Copyright © 1999, by the author(s). All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

LATENCY INSENSITIVE PROTOCOLS

by

Luca P. Carloni, Kenneth L. McMillan and Alberto L. Sangiovanni-Vincentelli

Memorandum No. UCB/ERL M99/11

1 February 1999



LATENCY INSENSITIVE PROTOCOLS

by

Luca P. Carloni, Kenneth L. McMillan and Alberto L. Sangiovanni-Vincentelli

Memorandum No. UCB/ERL M99/11

1 February 1999

ELECTRONICS RESEARCH LABORATORY

College of Engineering University of California, Berkeley 94720

Latency Insensitive Protocols

Luca P. Carloni[†]

Kenneth L. McMillan[‡]

Alberto L. Sangiovanni-Vincentelli[†]

[†] University of California at Berkeley Berkeley, CA 94720-1772 [‡] Cadence Berkeley Laboratories Berkeley, CA 94704-1103

Abstract

The theory of latency insensitive design is presented as the foundation of a new correct by construction methodology to design very large digital systems by assembling blocks of Intellectual Properties. Latency insensitive designs are synchronous distributed systems and are realized by assembling functional modules exchanging data on communication channels according to an appropriate protocol. The goal of the protocol is to guarantee that latency insensitive designs composed of functionally correct modules, behave correctly independently of the wire delays. A latency-insensitive protocol is presented that makes use of relay stations buffering signals propagating along long wires. To guarantee correct behavior of the overall system, modules must satisfy weak conditions. The weakness of the conditions makes our method widely applicable.

1 Introduction

The level of integration available today with Deep Sub-Micron (DSM) technologies $(0.1\mu m)$ and below) is so high that entire systems can now be implemented on a single chip. Designs of this kind expose problems that were barely visible at previous levels of integration: the dominance of wire delays on the chip and the strong effects created by the clock skew [1]. It is predicted that a signal will need more than five (and up to more than ten!) clock ticks to traverse the entire chip area. Thus it will be very important to limit the distance traveled by critical signals to guarantee the performance of the design. However, precise data on wire-lengths are available late in the design process and several costly re-designs may be needed to change the placement or the speed of the components of the design to satisfy performance and functionality constraints. We believe that, for deep sub-micron designs where millions of gates are customary, a design method that guarantees by construction that certain properties are satisfied is the only hope to achieve correct designs in short time. In particular, we focus on methods that allow a designer to compose pre-designed and verified components so that the composition formally satisfies certain properties.

In this paper, we present a theory for the design of digital systems that maintains the inherent simplicity of synchronous designs and yet does not suffer of the "long-wire" problem. According to our methodology, the system can be thought as completely synchronous, i.e. just as a collection of modules that communicate by means of channels having a latency of one clock cycle. Unfortunately, the final layout may require more than one clock cycle to transmit the appropriate signals. Our methodology does not require costly redesign cycles or to slow down the clock. The key idea is borrowed from pipelining: partition the long wires into segments whose lengths satisfy the timing requirements imposed by the clock by inserting logic blocks called relay stations, which have a function similar to the one of latches on a pipelined datapath. The timing requirements imposed by the real clock are now met by construction. However, the latency of a channel connecting two modules is generally equal to more than one real clock cycle. If the functionality of the design is based on the sequencing of the output signals and not on their exact timing, then this modification of the design does not change functionality provided that the components of the design are latency insensitive, i.e., the behavior of each module does not depend on the latency of the communication channels. We have essentially traded latency for throughput by not slowing down the clock and by inserting relay stations. In this paper, we introduce these concepts formally and we prove the properties outlined above. The framework used for the theory is the Tagged Signal Model of Lee and Sangiovanni-Vincentelli [3].

The paper is organized as follows: in Section 2 we give the foundation of latency insensitive design by presenting the notion of patient processes. In Section 3 we discuss how in a system of patient processes communication channels can be segmented by introducing *relay stations*. Section 4 illustrates the overall design methodology and discusses under which assumption a generic system can be transformed in a patient one.

2 Latency Insensitivity

To cast our methodology in a formal framework, we use the approach of Lee and Sangiovanni-Vincentelli to represent signals and processes [4].

2.1 The Tagged-Signal Model

Given a set of values V and a set of tags T, an event is a member of $V \times T$. Two events are synchronous if they have the same tag. A signal s is a set of events. Two signals are synchronous if each event in one signal is synchronous with an event in the other signal and vice versa. Synchronous signals must have the same set of tags.

The set of all N-tuples of signals is denoted S^N . A process P is a subset of S^N . A particular N-tuple $s \in S^N$ satisfies the process if $s \in P$. A N-tuple s that satisfies a process is called a *behavior* of the process. Thus a process is a set of possible behaviors ¹. A composition of processes (also called a system) $\{P_1, \ldots, P_M\}$, is a process defined as the intersection of their behaviors $P = \bigcap_{m=1}^M P_m$. Since processes can be defined over different signal sets, to form the composition we need to extend the set of signals over which each process is defined to contain all the signals of all processes. Note that the extension changes the behavior of the processes only formally.

Let $J = (j_1, \ldots, j_h)$ be an ordered set of integers in the range [1, N], the projection of a behavior $b = (s_1, \ldots, s_N) \in S^N$ onto S^h is $proj_J(b) = (s_{j_1}, \ldots, s_{j_h})$. The projection of a process $P \subseteq S^N$ onto S^h is $proj_J(P) = (s' | \exists s \in P \land proj_J(s) = s')$. A connection C is a particularly simple process where two (or more) of the signals in the N-tuple are constrained to be identical: for instance, $C(i, j, k) \subset S^N$: $(s_1, \ldots, s_N) \in C(i, j, k) \Leftrightarrow s_i = s_j = s_k$, with $i, j, k \in [1, N]$.

In a synchronous system every signal in the system is synchronous with every other signal. In a timed system the set T of tags, also called timestamps, is a totally ordered set. The ordering among the timestamps of a signal s induces a natural order on the set of events of s.

2.2 Informative Events and Stalling Events

As outlined in the Introduction, we start with the basic assumption that the designs follow the synchronous design methodology and then we proceed to modify it to allow for delays on long wires. Hence, the basic theory of [4] will now be used in this context.

A latency insensitive system is a synchronous timed system where the set of values \mathcal{V} is equal to $\Sigma \cup \{\tau\}$, where Σ is the set of *informative symbols* which are exchanged among modules and τ is a special symbol, representing the absence of an informative symbol. From now on, all signals are assumed to be synchronous. An event is called *informative* if it has an informative symbol ι_i as value². An event whose value is a τ symbol is said a *stalling event* (or τ event or, simply, *bubble*).

Definition 2.1 $\mathcal{E}(s)$ denotes the set of events of signal s while $\mathcal{E}_{\iota}(s)$ and $\mathcal{E}_{\tau}(s)$ are respectively the set of informative events and the set of stalling events of s. The k-th event (v_k, t_k) of a signal s is denoted $e_k(s)$. T(s) denotes the set of timestamps in signal s, while $T_{\iota}(s)$ is the set of timestamps corresponding to informative events.

Processes exchange "useful" data by sending and receiving informative events. Ideally only informative events should be communicated among processes. However, in a latency insensitive system, a process

¹For $N \ge 2$, processes may also be viewed as a relation between the N signals in s.

²We use subscripts to distinguish among the different informative symbols of Σ : $\iota_1, \iota_2, \iota_3, \ldots$

may not have data to output at a given timestamp thus requiring the output of a stalling event at that timestamp.

Definition 2.2 The set of all sequences of elements in $\Sigma \cup \{\tau\}$ is denoted by Σ_{lat} . The length of a sequence σ is $|\sigma|$ if it is finite, otherwise is infinity. The empty sequence is denoted as ϵ and, by definition, $|\epsilon| = 0$. The *i*-th term of a sequence σ is denoted σ_i .

Definition 2.3 The function $\sigma: S \times T^2 \to \Sigma_{lat}$ takes a signal $s = \{(v_0, t_0), (v_1, t_1), \ldots\}$ and an ordered pair of timestamps $(t_i, t_j), i \leq j$, and returns a sequence $\sigma_{[t_i, t_j]} \in \Sigma_{lat} s.t.^3: \sigma_{[t_i, t_j]}(s) = v_i, v_{i+1}, \ldots, v_j$. The sequence of values of a signal is denoted $\sigma(s)$. The infinite subsequence of values corresponding to an infinite sequence of events, starting from timestamp t_i is denoted $\sigma_{[t_i,\infty]}(s)$.

For example, considering the signal $s = \{(\iota_1, t_1), (\iota_2, t_2), (\tau, t_3), (\iota_2, t_4), (\iota_1, t_5), (\tau, t_6)\}$, we have

$$\sigma(s) = \iota_1 \ \iota_2 \ \tau \ \iota_2 \ \iota_1 \ \tau \qquad \sigma_{[t_2, t_4]}(s) = \iota_2 \ \tau \ \iota_2 \qquad \sigma_{[t_3, t_3]}(s) = \iota_1$$

and, respectively, $|\sigma(s)| = 6$, $|\sigma_{t_2,t_4}(s)| = 3$, $|\sigma_{t_5,t_5}(s)| = 1$. To manipulate sequences of values we define the following filtering operators.

