

Resource-Aware Automata and Games for Optimal Synthesis

Corina Cîrstea

University of Southampton, UK

cc2@ecs.soton.ac.uk

We consider quantitative notions of parity automaton and parity game aimed at modelling resource-aware behaviour, and study (memory-full) strategies for exhibiting accepting runs that require a minimum amount of initial resources, respectively for winning a game with minimum initial resources. We also show how such strategies can be simplified to consist of only two types of moves: the former aimed at increasing resources, the latter aimed at satisfying the acceptance condition.

1 Introduction

This paper studies quantitative versions of Büchi/parity automata on infinite words/trees, as well as of Büchi/parity games [10], aimed at the modelling, verification and synthesis of *resource-aware* systems. In the case of automata, our results concern synthesising strategies for exhibiting *accepting runs* that are (i) *resource-aware*, in the sense that at any point during the run, sufficient resources are available to take the next transition, and (ii) *optimal* in their requirement on initial resources. In the case of games, we are interested in winning strategies that require the smallest amount of initial resources.

Our approach builds on earlier work on a general approach to quantitative verification [7, 6, 8, 5]. This work uses *coalgebras* [11] to model systems with quantitative features and quantitative fixpoint logics to specify correctness properties, and extends the standard automata-based approach to verification to this setting. This general approach is parameterised by (i) a quantitative branching type (modelled via a partial semiring of quantities, e.g. probabilities or resources) and (ii) a qualitative behaviour type (e.g. infinite words or trees): each choice for these two ingredients gives rise to a specific model type together with an associated quantitative logic, for which the model checking problem reduces to computing the *extent* (generalising non-emptiness) of a certain quantitative parity automaton [8]. At the heart of our approach are nested fixpoints, akin to those used in the semantics of the modal μ -calculus.

Here we focus on a concrete instantiation for the universe of quantities, with the resulting quantitative automata being used to model resource-aware behaviour. The qualitative behaviour type remains a parameter, thus making our results applicable to automata over both *words* and *trees*. To maintain this level of generality, and at the same time be succinct in our presentation, we continue to rely on a coalgebraic formulation of key concepts; however, the use of coalgebraic machinery is minimised, to make our results accessible also to readers not familiar with coalgebra.

The notion of resource-aware automaton we consider is similar to that of [5], and more general than the one in [8], in that not only resource *use* but also resource *gain* is modelled. Unlike non-deterministic parity automata/games where strategies for exhibiting accepting runs / winning a game are memoryless [12], we show that, in the resource-aware case, *memory-full strategies* are needed even for automata. This is because the synthesised run must not only satisfy the parity condition, but also respect available resources; as a result, the transition to take in a state depends on the available resources, being aimed either at accumulating additional resources or at advancing towards an accepting state.

To study such memory-full strategies, we exploit a connection to standard parity games. Specifically, we associate to each resource-aware parity automaton/game a standard parity game, and show how memory-full strategies for exhibiting optimal accepting runs in the original automaton can be derived

from memoryless winning strategies in the associated parity game. Further, we show how to reduce these strategies to only *two* moves per state while maintaining their winning property. Finally, we show how the algorithm for computing the extent of a quantitative automaton [8] can be enhanced to compute such strategies.

While the majority of the paper concerns automata, we also consider a notion of quantitative parity game, parameterised by a semiring of quantities, and show how the notion of extent generalises from automata to games. Moreover, in the resource-aware case, we show that the construction of strategies for exhibiting optimal accepting runs also generalises – where this time, the objective is to synthesise optimal winning strategies for player \exists .

Our algorithms for computing optimal strategies apply to resource-aware *Büchi* automata/games on words/trees; a generalisation to *parity* automata is discussed briefly, with a full treatment deferred to future work. In the Büchi case, the computation of optimal strategies has complexity quadratic in the size of the automaton/game, *assuming extents have already been computed*. (The computation of extents has complexity $O(n^{k+1})$ with n the automaton size and k the number of parities.)

Our results yield techniques for the synthesis of behaviours within specified design spaces (automaton case), or the quantitative synthesis of strategies in non-cooperative settings (game case), in scenarios where both correctness and optimality are important. Similar combinations of qualitative and quantitative objectives have been studied previously. [4] considered games that combine parity objectives with mean-payoff ones, and studied algorithms for computing the values (for each of the two players) associated to such games along with strategies to achieve them. (Unlike parity or mean-payoff games, such strategies are not anymore memoryless; in fact they may require *infinite* memory.) [3] considered weighted parity games, with weights thought of as energy values and with objectives combining a parity condition with the requirement to maintain a positive energy level. An algorithm for solving such games with complexity exponential in the size of the game and linear in the largest weight was described, and it was shown that *finite* (exponential) memory suffices to represent optimal strategies. [9] combined parity/Streett conditions with quantitative constraints bounding the costs between requests (odd parity states) and responses (higher, even parity states), and showed that solving such games can be reduced to solving several instances of their purely qualitative variants. A generalisation from (positive) costs to weights that can be either positive or negative was described in [13], and shown to be polynomial-time equivalent to the energy parity games of [3]. Finally, [2] considered weighted *timed* games with energy constraints, where positive/negative weights can be associated to both locations and transitions, and showed that the existence of strategies ensuring that the accumulated weight in any finite prefix of the resulting play stays within a given interval is undecidable.

While our work addresses a similar problem, a key advantage compared to existing work is the applicability of our results to a *wider class of qualitative behaviours*: not just infinite words, but also infinite trees or a combination of word- and tree-like behaviour, including terminating behaviour. Such an extension allows modelling systems whose structure varies over time (e.g. new processes being created and other processes terminating while the system executes), where now the goal is to synthesise an optimal *system* w.r.t. the *overall* amount of resources required. Our use of coalgebras as models allows for an increased flexibility in the type of qualitative behaviour such systems can exhibit; whereas our use of coinduction on a technical level supports dealing with such general models of behaviour (where, for example, the standard notion of a cycle, as used e.g. in [3], becomes more complex).

Our setting is closest in spirit to the work on *energy games* [3], and our results on reduced-memory strategies (Section 4.3) are inspired by similar results in loc. cit. (although our proofs are different from those in [3]); we expand on the connection to [3] later in the paper (see Remark 5.6). The main *technical* difference between our approach and existing work is that we exploit a quantitative generalisation

of the nested fixpoints used to decide non-emptiness of parity automata / the winner of a parity game. This requires monotonicity/continuity assumptions, to allow a generalisation of the fixpoint approach to semiring-valued maps. Our choice to associate resource gains to automaton states and resource usage to automaton transitions (see Definition 3.1) ensures that the required assumptions are satisfied. Indeed, this particular choice of model structure is driven by the need to satisfy these assumptions.

The paper is structured as follows. After some basic preliminaries (Section 2), we recall the notion of *weighted parity automaton with offsetting* and the associated notion of *extent* [5], and show how one can associate a *value* to each run of such an automaton, which in the resource-aware instantiation measures the minimal initial resources required to execute the run (Section 3). The main contributions are:

1. a study of memory-full strategies for exhibiting optimal accepting runs in resource-aware *Büchi* automata over words/trees, including a simplified version of such strategies that uses only two moves for each state, along with an algorithm for computing such strategies (Section 4),
2. a notion of weighted parity game with offsetting, and a generalisation of the notion of extent to such games (Section 5),
3. a brief study of winning strategies for the resource-aware instance of weighted parity games with offsetting, outlining how the results of Section 4 generalise to the game setting (Section 6). (A full account of the game case is deferred to an extended version of the paper.)

We conclude by summarising our results and discussing future work (Section 7).

2 Preliminaries

Semirings and Monads We take as parameter a *commutative semiring* $(S, +, 0, \bullet, 1)$. Such a semiring induces a preorder \sqsubseteq on S given by $x \sqsubseteq y$ iff there exists $z \in S$ s.t. $x + z = y$, with 0 as bottom element.

Assumption 2.1. We assume that (S, \sqsubseteq) is a complete lattice with 1 as top element, and that both $+$ and \bullet preserve suprema of increasing chains and infima of decreasing chains, in each argument.

Example 2.2. The following examples of semirings satisfy our assumptions:

1. the *boolean semiring* $(\{0, 1\}, \vee, 0, \wedge, 1)$, with associated order \leq on $\{0, 1\}$,
2. the *tropical semiring* $\mathbb{N}^\infty = (\mathbb{N}^\infty, \min, \infty, +, 0)$, with order \geq on $\mathbb{N}^\infty = \mathbb{N} \cup \{\infty\}$,
3. for $B \in \mathbb{N}$, the *B-bounded variant of the tropical semiring* $\mathbb{N}_B^\infty = (\mathbb{N}_B^\infty, \min, \infty, +_B, 0)$, where $\mathbb{N}_B^\infty = \{n \in \mathbb{N} \mid 0 \leq n \leq B\} \cup \{\infty\}$ and where $m +_B n = \begin{cases} m + n, & \text{if } m + n \leq B \\ \infty, & \text{otherwise} \end{cases}$, with associated order \geq ,
4. the *tropical positive rationals* $\mathbb{Q}_{\geq 0}^\infty = (\mathbb{Q}^\infty, \min, \infty, +, 0)$ and its *B-bounded variant* for $B \in \mathbb{Q}^+$, with $\mathbb{Q}_{\geq 0}^\infty = \{q \in \mathbb{Q} \mid q \geq 0\} \cup \{\infty\}$ and structure similar to the tropical semiring.

