

The Queue Automaton Revisited

Jos C.M. Baeten¹ and Bas Luttik²

¹ CWI, Amsterdam, The Netherlands

² Eindhoven University of Technology, Eindhoven, The Netherlands

Abstract. We consider the computational model of the Queue Automaton. An old result is that the deterministic queue automaton is equally expressive as the Turing machine. In a previous paper, we introduced the Reactive Turing Machine, enhancing the Turing machine with a notion of interaction. The Reactive Turing Machine defines all executable processes. In this paper, we prove that the non-deterministic queue automaton is equally expressive as the Reactive Turing Machine. Together with finite automata, pushdown automata and parallel pushdown automata, queue automata form a nice hierarchy of executable processes, with stacks, bags and queues as central elements.

1 Introduction

Replacing, in a pushdown automaton, the (last-in first-out) stack memory by a (first-in first-out) queue memory yields the computational model of the queue automaton. This computational model, sometimes also called a Post machine or a pullup automaton, has not raised a lot of attention in the literature, but it is a known result that the deterministic queue automaton is equally expressive as the Turing machine of [24], so that it defines all computable languages and all computable functions. Implicitly, this result is already mentioned in [21], and further given in [25,18].

In this paper, we investigate the (non-deterministic) queue automaton. We do not define the language of a queue automaton directly, but instead, we define its process graph or transition system. A state of this process graph is given by the state of the queue automaton together with the contents of the queue and a transition is given by the label of the transition of the queue automaton. By considering the language equivalence class of the process graph, we obtain again the language, but we can also divide out other equivalence relations. Notable among these is branching bisimilarity (see [14]). By dividing out branching bisimilarity, we obtain the *process* of the queue automaton, incorporating a notion of interaction or communication. These notions of process and communication come from process theory or concurrency theory [20,3]. We prove that several variants of the queue automaton yield the same set of languages and the same set of processes, and prove that a queue automaton with two queues also yields the same set of languages and the same set of processes.

Also the Turing machine can be extended to incorporate processes and communication, but this is not so straightforward as we just sketched for the queue automaton. We achieved this, nonetheless, in [8], where we introduced the Reactive Turing Machine. Whereas the classical Turing machine defines the class of computable languages and computable functions, the reactive Turing machine also yields a process graph that can

be used to define the class of executable processes. In this paper, we prove that the non-deterministic queue automaton is equally expressive as the reactive Turing machine. This shows again that the notion of executability we introduced is robust: it is given by different computational models and also by the π -calculus, see [17].

The queue automaton has certain advantages over the reactive Turing machine: it is mathematically simpler, the extension with interaction is easier, and we get a better process hierarchy, as we explain now. In the queue automaton, we can make the interaction between the finite control and the queue memory explicit, by proving that every executable process is branching bisimilar to a regular process communicating with a queue. In earlier papers [7,1], we proved that every pushdown process is branching bisimilar to a regular process communicating with a stack, and every parallel pushdown process is branching bisimilar to a regular process communicating with a bag. Thus, the queue, stack and bag are the central elements in this Chomsky-Turing hierarchy.

This paper contributes to our ongoing project to integrate the theory of automata and formal languages on the one hand and concurrency theory on the other hand. The integration requires a more refined view on the semantics of automata, grammars and expressions. Instead of treating automata as language acceptors, and grammars and expressions as syntactic means to specify languages, we propose to view them both as defining process graphs. The great benefit of this approach is that process graphs can be considered modulo a plethora of behavioural equivalences [12]. One can still consider language equivalence and recover the classical theory of automata and formal languages. But one can also consider finer notions such as bisimilarity, which is better suited for interacting processes.

Note 1. This paper is dedicated to Sjouke Mauw, our friend and colleague, on the occasion of his 65th birthday. Sjouke was one of the first to see that information security would become an important area, in practice and research, and that formal methods can play an important role in this area. Sjouke does not stop at borders in order to realise his ambition.

As we said, this paper is part of our ongoing project to integrate automata theory and concurrency theory. In this project, we looked at context-free processes in, e.g., [6,7]. Here, we considered the place of context-free processes in the Chomsky-Turing hierarchy. Sjouke already considered context-free processes in [19].

2 Preliminaries

As a common semantic framework we use the notion of a *labelled transition system*.

Definition 1 A labelled transition system is a quadruple $(S, \mathcal{A}, \rightarrow, \downarrow)$, where

1. S is a set of states;
2. \mathcal{A} is a set of actions, $\tau \notin \mathcal{A}$ is the unobservable or silent action;
3. $\rightarrow \subseteq S \times (\mathcal{A} \cup \{\tau\}) \times S$ is an $\mathcal{A} \cup \{\tau\}$ -labelled transition relation; and
4. $\downarrow \subseteq S$ is the set of final, accepting or terminating states.

A process graph is a labelled transition system with a special designated root state or initial state \uparrow , i.e., it is a quintuple $(\mathcal{S}, \mathcal{A}, \rightarrow, \uparrow, \downarrow)$ such that $(\mathcal{S}, \mathcal{A}, \rightarrow, \downarrow)$ is a labelled transition system, and $\uparrow \in \mathcal{S}$. We write $s \xrightarrow{a} s'$ for $(s, a, s') \in \rightarrow$ and $s \downarrow$ for $s \in \downarrow$. We write \mathcal{A}_τ for $\mathcal{A} \cup \{\tau\}$.

For $w \in \mathcal{A}^*$ we define $s \xrightarrow{w} t$ inductively, for all states s, t, u : first, $s \xrightarrow{\varepsilon} s$, and then, for $a \in \mathcal{A}$, if $s \xrightarrow{a} t$ and $t \xrightarrow{w} u$, then $s \xrightarrow{aw} u$, and if $s \xrightarrow{\tau} t$ and $t \xrightarrow{w} u$, then $s \xrightarrow{w} u$. We see that τ -steps do not contribute to the string w . We write $s \rightarrow t$ for there exists $a \in \mathcal{A}_\tau$ such that $s \xrightarrow{a} t$. Similarly, we write $s \twoheadrightarrow t$ for “there exists $w \in \mathcal{A}^*$ such that $s \xrightarrow{w} t$ ” and say that t is *reachable* from s . Finally, we write $s \xrightarrow{(a)} t$ for “ $s \xrightarrow{a} t$ or $a = \tau$ and $s = t$ ”.

By considering language equivalence classes of process graphs, we recover language equivalence as a semantics, but we can also consider other equivalence relations. Notable among these is *bisimilarity*.

Definition 2 Let $(\mathcal{S}, \mathcal{A}, \rightarrow, \downarrow)$ be a labelled transition system. A symmetric binary relation R on \mathcal{S} is a strong bisimulation if it satisfies the following conditions for every $s, t \in \mathcal{S}$ such that $s R t$ and for all $a \in \mathcal{A}_\tau$:

1. if $s \xrightarrow{a} s'$ for some $s' \in \mathcal{S}$, then there is a $t' \in \mathcal{S}$ such that $t \xrightarrow{a} t'$ and $s' R t'$; and
2. if $s \downarrow$, then $t \downarrow$.

If there is a strong bisimulation relating s and t we write $s \dot{\leftrightarrow} t$.

Note that a non-symmetric relation is never a strong bisimulation.

Sometimes we can use the *strong* version of bisimilarity defined above, which does not give special treatment to τ -labelled transitions. In general, when we do give special treatment to τ -labelled transitions, we use (some form of) *branching bisimulation* [14].

Definition 3 Let $(\mathcal{S}, \mathcal{A}, \rightarrow, \downarrow)$ be a labelled transition system. A symmetric binary relation R on \mathcal{S} is a branching bisimulation if it satisfies the following conditions for every $s, t \in \mathcal{S}$ such that $s R t$ and for all $a \in \mathcal{A}_\tau$:

1. if $s \xrightarrow{a} s'$ for some $s' \in \mathcal{S}$, then there are states $t', t'' \in \mathcal{S}$ such that $t \xrightarrow{\varepsilon} t' \xrightarrow{(a)} t''$, $s R t''$ and $s' R t'$; and
2. if $s \downarrow$, then there is a state $t' \in \mathcal{S}$ such that $t \xrightarrow{\varepsilon} t'$ and $t' \downarrow$.

If there is a branching bisimulation relating s and t , we write $s \dot{\leftrightarrow}_b t$. If $s \xrightarrow{\tau} t$ and $s \dot{\leftrightarrow}_b t$, we say this τ -step is *inert*.

Again, a non-symmetric relation is never a branching bisimulation.

Theorem 1 Strong bisimilarity and branching bisimilarity are equivalence relations on labeled transition systems.

Proof. See [9] and [13].

