

Synchrony and Clocks in Kahn Process Networks ^a

Marc Pouzet

Univ. Paris-Sud 11

IUF

ISOR 2008, Algier

November 5, 2008

^aJoint work with Albert Cohen, Florence Plateau, Louis Mandel

Overview

- Real-time Systems and Synchronous Data-flow Languages
- Synchronous Kahn Process Networks
- Introducing logical time: *clocks*
- Checking synchrony with a dedicated type system: the *clock calculus*
- Relaxed synchrony through buffer communication
- Clock envelopes and a relaxed clock calculus

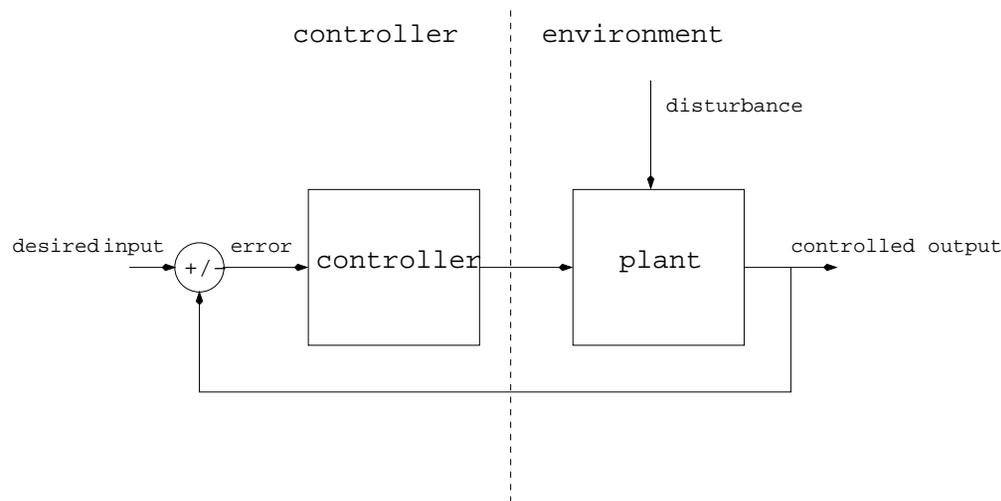
Real-time Systems

Focus on systems which continuously interact with each others.

- with a **physical environment** (e.g., fly-by-wire command, control-engine)
- or **other digital devices** (e.g., phone, TV boxes)

Real time is always **related to the environment** and is not an absolute notion. To ensure safety, think of **“what is the worst case”**?

The environment is often not precisely known: most systems run in **closed-loop**



How can we program those systems, focusing first on the **functionality**, abstracting some implementation details?

The need for High-level Programming Languages

Conciliate three notions:

- a **formal** (and computable) model of time
 - express deadlines, simultaneous events, etc.
- **parallelism** to describe complex systems from simpler ones
 - control **at the same time** rolling and pitching
 - **closed-loop** systems (the controller and the plant run in parallel)
- **statically guaranty safety properties**
 - **determinism**, dead-lock freedom
 - execution in **bounded time and memory**

Safety is important:

- critical systems: fly-by-wire, braking, airbags, etc.
- properties must be guaranteed statically: **“dynamic” = “too late”**
- build the language on a strong mathematical basis to simplify verification/validation tasks

Synchronous Data-flow Languages

Invented in the 80's to model/program critical embedded software.

The idea of Lustre:

- directly write equations over sequences as **executable specifications**
- provide a **compiler** and static analysis tools to generate code

E.g, the linear filter defined by:

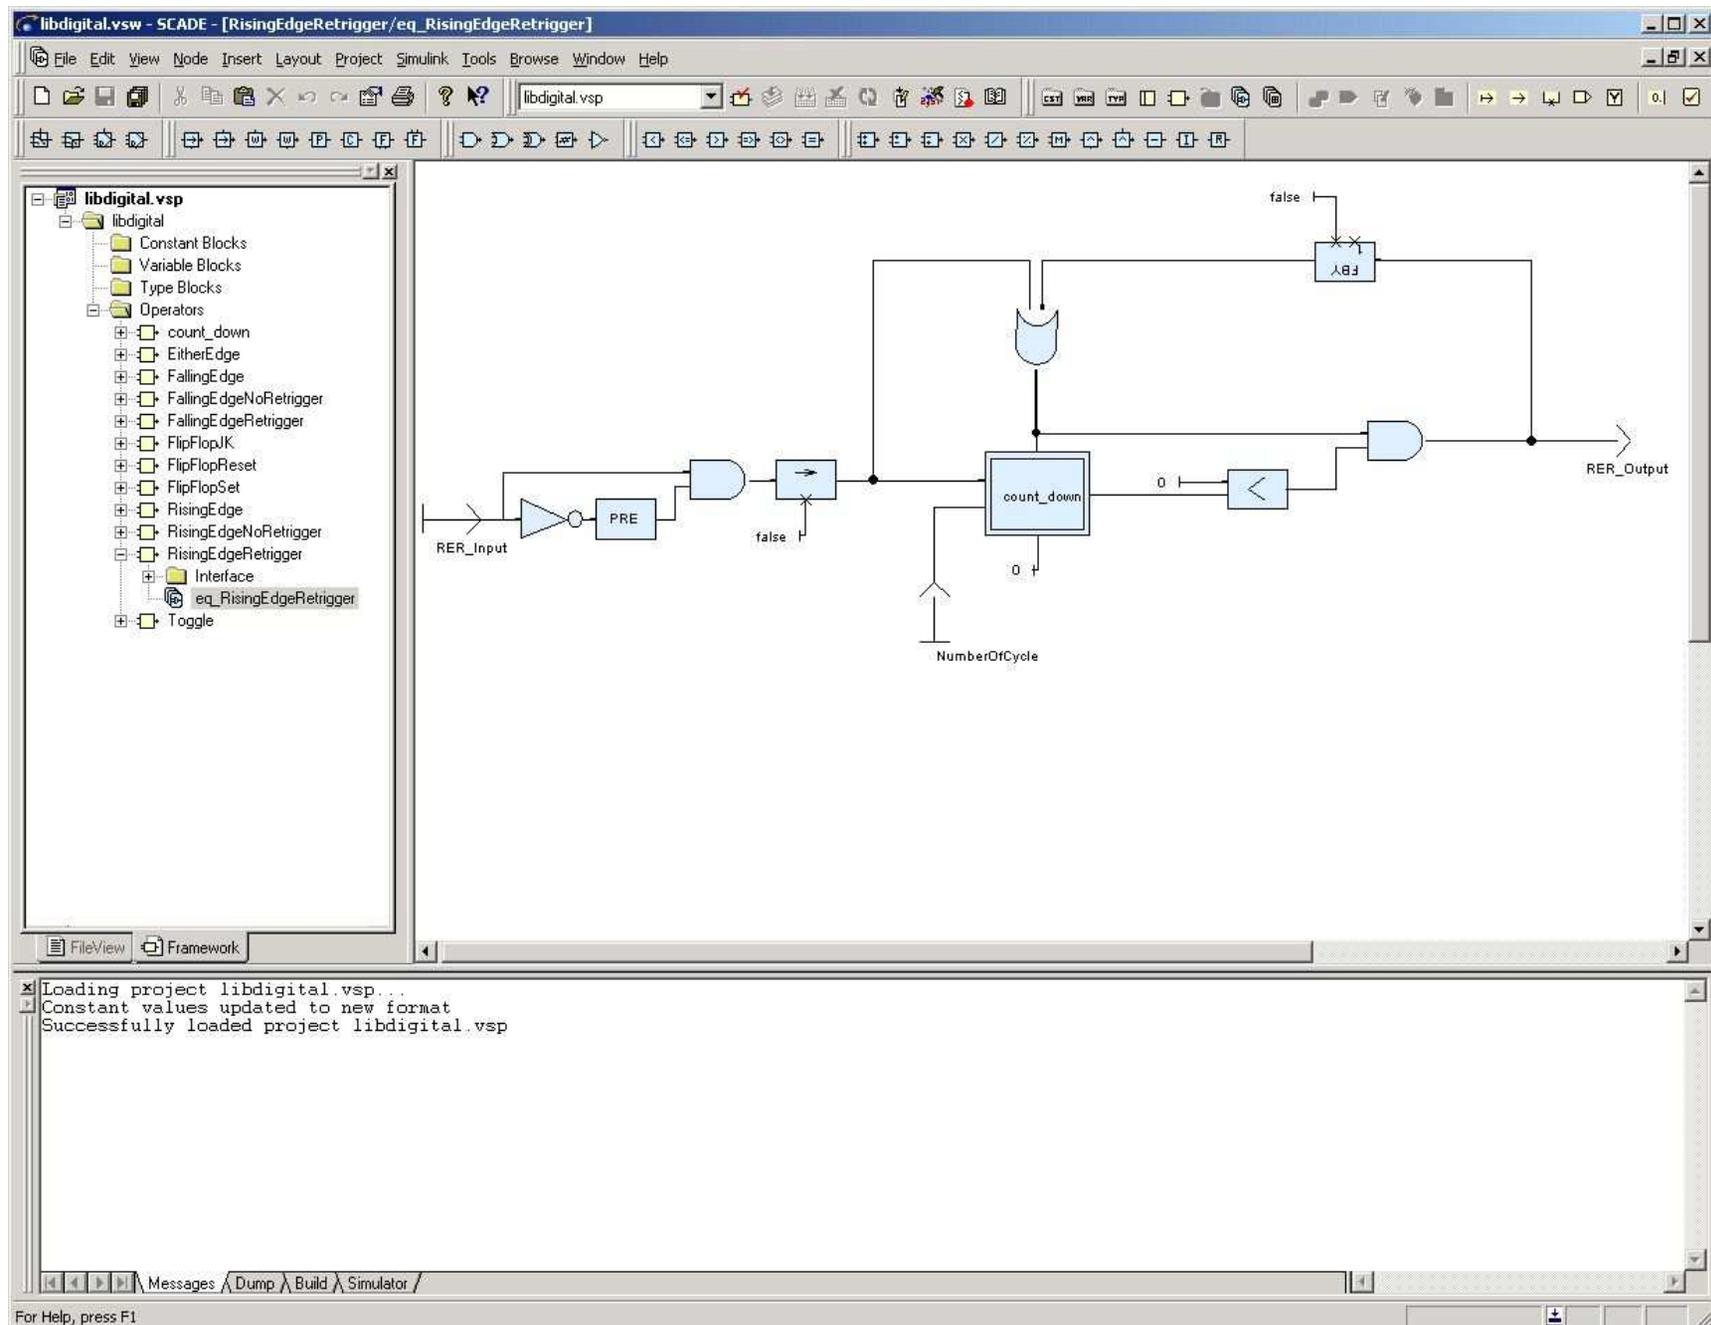
$$Y_0 = bX_0 , \forall n Y_{n+1} = aY_n + bX_{n+1}$$