A commutative semiring S as above induces a *semiring monad* (T_S, η, μ) , with $T_S(X) = \{\varphi : X \rightarrow S \mid \text{supp}(\varphi) \text{ is finite}\}$, and *unit* and *multiplication* given by

$$\eta_X(x)(y) = \begin{cases} 1 & \text{if } y = x \\ 0 & \text{otherwise} \end{cases}, \quad \mu_X(\Phi)(x) = \sum_{\varphi \in \text{supp}(\Phi)} \Phi(\varphi) \bullet \varphi(x)$$

where $\text{supp}(\varphi) = \{x \in X \mid \varphi(x) \neq 0\}$ is the *support* of φ . We use formal sum notation $\sum_{i \in I} c_i x_i$ with I finite to denote the element of $T_S(X)$ mapping x_i to c_i for $i \in I$, and $x \notin \{x_i \mid i \in I\}$ to 0. We also note that $T_S(1) \simeq S$, where 1 denotes a one-element set. In what follows we use $T_S(1)$ and S interchangeably.

Equational Systems Given a complete lattice (L, \sqsubseteq) , a *nested equational system* (cf. [1]) has the form:

$$\left[\begin{array}{l} u_1 =_{\mu} f_1(u_1, \dots, u_n) \\ u_2 =_{\nu} f_2(u_1, \dots, u_n) \\ \vdots \\ u_n =_{\eta} f_n(u_1, \dots, u_n) \end{array} \right] \quad \text{with } \eta = \mu \text{ (} \eta = \nu \text{) if } n \text{ is odd (resp. even), with } u_1, \dots, u_n \text{ ranging over } L. \text{ The reader is referred to [1, Section 1.4.4] for a formal definition of the solution of a nested equational system. Here we only give an informal description. We assume that the most significant equation is the } \textit{last} \text{ one.}$$

The solution is obtained as follows: a solution for u_1 parameterised by values for u_2, \dots, u_n is obtained from the first equation and substituted into the remaining ones; a solution for u_2 parameterised by values for u_3, \dots, u_n is obtained from the second equation and substituted into the remaining ones; and so on until a closed solution for u_n is obtained; this is substituted back to give closed solutions for u_{n-1}, \dots, u_1 .

Two-Player Games A *game graph* is given by a pair (Q, E) with Q a finite set of *states* partitioned as $Q = Q_{\exists} \sqcup Q_{\forall}$, and $E \subseteq Q \times Q$ a relation describing possible *moves*. We write $E_{\exists} := \{(q, q') \in E \mid q \in Q_{\exists}\}$ for the set of \exists moves and similarly for E_{\forall} . A *two-player game* played on a game graph (Q, E) by players \exists and \forall starts in an *initial state* $q_0 \in Q$ and proceeds by either \exists or \forall choosing a successor state, depending on whether the current state belongs to Q_{\exists} or Q_{\forall} . A (*complete*) *play* is a finite or infinite path $q_0 q_1 \dots$ with $(q_i, q_{i+1}) \in E$ for all i , where if the path ends in q_n then there is no $q \in Q$ with $(q_n, q) \in E$. A *parity game* is a game graph (Q, E) together with a parity map $\Omega : Q \rightarrow \mathbb{N}$ with finite codomain. A complete play in a parity game is won by \exists iff it is either a finite play $q_0 q_1 \dots q_n$ with $q_n \in \forall$, or an infinite play $q_0 q_1 \dots$ with $\max\{n \mid n = \Omega(q_i) \text{ for infinitely many } i\}$ even; otherwise the play is won by \forall .

A *strategy* for \exists in such a game maps partial plays $q_0 q_1 \dots q_i$ with $q_i \in Q_{\exists}$ to a next move $(q_i, q) \in E$. Strategies for \forall are defined similarly. A strategy is *winning for* \exists iff \exists wins all plays which conform to this strategy. A strategy for \exists is said to be *finite memory* if it can be encoded using a finite set M along with two update functions, $\sigma_s : M \times Q_{\exists} \rightarrow Q$ and $\sigma_m : M \times (Q_{\exists} + E_{\forall}) \rightarrow M$, with $(q, \sigma_s(m, q)) \in E$. Assuming an initial memory value $m_0 \in M$, such a pair (σ_s, σ_m) yield a strategy that in state $q \in Q_{\exists}$ with memory value m prescribes a move to $\sigma_s(m, q)$; moreover, following the prescribed \exists move from q (an arbitrary \forall move $e \in E$), the memory is updated to $\sigma_m(m, q)$ (respectively $\sigma_m(m, e)$).

Coalgebras A *coalgebra* for a functor $F : \text{Set} \rightarrow \text{Set}$ is given by a pair (C, γ) with C a set (of *states*) and $\gamma : C \rightarrow F(C)$ a function (the *transition map*). A *pointed coalgebra* (C, γ, c_0) additionally provides an initial state $c_0 \in C$. The intuition is that, for a state $c \in C$, $\gamma(c)$ describes the one-step observations that one can make of c , structured according to F (see [11, Chapters 1, 2] for an introduction to coalgebras).

Example 2.3. Let $T_S : \text{Set} \rightarrow \text{Set}$ be as in Section 2. A T_S -coalgebra (Q, γ) is the same as a finitely-branching transition system with set of states Q and transitions labelled by quantities in S . Given such a quantitative transition system, the corresponding γ is given by $\gamma(q)(q') = s$ whenever the transition $q \rightarrow q'$ carries the quantity s , and $\gamma(q)(q') = 0_S$ whenever there is not transition $q \rightarrow q'$.

[8] consider coalgebras of type $T_S \circ F$, where $T_S : \text{Set} \rightarrow \text{Set}$ is as above and $F : \text{Set} \rightarrow \text{Set}$ is given by $F(X) = \coprod_{\lambda \in \Lambda} X^{\text{ar}(\lambda)}$, with \coprod denoting disjoint union and Λ a set of operation symbols with finite arities. Thus, an F -coalgebra (C, γ) (used in what follows to model *runs* of a parity automaton) consists of a set of states C together with, for each state $c \in C$, a transition of type $c \rightarrow (\lambda, c_1, \dots, c_{\text{ar}(\lambda)})$ for some $\lambda \in \Lambda$. Taking $F(X) = A \times X \simeq \coprod_{a \in A} X$ will allow us to cover (automata over) infinite words over A , while taking $F(X) = A \times X \times X \simeq \coprod_{a \in A} X \times X$ will cover infinite trees with nodes labelled by elements of A . More complex choices of F are also possible: e.g. $F(X) = X + X \times X + \{*\}$ models runs which combine word-like structure with tree-like structure as well as termination: in each state of the run, there is a *choice* of continuing in another state (first term), splitting into two states (second term) or terminating (last term). Finally, the presence of T_S in the coalgebra type will model the existence of *several* transitions from each state of an automaton, with each transition being assigned a quantity in S .

3 Weighted Parity Automata with Offsetting

Throughout this section, $(S, +, 0, \bullet, 1)$ is a commutative semiring satisfying the assumptions in Section 2. We also assume $F : \text{Set} \rightarrow \text{Set}$ is given by $F(X) = \coprod_{\lambda \in \Lambda} X^{\text{ar}(\lambda)}$ with Λ a set of operation symbols with finite arities. The functor F specifies the structure of individual runs of a quantitative automaton. While an elementary definition of weighted parity automata is possible for each F , the following coalgebraic definition, parameterised by F , captures all such instances.

Definition 3.1 ([5]). *A weighted parity automaton with offsetting (WPAO) is given by a $S \times (\text{T}_S \circ F)$ -coalgebra $(Q, \langle r, \gamma \rangle)$ together with a parity map $\Omega : Q \rightarrow \mathbb{N}$ assumed to have finite range. We call $r : Q \rightarrow S$ the offset map and $\gamma : Q \rightarrow \text{T}_S(F(Q))$ the transition map. For $n \in \text{ran}(\Omega)$, we write $Q_n = \{q \in Q \mid \Omega(q) = n\}$. If $\text{ran}(\Omega) = \{1, 2\}$, we call \mathcal{A} a weighted Büchi automaton with offsetting (WBAO).*

WPAOs generalise a standard parity automata in several ways: transitions carry weights from S , states carry offsets, also from S , and the structure of transitions strictly subsumes automata over words/trees.

Taking $FX = \text{Id}$ (the identity functor), $S = (\{0, 1\}, \vee, 0, \wedge, 1)$ and ignoring offsets yields the standard notion of parity automaton with the alphabet a singleton set. Taking $S = (\mathbb{N}^\infty, \min, \infty, +, 0)$ and $F = \text{Id}$ yields a variant of parity automata over infinite words where both transitions and states are assigned values in S , to be viewed as resources consumed by transitions, and respectively resources gained upon visiting states. In the latter case, we will use an offsetting operation $\ominus : \mathbb{N}^\infty \times \mathbb{N}^\infty \rightarrow \mathbb{N}^\infty$ to capture the fact that resources gained can only be used for future computation. This operation is defined by $n \ominus m = \begin{cases} \max(n - m, 0), & \text{if } m \neq \infty \text{ or } n \neq \infty \\ \infty, & \text{otherwise} \end{cases}$, with \max the standard one on \mathbb{N}^∞ . Thus, \ominus is a subtraction operation capped below at $0 \in \mathbb{N}$; this matches the intuition that m acts as an offset value.