Now suppose we are given a labelled transition system and two states s, t in this labelled transition system. These states give rise to two process graphs with the given states as initial states. We say the two process graphs are strongly bisimilar or branching bisimilar if there is a strong bisimulation or a branching bisimulation on the labelled transition system that relates s and t .

A *process* is a branching bisimilarity equivalence class of process graphs. We say a process is *regular* if its branching bisimilarity equivalence class contains an element with finitely many states and finitely many transitions.

Finally, we define when a labelled transition system is *deterministic*.

Definition 4 A labelled transition system $(\mathcal{S}, \mathcal{A}, \rightarrow, \downarrow)$ is deterministic iff for all $s \in \mathcal{S}$ and for all $a \in \mathcal{A}_\tau$ there is at most one $t \in \mathcal{S}$ with $s \xrightarrow{a} t$. Moreover, whenever $s \xrightarrow{\tau} t$, there is no $a \in \mathcal{A}$ and $u \in \mathcal{S}$ with $s \xrightarrow{a} u$.

A process graph $(\mathcal{S}, \mathcal{A}, \rightarrow, \uparrow, \downarrow)$ is deterministic if the labelled transition system $(\mathcal{S}, \mathcal{A}, \rightarrow, \downarrow)$ is deterministic.

3 Queue Automata

We define queue automata, prove that their type of transitions can be restricted without losing expressiveness, and prove that having two queues instead of one does not increase their expressiveness.

Definition 5 (Queue automaton) A queue automaton Q is a sextuple $(\mathcal{S}, \mathcal{A}, \mathcal{D}, \rightarrow, \uparrow, \downarrow)$ where:

1. \mathcal{S} is a non-empty finite set of states,
2. \mathcal{A} is a non-empty finite action alphabet,
3. \mathcal{D} is a non-empty finite data alphabet,
4. $\rightarrow \subseteq \mathcal{S} \times \mathcal{A}_\tau \times (\mathcal{D} \cup \{\varepsilon, *\}) \times \mathcal{D}^* \times \mathcal{S}$ is a finite set of transitions or steps,
5. $\uparrow \in \mathcal{S}$ is the root state or initial state,
6. $\downarrow \subseteq \mathcal{S}$ is the set of final, accepting or terminating states.

If $(s, a, d, \delta, t) \in \rightarrow$ with $d \in \mathcal{D}$, we write $s \xrightarrow{a[d/\delta]} t$, and this means that the machine, when it is in state s and d is the head element of the queue, can execute action a , dequeue this d and enqueue the string δ and thereby move to state t . Likewise, writing $s \xrightarrow{a[\varepsilon/\delta]} t$ means that the machine, when it is in state s and the queue is empty, can execute action a , enqueue the string δ and thereby move to state t . Writing $s \xrightarrow{a[*/\delta]} t$ means that the machine, when it is in state s , can execute action a , enqueue δ and thereby move to state t , irrespective of the contents of the queue and without dequeuing anything. In steps $s \xrightarrow{\tau[d/\delta]} t$, $s \xrightarrow{\tau[\varepsilon/\delta]} t$ and $s \xrightarrow{\tau[*/\delta]} t$, no action is executed, only the queue is modified. The semantic interpretation of the elements of the action alphabet can be left unspecified, but often they stand for some kind of interaction with the automaton.

In definitions of queue automata appearing in the literature, the set of transitions is sometimes defined a little differently, but all of them are equally expressive, in the sense

that they all give rise to the same set of languages and the same set of processes. For instance, in [16], the $a[*/\delta]$ -labeled transitions do not occur, but there are two variants of the $a[d/\delta]$ -labeled transitions: one where d is dequeued, and one where d is not dequeued. Moreover, instead of using general sequences δ , only singleton sequences or empty sequences are used. We show some of these expressiveness results in the sequel.

We believe the present definition fits in best with the definitions of a pushdown automaton (in [7]) and parallel pushdown automaton (in [1]).

The notion of queue automaton is very similar to the classical notion of a pushdown automaton (see, e.g., [15]), only there, a stack is used instead of a queue, and in a stack, push and pop (as enqueue and dequeue are called) occur according to the last-in first-out principle. Therefore, in a pushdown automaton, we do not need the separate $s \xrightarrow{a[*/\delta]} t$ transitions, as they can be replaced by a $s \xrightarrow{a[\varepsilon/\delta]} t$ transition in combination with $s \xrightarrow{a[d/\delta d]} t$ transitions for all $d \in \mathcal{D}$. Further on, we will see that the expressivity of the queue automaton does not change when we omit these transitions, using a recycle operation and working modulo branching bisimulation.

Now we could proceed to define the language of a queue automaton, but instead, we take an intermediate step and first define the process graph of a queue automaton. By considering the language of this process graph, we find the language of the queue automaton again, and by considering its branching bisimulation equivalence class, we obtain the process of the queue automaton.

Definition 6 Let $Q = (\mathcal{S}, \mathcal{A}, \mathcal{D}, \rightarrow, \uparrow, \downarrow)$ be a queue automaton. The process graph of Q is defined as follows, for all $\delta \in \mathcal{D}^*$:

1. the set of states is $\{(s, \delta) \mid s \in \mathcal{S}, \delta \in \mathcal{D}^*\}$;
2. the set of actions is \mathcal{A} ;
3. the transition relation is generated by the following clauses:
 - if $s \xrightarrow{a[\varepsilon/\delta]} t$ then $(s, \varepsilon) \xrightarrow{a} (t, \delta)$;
 - if $s \xrightarrow{a[d/\delta]} t$ then $(s, \zeta d) \xrightarrow{a} (t, \delta \zeta)$, for all $\zeta \in \mathcal{D}^*$;
 - if $s \xrightarrow{a[*/\delta]} t$ then $(s, \zeta) \xrightarrow{a} (t, \delta \zeta)$, for all $\zeta \in \mathcal{D}^*$;
4. the initial state is (\uparrow, ε) ; and
5. $(s, \delta) \downarrow$ if $s \downarrow$.

Usually, we consider only those states (s, δ) that are reachable from the initial state. According to this definition, a state (s, δ) can be final also when the queue contents δ is non-empty. Defining that only states (s, ε) can be final is more limiting, and yields a smaller set of processes that are the process of a queue automaton. We can still code in a queue automaton that only states of the form (s, ε) can be final, by only allowing to enter such a state by means of a transition labeled by $a[\varepsilon/\varepsilon]$. This is illustrated in Examples 1 and 3. For an extensive treatment of termination conditions for pushdown automata and parallel pushdown automata, see [23].

Definition 7 Let $Q = (\mathcal{S}, \mathcal{A}, \mathcal{D}, \rightarrow, \uparrow, \downarrow)$ be a queue automaton. The language accepted by Q , $\mathcal{L}(Q)$, is the language of its transition system, i.e.

$$\mathcal{L}(Q) = \{w \in \mathcal{A}^* \mid \exists s \in \mathcal{S} \quad \exists \delta \in \mathcal{D}^* \text{ such that } s \downarrow \text{ and } (\uparrow, \varepsilon) \xrightarrow{w} (s, \delta)\}.$$

The process of Q , $\mathcal{P}(Q)$, is the branching bisimulation equivalence class of its transition system, often represented by its minimal element (identifying all branching bisimilar states).

Example 1 Let us construct a queue automaton for the language $\{ww \mid w \in \{a, b\}^*\}$. This language is not a pushdown language and also not a parallel pushdown language. We use a, b also as data symbols, so $\mathcal{D} = \{a, b\}$. In the initial state, a string can be read and enqueued. At some point (non-deterministically) it will switch to dequeue elements of the string again by moving to the second state, using the same order (the queue is first-in-first-out). Acceptance or termination takes place when the queue is empty again. See Figure 1. In the figure, we represent the states by circles, the initial state by a small incoming arrow and a final state by a double circle. An arrow that is labeled with multiple labels means that there is such a transition for each of these labels.




Fig. 1. Queue automaton for the language $\{ww \mid w \in \{a, b\}^*\}$.

Example 2 Figure 2 shows a queue automaton for the language $\{a^n b^n c^n \mid n > 0\}$; it uses data symbols 1, 2. The language is not a pushdown language and also not a parallel pushdown language.




Fig. 2. Queue automaton for the language $\{a^n b^n c^n \mid n > 0\}$.

Example 3 We can define the behaviour of the (first-in first-out) queue itself. Given a finite data set \mathcal{D} , the queue process can execute the following actions:

- $i?d$, **enqueue** data element d (input at port i);
- $o!d$, **dequeue** data element d , if this is the element at the head of the queue (output at port o);
- $o!\varepsilon$, show that the queue is empty.

