Computational Complexity of Interactive Behaviors

Ugo Dal Lago¹, Tobias Heindel², Damiano Mazza³, and Daniele Varacca⁴

¹ Dip. di Scienze dell'Informazione – Univ. di Bologna, Italy – dallago@cs.unibo.it
² CEA – LIST, Gif-sur-Yvette, France – Tobias.Heindel@cea.fr

³ LIPN – CNRS and Université Paris 13, France – mazza@lipn.univ-paris13.fr
⁴ PPS – CNRS and Université Paris Diderot, France – varacca@pps.jussieu.fr

Abstract. The theory of computational complexity focuses on functions and, hence, studies programs whose interactive behavior is reduced to a simple question/answer pattern. We propose a broader theory whose ultimate goal is expressing and analyzing the intrinsic difficulty of fully general interactive behaviors. To this extent, we use standard tools from concurrency theory, including labelled transition systems (formalizing behaviors) and their asynchronous extension (providing causality information). Behaviors are implemented by means of a multiprocessor machine executing CCS-like processes. The resulting theory is shown to be consistent with the classical definitions: when we restrict to functional behaviors (i.e., question/answer patterns), we recover several standard computational complexity classes.

Keywords: Computational complexity, Interactive computation, Process algebras, Asynchronous transition systems.

1 Introduction

In the early days, computers were considered as oracles: one would have a question and the computer would provide the answer. For instance, one day the American army had just launched a rocket to the Moon, and the four star General typed in two questions to the computer: (1) Will the rocket reach the Moon? (2) Will the rocket return to the Earth? The computer did some calculations for some time and then ejected a card which read: "Yes." The General was furious; he didn't know whether "Yes" was the answer to the first question, the second or both. Therefore he angrily typed in "Yes, what?". The computer did some more calculations and then printed on a card: "Yes, Sir."⁵

That every computation may eventually be reduced to the input/output pattern is an assumption underlying most of classical computability theory. The theory of computational complexity is an excellent example: it studies the intrinsic difficulty of *problems*, which are nothing but "yes or no" questions. Accordingly, the classical methods that measure the complexity of a program ignore the possibility that it may interact with its environment between the initial request and the final answer; even when a more complex interaction pattern is considered (*e.g.*, in interactive proofs [6]), it often is seen as yet another way to solve problems (viz. the class **IP**, which is a class of problems).

⁵ Adapted from Raymond Smullyan: What is the name of this book?

Nowadays, we live in a world of ubiquitous computing systems that are highly interactive and communicate with their environments, following possibly complicated protocols. These computing systems are fundamentally different from those that just provide answers to questions without any observable intermediate actions. To study this phenomenon of *interactive computation* [4], theoretical computer scientists have developed several formalisms and methodologies, such as process calculi and algebras.

However, little has been done so far to tackle the computational complexity of interactive systems (one of the few examples being the competitive analysis of online algorithms [3]). Note that, as mentioned above, this issue is beyond the classical theory of computational complexity. Thus, we set out to provide grounds for a revised theory of computational complexity that is capable of gauging the efficiency of genuinely interactive behaviors.

The first conceptual step is the formalization of *behaviors*. Our approach follows standard lore of concurrency theory: a behavior is an equivalence class of *labelled transition systems* (LTS), which in turn are usually specified using process calculi, such as Milner's CCS [10]. In this paper, we use a variation of bisimilarity as behavioral equivalence; however, other equivalences that have been proposed in the literature, such as coupled similarity and testing, work equally well, as long as certain minimal requirements are satisfied. Shifting the focus towards behaviors, the fundamental question of classical computational complexity "What is the cost of solving a problem (or implementing a function)?" becomes "What is the cost of implementing a behavior?".

A suitable cost model is not as easily found as in the functional case where we just measure the resources (time, space) required to compute the answer as a function of the size of the question. Of course, the resources depend on the chosen computational model (such as Turing machines), but the general scheme does not depend on the specific model. We propose a notion of cost for general interactive behaviors that abstracts away from a specific model. Costs are attributed to events in *weighted asynchronous* LTS (or WALTS): asynchrony is a standard feature that is added to transition systems to represent causal dependencies [14], which we need to generalize the trivial dependency between questions and answers; weights are used to specify additional quantitative information about space and time consumption.

Finally, we introduce a computational model, the *process machine*, which implements behaviors (just as Turing machines implement functions) by executing concurrent programs written in a CCS-based language. Such a machine has an unbounded number of processors each equipped with a private memory and capable of performing basic string manipulation and communicating asynchronously with other processors or the external environment. The process machine admits a natural semantics in terms of WALTSS and thus provides us with a non-trivial, paradigmatic instance of our abstract framework for measuring the complexity of behaviors.

Complexity classes are then defined as sets of behaviors that can be implemented by a process running within given time and space bounds on the process machine. We conclude by showing that if we restrict to functional behaviors (*i.e.*, trivial input/output patterns) we obtain several standard complexity classes; thus, at least in many paradigmatic cases, we have in fact a consistent extension of complexity theory into the realm of interactive computation.

2 Behaviors

In this section we formally define behaviors as equivalence classes of labelled transition systems. Such systems can receive messages on some *input channels*, send messages on some *output channels*, and perform internal, invisible, computation. With the aim of being as concrete as possible, we consider messages to be binary strings. We denote by $\mathbb{W} = \{0, 1\}^*$ the set of such strings, with ε denoting the empty string. We also fix two disjoint sets \mathcal{I}, \mathcal{O} of input and output channel names.

Definition 1 (Labelled transition system). An input action (*resp.* output action) is an element of $\mathcal{I} \times \mathbb{W}$ (*resp.* $\mathcal{O} \times \mathbb{W}$); together, they form the set of visible actions, denoted by \mathcal{A}_v . The set of actions is $\mathcal{A} = \mathcal{A}_v \cup \{\tau\}$, where τ is the internal action.

A labelled transition system (LTS for short) is a triple $S = (|S|, s_0, \operatorname{trans}_S)$, where |S| is a set, whose elements are called states, $s_0 \in |S|$ is the initial state, and $\operatorname{trans}_S \subseteq |S| \times A \times |S|$ is the transition relation.

Given an LTS S, we write $s \xrightarrow{\alpha} s'$ when $(s, \alpha, s') \in \text{trans}_S$. Since internal computation is invisible, it is standard practice to consider several internal steps as one single, still invisible step. We denote by \Longrightarrow the reflexive-transitive closure of $\xrightarrow{\tau}$ and, given $\alpha \in \mathcal{A}_v$, we write $s \xrightarrow{\alpha} t$ just if there exist s', t' such that $s \Longrightarrow s' \xrightarrow{\alpha} t' \Longrightarrow t$.

The standard notion of equivalence of transition systems is bisimilarity.

Definition 2 (Bisimilarity). Let S, T be LTSs, with initial states s_0, t_0 , respectively. A simulation from S to T is a relation $\mathcal{R} \subseteq |S| \times |T|$ such that $(s_0, t_0) \in \mathcal{R}$ and, for all $(s, t) \in \mathcal{R}$, we have:

1. if $s \xrightarrow{\alpha} s'$ with $\alpha \in A_v$, then there exists t' such that $t \xrightarrow{\alpha} t'$ and $(s', t') \in \mathcal{R}$; 2. if $s \xrightarrow{\tau} s'$, then there exists t' such that $t \Longrightarrow t'$ and $(s', t') \in \mathcal{R}$.

