

Modelling Operating System Structures by Timed Stream Processing Functions *)

Manfred Broy, Claus Dendorfer
Institut für Informatik
Technische Universität München
Postfach 20 24 20, D-8000 München 2

August 1991

Abstract

Some extensions of the basic formalism of stream processing functions are useful to specify complex structures such as operating systems. In this paper we give the foundations of higher order stream processing functions. These are functions which send and accept not only messages representing atomic data, but also complex elements such as functions. Some special notations are introduced for the specification and manipulation of such functions. A representation of time is outlined, which enables us to model time dependent behaviour. Finally, we demonstrate how characteristic operating system structures can be modelled by timed higher order stream processing functions.

*) This work was supported by the Sonderforschungsbereich 342
Werkzeuge und Methoden für die Nutzung paralleler Architekturen.

1. Introduction

In computer science formal techniques show their full strength when applied to complex system structures. This is because such structures often exhibit peculiar characteristics due to technical and application dependent reasons. For instance, when specifying operating systems, concepts like time, resources, individual processes and their scheduling have to be modelled explicitly. Moreover, programs and processes occur both as data to be handled and as algorithms to be executed.

Higher order agents are communicating entities that do not only exchange atomic data, but also send and receive agents. A number of suggestions have been made to incorporate the passing of processes as messages into the well-known calculi of communicating systems such as CSP and CCS [AR87, HC90, MPW89, Nie89, Tho89]. We prefer to follow the concept of functional system models as outlined in [Bro90] and extend this approach to ‘higher order messages’. Higher order functions are a common concept in functional programming.

Higher order functions can be used to specify communicating agents succinctly, and they provide appropriate models for particular system structures which can be found, for instance, in operating systems. Operating systems show typical characteristics of distributed systems because both consist of several processes, which act in parallel and are coordinated by the exchange of messages. An operating system receives not only first order messages (like signals, data, etc.) but also programs to be executed. This can be modelled by higher order stream processing functions. Since an operating system often exhibits time dependent behaviour, the formalism used must be powerful enough to express timing aspects.

Modelling operating system structures by stream processing functions is certainly not a new idea. See [Kah74] for an early example, which anticipates many later developments. In the world of functional programming, the advent of lazy evaluators [FW76, HM76] made stream processing and interaction possible. The functional implementation of operating systems is described in [Kar81, Hen82, Jon84]. The work of Jones and Sinclair in [JS89] contains an overview of these and other approaches.

We acknowledge that full scale operating systems have already been implemented in functional programming languages. In this paper, however, we are rather interested in the *specification* of structures that typically appear in operating systems or other distributed interactive systems. We start by outlining the basic theory of higher order streams and higher order stream processing functions. Then we introduce auxiliary notations to manipulate and specify such functions. Next we deal with the functional modelling of time. Based on the introduced notions we give some small examples of how to model characteristic operating system structures.

$\langle \rangle$	empty sequence
$\langle m \rangle$	one element sequence for $m \in M$
$\langle m_1, \dots, m_n \rangle$	sequence containing n elements, for $n \in \mathbb{N}$ and $m_1, \dots, m_n \in M$
m^n	sequence $\langle m, \dots, m \rangle$ of length n , for $m \in M$ and $n \in \mathbb{N}$
$\#s$	length of a stream $s \in M^\omega$, where $\#s \in \mathbb{N} \cup \{\infty\}$
$m:s$	element $m \in M$ prefixed to a stream $s \in M^\omega$, and $\perp:s = \langle \rangle$

Figure 1: Fundamental Notations and Functions

2. Streams and Stream Processing Functions

Let M be a set with a partial order \sqsubseteq . We assume $\perp \notin M$ and write M^\perp for $M \cup \{\perp\}$. Note that \perp represents a pseudo value standing for ‘no actual value’. The order \sqsubseteq on M is extended to an order on M^\perp by defining for all $m \in M$:

$$\perp \sqsubseteq \perp \wedge \perp \sqsubseteq m \wedge m \not\sqsubseteq \perp$$

A stream over M is denoted by M^ω . It consists of the finite sequences M^* and the infinite sequences M^∞ , which can be understood as mappings from the natural numbers \mathbb{N} to M :

$$M^\omega = M^* \cup M^\infty$$

We use a few fundamental notations and functions, which are given in Figure 1. Here, the law $\perp:s = \langle \rangle$ should especially be noted. For example, the stream $\langle 1 \rangle$ can be written as $1:\langle \rangle$ and also as $1:\perp:s$ for arbitrary finite or infinite streams s . Note that this is a property of the function $\perp: _$, which is a constructor in the sense of functional programming only with respect to elements from M , and not with respect to \perp . The special element \perp must not appear in a stream. This is the reason why we need both finite and infinite streams.

The set M^ω is partially ordered by the *refined prefix order* \sqsubseteq , which is defined by the following axiom for all $x, y \in M^\omega$:

$$x \sqsubseteq y \equiv \forall i \in \mathbb{N} : x.i \sqsubseteq y.i$$

With these definitions, M^ω forms a domain, i.e., a complete partially ordered set with least element $\langle \rangle$, provided that M^\perp with the order \sqsubseteq forms a domain.

We write $M \rightarrow N$ to denote the set of all functions from a set M to a set N , and we write $M \rightarrowtail N$ to denote the set of continuous functions between two domains M and N . Both arrow operators associate to the right. The set $M \rightarrowtail N$ is a domain if we use a partial order \sqsubseteq , such that for all $f, g \in M \rightarrowtail N$:

$$f \sqsubseteq g \equiv \forall x \in M : f.x \sqsubseteq g.x$$

$_ \# _ :: M^\omega \rightarrow M^\omega \rightarrow M^\omega$	concatenation of two streams (infix notation)
$\langle \rangle \# x = x$	
$(m : x) \# y = m : (x \# y)$	for $m \in M$
$hd :: M^\omega \rightarrow M^\perp$	first element of a stream
$hd.(m : x) = m$	for $m \in M^\perp$ (this implies $hd.\langle \rangle = \perp$ for $m = \perp$)
$tl :: M^\omega \rightarrow M^\omega$	rest of a stream without first element
$tl.(m : x) = x$	for $m \in M$
$last :: M^\omega \rightarrow M^\perp$	last element of a finite non-empty stream
$last.x = \perp$	if $x = \langle \rangle \vee \#x = \infty$
$last.\langle m \rangle = m$	for $m \in M$
$last.(m : x) = last.x$	for $m \in M$, if $\#x < \infty$
$_ \odot _ :: \wp(M) \rightarrow M^\omega \rightarrow M^\omega$	filter operator (infix notation)
$S \odot (m : x) = m : (S \odot x)$	for $m \in S^\perp$ (this implies $S \odot \langle \rangle = \langle \rangle$ for $m = \perp$)
$S \odot (m : x) = S \odot x$	for $m \in M \setminus S$

Figure 2: Standard Functions on Streams

Figure 2 presents a collection of well-known functions on streams, which we will use frequently. The application of a function f to an argument x is denoted by $f.x$. Function application associates to the left, so that $f.g.x$ stands for $(f.g).x$. Function application has a higher binding power than all other infix operators, therefore $f.x \# g.y$ means $(f.x) \# (g.y)$. With the exception of $last$ and $_ \odot _$, the functions defined so far are continuous, and therefore also monotonic, with respect to the order \sqsubseteq on M^\perp and M^ω , respectively. Concatenation is continuous in its second argument only.