is programmed by writing the equation:

$$Y = (0 \rightarrow a * \text{pre}(Y)) + b * X$$

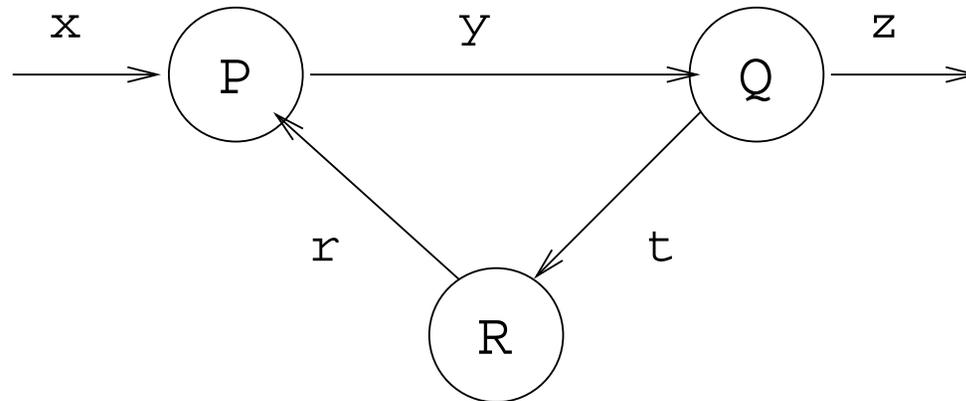
that is, we write **invariants**

An example of a SCADE specification



Kahn Process Networks [IFIP 74]

Kahn answered the following question: What is the semantics of a set of sequential processes communicating through FIFOs (e.g., Unix pipe, sockets)?



- message-based asynchronous communication (`send/wait`) through FIFOs
- reliable channels, bounded communication delays
- waiting on a single channel only. The program:

```
if (A is present) or (B is present) then ...
```

is **forbidden**

Semantics

Domain:

- $V^\infty = V^* + V^\omega$, set of finite and infinite sequences of elements in V .
- V^∞ contains the empty sequence ϵ (bottom element)
- prefix order \leq_p : for all $x \in V^\infty$, $\epsilon \leq_p x$ and for all $v \in V$, $x, y \in V^\infty$, $x \leq_p y$ iff $v.x \leq_p v.y$
- $(V^\infty, \leq_p, \epsilon)$ is a CPO.

Kahn Principle:

- **a channel** = an history of values $X = x_1, \dots, x_n, \dots \in V^\infty$
- a process = a function from an *history* of inputs to an *history* of outputs
- **causality**: a process is a **continuous function** ($f(\cup_{i=0}^\infty (x_i)) = \cup_{i=0}^\infty (f(x_i))$)

Interest/Weakness of the model

(+): Simple semantics: a process defines a function (a deterministic system); composition is functional composition; Kleene's fix-point theorem gives meaning to feedback loops

(+): Modularity: a network defines a continuous function; closed by composition and feedback

(+): Time invariance: no explicit time; semantics is invariant through slow-down/speed-up

(+): Distributed asynchronous execution: no need for a centralised scheduler

x	=	x_0		x_1		x_2		x_3		x_4		x_5		...
$f(x)$	=	y_0		y_1		y_2		y_3		y_4		y_5		...
$f(x)$	=	y_0		y_1	y_2			y_3		y_4		y_5		...

A natural model for video streaming applications (TV boxes): Sally (Philips NatLabs), StreamIt (MIT), Xstream (ST-micro) and restricted models *à la SDF* (Ptolemy)

A Small Data-flow Kernel

Consider a small language kernel with basic data-flow primitives

$$\begin{aligned} e \quad ::= \quad & e \text{ fby } e \mid op(e, \dots, e) \mid x \mid i \\ & \mid \text{merge } e \ e \ e \mid e \text{ when } e \\ & \mid \lambda x.e \mid e(e) \mid \text{rec } x.e \\ op \quad ::= \quad & + \mid - \mid \text{not} \mid \dots \end{aligned}$$

- functions $(\lambda x.e)$, application $(e(e))$, fix-point $(\text{rec } x.e)$
- constant i and variables (x)
- data-flow primitives: $x \text{ fby } y$ is the initialized delay; $op(e_1, \dots, e_n)$ the point-wise application; sampling operators (when/merge).

Data-flow Primitives

x	x_0	x_1	x_2	x_3	x_4	x_5
y	y_0	y_1	y_2	y_3	y_4	y_5
$x + y$	$x_0 + y_0$	$x_1 + y_1$	$x_2 + y_2$	$x_3 + y_3$	$x_4 + y_4$	$x_5 + y_5$
$x \text{ fby } y$	x_0	y_0	y_1	y_2	y_3	y_4
h	1	0	1	0	1	0
$x \text{ when } h$	x_0		x_2		x_4	
z		z_0		z_1		z_2
$\text{merge } h \ x \ z$	x_0	z_0	x_2	z_1	x_4	z_3

Sampling:

- if h is a boolean sequence , $x \text{ when } h$ produces a sub-sequence of x
- $\text{merge } h \ x \ z$ combines two sub-sequences

Kahn Semantics

Define a stream semantics for each data-flow primitive. E.g., if $x \mapsto s_1$ and $y \mapsto s_2$ then the value of $x + y$ is $+^\# (s_1, s_2)$

$$i^\# = i.i^\#$$

$$+^\# (s_1, s_2) = \epsilon \text{ if } s_1 = \epsilon \text{ or } s_2 = \epsilon$$

$$+^\# (x.s_1, y.s_2) = (x + y).+^\# (s_1, s_2)$$

$$\epsilon \text{ fby}^\# y = \epsilon$$

$$(x.s_1) \text{ fby}^\# s_2 = x.s_2$$

$$s_1 \text{ when}^\# s_2 = \epsilon \text{ if } s_1 = \epsilon \text{ or } s_2 = \epsilon$$

$$x.s \text{ when}^\# 1.c = x.(s \text{ when}^\# c)$$

$$x.s \text{ when}^\# 0.c = s \text{ when}^\# c$$

$$\text{merge}^\# c s_1 s_2 = \epsilon \text{ if } s_i = \epsilon$$

$$\text{merge}^\# 1.c x.s_1 s_2 = x.\text{merge}^\# c s_1 s_2$$

$$\text{merge}^\# 0.c s_1 y.s_2 = y.\text{merge}^\# c s_1 s_2$$

Property: Data-flow operators are continuous functions; a program is a continuous functions

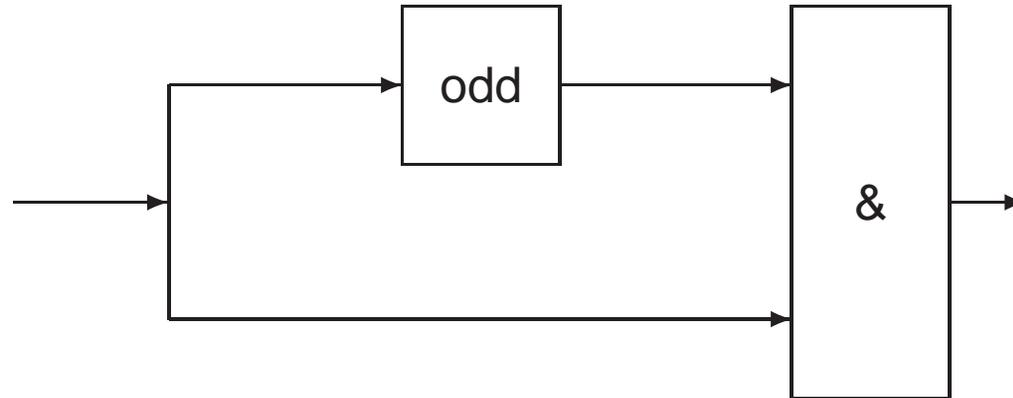
Derived operators:

- $\text{if } c \text{ then } x \text{ else } y = \text{merge } c (x \text{ when } c) (x \text{ when not } c)$

Final remark: Up to syntactic details, we can write most Lustre programs.