Definition 2.4
$$\mathcal{F}_{\iota}: \Sigma_{lat} \to \Sigma^{\star}$$
 returns a sequence $\sigma' = \mathcal{F}_{\iota}[\sigma]$ s.t. $\sigma'_{i} = \begin{cases} \sigma_{[t_{i},t_{i}]}(s) & \text{if } \sigma_{[t_{i},t_{i}]}(s) \in \Sigma \\ \epsilon & \text{otherwise} \end{cases}$

Definition 2.5 $\mathcal{F}_{\tau}: \Sigma_{lat} \to \{\tau\}^{\star}$ returns a sequence $\sigma' = \mathcal{F}_{\tau}[\sigma]$ s.t. $\sigma'_{i} = \begin{cases} \sigma_{[t_{i},t_{i}]}(s) & \text{if } \sigma_{[t_{i},t_{i}]}(s) = \tau \\ \epsilon & \text{otherwise} \end{cases}$

For instance, if $\sigma(s) = \iota_1 \iota_2 \tau \iota_2 \iota_1 \tau$, then $\mathcal{F}_{\iota}[\sigma(s)] = \iota_1 \iota_2 \iota_1$ and $\mathcal{F}_{\tau}[\sigma(s)] = \tau \tau$. Obviously, $|\sigma(s)| = |\mathcal{F}_{\iota}[\sigma(s)]| + |\mathcal{F}_{\tau}[\sigma(s)]|$. Latency insensitive systems are assumed to have a finite horizon over which informative events appear, i.e., for each signal s there is a greatest timestamp $T \in \mathcal{T}_{\iota}(s)$ which corresponds to the "last" informative event. However, to build our theory we need to extend the set of signals of a latency insensitive system over an infinite horizon by adding a set of timestamps such that all events with timestamp greater than T have τ values. The set of timestamps is assumed to be in one-to-one correspondence with the set \mathbb{N} of natural numbers.

Definition 2.6 A signal s is strict iff all informative events precede all stalling events, i.e., iff there exists a $k \in \mathbb{N}$ s.t. $|\mathcal{F}_{\tau}[\sigma_{[t_0,t_k]}(s)]| = 0$ and $|\mathcal{F}_{\iota}[\sigma_{[t_k,t_{\infty}]}(s)]| = 0$. A signal which is not strict is said to be delayed (or stalled).

2.3 Latency Equivalence

Two signals are latency equivalent if they present the same sequence of informative events, i.e., they are identical except for different delays between two successive informative events. Formally:

Definition 2.7 Two signals s_1 , s_2 are latency equivalent $s_1 \equiv_{\tau} s_2$ if and only if (iff) $\mathcal{F}_{\iota}[\sigma(s_1)] = \mathcal{F}_{\iota}[\sigma(s_2)]$.

The reference signal s_{ref} of a class of latency equivalent signals is a strict signal obtained by assigning the sequence of informative values that characterizes the equivalence class to the first $|\mathcal{F}_{\iota}[\sigma(s_1)]|$ timestamps. For instance, signals s_1 and s_2 presenting the following sequences of values

$$\begin{aligned} \sigma(s_1) &= \iota_1 \ \iota_2 \ \tau \ \iota_1 \ \iota_2 \ \iota_3 \ \tau \ \iota_1 \ \iota_2 \ \tau \ \tau \ \tau \ \ldots \\ \sigma(s_2) &= \iota_1 \ \iota_2 \ \tau \ \tau \ \iota_1 \ \tau \ \iota_2 \ \iota_3 \ \tau \ \iota_1 \ \tau \ \iota_2 \ \tau \ldots \end{aligned}$$

are latency equivalent. Their reference signal s_{ref} is characterized by the sequence

$$\sigma(s_{ref}) = \iota_1 \ \iota_2 \ \iota_1 \ \iota_2 \ \iota_3 \ \iota_1 \ \iota_2 \ \tau \ \tau \ \tau \ldots$$

Latency equivalent signals contain the same sequences of informative values, but with different timestamps. Hence, it is useful to identify their informative events with respect to the common reference signal: the *ordinal* of an informative event coincides with its position in the reference signal.

³Notice that $\sigma_{[t_1,t_1]}(s)$ denotes the value of the event at timestamp t_i .

Definition 2.8 The ordinal of informative event $e_k = (v_k, t_k) \in \mathcal{E}_\iota(s)$ is $ord(e_k) = |\mathcal{F}_\iota[\sigma_{[t_0, t_k]}](s)| - 1$. Let s_1 and q_1 be two latency equivalent signals: two informative events $e_k(s_1) \in \mathcal{E}_\iota(s_1)$ and $e_l(q_1) \in \mathcal{E}_\iota(q_1)$ are said corresponding events iff $ord(e_k(s_1)) = ord(e_l(q_1))$. The slack between two corresponding events is defined as $slack(e_k(s_1), e_l(q_1)) = |k - l|$.

We extend the notion of latency equivalence to behaviors, in a component-wise manner:

Definition 2.9 Two behaviors (s_1, \ldots, s_N) and (s'_1, \ldots, s'_N) are equivalent iff $\forall i \ (s_i \equiv_{\tau} s'_i)$. A behavior $b = (s_1, \ldots, s_N)$ is strict iff every signal $s_i \in b$ is strict. Every class of latency equivalent behaviors contains only one behavior which is strict: this is called the reference behavior.

Definition 2.10 Two processes P_1 and P_2 are latency equivalent, $P_1 \equiv_{\tau} P_2$, if every behavior of one is latency equivalent to some behavior of the other. A process P is strict iff every behavior $b \in P$ is strict. Every class of latency equivalent processes contains only one process which is strict: this is called the reference process.

Definition 2.11 A signal s_1 is latency dominated by s_2 , $s_1 \leq_{\tau} s_2$ iff $s_1 \equiv_{\tau} s_2$ and $T_1 \leq T_2$, where T_i is the greatest timestamp with an informative value for s_i , i = 1, 2.

Hence, referring to the previous example, signal s_1 is dominated by signal s_2 since $T_1 = 9$ while $T_2 = 12$. Notice that a reference signal is latency dominated by every signal belonging to its equivalence class. Latency dominance is extended to behaviors and processes as in the case of latency equivalence. A total order among events of a behavior is necessary to develop our theory. In particular, we introduce an ordering among events that is motivated by causality: events that have smaller ordinal are ordered before the ones with larger ordinal (think of a strict process where the ordinal is related to the timestamp; the order implies that past events do not depend on future events). In addition, to avoid cyclic behaviors created by processing events with the same ordinal, we assume that there is an order among signals. This order in real-life designs corresponds to input-output dependencies. We cast this consideration in the most general form possible to extend maximally the applicability of our method.

Definition 2.12 Given a behavior $b = (s_1, \ldots, s_N)$, \leq_c denotes a well-founded order on its set of signals. The well-founded order induces a lexicographic order \leq_{lo} over the set of informative events of b, s.t. for all pairs of events (e_1, e_2) with $e_1 \in \mathcal{E}_{\iota}(s_i)$ and $e_2 \in \mathcal{E}_{\iota}(s_j)$

 $e_1 \leq_{lo} e_2 \iff [(ord(e_1) < ord(e_2)) \lor ((ord(e_1) = ord(e_2)) \land (s_i \leq_c s_i))]$

The following function returns the first informative event (in signal s_j of behavior b) which follows a given event $e \in b$ with respect to the lexicographic order \leq_{lo} .

Definition 2.13 Given a behavior $b = (s_1, \ldots, s_N)$ and an informative event $e(s_i) \in \mathcal{E}_{\iota}(s_i)$, the function nextEvent is defined as: nextEvent $(s_j, e(s_i)) = \min_{e_k(s_j) \in \mathcal{E}_{\iota}(s_j)} \{e(s_i) \leq_{lo} e_k(s_j)\}$

A stall move is used to postpone an informative event of a signal of a given behavior by one timestamp. The stall move is used to account for long delays along wires and to add delays where needed to guarantee functional correctness of the design.

Definition 2.14 Given a behavior $b = (s_1, \ldots, s_j, \ldots, s_N)$ and an informative event $e_k(s_j) = (v_k, t_k)$, a stall move returns a behavior $b' = stall(e_k(s_j)) = (s_1, \ldots, s'_j, \ldots, s_N)$, s.t. for $l \in \mathbb{N}$:

$$\sigma_{[t_0,t_{k-1}]}(s'_j) = \sigma_{[t_0,t_{k-1}]}(s_j), \qquad \sigma_{[t_k,t_k]}(s'_j) = \tau, \qquad \sigma_{[t_{k+l+1},t_{k+l+1}]}(s'_j) = \sigma_{[t_{k+l},t_{k+l}]}(s_j)$$

A procrastination effect represents the "effect" of a stall move $stall(e_k(s_j))$ on other signals of behavior b in correspondence of events following $e_k(s_j)$ in the lexicographic order. The processes will "respond" to the insertion of stalls in some of their signals "delaying" other signals that are causally related to the stalled signals.