The operation \ominus on \mathbb{N}^∞ generalises to an arbitrary, and even *partial* semiring $(S, +, 0, \bullet, 1)$ [5]. Specifically, the binary operation $\oslash : S \times S \rightarrow S$ is defined by $s \oslash t = \inf\{u \mid u \bullet t \sqsupseteq s\}$. The notation used reflects that \oslash is *almost* an inverse to the semiring multiplication. Well-definedness follows by Assumption 2.1. When $S = (\{0, 1\}, \vee, 0, \wedge, 1)$, $\oslash : S \times S \rightarrow S$ instantiates to the first projection.

Proposition 3.2 ([5]). *The following hold for any $(S, +, 0, \bullet, 1)$ satisfying the assumptions in Section 2:*

1. $(a \oslash c) + (b \oslash c) = (a + b) \oslash c$,
2. $\sup_{i \in \omega} (a_i \oslash c) = (\sup_{i \in \omega} a_i) \oslash c$, for $a_0 \sqsubseteq a_1 \sqsubseteq \dots$,
3. $\inf_{i \in \omega} (a_i \oslash c) = (\inf_{i \in \omega} a_i) \oslash c$, for $a_0 \sqsupseteq a_1 \sqsupseteq \dots$

Next, we define the notion of an (*accepting*) run of a WPAO (which generalises the standard notions from parity word/tree automata), and use transition weights and state offsets to associate *values* to runs.

Definition 3.3. *Given a set Q (of states), an F -run with states in Q is a possibly infinite tree with nodes labelled by pairs (q, λ) with $q \in Q$ and $\lambda \in \Lambda$, and with each node (q, λ) having exactly $\text{ar}(\lambda)$ children $(q_1, \lambda_1), \dots, (q_n, \lambda_n)$. For a WPAO $\mathcal{A} = (Q, \langle r, \gamma \rangle)$, a run of \mathcal{A} additionally satisfies $\gamma(q)(\iota_\lambda(q_1, \dots, q_n)) \neq 0_S$ for each node (q, λ) with children $(q_1, \lambda_1), \dots, (q_n, \lambda_n)$ ¹. We write *Runs* for the set of F -runs with states in Q , and $\text{Runs}_{\mathcal{A}}$ for the set of runs of \mathcal{A} . A run is *accepting* iff for each branch of the associated tree, the maximum parity occurring infinitely often is even.*

Thus, at each state q , a run of \mathcal{A} selects an available transition, say of type λ , and proceeds with a number of successors equal to $\text{ar}(\lambda)$ – this can be 0 (termination), 1 (linear structure) or ≥ 2 (tree structure).

Definition 3.4. *For a WPAO $\mathcal{A} = (Q, \langle r, \gamma \rangle)$, the function $\text{val} : \text{Runs} \rightarrow S$ is the greatest fixpoint of the operator on S^{Runs} taking $v : \text{Runs} \rightarrow S$ to $v' : \text{Runs} \rightarrow S$, where for $z \in \text{Runs}$ with root (q, λ) and children z_1, \dots, z_n labelled by q_1, \dots, q_n , respectively:*

$$v'(z) = (\gamma(q, \iota_\lambda(q_1, \dots, q_n)) \bullet v(z_1) \bullet \dots \bullet v(z_n)) \oslash r(q)$$

¹That is, a transition from q to $(\lambda, q_1, \dots, q_n)$ exists in \mathcal{A} .

The above coinductive definition formalises the idea that, in order to obtain the value of a run z with root (q, λ) and children z_1, \dots, z_n , one computes increasingly finer approximations for $\text{val}(z)$ by using, at each step, the current approximations for $\text{val}(z_1), \dots, \text{val}(z_n)$; at each state q along the run, the value $r(q)$ is used to offset the "future" value. As a result, the function val maps any $z \in \text{Runs}$ which is not a run of \mathcal{A} to 0_S . The existence of the required fixpoint follows by Kleene's theorem, from (S, \sqsubseteq) being a complete lattice together with the operator involved being co-continuous (by Assumption 2.1 and Proposition 3.2).

Example 3.5. When $S = (\mathbb{N}^\infty, \min, \infty, +, 0)$, the value of a run measures the minimal resources required in the initial state to ensure that the run can be executed (assuming that resources are increased using offset values and decreased using transition costs). The use of \bullet in Definition 3.4 ($+$ on \mathbb{N}^∞) amounts to adding the costs of all transitions across all branches of a run, whereas the use of \ominus (\ominus on \mathbb{N}^∞) ensures that the resource gain in each visited state offsets the cost of *future* transitions. When $S = (\{0, 1\}, \vee, 0, \wedge, 1)$ and assuming trivial offset values of 1, a run has value 1 iff it is a run of \mathcal{A} in the standard sense.

We now recall the notion of *extent* of a WPAO, which generalises non-emptiness of parity automata [5]. The extent associates a value in S to each automaton state, with the intuition that this value "amalgamates" (via the semiring addition) the values of *all* accepting runs from that state. Crucially, this amalgamation is performed in a step-wise fashion and does not require the values of individual runs to be known. The coalgebraic formulation, if somewhat abstract, allows for a generic and concise definition.

Definition 3.6 (Extent of WPAO [5]). *Let $\mathcal{A} = (Q, \langle r, \gamma \rangle, \Omega)$ be a WPAO with $\text{ran}(\Omega) = \{1, \dots, n\}$. For $k \in \text{ran}(\Omega)$, let $\gamma_k : Q_k \rightarrow T_S(F(Q))$ and $r_k : Q_k \rightarrow S$ denote the restrictions of γ and respectively r to $Q_k := \{q \in Q \mid \Omega(q) = k\}$. The extent $\text{ext}_{\mathcal{A}} = [e_1, \dots, e_n] : Q \rightarrow S$ of \mathcal{A} is the solution of the following nested equational system, with the most significant equation being the last one:*

$$\begin{cases} u_1 =_{\mu} (\mu_1 \circ T_S(\bullet_F) \circ T_S F[u_1, \dots, u_n] \circ \gamma_1) \ominus r_1 \\ u_2 =_{\nu} (\mu_1 \circ T_S(\bullet_F) \circ T_S F[u_1, \dots, u_n] \circ \gamma_2) \ominus r_2 \\ \vdots \\ u_n =_{\eta} (\mu_1 \circ T_S(\bullet_F) \circ T_S F[u_1, \dots, u_n] \circ \gamma_n) \ominus r_n \end{cases} \quad (1)$$


with $\eta = \mu$ ($\eta = \nu$) if n is odd (resp. even), with variables u_k ranging over the poset (S^{Q_k}, \sqsubseteq) (and thus $[u_1, \dots, u_n] : Q \rightarrow S$), and with the first operands in the rhss pictured below:

$$Q_k \xrightarrow{\gamma_k} T_S F Q \xrightarrow{T_S F[u_1, \dots, u_n]} T_S F S \xrightarrow{T_S(\bullet_F)} T_S S = T_S^2 1 \xrightarrow{\mu_1} T_S 1 = S$$

In the above, $\bullet_F : FS \rightarrow S$ is given by $\bullet_F(l_\lambda(s_1, \dots, s_{\text{ar}(\lambda)})) = s_1 \bullet \dots \bullet s_{\text{ar}(\lambda)}$ for $\lambda \in \lambda$, and $\mu_1 : T_S S = T_S T_S 1 \rightarrow T_S 1 = S$ is the monad multiplication.

Take $S = (\mathbb{N}^\infty, \min, \infty, +, 0)$ in Definition 3.6. By (1), the extent value associated to $q \in Q_k$ is the minimum (use of monad multiplication μ_1) across all transitions from q (use of γ_k) of the combined extents (use of \bullet , in this case $+$) of the successors of q in that transition, and where offsets (use of r_k) are applied to account for resource gain in q . The existence of the nested fixpoint follows again by Kleene's theorem.

Example 3.7. Consider the resource-aware automaton on the left below ($F = \text{Id}$), with offsets of 0 in x and y , 2 in y_1 , and 4 in y_2 . The associated equational system, on the right below, has solution $(1, 1, 0, 0)$:



$$\begin{cases} e_x =_{\nu} e_y \\ e_y =_{\mu} \min(e_x + 5, e_{y_1} + 1, e_{y_2} + 2) \\ e_{y_1} =_{\mu} e_y \ominus 2 \\ e_{y_2} =_{\mu} e_y \ominus 4 \end{cases}$$

This matches the intuition that an initial resource value of 1 suffices in both x and y to visit the accepting state x infinitely often, while never running out of resources: the transitions $x \rightarrow y$ and $y \rightarrow y_1$, together consuming 1 resource, can be used to reach y_1 , and no additional resources are required from y_1 or y_2 , given the resource gains in these states: one can loop through y_1 , y and y_2 until the available resources in y reach 6 and the transition to x can be taken. On the other hand, the run $xyxy \dots$ has value ∞ .

As expected, no *accepting* run of a WPAO can have a value that is strictly above (w.r.t. \sqsubseteq) the extent of its initial state. The proof of this result (omitted due to space limitations) uses coinduction.