Synchronisation Issues

What happens when streams are sampled and composed together?



If $x = (x_i)_{i \in \mathbb{N}}$ then $\text{odd}(x) = (x_{2i})_{i \in \mathbb{N}}$ and $x \& \text{odd}(x) = (x_i \& x_{2i})_{i \in \mathbb{N}}$.

Execution with unbounded FIFOs!

Remarks:

- These programs must be detected and rejected
- each operator is finite-memory through the composition is not: all the complexity (here synchronisation) is hidden in the communication channels
- The Kahn semantics is unable to deal with time, e.g., specify that two events arrive **at the same time**

Synchronous Streams

complete the set of values with an explicit absent value abs . A signal s is a stream.

$$s : (V^{abs})^\infty$$

Clock: the clock of a stream x is a boolean stream indicating the instant where x is present

$$\mathcal{B} = \{0, 1\}$$

$$\mathcal{CLOCK} = \mathcal{B}^\infty$$

$$\text{clock } \epsilon \quad = \quad \epsilon$$

$$\text{clock } (abs.x) \quad = \quad 0.\text{clock } x$$

$$\text{clock } (v.x) \quad = \quad 1.\text{clock } x$$

Clocked Streams:

$$ClStream(V, cl) = \{s \mid s \in (V^{abs})^\infty \wedge \text{clock } s \leq_p cl\}$$

Data-flow Primitives

Constant generator:

$$i^\#(\epsilon) = \epsilon$$

$$i^\#(1.cl) = i.i^\#(cl)$$

$$i^\#(0.cl) = abs.i^\#(cl)$$

Pointwise application:

Arguments must be synchronous, i.e., they should have the same clock

$$+^\#(s_1, s_2) = \epsilon \text{ if } s_i = \epsilon$$

$$+^\#(abs.s_1, abs.s_2) = abs.+^\#(s_1, s_2)$$

$$+^\#(v_1.s_1, v_2.s_2) = (v_1 + v_2).+^\#(s_1, s_2)$$

Partial Definitions

As such, these functions are not total. What does it mean when one element is present and the other is absent?

Restrict the domain:

$$(+): \forall cl : \mathcal{CLOCK}. ClStream(\mathbf{int}, cl) \times ClStream(\mathbf{int}, cl) \rightarrow ClStream(\mathbf{int}, cl)$$

that is $(+)$ is a function which expect two integer inputs with the same clock cl and return an output with the same clock cl .

These extra conditions are **types**: programs which do not conform to these constraints are rejected.

Remark: Regular types and clock types can be specified separately:

- $(+) : \mathbf{int} \times \mathbf{int} \rightarrow \mathbf{int} \quad \leftarrow \text{its type signature}$
- $(+) :: \forall cl. cl \times cl \rightarrow cl \quad \leftarrow \text{its clock signature}$

In the sequel, we only write the clock signature.

Delays

$$\begin{aligned}\epsilon \text{ fby}^\# s &= \epsilon \\ (abs.s_1) \text{ fby}^\# (abs.s_2) &= abs.(s_1 \text{ fby}^\# s_2) \\ (v.s_1) \text{ fby}^\# (w.s_2) &= v.(fby_1^\# w s_1 s_2) \\ \text{fby}_1^\# v \in s &= \epsilon \\ \text{fby}_1^\# v (abs.s_1) (abs.s_2) &= abs.(fby_1^\# v s_1 s_2) \\ \text{fby}_1^\# v (w.s_1) (v'.s_2) &= v.(fby_1^\# v' s_1 s_2)\end{aligned}$$

As a consequence:

$$\text{fby} : \forall cl.cl \times cl \rightarrow cl$$

Sampling

$$s_1 \text{ when}^\# s_2 = \epsilon \text{ if } s_1 = \epsilon \text{ or } s_2 = \epsilon$$

$$(abs.s) \text{ when}^\# (abs.c) = abs.s \text{ when}^\# c$$

$$(v.s) \text{ when}^\# (1.c) = v.s \text{ when}^\# c$$

$$(v.s) \text{ when}^\# (0.c) = abs.x \text{ when}^\# c$$

$$\text{merge } c \ s_1 \ s_2 = \epsilon \text{ if one of the } s_i = \epsilon$$

$$\text{merge } (abs.c) \ (abs.s_1) \ (abs.s_2) = abs.\text{merge } c \ s_1 \ s_2$$

$$\text{merge } (1.c) \ (v.s_1) \ (abs.s_2) = v.\text{merge } c \ s_1 \ s_2$$

$$\text{merge } (0.c) \ (abs.s_1) \ (v.s_2) = v.\text{merge } c \ s_1 \ s_2$$

Examples

$base = (1)$	1	1	1	1	1	1	1	1	1	1	1	1	...
x	x_0	x_1	x_2	x_3	x_4	x_5	x_6	x_7	x_8	x_9	x_{10}	x_{11}	...
$h = (10)$	1	0	1	0	1	0	1	0	1	0	1	0	...
$y = x$ when h	x_0		x_2		x_4		x_6		x_8		x_{10}	x_{11}	...
$h' = (100)$	1		0		0		1		0		0	1	...
$z = y$ when h'	x_0						x_6					x_{11}	...
k			k_0		k_1				k_2		k_3		...
merge $h' z k$	x_0		k_0		k_1		x_6		k_2		k_3		...

Sampling and Clocks

- in x when[#] y , x and y must have the same clock cl
- the clock of x when[#] c is noted cl on c : it means that c moves at the pace cl

$$s \text{ on } c = \epsilon \text{ if } s = \epsilon \text{ or } c = \epsilon$$

$$(1.cl) \text{ on } (1.c) = 1.cl \text{ on } c$$

$$(1.cl) \text{ on } (0.c) = 0.cl \text{ on } c$$

$$(0.cl) \text{ on } (abs.c) = 0.cl \text{ on } c$$

We get:

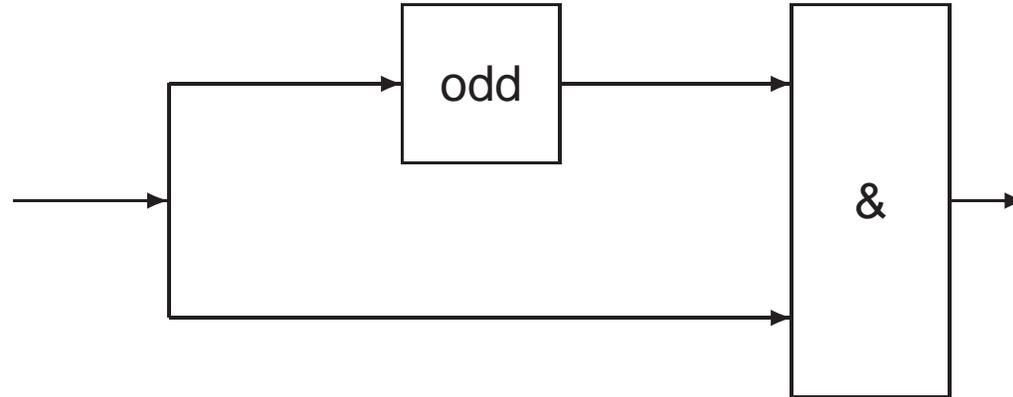
$$\text{when} : \forall cl. \forall x : cl. \forall c : cl. cl \text{ on } c$$

$$\text{merge} : \forall cl. \forall c : cl. \forall x : cl \text{ on } c. \forall y : cl \text{ on } \text{not } c.cl$$

For any clock cl , if the first input x has clock cl and the second input c has clock cl then x when c has clock cl on c .

Checking Synchrony

The previous programs is now statically rejected by the compiler.