Definition 2.15 A procrastination effect is a point-to-set map which takes a behavior $b' = (s'_1, \ldots, s'_N) = stall(e_k(s_j))$ resulting from the application of a stall move on event $\epsilon_k(s_j)$ of behavior $b = (s_1, \ldots, s_N)$ and returns a set of behaviors $\mathcal{PE}[stall(e_k(s_j))]$ s.t. $b'' = (s''_1, \ldots, s''_N) \in \mathcal{PE}[b']$ iff

- $s''_{j} = s'_{j};$
- $\forall i \in [1, N], i \neq j, s_i'' \equiv_{\tau} s_i' \text{ and } \sigma_{[t_0, t_{l-1}]}(s_i'') = \sigma_{[t_0, t_{l-1}]}(s_i'), \text{ where } t_l \text{ is the timestamp of event } e_l(s_i) = nextEvent(s_i, e_k(s_j));$
- $\exists K \text{ finite s.t. } \forall i \in [1, N], i \neq j, \exists k_i \leq K, \sigma_{[t_{i+k_i}, \infty]}(s''_i) = \sigma_{[t_i, \infty]}(s'_i).$

Each behavior in $\mathcal{PE}[b']$ is obtained from b' by possibly inserting other stalling events in any signal of b', but only at "later" timestamps, i.e. to postpone informative event which follow $e_k(s_j)$ with respect to the lexicographic order \leq_{lo} . Observe that a procrastination effect returns a behavior that latency dominates the original behavior.

2.4 Patient Processes

We are now ready to define the notion of patient process: a patient process can take stall moves on any signal of its behaviors by reacting with the appropriate procrastination effects. Patience is the key condition for the IP blocks to be combinable according to our method.

Definition 2.16 A process P is patient iff

 $\forall b = (s_1, \dots, s_N) \in P, \forall j \in [1, N], \forall e_k(s_j) \in \mathcal{E}_\iota(s_j), (\mathcal{PE}[stall(e_k(s_j))] \cap P \neq \emptyset)$

Hence, the result of a stall move on one of the events of a patient process may not satisfy the process, but one of the behaviors of the procrastination effect corresponding to the stall move does satisfy the process, i.e., if we stall a signal on an input to a functional block, the block will be forced to delay some of its outputs or if we request an output signal to be delayed then an appropriate delay has to be added to the inputs.

Lemma 2.1 Let P_1 and P_2 be two patient processes. Let $b_1 \in P_1$, $b_2 \in P_2$ be two behaviors with the same lexicographic order s.t. $b_1 \equiv_{\tau} b_2$. Then, there exists a behavior $b' \in (P_1 \cap P_2)$, $b_1 \equiv_{\tau} b' \equiv_{\tau} b_2$.

Proof. The proof is constructive. Let $b_1 = (r_1, \ldots, r_N) \in P_1$ and $b_2 = (q_1, \ldots, q_N) \in P_2$ be the two behaviors with the same lexicographic order. Since b_1 and b_2 are latency equivalent, each event in b_1 has a corresponding event in b_2 and vice versa (see definition 2.8). Let $r_1 \equiv_{\tau} q_1, \ldots, r_N \equiv_{\tau} q_N$. Let $W = \{w \mid \exists k \in \mathbb{N}, \exists l \in \mathbb{N} \ (k \neq l \land ord(e_k(r_j)) = ord(e_l(q_j)) = w)\}$ be the set of ordinals associated to pairs of corresponding events of b_1 and b_2 whose timestamps differ. Define the distance between behaviors b_1, b_2 as

$$d(b_1, b_2) = \begin{cases} \max_{w \in W} \{\frac{1}{2^w}\} & \text{if } W \neq \emptyset \\ 0 & \text{otherwise} \end{cases}$$

This distance is reminiscent of the Cantor metric. Thus, b_1 and b_2 have distance equal to zero if all pairs of corresponding events are aligned ⁴. In this case, b_1 and b_2 are identical, i.e. they are the same behavior that belongs to $(P_1 \cap P_2)$. Now suppose that $d(b_1, b_2) = \frac{1}{2^{w_1}} \neq 0$: in this case, w_1 is the smallest ordinal among those which are associated to unaligned pairs of corresponding events. Without loss of generality, let $p_1 = (e_k(r_j), e_l(q_j))$ be the pair of corresponding events whose ordinal is equal to w_1 and let l > k. Apply a stall move to $e_k(r_j)$ to obtain a new behavior $b'_1 = (s'_1, \ldots, s'_N) = stall(e_k(r_j)) \equiv_{\tau} b_1$. Obviously, $slack(e_{k+1}(s'_j), e_l(q_j)) = slack(e_k(r_j), e_l(q_j)) - 1$. Note that b'_1 is not necessarily a behavior of P_1 . However, since P_1 is patient, there exists $b''_1 = (s''_1, \ldots, s''_N) \equiv_{\tau} b_1$ s.t. $b'' \in \mathcal{PE}(stall(e_k(r_j)) \cap P_1$. Since, by definition of procrastination effect, $s''_j = s'_j$, then also $slack(e_{k+1}(s''_j), e_l(q_j)) = slack(e_k(r_j), e_l(q_j)) - 1$. Since the procrastination effect may postpone only events following $e_k(r_j)$ in the lexicographic order \leq_{lo} , then all the pairs of corresponding events of b''_1 and b_2 with ordinal smaller than w are still aligned.

⁴A pair of corresponding event is said *aligned* if the events are synchronous, or, according to def. 2.8, if their slack is 0.

Now, there are two possibilities: if $slack(e_{k+1}(s''_j), e_l(q_j)) = 0$, then one more pair has been aligned and $d(b''_1, b_2) < d(b_1, b_2)$; otherwise, we can reduce by 1 this slack by repeating the same procedure starting from behavior b''_1 . In any case, after l - k steps of the procedure outlined above, we obtain a behavior $b_1^* \equiv_{\tau} b_1$ that satisfies P_1 and s.t. $d(b_1^*, b_2) < d(b_1, b_2)$, because one more pair of corresponding events has been aligned. We say that we have performed an *alignment step*.

Now, if $d(b_1^*, b_2) = 0$ then there are no more unaligned pairs, the two behaviors are identical and the lemma is proven since $b_1^* \equiv_{\tau} b_1$. Instead, if $d(b_1^*, b_2) = \frac{1}{2^{w_2}} \neq 0$ then we consider the next unaligned pair p_2 of corresponding events and we execute a second alignment step. Note that at the *m*-th step, while aligning pair p_m with ordinal w_m , we may increase the slack of some of the pairs following p_m in the lexicographic order, but we keep aligned all the pairs preceding p_m . During this sequence of alignment steps we obtain two sequences of behaviors (one of behaviors in P_1 latency equivalent to b_1 and one of behaviors in P_2 latency equivalent to b_2), whose distance is decreasing monotonically. Since both b_1 and b_2 contain the same finite ⁵ number of informative events the set U of pairs of unaligned corresponding events is also finite. The slack of each of these pair is also a finite number. At the m-th step, we have at most h_m sub-steps to align p_m , where h_m is the starting slack for p_m . In the worst case, each behavior b^* obtained during the sub-steps of the alignment step may have slacks of all the remaining unaligned pairs increased by at most K (see definition 2.15). Hence, at the end of the m-th step |U| has been decreased by one, while all the slacks of its remaining elements have been increased by at most $h_m \cdot K$, a finite number. Thus for $|U| \ge l > m$, the new slacks for the remaining unaligned pairs is $h'_l \le h_l + h_m \cdot K$. Globally, we perform in the worst case |U| alignment steps and for each of them we have a finite number of sub-steps. Hence, the two sequences of behaviors are also finite and the last elements of these sequences do not have unaligned pairs, and, therefore, have distance equal to zero.

Theorem 2.1 If P_1 and P_2 are patient processes then $(P_1 \cap P_2)$ is a patient process.

Proof. Let $b = (s_1, \ldots, s_N)$ be a behavior in $P_1 \cap P_2$. Consider behaviors $b_1 = (r_1, \ldots, r_N) \in P_1$ and $b_2 = (q_1, \ldots, q_N) \in P_2$, s.t. $b_1 = b_2 = b$. For all $j \in [1, N]$ and for all $k \in \mathbb{N}$, let $e_k(s_j) \in \mathcal{E}_\iota(s_j)$. Since $b_1 = b_2 = b$, then $e_k(r_j) \in \mathcal{E}_\iota(r_j)$ and $e_k(q_j) \in \mathcal{E}_\iota(q_j)$. Let $b' = stall(e_k(s_j)) \equiv_\tau b$. Similarly, $b'_1 = stall(e_k(r_j)) \equiv_\tau b_1$ and $b'_2 = stall(e_k(s_j)) \equiv_\tau b_2$. Since P_1 is patient there exists a behavior $b''_1 \equiv_\tau b_1$ s.t. $b''_1 \in \mathcal{PE}[b'_1] \cap P_1$ and since P_2 is patient there exists a behavior $b''_2 \equiv_\tau b_2$ s.t. $b''_2 \in \mathcal{PE}[b'_2] \cap P_2$. Notice that $b_1 = b_2$ implies that $b''_1 \equiv_\tau b''_2$, however, it is not necessarily the case that $b''_1 = b''_2$. In fact, procrastination effects may have misaligned pairs of corresponding informative signals which come after $(e_k(r_j), e_k(s_j))$ with respect to lexicographic order \leq_{lo} . Since $b_1 = b_2$ share the same lexicographic order, by lemma 2.1, there exists a behavior $b'' \equiv_\tau b''_1 \equiv_\tau b''_2$ s.t. $b'' \in P_1 \cap P_2$. The construction of b'' given in the proof of lemma 2.1 involves only unaligned pairs of corresponding events between b''_1 and b''_2 and all these unaligned pairs correspond to informative events which come after $e_k(s_j)$ with respect to lexicographic order \leq_{lo} . Further, since the number of informative events is finite, the number of unaligned pairs is also finite. Hence, each signal s''_i of b'' is obtained by inserting a finite number of stalling events not earlier than timestamp t_l , with $e_l(s_l) = nextEvent(s_i, e_k(s_j))$. Therefore, by definition 2.15, $b'' \in \mathcal{PE}(stall(e_k(s_j))$. Since we have already seen that $b'' \in P_1 \cap P_2$, then $(P_1 \cap P_2)$ is a patient process.

