



F E R M A T

Formal Engineering Research using
Methods, Abstractions and Transformations

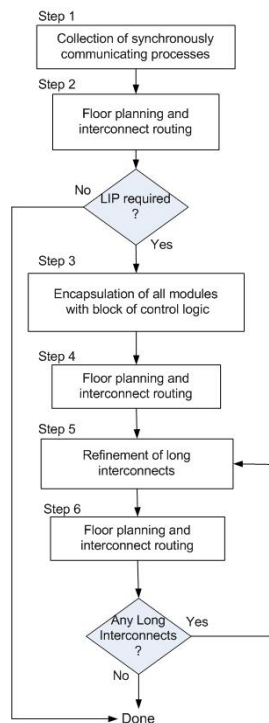
Technical Report No: 2005-03

A Functional Programming Framework for Latency Insensitive Protocol Validation

Syed Suhaib, Deepak Mathaikutty, David Berner,
Jean-Pierre Talpin, Sandeep Shukla

ssuhaib.damathai.shukla@vt.edu
{david.berner,talpin}@irisa.fr

Abstract -- Latency insensitive protocols (LIPs) have been proposed as a viable means to connect synchronous IP blocks via long interconnects in a system-on-chip. The reason why one needs to implement LIPs on long interconnects stems from the fact that with increasing clock frequencies, the signal delay on some interconnects exceeds the clock period. Correctness of a system composed of synchronous blocks communicating via LIPs is established by showing latency equivalence between a completely synchronous composition of the blocks, and the LIP based composition. A design flow based on a synchronous composition specification, and stepwise refinement to LIP composition can be easily conceived, and a proof obligation to show latency equivalence between the synchronous specification and the refinement needs to be discharged. In this work, we propose a functional programming based framework for modeling and simulating LIP, and implement the semantics of various refinement steps in the programming model, so we can validate the LIP model against the original system within this functional programming framework. Such validation becomes easier due to the inherent denotational model of functional languages. We specifically use Standard ML to model the original system implementation as well as its latency insensitive version and compare the two by creating a model that contains both, giving them the same inputs and checking their outputs to be latency equivalent.



A Functional Programming Framework for Latency Insensitive Protocol Validation

Syed Suhaib, Deepak Mathaikutty, David Berner,
Jean-Pierre Talpin, Sandeep Kumar Shukla

Virginia Tech, USA
INRIA, France

{ssuhaib,damathai,shukla}@vt.edu
{dberner,talpin}@irisa.fr

Contents

1	Introduction	2
2	Related Work	4
3	Background and Preliminary Definitions	5
3.1	Functional Programming	5
3.2	Preliminary Definitions	6
4	LI Refinement Steps	8
4.1	SML based LIP description	10
4.2	Check for Correctness	17
4.3	Case Study	18
5	Multiclock extension to LIP	20
6	Conclusion and Future Work	21

List of Figures

1	Refinement steps to LI implementation	9
2	Bridge	13
3	Comparing synchronous system with its LI implementation	18
4	Adaptive Modulator	19
5	LI based Adaptive Modulator	19

List of Tables

Abstract

Latency insensitive protocols (LIPs) have been proposed as a viable means to connect synchronous IP blocks via long interconnects in a system-on-chip. The reason why one needs to implement LIPs on long interconnects stems from the fact that with increasing clock frequencies, the signal delay on some interconnects exceeds the clock period. Correctness of a system composed of synchronous blocks communicating via LIPs is established by showing latency equivalence between a completely synchronous composition of the blocks, and the LIP based composition. A design flow based on a synchronous composition specification, and stepwise refinement to LIP composition can be easily conceived, and a proof obligation to show latency equivalence between the synchronous specification and the refinement needs to be discharged. In this work, we propose a functional programming based framework for modeling and simulating LIP, and implement the semantics of various refinement steps in the programming model, so we can validate the LIP model against the original system within this functional programming framework. Such validation becomes easier due to the inherent denotational model of functional languages. We specifically use Standard ML to model the original system implementation as well as its latency insensitive version and compare the two by creating a model that contains both, giving them the same inputs and checking their outputs to be latency equivalent.

1 Introduction

In today's embedded systems, clock speeds keep rising. The signal propagation speed however is not increasing, hence a growing number of designs hit a limit where some wires on the chip are as long as the distance a signal covers during one clock cycle. Since, the number of gates reachable in a single cycle does not change significantly, the percentage of the chip reachable within a single clock cycle is decreasing, and as a result we have reached a point where more gates fit on a chip than can be communicated in a single clock cycle [1]. In order to go past this limit, latency insensitive protocols (LIP) provide means to let components with a multiple clock cycle distance still communicate correctly [2].

In current SoC based design methodologies, reduced time-to-market dictates efficient reuse of complex components. This has led to the idea of developing libraries of Intellectual Property (IP) components. The integration of such complex IPs on SoC and communication between them has shifted the performance bottleneck of the system from computation within them to communication between them.

The impact of this shift of problem domains can be seen in state-of-the-art microprocessor designs. For example, the design of the hyperpipelined Netburst microarchitecture of Intel's Pentium 4 processor uses one of the first pipelines containing pipeline stages designed exclusively to handle wire delays. A *drive-stage* is dedicated to specifically handle the signal propagation without performing any computation [3]. Furthermore, recent studies using cache analysis tools predict that in a 35-nm design running at 10GHz, accessing a 4-Kbyte Level-1 cache requires about three clock cycles [4]. Increased relative interconnect length affects current memory-oriented microarchitectures that strongly rely on the low communication latency assumption [5]. Therefore, it is important to equip such models with protocols that make them latency insensitive (LI) and ensure proper functioning also for distances beyond one clock cycle.