This is essentially a **typing problem**:

```
let odd x = x when half
let non_synchronous x = x & (odd x)
                        ^^^^^^^
```

This expression has clock 'a on half,
but is used with clock 'a

In synchronous languages, we only consider **clock equality**

From pure synchrony to N -synchrony

- The comparison of clocks is limited to clock equality, i.e., “two streams are synchronous or not”
- What about comparing streams which are not exactly synchronous but “not far”?
- How to account for possible “gittering” in the system as found in video applications?
- How to model execution time?

A typical example: the Downscaler

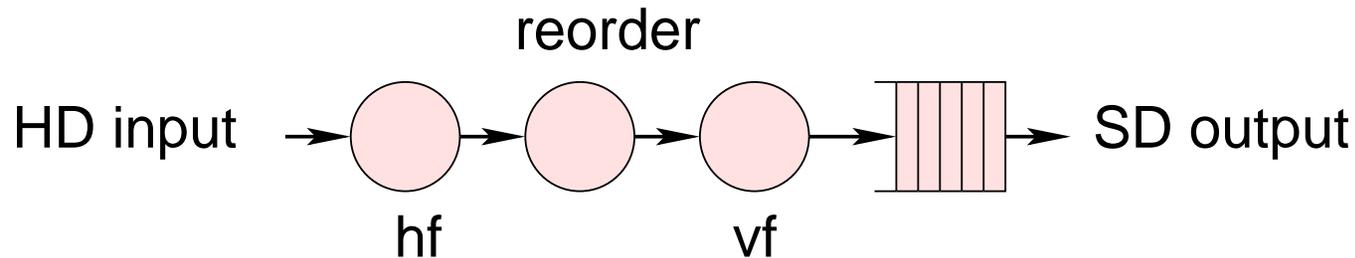
high definition (HD) → standard definition (SD)

1920×1080 pixels

720×480

horizontal filter: number of pixels in a line from 1920 pixels down to 720 pixels,

vertical filter: number of lines from 1080 down to 480



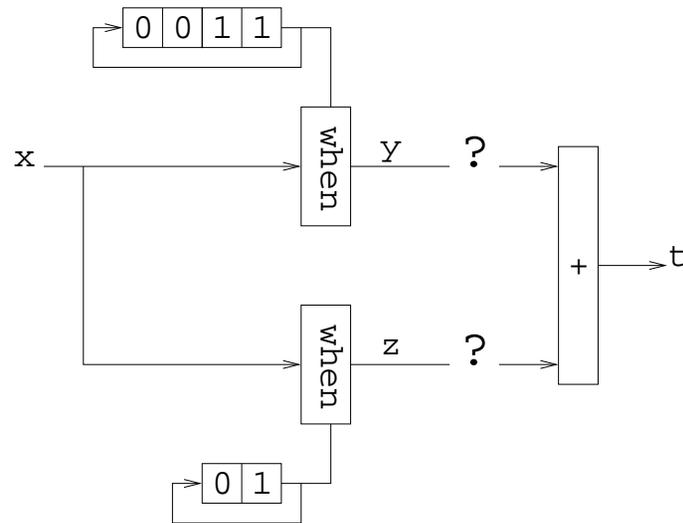
Real-Time Constraints

the input and output processes: 30Hz.

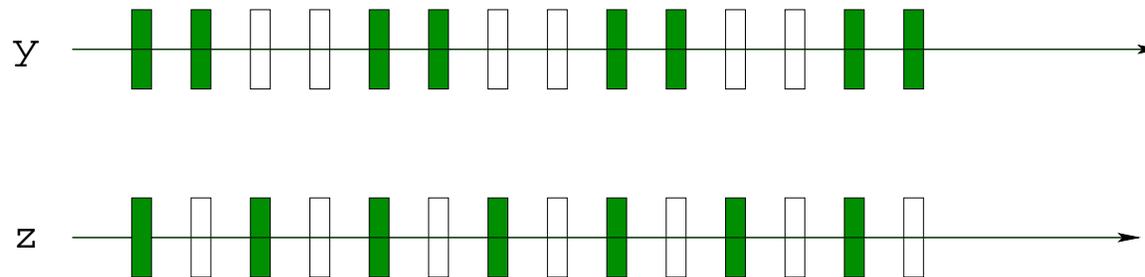
HD pixels arrive at $30 \times 1920 \times 1080 = 62,208,000 Hz$

SD pixels at $30 \times 720 \times 480 = 10,368,000 Hz$ (6 times slower)

But too restrictive for our video applications



- streams must be synchronous when composed ($y+z$ is rejected by the clock calculus)

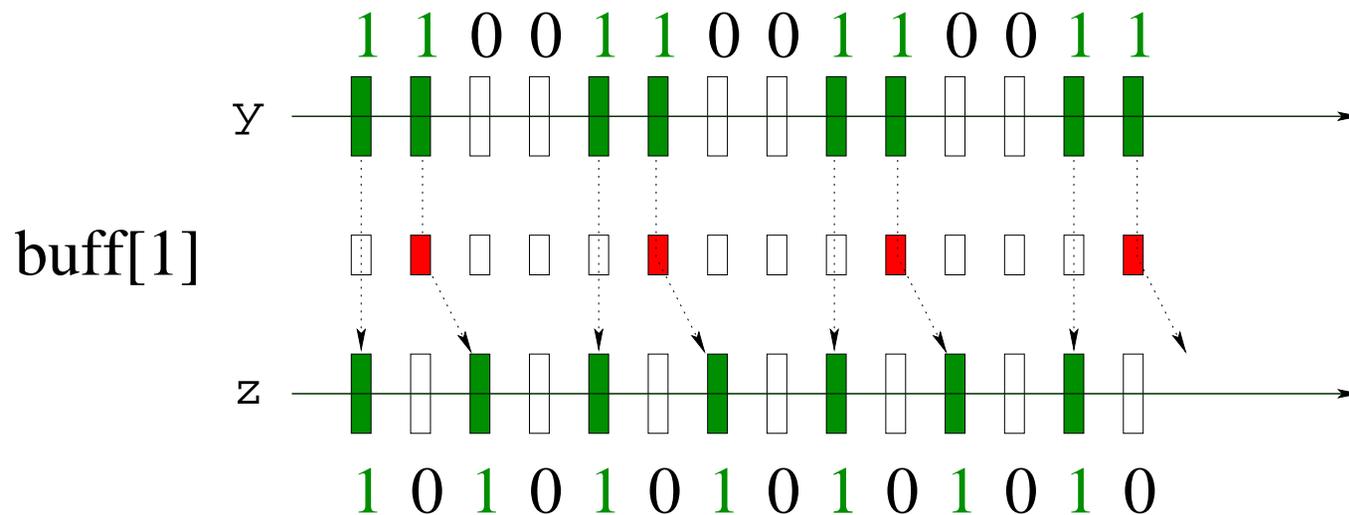


- adding buffer code (by hand) is feasible but hard and error-prone
- can we compute it automatically and obtain regular synchronous code?

we need a relaxed model of synchrony and relaxed clock calculus

N-Synchronous Kahn Networks

- propose a programming model based on a relaxed notion of synchrony
- yet compilable to some synchronous code
- allows to compose programs as soon as they can be made synchronous through the insertion of a bounded buffer



- based on the use of *infinite ultimately periodic clocks*
- a precedence relation between clocks $ck_1 <: ck_2$

Infinite Ultimately Periodic Clocks

Introduce \mathbb{Q}_2 as the set of infinite periodic binary words. Coincides with rational 2-adic numbers

$$(01) = 01\ 01\ 01\ 01\ 01\ 01\ 01\ 01\ 01\ 01\ \dots$$

$$0(1101) = 0\ 1101\ 1101\ 1101\ 1101\ 1101\ 1101\ 1101\ 1101\ \dots$$

- 1 stands for the presence of an event
- 0 for its absence

Definition:

$$w ::= u(v) \quad \text{where } u \in (0 + 1)^* \text{ and } v \in (0 + 1)^+$$

Causality order and Synchronisability

Precedence relation: $w_1 \preceq w_2$

- “1s from w_1 arrive before 1s from w_2 ”
- \preceq is a partial order which abstracts the causality order between streams
- $(\mathbb{Q}_2, \preceq, \sqcup, \sqcap)$ is a lattice