The following theorem shows that, for patient processes, the notion of latency equivalence of processes is compositional. Figure 1 illustrates the proof for the case when the two behaviors are just 1-tuple signals.

Theorem 2.2 For all patient processes P_1, P_2, P'_1, P'_2 , if $P_1 \equiv_{\tau} P'_1$ and $P_2 \equiv_{\tau} P'_2$ then

$$(P_1 \cap P_2) \equiv_{\tau} (P_1' \cap P_2')$$

Proof. Let $b = (s_1, \ldots, s_N)$ be a behavior in $P_1 \cap P_2$. Latency equivalence implies that there must be behaviors $b_1 = (r_1, \ldots, r_N) \in P'_1$ and $b_2 = (q_1, \ldots, q_N) \in P'_2$ such that $b_1 \equiv_{\tau} b \equiv_{\tau} b_2$. Since b_1 and b_2 are latency equivalent and P'_1 and P'_2 are patient, lemma 2.1 guarantees that there must be a latency equivalent behavior $b' \in (P'_1 \cap P'_2)$. The other direction of the proof is symmetric.

Therefore, we can replace any process in a system of patient processes by a latency equivalent process, and the resulting system will be latency equivalent. A similar theorem holds for the replacement of strict processes with patient processes.

⁵Recall that the number of informative events for every behavior considered in latency insensitive designs is finite.



Figure 1: Sketch for proof on compositionality of latency equivalence.

Theorem 2.3 For all strict processes P_1 , P_2 and patient processes P'_1 , P'_2 , if $P_1 \equiv_{\tau} P'_1$ and $P_2 \equiv_{\tau} P'_2$ then

 $(P_1 \cap P_2) \equiv_{\tau} (P_1' \cap P_2')$

Proof. The argument that every behavior in $(P_1 \cap P_2)$ has an equivalent in $(P'_1 \cap P'_2)$ is as in theorem 2.2. For the other direction, let b' be a behavior in $P'_1 \cap P'_2$. Latency equivalence implies that there must be behaviors $b_1 \in P_1$ and $b_2 \in P_2$ such that $b_1 \equiv_{\tau} b' \equiv_{\tau} b_2$. Since P_1 and P_2 are strict, b_1 and b_2 are also strict. Being latency equivalent, they must therefore be equal. Thus $b_1 \in (P_1 \cap P_2)$.

This means that we can replace all processes in a system of strict processes by corresponding patient processes, and the resulting system will be latency equivalent. This is the core of our idea: take a design based on the assumption that computation in one functional block and communication among blocks "take no time" (synchronous hypothesis) ⁶, i.e., the processes corresponding to the functional blocks and their composition are strict, and replace it with a design where communication does take time (more than one clock cycle) and, as a result, signals are delayed, but without changing the sequence of informative events observed at the system level, i.e., with a set of patient processes.

3 Latency Insensitive Design

As explained in Section 1, one of the goal of the latency insensitive design methodology is to be able to "pipeline" a communication channel by inserting an arbitrary amount of memory elements. In the framework of our theory, this operation corresponds to adding some particular processes, called *relay stations*, to the given system. In this section, we first show how patient systems (i.e. systems of patient processes) are insensitive to the insertion of relay stations and, then, we discuss under which assumption a generic system can be transformed into a patient system.

3.1 Channels and Buffers

A channel is a connection ⁷ constraining two signals to be identical.

Definition 3.1 A channel $C(i, j) \subset S^N$, $i, j \in [1, N]$ is a process s.t. $b = (s_1, ..., s_N) \in C(i, j) \Leftrightarrow s_i = s_j$.

Lemma 3.1 A channel $C(i, j) \subset S^N$ is not a patient process.

Proof: Let $b = (s_1, ..., s_N)$ be a behavior of a channel C(i, j) and, without loss of generality, suppose that $s_i \leq_c s_j$. Consider a pair of corresponding informative events in s_i and s_j : $e_k(s_i) = (v_1, t_k)$ and $e_k(s_j) = (v_2, t_k)$. Since $b \in C(i, j)$ then $s_i = s_j$ and, therefore, $v_1 = v_2 \neq \tau$. Moreover, $s_i = s_j$ implies that $ord(e_k(s_i)) = ord(e_k(s_j))$ and, since $s_i \leq_c s_j$, $e_k(s_j) = nextEvent(s_j, e_k(s_i))$. Without loss of generality, suppose that $e_k(s_i)$ and $e_k(s_j)$ are followed by (l-1) > 0 stalling events, i.e., formally, that for l > 1, $|\mathcal{F}_t[\sigma_{[t_k, t_{k+(l-1)}]}(s_i)] | = |\mathcal{F}_t[\sigma_{[t_k, t_{k+(l-1)}]}(s_i)] | = 0$. Then, consider informative event $e_{k+l}(s_i) = (v_{k+l}, t_{k+l})$. By definition 2.13, $e_{k+l}(s_i) = nextEvent(s_i, e_k(s_j))$. Now, let $b' = (s'_1, ..., s'_N) = stall(e_k(s_j))$ be the behavior obtained by applying a stall move on $e_k(s_j)$. At timestamp t_k, s'_j presents

⁶In other words, communication and computation are completed in one clock cycle.

⁷See section 2.1 for the definition of connection.

a stalling event, while the event of s'_j corresponding to $e_k(s_j)$ is $e_{k+1}(s'_j) = (v_2, t_{k+1})$, which occurs at timestamp t_{k+1} . Then, consider any behavior $b'' = (s''_1, ..., s''_N) \in \mathcal{PE}[b']$. By definition 2.15, since $e_{k+l}(s_i) = nextEvent(s_i, e_k(s_j))$, then $\sigma_{[t_0, t_k]}(s''_i) = \sigma_{[t_0, t_k]}(s_i) = \sigma_{[t_0, t_k]}(s_i)$. In particular, $\sigma_{[t_k, t_k]}(s''_i) = \sigma_{[t_k, t_k]}(s_i) \neq \tau$ and therefore, $\sigma_{[t_k, t_k]}(s''_i) \neq \sigma_{[t_k, t_k]}(s''_j)$, which, finally, implies that $s''_i \neq s''_j$. Hence, $\forall b'' \in \mathcal{PE}[b']$ (b'' $\notin C(i, j)$) and, by definition 2.16, C(i, j) is not patient.

Definition 3.2 A buffer $B_{l_f,l_b}^c(i,j)$ with capacity $c \ge 0$, minimum forward latency $l_f \ge 0$ and minimum backward latency $l_b \ge 0$ is a process s.t. $\forall i, j \in [1, N]$: $b = (s_1, ..., s_N) \in B_{l_f,l_b}^c(i, j)$ iff $(s_i \equiv_{\tau} s_j)$ and $\forall k \in \mathbb{N}$

$$0 \leq |\mathcal{F}_{\iota}\left[\sigma_{\left[t_{0},t_{k-l_{f}}\right]}\left(s_{i}\right)\right]| - |\mathcal{F}_{\iota}\left[\sigma_{\left[t_{0},t_{k}\right]}\left(s_{j}\right)\right]|$$
(1)

$$c \geq |\mathcal{F}_{\iota}\left[\sigma_{\left[t_{0},t_{k}\right]}\left(s_{i}\right)\right]| - |\mathcal{F}_{\iota}\left[\sigma_{\left[t_{0},t_{\left(k-t_{k}\right)}\right]}\left(s_{j}\right)\right]|$$

$$(2)$$

By definition, given a pair of indexes $i, j \in [1, N]$, for all $l_b, l_f, c \ge 0$, all buffers $B_{l_f, l_b}^c(i, j)$ are latency equivalent. Observe also that buffer $B_{0,0}^0(i, j)$ coincides with channel C(i, j). In particular, we are interested in buffers having unitary latencies and we want to establish under which conditions such buffers are patient processes.