We have seen two domain constructing operations, namely $_^\omega$ and $_ \rightarrow _$, which may be applied repeatedly. Let M and N be domains. A *stream processing function* is an element of $M \rightarrow N$, such that $_^\omega$ has been used at least once in the construction of both M and N . For a *higher order stream processing function* we have the additional requirement that $_ \rightarrow _$ has been used at least once in the construction of either M or N , or both.

As an example of a higher order stream processing function, consider the function f , which takes a stream of stream processing functions and concatenates their initial outputs, i.e., their respective first outputs for the empty input stream:

$$f :: (M^\omega \rightarrow M^\omega)^\omega \rightarrow M^\omega$$

$$f.(g : x) = hd.(g.\langle \rangle) : f.x$$

Similar to typed and untyped λ -calculus, we distinguish between nonreflexive and reflexive function domains. A nonreflexive function domain is a domain of the form $M \rightarrow N$ that has been built by finitely many applications of domain constructing operations. The messages of a function from such a domain always have a strictly smaller order of type than the function itself. This is not so for functions from reflexive domains. As an example, consider the reflexive function domain FUN , which is defined as the least solution of the following equation. Note that we used a third domain constructing operation here, $_ \cup _$, which is similar to set union:

$$FUN = (M \cup FUN)^\omega \rightarrow (M \cup FUN)^\omega$$

FUN denotes the set of those stream processing functions which take and produce streams that may contain elements from a basic message set M and functions from FUN . For instance, the function f defined above is an element of FUN . Moreover, reflexive domains open the possibility of self-application. The use of reflexive domains, however, increases the complexity of the semantic model. We will restrict ourselves to nonreflexive function domains for the rest of this paper since these seem to be sufficient for all practical applications.

3. Operations on Stream Processing Functions

In this section we introduce a number of useful operations on functions. We start by describing how functions can generally be defined. Then we introduce notations to specify stream processing functions. Finally we treat the class of those stream processing functions which operate on labelled streams.

3.1. General Operations on Functions

We view a function f as a mathematical object for which the two sets $DOMAIN.f$ (the *domain* of f) and $RANGE.f$ (the *range* of f) are defined, such that for every $m \in DOMAIN.f$ the function application $f(m)$ yields an element of $RANGE.f$. Note that here we use the word “domain” in a different sense than in the last section.

The simplest possible function is the unique *empty function* ν described by:

$$\nu :: \emptyset \rightarrow \emptyset$$

Another simple function, which we may find useful, is the mapping from a singleton set into another one. For all a and b , the function $a \mapsto b$ is defined as follows:

$$\begin{aligned} a \mapsto b &:: \{a\} \rightarrow \{b\} \\ (a \mapsto b).a &= b \end{aligned}$$

Let M and N be arbitrary sets. For the rest of this section we will assume that a function f from M to N is given:

$$f :: M \rightarrow N$$

Two functions with disjoint domains can be joined ‘in parallel’. Let $f' :: M' \rightarrow N'$ such that $M \cap M' = \emptyset$ (or, more liberally, $f.m = f'.m$ for all $m \in M \cap M'$). Then we write $f \cup f'$ to denote the *join* of f and f' :

$$\begin{aligned} f \cup f' &:: (M \cup M') \rightarrow (N \cup N') \\ (f \cup f').m &= f.m \quad \text{for } m \in M \\ (f \cup f').m &= f'.m \quad \text{for } m \in M' \end{aligned}$$

Another operation which combines two functions is functional (sequential) composition: given a function g , where $RANGE.f = DOMAIN.g$:

$$g :: N \rightarrow R$$

Then we write $f \circ g$ for the *functional (sequential) composition* of f and g :

$$\begin{aligned} f \circ g &:: M \rightarrow R \\ (f \circ g).m &= g.f.m \end{aligned}$$

Note the order of application. Functional composition is associative.

Given a set $M' \subseteq M$, we write $f|_{M'}$ to denote the *restriction* of f to M' , i.e., the function which has domain M' and, within this domain, behaves like f :

$$\begin{aligned} f|_{M'} &:: M' \rightarrow N \\ (f|_{M'}).m &= f.m \quad \text{for } m \in M' \end{aligned}$$

Up to now we have seen three operations on functions, namely join, functional composition and restriction. These operations take one or more functions and construct another function as a result. Here we consider it essential that the domain and range of the resulting function only depend on the domains and ranges of the argument functions, and not on their graphs.

We now present an example of an operation which changes only the graph of a function but leaves its domain and range invariant. Let $m \in M$, $n \in N$. Then the function resulting from a pointwise change of f is defined by:

$$\begin{aligned}
f[m \mapsto n] &:: M \rightarrow N \\
(f[m \mapsto n]).x &= f.x \quad \text{if } x \neq m \\
(f[m \mapsto n]).m &= n
\end{aligned}$$

Note that we could also have defined $f[m \mapsto n]$ in terms of the operators given so far:

$$f[m \mapsto n] = f|_{M \setminus \{m\}} \cup (m \mapsto n)$$

3.2. Operations on Simple Stream Processing Functions

The operations introduced in the last section applied to all kinds of functions. Now, we treat concepts tailored to stream processing functions.

Consider arbitrary domains M^\perp and N^\perp , and let f be a continuous stream processing function:

$$f :: M^\omega \rightarrow N^\omega$$

For all $m \in M^\perp$ the function $f \ll m$ (' f after m ') behaves like f after consuming element m :

$$\begin{aligned}
f \ll m &:: M^\omega \rightarrow N^\omega \\
(f \ll m).x &= f.(m : x)
\end{aligned}$$

For all $n \in N^\perp$ the function $n \ll f$ (' n then f ') outputs n and then behaves like f :

$$\begin{aligned}
n \ll f &:: M^\omega \rightarrow N^\omega \\
(n \ll f).x &= n : f.x
\end{aligned}$$

As a further overloading, $s \ll f$ will be defined for all streams $s \in N^\omega$. Note that in this case the function $_ \ll _$ is not monotonic in its first argument.

$$\begin{aligned}
s \ll f &:: M^\omega \rightarrow N^\omega \\
(s \ll f).x &= s \# f.x
\end{aligned}$$

With these operations we can specify a number of well-known functions in a convenient way:

- The *map*-function, which performs pointwise application of a function to the elements of a stream. The middle line of the declaration below can be read: "The function $map.f$, after taking in a message m , outputs $f.m$ and then behaves like $map.f$ again".

$$\begin{aligned}
\text{map} &:: (M \rightarrow M) \rightarrow M^\omega \rightarrow M^\omega \\
\text{map}.f \ll m &= f.m \ll \text{map}.f \quad \text{for } m \in M \\
\text{map}.f \ll \perp &= \perp \ll \text{map}.f
\end{aligned}$$

- A simple interactive store, which always records the last data item received. The special symbol $?$ denotes a query operation.

$$\begin{aligned}
f &:: D \rightarrow (D \cup \{?\})^\omega \rightarrow D^\omega \\
f.d \ll ? &= d \ll f.d \\
f.d \ll d' &= f.d' \quad \text{for } d' \in D \\
f.d \ll \perp &= \perp \ll f.d
\end{aligned}$$

A nondeterministic interactive system is specified as a set of deterministic systems, i.e., as a predicate on stream processing functions. Every choice of a function that fulfils this predicate represents a possible behaviour of the system. The well-known merge anomaly [BA81], which arises if nondeterministic systems are specified by predicates on their input and output histories, is avoided in our approach. This has been shown in [Bro88].