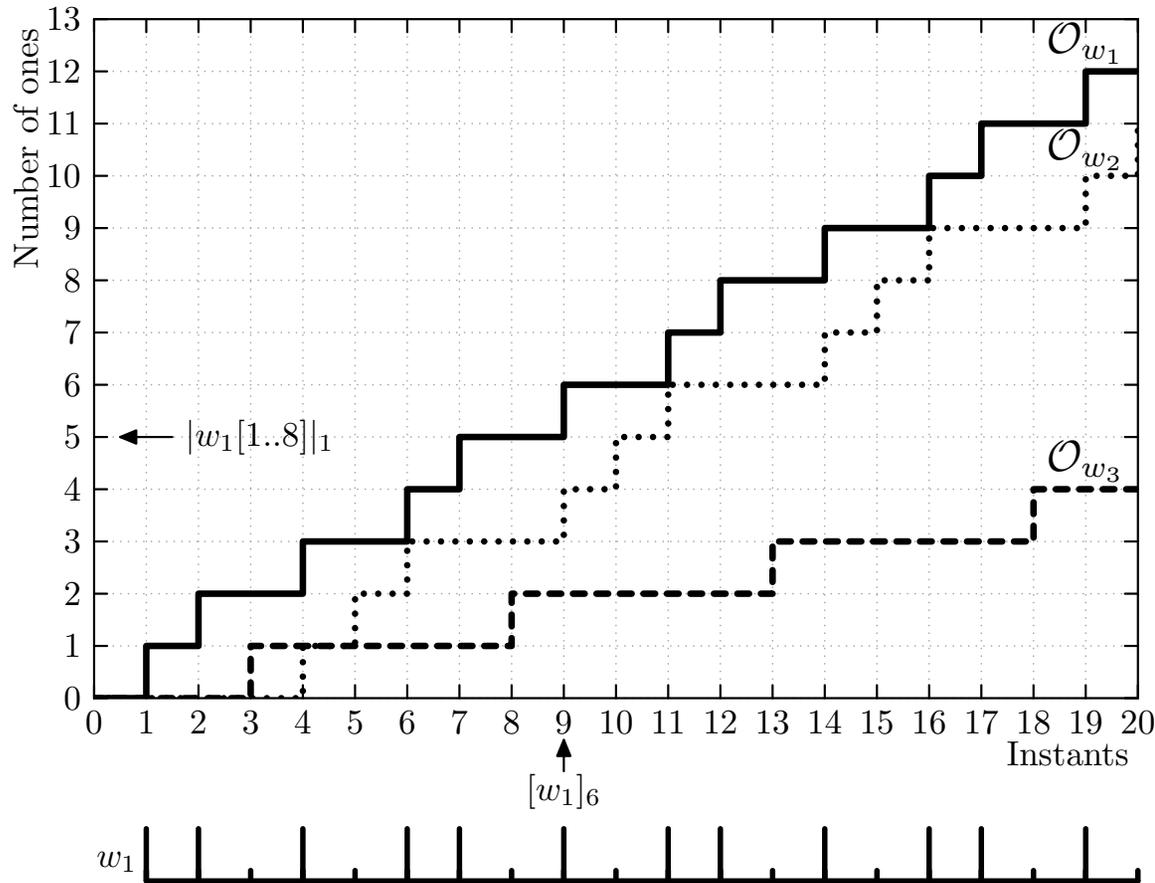
Synchronisability:

Two infinite periodic binary words w and w' are *synchronisable*, noted $w \bowtie w'$ iff it exists $d \in \mathbb{N}$ such that $w \preceq 0^d w'$ and $d' \in \mathbb{N}$ such that $w' \preceq 0^{d'} w$.

1. 11(01) and (10) are synchronisable
2. (010) and (10) are not synchronisable since there are too much reads or too much writes (infinite buffers)

Subsumption (sub-typing): $w_1 <: w_2 \iff w_1 \bowtie w_2 \wedge w_1 \preceq w_2$

Clocks represented graphically



Notations:

$w[i]$: element at index i

$[w]_j$: position of the j^{th} 1

$\mathcal{O}_w(i)$: number of 1s seen in w until index i .

buffer size $\text{size}(w_1, w_2) = \max_{i \in \mathbb{N}} (\mathcal{O}_{w_1}(i) - \mathcal{O}_{w_2}(i))$

precedence $w_1 \preceq w_2 \Leftrightarrow \forall i, \mathcal{O}_{w_1}(i) \geq \mathcal{O}_{w_2}(i)$

synchronizability $w_1 \bowtie w_2 \Leftrightarrow \exists b_1, b_2 \in \mathbb{Z}, \forall i, b_1 \leq \mathcal{O}_{w_1}(i) - \mathcal{O}_{w_2}(i) \leq b_2$

Multi-sampled Systems (clock sampling)

$$c ::= w \mid c \text{ on } w \quad w \in (0 + 1)^\omega$$

$c \text{ on } w$ denotes a *subsamped clock*.

$c \text{ on } w$ is the clock obtained in advancing in w at the pace of clock c . E.g.,

$$1(10) \text{ on } (01) = (0100).$$

base	1 1 1 1 1 1 1 1 1 1 1 ...	(1)
p_1	1 1 0 1 0 1 0 1 0 1 ...	1(10)
base on p_1	1 1 0 1 0 1 0 1 0 1 ...	1(10)
p_2	0 1 0 1 0 1 ...	(01)
(base on p_1) on p_2	0 1 0 0 0 1 0 0 0 1 ...	(0100)

Computing with Periodic Clocks

In the case of infinite periodic binary words, precedence relation, synchronizability, equality can be decided in bounded time

Synchronizability: Two infinite periodic binary words $u(v)$ and $u'(v')$ are *synchronizable*, noted $u(v) \bowtie u'(v')$ iff they have the same rate, i.e., $\frac{|v|_1}{|v'|_1} = \frac{|v|}{|v'|}$.

Equality: Let $w = u(v)$ and $w' = u'(v')$. We can always write $w = a(b)$ and $w' = a'(b')$ with $|a| = |a'| = \max(|u|, |u'|)$ and $|b| = |b'| = \text{lcm}(|v|, |v'|)$

Delays and Buffers: can be computed practically after normalisation

The set of infinite periodic binary words is closed by sampling (`on`), delaying (`pre`) and point-wise application of a boolean operation

$$w ::= u(v)$$

$$c ::= w \mid c \text{ on } w \mid \text{not } c \mid \text{pre } (c) \mid \dots$$

From Pure-Synchrony to N -Synchrony

Pure-Synchrony:

- Synchrony can be checked using standard type system
- only need clock equality (and clocks are not necessarily periodic)

$$\frac{H \vdash e_1 : ck \quad H \vdash e_2 : ck}{H \vdash op(e_1, e_2) : ck}$$

N-Synchrony:

- extend the basic clock calculus with a **sub-typing rule**:

$$\text{(SUB)} \quad \frac{H \vdash e : ck \text{ on } w \quad w <: w'}{H \vdash e : ck \text{ on } w'}$$

- defines the synchronisation points where buffer code should be inserted

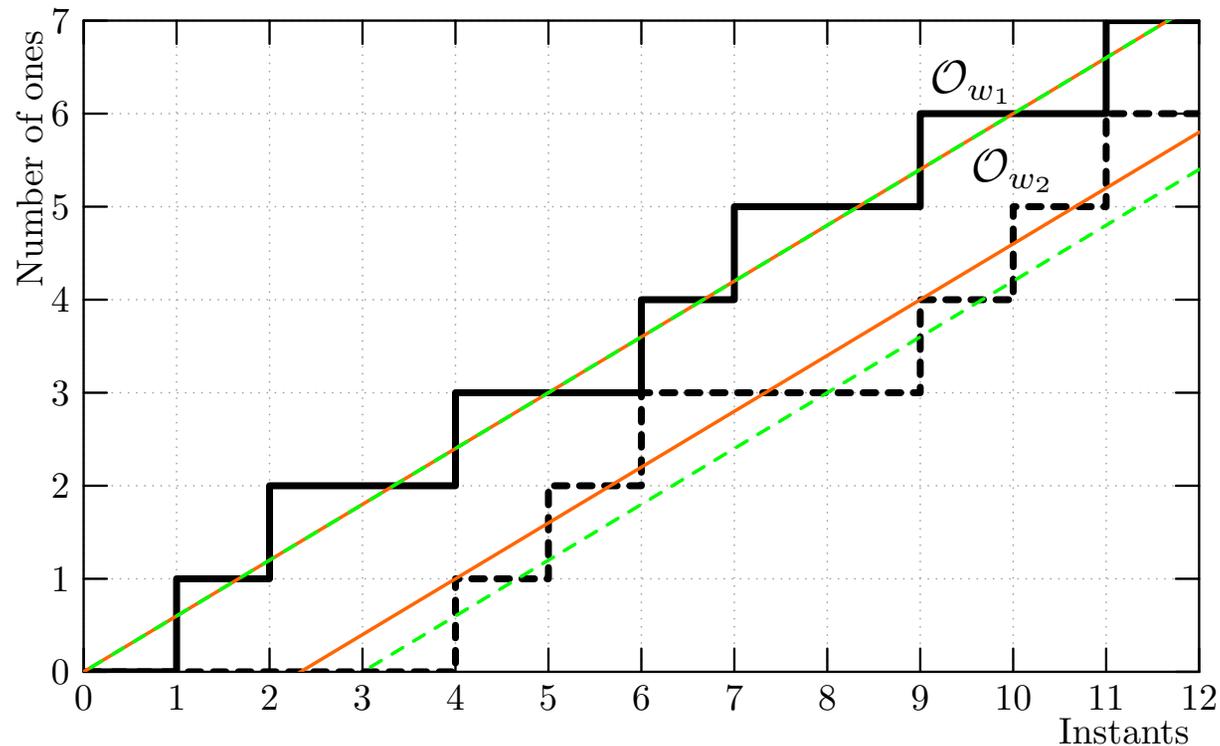
Going further: what about non periodic systems?