A simulation \mathcal{R} from \mathcal{S} to \mathcal{T} is a bisimulation if $\mathcal{R}^{\text{op}} = \{(t, s) \in |\mathcal{T}| \times |\mathcal{S}| \mid (s, t) \in \mathcal{R}\}$ is a simulation from \mathcal{T} to \mathcal{S} . We define $\mathcal{S} \approx \mathcal{T}$ iff there exists a bisimulation between \mathcal{S} and \mathcal{T} . This relation is called bisimilarity.

Bisimilarity can be shown to be an equivalence relation.

For our purposes we furthermore require that the equivalence does not *introduce* divergence. Given $s \in |\mathcal{S}|$, we say that there is a divergence at s (denoted as $s \uparrow$) if there exists an infinite sequence $s \xrightarrow{\tau} s_1 \xrightarrow{\tau} s_2 \xrightarrow{\tau} \cdots$.

Definition 3 (Divergence-sensitive bisimilarity). We say that a (bi)simulation \mathcal{R} between S and \mathcal{T} does not introduce divergence if, for all $(s,t) \in \mathcal{R}$, $t \uparrow implies s \uparrow$. We define divergence-sensitive bisimilarity, denoted by \approx_d , by requiring the existence of a bisimulation not introducing divergence.

For our purposes, weaker equivalences (such as *coupled simulation* [11, 12]) suffice, and might actually even be desirable. Whatever equivalence is chosen, the essential point is that it does not introduce divergence.

Definition 4 (Behavior). A behavior is $a \approx_d$ -equivalence class.

In the sequel, it will be useful to have a compact notation for describing LTS's. For this, we shall use a notation similar to the syntax of Milner's CCS [10]. For instance, if $f : \mathbb{W} \to \mathbb{W}$ is a function, $i(x).\overline{o}\langle f(x) \rangle$ denotes the LTS whose states are $\{s_0\} \cup$ $\bigcup_{\xi \in \mathbb{W}} \{s_{\xi}, t_{\xi}\} \text{ and whose transitions are } s_0 \xrightarrow{i(\xi)} s_{\xi} \text{ and } s_{\xi} \xrightarrow{\overline{o}\langle f(\xi) \rangle} t_{\xi}, \text{ for all } \xi \in \mathbb{W}.$ This kind of LTS is used to define the behaviors that correspond to classical input/output computations.

Definition 5 (Functional behavior). In the following, we fix two channels $i \in \mathcal{I}$ and $o \in \mathcal{O}$. Let $f : \mathbb{W} \to \mathbb{W}$ be a function. The functional behavior induced by f, denoted by $\mathfrak{b}_{\mathfrak{f}}$, is the equivalence class of $i(x).\overline{o}\langle f(x) \rangle$. We denote by \mathcal{FUN} the set of all functional behaviors.

Lemma 1. Let $f, g : \mathbb{W} \to \mathbb{W}$. Then, f = g iff $\mathfrak{b}_f = \mathfrak{b}_g$.

Definition 6 (Language of a functional behavior). By Lemma 1, every functional behavior $\mathfrak{b} \in \mathcal{FUN}$ determines a unique function funb on \mathbb{W} such that $\mathfrak{b} = \mathfrak{b}_{funb}$. This *induces a language (i.e., a subset of* \mathbb{W}) lang $\mathfrak{b} = \{\xi \in \mathbb{W} \mid \operatorname{fun}\mathfrak{b}(\xi) = \varepsilon\}.$

3 **Abstract Cost Models for Interactive Computation**

In order to define the complexity of behaviors, we need to add concurrency and causality information to keep track of the dependencies of outputs on relevant, "previous" inputs and to identify independent "threads" of computation in a parallel algorithm. There are several models of concurrency in the literature (see [14] for an overview of standard approaches). Asynchronous transition systems [14], which are an extension of the well known model of Mazurkiewicz traces [9], are sufficiently expressive for our purposes. In order to speak about complexity, we shall add a notion of weight: on transitions, for time complexity, and on states, for space complexity. This justifies our choice of asynchronous transition systems, which have an explicit notion of state, over the *a priori* simpler model of Mazurkiewicz traces.

Definition 7 (Asynchronous LTS [14]). An asynchronous LTS (ALTS) is a tuple S = $(|\mathcal{S}|, s_0, \mathcal{E}(\mathcal{S}), \operatorname{trans}_{\mathcal{S}}, \gamma_{\mathcal{S}})$ where $|\mathcal{S}|$ is a set of states, $s_0 \in |\mathcal{S}|$ is the initial state, $\mathcal{E}(\mathcal{S})$ is a set of event types, trans $_{\mathcal{S}} \subseteq |\mathcal{S}| \times \mathcal{E}(\mathcal{S}) \times |\mathcal{S}|$ is the transition relation and γ_{S} is an antireflexive, symmetric relation on $\mathcal{E}(S)$, called independence relation, such that (using the notations of Definition 1):

- 1. $a \in \mathcal{E}(\mathcal{S})$ implies $s \xrightarrow{a} t$ for some $s, t \in |\mathcal{S}|$;
- 2. $s \xrightarrow{a} s'$ and $s \xrightarrow{a} s''$ implies s' = s''; 3. $a_1 \frown a_2$ and $s \xrightarrow{a_1} s_1$, $s \xrightarrow{a_2} s_2$ implies $\exists t \in |S|$ s.t. $s_1 \xrightarrow{a_2} t$ and $s_2 \xrightarrow{a_1} t$; 4. $a_1 \frown a_2$ and $s \xrightarrow{a_1} s_1 \xrightarrow{a_2} t$ implies $\exists s_2 \in |S|$ s.t. $s \xrightarrow{a_2} s_2 \xrightarrow{a_1} t$.

In complete analogy to the definitions for Mazurkiewicz traces, we have trace equivalence classes of transition sequences in ALTSs and we define events with the expected causality relation that relates them.

Definition 8 (**Run, trace equivalence, event, causal order**). A run *in an* ALTS *S is a finite, possibly empty sequence of consecutive transitions* $\varphi = s \xrightarrow{a_1} \cdots \xrightarrow{a_n} t$, which we denote by $s \xrightarrow{\varphi} t$. Concatenation of runs is denoted by juxtaposition. Trace equivalence, denoted by \sim , is the smallest equivalence relation on runs such that, for all $a_1 \cap a_2$, if $\varphi = s' \xrightarrow{\varphi'} s \xrightarrow{a_1} s_1 \xrightarrow{a_2} t \xrightarrow{\varphi''} t'$ and $\psi = s' \xrightarrow{\varphi'} s \xrightarrow{a_2} s_2 \xrightarrow{a_1} t \xrightarrow{\varphi''} t'$ with s, s_1, s_2, t as in point 3 of Definition 7, then $\varphi \sim \psi$. We define a preorder between runs by $\varphi \lesssim \psi$ iff $\psi \sim \varphi\varphi'$ for some run φ' . A run φ is essential if it is of the form $s_0 \xrightarrow{\varphi'} s \xrightarrow{a} t$, with s_0 the initial state of *S*, and for all $\psi \sim \varphi$, we have $\psi = s_0 \xrightarrow{\psi'} s \xrightarrow{a} t$ with $\psi' \sim \varphi'$.