Lemma 3.2 If s_i, s_j are two signals s.t. $s_i \equiv_{\tau} s_j$ and $s_i \leq_c s_j$, then

- $\forall g \in \mathbb{N} \ s.t. \ e_g(s_i) \in \mathcal{E}_\iota(s_i), \ nextEvent(s_j, e_g(s_i)) \ is \ the \ corresponding \ event \ of \ e_g(s_i) \ in \ s_j.$
- $\forall h \in \mathbb{N} \text{ s.t. } e_h(s_j) \in \mathcal{E}_i(s_j), \text{ nextEvent}(s_i, e_h(s_j)) = nextEvent(s_i, e_g(s_i)), \text{ where } e_g(s_i) \text{ is the corresponding event of } e_h(s_j) \text{ in } s_i.$

Theorem 3.1 Let $l_b = l_f = 1$. For all $c \ge 1$, $B_{1,1}^c(i, j)$ is a patient process iff $s_i \le_c s_j$,

Proof.: First, if $l_b = l_f = 1$ then inequalities (1) and (2) become:

$$0 \leq |\mathcal{F}_{\iota}[\sigma_{[t_0,t_{(k-1)}]}(s_i)]| - |\mathcal{F}_{\iota}[\sigma_{[t_0,t_k]}(s_j)]|$$
(3)

$$c \geq |\mathcal{F}_{\iota}[\sigma_{[t_0,t_k]}(s_i)]| - |\mathcal{F}_{\iota}[\sigma_{[t_0,t_{(k-1)}]}(s_j)]|$$
(4)

["only if" part]: By contradiction: we prove that if $s_i \not\leq_c s_j$ then $B_{1,1}^c(i,j)$ is not a patient process. Suppose $s_i \not\leq_c s_j$. For all $c \geq 1$, let $b = (s_1, \ldots, s_N)$ be a behavior of $B_{1,1}^c(i,j)$ s.t. $\sigma(s_i) = \iota_0 \iota_1 \tau \tau \tau \ldots$ and $\sigma(s_j) = \tau \iota_0 \iota_1 \tau \tau \tau \tau \ldots$ Let $b' = (s'_1, \ldots, s'_N) = stall(e_0(s_i))$, with $e_0(s_i) = (\iota_0, t_0)$. Clearly, $b' \notin B_{1,1}^c(i,j)$, because inequality (3) does not hold for k = 1 since $s'_j = s_j$. Further, for all $b'' = (s''_1, \ldots, s''_N) \in \mathcal{PE}(stall(e_0(s_i)))$ we can prove that $b'' \notin B_{1,1}^c(i,j)$. In fact, since $s''_i = s'_i$, $b'' \in B_{1,1}^c(i,j)$ iff $\sigma_{[t_1,t_1]}(s''_j) = \tau$. But, consider that $ord(e_0(s_i)) = ord(e_1(s_j))$ and, since $s_i \not\leq_c s_j$, then $e_0(s_i) \not\leq_{lo} e_1(s_j)$. Further, $ord(e_0(s_i)) = ord(e_2(s_j)) - 1$. Therefore, $e_2(s_j) = nextEvent(s_j, e_0(s_i))$. Recall that, by definition 2.15 of procrastination effect, $\sigma_{[t_0,t_{l-1}]}(s''_j) = \sigma_{[t_0,t_{l-1}]}(s'_j)$. Since $s'_j = s_j$, $\sigma_{[t_1,t_1]}(s''_j) = \tau_0 \neq \tau$. This implies that $b'' \notin B_{1,1}^c(i,j)$. Hence, $\mathcal{PE}[stall(e_0(s_i))] \cap B_{1,1}^c(i,j) = \emptyset$ and $B_{1,1}^c(i,j)$ is not patient.

["if" part]: We prove that if $s_i \leq_c s_j$ then $B_{1,1}^c(i,j)$ is patient.

For all $c \ge 1$, let $b = (s_1, \ldots, s_N)$ be a behavior of $B_{1,1}^c(i, j)$. For all $g \in \mathbb{N}$, such that $e_g(s_i) \in \mathcal{E}_\iota(s_i)$, let $b' = (s'_1, \ldots, s'_N) = stall(e_g(s_i)) \equiv_\tau b$. Since $s'_j = s_j$, $b' \notin B_{1,1}^c(i, j)$ iff inequality (3) does not hold for some $k \in \mathbb{N}$. In fact, b' satisfies the other two conditions of definition 3.2, because $b' \equiv_\tau b$ and to insert a stalling event on s_i (while s_j remains the same) can not induce a violation of inequality (4). Now, suppose first that b' satisfies also inequality (3) for all $k \in \mathbb{N}$: then, there exists at least a behavior which belongs to $\mathcal{P}\mathcal{E}[stall(e_g(s_i))] \cap B_{1,1}^c(i, j)$ and this behavior is b', because, $\forall g \forall i, stall(e_g(s_i)) \in \mathcal{P}\mathcal{E}[stall(e_g(s_i))]$. A more interesting case is when inequality (3) does not hold: in this case $b' \notin B_{1,1}^c(i, j)$. Then, consider a behavior $b'' = (s''_1, \ldots, s''_N)$ s.t. $\forall n \in [1, N], n \neq j, (s''_n = s'_n)$, while s''_j is obtained from s'_j by inserting a stalling event at timestamp t_h , where t_h is also the timestamp of event $e_h(s_j) = nextEvent(s_j, e_g(s_i))$. Clearly, this construction guarantees that $b'' \in \mathcal{P}\mathcal{E}[stall(e_g(s_i))]$. It remains to be proven that $b'' \in B_{1,1}^c(i, j)$. First, by construction, $b'' \equiv_\tau b$. Then, check whether s''_i, s''_j satisfy inequalities (3) and (4) for all $k \in \mathbb{N}$. First, since $s_i \equiv_{\tau} s_j$ and $s_i \leq_c s_j$, by lemma 3.2, $e_h(s_j) = nextEvent(s_j, e_g(s_i))$ is the corresponding event of $e_g(s_i)$ in s_j . Hence $ord(e_g(s_i)) = ord(e_h(s_j)) = ord(e_g(s''_i)) = ord(e_h(s''_j))$ and, recalling definition 2.8, $| \mathcal{F}_i[\sigma_{[t_0,t_g]}(s_i)| = | \mathcal{F}_i[\sigma_{[t_0,t_h]}(s_j)|$. Since, by hypothesis, s_i, s_j satisfy inequality (3) for all $k \in \mathbb{N}$. then g < h. Compare s''_i and s''_j respectively with s_i and s_j : s''_i has been derived by s_i inserting a τ at t_g , while s''_j has been derived by s_j inserting a τ at t_h . Hence, we can derive the following 4 equations. Further each term in these equations can be bounded using the fact that s_i, s_j satisfy inequalities (3) and (4) for all $k \in \mathbb{N}$:

$$\forall k \in [0, g-1], \quad |\mathcal{F}_{\iota}[\sigma_{[t_0, t_k]}(s''_{\iota})| = |\mathcal{F}_{\iota}[\sigma_{[t_0, t_k]}(s_{\iota})| \leq |\mathcal{F}_{\iota}[\sigma_{[t_0, t_{k-1}]}(s_{\iota})| + c \tag{5}$$

$$\forall k \in [0, h-1], \quad |\mathcal{F}_{i}[\sigma_{[t_{0}, t_{k}]}(s_{i}^{*})| = |\mathcal{F}_{i}[\sigma_{[t_{0}, t_{k-1}]}(s_{i})| \leq |\mathcal{F}_{i}[\sigma_{[t_{0}, t_{k-2}]}(s_{j})| + c$$

$$\forall k \in [0, h-1], \quad |\mathcal{F}_{i}[\sigma_{[t_{0}, t_{1}]}(s_{i}^{*})| = |\mathcal{F}_{i}[\sigma_{[t_{0}, t_{1}]}(s_{i})| \leq |\mathcal{F}_{i}[\sigma_{[t_{0}, t_{1}]}(s_{i})|$$

$$(7)$$

$$\forall k \in [h, \infty[, |\mathcal{F}_{\iota}[\sigma_{[t_0, t_k]}(s''_{j})| = |\mathcal{F}_{\iota}[\sigma_{[t_0, t_k-1]}(s_{j})| \leq |\mathcal{F}_{\iota}[\sigma_{[t_0, t_k-2]}(s_{i})|$$
(8)

Now, keeping in mind that g < h, it is easy to prove that:

- using inequality (7) and equation (5), s''_i, s''_j satisfy inequality (3), $\forall k \in [0, g-1]$.
- using ⁸ inequality (7) and equation (6), s''_i, s''_j satisfy inequality (3), $\forall k \in [g, h-1]$.
- using inequality (8) and equation (6), s''_i, s''_i satisfy inequality (3), $\forall k \in [h, \infty[$.
- using inequality (5) and equation (7), s''_i, s''_j satisfy inequality (4), $\forall k \in [0, g-1]$.
- using inequality (6) and equation (8), s''_i, s''_j satisfy inequality (4), $\forall k \in [g, h-1[$.
- using inequality (6) and equation (8), s''_i, s''_j satisfy inequality (4), $\forall k \in [h, \infty[$.

Therefore, $b'' \in B_{1,1}^c(i,j)$.