Proposition 3.8. *Let $\mathcal{A} = (Q, \langle r, \gamma \rangle, \Omega)$ be a WPAO, and let $z_0 \in \text{Runs}_{\mathcal{A}}$ be an accepting run of \mathcal{A} with root q_0 . Then, $\text{ext}_{\mathcal{A}}(q_0) \sqsupseteq \text{val}(z_0)$.*

When $S = (\{0, 1\}, \vee, 0, \wedge, 1)$ and with a trivial offset function, the extent associates a value of 1 to an automaton state iff there exists an accepting run from that state. Also, when $S = (\mathbb{N}^{\infty}, \min, \infty, +, 0)$, the extent associates to a state q the minimum value of a run (Definition 3.3), calculated across all accepting runs from q : that the extent is \leq the minimum value follows from Proposition 3.8; the existence of a run with the same value as the extent will follow from Theorem 4.10.

[8] describes an algorithm for computing extents of weighted parity automata (no offsetting), under the additional assumption that the order \sqsubseteq admits no infinite strictly ascending/descending chains. The algorithm, specialised to the resource-aware setting (B -bounded variant of the tropical semiring, which satisfies the above assumption), is given in Figure 1. If $N_i = |Q_i|$ and E_i is the number of edges from states in Q_i (where an edge of type λ is counted $\text{ar}(\lambda)$ times), for $i \in \text{ran}(\Omega)$, and if $k = |\text{ran}(\Omega)|$, the algorithm has complexity $O(N_k \times (N_k + E_k + N_{k-1} \times (N_{k-1} + E_{k-1} + \dots + N_1 \times (N_1 + E_1)))$. Thus, if $n = \sum_{i \in \text{ran}(\Omega)} (N_i + E_i)$ is the size of the automaton, the complexity is $O(n^{k+1})$. The actual upper bound for the time complexity is in fact proportional to $B^{|\text{ran}(\Omega)|}$, since for each $i \in \text{ran}(\Omega)$, the **repeat ... until** statement in $\text{Extent}(i)$ is executed $B \times N_i$ times in the worst case.

Figure 1: Algorithm for computing extents of resource-aware WPAOs [8]

<p>Input: resource-aware WPAO $\mathcal{A} = (Q, \langle r, \gamma \rangle, \Omega)$</p> <p>Output: extent $e : Q \rightarrow \mathbb{N} \cup \{\infty\}$ of \mathcal{A}</p> <ol style="list-style-type: none"> 1. $\text{Extent}(\max(\text{ran}(\Omega)))$ <p>Procedure $\text{Extent}(n \in \mathbb{N})$</p> <ol style="list-style-type: none"> 1. if $n = 0$ then return endif 2. for $q \in Q_n$ do 3. let $e(q) := \begin{cases} 0, & \text{if } n \text{ is even} \\ \infty, & \text{if } n \text{ is odd} \end{cases}$ 4. endfor 5. repeat 6. let $old := e$ 7. $e := \text{Extent}(n-1)$ 8. for $q \in Q_n$ do 9. $e(q) \leftarrow \min_{i \in I} ((v_i + old(q_1) + \dots + old(q_{\text{ar}(\lambda)})) \ominus r(q))$ where $\gamma(q) = \sum_{i \in I} v_i \iota_{\lambda}(q_1, \dots, q_{\text{ar}(\lambda)})$ 10. endfor 11. until $e = old$ 	<ul style="list-style-type: none"> • Lines 8–10 of the <i>Extent</i> procedure compute a better approximation of the extent for automaton states with the current parity n, based on a one-step unfolding of the transition structure. • Line 7 computes the extent of states with immediately lower parity, relative to the current values for states with parity n, through a recursive call to <i>Extent</i> (which may involve further recursive calls). • Recursive calls to <i>Extent</i> update the same copy of e, and only make an additional copy to remember values from the previous step.
--	--

4 Strategies in Resource-Aware Büchi Automata

For the remainder of the paper we focus on *resource-aware WBAOs*, that is, we fix $S = (\mathbb{N}_B^{\infty}, \min, \infty, +, 0)$ with $B \in \mathbb{N}$ and assume $\text{ran}(\Omega) = \{1, 2\}$. While some of our basic results apply more generally to

resource-aware *parity* automata, extensions of our main results in this section (Theorems 4.13 and 4.16) to parity automata are only sketched, with a full account being left as future work.

For the above choice of semiring, the generic order \sqsubseteq instantiates to the \geq order on \mathbb{N}_B^∞ . When referring to least/greatest fixpoints, we will implicitly assume the generic order \sqsubseteq . Yet, depending on context, we will also talk about least upper bounds/greatest lower bounds w.r.t. the order \leq on \mathbb{N}_B^∞ .

Our results concern the construction of strategies for exhibiting *optimal* runs.

Definition 4.1. *A run from q_0 in a WPAO \mathcal{A} is optimal if it is accepting and has value $\text{ext}_{\mathcal{A}}(q_0)$.*

Unlike standard parity automata on infinite words/trees, strategies that exhibit optimal runs in resource-aware automata are not memoryless; intuitively, this is because proceeding from a non-accepting state q with available resources equal to $\text{ext}_{\mathcal{A}}(q)$ may require accumulating additional resources in q before moving towards an accepting state. We tackle the problem of finding strategies for exhibiting optimal runs in a WPAO by moving to a standard parity game called the *resource-aware game* (Section 4.1). In this game, available resources are recorded within states, \exists moves correspond to transitions that can be taken with the currently available resources, and \forall moves arise from the presence of branching in runs (e.g. in the case of WBAOs on trees). We show here how the algorithm in Figure 1 can be extended to compute a memoryless winning strategy for \exists in the resource game, and use this to derive a memory-full strategy for exhibiting an optimal run in \mathcal{A} (Section 4.2). We then show that in fact simpler strategies, combining a *good-for-energy* strategy with an *attractor* strategy, exist (Section 4.3).

4.1 Resource Games for WPAOs


Definition 4.2 (Resource Game). *Let $\mathcal{A} = (Q, \langle r, \gamma \rangle, \Omega)$ be a resource-aware WPAO. The resource game $\mathcal{R}_{\mathcal{A}}$ of \mathcal{A} is a standard parity game with states $Q_{\exists} = \{(q, n) \in Q \times \mathbb{N} \mid n \geq \text{ext}_{\mathcal{A}}(q)\}$ and $Q_{\forall} = F(Q_{\exists})$, parities inherited from \mathcal{A} on \exists states and equal to $\min\{\Omega(q) \mid q \in Q\}$ on \forall states, and moves given by:*

$$(q, n) \xrightarrow{\exists} (\lambda, (q_1, n_1), \dots, (q_m, n_m)) \quad \text{whenever } n + r(q) \geq \gamma(q)(\iota_{\lambda}(q_1, \dots, q_m)) + n_1 + \dots + n_m,$$

$$(\lambda, (q_1, n_1), \dots, (q_m, n_m)) \xrightarrow{\forall} (q_i, n_i) \quad \text{for all } i \in \{1, \dots, m\}.$$

The reason for requiring $n \geq \text{ext}_{\mathcal{A}}(q)$ in the definition of Q_{\exists} is that, by Proposition 3.8, at least $\text{ext}_{\mathcal{A}}(q)$ resources are required to exhibit an *accepting* run from q . A transition from q to $(\lambda, q_1, \dots, q_m)$ in \mathcal{A} with weight $w = \gamma(q)(\iota_{\lambda}(q_1, \dots, q_m))$ gives rise to one \exists move from (q, n) in $\mathcal{R}_{\mathcal{A}}$ for each choice to distribute the available resources following the transition (that is, after gaining $r(q)$ and consuming w) to acceptable resources for the states q_1, \dots, q_m (that is, $n_i \geq \text{ext}_{\mathcal{A}}(q_i)$ for $i \in \{1, \dots, m\}$). The existence of at least one such move is guaranteed by the definition of extents together with the assumption that extents are finite for all \exists states in $\mathcal{R}_{\mathcal{A}}$. Finally, note that assigning the lowest parity to \forall states results in these parities not affecting the winner of a play, since Q_{\exists} and Q_{\forall} states alternate strictly in any play.

Example 4.3. Part of the game model associated to the WBAO in Example 3.7 is depicted below. Since $F = \text{Id}$, \forall moves are fully determined and therefore not shown.



Take, for example, the two transitions from state $(y, 6)$. The transition leading to $(x, 1)$ consumes the 5 resources required by $y \rightarrow x$, whereas the transition to $(y_1, 5)$ consumes a single resource as required by $y \rightarrow y_1$. On the other hand, the transition from $(y_2, 2)$ to $(y, 6)$ consumes no resources, while additionally gaining 4 resources as a result of visiting state y_2 .

Next, we define runs of \mathcal{A} that never run out of resources.

Definition 4.4. A resource-aware run of a WPAO \mathcal{A} is a run of \mathcal{A} (Definition 3.3) whose states can be annotated by values in \mathbb{N} in such a way that paths through the annotated tree correspond to plays in $\mathcal{R}_{\mathcal{A}}$.

The characterisation of the value of a run as a greatest fixpoint w.r.t. \sqsubseteq (least fixpoint w.r.t. \leq) can be used to show that the value of a resource-aware run with root q_0 annotated by r_0 is at most r_0 .

Proposition 4.5. If z is a resource-aware run of a WPAO \mathcal{A} , with root annotated by r_0 , then $\text{val}(z) \leq r_0$.

By Propositions 3.8 and 4.5, an *accepting* run can only exist from state (q, r) in $\mathcal{R}_{\mathcal{A}}$ if $r \geq \text{ext}_{\mathcal{A}}(q)$. To show that an accepting run from $(q, \text{ext}_{\mathcal{A}}(q))$ does indeed exist, we will construct a strategy to exhibit it.