An event is a \sim -equivalence class of essential runs. We denote by Ev(S) the set of events of S; it is a poset under the quotient relation \leq /\sim , which we denote by \leq and call causal order. Note that, if $e \in Ev(S)$, all $\varphi \in e$ "end" with the same transition; we denote by evtype(e) the event type of this transition.

Finally we can add data for "time consumption" of event types and the "size" of states, which allow to define the time and space cost of events.

Definition 9 (Weights). A weighted ALTS (WALTS) is a triple (S, w_t, w_s) , where S is an ALTS, $w_t : \mathcal{E}(S) \to \mathbb{N}$ is the time weight, and $w_s : |S| \to \mathbb{N}$ is the space weight.

Let $\varphi = s_0 \xrightarrow{a_1} \cdots \xrightarrow{a_n} s_n$ be a run. Its space cost is $\operatorname{space}(\varphi) = \max_{0 \le i \le n} w_s(s_i)$. The space cost of an event $e \in \operatorname{Ev}(\mathcal{S})$ is $\operatorname{space}(e) = \max_{\varphi \in e} \operatorname{space}(\varphi)$.

Let $e \in Ev(S)$. We denote by tot(e) the set of chains of events, i.e., totally ordered subsets of $(Ev(S), \leq)$, whose maximum is e. The time cost of e is

$$\operatorname{time}(e) = \max_{X \in \operatorname{tot}(e)} \sum_{d \in X} w_t(\operatorname{evtype}(d)).$$

Roughly speaking, the space cost of events is independent of their scheduling; however, for the time cost of an event we assume an "ideal" scheduler that fully exploits all concurrency of the WALTS.

4 The Process Machine

We start by defining *string expressions* and *Boolean expressions*, which are generated by the following grammar:

$$E, F ::= x \mid \xi \mid 0(E) \mid \mathbf{1}(E) \mid tail(E)$$
$$B ::= tt \mid ff \mid 0_{?}(E) \mid \varepsilon_{?}(E),$$

where x ranges over a denumerably infinite set of *variables*, and ξ ranges over \mathbb{W} . *Processes* are defined by the following grammar:

$$P,Q ::= \mathbf{0} \mid A\langle E_1, \dots, E_n \rangle \mid \overline{O}\langle E \rangle P \mid I(x) P \mid B.(P,Q) \mid P \mid Q.$$

Table 1. The transitions of the process machine.

where O stands for either an output channel $o \in O$ or a string expression, I stands for either an input channel $i \in \mathcal{I}$ or a string expression, E, E_1, \ldots, E_n range over string expressions, and A ranges over a denumerably infinite set of *process identifiers*, each coming with an *arity* $n \in \mathbb{N}$ and a *defining equation* of the form

$$A(x_1,\ldots,x_n) \stackrel{\mathrm{def}}{=} P$$

where P is a process whose free variables are included in x_1, \ldots, x_n . As usual in process calculi, the free variables of a process (denoted by FV(P)) are defined to be the variables not in the scope of an input prefix I(x), which *binds* x. A process P is *closed* if $FV(P) = \emptyset$. In the following, all bound variables of a process are supposed to be pairwise distinct.

To assign values to expressions, we use *environments*, *i.e.*, finite partial functions from variables to \mathbb{W} . If E is a string expression whose variables are all in the domain of an environment M, we define its *value* E^M by induction: $x^M = M(x)$; $\xi^M = \xi$; $O(E)^M = OE^M$; $1(E)^M = 1E^M$; and $tail(E)^M = \xi$ if $E^M = b\xi$, with $b \in \{0, 1\}$. Similarly, we define the *value* of Boolean expressions: $tt^M = tt$; $ff^M = ff$; $O_2(E)^M =$ tt if $E^M = 0\xi$, otherwise it is ff; and $\varepsilon_2(E)^M = tt$ if $E^M = \varepsilon$, otherwise it is ff.

Definition 10 (Machine configurations, transitions). A processor state is a triple $(P, M)_p$ where P is a process, M is an environment whose domain includes FV(P), and p is a binary string, the processor tag.

A queue function is a function Θ from \mathbb{W} to finite lists of \mathbb{W} , which is almost everywhere equal to the empty list. In the following, lists of words are ranged over by q, and we denote by \cdot their concatenation.

A configuration C is a pair $[\Gamma]\Theta$, where Γ is a set of processor states whose processor tags are pairwise incompatible in the prefix order (i.e., no processor tag is the prefix of another), and Θ is a queue function.

Definition 11 (LTS of a process). Let P be a closed process. We define [P] to be the LTS generated by Table 1 with the initial state $[(P, \emptyset)_{\varepsilon}]\epsilon$ (empty environment, tag and queue function).

The reader acquainted with process algebras will note how, in spite of the presence of output prefixes in the syntax of processes, the machine treats outputs asynchronously: strings are sent (internally or externally) without waiting to synchronize with a receiver.

Given a deterministic Turing machine computing the function $f : \mathbb{W} \to \mathbb{W}$, it is possible to exhibit a closed process P such that $[P] \in \mathfrak{b}_f$; moreover, the execution of this process on the machine uses only one processor. Many more standard, "functional" models of computation can be simulated by our process machine (see Appendix A). However, the process machine is obviously richer, in the sense that it may implement more complex, "non-functional" interactive behaviors.

As announced, each transition will be given a weight, and each configuration a size. We fix a positive integer constant k and, for every string expression E and Boolean expression B, given an environment M whose domain contains the variables of E and B, we fix positive integers time_M(E) and time_M(B), representing the time it takes for a processor with environment M to compute the string E^M and the Boolean B^M . In the following we denote by $|\xi|$ the length of $\xi \in \mathbb{W}$.

Definition 12 (Weight of transitions and size of configurations). *The* weight *of a machine transition t, denoted by \$t, is defined as follows, with reference to Table 1:*

The size of an environment M is $|M| = \sum_{x \in \text{dom}(M)} |M(x)|$. If q is a list of strings, its size |q| is the sum of the lengths of the strings appearing in q; then, the size of a queue function Θ is $|\Theta| = \sum_{\xi \in \text{dom}(\Theta)} |\Theta(\xi)|$. Finally, the size of a configuration $C = [(P_1, M_1), \dots, (P_n, M_n)]\Theta$ is $|C| = |\Theta| + \sum_{i=1}^n |M_i|$.