The notations presented above are useful in expressing specifications. Consider, for example, a system that inserts an arbitrary but finite number of special elements \surd into a stream. This system can be characterized by the most liberal predicate which fulfils the following equation. Note that this equation has a uniquely defined weakest solution. One can read the specification as follows: “A function represents a possible behaviour of the system if it takes an input m , then outputs some elements \surd , then outputs m , and then represents another possible behaviour of the system”.

$$\begin{aligned}
P &:: (M^\omega \rightarrow (M \cup \{\surd\})^\omega) \rightarrow \mathbb{B} \\
P.f &\equiv \exists n \in \mathbb{N}; g : P.g \wedge (\forall m \in M^\perp : f \ll m = \surd^n \ll m \ll g)
\end{aligned}$$

3.3. Operations on Labelled Stream Processing Functions

Often it is convenient to consider stream processing functions with labelled input and output streams. The labels can be understood as channel or port names. Let M^\perp be a domain, which is assumed to be fixed for the rest of this section. Then a *labelled stream processing function* is an element of the domain:

$$(I \rightarrow M^\omega) \rightarrow (O \rightarrow M^\omega)$$

where I and O are two arbitrary and not necessarily disjoint sets of labels for the input and output streams, respectively. Since the arrow associates to the right we

also write $(I \rightarrow M^\omega) \rightarrow O \rightarrow M^\omega$ for the above domain. Let f be a labelled stream processing function. We write:

$$\begin{aligned} IN.f & \quad \text{for } I \\ OUT.f & \quad \text{for } O \end{aligned}$$

Moreover, we use the ternary operators $_{-}\ll_{-}?$ and $!_{-}\ll_{-}$, which are defined as follows for all input labels $i \in IN.f$, output labels $o \in OUT.f$, and data items $m \in M^\perp$:

$$\begin{aligned} f \ll i?m & \quad \text{is the function defined by } (f \ll i?m).s = f.(s[i \mapsto (m:s.i)]) \\ o!m \ll f & \quad \text{is the function defined by } (o!m \ll f).s = (f.s)[o \mapsto (m:(f.s).o)] \end{aligned}$$

These simple operations on labelled stream processing functions considerably support the process of writing specifications. For example, an agent g which zips two streams together can be defined as follows:

$$\begin{aligned} g & :: (\{x, y\} \rightarrow M^\omega) \rightarrow \{r\} \rightarrow M^\omega \\ g & \ll x?m \ll y?m' = r!m \ll r!m' \ll g \end{aligned}$$

Note that an expression like $f \ll x?m \ll y?m'$ does *not* impose any order on the processing of messages which arrive at different input ports. For any labelled function f with $\{x, y\} \subseteq IN.f$ and $x \neq y$, we have:

$$f \ll x?m \ll y?m' = f \ll y?m' \ll x?m$$

Proof: Let $s \in (IN.f \rightarrow M^\omega)$ be an arbitrary input channel state and $\{x, y\} \subseteq IN.f$ with $x \neq y$. We calculate:

$$\begin{aligned} & (f \ll x?m \ll y?m').s \\ & = f.((s[x \mapsto (m:s.x)])[y \mapsto (m':s.y)]) \quad [\text{definition of } _{-}\ll_{-}?\text{- (twice)}] \\ & = f.((s[y \mapsto (m':s.y)])[x \mapsto (m:s.x)]) \quad [\text{since } x \neq y] \\ & = (f \ll y?m' \ll x?m).s \quad [\text{definition of } _{-}\ll_{-}?\text{- (twice)}] \end{aligned}$$

Therefore the above theorem holds by extensionality. \square

Again, nondeterministic behaviour can be expressed by giving a predicate on deterministic labelled stream processing functions. Consider, for instance, the predicate *MERGE*, which characterizes the set of deterministic merge functions. *MERGE* is the weakest predicate that fulfils the following equation:

$$\begin{aligned} MERGE & :: ((\{x, y\} \rightarrow M^\omega) \rightarrow \{r\} \rightarrow M^\omega) \rightarrow \mathbb{B} \\ MERGE.f & \equiv \exists g : MERGE.g \wedge \\ & \quad (\forall m \in M^\perp : f \ll x?m = r!m \ll g \vee f \ll y?m = r!m \ll g) \end{aligned}$$

As an example of a function which fulfils the predicate *MERGE*, consider the zip function g defined above.

A basic operation on labelled stream processing functions is the renaming of port labels. Let f be a labelled stream processing function. For $i \in IN.f$ and $i' \notin IN.f$ we define $[i'/i]f$ as the function which acts like f , but with input label i renamed to i' :

$$\begin{aligned} [i'/i]f &:: ((IN.f \setminus \{i\}) \cup \{i'\}) \rightarrow M^\omega \rightarrow OUT.f \rightarrow M^\omega \\ ([i'/i]f).x &= f.(x|_{IN.f \setminus \{i\}} \cup (i \mapsto x.i')) \end{aligned}$$

We will use a symmetric notation to rename output labels, where $o \in OUT.f$ and $o' \notin OUT.f$:

$$\begin{aligned} f[o'/o] &:: (IN.f \rightarrow M^\omega) \rightarrow ((OUT.f \setminus \{o\}) \cup \{o'\}) \rightarrow M^\omega \\ (f[o'/o]).x &= (f.x)|_{OUT.f \setminus \{o\}} \cup (o' \mapsto (f.x).o) \end{aligned}$$

Again let f be a labelled stream processing function, and choose $I' \subseteq IN.f$ and $O' \subseteq OUT.f$. We assume that the result of f does not depend on any input channel from $IN.f \setminus I'$, i.e., that $f.x = f.y$ if only $x|_{I'} = y|_{I'}$. Then $f|_{O'}^{I'}$ is the restriction of f to input ports from I' and output ports from O' :

$$\begin{aligned} f|_{O'}^{I'} &:: (I' \rightarrow M^\omega) \rightarrow O' \rightarrow M^\omega \\ (f|_{O'}^{I'}).x &= (f.x')|_{O'} \quad \text{where } x' :: IN.f \rightarrow M^\omega \quad \text{such that } x'|_{I'} = x \end{aligned}$$

Port names are not only helpful for the specification of individual agents, they can also be used to describe the connections in a whole network. For functions with port names, a uniform composition operator $\|$ (called ‘connect’) can be defined. Our intuitive view of this operation is that ports with identical names should be connected. We require that there are not two output ports with the same name. Input ports which are connected to an output port can no longer be influenced by messages arriving from the environment. Assume that two labelled stream processing functions f and g are given, such that $OUT.f \cap OUT.g = \emptyset$. We write IN for $IN.f \cup IN.g$, and OUT for $OUT.f \cup OUT.g$. Then $f\|g$ is defined as the least solution of the equations:

$$\begin{aligned} f\|g &:: (IN \rightarrow M^\omega) \rightarrow OUT \rightarrow M^\omega \\ (f\|g).x &= f.(x'|_{IN.f}) \cup g.(x'|_{IN.g}) \\ &\quad \text{where } x' = x|_{IN \setminus OUT} \cup ((f\|g).x)|_{IN \cap OUT} \end{aligned}$$