- Introducing clock relations gives more flexibility with as much guarantees as in synchronous model. No deadlocks, no buffer overflows.
- Subtyping relation can be checked provided clocks are periodic.
- Computing with exact period is unfeasible in practice. E.g.,
 $(10100100) \text{ on } 0^{3600}(1) \text{ on } (101001001) =$
 $0^{9600}(10^4 10^7 10^7 10^2)$
- Motivations:
 1. dealing with long patterns in periodic clocks. Avoid exact computation.
 2. specify/model jittering, i.e., how to deal with “almost periodic” clocks ? For instance
 $\alpha \text{ on } w \text{ with } w = 00.((10) + (01))^*$
(e.g. $w = 00 01 10 01 01 10 01 10 \dots$)

Idea: Manipulate sets of clocks instead of clocks. Transform the synchronisation problem into a linear problem with rational numbers.

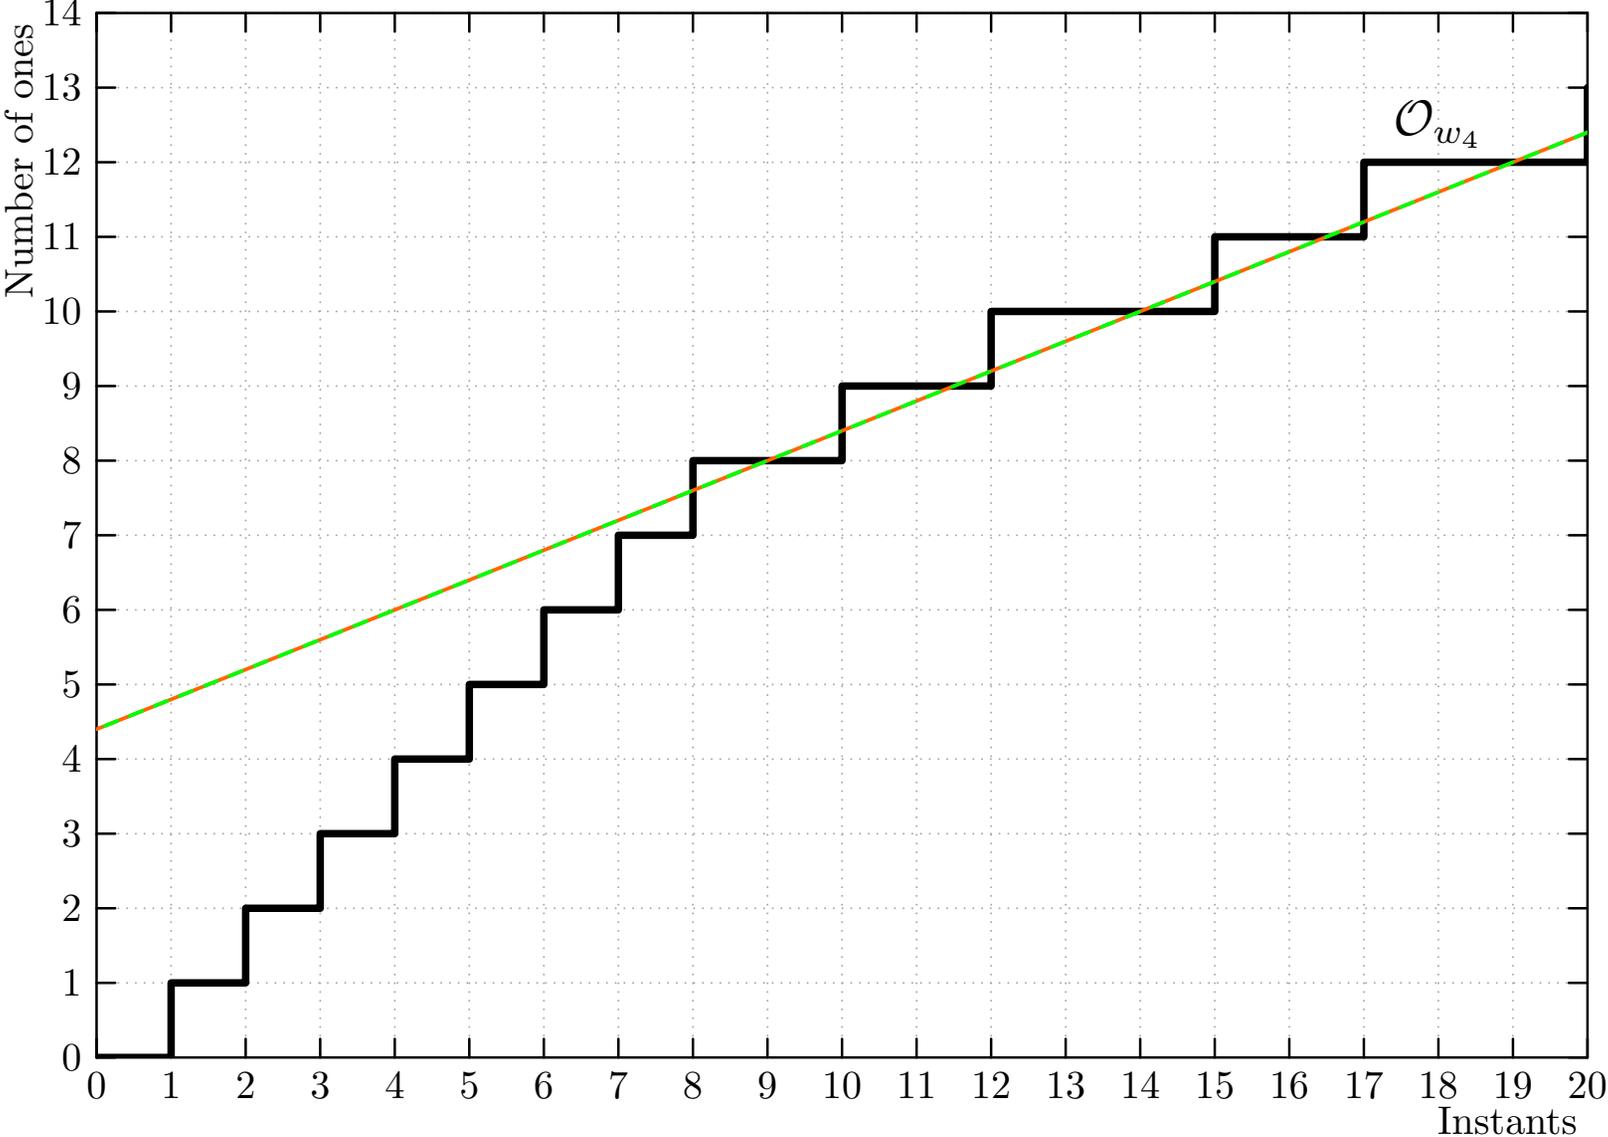
Clock abstraction (work in progress)

$$\text{concr}((b^0, b^1, r)) \stackrel{\text{def}}{\iff} \left\{ \begin{array}{l} w, \forall i \geq 1, \\ \wedge \left. \begin{array}{l} w[i] = 1 \implies \mathcal{O}_w(i-1) < r \times i + b^1 \\ w[i] = 0 \implies \mathcal{O}_w(i-1) \geq r \times i + b^0 \end{array} \right\} \right.$$