4.2 Memory-full Strategies in WBAOs

From now on we assume, for the most part, that \mathcal{A} is a WBAO. The extension to parity automata is sketched at the end of Section 4.3. We first consider a sub-game of $\mathcal{R}_{\mathcal{A}}$ derived from the execution of the algorithm in Figure 1. Specifically, we assume that the values $e(q)$ with $q \in Q_2$ are already computed, and consider the last recursive call to *Extent*(1) and the updates made to $e(q)$ with $q \in Q_1$ within this call. An assignment to $e(q)$ qualifies as an update if the value of $e(q)$ is *strictly* decreased.

Definition 4.6. Let $\mathcal{A} = (Q, \langle r, \gamma \rangle, \Omega)$ be a WBAO. The resource sub-game $\mathcal{R}'_{\mathcal{A}}$ of $\mathcal{R}_{\mathcal{A}}$ has states

$$Q'_{\exists} = \{(q, n) \in Q_1 \times \mathbb{N} \mid \text{an update } e(q) \leftarrow n \text{ is performed}\} \cup \{(q, \text{ext}_{\mathcal{A}}(q)) \mid q \in Q_2\}, \quad Q'_{\forall} = FQ'_{\exists}$$

and moves inherited from $\mathcal{R}_{\mathcal{A}}$.

As a result, $\mathcal{R}'_{\mathcal{A}}$ includes all the moves of $\mathcal{R}_{\mathcal{A}}$ which witness updates to $e(q)$, with $q \in Q_1$, as well as moves from $(q, \text{ext}_{\mathcal{A}}(q))$ which witness the value of $(q, \text{ext}_{\mathcal{A}}(q))$, with $q \in Q_2$. Clearly, Q'_{\exists} is finite and contains all states $(q, \text{ext}_{\mathcal{A}}(q))$ with $q \in Q$ and $\text{ext}_{\mathcal{A}}(q)$ finite. Moreover, we have:

Proposition 4.7. Let $\mathcal{A} = (Q, \langle r, \gamma \rangle, \Omega)$ be a WBAO. Then, \exists has a memoryless winning strategy from any state $(q, \text{ext}_{\mathcal{A}}(q))$ in $\mathcal{R}'_{\mathcal{A}}$ (and hence also from any state $(q, \text{ext}_{\mathcal{A}}(q))$ in $\mathcal{R}_{\mathcal{A}}$).

The proof (omitted) constructs a memoryless strategy σ for \exists in $\mathcal{R}'_{\mathcal{A}}$ that in accepting states $(q, \text{ext}_{\mathcal{A}}(q))$ proceeds according to a transition witnessing the value of $\text{ext}_{\mathcal{A}}(q)$ – such a transition is always of the form $(q, \text{ext}_{\mathcal{A}}(q)) \rightarrow (\lambda, (q_1, \text{ext}_{\mathcal{A}}(q_1)), \dots, (q_{\text{ar}(\lambda)}, \text{ext}_{\mathcal{A}}(q_{\text{ar}(\lambda)})))$, and in non-accepting states (q, n) selects a transition witnessing the update of $e(q)$ to n . The resulting strategy provides a way to either terminate or reach an accepting state $(x, \text{ext}_{\mathcal{A}}(x))$ from any state $(q, \text{ext}_{\mathcal{A}}(q))$ in $\mathcal{R}'_{\mathcal{A}}$ via one or more transitions. Moreover, by construction σ has *no redundancy*, in that a σ -conform play can not visit a non-accepting state (y, n) and subsequently a state (y, n') with $n' < n$ without visiting an accepting state in-between. This property, together with the finiteness of σ , are captured by the following definition.

Definition 4.8. A skeleton strategy for a WPAO \mathcal{A} is a strategy for \exists in a finite sub-game \mathcal{R} of $\mathcal{R}_{\mathcal{A}}$ s.t.:

- \mathcal{R} includes all states $(q, \text{ext}_{\mathcal{A}}(q))$ with $\text{ext}_{\mathcal{A}}(q) \neq \infty$, the \exists moves in \mathcal{R} are those in $\mathcal{R}_{\mathcal{A}}$ with source and target in \mathcal{R} , and if a \forall state belongs to \mathcal{R} then so are all $\mathcal{R}_{\mathcal{A}}$ -moves from that state;
- σ has no redundancy.

Given the constraint regarding \forall moves, a *winning* skeleton strategy for \exists in \mathcal{R} yields a winning strategy for \exists in $\mathcal{R}_{\mathcal{A}}$ from $(q, \text{ext}_{\mathcal{A}}(q))$. Note also that \mathcal{R} can be taken to be $\mathcal{R}'_{\mathcal{A}}$. We show next that such a winning skeleton strategy induces a memory-full strategy for exhibiting an optimal run of \mathcal{A} from any $q \in Q$ with $\text{ext}_{\mathcal{A}}(q) \neq \infty$. The induced strategy uses as memory the value of available resources.

Definition 4.9 (Memory-full strategy induced by skeleton). Let σ be a skeleton strategy for a WPAO \mathcal{A} . The memory-full strategy ρ induced by σ uses values in \mathbb{N}_B as memory and is defined by:

- in state $q \in Q$ with memory $s \geq \text{ext}_{\mathcal{A}}(q) \in \mathbb{N}_B$, if $\sigma(q, s) = (\lambda, (q_1, s_1), \dots, (q_n, s_n))$ then ρ prescribes parallel moves to each q_i , with associated memory value s_i , for $i \in \{1, \dots, n\}$.

Thus, ρ is defined in state q with memory value s precisely when (q, s) belongs to \mathcal{R} . As a result, playing ρ from state $q \in Q$ with initial memory value $\text{ext}_{\mathcal{A}}(q)$ will always lead to states where ρ is defined.

Theorem 4.10. *Let σ and ρ be as in Definition 4.9. If σ is winning for \exists in \mathcal{R} , the run from $q \in Q$ obtained by playing ρ (in parallel) with initial memory $\text{ext}_{\mathcal{A}}(q)$ is accepting and has value $\text{ext}_{\mathcal{A}}(q)$.*

Proof. That any ρ -conform run from $q \in Q$ with initial memory $\text{ext}_{\mathcal{A}}(q)$ is accepting follows immediately from the definition of ρ . Also, by Proposition 4.5, the value of an $\mathcal{R}_{\mathcal{A}}$ -run starting in $(a, \text{ext}_{\mathcal{A}}(a))$ cannot exceed $\text{ext}_{\mathcal{A}}(a)$. Since by Proposition 3.8 there can not exist *any* accepting run with value strictly below $\text{ext}_{\mathcal{A}}(a)$, it follows that the value of the run is precisely $\text{ext}_{\mathcal{A}}(a)$. \square

In particular, Theorem 4.10 applies to the strategy σ of Proposition 4.7. Moreover, the algorithm in Figure 1 can be extended to compute σ , once extents have been computed. We only outline the algorithm here, as an improved version is detailed in Section 4.3. The algorithm is an enhanced version of the call to *Extent*(1) in the algorithm in Figure 1, which assumes that $\text{ext}_{\mathcal{A}}(x)$ is known for $x \in Q_2$. In addition to (re-)computing $\text{ext}_{\mathcal{A}}(y)$ for $y \in Q_1$, the algorithm computes a partial function $\sigma : Q_1 \times \mathbb{N} \rightarrow Q$ with domain $\text{dom}(\sigma) \supseteq \{(q, s) \in Q_1 \times \mathbb{N} \mid s = \text{ext}_{\mathcal{A}}(q)\}$ as follows: every time the value of $e(q)$ with $q \in Q_1$ is updated (line 9), and this is witnessed by an \mathcal{A} -transition $q \rightarrow (\lambda, q_1, \dots, q_{\text{ar}(\lambda)})$, set $\sigma(q, e(q))$ to $(\lambda, q_1, \dots, q_{\text{ar}(\lambda)})$ and update the memory accordingly (so as to proceed with memory value $\text{old}(q_i)$ from q_i , for $i \in \{1, \dots, n\}$). Finally, for $q \in Q_2$, define $\sigma(q, \text{ext}_{\mathcal{A}}(q))$ based on *any* transition that witnesses the value of $\text{ext}_{\mathcal{A}}(q)$, and again, update the memory accordingly.

Example 4.11. For the automaton in Example 3.7, removing the move $(y, 6) \rightarrow (y_1, 5)$ from the moves in Example 4.3 yields a winning skeleton strategy.

4.3 Reduced-Memory Strategies for \mathcal{A}

A drawback of the memory-full strategy ρ induced by a skeleton strategy σ for \exists is the amount of memory required to store the values triggering different moves for each $q \in Q_1$. For instance, the strategy in Example 4.11 prescribes moves to each of y_1, y_2 and x in state y , depending on the current memory value. In general, the number of moves is only bounded by the number of states. We now show that substantially simpler strategies exist. Specifically, we consider strategies wherein only *two* moves are prescribed for each $q \in Q_1$: one aimed at reaching an accepting state (provided sufficient resources are available), and one aimed at increasing the available resources in q . We obtain such a strategy by discarding certain moves from a skeleton strategy. However, for the resulting strategy to remain winning, its definition must be modified so as to allow unused resources to be carried over.