Dynamic test and validation can be used to show that a system using LI tech-

niques is functionally equivalent to the same system assuming zero delay communication. These tests are not reliable, since these tests only cover certain input vectors. If the LI protocols can be validated using formal verification techniques, the system is correct for all possible input vectors. In order to formally verify such protocols, they have to be specified in a formal language and checked for formal properties. One restriction of formal verification is that the state space of models is growing exponentially with their complexity, but the number of states that can be verified is limited [6]. Another way to confirm the correctness of such an implementation is to mathematically formalize it, but mathematically proving the equivalence of two systems is a challenging task. It requires complex mathematical proofs, that are not straightforward to follow by others who want to confirm them.

In this work, we propose a functional programming framework to validate such systems. A functional program is a function that receives the programs input as argument and delivers the programs output as result. It has no internal state, which makes it free from any side effects. Most other LIP approaches [2, 7] lack the validation of functionality. [1] shows a formal proof of correctness of the LIP, but is difficult to follow. In [8] we formally verify the protocol for a simple example, but to verify larger systems would be difficult due to the state space explosion. Here, the validation of the protocol is easy to formulate and to comprehend, and can also be applied to large systems. We specifically use Standard ML (SML) [9] to model the original system implementation as well as its LI version and check the output of both models to be latency equivalent.

The paper is organized as follows: In Section 2, we show the related work done using LIPs. In Section 3, we introduce the preliminary definitions and notations used in the paper, followed by the LIP refinement methodology illustrated in Section 4. In this section, we described the components needed for the LIP refinement and show how it is implemented in SML along with a case study. In Section 5, we show the LIP implementation for multiclock systems followed by

the conclusion in Section 6.

2 Related Work

LIP for systems with long interconnection delays (i.e. greater than one clock cycle) were initially proposed by Carloni et al [10, 1, 11] for single-clock SoCs. In their approach, all processes are encapsulated in a wrapper to derive a process that is latency equivalent¹ to the actual process, without having to modify the internals of the original IP. Relay stations are added along the long interconnections. They act like pipeline blocks to store and forward data, and contain at least two registers and a control logic. The insertion of these relay stations increases the number of elements to route and requires additional space on the chip for their placement. Once it is determined where relay stations have to be added based on the length of the interconnects, placement and routing of the entire chip design now including the relay stations has to be redone. Several iterations for placement and routing are needed in order to get a configuration that satisfies all interconnection constraints.

All components of such LI designs are assumed to operate with the same clock. Singh and Theobald generalize the LI theory for Globally Asynchronous and Locally Synchronous (GALS) systems [7]. In their approach, complex FSMs implemented in the wrapper control all input and output signals. The communication network is implemented as an asynchronous system to connect modules with different clocks. Overall this approach is associated with heavy penalties in terms of implementation costs and performance.

Casu and Macchiarulo show how to reduce chip area compared to Carloni's approach [2]. They use a smart scheduling algorithm for the functional block activation and substitute relay stations with simple flip-flops. One disadvantage of this approach is that the schedule has to be computed a priori and depends on the computation in the process. If any change is made in any process, it may result in a change of the flow of tokens. In this case, the schedule has to be recalculated,

¹We define the notion of latency equivalence in the preliminary definitions section

which is expensive.

In [8] we propose another modification of Carloni’s approach, which involves removing the relay stations along the long interconnects and inserting extra wiring logic using a splitter and a merger process. This solution is generalized for multi-clock systems where communication is done based on a global clock and the process wrapper links the processes to the environment, irrespective of the local clock of the process. Since there are no relay stations, there is no requirement to place the relay stations along the wires, whereas the splitter and merger processes are placed on the interface of the process. Here, we use this approach and show how to validate it in a functional programming framework.

3 Background and Preliminary Definitions

3.1 Functional Programming

Functional programming is seen to be highly relevant to the understanding of reactive and interactive systems. A computation is expressed as a function and its interaction with the outside world is modeled as inputs given to the function. These interactions are considered infinite in nature, and are modeled using streams [12]. Functional programs contain implicit parallelism, which is very useful when dealing with embedded system applications, since they typically have a considerable amount of built-in parallelism, whereas imperative languages such as C++ are inherently sequential. While functional languages fit naturally for data flow applications, they provide a rich variety of control constructs, making them more suitable for control dominated applications. Furthermore, functional languages provide a clean and simple semantic model, which performs all computation by function application, thereby providing a more abstract notation to express computation.

We use Standard ML (SML) [9] to model the original system as well as its LI version and compare both the models. SML offers an excellent ratio of expressiveness to language complexity, and provides competitive efficiency. SML

manages to combine safety, security, and robustness with a great deal of flexibility because of its type and module system. Other features of the SML based framework include:

- SML provides good expressiveness with its ability to treat functions as first-class values, and its usage of higher-order functions. The availability of imperative constructs provide great expressive power within a simple and uniform conceptual framework.
- SML provides a high-level model which makes programming more efficient and more reliable by automating memory management and garbage collection.
- SML does static type checking which detects many errors at evaluation time. Error detection is enhanced by the use of pattern matching and by the exception mechanism.
- The SML module system is an extension of the underlying polymorphic type system thereby providing separation of interface specification and implementation. These facilities are very effective in structuring large programs and defining generic, reusable software components.