Let us consider a few special cases to obtain a clearer understanding of the connect operator. First we assume that f and g are independent and feedback free, i.e.,

that all four sets $IN.f$, $IN.g$, $OUT.f$ and $OUT.g$ are pairwise disjoint. Then $f \parallel g$ is simply the parallel composition of f and g , which can be expressed by a non-recursive equation:

$$(f \parallel g).x = f.(x|_{IN.f}) \cup g.(x|_{IN.g})$$

Next we consider the case that (some of) the output channels of f are connected to g , but that there are no feedback loops. Let $IN.f$, $IN.g$, $OUT.f$ and $OUT.g$ be pairwise disjoint, with the exception that $OUT.f \cap IN.g \neq \emptyset$. Then $f \parallel g$ can be seen as the sequential composition of the agents f and g . Again, a non-recursive equation, which is derived from the definition of $_ \parallel _$ by unfolding, is sufficient to express this:

$$(f \parallel g).x = f.(x|_{IN.f}) \cup g.((f.(x|_{IN.f}))|_{IN.g \cap OUT.f} \cup x|_{IN.g \setminus OUT.f})$$

If we furthermore require $OUT.f = IN.g$, then we even obtain:

$$(f \parallel g)|_{OUT.g}^{IN.f} = f \circ g$$

Finally, the special case of a simple (one-channel) feedback loop is treated. Let $IN.f \cap OUT.f = \emptyset$, and take $i \in IN.f$, $o \in OUT.f$. We write $f[i \leftarrow o]$ to denote the agent obtained from f by connecting output port o to input port i . The definition of $f[i \leftarrow o]$ uses the empty function ν as a dummy argument for the binary connect operator:

$$\begin{aligned} f[i \leftarrow o] &:: ((IN.f \setminus \{i\}) \rightarrow M^\omega) \rightarrow OUT.f \rightarrow M^\omega \\ f[i \leftarrow o] &= ([o/i]f \parallel \nu)|_{OUT.f}^{IN.f \setminus \{i\}} \end{aligned}$$

As an alternative, $f[i \leftarrow o]$ can also be described as the least function which fulfils the following system of equations:

$$(f[i \leftarrow o]).x = y \quad \text{where} \quad y = f.(x \cup (i \mapsto y.o))$$

An example of a network of labelled stream processing functions will be given in section 5.

4. Timed Stream Processing Functions

The modelling of time in an interactive system requires a number of design decisions on the particular aspects which should be expressed in the model. Our approach is very simple, but it nevertheless is powerful enough to support the description of a number of interesting time dependent interactive systems.

4.1. Basic Definitions

In our model we assume a global discrete time, where in every time interval at most one message can be sent or received. This may be understood as a restriction of the maximum data transfer speed of the channels. Each element of a timed stream represents the (single) communication event on a channel during one time interval. If the i -th element is a message $m \in M$, then m has been sent onto the channel during the i -th time interval. If the i -th element is the special symbol \surd (called ‘tick’), then no proper message has been sent. We require $\surd \notin M$ and write M^\surd for $M \cup \{\surd\}$. A *timed stream* is an element of $(M^\surd)^\omega$, for which we write M^τ . Such a stream represents the (finite or infinite) timed history of the communication on some channel.

For the rest of this section we will assume that the partial order \sqsubseteq on M^\surd is the identity relation, i.e., that there are not two elements in M^\surd which are non-equal but comparable. This implies that $(M^\surd)^\perp$ is a flat domain, and that the refined prefix order \sqsubseteq on M^τ actually is the ‘standard’ prefix relation.

Since our model reflects a global notion of time we require that the output history is fixed for a time interval of length n as soon as the input history is known for this interval. Hence a *timed stream processing function* is a continuous function f which operates on timed streams and has the *time progress property*:

$$\begin{aligned} f &:: M^\tau \rightarrow N^\tau \\ \forall x \in M^\tau &: \#f.x \geq \#x \end{aligned}$$

We call f a *synchronous function* if in every time interval the output for exactly this interval is determined, i.e., if input and output always have the same length:

$$\forall x \in M^\tau : \#f.x = \#x$$

Similarly, f is called an *output delayed function* if at least a bit of its future behaviour is known in advance for all finite input streams:

$$\forall x \in M^\tau : \#f.x > \#x \vee \#f.x = \infty$$

Our intuitive understanding is that output delayed functions model systems which have a certain delay or propagation time.

The above concepts will also be applied to *timed labelled stream processing functions*. For this class of functions, the time progress property has the following form. Note that infinite output must be produced by a timed function f with $IN.f = \emptyset$. This is an immediate consequence of the definition below if we state that $\min \emptyset = \infty$.

$$\forall x \in (IN.f \rightarrow M^\tau); o \in OUT.f : \#(f.x).o \geq \min\{\#x.i \mid i \in IN.f\}$$

A timed labelled stream processing function is synchronous if each output stream has exactly the length of the shortest input stream:

$$\forall x \in (IN.f \rightarrow M^\tau); o \in OUT.f : \#(f.x).o = \min\{\#x.i \mid i \in IN.f\}$$

Finally, a timed labelled stream processing function is output delayed if each output stream is strictly longer than the shortest input stream, or infinite:

$$\forall x \in (IN.f \rightarrow M^\tau); o \in OUT.f : \#(f.x).o > \min\{\#x.i \mid i \in IN.f\} \vee \#(f.x).o = \infty$$

When specifying timed systems, we have to be careful to establish the time progress property. For instance, the standard function tl is not a timed function since it decreases the length of its input stream. A timed version of tl would replace the input stream's first element with a time tick. Our intuitive understanding is that ttl is idle during the first time interval, and from then on copies its input to the output channel:

$$\begin{aligned} ttl &:: M^\tau \rightarrow M^\tau \\ ttl &= \surd \ll tl \end{aligned}$$

Sometimes we want to write a specification in which some output is produced step by step. Then the following operations \mathfrak{S} (for ‘initial’) and \mathfrak{R} (for ‘remainder’) can be employed:

$$\begin{aligned} \mathfrak{S}.f &= hd.(f.\langle \rangle) \\ \mathfrak{R}.f &= f \circ tl \end{aligned}$$

Note that $\mathfrak{S}.f \ll \mathfrak{R}.f = f$ for all output delayed stream processing functions f . See section 5 for an example of the use of these operations.

4.2. Connecting Timed Functions

Shifting our attention from single functions to networks composed of functions, we now investigate the various composition operations. It is easy to see that sequential composition of synchronous and output delayed functions yields functions of the same respective class. Parallel composition preserves both the time progress and the output delay property, but not synchronicity, in general. The feedback of an output delayed function is output delayed again. This is an important observation, since it guarantees that a whole network of functions is output delayed if only the

component functions are. Hence output delayed networks can easily be specified in a modular way.