Definition 13 (WALTS of a process). We define the set of operations as $Op = {Nil, Rec, Snd, Rcv, Out, Inp, Cnd, Spn}$. Let P be a closed process. We define a WALTS [P] as follows:

- |[P]| = |[P]|;
- the initial state is $[(P, \emptyset)_{\varepsilon}]\epsilon$;
- *E*([[*P*]]) is the set of all (*p*, *l*, *n*) ∈ W × Op × N s.t. in [*P*] there is a transition t of type *l* performed by a processor whose tag is *p* and s.t. \$t = n;
- the independence relation is the smallest symmetric relation s.t. (p, l, n) ∩ (p', l', n') holds as soon as p ≠ p' and one of the following conditions is met:
 - $l \notin \{$ Snd, Rcv, Out, Inp $\};$
 - $l \in {$ **Snd**, **Rcv** $}$ *and* $l' \in {$ **Out**, **Inp** $};$
 - $l, l' \in {$ **Snd**, **Rcv** $}$ *and the transitions concern different queues;*
 - *l*, *l'* ∈ {Out, Inp} and either *l* ≠ *l'* or the transitions concern different external channels.
- trans_[P] = { $(C, (p, l, n), C') \mid \forall (C, \alpha, C') \in \text{trans}_{[P]}$ performed by processor p of type l and weight n};

• the time weight is $w_t((p, l, n)) = n$, and the space weight is $w_s(C) = |C|$.

Note that two **Snd/Rcv** transitions on the same queue are never independent. This amounts to forbidding concurrent access to a queue, even when this could be safe. We could consider queues with concurrent access at the price of some technical complications. In this extended abstract, we prefer not to address such an arguably minor detail.

5 Complexity Classes

We now propose our definition of complexity classes of behaviors. We essentially measure the cost of producing an output as a function of all the inputs that are below it in the causal order.

Definition 14 (Input and output events, input size). Let P be a closed process. An input event (resp. output event) of $\llbracket P \rrbracket$ is an event $d \in \operatorname{Ev}(\llbracket P \rrbracket)$ s.t. $\operatorname{evtype}(d)$ is an input (resp. output) on an external channel. In the input case, if the string read is ξ , we set $|d| = |\xi|$. Let e be an output event, and let $\operatorname{Inp}(e)$ be the set of input events below e (w.r.t. the causal order). We define the input size of e as $||e|| = \sum_{d \in \operatorname{Inp}(e)} |d|$.

Definition 15 (Cost of a process). Let $f, g : \mathbb{N} \to \mathbb{N}$. We say that P works in time f and space g if for every output event e of $[\![P]\!]$, time $(e) \leq f(|\!|e|\!|)$ and space $(e) \leq g(|\!|e|\!|)$.

Definition 16 (Complexity class). Let $f, g : \mathbb{N} \to \mathbb{N}$. We define $\mathbf{BTS}(f, g)$ to be the set of behaviors \mathfrak{b} such that there exists a process P such that $[P] \in \mathfrak{b}$ and P works in time f and space g.

As sanity check we show that, in the case of functional behaviors, we essentially recover the standard complexity classes

Definition 17 (Functional complexity). Let $f, g : \mathbb{N} \to \mathbb{N}$. We define the set of languages $\mathbf{FUNTS}(f,g) = \log(\mathbf{BTS}(f,g) \cap \mathcal{FUN})$.

In the following, TIME(f) and ATIME(f) denote the standard time complexity classes (languages decidable by a deterministic and alternating Turing machine in at most f(n) steps, respectively).

Theorem 1. Let f, g be functions from \mathbb{N} to \mathbb{N} .

- 1. **TIME** $(f(n)) \subseteq$ **FUNTS**(f(n), f(n));
- 2. **FUNTS** $(f(n), g(n)) \subseteq$ **TIME** $(\mathcal{O}(f(n)g(n)^h))$ for a constant integer h > 0;
- 3. ATIME $(f(n)) \subseteq$ **FUNTS** $(f(n), 2^{\mathcal{O}(f(n))})$.

Proof. Points (1) and (3) are proved by efficiently encoding Turing machines and alternating Turing machines in the process machine (see Appendix A.1 and Appendix A.2). For point (2), we simulate with a deterministic Turing machine the execution of a process P implementing a functional behavior. This is possible because, by the properties

of \approx_d , the non-determinism that may be present during the execution of P is actually vacuous: when facing a configuration with more than one active processor, the Turing machine may simulate any one of them, without worrying about influencing the outcome or falling into infinite computations. Simulating a single transition of the process machine may be assumed to require at most $c \cdot g(n)^{h'}$ Turing machine steps, where c, h' are constant. Now, a simple combinatorial argument based on the maximum length of runs (which is f(n)) and the maximum number of active processors (which is g(n)) gives that the Turing machine halts after simulating at most f(n)g(n) transitions, yielding the desired bound. The details are given in Appendix B.

Corollary 1. Every standard polynomial or superpolynomial deterministic complexity class may be reformulated in terms of $\mathbf{FUNTS}(f, g)$. For instance:

$$\mathbf{P} = \bigcup_{k < \omega} \mathbf{FUNTS}(n^k, n^k), \qquad \mathbf{EXP} = \bigcup_{k < \omega} \mathbf{FUNTS}(2^{n^k}, 2^{n^k}).$$

Thanks to the well know equality **PSPACE** = **AP**, Theorem 1 also immediately implies the inclusion **PSPACE** $\subseteq \bigcup_{k < \omega} \mathbf{FUNTS}(n^k, 2^{n^k})$ which shows, for instance, that **NP**-complete problems may be solved in polynomial time if we allow an exponential number of processors working in parallel, as expected.

6 Discussion

The idea to revise and extend the theory of computability (and formal languages) by replacing functions with behaviors is not new [5, 8, 1, 2]. As an example, the latter works introduce reactive Turing machines (RTMs), which are ordinary Turing machines with an additional action (*i.e.*, an element of A as in Definition 1) for each transition between configurations. Each such RTM induces an LTS, which then is *executable* by definition; two RTMs execute the same behavior if their LTSs are related by a certain behavioral equivalence. Finally, so-called *effective* LTSs, i.e. LTSs with recursively enumerable transition relations, coincide with executable ones (up to behavioral equivalence). As one might expect, it is easy to construct for each RTM (without final states) a corresponding process such that their LTSs are weakly bisimilar (see Appendix C).

Not much has been said about interactive *complexity*, however. A notable exception is Japaridze's system of Clarithmetic [7], whose focus however is on logic rather than complexity theory. Another example are lineage of automata [13], whose nature is very finitistic contrarily to the one of our model.

Concerning our own work, in this extended abstract we described merely the first steps of a proposal which, at least in the case of polynomial and superpolynomial deterministic time complexity classes, has the good taste of not being inconsistent with the standard definitions. Starting from here, we have of course a great number of open questions and directions for further investigation.

First of all, in light of Theorem 1, we may ask how standard *space* complexity classes (*e.g.* **PSPACE**) may be recovered from our definitions. In this respect, we already know that, in perfect analogy with one-tape Turing machines, the sequential treatment of input strings in the current definition of the process machine prevents us

from capturing "low" complexity classes, such as L and NC. To deal with these, random access to the bits of an input string must be allowed (see Appendix A.4 and A.5).