3.2 Preliminary Definitions

In this section, we show some of the definitions we use in the rest of the paper. Let V be the set of data values and, T be a countable set of time stamps. Unless otherwise specified, in this paper, we assume $T = \mathbb{N}$ = set of natural numbers. An event $e \in V \times T$ is an occurrence of a data value with a particular time stamp. However, in the systems we consider, a special event called *absent event* denoted by τ may occur². Therefore, the set of all events is denoted by E , where $\tau \in E$ and for all other $e \in E$, $e \in V \times T$. When $e \in V \times T$ it is called an *informative event*.

²It may be caused due to lack of valid data in the producer or due to the consumer's request to delay a transmission

A *signal* s is defined to be a sequence of events, often denoted as $e_1e_2e_3 \dots$ where $e_i \in E$.

For the preliminary definitions, if s is a signal, $s[i]$ denotes the i^{th} event, hence either $s[i] \in V \times T$ or $s[i] = \tau$. The set of all signals is denoted by S . The signals can be either input signals or output signals of a process. We also distinguish *Stall signals* from all signals in the system. A stall signal st is a sequence of boolean events, i.e., $st[i] \in Bool \times T$. The set of all stall signals is denoted by S_T . In our system, IPs are hardware modules that map input signals to output signals, therefore in this paper we refer to them as processes. A process p is a function $S^n \rightarrow S^m$ where n, m are natural numbers. A synchronous system consists of these processes where zero-delay communication and zero-time computation among these processes happen at the global clock edge.

In the remainder of this section, we define a few terms and notations that are used throughout the paper.

Definition 1 Given one tuple of m elements and another of n elements, \odot creates a tuple of $m+n$ elements.

$$\langle a_1, a_2, \dots, a_n \rangle \odot \langle b_1, b_2, \dots, b_m \rangle = \langle a_1, a_2, \dots, a_n, b_1, b_2, \dots, b_m \rangle$$

Definition 2 Given two tuples of n events and n signals respectively, \oplus creates a tuple of n signals with the n events appended to the n signals.

$$\langle e_1, e_2, \dots, e_n \rangle \oplus \langle s_1, s_2, \dots, s_n \rangle = \langle e_1 \oplus s_1, e_2 \oplus s_2, \dots, e_n \oplus s_n \rangle \text{ where } e_1 \oplus s_1 = s'' \text{ is a signal such that } s''_1 = e_1s_1$$

Definition 3 $Vect_{i=1}^n(\text{exp}(i)) = \langle \text{exp}(1), \text{exp}(2), \dots, \text{exp}(n) \rangle$ where, $\text{exp}(i)$ is an expression s.t. $\text{exp}(k)$ is a textual replacement of i by k .

Definition 4 Latency Equivalence: Two signals s_1 and s_2 are said to be latency equivalent, $s_1 \equiv_e s_2 \Leftrightarrow \mathcal{F}(s_1) = \mathcal{F}(s_2)$, where

$\mathcal{F} : S \rightarrow S$ be defined as, $\mathcal{F}(s) = f(s, 1, n)$ and,

$$f(s, i, n) = \begin{cases} f(s, i+1, n), & \text{if } s[i] = \tau \\ s[i], & \text{if } (i = n) \\ s[i] \oplus f(s, i+1, n), & \text{otherwise} \end{cases}$$

\mathcal{F} takes a signal as input and outputs a signal dropping all τ events, but preserving all informative events. The above definition assumes that n is the length of the signal and signal is assumed to be finite. If s is infinite, then the above definition changes slightly.

If the output signals of two systems are latency equivalent given the same input signals, then these systems are said to be latency equivalent.

Definition 5 Sequential composition: Let p_1 and p_2 be two processes where, $p_1 : S \rightarrow S$ and $p_2 : S \rightarrow S$ then $p_2 \circ p_1(s) = s'$ is called a sequential composition.

Definition 6 Feedback composition [13]: Given a process $p: (S \times S) \rightarrow (S \times S)$ with two input signals and two output signals, the feedback process $FB_p(p) : S \rightarrow S$ is defined as $FB_p(p)(s_i) = s_k$ where $p(s_i, s_j) = (s_j, s_k)$. The behavior of the feedback is defined by the fixed point semantics.

4 LI Refinement Steps

In this section, we show a transformation procedure to design a LI system. The transformation of a synchronous system to a LI system is shown in Figure 1.

The steps to LIP refinement are as follows:

1. We start with a collection of synchronously communicating components. These components can be custom-made modules or IP cores.
2. Approximate floor planning and interconnection routing are done by design engineers to check for long interconnects. If all communication can be done in a clock cycle, then there is no need for LIP refinement.