On the other hand, the standard feedback operator applied to a synchronous function yields as result a function which does never produce any output, so that it certainly does not fulfil the time progress property. This problem can be remedied by using explicitly ‘delayed’ operators. For example, consider an alternative feedback operator which ‘initializes’ the feedback loop with a time tick. We define $f[i \checkleftarrow o]$ for $IN.f \cap OUT.f = \emptyset$, $i \in IN.f$, $o \in OUT.f$ as the least function which solves the following equations:

$$\begin{aligned} f[i \checkleftarrow o] &:: ((IN.f \setminus \{i\}) \rightarrow M^\tau) \rightarrow OUT.f \rightarrow M^\tau \\ (f[i \checkleftarrow o]).x &= y \quad \text{where } y = f.(x \cup (i \mapsto (\checkleftarrow{:}y.o))) \end{aligned}$$

Let the function g be defined by $g(y) = f.(x \cup (i \mapsto (\checkleftarrow{:}y.o)))$. The delayed feedback operator has some interesting properties:

- (1) $f[i \checkleftarrow o]$ is synchronous for any synchronous function f .
- (2) $f[i \checkleftarrow o]$ is output delayed for any output delayed function f .
- (3) $f[i \checkleftarrow o]$ is timed (has the time progress property) for any timed function f .
- (4) The function g has a unique fixed point for any synchronous function f .

Proof of (1): Let y' be an arbitrary fixed point of g . For all $o' \in OUT.f$ we calculate:

$$\begin{aligned} \#y'.o' &= \#(f.(x \cup (i \mapsto (\checkleftarrow{:}y'.o)))) \cdot o' && \text{[definition of } g\text{]} \\ &= \min\{\#(x \cup (i \mapsto (\checkleftarrow{:}y'.o))).i' \mid i' \in IN.f\} && \text{[} f \text{ is synchronous]} \\ &= \min(\{\#x.i' \mid i' \in IN.f \setminus \{i\}\} \cup \{1 + \#y'.o\}) && \text{[range splitting]} \\ &= \min\{\#x.i' \mid i' \in IN.f \setminus \{i\}\} && \text{[see below]} \end{aligned}$$

The last step is valid because the assumption $1 + \#y'.o < \min\{\#x.i' \mid i' \in IN.f \setminus \{i\}\}$ leads to a contradiction: let $o' = o$. Then $\#y'.o = 1 + \#y'.o$ by the first three steps of the calculation above and the assumption. This implies $1 + \#y'.o = \infty$, which is a contradiction because ∞ can not be strictly smaller than any value from the range of *min*. Thus we have shown that all output streams have the same length as the shortest input stream, which implies (1) by the definition of synchronicity. Note that we took y' to be *any* fixed point of g , not necessarily the least one. This will be important in the proof of (4). \square

Proof of (2) and (3): Analog to the above. \square

Proof of (4): Note that the functions used in the definition of g are continuous, hence a least fixed point y exists. By the proof of (1) we know that $\#y.o' = \#y'.o'$ for all fixed points y' of g and all $o' \in OUT.f$. If we had $y \neq y'$ then y and y' would not be comparable, so that y could not be the least fixed point. Hence $y = y'$ must hold. It should be noted that proposition (4) is only valid because we assumed that $(M^\vee)^\perp$ is a flat domain. \square

A property corresponding to (4) does not hold for output delayed functions (and therefore neither for general timed functions). Consider, as a counterexample, the output delayed function f defined as follows:

$$\begin{aligned} f &:: (\{i, i'\} \rightarrow M^\tau) \rightarrow \{o\} \rightarrow M^\tau \\ f.x &= (o \mapsto \langle \surd \rangle) && \text{if } \#x.i \leq 2 \wedge \#x.i' = 0 \\ f.x &= (o \mapsto (\surd : x.i)) && \text{otherwise} \end{aligned}$$

Now we have that $(f[i \xleftarrow{\surd} o]).(i' \mapsto \langle \rangle) = (o \mapsto \langle \surd \rangle)$, but the infinite channel state $(o \mapsto \langle \surd, \surd, \surd, \dots \rangle)$ is a fixed point of g , too.

The general connect operator $_||_$ can also be redefined to insert explicit time ticks both into sequential composition and feedback. It depends on the physical characteristics of the modelled system whether or not the insertion of a time tick between two sequentially composed functions is appropriate. Such an operator, however, does not maintain synchronicity, i.e., it will generally turn synchronous functions into non-synchronous ones.

The exclusive use of output delayed functions seems to be a convenient way to avoid these specially defined operators. If the time scale is chosen fine enough, the causality between input and output is modelled appropriately, and more information is maintained. On the other hand, one could argue that the use of output delayed functions introduces unnecessary details (namely, explicit delay times) into a specification. This is a valid point. We will investigate how one can abstract from detailed timing information in the next section.

There is another problem that arises when composing timed stream processing functions: our basic assumption was that every element of a timed stream represented a certain fixed time interval. However, it is well possible that two timed functions have been designed to work with ‘different speeds’, i.e., on different time scales. (Note that this does not necessarily mean that the length of an interval is given in, say, microseconds). It is obvious that we can not just connect these functions without first converting them to a uniform time scale.

We will first tackle the problem of adapting single streams to different time scales.

A timed stream is easily converted to a finer time scale by inserting a certain number of ticks before each element. For every $i \in \mathbb{N}$, the function $expand.i$ performs this scaling by a factor of $i+1$:

$$\begin{aligned} expand &:: \mathbb{N} \rightarrow M^\tau \rightarrow M^\tau \\ expand.i \ll m &= \sqrt{i} \ll m \ll expand.i \end{aligned}$$

The converse operation is a bit more complicated: for every time interval represented by $i+1$ elements on the finer time scale, any one message which occurred during this interval should be produced, if at least one such message exists. We make an arbitrary decision by defining that $compress$ takes the *first* proper message from every time interval, or delivers a time tick if no message occurred in a complete interval:

$$\begin{aligned} compress &:: \mathbb{N} \rightarrow M^\tau \rightarrow M^\tau \\ compress.i.x &= \langle \rangle \quad \text{if } \#x \leq i \wedge M \odot x = \langle \rangle \\ compress.i.x &= \langle hd.(M \odot x) \rangle \quad \text{if } \#x \leq i \wedge M \odot x \neq \langle \rangle \\ compress.i.(\langle m_0, \dots, m_i \rangle \# x) & \\ &= hd.((M \odot \langle m_0, \dots, m_i \rangle) \# \langle \sqrt{\rangle}) : compress.i.x \end{aligned}$$

The following proposition can easily be proved by induction on the structure of x :

$$\forall i \in \mathbb{N}; x \in M^\tau : (expand.i \circ compress.i).x = x \quad (*)$$

It is clear that compressing a stream, i.e., adapting it to a coarser time scale, may cause many data elements to be lost if there is not enough ‘time’ between the individual messages. A sufficient (and more than necessary) condition that compressing a stream by a factor of $i+1$ will not lose any messages is that the minimal space between two messages is at least i . We define the function $minspace$ such that for a timed stream $x \in M^\tau$, $minspace(x)$ is the shortest interval of ticks between two messages in x :

$$\begin{aligned} minspace &:: M^\tau \rightarrow \mathbb{N} \cup \{\infty\} \\ minspace.x &= \min\{i \in \mathbb{N} \mid \\ &\quad \exists y \in (M^\vee)^*; m, m' \in M : y \# \langle m \rangle \# \sqrt{i} \# \langle m' \rangle \sqsubseteq x\} \end{aligned}$$

In the definition above we used the refined prefix order \sqsubseteq just as a convenient substitute for the standard prefix relation. No data will be lost by compressing a sufficiently sparse stream:

$$\forall i \in \mathbb{N}; x \in M^\tau : minspace.x \geq i \Rightarrow M \odot compress.i.x = M \odot x$$

Having defined how streams are converted between different time scales, it is now easy to state operations which perform the corresponding adjustment on functions. Here, the operator $_ \uparrow _$ (read ‘scale up by’) adapts a function to an environment which works on a finer time scale, and $_ \downarrow _$ (read ‘scale down by’) is the converse operation:

$$\begin{aligned} f \uparrow i &= \text{compress}.i \circ f \circ \text{expand}.i \\ f \downarrow i &= \text{expand}.i \circ f \circ \text{compress}.i \end{aligned}$$

With the up- and downscaling operations we may combine timed functions which are based on different time scales by speeding them up or slowing them down in an appropriate way. As an immediate consequence of (*), we obtain the proposition:

$$\forall i \in \mathbb{N}; x \in M^T : (f \uparrow i) \downarrow i = f \quad (**)$$

Note that the converse of (**), namely $(f \downarrow i) \uparrow i = f$, does not hold in general. This problem will be addressed again in the next section.