To simplify presentation, we only consider *word* automata, i.e. $\text{ar}(\lambda) \leq 1$ for $\lambda \in \Lambda$. In this case, resource-aware accepting runs of \mathcal{A} are essentially plays in $\mathcal{R}_{\mathcal{A}}$ won by \exists – this is because \forall moves are fully determined, effectively resulting in a *one-player* game with skeleton strategies having type $Q \times \mathbb{N}_B \rightarrow Q \times \mathbb{N}_B$. We explain how our results generalise to tree automata at the end of the section.

The next definition is a variation of Definition 4.9 which allows unused resources to be carried over.

Definition 4.12 (Carry-over memory-full strategy induced by skeleton). *Let σ be a skeleton strategy for \mathcal{A} . The carry-over memory-full strategy $\bar{\rho}$ induced by σ is defined as follows:*

- in state $q \in Q$ with memory $s \in \mathbb{N}$, let $n = \max\{m \mid m \leq s, (q, m) \in \text{dom}(\sigma)\}$ and let $\sigma(q, n) = (q', s')$; then $\bar{\rho}$ prescribes a move to q' and updates the memory to $s' + (s - n)$.

Thus, in state q with memory s , $\bar{\rho}$ proceeds according to the σ -move in q that requires the *most* resources, n , not exceeding s . When not all the resources s are required by the move (that is, $s > n$), the unused resources ($s - n$) are carried over. We can then prove a result similar to Theorem 4.10.

Theorem 4.13. *Let \mathcal{A} be a WBAO. If σ and $\bar{\rho}$ are as in Definition 4.12 and σ is winning, then the run from $q \in Q$ obtained by playing $\bar{\rho}$ with initial memory $\text{ext}_{\mathcal{A}}(q)$ is accepting and has value $\text{ext}_{\mathcal{A}}(q)$.*


Proof. The only difference between the $\bar{\rho}$ here and the ρ in Theorem 4.10 is that $\bar{\rho}$ carries over unused resources. This may result in a *different* path to an accepting state (x, n) with $n \geq \text{ext}_{\mathcal{A}}(x)$, in case the resources carried over result in $\bar{\rho}$ making use of a σ -move $(q, n) \rightarrow (q', n')$ in a situation when ρ would have used a σ -move $(q, m) \rightarrow (q', m')$ with $m < n$. However, the fact that σ has no redundancy ensures that any such "jumps" do not result in cycling forever through non-accepting states – no $\bar{\rho}$ -conform play will ever decrease the resources available in a state q , without first visiting an accepting state. \square

In particular, Theorem 4.13 applies to the skeleton σ from Proposition 4.7. We show next how to simplify the induced strategy $\bar{\rho}$ while maintaining its winning property.


Definition 4.14 (Reduced-memory strategy induced by skeleton). *Let σ be a skeleton strategy for \mathcal{A} . For $q \in Q$, the value $\theta(q) = \max\{m \mid (q, m) \in \text{dom}(\sigma)\}$ is called threshold for q . A σ -move from (q, n) is an acceptor move if $n = \theta(q)$, and a base move if it is not an acceptor move and $n = \text{ext}_{\mathcal{A}}(q)$. The reduced-memory strategy $\underline{\rho}$ induced by σ is obtained by removing any non-acceptor, non-base moves from σ and taking the carry-over memory-full strategy induced by the resulting skeleton.*

Thus, $\underline{\rho}$ -moves aimed at reaching accepting states are as prescribed by $\bar{\rho}$, however, moves aimed at increasing the available resources always use the moves prescribed by $\bar{\rho}$ when the *smallest* possible amount of resources (that is, $\text{ext}_{\mathcal{A}}(q)$) is available in q . The definition of $\underline{\rho}$ is inspired by the strategies of [3], which in non-accepting states combine an *attractor strategy* (here the acceptor moves) with a *good-for-energy strategy* (here the base moves). Our next result states that, if $\bar{\rho}$ can be used to exhibit an optimal run in \mathcal{A} from each $(q, \text{ext}_{\mathcal{A}}(q))$, then so can $\underline{\rho}$. We use an example to illustrate the proof idea.

Example 4.15. Example 4.11 describes a skeleton strategy σ which contains non-base, non-acceptor moves from states $(y, 2)$ and $(y, 4)$. Removing such moves from σ yields the following skeleton:



with moves labelled so as to indicate their type (acceptor/base). In state $(y, 1)$, a $\underline{\rho}$ -conform play proceeds like a $\bar{\rho}$ -conform play until $(y, 2)$ is reached. Since $\text{ext}_{\mathcal{A}}(y) < 2 < \theta(y)$, a $\underline{\rho}$ -conform play from $(y, 2)$ will repeat the base move for y followed by the acceptor move for y_1 , carrying over unused resources, until the acceptor move for y is enabled and $(x, 1)$ is reached:



To prove that $\underline{\rho}$ leads to an accepting state from any $(q, \text{ext}_{\mathcal{A}}(q))$, we remove non-base, non-acceptor moves from σ one by one in a particular order, and show that each removal maintains the winning property of the induced $\bar{\rho}$. Specifically, move m from state q is only removed if a $\bar{\rho}$ -conform partial play starting in $(q, \text{ext}_{\mathcal{A}}(q))$, which only visits non-accepting states, will not contain non-base, non-acceptor moves prior to m . For example, the move from $(y, 4)$ is removed *after* the move from $(y, 2)$ is removed.

Theorem 4.16. *Let σ be a winning skeleton strategy for a WBAO \mathcal{A} , and let $\underline{\rho}$ be as in Definition 4.14. Then, any $\underline{\rho}$ -conform run from q with initial memory $\text{ext}_{\mathcal{A}}(q)$ is accepting and has value $\text{ext}_{\mathcal{A}}(q)$.*

Figure 2: Algorithm for computing attractor and base strategies for resource-aware WBAOs

```

Input: resource-aware WBAO  $\mathcal{A} = (Q, \langle r, \gamma \rangle, \Omega)$ 
Outputs: strategy  $\sigma : Q_2 \rightarrow Q \cup \{\perp\}$ ; threshold  $\theta : Q_1 \rightarrow \mathbb{N} \cup \{\perp\}$ ;
            attractor strategy  $\sigma_a : Q_1 \rightarrow Q \cup \{\perp\}$ ; good-for-energy strategy  $\sigma_e : Q_1 \rightarrow Q \cup \{\perp\}$ 
1. let  $e := \text{Extent}(2)$ 
2. for  $q \in Q_2$  do
3.   let  $\sigma(q) := \mathbf{any}(\lambda, q')$  s.t.  $(\gamma(q)(\iota_\lambda(q')) \neq \infty) \wedge (e(q) = (\gamma(q)(\iota_\lambda(q')) + e(q')) \ominus r(q))$ 
4. endfor
5. for  $q \in Q_1$  do
6.   let  $e(q) := \infty$ ;
7.   let  $\sigma_a(q) := \perp$ 
8. endfor
9. repeat
10.  let  $old := e$ 
11.  for  $q \in Q_1$  do
12.    let  $e(q) := \min\{(\gamma(q)(\iota_\lambda(q')) + old(q')) \ominus r(q) \mid \gamma(q)(\iota_\lambda(q')) \neq \infty\}$ 
13.    let  $q_1 := \mathbf{any} q'$  s.t.  $(\gamma(q)(\iota_\lambda(q')) \neq \infty) \wedge e(q) = (\gamma(q)(\iota_\lambda(q')) + old(q')) \ominus r(q)$ 
14.    if  $e(q) \neq old(q)$  then
15.      if  $\sigma_a(q) = \perp$  then
16.        let  $\sigma_a(q) := q_1$  /* set attractor strategy */
17.        let  $\theta(q) := e(q)$ 
18.      else
19.        let  $\sigma_e(q) := q_1$  /* set good-for-energy strategy */
20.      endif
21.    endif
22.  endfor
23. until  $e = old$ 

```

Theorem 4.16 applies to the strategy computed by the algorithm outlined in Section 4.2. In fact, this algorithm can be adapted to *only* compute the resulting attractor and good-for-energy strategies – see Figure 2 (where for succinctness we further assume that $\text{ar}(\lambda) = 1$ for $\lambda \in \Lambda$). The correctness of the algorithm follows from Theorems 4.13 and 4.16. We now analyse the complexity of computing the attractor strategy σ_a and the good-for-energy strategy σ_e , assuming extents are already computed. If $N_i = |Q_i|$ and E_i is the number of edges from states in Q_i , for $i \in \{1, 2\}$, then (i) steps 2–8 take time $O(N_2 + E_2 + N_1)$, while (ii) steps 10–22 take time $O(N_1 + E_1)$ and are repeated at most $N_1 \times B$ times (in the worst case, each iteration changes only one $e(q)$). This gives an upper bound of $O((N_1 + E_1) \times N_1 \times B + N_2 + E_2)$ for the overall time complexity. While this is quadratic in the size of the automaton, the constant B should not be ignored.