And then, of course, there is a plethora of questions regarding non-functional behaviors, the main motivation behind our work. We did not speak of this here, because we are more concerned with establishing the main definitions and their consistency (at least in some remarkable cases), but we do have several examples of non-functional behaviors to which our definitions may be applied, including server-like behaviors and behaviors corresponding to functions on streams (such as those implemented by online algorithms). This preliminary investigation reveals that, in general, the causal dependency of WALTSS may not be enough. Instead, the behavior itself must contain additional dependency information which, for each output, specifies which inputs its cost is to be considered a function of. This is not needed in functional behaviors, because there is no choice: each output depends on the only input preceding it. For more complex behaviors, Definition 14 (the input size) is too naive and the dependency information seems to be required for interesting non-functional complexity classes to exist.

References

- Baeten, J.C.M., Luttik, B., van Tilburg, P.: Computations and interaction. In: Natarajan, R., Ojo, A.K. (eds.) ICDCIT. LNCS, vol. 6536, pp. 35–54. Springer (2011)
- Baeten, J.C.M., Luttik, B., van Tilburg, P.: Reactive Turing machines. In: Owe, O., Steffen, M., Telle, J.A. (eds.) Proceedings of FCT 2011. LNCS, vol. 6914, pp. 348–359. Springer-Verlag (2011)
- Borodin, A., El-Yaniv, R.: Online Computation and Competitive Analysis. Cambridge University Press (1998)
- Goldin, D., Smolka, S., Wegner, P.: Interactive Computation: The New Paradigm. Springer-Verlag (2006)
- Goldin, D.Q., Smolka, S.A., Wegner, P.: Turing machines, transition systems, and interaction. Electr. Notes Theor. Comput. Sci. 52(1), 120–136 (2001)
- Goldwasser, S., Micali, S., Rackoff, C.: The knowledge complexity of interactive proofsystems. In: Proceedings of STOC. pp. 291–304 (1985)
- 7. Japaridze, G.: Introduction to clarithmetic I. Inf. Comput. 209(10), 1312–1354 (2011)
- 8. van Leeuwen, J., Wiedermann, J.: On algorithms and interaction. In: MFCS. pp. 99–113 (2000)
- 9. Mazurkiewicz, A.: Trace theory. In: Petri Nets: Applications and Relationships to Other Models of Concurrency, LNCS, vol. 255, pp. 279–324. Springer (1986)
- 10. Milner, R.: Communicating and Mobile Systems: The π -calculus. Cambridge University Press (1999)
- Parrow, J., Sjodin, P.: Multiway synchrinizaton verified with coupled simulation. In: Cleaveland, R. (ed.) Proceedings of CONCUR '92. LNCS, vol. 630, pp. 518–533. Springer-Verlag (1992)
- van Glabbeek, R.J.: The linear time branching time spectrum II: The semantics of sequential systems with silent moves. In: Best, E. (ed.) Proceedings of CONCUR '93. LNCS, vol. 715, pp. 66–81. Springer-Verlag (1993)
- Verbaan, P., van Leeuwen, J., Wiedermann, J.: Complexity of evolving interactive systems. In: Theory Is Forever. pp. 268–281 (2004)
- Winskel, G., Nielsen, M.: Models for concurrency. In: Handbook of logic in Computer Science, vol. 4. Clarendon Press (1995)

A Encoding standard computational models

A.1 Turing machines

1 0

We show how the process machine can simulate a deterministic Turing machine with a constant slowdown.

Let δ be a function from \mathbb{W} to processes with free variables among \vec{x} , whose domain is *finite*. Such function may be represented by a process (identifier) $F_{\delta}(s, \vec{x})$ which progressively "explores" *s* and returns the appropriate process, returning some default process (for instance 0) in case *s* does not belong to the domain of δ . For example, if δ is defined only on ε , 0 and 1, and is a closed process in all cases, we have

$$F_{\delta}(s) \stackrel{\text{der}}{=} \varepsilon_{?}(s).(\delta(\varepsilon), \mathbf{0}_{?}(s).(\varepsilon_{?}(tail(s)).(\delta(0), \mathbf{0}), \varepsilon_{?}(tail(s)).(\delta(1), \mathbf{0}))).$$

Let now M be a deterministic Turing machine with alphabet $\{0, 1\}$ and one semiinfinite tape (*i.e.*, it is finite to the "left" and extends indefinitely to the "right"). We assume that the input string is written in the leftmost cells of the tape, the rest of the tape being covered by blank symbols. We also assume that the machine halts with a failure if it attempts to read to the left of the leftmost cell.

The configurations of the machine may be represented by three binary strings s, l, r, representing the current state, the contents to the left of the head, in reverse order and the contents of the tape to the right of the head (including the head as the first symbol of r). Then, the transition function of M induces three functions of s, depending on whether r is empty (the head "wandered off" to the right), starts with a 0, or starts with a 1. These, in turn, induce three functions δ_{ε} , δ_0 , δ_1 , from binary strings to processes, all of finite domain, which we describe as follows: let T be a process identifier of arity 3. Suppose the state of M is s, and that the current symbol is b (which may be blank). Suppose that, from this information, the transition function of M goes to state q, writes c, and moves to the right. Then, we have

$$\delta_b(s) = T\langle q, c(l), tail(r) \rangle.$$

Had the transition function of M decreed a movement to the left instead, we would have

$$\delta_b(s) = \varepsilon_?(l).(\mathbf{0}, \mathbf{0}_?(l).(T\langle q, tail(l), \mathbf{0}(c(tail(r))) \rangle, T\langle q, tail(l), \mathbf{1}(c(tail(r))) \rangle))$$

Finally, if s is a halting state, the machine may output the result, which we stipulate to be written to the right of the head (including the current position), so we have $\delta_b(s) = \overline{o} \langle r \rangle$.

But we still have to define the behavior of the process identifier T: the defining equation for it is

$$T(s,r,l) \stackrel{\text{def}}{=} \varepsilon_{?}(r).(F_{\delta_{\varepsilon}}\langle s,l,r\rangle,\mathsf{O}_{?}(r).(F_{\delta_{0}}\langle s,l,r\rangle,F_{\delta_{1}}\langle s,l,r\rangle)).$$

The Turing machine M may be represented by the process

$$i(x).T\langle \xi_0,\varepsilon,x\rangle,$$

where ξ_0 is the initial state of M.

Note how the parallel operator is never used by processes representing Turing machines. This implies that, when such processes are executed on the process machine, only one processor is used, and the causal structure of events is purely sequential. Note also that we may encode in the same way any computational model based on states whose transitions are described by a function of finite domain.

A.2 Alternating Turing machines

It is immediate to define processes corresponding to binary logical operators:

$$And(a, b, c) = a(x).b(y).\mathbf{0}_{?}(x).(\overline{c}\langle 0 \rangle, \mathbf{0}_{?}(y).(\overline{c}\langle 0 \rangle, \overline{c}\langle 1 \rangle));$$

$$Or(a, b, c) = a(x).b(y).\mathbf{0}_{?}(x).(\mathbf{0}_{?}(y).(\overline{c}\langle 0 \rangle, \overline{c}\langle 1 \rangle), \overline{c}\langle 1 \rangle).$$

Consider now an alternating Turing machine M. We suppose that at every step, M nondeterministically branches in two computations; some states will be existential (*i.e.*, will accept if the result of *one* of the two branches is accepting), while others will be universal (*i.e.*, will accept if the result of *both* of the two branches is accepting). Then, the transition function of M induces 6 functions of finite domain from strings to processes, which we call $\delta_b^i(s)$, with $b \in \{0, 1, \varepsilon\}$ and $i \in \{0, 1\}$. Intuitively, $\delta_b^i(s)$ corresponds to the behavior of M when, given that the current state and symbol are sand b, the branch i is chosen.