Obviously, both the upscaling and the downscaling operations preserve the time progress property. Moreover, an output delayed function remains output delayed under upscaling, and a synchronous function remains synchronous under downscaling. Note that upscaling in general does not preserve synchronicity, and downscaling does not preserve the output delay property.

4.3. Time Scale Conversion as a Means of Abstraction

An important feature of a good specification formalism is that it allows us to view a system at different levels of abstraction. One method of abstraction is to convert the behaviour of a timed function to a coarser time scale. Hence the downscaling operator $_ \downarrow _$ can be seen as an abstraction operator, too.

When applying the abstraction operator to some function f , we have to be careful that no essential information is lost. In the following we will state conditions on f , which ensure that f is a proper candidate for time abstraction.

Remember that $\text{compress}.i$ only takes the first message from an interval of length $i+1$, and discards subsequent messages. Since we do not want to lose any messages, the first informal requirement is:

- (1) If supplied with a sufficiently sparse input stream, f produces a sufficiently sparse result stream.

Now assume that the abstracted version of f receives some message m as an input. On the finer time scale, we really should consider *all* of the $i+1$ intervals (of length $i+1$) which contain i ticks and the message m , thereby obtaining a possibly nondeterministic abstraction. However, our deterministic downscaling operator just selects one of these $i+1$ intervals, namely the one which starts with m . Thus we would again lose information if this arbitrary selection actually had any effect observable in the coarser time modelling, or, in other words:

- (2) If the input stream is sufficiently sparse, then the behaviour of f , as seen on the coarser time scale, only depends on those properties of the input which can be observed on the coarser time scale, too.

We start to formalize these requirements by defining an equivalence relation \sim_i (called ‘ i -space equivalence’). Two streams $x, y \in M^T$ are i -space equivalent iff they are equal on a time scale which is coarser by a factor of $i+1$:

$$x \sim_i y \equiv \text{compress}.i.x = \text{compress}.i.y$$

A function f is called i -stable iff the following formal translation of the requirements (1) and (2) holds for all streams $x, y \in M^T$:

$$\text{minspace}.x \geq i \Rightarrow \text{minspace}.(f.x) \geq i \quad \wedge \quad (1)$$

$$(\text{minspace}.x \geq i \wedge \text{minspace}.y \geq i \wedge x \sim_i y) \Rightarrow f.x \sim_i f.y \quad (2)$$

It is easy to see that, in the context of time abstraction, f does indeed always receive sufficiently sparse input streams due to the properties of *expand*. Thus, for an i -stable function f , the downscaling $f \downarrow i$ can be understood as a proper deterministic time abstraction.

The notions introduced in this section to support time abstraction can also be used to express converse versions of last section’s propositions (*) and (**). Let f be an i -stable function. Then, for all $x \in M^T$ and $i \in \mathbb{N}$, these implications hold:

$$\text{minspace}.x \geq i \Rightarrow (\text{compress}.i \circ \text{expand}.i).x \sim_i x$$

$$\text{minspace}.x \geq i \Rightarrow ((f \downarrow i) \uparrow i).x \sim_i f.x$$

The possibility of switching back and forth between different time scales is vital both for the composition of timed functions and for the abstraction of lower level details. As we already mentioned in the last section, downscaling/abstraction need not preserve the output delay property, so that the already mentioned problems with feedback arise again. It is an interesting open question to find an abstraction method which guarantees that abstracted functions maintain the time progress property even under feedback.

4.4. The Stability of Functions

In the last two sections we saw that agents operating at different (but constant) speeds could easily be connected, and we introduced the notion of stability to denote a function which showed a similar reaction to a whole class of possible input streams. We will now merge the two concepts to characterize agents that are ‘well-behaved’ even for changing input speeds, if only certain timing constraints are met.

Consider, for example, an agent receiving keyboard input from a human operator. This agent will normally just echo its input. However, there should be some idle time after each arriving message to mask key bounces. Also a special timeout message m_t should be sent to, say, a screen saver if no proper input has arrived for some time. The function $f.1$ defined below represents such an agent with an idle time of one unit and a timeout threshold of nine units:

$$\begin{aligned}
 f &:: \mathbb{N} \rightarrow M^\tau \rightarrow M^\tau \\
 f.n \ll \surd &= \surd \ll f.(n+1) && \text{if } n \neq 9 \\
 f.n \ll \surd &= m_t \ll f.(n+1) && \text{if } n = 9 \\
 f.n \ll m &= \surd \ll f.1 && \text{for } m \in M, \text{ if } n \neq 1 \\
 f.n \ll m &= m \ll f.1 && \text{for } m \in M, \text{ if } n = 1 \\
 f.n \ll \perp &= \perp \ll f.n
 \end{aligned}$$

What does it mean that a function behaves ‘similarly’ for two input streams? The first requirement (1) is that there are the same (proper) messages in both result streams. However, this is not enough because we want to distinguish, for instance, a function that immediately outputs a message from a function which outputs the same message only as a reaction to some input. Therefore, the two input streams will be partitioned into time intervals, such that every time interval (except the first one) starts with a message, and from then on contains ticks only. We state as a second requirement (2) that the same messages are produced during corresponding time intervals. Given a function f , two streams $x, y \in M^\tau$ are *f-equivalent* (written $x \sim_f y$) iff the two conditions are satisfied:

$$\begin{aligned}
 x \sim_f y &\equiv M \odot f.x = M \odot f.y \wedge && (1) \\
 &\forall x', y' \in (M^\surd)^*; m, m' \in M : \\
 &(\#(M \odot x') = \#(M \odot y') \wedge \\
 &x' \# \langle m \rangle \sqsubseteq x \wedge y' \# \langle m' \rangle \sqsubseteq y) \Rightarrow M \odot f.x' = M \odot f.y' && (2)
 \end{aligned}$$

We want to describe functions that show a similar reaction to input streams which

have at least a spacing of i and at most a spacing of k . Here, i may be seen as the minimum required distance between two messages, and k as the maximum allowed distance. The auxiliary function *maxspace* returns the maximum number of consecutive ticks in a finite stream:

$$\begin{aligned} \text{maxspace} &:: (M^\vee)^* \rightarrow \mathbb{N} \\ \text{maxspace}.x &= \max\{i \in \mathbb{N} \mid \exists y \in (M^\vee)^* : y \# \sqrt^i \sqsubseteq x\} \end{aligned}$$

A function f is called *i,k -stable* for natural numbers i and k iff all finite streams which carry the same messages and fulfil the timing requirements are f -equivalent, i.e., iff the following implication holds for all $x, y \in (M^\vee)^*$. Here we assume that m is an arbitrary element of M . The use of m is merely a coding trick to ensure that *minspace* will consider initial and trailing sequences of ticks, too.

$$\begin{aligned} (M \odot x = M \odot y \wedge \\ \text{minspace}(\langle m \rangle \# x \# \langle m \rangle) \geq i \wedge \text{maxspace}(x) \leq k \wedge \\ \text{minspace}(\langle m \rangle \# y \# \langle m \rangle) \geq i \wedge \text{maxspace}(y) \leq k) \Rightarrow x \sim_f y \end{aligned}$$

Note that we defined *i,k -stability* using only finite streams to avoid some technical problems. This is sufficient as long as we restrict ourselves to continuous functions since such a function is ‘ *i,∞ -stable*’ iff it is *i,k -stable* for all $k \in \mathbb{N}$. There are some obvious consequences of the above definition:

- (1) Every function is *i,i -stable* for all $i \in \mathbb{N}$.
- (2) An *i,k -stable* function is also *i',k' -stable* for $i \leq i'$ and $k' \leq k$.