For our results to extend to *tree* automata, Definition 4.12 must redistribute unused resources following a transition $q \rightarrow (\lambda, q_1, \dots, q_n)$, so that an amount > 0 is assigned to each q_i . (Without this, removing non-base, non-acceptor moves from σ may lead to some branches of the resulting run never reaching accept states.) This can be done by moving to the (bounded) tropical rationals semiring, but *only* once σ is computed. This allows even a resource of 1 to be redistributed among *several* (finitely many) successors, ensuring that when a sequence of base/acceptor moves is repeated with *more* resources, and this leads (on a branch) to a state seen before where the base move applies again, the resources have strictly increased,

and a repetition of the same moves will again increase resources by at least the same amount. Then, the proof of Theorem 4.16 generalises: the only difference is that instead of a single play from $(y, \text{ext}_{\mathcal{A}}(y))$ one must consider all plays arising from \forall choices. The complexity in the case of tree automata remains similar, provided that E_1 now counts an "edge" $q \rightarrow (\lambda, q_1, \dots, q_n)$ a number of times equal to $\text{ar}(\lambda)$.

An extension to *parity* automata can be obtained by observing that optimal accepting runs of a special form always exist in \mathcal{A} . Such runs can be described using (i) a *lasso*, that is, a partial run of $\mathcal{R}_{\mathcal{A}}$ with root $(q_0, \text{ext}_{\mathcal{A}}(q_0))$ and leaves of the form $(q, \text{ext}_{\mathcal{A}}(q))$, and (ii) a *loop*, that is, a winning strategy σ in a *sub-model* of $\mathcal{R}_{\mathcal{A}}$ that contains all the leaves in the lasso, with the additional property that a σ -conform play will *never* visit an odd-parity state without later visiting a higher, even parity state. The lasso leaves thus identify points in the run from which states with odd parity larger than the highest infinitely occurring even parity on *any* given branch do not occur. Once extents are known, a *maximal* loop (i.e. one which includes as many states of the form $(q, \text{ext}_{\mathcal{A}}(q))$ as possible) can be computed as a greatest fixpoint, whereas a lasso from each of the remaining states of the form $(q, \text{ext}_{\mathcal{A}}(q))$, with leaves belonging to the loop, can be computed as a least fixpoint. For the maximal loop, the associated strategy is computed similarly to the Büchi case, additionally recording a lower bound on the even parities guaranteed to be seen in the future on *any* branch of the computation. However, this time several iterations are needed to successively discard states of the form $(q, \text{ext}_{\mathcal{A}}(q))$ that cannot be part of a loop. For lassos, a backwards computation suffices. As in the Büchi case, the resulting strategies (either lasso or loop for any given state) can be simplified to consist of two moves only. The details of this extension, including a complexity analysis, are left for future work.

5 Weighted Parity Games with Offsetting

We now define a game version of WPAOs, along with a suitable notion of extent. The definition below is parameterised by a semiring $(S, +, 0, \bullet, 1)$ subject to our earlier assumptions.

Definition 5.1. A weighted parity game with offsetting (WPGO) is given by a $S \times (\mathcal{P}_{\omega} \circ T_S \circ F)$ -coalgebra $(Q, \langle r, \gamma \rangle)$ together with a parity map $\Omega : Q \rightarrow \mathbb{N}$ with finite range. If $\text{ran}(\Omega) = \{1, 2\}$, we call $\mathcal{G} = (Q, \langle r, \gamma \rangle, \Omega)$ a weighted Büchi game with offsetting (WBGO).

Thus, WPAOs are special cases of WPGOs, with each $\gamma(q)$ a singleton. The presence of the finite powerset functor $\mathcal{P}_{\omega} : \text{Set} \rightarrow \text{Set}$ gives \forall additional power to increase the amount of initial resources required, by making a choice (in $\gamma(q)$) at each step of the computation.

Remark 5.2. While standard weighted parity games as used e.g. in [3, 9, 13] make no restrictions on the alternation of \exists and \forall moves, nor on the weights on \forall moves or the parities of \forall states, any weighted parity game which does not satisfy our constraints can be transformed into one that does, by adding intermediary states/transitions that offer no real choice / carry trivial weights.

Example 5.3. When $S = (\{0, 1\}, \vee, 0, \wedge, 1)$ (so $T_S = \mathcal{P}_{\omega}$) and $F = \text{Id}$, ignoring offsets yields a (special type of) standard parity game, with \forall states given by elements of Q and inheriting parities from Q , \exists states given by elements of $\gamma(q)$ ($q \in Q$) and having minimum parity, and \forall and \exists moves alternating.

Definition 5.4 (Extent of WPGO). Let $\mathcal{G} = (Q, \langle r, \gamma \rangle, \Omega)$ be a WPGO with $\text{ran}(\Omega) = \{1, \dots, n\}$. For $k \in \text{ran}(\Omega)$, let $\gamma_k : Q_k \rightarrow \mathcal{P}_{\omega} T_S F Q$ and $r_k : Q_k \rightarrow S$ denote the restrictions of γ and respectively r to $Q_k := \{q \in Q \mid \Omega(q) = k\}$. The extent $\text{ext}_{\mathcal{G}} = [e_1, \dots, e_n] : Q \rightarrow S$ of \mathcal{G} is the solution of the following nested equational system, with the most significant equation being the last one:

$$\begin{cases} u_1 =_{\mu} & ((\mu_1 \circ T_S(\bullet_F) \circ T_S F[u_1, \dots, u_n])^{\#} \circ \gamma_1) \otimes r_1 \\ u_2 =_{\nu} & ((\mu_1 \circ T_S(\bullet_F) \circ T_S F[u_1, \dots, u_n])^{\#} \circ \gamma_2) \otimes r_2 \\ & \vdots \\ u_n =_{\eta} & ((\mu_1 \circ T_S(\bullet_F) \circ T_S F[u_1, \dots, u_n])^{\#} \circ \gamma_n) \otimes r_n \end{cases} \quad (2)$$

with $\eta = \mu$ ($\eta = \nu$) if n is odd (resp. even), with variables u_k ranging over the poset $(S^{\mathcal{Q}^k}, \sqsubseteq)$ (and therefore $[u_1, \dots, u_n] : \mathcal{Q} \rightarrow S$), and with the first operands in the rhs pictured below:

$$\begin{aligned} \mathcal{Q}_k &\xrightarrow{\gamma_k} \mathcal{P}_\omega \mathbb{T}_S F \mathcal{Q} \xrightarrow{(\mathbb{T}_S F[u_1, \dots, u_n]; \mathbb{T}_S(\bullet_F); \mu_1)^\sharp} S \\ \mathbb{T}_S F \mathcal{Q} &\xrightarrow{\mathbb{T}_S F[u_1, \dots, u_n]} \mathbb{T}_S F S \xrightarrow{\mathbb{T}_S(\bullet_F)} \mathbb{T}_S S = \mathbb{T}_S^2 1 \xrightarrow{\mu_1} \mathbb{T}_S 1 = S \end{aligned}$$


Here, for $f : X \rightarrow S$, $f^\sharp : \mathcal{P}_\omega X \rightarrow S$ takes $Y \subseteq X$ to $\inf_{y \in Y} f(y)$.

Thus, the only difference w.r.t. WPAOs is the use of infima w.r.t. \sqsubseteq to account for the worst \forall move.

Example 5.5. For $S = (\{0, 1\}, \vee, 0, \wedge, 1)$, $F = \text{Id}$ and no offsetting, a state has extent 1 iff \exists has a winning strategy from it in the associated (standard) parity game. For $S = (\mathbb{N}_B^\infty, \min, \infty, +, 0)$ and $F = \text{Id}$, the extent gives the minimum resources required to exhibit an accepting run, irrespective of how \forall moves.

Remark 5.6. The energy parity games of [3], which use integer weights and no offsetting, can be modelled as WPGOs over $(\mathbb{N}^\infty, \min, \infty, +, 0)$ with $F = \text{Id}$: transitions with negative weights become positively weighted, while positive weights are turned into offsets by suitably duplicating their target states along with any outgoing transitions from such states. Then, the *initial credit problem* for energy parity games [3] (of computing the minimum initial credit required for an accepting run whose energy level never drops below 0) is equivalent to the problem of computing the extents of states in the associated WPGO.

Example 5.7. Consider the WBGO on the left ($F = \text{Id}$), where unlabelled transitions correspond to \forall moves leading to different elements of $\mathbb{T}_S \mathcal{Q}$ (namely f and g). The state space is $\mathcal{Q} = \{x, y_1, y_2\}$ and the offset function maps x to 0, y_1 to 1 and y_2 to 4. In states x and y_1 , \forall can move to either f or g . In state f , \exists can move to y_1 or y_2 (with the latter move increasing the available resources in f) or to the accepting state x . The associated equational system (given on the right) has solution $(2, 1, 0, 0, 0)$. Intuitively, this is because by starting in state x with 2 initial resources, in the worst case \forall moves to state f , where the only \exists move that increases available resources is to y_2 . Thereon \exists can control whether to further increase the available resources (by moving to y_2 again) or move to the accept state (when 6 resources are available in f). On the other hand, 1 initial resource is required in y_1 : together with the offset of 1, this allows a move to y_2 when \forall moves to f , or a move to x that leaves 2 available resources, when \forall moves to g .



Remark 5.8. Notions of *run* and *value* of a run can be defined similarly to WPAOs: a run selects a *single* \forall move in states $q \in \mathcal{Q}$, and explores *all* \forall moves in states $(\lambda, q_1, \dots, q_n)$. However, Proposition 3.8 does not generalise, since \forall moves in a *specific* play may require fewer resources than specified by the extents.