We consider two variants: either the queue can always terminate, or it can terminate only when empty. We can define queue automata for these two queues, see Figure 3: for all $d \in \mathcal{D}$, there are the edges shown. The queue automaton on the left has only one state and can always terminate, irrespective of the contents of the queue the behaviour of which it represents. The queue automaton on the right has an additional τ -transition, that can only be executed when the queue the behaviour of which it represents has become empty: for this queue automaton, termination can only take place when the queue is empty.




Fig. 3. Queue automata of the queue.

We stated before that the $a[* / \delta]$ -labeled transitions are not really needed in the definition of queue automata. We prove this in the following lemma.

Lemma 1 *Let $Q = (\mathcal{S}, \mathcal{A}, \mathcal{D}, \rightarrow, \uparrow, \downarrow)$ be a queue automaton. Then there is a queue automaton Q' that does not have any $a[* / \delta]$ -labeled transitions and has the same process as Q .*

Proof. We construct the queue automaton Q' as follows. The data set is \mathcal{D} plus an extra data element $\$ \notin \mathcal{D}$. The state set is \mathcal{S} plus fresh states s^* for each $s \in \mathcal{S}$. Q' has the same initial state and the same final states as Q , and the same $a[\varepsilon / \delta]$ -labeled and $a[d / \delta]$ -labeled transitions. Finally, whenever Q has a transition of the form $s \xrightarrow{a[* / \delta]} t$, Q' has a transition $s \xrightarrow{a[\varepsilon / \delta]} t$ and, for all $d \in \mathcal{D}$, transitions $s \xrightarrow{a[d / d\$]} s^* \xrightarrow{\tau[\$ / \delta]} t$ and $s^* \xrightarrow{\tau[d / d]} s^*$.

Now we show that the process graphs of Q and Q' are branching bisimilar. We start out from the identity relation on all common states. Note that, whenever $s \xrightarrow{a[* / \delta]} t$ in Q , then we have $(s, \varepsilon) \xrightarrow{a} (t, \delta)$ in the process graph of Q but also in the process graph

of Q' . For nonempty memory contents, say of the form ζd for some $\zeta \in \mathcal{D}^*$, $d \in \mathcal{D}$, we have $(s, \zeta d) \xrightarrow{a} (t, \delta \zeta d)$ in the process graph of Q , and $(s, \zeta d) \xrightarrow{a} (s^*, d\$ \zeta) \xrightarrow{\varepsilon} (s^*, \zeta d \$) \xrightarrow{\tau} (t, \delta \zeta d)$ in the process graph of Q' , so it is enough to relate s^* to t . Note that all of the added τ -steps are inert because whenever some state in the process graph has an outgoing τ -transition, then this τ -transition is the unique outgoing transition. Therefore, the source and target states of the τ -transition are branching bisimilar.

We use the $a[*/\delta]$ -labeled transitions, nonetheless, because they allow for concise descriptions of interesting processes such as the queue of Example 3. In the following proofs, it will be useful on occasion to restrict the transitions in a queue automaton to only singleton enqueues, and to separate enqueues and dequeues.

Definition 8 *A transition in a queue automaton is a singleton enqueue iff it has a label of the form $a[* / d]$ for some $d \in \mathcal{D}$; it is a separate dequeue iff it has a label of the form $a[\varepsilon / \varepsilon]$ or $a[d / \varepsilon]$ for some $d \in \mathcal{D}$.*

Notice that the queue automata of the queue of Example 3 have only singleton enqueues and separate dequeues.

Lemma 2 *Let $Q = (\mathcal{S}, \mathcal{A}, \mathcal{D}, \rightarrow, \uparrow, \downarrow)$ be a queue automaton. Then, there is a queue automaton Q' with only singleton enqueues and separate dequeues with the same process and language.*

Proof. Let N be the maximum length of a data sequence occurring in any transition of Q , and enumerate the transitions of Q in a sequence u_1, \dots, u_K . The queue automaton Q' uses an extra data element $\$ \notin \mathcal{D}$ and has the set of states of Q plus new states $\{s_{ik} \mid i \leq N, k \leq K\} \cup \{s_{1k}^d \mid d \in \mathcal{D}, k \leq K\}$. Q' has the same action set, initial state and final states as Q . The transitions of Q' are defined as follows:

1. For each transition in Q of the form $s \xrightarrow{a[* / \varepsilon]} t$, if this is transition u_k , then Q' has transitions $s \xrightarrow{a[* / \$]} s_{1k} \xrightarrow{\tau[\$ / \varepsilon]} t$, and transitions $s_{1k} \xrightarrow{\tau[d / \varepsilon]} s_{1k}^d \xrightarrow{\tau[* / d]} s_{1k}$ for all $d \in \mathcal{D}$;
2. For each transition in Q of the form $s \xrightarrow{a[* / \delta]} t$ with $\delta \neq \varepsilon$, if this is transition u_k and $\delta = d_1 \dots d_n$ with $n > 1$, then Q' has transitions $s \xrightarrow{a[* / d_n]} s_{1k} \xrightarrow{\tau[* / d_{n-1}]} \dots s_{(n-1)k} \xrightarrow{\tau[* / d_1]} t$;
3. For each transition in Q of the form $s \xrightarrow{a[\varepsilon / \delta]} t$ with $\delta \neq \varepsilon$, if this is transition u_k and $\delta = d_1 \dots d_n$ with $n \geq 1$, then Q' has transitions $s \xrightarrow{a[\varepsilon / \varepsilon]} s_{1k} \xrightarrow{\tau[* / d_n]} \dots s_{nk} \xrightarrow{\tau[* / d_1]} t$;
4. For each transition in Q of the form $s \xrightarrow{a[d / \delta]} t$ with $\delta \neq \varepsilon$, if this is transition u_k and $\delta = d_1 \dots d_n$ with $n \geq 1$, then Q' has transitions $s \xrightarrow{a[d / \varepsilon]} s_{1k} \xrightarrow{\tau[* / d_n]} \dots s_{nk} \xrightarrow{\tau[* / d_1]} t$.

In the process graph of Q' , all the added τ -steps are inert, so it is branching bisimilar to the process graph of Q .

We see that the notion of the queue automaton is quite robust, as different variants of the queue automaton yield the same set of processes and the same set of languages. As a final illustration of the robustness of the notion of a queue automaton, we show that we can code a memory of two queues into one queue. In order to show this, we first define what a queue automaton with two queues is, and then show that its behaviour can also be obtained by a queue automaton with one queue.

Definition 9 (Queue automaton with two queues) A queue automaton with two queues Q is a sextuple $(\mathcal{S}, \mathcal{A}, \mathcal{D}, \rightarrow, \uparrow, \downarrow)$ that is just like a queue automaton, only the transition relation \rightarrow is now a subset of $\mathcal{S} \times \mathcal{A}_\tau \times (\mathcal{D} \cup \{\varepsilon, *\})^2 \times (\mathcal{D}^*)^2 \times \mathcal{S}$, i.e. a pair of elements of $\mathcal{D} \cup \{\varepsilon, *\}$ and a pair of sequences from \mathcal{D}^* is considered. We write $s \xrightarrow{a[(d,e)/(\delta,\zeta)]} t$ for $(s, a, d, e, \delta, \zeta, t) \in \rightarrow$ (here, $s, t \in \mathcal{S}, a \in \mathcal{A}_\tau, d, e \in \mathcal{D} \cup \{\varepsilon, *\}, \delta, \zeta \in \mathcal{D}^*$).