We can characterize a function f by giving its *stability regions*, i.e., a partition of the natural numbers such that f is *i,k -stable* (for $i \leq k$) iff i and k are in the same stability region. For example, the (single) stability region of the identity function is \mathbb{N} . The keyboard input agent $f.1$ defined above has the three stability regions $\{0\}$, $\{1, \dots, 8\}$ and $\{9, \dots\}$, corresponding to its three kinds of behaviour (deletion of input during idle time, normal operation, timeout). If we changed the agent such that repeated timeouts were produced, we would get extra stability regions representing those input streams for which one, two, three etc. timeouts occur after every keyboard input. So the concept of stability allows us to investigate the time dependent properties of a function.

5. Modelling Operating System Structures

The modelling of operating system structures is a real test for a specification method. This is because many ‘low level’ concepts like time and shared resources have to be described. As a first example, we give the definition of a simple processor. This is an interactive system that takes a stream consisting of processes (modelled by functions) and input data. An arriving data element is fed into the current process, and one element of output is produced. An arriving new task replaces the current one. We will now define a function f , which represents such a multiplexing processor. Note that we expect f ’s first argument and all incoming functions to be output delayed, otherwise f may just deliver a partial result.

$$\begin{aligned}
 f &:: (M^\tau \rightarrow M^\tau) \rightarrow ((M^\tau \rightarrow M^\tau) \cup M^\vee)^\omega \rightarrow M^\tau \\
 f.g \ll d &= \mathfrak{S}.g \ll f.(\mathfrak{R}.g \ll d) \quad \text{for } d \in M^\vee \\
 f.g \ll h &= \surd \ll f.h \quad \text{for } h \in M^\tau \rightarrow M^\tau \\
 f.g \ll \perp &= \perp \ll f.g
 \end{aligned}$$

The arrival of a new task can be viewed as an interrupt. When this happens, both the old process and any pending outputs are lost. Since this is not a desirable behaviour, we will enhance our simple processor with a waiting queue for tasks. A newly arriving process is just added to the queue. Upon termination of the current task the next waiting process is started.

We will model the termination of a process by a predicate *finished* on tasks, which is true iff a task is completed:

$$finished :: (M^\tau \rightarrow M^\tau) \rightarrow \mathbb{B}$$

For instance, we could require that every task sent a special message m_c after completing its job, and define:

$$finished.g \equiv \mathfrak{S}.g = m_c$$

Then a processor with a queue of tasks is described by the function f , where:

$$\begin{aligned}
 f &:: (M^\tau \rightarrow M^\tau)^* \rightarrow ((M^\tau \rightarrow M^\tau) \cup M^\vee)^\omega \rightarrow M^\tau \\
 f.(g:q) &= f.q \quad \text{if } finished.g \\
 f.(g:q) \ll d &= \mathfrak{S}.g \ll f.((\mathfrak{R}.g \ll d):q) \quad \text{for } d \in M^\vee, \text{ if } \neg finished.g \\
 f.q \ll h &= \surd \ll f.(q \# \langle h \rangle) \quad \text{for } h \in M^\tau \rightarrow M^\tau \\
 f.\langle \rangle \ll d &= \surd \ll f.\langle \rangle \quad \text{for } d \in M^\vee \\
 f.q \ll \perp &= \perp \ll f.q
 \end{aligned}$$




Figure 3: A Processing Unit

Up to this point, we have restricted our attention to stream processing functions that were small enough to be defined and understood without a descriptive formal specification of their behaviour. We will now show how a function can be specified by a predicate on its input and output streams. Such a descriptive specification is useful because it allows us to talk about the behaviour of an agent without anticipating its implementation. Thus we separate the two concerns of *what* an agent should do versus *how* this is achieved.

Again we consider a multiplexing processing unit, which is now separated into the processor proper and a scheduler. For simplicity, we assume that a set T of tasks is given, which carry all input and output data with them, so that there is no need for additional data items. We require that $(T^\vee)^\perp$ is a flat domain. Furthermore, a predicate *finished* is needed, which tells us whether or not a task is completely evaluated:

$$finished :: T \rightarrow \mathbb{B}$$

The different priority levels of tasks are modelled by a total pre-order \succcurlyeq , such that for $t, t' \in T$, $t \succcurlyeq t'$ holds iff t has at least the same priority level as t' . This pre-order is extended onto T^\vee by defining $t \succcurlyeq \surd$ and $\surd \not\succeq t$ for all $t \in T$. This means that the special message \surd has a strictly lower priority level than any task.

$$_ \succcurlyeq _ :: T^\vee \times T^\vee \rightarrow \mathbb{B}$$

The whole processing unit is a timed labelled stream processing function, which is defined as a network of two other functions, namely the scheduler f and the processor g (see Figure 3):

$$\begin{aligned} unit &:: (\{new\} \rightarrow T^\tau) \rightarrow \{result\} \rightarrow T^\tau \\ unit &= (f \parallel g) \Big|_{\{\{result\}\}}^{\{\{new\}\}} \end{aligned}$$

We will not describe the processor g in detail. Our intuition is that g is an output delayed agent which accepts unfinished tasks, manipulates them in some way, and

produces either finished results or tasks that require some more processing:

$$g :: (\{process\} \rightarrow T^\tau) \rightarrow \{old\} \rightarrow T^\tau$$

Our aim is to give a descriptive specification of the scheduling function f :

$$f :: (\{new, old\} \rightarrow T^\tau) \rightarrow \{result, process\} \rightarrow T^\tau$$

The scheduler is a synchronous agent which receives tasks on its input channels, stores the tasks if necessary, and sends the finished ones to the result channel and the unfinished ones back to the processor. Tasks with higher priority levels are sent first. The following list describes these requirements in a more exact, but still informal way. The requirements (4) and (5) would have to be strengthened if non-synchronous scheduler functions were permitted.

- (1) Only those tasks are produced as output which have been received as input.
- (2) No finished tasks appear on the process channel.
- (3) No unfinished tasks appear on the result channel.
- (4) If there are pending finished tasks (i.e., finished tasks on the input channels which have not been output so far), then the last output on the result channel is a task which has at least the same priority level as these.
- (5) If there are pending unfinished tasks, then the last output on the process channel is a task which has at least the same priority level as these.