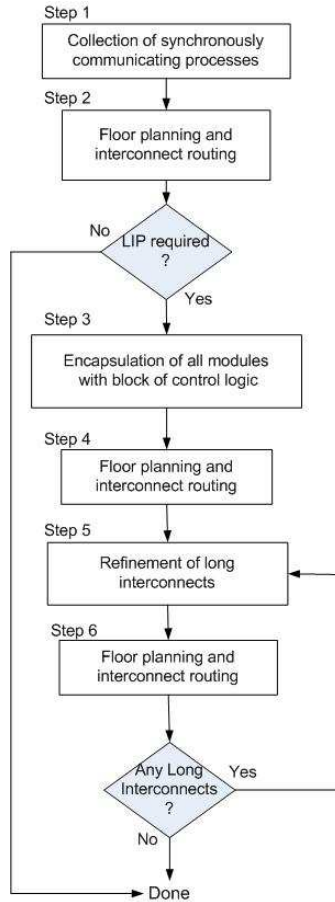


Figure 1. Refinement steps to LI implementation

3. All modules are encapsulated with a block of control logic. This encapsulation may include adding control logic that controls the flow of the events, buffers, control stations, repeater stations etc to enable correct transmission of data with the LIP. Once each process is encapsulated, verification is done to ensure its correctness, meaning that they behave similar to the original processes. By correctness here, we mean that if two processes are given the same set of input events, then the order of the informative events on its output signals are the same. We call this *latency equivalence*³.
4. Estimation using floor planning and interconnect routing is done again, this

³Two signals are said to be latency equivalent if both signals have the same order of informative events

time with the encapsulated processes to relocate and evaluate the delays on the long interconnects.

5. After finding the delays on the long interconnects, the designer can then segment those long interconnects with additional processes containing buffers, latches, forwarding stations, etc to ensure that data is properly communicated through the long interconnects. Depending on the delay of the interconnect, the events can be compared from the point they are placed on the signal to the point they leave the signal.
6. Floor planning and interconnect routing is done again to ensure that no long interconnects exist in the system.

4.1 SML based LIP description

In this section, we describe the components of the LI framework and its implementation in SML. A finite signal is modeled as generic list, whereas an infinite signal is written as delayed function application as shown in Listing 1.

Listing 1. Finite and Infinite Signal

```
(* Definition of a finite signal *)
2 datatype signal = nil | 'a :: 'a list

4 (* Definition of an infinite signal*)
datatype infseq = nil | cons of 'a * (unit -> infseq)
```

In SML, for our convenience we formulate an event to be a list of two elements, where the first element is the value and the second element identifies whether the event is an informative event or an absent event (eg. $e_j = [3,1]$ is the j^{th} event with 3 as the value and 1 as the identity of the event⁴). Hence, a signal can be formulated as a list of events. (eg: $s_i = [[1,1],[2,1],[3,0],\dots]$).

The refinement steps for transforming a synchronous system to a LI system can be thought of as a two stage operation. The first stage involves encapsulating

⁴1 corresponds to an informative event and 0 corresponds to an absent event

the synchronous components and the next stage involves refining the interconnects to make them consistent with the flow of events. These stages are shown in the steps in Figure 1. In the first stage, each module is encapsulated with an equalizer. An equalizer is a process instantiating template that given n input signals and a stall signal, it produces n output signals and n stall signals. The functionality of the equalizer can be divided into three modes:

1. *Disable mode*: In this mode, the equalizer gets a stall event on its input stall signal. In this case, the equalizer outputs absent events on all its output signals and enables all the output stall signals.
2. *Absent input mode*: In this mode, the equalizer receives an absent event on any of its input signals and a stall disabled value on its input stall signal. In this case, the equalizer outputs absent events on all its output signals and stalls only those processes from which it received an informative event.
3. *Valid mode*: The equalizer receives informative events on all its input signals and its stall input is false. At this time, the equalizer places informative events it received from its input signals to the output signals.

The formal model of the equalizer is shown below. The *if-else* conditions illustrate the functionality of the equalizer based on its modes.

Definition 7 An Equalizer function $\mathcal{E} : (S^n \times S_T) \longrightarrow (S^n \times S_T^n)$ is defined as:

$$\mathcal{E}(s_1, \dots, s_n, stall) = f(s_1, \dots, s_n, stall, 1, 1, \dots, 1) \text{ where,}$$

$$f(s_1, \dots, s_n, st_1 :: st, i_1, i_2, \dots, i_n) =$$

if($st_1 = F$) *then*

if($\exists_{j=1}^n (s_j[i_j]) = \tau$) *then* /* Absent input mode */

$$\langle \tau, \tau, \dots, \tau \rangle \odot Vect_{j=1}^n(exp_1(j)) \oplus f(s_1, \dots, s_n, st, Vect_{j=1}^n(exp_2(j)))$$

else /* Valid Mode */

$$Vect_{j=1}^n (s_j[i_j]) \odot \langle F, \dots, F \rangle \oplus f(s_1, \dots, s_n, st, Vect_{j=1}^n(i_j + 1))$$

else / Disable mode */*

$$\langle \tau, \tau, \dots, \tau \rangle \odot \langle T, \dots, T \rangle \oplus f(s_1, \dots, s_n, st, \text{Vect}_{j=1}^n(\text{exp}_2(j)))$$

exp₁(j) : if (s_j[i_j] = τ then F else T

exp₂(j) : if (s_j[i_j] = τ then i_j + 1 else i_j

The function f takes arguments i_1, i_2, \dots, i_n to index s_1, s_2, \dots, s_n to access the respective events. Here, $s = st_1::st$ means that st_1 is the first element and st is the rest of the signal. T and F denote the true and false values of the stall event.