From this definition, we get a process graph as expected: the states of the process graph are the triples (s, δ, ζ) reachable from initial state $(\uparrow, \varepsilon, \varepsilon)$ by means of the transition relation generated from the following clauses:

1. if $s \xrightarrow{a[(\varepsilon,\varepsilon)/(\delta,\delta')]} t$ then $(s, \varepsilon, \varepsilon) \xrightarrow{a} (t, \delta, \delta')$;
2. if $s \xrightarrow{a[(\varepsilon,*)/(\delta,\delta')]} t$ then $(s, \varepsilon, \zeta) \xrightarrow{a} (t, \delta, \delta'\zeta)$ for all $\zeta \in \mathcal{D}^*$;
3. if $s \xrightarrow{a[(*,\varepsilon)/(\delta,\delta')]} t$ then $(s, \zeta, \varepsilon) \xrightarrow{a} (t, \delta\zeta, \delta')$ for all $\zeta \in \mathcal{D}^*$;
4. if $s \xrightarrow{a[(*,*)/(\delta,\delta')]} t$ then $(s, \zeta, \zeta') \xrightarrow{a} (t, \delta\zeta, \delta'\zeta')$ for all $\zeta, \zeta' \in \mathcal{D}^*$;
5. if $s \xrightarrow{a[(d,\varepsilon)/(\delta,\delta')]} t$ then $(s, \zeta d, \varepsilon) \xrightarrow{a} (t, \delta\zeta, \delta')$ for all $\zeta \in \mathcal{D}^*$;
6. if $s \xrightarrow{a[(\varepsilon,d)/(\delta,\delta')]} t$ then $(s, \varepsilon, \zeta d) \xrightarrow{a} (t, \delta, \delta'\zeta)$ for all $\zeta \in \mathcal{D}^*$;
7. if $s \xrightarrow{a[(d,*)/(\delta,\delta')]} t$ then $(s, \zeta d, \zeta') \xrightarrow{a} (t, \delta\zeta, \delta'\zeta')$ for all $\zeta, \zeta' \in \mathcal{D}^*$;
8. if $s \xrightarrow{a[(*,d)/(\delta,\delta')]} t$ then $(s, \zeta, \zeta'd) \xrightarrow{a} (t, \delta\zeta, \delta'\zeta')$ for all $\zeta, \zeta' \in \mathcal{D}^*$;
9. if $s \xrightarrow{a[(d,e)/(\delta,\delta')]} t$ then $(s, \zeta d, \zeta'e) \xrightarrow{a} (t, \delta\zeta, \delta'\zeta')$ for all $\zeta, \zeta' \in \mathcal{D}^*$.

Theorem 2 Let Q be a queue automaton with two queues. Then there is a queue automaton with one queue M such that the process graphs of Q and M are branching bisimilar.

Proof. Let Q be a queue automaton with two queues. We define the queue automaton with one queue M in the following. As in the proof of Lemma 1, we use an extra data element $\$$ to be able to traverse the contents of a queue. A state in Q will contain the contents of two queues, two sequences δ and ζ . In M , this will be encoded as the contents of one queue: $\delta \ \checkmark \ \zeta$, where extra data element \checkmark is the separator between the two sequences. A state (s, δ, ζ) in Q will correspond to the state $(s, \delta \ \checkmark \ \zeta)$ in M . Enumerate all transitions of Q in a sequence u_1, \dots, u_K . For each transition u_k we add four new states $s_k^1, s_k^2, s_k^3, s_k^4$ in M , moreover, there is a new initial state \uparrow' in M . The further elements are as follows.

1. M has the same action set and the same final states as Q ;
2. M has the data set \mathcal{D} of Q and in addition the symbols $\checkmark, \$$;

3. M has initial state \uparrow' and a transition labelled $\tau[\varepsilon/\check{\emptyset}]$ to the initial state of Q ;
4. if transition $s \xrightarrow{a[(*,*)/(\delta,\zeta)]} t$ is transition u_k in Q , then there are transitions $s \xrightarrow{a[*/*\$]} s_k^1 \xrightarrow{\tau[\check{\emptyset}/\check{\emptyset}\zeta]} s_k^2 \xrightarrow{\tau[\$/\delta]} t$ in M , and for all $d \in \mathcal{D}$ transitions $s_k^1 \xrightarrow{\tau[d/d]} s_k^1$ and $s_k^2 \xrightarrow{\tau[d/d]} s_k^2$; note that in this case, in the process graph of M , there is always an a -step possible, irrespective of the contents of the queues in the state;
5. if transition $s \xrightarrow{a[(*,\varepsilon)/(\delta,\zeta)]} t$ is transition u_k in Q , then there are transitions $s \xrightarrow{a[\check{\emptyset}/\check{\emptyset}\zeta\$/]} s_k^1 \xrightarrow{\tau[\$/\delta]} t$ and for all $d \in \mathcal{D}$ transitions $s_k^1 \xrightarrow{\tau[d/d]} s_k^1$ in M ; in this case, in the process graph of M , there is only an a -step possible if the second queue is empty, which can be checked by seeing $\check{\emptyset}$ at the head of the queue;
6. if transition $s \xrightarrow{a[(\varepsilon,\varepsilon)/(\delta,\zeta)]} t$ is transition u_k in Q , then in M , there is only an a -labeled transition if both queues are empty. We can check whether the second queue is empty by seeing $\check{\emptyset}$ at the head of the queue, for the first queue we have to traverse the queue. The succeeding transitions in M (i.e., the transitions that correspond to the situation that the checks that both queues are empty are successful) are $s \xrightarrow{\tau[\check{\emptyset}/\check{\emptyset}\$/]} s_k^1 \xrightarrow{\tau[\$/\varepsilon]} s_k^2 \xrightarrow{a[\check{\emptyset}/\delta\check{\emptyset}\zeta]} t$ and the failing transitions (i.e., the transitions that correspond to the situation that the check revealed that one of the queues is not empty) are, for all $d \in \mathcal{D}$, $s_k^1 \xrightarrow{\tau[d/d]} s_k^3$ and $s_k^3 \xrightarrow{\tau[d/d]} s_k^3$ and, finally, $s_k^3 \xrightarrow{\tau[\$/\varepsilon]} s$;
7. if transition $s \xrightarrow{a[(\varepsilon,*)/(\delta,\zeta)]} t$ is transition u_k in Q , then again, we need to check whether the first queue is empty. In M , we add the succeeding transitions $s \xrightarrow{\tau[*/*\$/]} s_k^1 \xrightarrow{\tau[\check{\emptyset}/\varepsilon]} s_k^2 \xrightarrow{a[\$/\delta\check{\emptyset}\zeta]} t$ and for all $d \in \mathcal{D}$ transitions $s_k^1 \xrightarrow{\tau[d/d]} s_k^1$; moreover, we add in M the failing transitions $s_k^3 \xrightarrow{\tau[\$/\varepsilon]} s$ and for all $d \in \mathcal{D}$ transitions $s_k^2 \xrightarrow{\tau[d/d\check{\emptyset}]} s_k^3$ and $s_k^3 \xrightarrow{\tau[d/d]} s_k^3$;
8. if transition $s \xrightarrow{a[(*,d)/(\delta,\zeta)]} t$ is transition u_k in Q , then there are transitions $s \xrightarrow{a[d/\$/]} s_k^1 \xrightarrow{\tau[\check{\emptyset}/\check{\emptyset}\zeta]} s_k^2 \xrightarrow{\tau[\$/\delta]} t$ and for all $f \in \mathcal{D}$ transitions $s_k^1 \xrightarrow{\tau[f/f]} s_k^1$ and $s_k^2 \xrightarrow{\tau[f/f]} s_k^2$ in M ;
9. if transition $s \xrightarrow{a[(\varepsilon,d)/(\delta,\zeta)]} t$ is transition u_k in Q , then, in M , then there are succeeding transitions $s \xrightarrow{\tau[d/\$/]} s_k^1 \xrightarrow{\tau[\check{\emptyset}/\varepsilon]} s_k^2 \xrightarrow{a[\$/\delta\check{\emptyset}\zeta]} t$ and for all $f \in \mathcal{D}$ transitions $s_k^1 \xrightarrow{\tau[f/f]} s_k^1$; moreover, there are failing transitions $s_k^3 \xrightarrow{\tau[\$/\varepsilon]} s$ and, for all $f \in \mathcal{D}$, $s_k^2 \xrightarrow{\tau[f/f\check{\emptyset}]} s_k^3$ and $s_k^3 \xrightarrow{\tau[f/f]} s_k^3$;
10. if transition $s \xrightarrow{a[(d,\varepsilon)/(\delta,\zeta)]} t$ is transition u_k in Q , then, in M , there are succeeding transitions $s \xrightarrow{\tau[\check{\emptyset}/\$/]} s_k^1 \xrightarrow{a[d/\check{\emptyset}\zeta]} s_k^2 \xrightarrow{\tau[\$/\delta]} t$ and for all $f \in \mathcal{D}$ transitions $s_k^2 \xrightarrow{\tau[f/f]} s_k^2$; moreover, there are failing transitions $s_k^1 \xrightarrow{\tau[\$/\varepsilon]} s$ and $s_k^3 \xrightarrow{\tau[\$/\varepsilon]} s$ and, for all $e \in \mathcal{D}$, $e \neq d$, $s_k^1 \xrightarrow{\tau[e/e\check{\emptyset}]} s_k^3$ and, for all $f \in \mathcal{D}$, $s_k^3 \xrightarrow{\tau[f/f]} s_k^3$;
11. if transition $s \xrightarrow{a[(d,*)/(\delta,\zeta)]} t$ is transition u_k in Q , then, in M , there are succeeding transitions $s \xrightarrow{\tau[*/*\$/]} s_k^1 \xrightarrow{\tau[\check{\emptyset}/\varepsilon]} s_k^2 \xrightarrow{a[d/\check{\emptyset}\zeta]} s_k^3 \xrightarrow{\tau[\$/\delta]} t$ and for all $f \in \mathcal{D}$ transitions $s_k^1 \xrightarrow{\tau[f/f]} s_k^1$ and $s_k^3 \xrightarrow{\tau[f/f]} s_k^3$; moreover, there are failing transitions