The functions $\delta_b^i(s)$ are defined much like in the case of deterministic Turing machines, except that now we use another process identifier N, of arity 4. For instance, if state s, symbol b, and branch i give new state q, new symbol c, and movement to the right, we have

$$\delta_b^i(s) = N\langle q, c(l), tail(r), d \rangle,$$

and so on. An important difference with deterministic Turing machines is that, in case s is an accepting state, we set $\delta_b^i(s) = \overline{d}\langle 1 \rangle$, and in case it is a rejecting state, we set $\delta_b^i(s) = \overline{d}\langle 0 \rangle$, that is, the final decision (accept/reject) is output on an internal channel d, which is a parameter of N, instead of the external output channel o.

Then, we introduce two further process identifiers T_0, T_1 , both of arity 4, and define them mutually recursively with N:

$$\begin{split} T_{i}(s,l,r,d) &\stackrel{\text{def}}{=} \varepsilon_{?}(r).(F_{\delta_{\varepsilon}^{i}}\langle s,l,r,d\rangle,\mathbf{0}_{?}(r).(F_{\delta_{0}^{i}}\langle s,l,r,d\rangle,F_{\delta_{1}^{i}}\langle s,l,r,d\rangle));\\ N(s,r,l,d) &\stackrel{\text{def}}{=} T_{0}\langle s,l,r,\mathbf{0}(d)\rangle \mid T_{1}\langle s,l,r,\mathbf{1}(d)\rangle \mid (\mathbf{0}(d))(y).(\mathbf{1}(d))(z).F_{Op}\langle s,y,z,d\rangle \end{split}$$

where $i \in \{0,1\}$ and Op(s) is the finite-domain function yielding Or(y, z, d) or And(y, z, d) according to whether s is an existential or universal state, respectively.

At this point, the alternating Turing machine M may be represented by the process

$$i(x).(N\langle \xi_0, x, \varepsilon, \varepsilon \rangle \mid \xi(y).\overline{o}\langle y \rangle),$$

where ξ_0 is the initial state.

The reader may check that, upon reception of a string x on the external input channel i, the above process starts unfolding a parallel computation whose structure

is a binary tree of depth proportional to the depth of the computation of M. Each branch in the tree executes independently from the others; once a leaf is reached, the result (acceptance/rejection) is communicated to the parent, which computes a disjunction/conjunction of the two data received from its siblings, depending on its existential/universal nature, and passes the result to its parent, and so on. The last Boolean computed, which is the final answer of M for accepting or rejecting x, is sent on channel ξ , and is forwarded to the external world through the output channel o.

A.3 Random Access Machines

A memory cell may be represented by a process which waits on a channel ξ for a string v which is interpreted as follows:

• if $v = \varepsilon$, no action is taken;

1 0

- the string v = 0v' is interpreted as a read request, and the value stored in the cell is sent using channel v';
- the string v = 1v' is interpreted as a write request, so the value v' replaces the current value.

The above process is realized by the following recursive definition:

$$C(x,v) \stackrel{\text{der}}{=} x(y).\varepsilon_?(y).(C\langle x,v\rangle, \mathsf{0}_?(y).(\overline{tail(y)}\langle v\rangle.C\langle x,v\rangle, C\langle x,tail(y)\rangle)).$$

In RAMs, memory cells contain integers, and instructions too refer to integers. Here, we use a unary representation: n is represented by the string 0^n .

A random access memory made of infinitely many cells initially containing zero, located at addresses of the form 0^n , with n > 0, is generated by the process $M\langle 0 \rangle$, where the unary process identifier M has the following defining equation $M(c) \stackrel{\text{def}}{=} C\langle c, \varepsilon \rangle |$ $M\langle 0(c) \rangle$. However, such a process is divergent, so we cannot use it for implementing functional behaviors according to Definition 5. Then, we must define a process that only creates a finite number of memory cells at a time, as needed. We first give a couple of auxiliary definitions:

$$\begin{split} D[P,Q](m,n) &\stackrel{\text{def}}{=} \varepsilon_?(m).(P,\varepsilon_?(n).(Q,D[P,Q]\langle tail(m),tail(n)\rangle))\\ E[P,Q](m,n) &\stackrel{\text{def}}{=} \varepsilon_?(m).(\varepsilon_?(n).(P,Q),\varepsilon_?(n).(Q,E[P,Q]\langle tail(m),tail(n)\rangle)) \end{split}$$

These definitions are parametric in two arbitrary processes P, Q. Given two integers m, n represented as lists of zeros, the process $D[P,Q]\langle m,n\rangle$ (resp. $E[P,Q]\langle m,n\rangle$) evaluates to P if $m \leq n$ (resp. m = n) and to Q otherwise. Then, the memory process M may be defined as follows:

$$\begin{split} M'(m,n) &\stackrel{\text{def}}{=} C\langle m, \varepsilon \rangle \mid E[\mathbf{0}, M' \langle \mathbf{0}(m), n \rangle] \langle m, n \rangle \\ M(c) &\stackrel{\text{def}}{=} \mathbf{1}(x) \cdot D[M \langle c \rangle, M' \langle \mathbf{0}(c), x \rangle \mid M \langle x \rangle] \langle x, c \rangle. \end{split}$$

In other words, $M\langle c \rangle$ waits on channel 1 for an integer x, which corresponds to the address of a memory cell. If $x \leq c$, the process returns to its initial state $M\langle c \rangle$. If

x > c, the process goes to state $M\langle x \rangle$ and, in parallel, creates x - c memory cells, each initialized to zero, at the addresses going from c + 1 to x.

A RAM program is a finite sequence of instructions, which may be represented by mutually recursively defined process identifiers I_1, \ldots, I_n , I_j standing for the *j*th instruction. These processes access the memory M with the instructions allowed by the RAM (load/store operations, possibly with indirection), and do simple arithmetic operations (increment/decrement) on the contents of a special memory cell located at the channel ε , and called the *accumulator*. Of course, before accessing the memory cell at address c, each instruction must take care of sending c on channel 1, which has the effect of creating the cell if it does not exist (and has no effect otherwise). The HALT instruction corresponds to the process

$$\overline{\varepsilon}\langle 0\xi \rangle \mid \xi(v).\overline{o}\langle v \rangle$$

which reads the value stored in the accumulator and forwards it to the external world through the output channel o (the string ξ is arbitrary, as long as it is of length at least 2 and starts with 1 to avoid unwanted interferences).

Then, such a RAM program may be represented by the process

$$i(x).(I_1 \mid C\langle \varepsilon, x \rangle) \mid M\langle 0 \rangle.$$

Note that this encoding is not quite economic in terms of parallelism: a RAM is a sequential machine, whereas executing its encoding given above on the process machine will use several processors. However, it has the advantage of being easily generalized to PRAMs (see Sect. A.5).