Listing 2 shows the implementation of the *equalizer* process in SML. The equalizer reads one event from all the input signals of a process along with an event from the stall input. It then checks if all the events at a time are informative. The check for events is done through the *etypes* and *info* functions (lines 6-13). The functionality setting the stall values for *Disable mode* is done by the *stallon* function and the output is given by *e3* (line 21). The stall values when the equalizer is in *absent event mode* is set by *stallset* function and the output is given by *e2* (line 20). Finally, the *valid mode* output is given by *e1* (line 19).

Listing 2. Equalizer

```

1 fun equalizer () = fn s => fn st => f(s, st, indexstart(length(s)))
2
3 fun f([], st1::st, _) = [] | f(_, [], _) = [] | f(_, _, []) = [] |
4   f(s, st1::st, i) =
5   let
6     fun etype(x1::x2) = x2 | etype([]) = nil
7     fun etypes [] = [] | etypes(x1::x) = etype(x1) @ etypes(x)
8     fun info [] = false |
9       info(x1::[]) = if (x1 = 1) then true else false |
10    info(x1::x) = if (x1=1) then ( info(x)) else false
11    val allevents = e(s, i) (* Events from all the signals at a time *)
12    val allinfo = if info(etypes(allevents)) = true
13      then true else false (* True when all events are informative *)
14    fun stalloff(0) = [] | stalloff(n) = [1] @ stalloff(n-1)
15    fun stallon(0) = [] | stallon(n) = [0] @ stallon(n-1)
16    fun flipval(x) = if x=1 then 0 else 1
17    fun stallset([]) = [] | stallset(x1::x) = [flipval(x1)] @ stallset(x)
18

```

```

    val e1 = [ allevents , [ stalloff (length (allevents ))]]
20  val e2 = [ tauevents (length (s)) , [ stallset (tags (allevents ))]]
    val e3 = [ tauevents (length (s)) , [ stallon (length (allevents ))]]
22  in
    (case (st1) of
24  1 => (if allinfo = true
        then ([e1] @ f(s, st, incrementindex(i)))
26  else ([e2] @ f(s, st, incrementempty(i, etypes(allevents))))
        ) |
28  0 => ([e3] @ f(s, st, i)) |
    - => [])
30 end

```

The *equalizer* process is then sequentially composed with the synchronous process to form the shell of the process.

The next stage of the refinement methodology involves refining the long interconnects by inserting processes that not only ensure correct flow of events from one process to another, but also ensure that the delay in between the events is minimized. The long interconnects are refined by inserting *bridge* processes (Listing 3). A bridge is formed by sequential composition of a *splitter* and a *merger* process (Figure 2). Each *bridge* process has one input signal and one output signal. The delay on the bridge is modeled by the *Delayproc* process (line 1).




Figure 2. Bridge

Listing 3. *Bridge* process

```

fun Bridge(n) = fn s => Delayproc(n) (merger(n) (splitter(n) (s)))

```

The *splitter* and the *merger* process are connected by n interconnects where n is the delay on the long interconnect. Hence, *splitter* process has n output signals.

This process contains simple placement logic for the placement of events on these n signals. The splitter is implemented at the output of a process, and it transfers events on the corresponding signal. The splitter only places one event on one of the output interconnects and absent events are placed on the rest of the signals at a particular time stamp. Assuming that there are i events on the input signal of the splitter, at every cycle, the i^{th} event is placed on the n^{th} signal based on a rotational scheme. For example, if the delay on the interconnect is 3 cycles, then in the current cycle, the first element will be placed on the first signal and absent events will be placed on the other two signals. In the next cycle, the second event will be placed on the second signal and absent events will be placed on the first and third signals and for the third event it will follow the scheme. After the third event is placed, in the following cycle, the fourth event will be placed on the first cycle again. This rotation scheme will continue for the rest of the events. This functionality is illustrated by the formal definition shown below:

Definition 8 *Splitter \mathcal{H} is a process, s.t., $\mathcal{H}(s) = (s'_1, s'_2, \dots, s'_n)$ where*

$$(s'_1, s'_2, \dots, s'_n) = h(v, n, 1) \text{ and,}$$

$$g(n) = \begin{cases} \tau, & \text{if } n = 1 \\ \tau \odot g(n-1), & \text{otherwise} \end{cases}$$

$$f(x, n, i, j) = \begin{cases} x \odot g(n-j), & \text{if } i = 1 \\ \tau \odot f(x, n, i-1, j+1), & \text{otherwise} \end{cases}$$

$$h(x :: y, n, i) = \begin{cases} f(x, n, i, 1) \oplus h(y, n, 1), & \text{if } i = n \\ f(x, n, i, 1) \oplus h(y, n, i+1), & \text{otherwise} \end{cases}$$

Here, $s = x :: y$ means that x is the first event of the signal s and y is rest of the signal.

The SML implementation of the *splitter* process is shown in Listing 4. An input signal and the interconnect delay is given to the *splitter* process. One event is read from the input signal and *insertevent* function (line 6) places the event from

the input signal to one of the interconnects and absent events are placed on rest of the interconnects. The events are placed in the rotational scheme as illustrated earlier.