- $s_k^1 \xrightarrow{\tau[\$/\varepsilon]} s$ and $s_k^4 \xrightarrow{\tau[\$/\varepsilon]} s$ and, for all $e \in \mathcal{D}, e \neq d, s_k^2 \xrightarrow{\tau[e/e\check{Q}]} s_k^4$ and, for all $f \in \mathcal{D}, s_k^4 \xrightarrow{\tau[f/f]} s_k^4$;
12. if transition $s \xrightarrow{a[(d,e)/(\delta,\zeta)]} t$ is transition u_k in Q , then, in M , there are succeeding transitions $s \xrightarrow{\tau[e/\$]} s_k^1 \xrightarrow{\tau[\check{Q}/\varepsilon]} s_k^2 \xrightarrow{a[d/\check{Q}\zeta]} s_k^3 \xrightarrow{\tau[\$/\delta]} t$ and for all $f \in \mathcal{D}$ transitions $s_k^1 \xrightarrow{\tau[f/f]} s_k^1$ and $s_k^3 \xrightarrow{\tau[f/f]} s_k^3$; moreover, there are failing transitions $s_k^2 \xrightarrow{\tau[\$/\varepsilon]} s$ and $s_k^4 \xrightarrow{\tau[\$/\varepsilon]} s$ and, for all $g \in \mathcal{D}, g \neq d, s_k^2 \xrightarrow{\tau[g/g\check{Q}]} s_k^4$ and, for all $f \in \mathcal{D}, s_k^4 \xrightarrow{\tau[f/f]} s_k^4$.

Finally, it is straightforward to check that in the process graph of M , all the extra τ -steps are inert.

We see that the set of processes and the set of languages given by a queue automaton is not enlarged by adding another queue memory. Further on, we use this to show that the composition of two interacting processes both given by queue automata is again given by a queue automaton.

4 Comparison with Reactive Turing Machines

We use the definition of the Reactive Turing Machine (RTM) from [8].

Definition 10 (Reactive Turing machine) *A reactive Turing machine M is a sextuple $(\mathcal{S}, \mathcal{A}, \mathcal{D}, \rightarrow, \uparrow, \downarrow)$ where:*

1. \mathcal{S} is a non-empty finite set of states;
2. \mathcal{A} is a non-empty finite action alphabet;
3. \mathcal{D} is a non-empty finite data alphabet;
4. $\rightarrow \subseteq \mathcal{S} \times \mathcal{A}_\tau \times (\mathcal{D} \cup \{\square\}) \times (\mathcal{D} \cup \{\square\}) \times \{L, R\} \times \mathcal{S}$ is a finite set of transitions or steps;
5. $\uparrow \in \mathcal{S}$ is the initial state;
6. $\downarrow \subseteq \mathcal{S}$ is the set of final states.

The blank \square represents an empty tape cell. Henceforth, we will write \mathcal{D}_\square instead of $\mathcal{D} \cup \{\square\}$. If $(s, a, d, e, T, t) \in \rightarrow$, we write $s \xrightarrow{a[d/e]T} t$, and this means that the machine, when it is in state s and d is the data element read by the tape head, can execute action a , replace d by e , can move one position left (L) or right (R) and end up in state t .

It requires quite some notational overhead to define the transition relation and the process graph associated to an RTM. The states of the process graph are the configurations of the RTM, consisting of a state of the RTM, the contents of the tape, and the position of the read/write head on the tape. We represent the tape contents by an element of $(\mathcal{D}_\square)^*$, replacing exactly one occurrence of a tape symbol d by a marked symbol \check{d} , indicating that the read/write head is on this symbol. We denote by $\check{\mathcal{D}}_\square = \{\check{d} \mid d \in \mathcal{D}_\square\}$ the set of marked tape symbols; a *tape instance* is a sequence $\delta \in (\mathcal{D}_\square \cup \check{\mathcal{D}}_\square)^*$ containing exactly one element of $\check{\mathcal{D}}_\square$. Note that we do not use δ exclusively for tape instances;

we also use δ for sequences over \mathcal{D} . A tape instance thus is a finite sequence of symbols that represents the contents of a two-way infinite tape. Henceforth, we do not distinguish between tape instances that are equal modulo the addition or removal of extra occurrences of the blank symbol \square at the left or right extremes of the sequence. That is, we do not distinguish tape instances δ and ζ if $\square^\omega \delta \square^\omega = \square^\omega \zeta \square^\omega$.

A *configuration* of an RTM M is a pair (s, δ) where $s \in \mathcal{S}$ is a state of the RTM and δ is a tape instance.

We define an \mathcal{A}_τ -labelled transition system for each RTM such that a transition $s \xrightarrow{a[d/e]T} t$ corresponds to a transition $(s, \delta) \xrightarrow{a} (t, \zeta)$, where in δ some occurrence of d is marked, and in ζ this marked d is replaced by e , and the symbol to the left in ζ is marked (if $T = L$) or the symbol to the right in ζ is marked (if $T = R$). If necessary, a blank \square is added. For this, we use the following notation: if $\delta \in \mathcal{D}_\square^*$, then $\delta^< = \check{\square}$ if $\delta = \varepsilon$ and $\delta^< = \check{\zeta}d$ if $\delta = \zeta d$ for some $d \in \mathcal{D}_\square, \zeta \in \mathcal{D}_\square^*$. Likewise, if $\delta \in \mathcal{D}_\square^*$, then $\delta^> = \check{\square}$ if $\delta = \varepsilon$ and $\delta^> = d\check{\zeta}$ if $\delta = d\zeta$ for some $d \in \mathcal{D}_\square, \zeta \in \mathcal{D}_\square^*$.

Definition 11 Let $M = (\mathcal{S}, \mathcal{A}, \mathcal{D}, \rightarrow, \uparrow, \downarrow)$ be an RTM. The process graph associated with M is defined as follows:

1. the set of states is the set of configurations (s, δ) of M ;
2. the set of actions is \mathcal{A} ;
3. the transition relation \rightarrow is the least relation satisfying, for all $a \in \mathcal{A}_\tau, d, e \in \mathcal{D}_\square, \delta, \zeta \in \mathcal{D}_\square^*$:

$$(s, \delta\check{d}\zeta) \xrightarrow{a} (t, \delta^< e\zeta) \iff s \xrightarrow{a[d/e]L} t$$

and

$$(s, \delta d\check{\zeta}) \xrightarrow{a} (t, \delta e^> \zeta) \iff s \xrightarrow{a[d/e]R} t$$

4. the initial state is $(\uparrow, \check{\square})$;
5. the set of final states is $\{(s, \delta) \mid s \downarrow\}$.

A process is executable iff its branching bisimulation equivalence class of process graphs contains a process graph associated with an RTM. A language is computable iff it is the language of a process graph associated with an RTM.

Theorem 3 A process is executable if and only if it is the process of a queue automaton.

Proof. Let $M = (\mathcal{S}, \mathcal{A}, \mathcal{D}, \rightarrow, \uparrow, \downarrow)$ be an RTM, and suppose M has n transitions T_1, \dots, T_n . For each transition T_i , we have 5 new states $s_i^1, s_i^2, s_i^3, s_i^4, s_i^5$. In addition, we have an extra state \uparrow' . A configuration $(s, \delta\check{d}\zeta)$ of the RTM will correspond to the state $(s, \zeta^R \check{\zeta} \delta d)$ of the queue automaton to be constructed, where the symbol $\check{\zeta}$ is a separator and ζ^R is the reverse of the string ζ . Note that we have to treat \square as an extra data element, because the RTM uses blanks in this way. Now we define the queue automaton Q as follows:

1. the set of states is \mathcal{S} , extended with new states $\{s_i^1, s_i^2, s_i^3, s_i^4, s_i^5 \mid i \leq n\} \cup \{\uparrow'\}$;