Consider now $\dot{b}' = (s'_1, \ldots, s'_N) = stall(e_h(s_j)) \equiv_{\tau} b$, where for all $h \in \mathbb{N}$, $e_h(s_j) \in \mathcal{E}_\iota(s_j)$. Let $e_q(s_i) = nextEvent(s_i, e_h(s_j))$ and $e_p(s_i)$ be the corresponding event of $e_h(s_j)$ in s_i : then, since $s_i \equiv_{\tau} s_j$ and $s_i \leq_c s_j$, by lemma 3.2, $e_q(s_i) = nextEvent(s_i, e_p(s_i))$. Now, construct $b'' = (s''_1, \ldots, s''_N)$ in such a way that $\forall n \in [1, N], n \neq i, (s''_n = s'_n)$, while s''_i is obtained from s'_i by inserting a stalling event at timestamp t_g , where $g = \min_{k \in [h+1, \infty[} \{k \mid e_k(s_i) \in \mathcal{E}_\iota(s_i)\}$. Hence, if q > h then $e_g(s_i) = e_q(s_i)$ else $e_q(s_i) \leq_{lo} e_g(s_i)$. In both cases, this construction guarantees that $b'' \in \mathcal{PE}[stall(e_h(s_j))]$. It remains to be proven that $b'' \in B_{1,1}^c(i,j)$. First, by construction, $b'' \equiv_{\tau} b$. Then, check whether s''_i, s''_j satisfy inequalities (3) and (4) for all $k \in \mathbb{N}$. Compare s''_i and s''_j respectively with s_i and $s_j: s''_i$ has been derived by s_i inserting a τ at t_g , while s''_j has been derived by s_j inserting a τ at t_g . Now, keeping in mind that here h < g, it is easy to prove that:

- using inequality (7) and equation (5), s''_i, s''_j satisfy inequality (3), $\forall k \in [0, h-1]$.
- using inequality (8) and equation (5), s''_i, s''_j satisfy inequality (3), $\forall k \in [h, g-1]$.
- using inequality (8) and equation (6), s''_i, s''_j satisfy inequality (3), $\forall k \in [g, \infty[$.
- using inequality (5) and equation (7), s''_i, s''_j satisfy inequality (4), $\forall k \in [0, h-1]$.
- using 9 inequality (5) and equation (7), s''_i, s''_j satisfy inequality (4), $\forall k \in [h, g-1[$.
- using inequality (6) and equation (8), s''_i, s''_i satisfy inequality (4), $\forall k \in [g, \infty[$.

Therefore, in this case too $b'' \in B_{1,1}^c(i,j)$.

Finally, for all $n \in ([1, N]/\{i, j\})$ let $b' = (s'_1, \ldots, s'_N) = stall(e_h(s_n)) \equiv_{\tau} b$, where for all $h \in \mathbb{N}$, $e_h(s_j) \in \mathcal{E}_{\iota}(s_n)$. Then, trivially, $b' \in \mathcal{PE}[stall(e_k(s_n))] \cap B^c_{1,1}(i, j)$. In conclusion, we have that $\forall b = (s_1, \ldots, s_N) \in B^c_{1,1}(i, j), \forall n \in [1, N], \forall e_k(s_n) \in \mathcal{E}_{\iota}(s_n), (\mathcal{PE}[stall(e_k(s_n))] \cap B^c_{1,1}(i, j) \neq \emptyset)$. Hence, $B^c_{1,1}(i, j)$ is patient.

⁸Recall that $\sigma_{[t_k,t_k]}(s_i) = \tau$.

⁹Recall that $\sigma_{[t_h, t_h]}(s_j) = \tau$.

Consider a strict system $P_{strict} = \bigcap_{m=1}^{M} P_m$ with N strict signals s_1, \ldots, s_N . As explained in section 2.1, processes can be defined over different signal sets and to compose them we may need to formally extend the set of signals of each process to contain all the signals of all processes. However, without loss of generality, consider the particular case of composing M processes which are already defined on the same N signals. Hence, any generic behavior $b_m = (s_{m_1}, \ldots, s_{m_N})$ of P_m is also a behavior of P_{strict} iff for all $l \in [1, M], l \neq m$ process P_l contains a behavior $b_l = (s_{l_1}, \ldots, s_{l_N})$ s.t. $\forall n \in [1, N] (s_{l_n} = s_{m_n})$. In fact, we may assume to derive system P_{strict} by connecting the processes with $(M-1) \cdot N$ channel processes $C(l_n, (l+1)_n)$, where $l \in [1, (M-1)]$ and $n \in [1, N]$. Further, we may also assume to "decompose" any channel process $C(m_n, l_n)$ with an arbitrary number X of channel processes $C(m_n, x_1), C(x_1, x_2), \ldots, C(x_{X-1}, l_n)$, by adding X-1 auxiliary signal, each of them forced to be equal to $m_n = l_n$. The theory developed in section 2 guarantees that if we replace each process $P_m \in P_{strict}$ with a latency equivalent patient process and each channel C(i, j) with a patient buffer $B_{1,1}^1(i, j)$ we obtain a system P_{patient} which is patient and latency equivalent to P_{strict}. In fact, "having a patient buffer in a patient system is equivalent to having a channel in a strict system". Since "decomposing" a channel C(i, j) has no observable effect on a strict system, we are therefore free to add an arbitrary number of patient buffers into the corresponding patient system to replace this channel. Since we use patient buffers with unitary latencies, we can distribute them along that long wire on the chip which implements C(i, j), in such a way that the wire gets decomposed in segments whose physical lengths can be spanned in a single physical clock cycle.

3.2 Relay Stations

The following Lemma 3.3 proves that no behaviors in $B_{1,1}^1(i, j)$ may contain two informative events of s_i, s_j which are synchronous: this implies that the maximum throughput across such a buffer is 0.5, which may be considered suboptimal. Instead, buffer $B_{1,1}^2(i, j)$ is the minimum capacity buffer which is able to "transfer" one informative unit per timestamp (maximum throughput = 1).

Lemma 3.3 Buffer $B_{1,1}^2(i,j)$ is the minimum capacity buffer with $l_f = l_b = 1$ s.t.

$$\exists b^* = (s_1^*, \dots, s_N^*) \in B_{1,1}^2(i,j) \land \exists k \in \mathbb{N}, \ (e_k(s_i^*) \in \mathcal{E}_\iota(s_i^*) \land e_k(s_i^*) \in \mathcal{E}_\iota(s_i^*))$$
(9)

Proof.: Relation (9) says that $B_{1,1}^2(i,j)$ is the minimum capacity buffer with $l_f = l_b = 1$ containing a behavior b^* where s_i and s_j present at least a pair of synchronous informative event (i.e., the two informative event have the same timestamp t_k). Notice that the only buffer with $l_f = l_b = 1$ having capacity less than $B_{1,1}^2(i,j)$ is $B_{1,1}^1(i,j)$. We first show that $B_{1,1}^2(i,j)$ contains at least one behavior b^* satisfying relation (9) and then we prove that the same is not true for any behavior of $B_{1,1}^1(i,j)$. It is easy to construct an example of such a behavior: for instance, consider a behavior $b = (s_1, \ldots, s_N)$ s.t. $\sigma(s_i^*) = \iota_1 \iota_2 \tau \tau \tau \ldots$ and that $\sigma(s_j^*) = \tau \iota_1 \iota_2 \tau \tau \tau \ldots$. Clearly, $s_i^* \equiv_{\tau} s_j^*$ and inequalities (3) and (4) are satisfied for any $k \in \mathbb{N}$. Hence, $b \in B_{1,1}^2(i,j)$. Moreover, at timestamp t_1 both s_i^* and s_j^* present an informative event.

Now, consider $B_{1,1}^1(i,j)$. If c = 1, combining inequalities (3) and (4) we obtain that $\forall k \in \mathbb{N}$:

$$|\mathcal{F}_{\iota} \left[\sigma_{[t_0, t_{(k-1)}]} \left(s_j \right) \right] |+ 1 \ge |\mathcal{F}_{\iota} \left[\sigma_{[t_0, t_k]} \left(s_i \right) \right] |\ge |\mathcal{F}_{\iota} \left[\sigma_{[t_0, t_{(k+1)}]} \left(s_j \right) \right] | |\mathcal{F}_{\iota} \left[\sigma_{[t_0, t_{(k-1)}]} \left(s_i \right) \right] |\ge |\mathcal{F}_{\iota} \left[\sigma_{[t_0, t_k]} \left(s_j \right) \right] |\ge |\mathcal{F}_{\iota} \left[\sigma_{[t_0, t_{(k+1)}]} \left(s_i \right) \right] |- 1$$

Hence, for all behaviors $b = (s_1, \ldots, s_N) \in B_{1,1}^1(i, j)$, signals s_i, s_j are not only latency equivalent but also correlated according to a very regular pattern (see Figure 2) which can be summarized in two properties: (i) there are no two synchronous informative events in s_i, s_j ; (ii) for all timestamps, informative events appear alternately on s_i and on s_j . Property (i) is a negation of relation (9).

Definition 3.3 The buffer $B_{1,1}^2$ is called a relay station RS.

 $B_{1,1}^{1} \begin{cases} s_{1} = \iota_{1} \tau \iota_{2} \tau \iota_{3} \tau \iota_{4} \tau \tau \tau \iota_{5} \tau \iota_{6} \tau \iota_{7} \tau \iota_{8} \tau \iota_{9} \tau \tau \tau \iota_{10} \tau \dots \\ s_{2} = \tau \iota_{1} \tau \iota_{2} \tau \iota_{3} \tau \iota_{4} \tau \tau \tau \iota_{5} \tau \iota_{6} \tau \iota_{7} \tau \iota_{8} \tau \tau \tau \iota_{9} \tau \iota_{10} \tau \dots \\ B_{1,1}^{2} \begin{cases} s_{1} = \iota_{1} \iota_{2} \iota_{3} \tau \tau \iota_{4} \iota_{5} \iota_{6} \tau \tau \tau \iota_{7} \tau \iota_{8} \iota_{9} \iota_{10} \dots \\ s_{2} = \tau \iota_{1} \iota_{2} \iota_{3} \tau \tau \iota_{4} \iota_{5} \iota_{6} \tau \tau \tau \iota_{7} \tau \iota_{8} \iota_{9} \iota_{10} \dots \\ s_{2} = \tau \iota_{1} \iota_{2} \iota_{3} \tau \tau \iota_{4} \tau \tau \tau \iota_{5} \iota_{6} \iota_{7} \tau \iota_{8} \iota_{9} \iota_{10} \dots \end{cases} \end{cases}$

Figure 2: Comparing two possible behaviors of finite buffers $B_{1,1}^1$ and $B_{1,1}^2$.