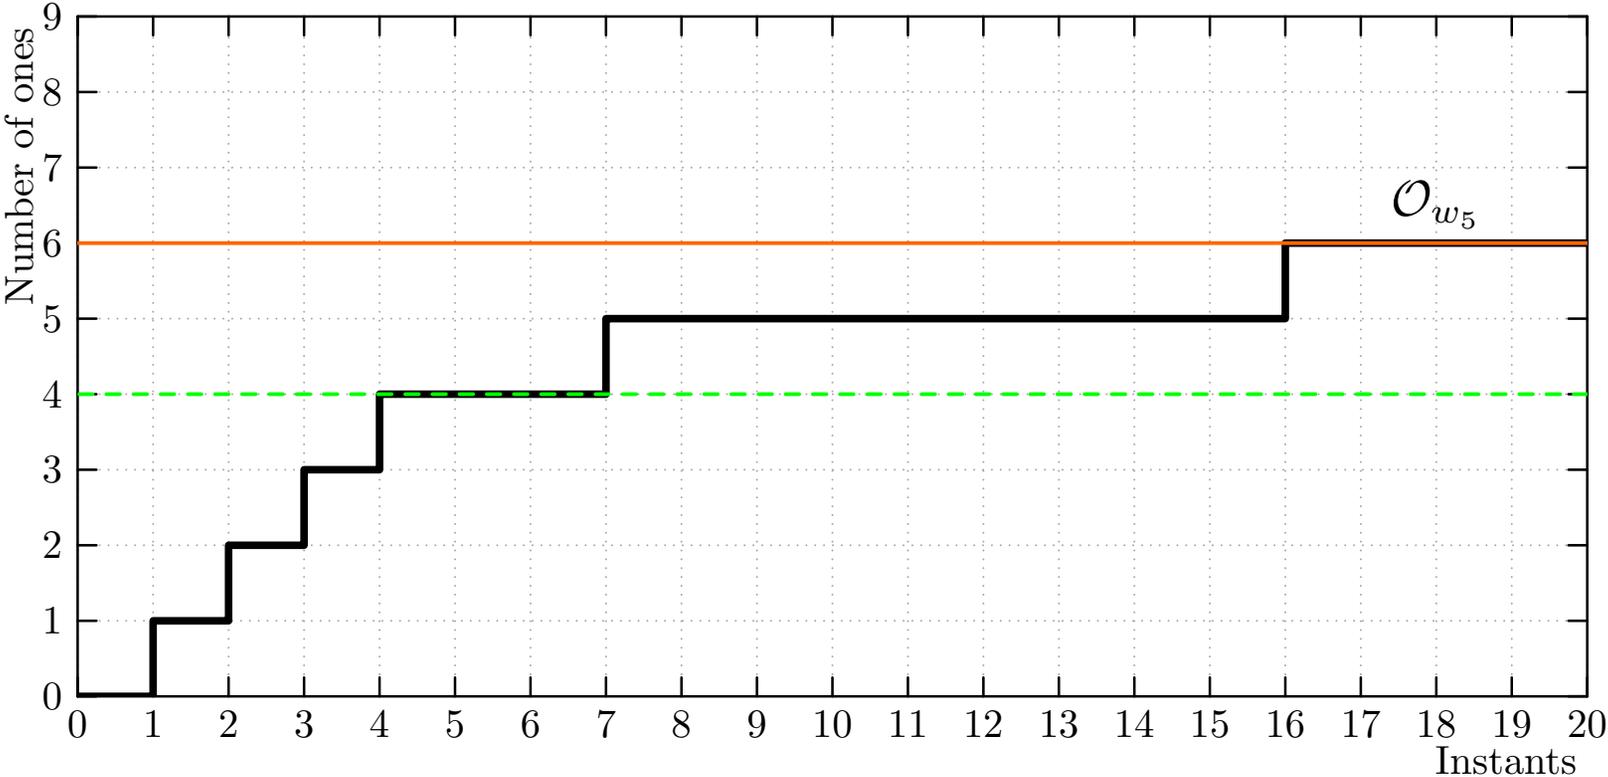


$$w_1 : (0, 0, \frac{3}{5}) \text{ and } w_2 : (-\frac{9}{5}, -\frac{7}{5}, \frac{3}{5})$$

- Initial sets of 1s are well abstracted.



- Clocks with a nul rate can be abstracted.



Properties

Definition 1 ($early_a, late_a$). Let $a = (b^0, b^1, r)$ be a clock envelope.

$$early_a = \sqcap \{w, \forall i \geq 1, w[i] = 1 \implies \mathcal{O}_w(i-1) < r \times i + b^1\}$$

$$late_a = \sqcup \{w, \forall i \geq 1, w[i] = 0 \implies \mathcal{O}_w(i-1) \geq r \times i + b^0\}$$

Proposition 1 (bounds of the envelope).

$$\forall w \in \text{concr}(a), (early_a \preceq w) \wedge (w \preceq late_a).$$

Proposition 2 (Empty concretisation). $\forall a, \text{concr}(a) = \emptyset \Leftrightarrow early_a \not\preceq late_a$.

Proposition 3 (Early and Late binary words).

$$\forall i, \mathcal{O}_{early_a}(i) = \max(0, \min(i, \lceil r \times i + b^1 \rceil))$$

$$\mathcal{O}_{late_a}(i) = \max(0, \min(i, \lceil r \times i + b^0 \rceil))$$

Proposition 4 (Non-empty envelope).

$$\forall a = (b^0, b^1, r), b^0 \leq b^1 \implies \text{concr}(a) \neq \emptyset.$$

Proposition 5 (Perfect Periodic Clock). $|\text{concr}(b^0, b^1, r)| = 1$.

Clock enveloppes as circuits (i.e., automata)

Given (b_0, b_1, r) , write a generator/acceptor of clocks within an envelope: this is indeed a synchronous circuit (here written in Lucid Synchrone syntax)

$<:$, $*:$, etc. are the classical operation lifted to rational.

```
type rat = { num: int; den: int }
```

```
let norm ( { num = n; den = 1 } , i , j ) =  
  if i >= 1 && j >= n then ( i - 1 , j - n ) else ( i , j )
```

```
let node check( ( b0 , b1 , r ) , clk ) = ok where  
  rec i , j = ( 1 , 0 ) fby norm( r , i+1 , if clk then j + 1 else j )  
  and ok = if clk  
    then ( rat_of_int j ) <: r * : ( rat_of_int i ) +: b1  
    else ( rat_of_int j ) >=: r * : ( rat_of_int i ) +: b0
```

We only need integer arithmetic. In the same way, we can implement a generator, which either non-deterministically produce a clock within an envelope or the *early* or *late* bounds.

Abstract Operators: not^{\sim} and on^{\sim}

Replace exact computation with not and on by abstract ones.

$$not^{\sim} ((b^0, b^1, r)) = (-b^1, -b^0, 1 - r)$$

Property: $a = not^{\sim} not^{\sim} a$

If $b^0_1 \leq 0$ and $b^0_2 \leq 0$:

$$(b^0_1, b^1_1, r_1) on^{\sim} (b^0_2, b^1_2, r_2) = (b^0_{12}, b^1_{12}, r_{12})$$

with: $r_{12} = r_1 \times r_2$, $b^0_{12} = b^0_1 \times r_2 + b^0_2$, $b^1_{12} = b^1_1 \times r_2 + b^1_2$

Abstraction of a sampled clock:

We are able to abstract a composed clock without computing the associated binary word.

$$abs(not\ w) \stackrel{def}{\Leftrightarrow} not^{\sim} abs(w)$$

$$abs(c_1\ on\ c_2) \stackrel{def}{\Leftrightarrow} abs(c_1) on^{\sim} abs(c_2)$$

Proposition: Those operations are correct, i.e., $c \in concr(abs(c))$

Abstract Relations: \bowtie^{\sim} , \preceq^{\sim} , $<:^{\sim}$

If the abstract relation is verified, the concrete one is verified on all elements of the respective concretization sets

$$(b^0_1, b^1_1, r_1) \bowtie^{\sim} (b^0_2, b^1_2, r_2) \Leftrightarrow r_1 = r_2$$

Proposition: $abs(c_1) \bowtie^{\sim} abs(c_2) \Leftrightarrow c_1 \bowtie c_2$

Checking precedence is checking an arithmetic inequality

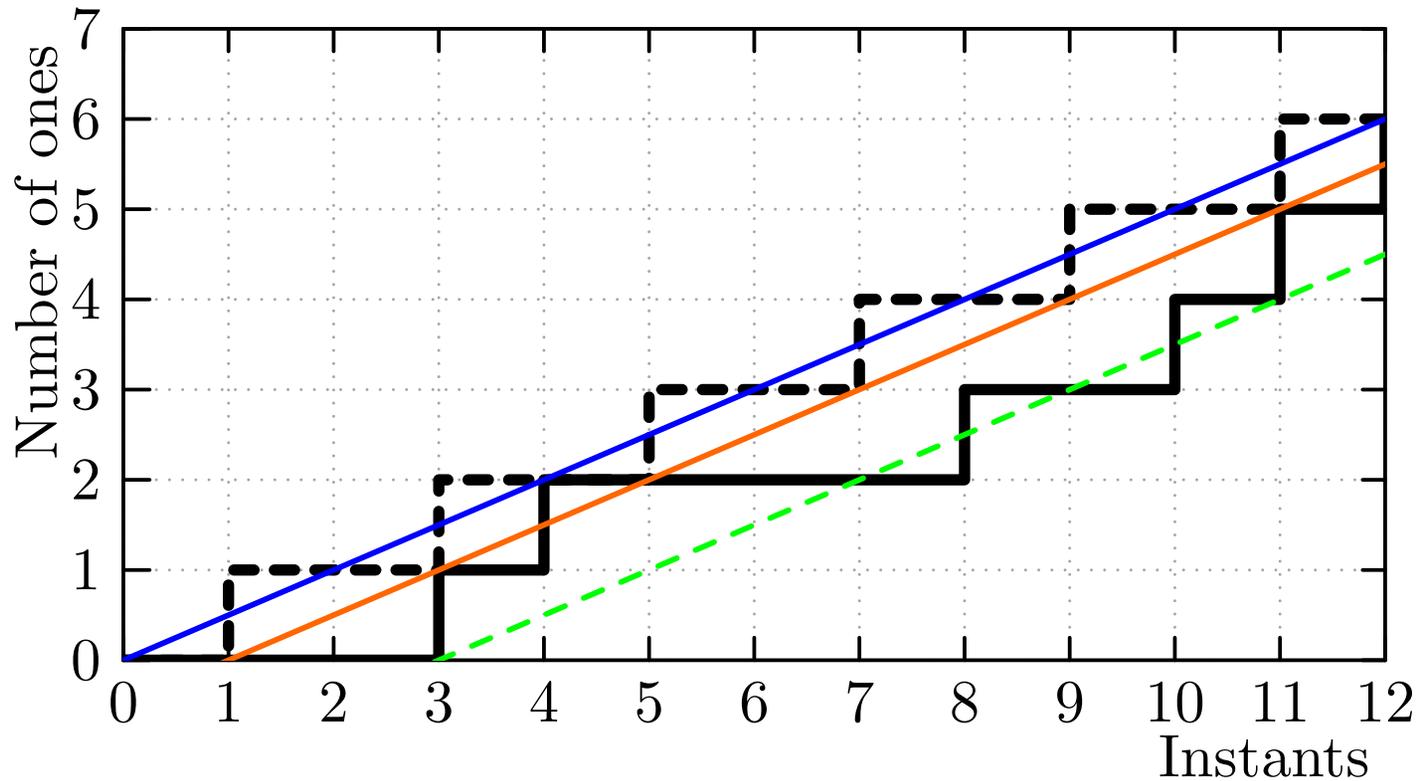
$$b^0_1 \geq b^1_2 \implies a_1 \preceq^{\sim} a_2$$