2. the set of actions is \mathcal{A} ;
3. the set of data is $\mathcal{D} \cup \{\square, \checkmark, \$\}$ ($\$$ is a new data element used as a bookmark);
4. if $s \xrightarrow{a[d/e]L} t$ is the transition T_i in M , then there are transitions $s \xrightarrow{a[d/\$]} s_i^1 \xrightarrow{\tau[\checkmark/e\checkmark\square]} s_i^2 \xrightarrow{\tau[\$/\varepsilon]} t$ and, for each $f \in \mathcal{D}_\square$, transitions $s_i^1 \xrightarrow{\tau[f/f]} s_i^1$ and $s_i^2 \xrightarrow{\tau[f/f]} s_i^2$ in Q ; moreover, if $s \xrightarrow{a[d/e]R} t$ is transition T_j in M , then there are transitions $s \xrightarrow{a[d/\$e]} s_j^1 \xrightarrow{\tau[\checkmark/\checkmark]} s_j^2 \xrightarrow{\tau[\$/\square\$]} s_j^3 \xrightarrow{\tau[\$/\varepsilon]} t$ and, for each $f, g \in \mathcal{D}_\square$, transitions $s_j^2 \xrightarrow{\tau[g/\varepsilon]} s_j^3 \xrightarrow{\tau[\$/g\$]} s_j^4, s_j^1 \xrightarrow{\tau[f/f]} s_j^1, s_j^3 \xrightarrow{\tau[f/f]} s_j^3$ and $s_j^4 \xrightarrow{\tau[f/f]} s_j^4$ in Q .
5. the initial state is \uparrow and add a transition $\uparrow \xrightarrow{\tau[\varepsilon/\checkmark\square]} \uparrow'$ in Q .
6. the set of final states is \downarrow .

Then, we can establish that the process graph of Q is branching bisimilar to the process graph associated with M .

For the other direction, suppose a process has a process graph given by a queue automaton $Q = (\mathcal{S}, \mathcal{A}, \mathcal{D}, \rightarrow, \uparrow, \downarrow)$. Without loss of generality, we can suppose Q has only singleton enqueues and separate dequeues. Suppose Q has n transitions T_1, \dots, T_n . For each transition T_i , we have a new state s'_i . A state $(s, \delta d)$ of the queue automaton will correspond to the configuration $(s, \square \delta d \square)$ of the RTM to be constructed, and state (s, ε) to configuration (s, \checkmark) . Now define an RTM M as follows:

1. the set of states is \mathcal{S} , extended with new states $\{s'_i, s''_i \mid i \leq n\}$;
2. the set of actions is \mathcal{A} , the set of data is \mathcal{D} ;
3. if $s \xrightarrow{a[d/\varepsilon]} t$ is a transition in Q , then $s \xrightarrow{a[d/\square]L} t$ is a transition in M ; further, if $s \xrightarrow{a[\varepsilon/\varepsilon]} t$ is a transition in Q , then $s \xrightarrow{a[\square/\square]L} t$ is a transition in M ; lastly, if $s \xrightarrow{a[* / d]} t$ is the transition T_i in Q , then $s \xrightarrow{a[\square / d]R} s'_i \xrightarrow{\tau[\square/\square]L} t$ are transitions in M and, for each $e \in \mathcal{D}$, $s \xrightarrow{a[e/e]L} s''_i$ and $s''_i \xrightarrow{\tau[e/e]L} s''_i$ and $s'_i \xrightarrow{\tau[e/\varepsilon]R} s'_i$ are also transitions in M ; finally, also $s''_i \xrightarrow{\tau[\square/d]R} s'_i$ is a transition in M ;
4. the initial state is \uparrow , and the set of final states is \downarrow .

Again, we can establish that the process graph of M is branching bisimilar to the process graph of Q .

Corollary 1 *A language is computable if and only if it is the language of a queue automaton.*

A Turing machine can also be used to define when a function is computable. In [8], we defined this also for the Reactive Turing Machine. Here, we do this for the queue automaton, and give a couple of examples. We designate an input port i and an output port o .

Definition 12 *We say a queue automaton performs a computation if*

1. *the process graph of the queue automaton is deterministic;*

2. every string in the language of the queue automaton consists of a sequence of inputs $i?d_1 \cdots i?d_n$ and a sequence of outputs $o!e_1 \cdots o!e_m$ interleaved (so not necessarily all outputs after all inputs) for some $n, m \geq 0, d_i, e_j \in \mathcal{D}$. Note that the length of input and output may differ.

In this case, we say the queue automaton computes the function f on a domain of data strings D ($D \subseteq \mathcal{D}^*$) if for all input $w \in D$ it has output $f(w) \in \mathcal{D}^*$.

Example 4 By an adaptation of Example 1, we can define the function $f(w) = ww$ on \mathcal{D}^* . See Figure 4.




Fig. 4. Queue automaton for the function $f(w) = ww$ on $\{a, b\}^*$.

Example 5 There is a queue automaton that compares quantities. Suppose we have two numbers in binary notation and we want to know whether the first number is larger than the second. For simplicity, we assume the numbers have an equal number of digits, adding leading zeroes if necessary. The input consists of the two numbers separated by a $>$ -sign, and $\mathcal{D} = \{0, 1, >, \text{yes}, \text{no}\}$. Figure 5 explains the rest. Note that no τ -steps are needed.

Thus, a queue automaton can be used to program a conditional branching in a program. In the same way, we can find queue automata for other program constructs, and other mathematical functions. Notice that there is a clear separation between input and output on the one hand, and memory use on the other hand. In this case, we only need to store the first number in memory, not the second one. Also notice that as soon as we have a difference between the two numbers, we can determine the output, and there is no need for the rest of the input.

In the following, we use some basic recursion theory. The reader is referred to, e.g., [22]. In [8], we characterised the set of executable processes as follows. We call a process graph *effective* if its transition relation and its set of final states are recursively enumerable (with some suitable encoding of these into natural numbers), see [10]. A process is effective if its branching bisimulation equivalence class contains an effective




Fig. 5. Queue automaton comparing quantities.

process graph. We proved in [8] that a process is effective if and only if it is executable. In this result, it is needed to abstract from divergencies (infinite sequences of inert τ -steps). Also note that the process graph of a queue automaton is always boundedly branching, but the minimal element of its branching bisimulation equivalence class can be infinitely branching. This is the case for the queue automaton shown in Figure 6.




Fig. 6. Queue automaton, of which the process graph is branching bisimilar to an infinitely branching process graph.

5 Process Algebra

In this section, we express the executable processes in a mathematical formalism. In the theory of automata and formal languages, such a notation is often called a *grammar*. In concurrency theory or process theory, such a notation is often called a *process algebra*. In this process algebra, we can concisely express communication between executable processes.

A queue automaton can be thought of as a finite automaton extended with queue memory. In this section, we shall formalise this idea using process algebra, by proving that every executable process can be specified as a regular process interacting with the queue process. Thereafter, we will also briefly consider universal queue automata.

In [5], it was proven that every executable process can be defined as the solution of a finite recursive specification over the process algebra ACP_τ ; in [3], this result was updated to the process algebra BCP_τ with standard communication. We proceed to present this process algebra here.

Definition 13 *The syntax of BCP_τ with standard communication has the following ingredients:*

1. *there is a constant $\mathbf{0}$ denoting inaction or deadlock: it denotes the finite automaton with a single state that is initial but not final and that has no transitions;*
2. *there is a constant $\mathbf{1}$ denoting termination or acceptance: it denotes the finite automaton with a single state that is initial and final and that has no transitions;*
3. *there is a finite set of constants \mathcal{X} of process identifiers or process variables; we use capital letters X, Y, \dots to range over \mathcal{X} ;*
4. *there is a set \mathcal{A} of actions, a finite set \mathcal{D} of data and a finite set \mathcal{C} of communication ports: for each $d \in \mathcal{D}, c \in \mathcal{C}$ there are actions $c!d$ (send or output d at c), $c?d$ (receive or input d at c) and $c(d)$ (communicate d at c); in addition, there is the unobservable action τ ; for each action $a \in \mathcal{A}_\tau$, there is the unary action prefix $a._;$*
5. *there is the binary operator $+$ denoting choice or alternative composition;*
6. *there is the binary operator \parallel denoting merge or parallel composition;*
7. *for each communication port c , there is the unary operator $\partial_c(_)$ denoting restriction or encapsulation of all actions $c?d, c!d$ for $d \in \mathcal{D}$;*
8. *there is the unary operator $\tau_C(_)$ denoting hiding or abstraction of all actions $c(d)$ for $d \in \mathcal{D}, c \in \mathcal{C}$.*

A recursive specification over BCP_τ is a mapping Γ from \mathcal{X} to the set of BCP_τ expressions. The idea is that the process expression p associated with a process identifier $X \in \mathcal{X}$ by Γ defines the behaviour of X . We prefer to think of Γ as a collection of *defining equations* $X \stackrel{\text{def}}{=} p$, exactly one for every $X \in \mathcal{X}$. We shall, in the sequel, presuppose a recursive specification Γ defining the process identifiers in \mathcal{X} , and we shall usually simply write $X \stackrel{\text{def}}{=} p$ for $\Gamma(X) = p$. Note that, by our assumption that \mathcal{X} is finite, Γ is finite too.