Listing 4. Splitter

```

1 fun splitter(n) = fn s => f(s,1,n)
2
3 fun f([],-,_) = [] |
4 f(x1::x , i , n) =
5   let
6     fun insertevent(_,j,0) = [] |
7       insertevent(y1,j,n) = (if n = j
8         then [y1] @ insertevent(y1,j,n-1)
9         else [[0,0]] @ insertevent(y1,j,n-1))
10    in
11      if (i = n)
12        then [insertevent(x1 , i , n)] @ f(x , 1 , n)
13        else [insertevent(x1 , i , n)] @ f(x , i+1 , n)
14    end

```

Contrary to the splitter, we implement a *merger* that takes n input signals and outputs one signal. The merger also extracts one event from the input signals based on the rotational scheme as illustrated earlier and places it on the output signal. The functionality of the *merger* is formally defined below:

Definition 9 *Merger* \mathcal{M} is a process, s.t. $\mathcal{M}(s_1, s_2, \dots, s_n) = s'$ where

$$s' = g((x_1 :: y_1, x_2 :: y_2, \dots, x_n :: y_n), n, 1) \text{ and,}$$

$$f(x :: y, n, i) = \begin{cases} x, & \text{if } i = n \\ f(y, n, i+1), & \text{otherwise} \end{cases}$$

$$g((x_1 :: y_1, x_2 :: y_2, \dots, x_n :: y_n), n, i) = \begin{cases} f((x_1, x_2, \dots, x_n), n, i) \oplus \\ g((y_1, y_2, \dots, y_n), n, 1), & i = n \\ f((x_1, x_2, \dots, x_n), n, i) \oplus \\ g((y_1, y_2, \dots, y_n), n, i+1), & \text{otherwise} \end{cases}$$

$$s_1 = x_1 :: y_1, s_2 = x_2 :: y_2, \dots, s_n = x_n :: y_n,$$

The SML representation of the merger is shown in Listing 5. The *extractevent* function extracts one event from all signals at a time (line 3). Extraction of events from the signals is done in similar way as they are placed on the interconnects by the splitter.

Listing 5. Merger

```

1 fun merger(n) = fn s => g(s,n,1)
2
3 fun g([], n, i) = [] | g(x1::x, n, i) =
4 let
5   fun extractevent([],n) = [] | extractevent(x1::x,n) =
6     (case (n) of
7       1 => x1 |
8       _ => extractevent(x, n-1))
9   in
10    if (i = n)
11      then [extractevent(x1,i)] @ g(x, n, 1)
12      else [extractevent(x1,i)] @ g(x, n, i+1)
13    end

```

After the refinement of all the components and the long interconnects of the synchronous system, all the components are composed together. The input sequence of the splitter and the output sequence of the merger are equivalent, since the order of events written by the splitter on the n output signals and the order of events read by the merger from its n input signals is the same. Therefore, the flow of events from the output of one shell across the long interconnect to the input of the corresponding shell is maintained. The stall signals connecting the components work on the feedback semantics. In functional languages, the feedback is defined by the fixed-point semantics. The fix point is computed for each event. Listing 6 shows the fixed point computation for the feedback semantics.

Listing 6. Feedback Process

```

1 fun fb(p) = fixpt(p,s,[],length(s)+1)
2 (* The fixpoint is computed on event basis *)
3 fun fixpt(q,s,sout,0) = sout | fixpt(q,s,sout,n) =
4   fixpt(q,s,(q s sout), n-1)

```

4.2 Check for Correctness

Once we have the LI system and the original synchronous system, we have to verify if they are latency equivalent in order to satisfy the proof obligation. We do this by checking if the outputs of the two systems are latency equivalent given the same input sequence. The comparator is modeled which is a reduced version of the equalizer. This *Eqcomparator* process compares the order of informative events output by the two systems. In the case when an absent event is seen on one of the output signals, it is discarded and the next event is considered on the same signal. The informative events on the two output signals are compared in sequence to ensure correct functionality. The LIP system satisfies the proof obligation if the output the two systems is latency equivalent, when given the same inputs. Figure 3 shows the setup of the problem. The SML code of the *Eqcomparator* is shown in Listing 7.

Listing 7. *Eqcomparator*

```
fun Eqcomparator () = fn s1 => fn s2 => compare(s1,s2,1,1);
2
fun compare([],[],_,_) = [] |
4  compare(_,[],_,_) = [] |
   compare(s1,s2,i,j) =
6  let
   val event1 = extractevent(s1,i);
8  val event2 = extractevent(s2,j);
   val valid = if event1 = [] orelse event2 = [] then false else true;
10 val valpresent =
   if valid = true
12   then if tag(event1) = [1] andalso tag(event2) = [1]
        then true else false
14  else false;
   fun comp(x,y) = if (x=y) then true else false;
16 in
   if valid = true andalso valpresent = true
18   then (if comp(value(event1),value(event2)) = true
        then [true] @ compare(s1,s2,i+1,j+1)
20        else [false])
   else if valid = true
22   then if (tag(event1) <> [0] andalso tag(event2) <> [0])
```

```

    then [ false ]
24  else if ( tag(event1) = [0] andalso tag(event2) = [0])
        then compare(s1, s2, i+1, j+1)
26  else if ( tag(event1) = [0] andalso tag(event2) = [1])
        then compare(s1, s2, i+1, j)
28  else if ( tag(event1) = [1] andalso tag(event2) = [0])
        then compare(s1, s2, i, j+1)
30  else [ true ]
    else [ true ]
32 end

```




Figure 3. Comparing synchronous system with its LI implementation

4.3 Case Study

We consider a case study of an adaptive modulator that consists of three IPs: regulator, convolutor and analyzer. The regulator module takes an input signal and a control signal and outputs based on the control signal by adding a threshold value. This output is then multiplied with a masking value by the convolutor module. The output of the system is given by the amplitude signal. The analyzer module outputs the control signal based on the input of the amplitude. The connections of these components are shown in Figure 4.