Note that a specification does not have to cover all possible input cases explicitly. For example, when thinking about the scheduler, we find it convenient to restrict our attention to finite input streams with equal length. We will later see that, as soon as the scheduler's behaviour is fixed for these inputs, its behaviour for *all* input streams is uniquely defined because of monotonicity and continuity (*).

A possible translation of the informal rules given above into a formal specification is that a scheduler is a synchronous function f such that the following implication holds for all $x \in (\{new, old\} \rightarrow T^\tau)$, and for all $t \in T$. Generally, we do *not* require that there is only one function which meets a given specification. In fact there exist many scheduler functions because our specification does not restrict the order in which two pending finished (or unfinished) tasks with the same priority level appear on the output channels.

$$\begin{aligned}
& \#x.new = \#x.old \wedge \#x.new < \infty \wedge \#x.old < \infty \\
& \implies \\
& \#(\{t\} \odot x.new) + \#(\{t\} \odot x.old) \geq \\
& \quad \#(\{t\} \odot f.x.result) + \#(\{t\} \odot f.x.process) \quad \wedge \tag{1} \\
& finished.t \Rightarrow \{t\} \odot f.x.process = \langle \rangle \quad \wedge \tag{2} \\
& \neg finished.t \Rightarrow \{t\} \odot f.x.result = \langle \rangle \quad \wedge \tag{3} \\
& (pending.t \wedge finished.t) \Rightarrow last.(f.x.result) \succcurlyeq t \quad \wedge \tag{4} \\
& (pending.t \wedge \neg finished.t) \Rightarrow last.(f.x.process) \succcurlyeq t \tag{5}
\end{aligned}$$

Here we used $pending.t$ as an abbreviation of the formula:

$$\begin{aligned}
& \#(\{t\} \odot x.new) + \#(\{t\} \odot x.old) > \\
& \quad \#(\{t\} \odot f.x.result) + \#(\{t\} \odot f.x.process)
\end{aligned}$$

Now the following formal version of proposition (*) will be shown: let f and g be two scheduler functions such that $\#x'.new = \#x'.old < \infty$ implies $f.x' = g.x'$ for all input channel states x' . Then $f = g$.

Proof: Take an arbitrary input channel state $x \in (\{new, old\} \rightarrow T^T)$, and let $n = \min\{\#x.new, \#x.old\}$. First we will consider the case $n < \infty$. Let $x' = (new \mapsto \langle x.new.1, \dots, x.new.n \rangle) \cup (old \mapsto \langle x.old.1, \dots, x.old.n \rangle)$. Note that $\#x'.new = \#x'.old < \infty$, so $f.x' = g.x'$ holds. It only remains to show that $f.x' = f.x$ (and, symmetrically, $g.x' = g.x$):

$$\begin{aligned}
& \#f.x'.new = n = \#f.x.new && [f \text{ is synchronous}] \\
& \#f.x'.old = n = \#f.x.old && [f \text{ is synchronous}] \\
& f.x' \sqsubseteq f.x && [f \text{ is monotonic, and } x' \sqsubseteq x] \\
& f.x' = f.x && [\text{from lines 1-3, since } (T^\vee)^\perp \text{ is a flat domain}]
\end{aligned}$$

Thus we have shown $f.x = g.x$ for all finite input channel states x . Since f and g are continuous functions, their behaviour for the case $n = \infty$ only depends on their behaviour for finite inputs. Hence $f = g$ holds by extensionality. \square

It is an interesting observation that, because of the explicit modelling of time, all progress requirements have already been captured by rules for the finite computations, together with continuity. Since liveness conditions on infinite behaviours are notoriously hard to deal with, this is quite a positive effect.

These examples demonstrate that typical concepts of operating systems can be modelled by higher order stream processing functions. Also the proof techniques developed for functional system specification can be applied.

6. Concluding Remarks

Higher order stream processing functions are a powerful and flexible instrument for the specification and modelling of all kinds of systems including operating systems. In some cases, the explicit inclusion of time is not only necessary, but can even lead to simpler specifications.

The functional treatment of operating system structures provides a solid foundation for their specification, modelling, simulation and verification. In a first case study, the inter-process communication system of the *Multiprocessor Multitasking Kernel (MMK)* developed at the Technische Universität München has been specified with the techniques described in this paper [Den91]. Larger, more complex case studies should be the next steps to obtain additional evidence about the practicability of these techniques.

References

- [AR87] Egidio Astesiano and Gianna Reggio: SMoLCS-driven concurrent calculi. In: H. Ehrig et al. (Eds.): Proc. Conf. on Theory and Practice of Software Development, pp. 169–201. Springer 1987. LNCS 249.
- [BA81] J. D. Brock and W. B. Ackermann: Scenarios: A model of non-determinate computation. In: J. Diaz and I. Ramos (Eds.): Formalization of Programming Concepts, pp. 252–259. Springer 1981. LNCS 107.
- [Bro88] Manfred Broy: Nondeterministic data flow programs: How to avoid the merge anomaly. *Science of Computer Programming*, Vol. 10, pp. 65–85. Elsevier Publishers 1988.
- [Bro90] Manfred Broy: Functional specification of time sensitive communicating systems. In: J. W. de Bakker et al. (Eds.): Stepwise Refinement of Distributed Systems, pp. 153–179. Springer 1990. LNCS 430.
- [Den91] Claus Dendorfer: Funktionale Modellierung eines Postsystems. Draft version. Technische Universität München 1991.
- [FW76] D. P. Friedman and D. S. Wise: Cons should not evaluate its arguments. In: Proc. 3rd Colloquium on Automata, Languages and Programming, pp. 257–284. Edinburgh University Press 1976.

- [HC90] Michael R. Hansen and Zhou Chao-Chen: Specification and verification of higher order processes. In: B. Rovan (Ed.): *Mathematical Foundations of Computer Science*, pp. 322–327. Springer 1990. LNCS 452.
- [Hen82] P. Henderson: Purely functional operating systems. In: Darlington et al. (Eds.): *Functional Programming and its Applications*. Cambridge University Press 1982.
- [HM76] P. Henderson and J. H. Morris: A lazy evaluator. In: *Proc. 3rd Conference on the Principles of Programming Languages*, pp. 95–103. ACM 1976.
- [Jon84] S. B. Jones: A range of operating systems written in a purely functional style. Technical Monograph PRG-42. Programming Research Group, Oxford University 1984.
- [JS89] S. B. Jones and A. F. Sinclair: Functional programming and operating systems. *The Computer Journal*, Vol. 32, pp. 162–174. British Computer Society 1989.
- [Kah74] Gilles Kahn: The semantics of a simple language for parallel programming. In: J. L. Rosenfeld (Ed.): *Information Processing 74*, pp. 471–475. Elsevier Publishers 1974.
- [Kar81] K. Karlsson: Nebula — A functional operating system. Laboratory for Programming Methodology, Chalmers University 1981.
- [MPW89] Robin Milner, Joachim Parrow and David Walker: A Calculus of Mobile Processes, Parts I and II. LFCS Reports 89-85 and 89-86. Edinburgh University 1989.
- [Nie89] Flemming Nielson: The typed λ -calculus with first-class processes. In: E. Odijk et al. (Eds.): *Parallel Architectures and Languages Europe*, pp. 357–373. Springer 1989. LNCS 366.
- [Tho89] Bent Thomsen: A calculus of higher order communicating systems. In: *Proc. 16th Symp. on Principles of Programming Languages*, pp. 143–154. ACM 1989.