A.4 Boolean Circuits

The definitions of And and Or given above may be easily adapted to encode gates, from which a Boolean circuit is implemented immediately. For what concerns the interface, a circuit with m inputs and n outputs will be represented by a process reading bits from the external input channels i_1, \ldots, i_m and sending bits to the external output channels o_1, \ldots, o_n .

Note that the behavior of a process representing a Boolean circuit as above is not functional (Definition 5). Therefore, although this encoding shows how circuits may be simulated on the process machine, it does not help extending Theorem 1 to sublinear classes such as **NC**.

A.5 Parallel Random Access Machines

A PRAM is composed of several RAM programs running in parallel, each with its own accumulator. They access the same memory, including the accumulators of all other programs. At each step, the current instruction of every program is executed, and the machine proceeds to the next step only when the execution of all instructions is complete; in other words, the parallel components share a *clock*. Concurrent access to memory is resolved on a first-come-first-served basis; it is the programmer's responsibility to ensure that the cooperation between the PRAM programs is consistent.

A PRAM program with n parallel components is implemented by a process of the following form:

$$i_1(x_1).(J_1^1 \mid C\langle 0, x_1 \rangle) \mid \dots \mid i_n(x_n).(J_1^n \mid C\langle 0^n, x_n \rangle) \mid M\langle 0^{n+1} \rangle \mid K_n$$

The idea is that we put the encodings of the *n* RAM programs in parallel, plus a clock process K_n . The *i*th program has an associated internal channel, say 1^i , with which it communicates with the clock.

If I_j is the process encoding the *j*th instruction of a RAM program, the same instruction in the *i*th component of the PRAM is encoded by a process J_i^i of the form

 $1^{i}.I',$

where the input prefix 1^i bounds a variable which does not appear in I', and where I' is I in which a bogus string is sent on channel 1^i upon completion of the instruction. The encoding of the HALT instruction is defined so that the *i*th program sends the contents of its accumulator through the external output channel o_i .

The clock may then be represented by the process defined as follows:

$$K_n \stackrel{\text{def}}{=} \overline{1}....1^n.1...1^n.K_n,$$

where an output action of the form $\overline{1^i}$ means that the data sent is irrelevant.

In contrast with the representation of Boolean circuits (Sect. A.4), the above encoding of PRAMs does yield functional behaviors. However, since the initial input instruction has a linear cost in the length of the input string, no process representing a PRAM runs in sublinear time. With the present definition of functional behavior, sublinear time classes (such as **NC**) may be captured only if we modify the process machine, for example allowing random access to the bits of the input string.

B Proof of Theorem 1

Points (1) and (3) are consequences of the encoding of Appendix A.1 and Appendix A.2.

For what concerns point (2), suppose there is a process P deciding a language in time f(n) and space g(n) on the process machine. First of all, observe that, since [P] is functional, the non-determinism that may be present in P is vacuous. Indeed, all choices made during the execution of P with a given input yield the same output; moreover, since \approx_d does not introduce divergence, the execution of P terminates no matter what choice is made. Therefore, a deterministic Turing machine may simulate the process machine executing P by simulating the transitions in any order (for instance, since the coding of a configuration $[\Gamma]\Theta$ will actually represent the set Γ as a list, we may choose to always execute the transition given by the first processor of the list).

We proceed to define the Turing machine simulating the execution of P. We start by fixing an encoding $(\cdot)^{\bullet}$ of configurations of the process machine as strings of a suitable (finite) alphabet. This may safely be supposed to satisfy, for every configuration C, $|C^{\bullet}| = k|C|$, where k is a positive constant. Moreover, we suppose that C^{\bullet} has a

distinguished processor among the ones active in C. The Turing machine is initialized with a binary string ξ on its input tape, and the string $([(P, \emptyset)_{\varepsilon}]\epsilon)^{\bullet}$ on its work tape, with the distinguished processor being the only active one.

Now, at each step, the Turing machine looks at the state $(Q, M)_p$ of the distinguished processor. Depending on the shape of Q, the Turing machine simulates the appropriate transition. If Q = i(x).R (by bisimulation, the channel must be i), the Turing machine assigns to x the string placed on its input tape. If $Q = \overline{o}\langle E \rangle.R$ (by bisimulation, the channel must be o) the Turing machine halts; it accepts iff $E^M = \varepsilon$ (remember the convention used in Definition 6). In all other cases, the Turing machine simulates the necessary operations and updates the encoding of the configuration accordingly. There are only two cases worth of attention. The first one is that in which Q = E(x).R and the queue E^M is empty. Then, the Turing machine simply selects a new distinguished processor (taking care of finding one which is not blocked), and simulate the next step from it. The second one is that in which $Q = R \mid S$. In that case, after the spawning is simulated, the distinguished processor is chosen to be the one executing R.

Thanks to the properties of \approx_d (in particular the fact that it does not introduce divergence), the above Turing machine is guaranteed to terminate with the correct state of acceptance/rejection w.r.t. the input string, which we suppose to be of length n. Note that, globally, the Turing machine actually simulates a run $C_0 \xrightarrow{\tau} C_1 \xrightarrow{\tau} \cdots \xrightarrow{\tau} C_t$. Consider now the tree defined as follows: the nodes at level $0 \leq i \leq t$ represent the active processors in C_i ; the root represents the only active processor in C_0 and, at each level, the siblings of a node are either one node (the processor did not spawn) or two nodes (the processor spawned), or none at all (the processor went idle). Now, obviously each transition in the run requires an active processor; therefore, t is bounded by the size of the above tree. But such a tree has height bounded by f(n) and width bounded by g(n), so $t \leq f(n)g(n)$. The simulation of a single transition $C_i \xrightarrow{\tau} C_{i+1}$ of the run may be assumed to require at most something like $c \cdot \max(|C_i|, |C_{i+1}|)^h$ Turing machine steps, where c, h are positive constants, which is bounded by $ck \cdot g(n)^{h+1}$.

C Reactive Turing machines

We base our discussion of reactive Turing machines [2] on the following two definitions from [1] because the latter work mentions explicitly the alphabets (while the former work leaves this information implicit).

Definition 18 (Reactive Turing machine). A Reactive Turing machine *is a six-tuple* $M = (S, \mathcal{A}, \mathcal{D}, \rightarrow, \uparrow, \downarrow)$ where:

- *1. S* is a finite set of states,
- 2. A is a finite action alphabet, A_{τ} also includes the silent step τ ,
- 3. \mathcal{D} is a finite data alphabet, we add a special symbol \Box standing for a blank and put $\mathcal{D}_{\Box} = \mathcal{D} \cup \{\Box\},\$
- 4. $\rightarrow \subseteq S \times A_{\tau} \times \mathcal{D}_{\Box} \times \mathcal{D}_{\Box} \times \{L, R\} \times S$ is a finite set of transitions or steps,
- 5. $\uparrow \in S$ is the initial state,

6. $\downarrow \subseteq S$ is the set of final states.

Intuitively, the machine starts at the initial state \uparrow with an empty tape; as usual a configuration consists of the state and the tape contents (an element of \mathcal{D}_{\Box}^*) and a position of the read-write head with a single symbol. The possible transitions of a configuration depend on the state and the symbol under the head. Each transition comes with a (possibly observable) action and a change on the tape as usual.