With these three IPs, we follow the refinement steps described in Figure 1 to construct the LI system. We have these three synchronous components. Early floor planning and interconnect routing is done to find the delays on the interconnects. We assume in this case that the regulated input signal is a long interconnect. Hence, we encapsulate the modules with the equalizer and repeat floor planning to find the delay on the interconnects. The long interconnect is then refined by

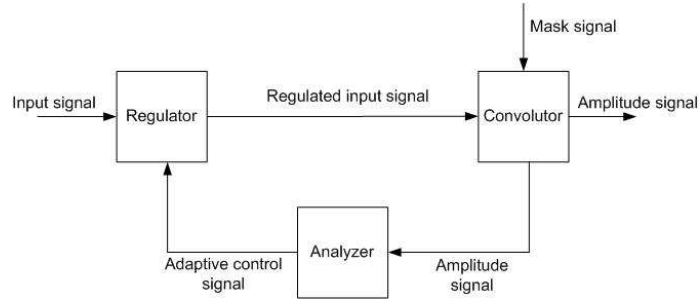


Figure 4. Adaptive Modulator

adding a bridge process. The new LI representation of the adaptive modulator is shown in Figure 5. The SML implementation is done using the components described in the previous section and is listed in the appendix. In order to check the correctness of the LI model, we create a model containing the LI implementation as well as the model with a zero communication assumption. We feed the same input sequence to both models and verify the latency equivalence of their outputs as described in Section 4.2.

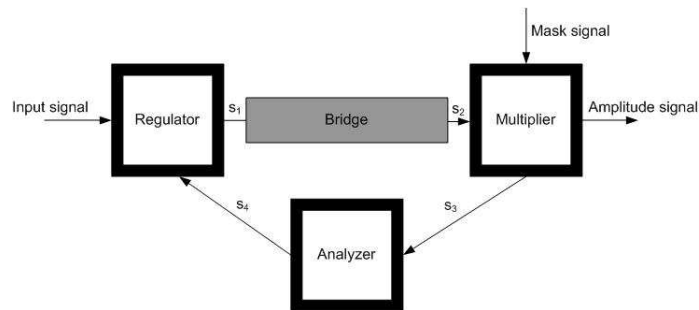


Figure 5. LI based Adaptive Modulator

The case study presented here is just an example to show how the LI refinement of a synchronous system can be done and validated against its original implementation. Any deterministic functionality can be integrated in a synchronous module and can be compared with its LI refinement. Since, changing the functionality is easy in a functional framework, many case studies can be analyzed successfully. On the other hand the functionality cannot be changed and verified easily when the system is being formally verified using a model checker.

5 Multiclock extension to LIP

The LI systems proposed earlier have been mainly targeting single clock synchronous systems where all components operate on the same clock. We now consider extending the existing LI implementation for multiclock systems where different components with different clocks are connected via arbitrarily long interconnects. The need for a system with components having different clocks arises when different IP blocks from different vendor are integrated in the same system. At this time, however, we are only permitting the use of components with defined clock relations. By clock relation, we mean that there is a known ratio of the evaluation cycle between different components. In the SML framework, the notion of clock is defined by the evaluation cycle of the processes. This approach therefore makes it possible to connect components with known clock ratios.

We modify our original refinement methodology for multiclock refinement. Before encapsulation of the processes, we add an *Insert* and a *Strip* process to each synchronous component of the system. The *Insert* process inserts n absent events for each event on the original incoming signal where n is the ratio of events on the incoming signal to the number of events evaluated by the process in each cycle. The output of the *Insert* process is then given to the original process. The formal definition of the *Insert* process is shown below:

Definition 10 *Insert* is a process, s.t. $I(s) = s'$ where

$$s' = g(y, n) \text{ and,}$$

$$f(n) = \begin{cases} \tau, & \text{if } n = 1 \\ \tau \odot f(n-1), & \text{otherwise} \end{cases}$$

$$g(x1 :: x, n) = (x1 \odot f(n)) \oplus g(x, n)$$

We also place a *Strip* process at the output of the synchronous component. This strip process removes the extra absent events inserted by the *Insert* process. The formal definition of the *Strip* process is given below:

Definition 11 *Strip* is a process, s.t. $\mathcal{W}(s) = s'$ where

$$s' = g(y, n) \text{ and, } t(x1 :: x) = x$$

$$f(s, n) = \begin{cases} t(s), & \text{if } n = 1 \\ f(t(s), n - 1), & \text{otherwise} \end{cases}$$

$$g(x1 :: x, n) = f(x1, n) \oplus g(x, n)$$

Once these processes are composed with the original component, we can then follow the refinement methodology. In SML, we can easily modify our aforementioned LI system to an LI system containing components with different evaluation cycles. The SML implementation of the two processes is show below:

Listing 8. Multiclock Interface

```

1 fun Insert(n) = fn s1 => h(s1, n)
2 fun h([], _) = [] | h(x1::x, n) =
   let
4   val sig1 = [x1] @ tausall(n)
   in
6   [sig1] @ h(x, n)
   end
8
   fun Strip(n) = fn s1 => f(s1, n)
10 fun f([], _) = [] |
    f(x1::x, n) =
12 let
    fun dr [] = [] | dr(x::xf) = xf
14 fun drop ([], _) = [] | drop(s, 1) = dr(s) |
    drop (s, i) = drop(dr(s), i-1)
16 in
    drop(x1, n) @ f(x, n)
18 end