6 Strategies in Resource-Aware Games

We now fix $S = (\mathbb{N}^\infty, \min, \infty, +, 0)$ and show how the results in Section 4 extend to *resource-aware WBGOs*: we define a resource game $\mathcal{R}_{\mathcal{G}}$ akin to the resource game $\mathcal{R}_{\mathcal{A}}$ and show that memoryless winning strategies for \exists in $\mathcal{R}_{\mathcal{G}}$ yield strategies in \mathcal{G} for exhibiting resource-aware accepting runs with minimum initial resources (where this time, \forall controls some of the choices made in defining the run).

Definition 6.1 (Resource Game). *The resource game $\mathcal{R}_{\mathcal{G}}$ associated to a WPGO $\mathcal{G} = (\mathcal{Q}, \langle r, \gamma \rangle, \Omega)$ is a standard parity game with states given by $\mathcal{Q}_\exists = \{(Y, s + r(q)) \in \mathbb{T}_S F \mathcal{Q} \times \mathbb{N}_B \mid Y \in \gamma(q) \text{ for some } q \in \mathcal{Q} \text{ with } s \geq \text{ext}_{\mathcal{G}}(q)\}$ and $\mathcal{Q}_\forall = \{(q, s) \in \mathcal{Q} \times \mathbb{N}_B \mid s \geq \text{ext}_{\mathcal{G}}(q)\} \cup F\{(q, s) \in \mathcal{Q} \times \mathbb{N}_B \mid s \geq \text{ext}_{\mathcal{G}}(q)\}$,*

parities inherited from \mathcal{G} on \forall states of the form (q, s) with $s \geq \text{ext}_{\mathcal{G}}(q)$, and equal to $\min\{\Omega(q) \mid q \in Q\}$ on all other states, and moves given by:


$$\begin{aligned} (q, s) &\xrightarrow{\forall} (Y, s + r(q)) && \text{whenever } Y \in \gamma(q) \\ (Y, s) &\xrightarrow{\exists} (\lambda, (q_1, s_1), \dots, (q_n, s_n)) && \text{whenever } s \geq Y(t_\lambda(q_1, \dots, q_n)) + s_1 + \dots + s_n, \\ (\lambda, (q_1, s_1), \dots, (q_n, s_n)) &\xrightarrow{\forall} (q_i, s_i) && \text{for all } i \in \{1, \dots, n\}. \end{aligned}$$

The difference w.r.t. Definition 4.2 is the addition of a new type of \forall move, which is also where the offset values are added to the available resources.

Definition 6.2. A resource-aware play of a WPGO \mathcal{G} is a play of $\mathcal{R}_{\mathcal{G}}$, viewed as a standard parity game. Such a play cycles through: \forall making a choice from those specified by γ ; \exists choosing a specific transition out of the available ones; and in case of tree-like structure, \forall choosing one particular branch resulting from the \exists move. Then any memoryless strategy σ for \exists in $\mathcal{R}_{\mathcal{G}}$ yields a memory-full strategy ρ for \exists in \mathcal{G} , defined as in Definition 4.9. Moreover, Theorem 4.10 generalises to WBGOs.

Theorem 6.3. Let $\mathcal{G} = (Q, \langle r, \gamma \rangle, \Omega)$ be a WPGO with associated resource game $\mathcal{R}_{\mathcal{G}}$. If $\sigma : Q_{\exists} \rightarrow Q_{\forall}$ is winning for \exists in $\mathcal{R}_{\mathcal{G}}$, the induced strategy ρ can be used to exhibit a resource-aware accepting run of \mathcal{G} from $q \in Q$ with initial memory $\text{ext}_{\mathcal{G}}(q)$. (Recall from Remark 5.8 that such a run requires input from \forall .)

Example 6.4. A strategy σ for the game in Example 5.7 is shown below. Note that this only prescribes moves (solid arrows) in states f and g . Dashed arrows correspond to \forall moves, and the values shown for y_1 and y_2 correspond to the *worst* \forall move from each state; a different \forall move may result in more resources available in f/g than shown in the target states (as is e.g. the case for $(y_1, 1) \rightarrow (g, 1)$).



Reduced-memory strategies for \exists can be defined for WPGOs just as in the automata case; they arise from *skeleton strategies*, also defined as before. For example, the reduced memory strategy induced by σ of Example 6.4 involves discarding the move from $(f, 4)$. Moreover, Theorem 4.16 generalises to games.

Theorem 6.5. Let σ be a winning skeleton strategy for a WBG O \mathcal{G} , inducing a reduced-memory strategy ρ . Then ρ is winning in \mathcal{G} (in the sense of Theorem 6.3) from any $q \in Q$ with initial memory $\text{ext}_{\mathcal{G}}(q)$.

The algorithm in Figure 1 generalises straightforwardly to WPGOs: the computation of the extent values $e(q)$ with $q \in Q$ (line 9) additionally quantifies over *all* \forall moves in q , taking the supremum (w.r.t. \leq) of the resulting values. With this change, the algorithm in Figure 2 generalises to WBG O s, updating either $\sigma_a(q)$ or $\sigma_e(q)$ whenever $e(q)$ is updated. The complexity of computing both extents and reduced-memory strategies is similar to WBAOs; the only difference is that E_i now counts all ways to reach a state in Q from a state in Q_i in one step (one application of γ) via an intermediary \forall -move, for $i \in \{1, 2\}$.

7 Conclusions and Future Work

We studied the optimal strategy synthesis problem for a notion of resource-aware parity automaton/game that models both resource gain and resource usage, and showed that memory-full strategies consisting of two moves only per state suffice. We focused on the Büchi case, leaving the extension to parity automata/games for future work. While our results are similar to existing ones (e.g. [3]), they are applicable to a wider class of qualitative behaviours which includes not only words but also trees. This is made possible by a coalgebraic treatment of quantitative automata/games, which also opens the road for further extensions of our approach. Future work will study more complex synthesis problems, applicable to systems with a variety of quantitative features (e.g. the synthesis of a resource-aware component operating in a probabilistic environment, with guarantees on the likelihood with which available resources suffice).

References

- [1] A. Arnold & D. Niwiński (2001): *Rudiments of μ -Calculus*. Studies in Logic and the Foundations of Mathematics, North-Holland.
- [2] P. Bouyer, U. Fahrenberg, K. G. Larsen, N. Markey & J Srba (2008): *Infinite Runs in Weighted Timed Automata with Energy Constraints*. In F. Cassez & C. Jard, editors: *Formal Modelling and Analysis of Timed Systems, Lecture Notes in Computer Science* 5215, Springer, pp. 33–47, doi:10.1007/978-3-540-85778-5_4.
- [3] K. Chatterjee & L. Doyen (2012): *Energy parity games*. *Theor. Comput. Sci.* 458, pp. 49 – 60, doi:10.1016/j.tcs.2012.07.038.
- [4] K. Chatterjee, T. A. Henzinger & M. Jurdzinski (2005): *Mean-payoff parity games*. In: *Proceedings - Symposium on Logic in Computer Science*, 20, pp. 178 – 187, doi:10.1109/LICS.2005.26.
- [5] C. Cîrstea: *Linear Time Logics – A Coalgebraic Perspective*. ArXiv:1612.07844.
- [6] C. Cîrstea (2014): *A Coalgebraic Approach to Linear-Time Logics*. In A. Muscholl, editor: *Foundations of Software Science and Computation Structures - 17th International Conference, FOSSACS 2014, Lecture Notes in Computer Science* 8412, Springer, pp. 426–440, doi:10.1007/978-3-642-54830-7_28.
- [7] C. Cîrstea (2017): *From Branching to Linear Time, Coalgebraically*. *Fundam. Inform.* 150(3-4), pp. 379–406, doi:10.3233/FI-2017-1474.
- [8] C. Cîrstea, S. Shimizu & I. Hasuo (2017): *Parity Automata for Quantitative Linear Time Logics*. In F. Bonchi & B. König, editors: *CALCO'17, Leibniz International Proceedings in Informatics (LIPIcs)* 72, Schloss Dagstuhl–Leibniz-Zentrum für Informatik, pp. 1–18, doi:10.4230/LIPIcs.CALCO.2017.7.
- [9] N. Fijalkow & M. Zimmermann (2014): *Parity and Streett Games with Costs*. *Logical Methods in Computer Science* 10, pp. 1–29, doi:10.2168/LMCS-10(2:14)2014.
- [10] E. Grädel, W. Thomas & T. Wilke, editors (2002): *Automata, Logics, and Infinite Games: A Guide to Current Research. Lecture Notes in Computer Science* 2500, Springer, doi:10.1007/3-540-36387-4.
- [11] B. Jacobs (2016): *Introduction to Coalgebra. Towards Mathematics of States and Observations*. Cambridge University Press, doi:10.1017/CBO9781316823187.
- [12] M. Jurdziński (2000): *Small Progress Measures for Solving Parity Games*. In H. Reichel & S. Tison, editors: *STACS 2000, Lecture Notes in Computer Science* 1770, Springer, pp. 290–301, doi:10.1007/3-540-46541-3_24.
- [13] S. Schewe, A. Weinert & M. Zimmermann (2018): *Parity Games with Weights*. In D. Ghica & A. Jung, editors: *27th EACSL Annual Conference on Computer Science Logic (CSL 2018), Leibniz International Proceedings in Informatics (LIPIcs)* 119, Schloss Dagstuhl–Leibniz-Zentrum für Informatik, pp. 36:1–36:17, doi:10.4230/LIPIcs.CSL.2018.36.