4 Latency Insensitive Design Methodology

In this section, we move towards the implementation of the theory introduced in the previous sections. To do so, we assume that:

- the pre-designed functional blocks are synchronous processes;
- there is a set of signals for each process that can be considered as inputs to the process and a set of signals that can be considered as outputs of the process, i.e., the processes are *functional*;
- the processes are strictly causal (a process is *strictly causal* if two outputs can only be different at timestamps that strictly follow the timestamps when the inputs producing these outputs show a difference ¹⁰).
- the processes belong to a particular class of processes called *stallable*, a weak condition to ask the processes to obey.

The basic ideas are as follows. Composing a set of pre-designed synchronous functional blocks in the most efficient way is fairly straightforward if we assume that the synchronous hypothesis holds. This composition corresponds to a composition of strict processes since there is a priori no need of inserting stalling events. However, as we argued in the introduction, it is very likely that the synchronous hypothesis be not valid for communication. If indeed the processes to be composed are patient, then adding an appropriate number of relay stations yields a process that is latency equivalent to the strict composition. Hence, if we use as the definition of correct behavior the fact that the sequences of informative events do not change, the addition of the relay stations solve the problem. However, requiring processes to be patient at the onset is quite strong. However, in practice, a patient system can be *derived* from a strict one as follows: first, we take each strict process P_m and we compose it with a set of auxiliary processes to obtain an equivalent patient process P'_m . To be able to do so, all processes P_m must satisfy a simple condition (the processes must be stallable) specified in the next section. Then, we put together all patient processes by connecting them with relay stations. The set of auxiliary processes implements a "queuing mechanism" across the signal of P_m in such a way that informative events are buffered and reordered before being passed to P_m : informative events having the same ordinal are passed to P_m synchronously.

In the sequel, we first introduce the formal definition of functional processes. Then, we present the simple notion of stallable processes and we prove that every stallable process can be encapsulated into a wrapper process which acts as an interface towards a latency insensitive protocol.

4.1 Stallable Processes

An *input* to a process $P \subseteq S^N$ is an externally imposed constraint $P_I \subseteq S^N$ such that $P_I \cap P$ is the total set of acceptable behaviors. Commonly, one considers processes having input signals and output signals: in this case, given process P, the set of signals can be partitioned into three disjoint subsets by partitioning the index set as $\{1, \ldots, N\} = I \cup O \cup R$, where I is the ordered set of indexes for the input signals of P, O is the ordered set of indexes for the output signals and R is the ordered set of indexes for the remaining signals (also called irrelevant signals with respect to P). A process is functional with respect to (I, O) if for every behavior $b \in P$ and $b' \in P$, where $proj_I(b) = proj_I(b')$, it follows that $proj_O(b) = proj_O(b')$.

¹⁰ For a more formal definition see [3].

$s_1 = \iota_1 \iota_3 \iota_1 \tau \iota_3 \tau \tau \ldots$	$s_4 = \tau \iota_1 \tau \iota_3 \iota_1 \tau \iota_3 \ldots$
$s_2 = \tau \iota_4 \tau \iota_7 \iota_8 \tau \iota_8 \ldots -$	$s_5 = \tau \iota_4 \tau \iota_7 \iota_8 \tau \iota_8 \ldots$
$s_3 = \tau \iota_5 \iota_5 \tau \iota_9 \tau \iota_6 \ldots$	$s_6 = \tau \iota_5 \tau \iota_5 \iota_9 \tau \iota_6 \ldots$

Figure 3: Example of a behavior of an equalizer E with $I = \{1, 2, 3\}$ and $O = \{4, 5, 6\}$.

In the sequel, we consider only strictly causal processes and for each of them we assume that the well founded order \leq_c of definition 2.12 subsumes the causality relations among its signals, i.e. formally: $\forall i \in I, \forall j \in O, (s_i \leq_c s_j)$.

Definition 4.1 A process P with $I = \{1, ..., Q\}$ and $O = \{Q + 1, ..., N\}$ is stallable iff for all $b = (s_1, ..., s_Q, s_{Q+1}, ..., s_N) \in P$ and for all $k \in \mathbb{N}$:

 $\forall i \in I \ (\sigma_{[t_k, t_k]}(s_i) = \tau) \ \Leftrightarrow \ \forall j \in O \ (\sigma_{[t_{k+1}, t_{k+1}]}(s_j) = \tau)$

Hence, while a patient process tolerates arbitrary distributions of stalling events among its signals (as long as causality is preserved), a stallable process demands more regular patterns: τ symbols can only be inserted synchronously (i.e., with the same timestamp) on all input signals and this insertion implies the synchronous insertion of τ symbols on all output signals at the following timestamp. To assume that a functional process is stallable is quite reasonable with respect to a practical implementation. In fact, most hardware systems can be stalled: for instance, any sequential logic block that has a gated clock or a Moore finite state machine M with an extra input, that, if equal to τ , forces M to stay in the current state and to emit τ at the next cycle.

4.2 Encapsulation of Stallable Processes

Now, our goal is to define a group of functional processes that can be composed with a stallable process P to derive a patient process which is latency equivalent to P. We start considering a process that aligns all the informative events across a set of channels.

Definition 4.2 An equalizer E is a process, with $I = \{1, \ldots, Q\}$ and $O = \{Q+1, \ldots, 2 \cdot Q\}$, s.t. for all behaviors $b = (s_1, \ldots, s_Q, s_{Q+1}, \ldots, s_{2 \cdot Q}) \in E$: $\forall i \in I, (s_i \equiv_{\tau} s_{Q+i})$ and $\forall k \in \mathbb{N}$:

$$\begin{aligned} \forall i, j \in O \ (\ (\sigma_{[t_k, t_k]}(s_i) = \tau) \ \Rightarrow \ (\sigma_{[t_k, t_k]}(s_j) = \tau) \) \\ \min_{i \in I} \ \{ \ | \ \mathcal{F}_\iota \ [\sigma_{[t_0, t_k]}(s_i)] \ | \ \} \ - \ \max_{i \in O} \ \{ \ | \ \mathcal{F}_\iota \ [\sigma_{[t_0, t_k]}(s_j)] \ | \ \} \ \ge 0 \end{aligned}$$

The first relation forces the output signals to have stalling events only synchronously, while the second guarantees that at every timestamp the number of informative events occurred at any input is always greater than the number of informative events occurred at any output. In particular, the presence of a stalling event at any input at a given timestamp forces the presence of a stalling event on all outputs at the same timestamp. Figure 3 illustrates a possible behavior of an equalizer.

Definition 4.3 An extended relay station \mathcal{ERS} is a process with $I = \{i\}$ and $O = \{j, l\}$, $i \neq j \neq l$ s.t. signals s_q, s_2 are related by inequalities (1) and (2) of definition 3.2 (with $l_f = l_b = 1$ and c = 2) and $\forall k \in \mathbb{N}$:

$$\sigma_{[t_k,t_k]}(s_l) = \begin{cases} 1 & if \mid \mathcal{F}_{\iota}\left[\sigma_{[t_0,t_k]}(s_i)\right] \mid - \mid \mathcal{F}_{\iota}\left[\sigma_{[t_0,t_{k-1}]}(s_j)\right] \mid = 2\\ 0 & otherwise \end{cases}$$

Definition 4.4 A stalling signal generator SSG is a process with $I = \{1, \ldots, Q\}$ and $O = \{Q+1\}$ s.t. $\forall b = (s_1, \ldots, s_{Q+1}), \forall k \in \mathbb{N}, \forall i \in [1, Q], (\mathcal{F}_i[\sigma_{[t_k, t_k]}(s_i)] \in [0, 1])$ and

$$\sigma_{[t_k,t_k]}(s_{Q+1}) = \begin{cases} \tau & \text{if } \exists j \in [1,Q] \ (\mathcal{F}_{\iota} \ [\sigma_{[t_k,t_k]}(s_j)] = 1 \) \\ 0 & \text{otherwise} \end{cases}$$

As illustrated in Figure 4, any stallable process P can be composed with an equalizer, a stalling signal generator and some extended relay stations to derive a patient process which is latency equivalent to P.

Definition 4.5 Let P be a stallable process with $I_P = \{p'_1, \ldots, p'_M\}$ and $O_P = \{q'_1, \ldots, q'_N\}$. A wrapper process (or, shell process) W(P) of P is the process with $I_W = \{p_1, \ldots, p_M\}$ and $O_W = \{q_1, \ldots, q_N\}$ which is obtained composing P with the following processes:

- an equalizer E with $I_E = \{p_1, \ldots, p_M, p_{M+1}\}$ and $O_E = \{p'_1, \ldots, p'_M, p'_{M+1}\},\$
- N extended relay stations $\mathcal{ERS}_1, \mathcal{ERS}_2, \dots, \mathcal{ERS}_N$ s.t. $I_j = \{q'_i\}$ and $O_j = \{q_j, r_j\}$, with $j \in [1, N]$
- a stalling signal generator SSG with $I_G = \{r_1, \ldots, r_N\}$ and $O_G = \{p_{M+1}\}$.

