ of π such that γ_{alt} is 0 on τ_0 , γ_{alt} is strictly greater than 0 on τ , and if q_j appears, then it only occurs as the final move in τ . In other words, τ_0 represents part of a play in \mathcal{A}_j and τ represents the continuation of this play in some \mathcal{A}_i for $1 \leq i \leq j-1$. Note that τ may be infinite if the play never returns to state q_j after τ_0 . Let σ' denote the strategy obtained by restricting each play $\pi \in \sigma$ as described above.

Let X'_j be the set of positions u such that there is some $\pi' = \tau_0 \tau \in \sigma'$ such that τ_0 ends in position u in state q_i . Let X'_j denote the set of positions u such that there is some $\pi' = \tau_0 \tau \in \sigma'$ such that τ ends in state q_j in position u .

The strategy σ' witnesses $\llbracket \mathcal{A}_j \rrbracket(t, v, X'_j, \dots, X'_1) \leq n$ and $\llbracket \mathcal{A}_i \rrbracket(t, u, X'_j) \leq n$ for all $u \in X'_j$, so we have $\llbracket \psi_j(x_j, X_j, \dots, X_1) \rrbracket(t, v, X'_j, \dots, \alpha_w(n)) \leq \alpha_j^n(n)$ and $\llbracket \psi_i \rrbracket(t, u, X'_j) \leq \alpha_{j-1}(n)$.

Moreover, for all $u \in X'_j$, it must be the case that $\llbracket \mathcal{A} \rrbracket(t_u) \leq n$ and has value at most $m-1$ for γ_{alt} (otherwise, it would contradict the initial assumption that $\llbracket \mathcal{A} \rrbracket(t_w) \leq n$ and has value at most m for γ_{alt}). By the inductive hypothesis, this means that $X'_j \subseteq X_{\psi', t, \alpha_j^m(n)}^m$.

Putting this together, $\llbracket \psi'(x_j, X_j) \rrbracket(t, v, X'_j) \leq \alpha_w(n) + \alpha_{j-1}(n) \leq \alpha_j^n(n)$, so by monotonicity, $\llbracket \psi'(x_j, X_j) \rrbracket(t, v, X_{\psi', t, \alpha_j^n(n)}^m) \leq \alpha_j^n(n)$, so $v \in X_{\psi', t, \alpha_j^n(n)}^{m+1}$ as desired.

$\llbracket \mathcal{A} \rrbracket(t_v) \preccurlyeq_{\alpha_j} \llbracket \psi(x) \rrbracket(t, v)$. We prove a slightly stronger result: if $v \in X_{\psi', t, n}^m$ for $n < \infty$ and $m \leq n$, then $\llbracket \mathcal{A} \rrbracket(t_v) \leq \alpha_j^m(n)$. The result follows since $\llbracket \psi(x) \rrbracket(t, v) \leq n < \infty$ implies that $v \in X_{\psi', t, n}^m$ and thus $\llbracket \mathcal{A} \rrbracket(t_v) \leq \alpha_j^n(n) = \alpha_j(n)$.

We proceed by induction on m .

Assume $v \in X_{\psi', t, n}^1$. Then it must be the case that $\llbracket \psi' \rrbracket(t, v, \emptyset) \leq n$. But this means there must be some X'_1, \dots, X'_{j-1} such that $\llbracket \psi_j(t, v, \emptyset, X'_{j-1}, \dots, X'_1) \rrbracket \leq n$ and $\llbracket \psi_i(y, X_j) \rrbracket(u, \emptyset) \leq n$ for all $u \in X'_i$. Hence, $\llbracket \mathcal{A}_j \rrbracket(t, v, \emptyset, X'_{j-1}, \dots, X'_1) \leq \alpha_w(n)$ and $\llbracket \mathcal{A}_i \rrbracket(t, u, \emptyset) \leq \alpha_{j-1}(n)$ for all $u \in X'_i$. Because $X_j = \emptyset$, we know that \mathcal{A}_i never returns to state q_j (otherwise, the value would be ∞). This means that any suffix of a finite play from $(\mathcal{A}_j, t, v, \emptyset, X'_{j-1}, \dots, X'_1)$ starting in position v and ending in position u in state q_i , can be extended with a play in (\mathcal{A}_i, t, u) to yield a play in (\mathcal{A}, t_v) of value at most $\alpha_w(n) + \alpha_{j-1}(n) \leq \alpha_j^1(n)$. Hence, there is a strategy for \mathcal{A} of value at most $\alpha_j^1(n)$ on t_v .

Now assume $v \in X_{\psi', t, n}^{m+1}$, so $\llbracket \psi' \rrbracket(t, v, X_{\psi', t, n}^m) \leq n$. As in the previous case, we can construct an initial part of a strategy for \mathcal{A} by extending plays from \mathcal{A}_j with plays from some \mathcal{A}_i . The difference here is that \mathcal{A}_i may return to state q_j at a position u . If it does, then it must be the case that $u \in X_{\psi', t, n}^m$ by definition of \mathcal{A}_i (otherwise, it would contradict $\llbracket \psi' \rrbracket(t, v, X_{\psi', t, n}^m) < \infty$). Thus, in order to construct a full strategy for \mathcal{A} , we continue plays from u by using the strategy given by the inductive hypothesis that witnesses for all $u \in X_{\psi', t, n}^m$, $\llbracket \mathcal{A} \rrbracket(t_u) \leq \alpha_j^m(n)$. This results in plays of value at most $\alpha_w(n) + \alpha_{j-1}(n) + \alpha_j^m(n) \leq \alpha_j^{m+1}(n)$ for \mathcal{A} on t_v . \square