Definition 19 (LTS of an RTM). Let $M = (S, \mathcal{A}, \mathcal{D}, \rightarrow, \uparrow, \downarrow)$ be an RTM. The labeled transition system of M, denoted by $\mathcal{T}(M)$, is defined as follows.

- 1. The set of states is the set of configurations $\{(s, \delta) \mid s \in S, \delta \text{ a tape instance}\}$.
- The transition relation → is the least relation that satisfies the following two properties for all a ∈ A_τ, d, e ∈ D_□, and δ, ζ ∈ D^{*}_□.
 - $(s, \delta \bar{d}\zeta) \xrightarrow{a} (t, (\delta) e\zeta) \text{ iff } s \xrightarrow{a[d/e]L} t,$
 - $(s, \delta \bar{d}\zeta) \xrightarrow{a} (t, \delta e(\zeta)) \text{ iff } s \xrightarrow{a[d/e]R} t.$
- *3. The initial state is* (\uparrow, \Box) *.*
- 4. $(s, \delta) \downarrow iff s \downarrow$.

We have not introduced all notation as we have the usual use of a "marker" on tape symbols, i.e., the marked symbol \bar{d} for each $d \in \mathcal{D}$.

For the encoding of RTMs into processes, we assume that the action alphabet \mathcal{A} has a suitable coding into words, i.e., that each $a \in \mathcal{A}$ has a binary code $\underline{a} \in \mathbb{W} \setminus \{\varepsilon\}$ such that \underline{a} is a palindrome⁶. Now, we shall code each action $a \in \mathcal{A}$ by the prefix $\overline{o}\langle \underline{a} \rangle$. First, we ignore termination, which is a minor point as discussed in Remark 1. Similarly, states and data symbols are assumed to have binary codes.

The encoding will use of the following auxiliary definition of internal choice.

Definition 20 (Choice at a channel). Let $\xi \in \mathbb{W}$ be a word, let $n \ge 0$ be a natural number, and let P_1, \ldots, P_n be a finite family of processes. Now, we define the process $\sum_{1\le i\le n}^{\xi} P_i$ inductively as follows.

$$\sum_{1 \le i \le 0}^{\xi} P_i = \mathbf{0}$$

$$\sum_{1 \le i \le m+1}^{\xi} P_i = \overline{\xi} \langle \mathbf{0} \rangle \cdot \mathbf{0} \mid \overline{\xi} \langle \mathbf{1} \rangle \cdot \mathbf{0} \mid \xi(x) \cdot \xi(y) \cdot [\mathbf{0}_?(y)] \cdot \left(\sum_{1 \le i \le m}^{\xi} P_i, P_{m+1}\right)$$

If ξ is omitted, it is automatically ε , i.e., $\sum_{1 \le i \le m} P_i$ stands for $\sum_{1 \le i \le m}^{\varepsilon} P_i$. We write $P_1 + \xi P_2$ for $\sum_{1 \le i \le 2}^{\xi} P_i$ and $P_1 + P_2$ for $\sum_{1 \le i \le 2} P_i$. The encoding of an RTM will be a recursive process with six string parameters:

The encoding of an RTM will be a recursive process with six string parameters: besides the state s, the symbol on the tape under the head d, and the left and right tape contents ϕ and ψ , we will use two auxiliary strings $\sigma, \rho \in W$, which "divide" ϕ and ψ into cells of suitable length. We shall also have an auxiliary processes for the

⁶ This makes the encoding simpler and causes only a constant factor in size and time.

manipulation of pairs of these strings, each of which is used to represent a stack: the first string of a pair contains the actual stack content and the other string is the 0-separated list of the unary encodings of the lengths of the elements in the stack. For example the pair $\langle 01100000, 1111011110 \rangle$ is a stack with element 0110 on top of element 0000. The complete encoding of RTMs is now as follows.

Definition 21 (Encoding RTMs without termination). Given a reactive Turing machine $M = (S, \mathcal{A}, \mathcal{D}, \rightarrow, \uparrow, \downarrow)$ with $\downarrow = \emptyset$, its encoding is the process $T\langle \underline{s}, \varepsilon, \varepsilon, \underline{\Box}, \varepsilon, \varepsilon \rangle$, with the following definitions.

$$T(s,\phi,\sigma,d,\psi,\rho) = \sum_{\substack{s \stackrel{a[d/e]_R}{\longrightarrow}t}} \overline{o}\langle\underline{a}\rangle.(Pop\langle\psi,\rho,\varepsilon\rangle \mid 0(d').0(\psi').0(\rho').T\langle\underline{t},\underline{e}\phi,\mathbf{1}^{|\underline{e}|}\mathbf{0}\sigma,d',\psi',\rho'\rangle) + \sum_{\substack{s \stackrel{a[d/e]_L}{\longrightarrow}t}} \overline{o}\langle\underline{a}\rangle.(Pop\langle\phi,\sigma,\varepsilon\rangle \mid 0(d').0(\phi').0(\sigma').T\langle\underline{t},\phi',\sigma',d',\underline{e}\psi,\mathbf{1}^{|\underline{e}|}\mathbf{0}\rho\rangle)$$

 $Pop(\xi, \theta, \upsilon) = [\varepsilon_?(\theta)]. \left(\overline{\mathsf{0}} \langle \Box \rangle . \overline{\mathsf{0}} \langle \varepsilon \rangle . \overline{\mathsf{0}} \langle \varepsilon \rangle . \mathbf{0}, [\mathsf{0}_?(\theta)]. \left(\overline{\mathsf{0}} \langle \upsilon \rangle . \overline{\mathsf{0}} \langle \xi \rangle . \overline{\mathsf{0}} \langle tail(\theta) \rangle . \mathbf{0}, P \right)\right)$

where P is

$$P = [\mathbf{0}_{?}(\theta)]. \left(Pop\langle tail(\xi), tail(\theta), \mathbf{0}\upsilon \rangle, Pop\langle tail(\xi), tail(\theta), \mathbf{1}\upsilon \rangle \right).$$

The *Pop*-process takes a pair of words – representing a stack – and an "accumulator" for the bit-wise compilation of the top element; it "returns" the top symbol and the updated "stack" via consecutive outputs on channel 0.

Proposition 1 (Simulation of RTMs). Every RTM with the empty set of final states can be simulated by a process.

Proof (idea). All non-determinism of the RTM is captured at the root of the definition of the process identifier T of Definition 21. Now, it remains only to verify that the rest of the process actually encodes the reactive Turing machine as expected. Note that there will never be any "garbage" sending actions on the "channels" 0 and ε .

Remark 1 (Termination). As for termination, without changing much, we could add a new process constant 1 to our process syntax; it would essentially behave as 0 with the sole difference that if all processes in a configuration have reduced to 1 (in an arbitrary environment and with arbitrary queues), then the configuration is a final configuration.

Effective transition systems An LTS is effective, according to [2, Definition 5], if the transition relation is recursively enumerable. Note that we again ignore termination for the sake of simplicity (cf. Remark 1). Clearly, the LTSs of processes are recursively enumerable.