We associate behaviour with process expressions by defining, on the set of process expressions, a unary acceptance predicate \downarrow (written postfix) and, for every $a \in \mathcal{A}_\tau$, a binary transition relation \xrightarrow{a} (written infix), by means of the transition system specification presented in Figure 7.

By means of these rules, the set of process expressions turns into a labelled transition system, so we have strong bisimilarity and branching bisimilarity on process expressions. Suppose p is a BCP_τ expression (possibly containing variables from \mathcal{X}). Then the *process* of p is the branching bisimulation equivalence class of the process graph generated by the operational rules.

We explicitly state the result in [3]:

Theorem 4 *The process of every BCP_τ expression is executable. For every executable process, there is a BCP_τ expression with this process.*

$$\begin{array}{c}
\frac{}{\mathbf{1} \downarrow} \quad \frac{}{a.p \xrightarrow{a} p} \\
\frac{p \downarrow}{(p+q) \downarrow} \quad \frac{q \downarrow}{(p+q) \downarrow} \quad \frac{p \xrightarrow{a} p'}{p+q \xrightarrow{a} p'} \quad \frac{q \xrightarrow{a} q'}{p+q \xrightarrow{a} q'} \\
\frac{p \downarrow \quad q \downarrow}{p \parallel q \downarrow} \quad \frac{p \xrightarrow{a} p'}{p \parallel q \xrightarrow{a} p' \parallel q} \quad \frac{q \xrightarrow{a} q'}{p \parallel q \xrightarrow{a} p \parallel q'} \\
\frac{p \xrightarrow{c!d} p' \quad q \xrightarrow{c?d} q'}{p \parallel q \xrightarrow{c(d)} p' \parallel q'} \quad \frac{p \xrightarrow{c?d} p' \quad q \xrightarrow{c!d} q'}{p \parallel q \xrightarrow{c(d)} p' \parallel q'} \\
\frac{p \downarrow}{\partial_c(p) \downarrow} \quad \frac{p \xrightarrow{a} p' \quad a \neq c!d, c?d}{\partial_c(p) \xrightarrow{a} \partial_c(p')} \\
\frac{p \downarrow}{\tau_C(p) \downarrow} \quad \frac{p \xrightarrow{c(d)} p'}{\tau_C(p) \xrightarrow{\tau} \tau_C(p')} \quad \frac{p \xrightarrow{a} p' \quad a \neq c(d)}{\tau_C(p) \xrightarrow{a} \tau_C(p')} \\
\frac{p \downarrow \quad X \stackrel{\text{def}}{=} p}{X \downarrow} \quad \frac{p \xrightarrow{a} p' \quad X \stackrel{\text{def}}{=} p}{X \xrightarrow{a} p'}
\end{array}$$

Fig. 7. Operational semantics for BCP_τ with standard communication ($a \in \mathcal{A}_\tau, c \in \mathcal{C}, d \in \mathcal{D}, X \in \mathcal{X}$).

Corollary 2 *Let p, q be two BCP_τ expressions, and $c \in \mathcal{C}$ a communication port. Then $p \parallel q$, $\partial_c(p \parallel q)$ and $\tau_C(\partial_c(p \parallel q))$ denote executable processes.*

In this way, we can express the communication of two processes. Using Theorem 2, we can also obtain this result directly for queue automata.

Example 6 *We give the following finite specification of the queue process of Example 3, adapted from [11]:*

$$\begin{aligned}
Q^{io} &\stackrel{\text{def}}{=} \mathbf{1} + o!\varepsilon.Q^{io} + \sum_{d \in \mathcal{D}} i?d.\tau_C(\partial_\ell(Q^{i\ell} \parallel (\mathbf{1} + o!d.Q^{\ell o}))) \\
Q^{i\ell} &\stackrel{\text{def}}{=} \mathbf{1} + \ell!\varepsilon.Q^{i\ell} + \sum_{d \in \mathcal{D}} i?d.\tau_C(\partial_o(Q^{io} \parallel (\mathbf{1} + \ell!d.Q^{oi}))) \\
Q^{\ell o} &\stackrel{\text{def}}{=} \mathbf{1} + o!\varepsilon.Q^{\ell o} + \sum_{d \in \mathcal{D}} \ell?d.\tau_C(\partial_i(Q^{\ell i} \parallel (\mathbf{1} + o!d.Q^{io}))) \\
Q^{o\ell} &\stackrel{\text{def}}{=} \mathbf{1} + \ell!\varepsilon.Q^{o\ell} + \sum_{d \in \mathcal{D}} o?d.\tau_C(\partial_i(Q^{oi} \parallel (\mathbf{1} + \ell!d.Q^{i\ell}))) \\
Q^{\ell i} &\stackrel{\text{def}}{=} \mathbf{1} + i!\varepsilon.Q^{\ell i} + \sum_{d \in \mathcal{D}} \ell?d.\tau_C(\partial_o(Q^{\ell o} \parallel (\mathbf{1} + o!d.Q^{oi})))
\end{aligned}$$

$$Q^{oi} \stackrel{def}{=} \mathbf{1} + i!\varepsilon.Q^{oi} + \sum_{d \in \mathcal{D}} o?d.\tau_C(\partial_\ell(Q^{o\ell} \parallel (\mathbf{1} + i!d.Q^{\ell i})))$$

This specification uses data set $\mathcal{D} \cup \{\varepsilon\}$.

Recall that a regular process has a process graph with finitely many states and transitions.

Theorem 5 *Let p be an executable process. Then there is a regular process q such that $p \stackrel{b}{\leftrightarrow} \tau_C(\partial_{io}(q \parallel Q^{io}))$.*

Proof. Let p be an executable process. Then there is a queue automaton $Q = (\mathcal{S}, \mathcal{A}, \mathcal{D}, \rightarrow, \uparrow, \downarrow)$ that defines p . By Lemma 2, we can assume without loss of generality that Q has singleton enqueues and separate dequeues. In order to define q , we have to remember the head of the queue or remember that the queue is empty, in order to be able to know which following step is possible.

1. for each state $s \in \mathcal{S}$ and $d \in \mathcal{D} \cup \{\varepsilon\}$, q has states s_d, s_d^1, s_d^2, s_d^3 ;
2. the queue Q^{io} uses data set $\mathcal{D} \cup \{\$\}$;
3. the initial state of q is \uparrow_ε , and state s_d is final whenever s is final in Q ;
4. whenever Q has $s \xrightarrow{a[\varepsilon/\varepsilon]} t$, then q has $s_\varepsilon \xrightarrow{a} t_\varepsilon$;
5. whenever Q has $s \xrightarrow{a[* / d]} t$ for some $d \in \mathcal{D}$, then q has $s_e \xrightarrow{a} t_e^1 \xrightarrow{i!d} t_e$ for all $e \in \mathcal{D}$ and $s_e \xrightarrow{a} t_d^1 \xrightarrow{i!d} t_d$;
6. whenever Q has $s \xrightarrow{a[d/\varepsilon]} t$ for some $d \in \mathcal{D}$, then q has $s_d \xrightarrow{a} s_d^1 \xrightarrow{o?d} s_d^2 \xrightarrow{o?\varepsilon} t_\varepsilon$; moreover, for all $e \in \mathcal{D}$ q has steps $s_d^2 \xrightarrow{o?e} t_e^3 \xrightarrow{i!\$} t_e^2 \xrightarrow{i!e} t_e^1 \xrightarrow{o?\$} t_e$ and for all $f \in \mathcal{D}$ steps $t_e^1 \xrightarrow{o?f} t_e^1$.

All input and output steps in q will successfully communicate with the queue, and be turned into τ -steps by the abstraction operator. All will turn out to be inert.

Note that the converse of this theorem is also true, as every process defined by a finite specification over BCP_τ with standard communication is executable.

Just like we did in [8], we can prove a universal queue automaton U exists. For an arbitrary queue automaton M , let \overline{M} be the deterministic queue automaton that outputs the Gödel number of M (in some appropriate representation) along a special communication port u , then terminates with empty queue, and has no other behaviour. Then a *universal* queue automaton U is such that $\tau_C(\partial_u(\overline{M} \parallel U))$ has the same process as M for all queue automata M (here, we use the BCP_τ -expressions of these queue automata to define this process).

6 A hierarchy

A computer shows interaction between the finite control and the memory. The finite control can be represented by a regular process (a finite automaton). In this article, we considered a memory in the form of a queue.