```

6 Conclusion and Future Work

We propose a functional programming based framework using SML for the validation of LI systems against their original system implementations. The in-

herent denotational model of functional languages makes them well suited to formalize such complex protocols. In this framework, computation within the blocks can be changed without much additional effort whereas in formal verification any change of the model results in time consuming verification runs. We show a refinement methodology that defines how to transform a system consisting of synchronous blocks assuming zero delay communication to a corresponding LI system with long interconnects. An *Eqcomparator* comparator process is modeled that does a latency equivalence check between the outputs of the original system and its LI version given the same input sequence. We model a set of LI components, an *equalizer*, a *splitter* and a *merger* process, with which any deterministic synchronous system can be implemented. We extend this with processes *insert* and *strip* to multiclock systems where IPs with known clock ratios are also allowed.

Another possible extension for this is to allow unknown clock ratios. This would be another step towards being able to handle GALS systems where components can have clocks that are entirely unrelated.

References

- [1] L. Carloni, K. McMillan, and A. Sangiovanni-Vincentelli. The theory of latency insensitive design. *IEEE Transactions on Computer Aided Design of Integrated Circuits and System*, 20(9):1059–1076, 2001.
- [2] M. Casu and L. Macchiarulo. A new approach to latency insensitive design. In *Design Automation Conference*, 2004.
- [3] P. Glaskowski. Pentium 4 (partially) previewed. *Microprocessor Report*, 14(8):10–13, 2000.
- [4] Vikas Agarwal, M. S. Hrishikesh, Stephen W. Keckler, and Doug Burger. Clock rate versus ipc: the end of the road for conventional microarchitectures. *SIGARCH Comput. Archit. News*, 28(2):248–259, 2000.

- [5] L.P. Carloni and A.L. Sangiovanni-Vincentelli. Coping with latency in SoC design. *IEEE Micro, Special Issue on Systems on Chip*, 22(5):12, October 2002.
- [6] Edmund Clarke, Orna Grumberg, and Doron Peled. *Model Checking*. The MIT Press, 2000.
- [7] M. Singh and M. Theobald. Generalized latency-insensitive systems for single-clock and multi-clock architectures. In *Design, Automation and Test in Europe (DATE'04)*, 2004.
- [8] Syed Suhaib, David Berner, Deepak Mathaikutty, Jean-Pierre Talpin, and Sandeep Shukla. Presentation and formal verification of a family of protocols for latency insensitive design. Technical report, Virginia Tech, 2005.
- [9] R. Milner, M. Tofte, R. Harper, and D. MacQueen. *The Definition of Standard ML - Revised*. MIT Press, 1997.
- [10] L. Carloni, K. McMillan, A. Saldanha, and A. Sangiovanni-Vincentelli. A methodology for correct-by-construction latency insensitive design. In *In Proc. International Conf. Computer Aided Verification*, pages 309–315, November 1999.
- [11] L.P. Carloni, K. L. McMillan, and A. L. Sangiovanni-Vincentelli. Latency insensitive protocols. In *11th International Conference on Computer-Aided Verification*, volume 1633, pages 123–133, Trento, Italy, 07 1999. Springer Verlag.
- [12] E. Lee and T. Parks. Dataflow process networks, May 1995.
- [13] Axel Jantsch. *Modeling Embedded Systems and SoCs - Concurrency and Time in Models of Computation*. Morgan Kaufmann, 2001.

APPENDIX

Listing 9. SML Implementation for Adaptive Modulator

```
fun regulate([],_) = [] | regulate(-,[]) = [] |
2   regulate(x1::x, y1::y) =
   if tag(y1) = [1] andalso value(y1) = 1
4   then [[value(x1) - 10,1]] @ regulate(x,y)
   else [x1] @ regulate(x,y)
6
(* #1 and #2 corresponds to the first signal and second signal *)
8 fun Reg() = fn s1 => fn s2 => regulate(s1,s2)
   fun Regulate() = fn s1 => fn s2 => fn st => Reg() (#1(Punzip()
10  (Equalizer() s1 s2 st))) (#2(Punzip() (Equalizer() s1 s2 st)))

12 fun mul(-,[]) = [] | mul([],_) = [] | mul(x::xf, y1::y) =
   if tag(x) = [1] andalso tag(y1) = [1]
14  then ([[value(x) * value(y1),1]] @ mul(xf,y))
   else ([[0,0]] @ mul(xf,y));
16
   fun Con() = fn s1 => fn s2 => mul(s1,s2);
18 fun Convolute() = fn s1 => fn s2 => fn st => Con() (#1(Punzip()
   (EqualizerTwo() s1 s2 st))) (#2(Punzip() (EqualizerTwo() s1 s2 st)))
20
   fun comp(a,b) = if a > b then [1,1] else [0,1]
22 fun compsig([],_) = [] | compsig(x1::x,y1) = if tag(x1) = [1]
   then ([comp(value(x1),y1)] @ compsig(x,y1))
24  else ([[0,0]] @ compsig(x,y1));
   fun Alt(a) = fn s1 => compsig(s1,a)
26 fun Alternate(a) = fn s1 => fn st =>
   Alt(a) (#1(Punzip2() (EqualizerOne() s1 st)))
```