Proposition: $abs(c_1) \preceq^{\sim} abs(c_2) \implies c_1 \preceq c_2$

$$a_1 <:^{\sim} a_2 \Leftrightarrow a_1 \bowtie^{\sim} a_2 \wedge a_1 \preceq^{\sim} a_2$$

\implies Subtyping can be checked in constant time.

Modelizing Execution Time



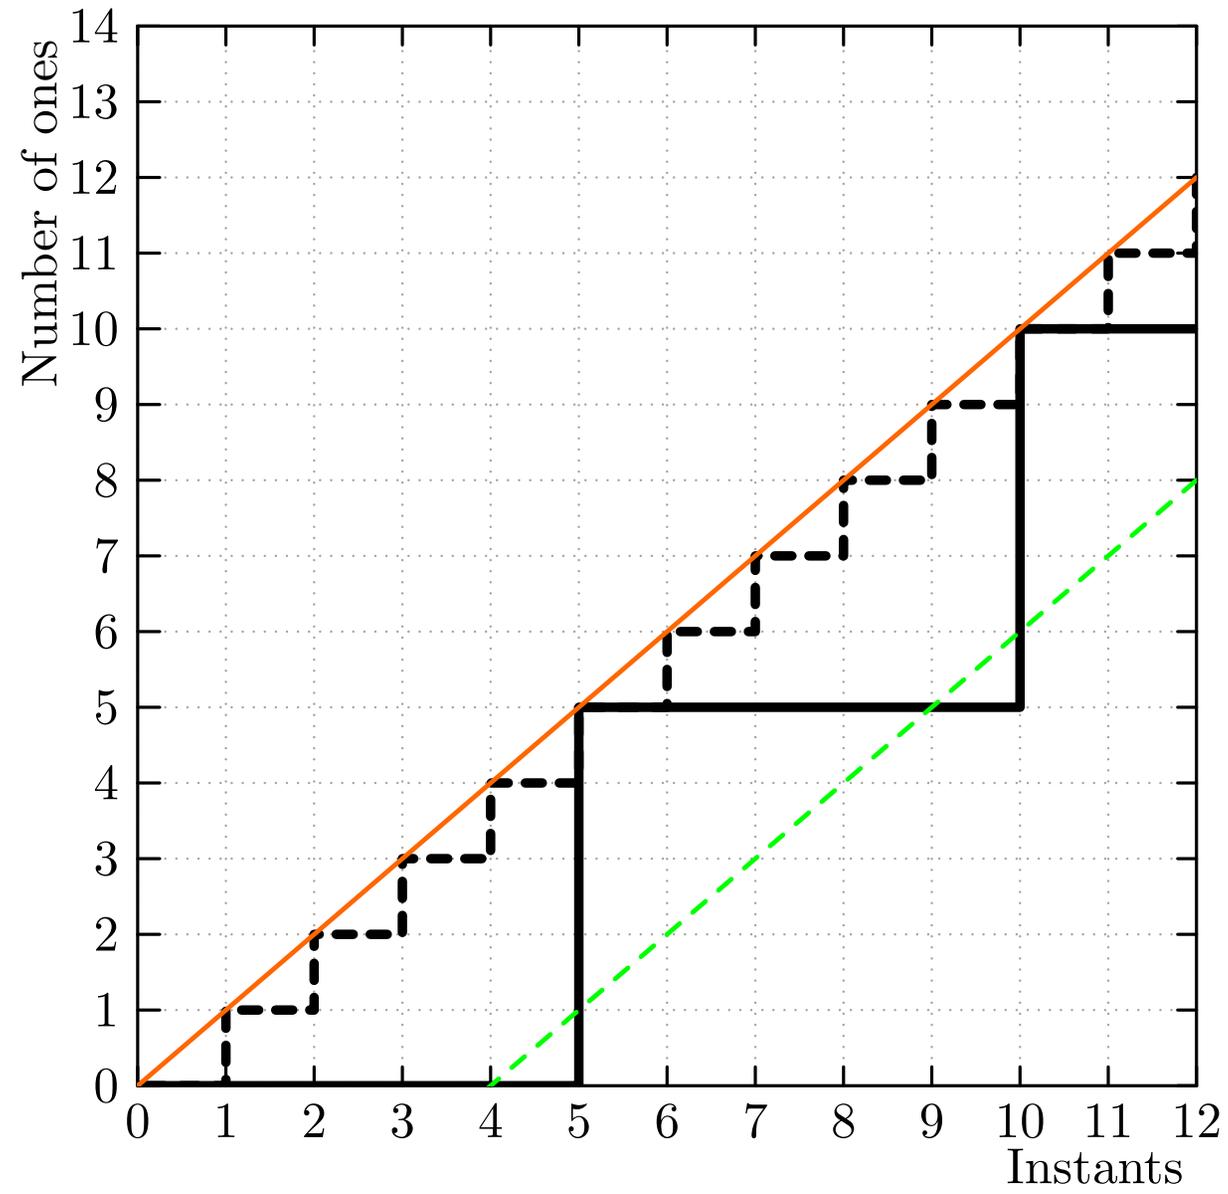
$f :: \forall \alpha. \alpha \text{ on}^{\sim} (0, 0, \frac{1}{2}) \rightarrow \alpha \text{ on}^{\sim} (-\frac{3}{2}, -\frac{1}{2}, \frac{1}{2})$ i.e.

$\alpha \text{ on}^{\sim} (0, 0, \frac{1}{2}) \rightarrow \alpha \text{ on}^{\sim} (-3, -1, 1) \text{ on}^{\sim} (0, 0, \frac{1}{2})$

f must be executed every 2 cycles and that its computation takes between one and three cycles

Composed twice: $f \circ f :: \forall \alpha. \alpha \text{ on}^{\sim} (0, 0, \frac{1}{2}) \rightarrow \alpha \text{ on}^{\sim} (-\frac{6}{2}, -\frac{2}{2}, \frac{1}{2})$

Modelizing Several Reads (or writes) at the Same Instant



Conclusion

- Synchronous data-flow as a sub-set of Kahn Process Networks
- Synchrony means the existence of a common time scale between two communicating processes
- Checking synchrony is mainly a typing problem
- Relaxing synchrony to model a larger class of systems, yet ensuring bounded buffering communication
- algebraic properties on clock sequences (e.g., synchronization, clock enveloppes) have been formalized and proof in the proof assistant Coq (5000 lines)
- We are currently developing a new language to incorporate those clocks

References

- [1] Paul Caspi and Marc Pouzet. Synchronous Kahn Networks. In *ACM SIGPLAN International Conference on Functional Programming*, Philadelphia, Pennsylvania, May 1996.
- [2] Albert Cohen, Marc Duranton, Christine Eisenbeis, Claire Pagetti, Florence Plateau, and Marc Pouzet. N -Synchronous Kahn Networks: a Relaxed Model of Synchrony for Real-Time Systems. In *ACM International Conference on Principles of Programming Languages (POPL'06)*, Charleston, South Carolina, USA, January 2006.
- [3] Albert Cohen, Louis Mandel, Florence Plateau, and Marc Pouzet. Abstraction of Clocks in Synchronous Data-flow Systems. In *The Sixth ASIAN Symposium on Programming Languages and Systems (APLAS)*, Bangalore, India, December 2008.
- [4] Marc Pouzet. *Lucid Synchrone, version 3. Tutorial and reference manual*. Université Paris-Sud, LRI, April 2006. Distribution available at:
www.lri.fr/~pouzet/lucid-synchrone.