In [7], we considered a memory in the form of a stack, and we established that a pushdown process can be characterised as a regular process communicating with a stack. This work was continued in [6] and [2] to find the process algebra $TSP_{;sc}$ that is associated with pushdown automata. $TSP_{;sc}$ is obtained from the process algebra BCP_τ of the previous section by leaving out parallel composition, encapsulation and abstraction, and adding sequencing with sequential value passing.

In [1], we considered a memory in the form of a bag, and we established that a parallel pushdown process can be characterised as a regular process communicating with a bag. We found the process algebra associated with parallel pushdown automata. This process algebra is obtained from the process algebra BCP_τ of the previous section by leaving out the abstraction operator and adding the priority operator of [4].

We see that the queue is the prototypical executable process, as all executable processes can be realised as a regular process communicating with a queue. Likewise, the bag is the prototypical parallel pushdown process. A bag is an executable process, but not a pushdown process. Further, the stack is the prototypical pushdown process, but not a parallel pushdown process.

A counter can be realised as a stack with a singleton data set, and also as a bag with a singleton data set. Thus, it is in the intersection of pushdown processes and parallel pushdown processes. It is not a regular process, as it has infinitely many different states that are not bisimilar. We conjecture that every process in the intersection of pushdown processes and parallel pushdown processes can be realised as a regular process communicating with a counter, but have no proof of this yet. Figure 8 provides a complete picture.




Fig. 8. Classification of Executable, Pushdown, Parallel Pushdown, and Regular processes and the prototypical processes Queue, Bag, Stack and Counter.

7 Conclusion

We considered the computational model of the Queue Automaton, and have proved that it is equally expressive as the Reactive Turing Machine of [8]. Thus, a process is executable if and only if it is the process of a queue automaton, and a language or function is computable if and only if it is the language or function of a queue automaton. Every executable process can be defined as a regular process communicating with a queue. This fits in very well with earlier results, that every pushdown process can be defined as a regular process communicating with a stack, and a parallel pushdown process can be defined as a regular process communicating with a bag. We think that a pushdown automaton can be better called a stack automaton, and a parallel pushdown automaton a bag automaton, in order to emphasise the relation to the queue automaton. As grammar for executable processes we use the process algebra BCP_{τ} , Basic Communicating Processes with abstraction and standard communication.

References

1. Jos C. M. Baeten & Bas Luttik (2023): *Parallel Pushdown Automata and Commutative Context-Free Grammars in Bisimulation Semantics (Extended Abstract)*. In Claudio Antares Mezzina & Georgiana Caltais, editors: *Proceedings Combined 30th International Workshop on Expressiveness in Concurrency and 20th Workshop on Structural Operational Semantics, EXPRESS/SOS 2023, and 20th Workshop on Structural Operational Semantics Antwerp, Belgium, 18th September 2023, EPTCS 387*, pp. 114–131, <https://doi.org/10.4204/EPTCS.387.9>.
2. Jos C. M. Baeten & Bas Luttik (2024): *Sequential Value Passing Yields a Kleene Theorem for Processes*. In Venanzio Capretta, Robbert Krebbers & Freek Wiedijk, editors: *Logics and Type Systems in Theory and Practice - Essays Dedicated to Herman Geuvers on The Occasion of His 60th Birthday, Lecture Notes in Computer Science 14560*, Springer, pp. 1–16, https://doi.org/10.1007/978-3-031-61716-4_1.
3. Jos C.M. Baeten, Twan Basten & Michel Reniers (2010): *Process algebra: equational theories of communicating processes*. 50, Cambridge university press.
4. Jos C.M. Baeten, Jan A. Bergstra & Jan Willem Klop (1986): *Syntax and Defining Equations for an Interrupt Mechanism in Process Algebra*. *Fundamenta Informaticae* 9, pp. 127–168.
5. Jos C.M. Baeten, Jan A. Bergstra & Jan Willem Klop (1987): *On the consistency of Koomen’s fair abstraction rule*. *Theoretical Computer Science* 51(1-2), pp. 129–176.
6. Jos C.M. Baeten, Cesare Carissimo & Bas Luttik (2023): *Pushdown Automata and Context-Free Grammars in Bisimulation Semantics*. *Logical Methods in Computer Science* 19, pp. 15:1–15.32, [https://doi.org/10.46298/LMCS-19\(1:15\)2023](https://doi.org/10.46298/LMCS-19(1:15)2023).
7. Jos C.M. Baeten, Pieter Cuijpers & Paul van Tilburg (2008): *A context-free process as a pushdown automaton*. In F. van Breugel & M. Chechik, editors: *Proceedings CONCUR’08, Lecture Notes in Computer Science 5201*, pp. 98–113.
8. Jos C.M. Baeten, Bas Luttik & Paul van Tilburg (2013): *Reactive Turing machines*. *Information and Computation* 231, pp. 143 – 166, <https://doi.org/https://doi.org/10.1016/j.ic.2013.08.010>. *Fundamentals of Computation Theory*.
9. Twan Basten (1996): *Branching Bisimulation is an Equivalence Indeed!* *Information Processing Letters* 58, pp. 141–147.

10. Gérard Boudol (1985): *Notes on Algebraic Calculi of Processes*. In Krzysztof R. Apt, editor: *Logics and Models of Concurrent Systems*, Springer Verlag, pp. 261–303.
11. Rob J. van Glabbeek & Frits W. Vaandrager (1993): *Modular Specification of Process Algebras*. *Theor. Comput. Sci.* 113(2), pp. 293–348, [https://doi.org/10.1016/0304-3975\(93\)90006-F](https://doi.org/10.1016/0304-3975(93)90006-F).
12. Rob J. van Glabbeek (1993): *The Linear Time – Branching Time Spectrum II; The semantics of sequential systems with silent moves (extended abstract)*. In E. Best, editor: *Proceedings CONCUR'93, 4th International Conference on Concurrency Theory*, Hildesheim, Germany, August 1993, *Lecture Notes in Comp. Sci.* 715, Springer Verlag, pp. 66–81, https://doi.org/10.1007/3-540-57208-2_6.
13. Rob van Glabbeek, Bas Luttik & Nikola Trcka (2009): *Branching Bisimilarity with Explicit Divergence*. *Fundamenta Informaticae* 93(4), pp. 371–392, <https://doi.org/10.3233/FI-2009-109>.
14. Rob van Glabbeek & Peter Weijland (1996): *Branching time and abstraction in bisimulation semantics*. *Journal of the ACM* 43(3), pp. 555–600, <https://doi.org/10.1145/233551.233556>.
15. John E. Hopcroft, Rajeev Motwani & Jeffrey D. Ullman (2007): *Introduction to Automata Theory, Languages, and Computation*. Pearson Education.
16. Martin Kutrib, Andreas Malcher & Matthias Wendlandt (2018): *Queue Automata: Foundations and Developments*. In Andrew Adamatzky, editor: *Reversibility and Universality: Essays Presented to Kenichi Morita on the Occasion of his 70th Birthday*, Springer International Publishing, pp. 385–431, https://doi.org/10.1007/978-3-319-73216-9_19.
17. Bas Luttik & Fei Yang (2021): *The π -Calculus is Behaviourally Complete and Orbit-Finitely Executable*. *Log. Methods Comput. Sci.* 17(1). Available at <https://lmcs.episciences.org/7165>.
18. Zohar Manna (1974): *Mathematical Theory of Computation*. McGraw-Hill.
19. Sjouke Mauw & Hans Mulder (1994): *Regularity of BPA-Systems is Decidable*. In Bengt Jonsson & Joachim Parrow, editors: *CONCUR '94, Concurrency Theory, Uppsala, Sweden, 1994, Proceedings, Lecture Notes in Computer Science* 836, Springer, pp. 34–47, https://doi.org/10.1007/978-3-540-48654-1_4.
20. Robin Milner (1989): *Communication and Concurrency*. Prentice Hall, Englewood Cliffs.
21. Emile L. Post (1943): *Formal reductions of the classical combinatorial decision problem*. *Am. J. Math.* 65, pp. 197–215.
22. Hartley Rogers, Jr. (1967): *Theory of Recursive Functions and Effective Computability*. McGraw-Hill.
23. Paul van Tilburg (2011): *From Computability to Executability*. Ph.D. thesis, Eindhoven University of Technology.
24. Alan M. Turing (1936): *On Computable Numbers, with an Application to the Entscheidungsproblem*. *Proceedings of the London Mathematical Society* 2(42), pp. 230–265.
25. Roland Vollmar (1970): *Über einen Automaten mit Pufferspeicherung*. *Computing* 5, pp. 57–70.