Theorem 4.1 Let W(P) be the wrapper process of definition 4.5. Process $W = \operatorname{proj}_{I_W \cup O_W}(W(P))$ is a patient process that is latency equivalent to P.

Proof: Throughout the proof we follow the index notation of definition 4.5.

[" $W \equiv_{\tau} P$ " part]: We first prove that $W \equiv_{\tau} P$. Let $b' = (s_{p_1'}, \ldots, s_{p_M'}, s_{q_1'}, \ldots, s_{q_N'})$ be a behavior of P and $b = (s_{p_1}, \ldots, s_{p_{M+1}}, s_{q_1}, \ldots, s_{q_N}, s_{p_1'}, \ldots, s_{p_{M+1}'}, s_{q_1'}, \ldots, s_{q_N'}, s_{r_1}, \ldots, s_{r_N})$ a behavior of W(P). Further, let $b_W = proj_{I_W \cup O_W}(b) = (s_{p_1}, \ldots, s_{p_M}, s_{q_1}, \ldots, s_{q_N})$ be the corresponding behavior of W. Then, by definition 4.2 of equalizer, $\forall i \in I_E$, $(s_i \equiv_{\tau} s_{i}')$, and, by definition of relay station, $\forall i \in [1, N]$, $(s_{q_i} \equiv_{\tau} s_{q_i'})$. Therefore, $b_W \equiv_{\tau} b$.

["W patient" part]: Recalling definition 2.16, we need to prove that: $\forall b = (s_{p_1}, \ldots, s_{p_M}, s_{q_1}, \ldots, s_{q_N}) \in W, \ \forall j \in I_W \cup O_W, \ \forall e_k(s_j) \in \mathcal{E}_\iota(s_j), \ (\ \mathcal{PE}[stall(e_k(s_j))] \cap W \neq \emptyset)$. Consider first stalling any input signal of W: for all $s_{p_1}, i \in [1, M]$ and all $e_g(s_{p_1}) \in \mathcal{E}_\iota(s_{p_1})$, let $b' = (s'_{p_1}, \ldots, s'_{p_M}, s'_{q_1}, \ldots, s'_{q_N}) = stall(e_g(s_{p_i}))$. By analyzing the interrelationships among the components of W, it is easy to verify that $b' \notin W$. In fact, the insertion of a stalling event on input s_{p_1} at t_g implies that:

- 1. all the output signals of equalizer E have a stalling event at t_g ,
- 2. all the output signals of P have a stalling event at t_{g+1} and
- 3. all the output signals s_{q_1}, \ldots, s_{q_N} have a stalling event at t_{g+2} .

Hence, $\forall j \in N, e_{g+2}(s_{q_j}) \in \mathcal{E}_{\iota}(s_{q_j})$ must be also stalled. Then, since move $stall(e_k(s_{p_i}))$ does not affect any other signal but $s_{p_i}, b' \notin W$. However, since $ord(e_{g+2}(s_{q_j})) = nextEvent((s_{q_j}), e_g(s_{p_i}))$ the insertion of one stalling event on each of the wrapper outputs at t_{k+2} is compatible with the definition of procrastination effect and, therefore, $\mathcal{PE}[stall(e_k(s_{p_i}))] \cap W \neq \emptyset$.

Next, consider stalling any output signal of W: for all s_{q_j} , $j \in [1, N]$ and all $e_h(s_{q_j}) \in \mathcal{E}_t(s_{q_j})$, let $b' = (s'_{p_1}, \ldots, s'_{p_M}, s'_{q_1}, \ldots, s'_{q_N}) = stall(e_h(s_{q_j}))$. By definition of stall move, $\forall m \in [1, M]$, $(s'_{p_m} = s_{p_m})$ and $\forall n \in [1, N]$, $n \neq j$, $(s'_{q_n} = s_{q_n})$. Hence, again, $b' \notin W$. In fact, the insertion of a stalling event on signal s_{q_j} at t_h has an impact on signal s_{r_j} of \mathcal{ERS}_j , set to 1 at t_h and signal $s_{q'_j}$, constrained to stall event $e_{h+l}(s_{q'_j}) \in \mathcal{E}_t(s_{q'_j})$. Note that, by definition 3.3, $e_{h+l}(s_{q'_j})$ must be stalled even though $\forall k \in [h+1, h+l-1], (\sigma_{[t_k,t_k]}(s_{q'_j}) = \tau)$. Therefore, without loss of generality, we may assume l = 1. Then, $s_{r_j} = 1$ forces the output $s_{p_{M+1}}$ of SSG to be a τ symbol at t_h : this implies that all the equalizer outputs have a stalling event at t_h (but, no stalling events are forced on the remaining inputs of SSG) and, finally, all the outputs of the stallable process must have a stalling event at t_{h+l} according to constraint $\mathcal{F}_t [\sigma_{[t_{h+l},t_{h+l}]}(s_{q'_j})] = \tau$. While no other stalling events are forced on $s_{q'_j}$ of \mathcal{ERS}_j at t_{h+l+1} , all the remaining relay stations $\mathcal{ERS}_r, r \in [1, N]/\{j\}$ must stall their $e_{h+l+1}(s_{q_r}) \in \mathcal{E}_t(s_{q_r})$. Again, since $stall(e_h(s_{q_j}))$ does not affect any other signal but s_{q_j} , $b' \notin W$. However, $\forall r \in ([1, N]/\{j\})$, since $ord(e_{h+l+1}(s_{q_r})) = nextEvent(s_{q_r}, e_g(s_{q'_r}))$, where $e_g(s_{q'_r}) = nextEvent(s_{q'_r}, e_h(s_{q_j}))$, then $e_h(s_{q_j}) \leq l_0$ $e_{h+l+1}(s_{q_r})$. Hence, the insertion of procrastination effect. Therefore, $\mathcal{PE}[stall(e_h(s_{q_i}))] \cap W \neq \emptyset$.

Our latency insensitive design methodology can then be summarized as follows:



Figure 4: Encapsulation of a stallable process P into a wrapper W(P).

- 1. Begin with a system of M stallable processes and N channels.
- 2. Encapsulate each stallable process to yield a wrapper process.
- 3. Using relay stations decompose each channel in segments whose physical length can be spanned in a single physical clock cycle.

This approach clearly "orthogonalizes" computation and communication: in fact, we can build systems by putting together hardware cores (which can be arbitrarily complex as far as they satisfy the stalling assumption) and wrappers interfacing them with the channels, by "speaking" the latency insensitive protocol. While the specific functionality of the system is distributed in the cores, the wrappers can be automatically generated around them ¹¹. Finally, the validation of the system can now be efficiently decomposed based on assume-guarantee reasoning [5, 2]: each wrapper is verified assuming a given protocol, and the protocol is verified separately.

5 Conclusions and Future Work

A new design methodology for large digital systems implemented in DSM technology has been presented. The methodology is based on the assumption that the design is built by assembling blocks of Intellectual Properties (IPs) that have been designer and verified previously. The main goal is to develop a theory for the composition of the IP blocks that *ensures* the correctness of the overall design. The focus is on timing properties since DSM designs suffer (and will continue to suffer even more for the foreseeable future) from delays on long wires that often cause costly redesigns. Designs carried out with our methodology are called latency insensitive design. Latency insensitive designs are synchronous distributed systems and are realized by assembling functional modules exchanging data on communication channels according to a latencyinsensitive protocol. The protocol guarantees that latency insensitive designs composed of functionally correct modules, behave correctly independently of the wire delays. This allow us to pipeline long wires by inserting special memory elements called relay stations. The protocol works on the assumption that the functional block satisfy certain weak properties.

The method trades-off latency for throughput, hence it is important to optimize the amount of latency that we must allow to obtain correct designs. This optimization leads to the concept of speculative latency insensitive protocols which will be the subject of a future paper.

We are in the process of applying the method to industrial strength designs to demonstrate its applicability and its properties to the digital design community. Preliminary results indicate that the method yields highly efficient designs and fully delivers on its promises.

¹¹This is the reason why wrappers are also called *shells*: they just "protect" the intellectual property (*the pearl*) they contain from the "troubles" of the external communication architecture.

6 Acknowledgments

We wish to acknowledge the discussions with Luciano Lavagno and Alex Saldanha that led to the theory of latency insensitive designs. Patrick Scaglia gave us strong support based on his experience as a designer of highly complex digital systems and continuous encouragement. Finally this research has been partially sponsored by Cadence Design Systems, SRC and by CNR.

References

- D. Matzke. Will Physical Scalability Sabotage Performance Gains? IEEE Computer, 8(9):37-39, September 1997.
- [2] T.A. Henzinger, S. Qadeer, and R.K. Rajamani. You Assume, We Guarantee: Methodology and Case Studies. In Proceedings of the 10th International Conference on Computer-Aided Verification, Vancouver, Canada, July 1998.
- [3] E. A. Lee and A. Sangiovanni-Vincentelli. Comparing Models of Computation. In Proc. Intl. Conf. on Computer-Aided Design, pages 234-241. IEEE, November 1996.
- [4] E. A. Lee and A. Sangiovanni-Vincentelli. A Framework for Comparing Models of Computation. *IEEE Transactions on Computer-Aided Design*, 17(12):1217-1229, December 1998.
- [5] K. L. McMillan. A Compositional Rule for Hardware Design Refinement. In Proceedings of the 9th International Conference on Computer-Aided Verification, Haifa, Israel, July 1997.