

# Design, Implementation, and Validation of a New Class of Interface Circuits for Latency-Insensitive Design

Cheng-Hong Li, Rebecca Collins, Sampada Sonalkar, and Luca P. Carloni  
Department of Computer Science - Columbia University in the City of New York

**Abstract**—With the arrival of nanometer technologies wire delays are no longer negligible with respect to gate delays, and timing-closure becomes a major challenge to System-on-Chip designers. Latency-insensitive design (LID) has been proposed as a “correct-by-construction” design methodology to cope with this problem. In this paper we present the design and implementation of a new class of interface circuits to support LID that offers substantial performance improvements with limited area overhead with respect to previous designs proposed in the literature. This claim is supported by the experimental results that we obtained completing semi-custom implementations of the three designs with a 90nm industrial standard-cell library. We also report on the formal verification of our design: using the NuSMV model checker we verified that the RTL synthesizable implementations of our LID interface circuits (relay stations and shells) are correct refinements of the corresponding abstract specifications according to the theory of LID.

## I. INTRODUCTION

One of the most critical issues in designing Systems-on-Chip (SOC) with nanometer technology processes is the increasing impact of global wire delays: as more and smaller processing cores are accommodated on a chip, global (inter-core) wires do not scale in delay as local (intra-core) wires do because they need to span physical distances that represent significant proportions of the die [1], [2]. As the delays of global wires are no longer negligible compared to gate delays, the chip becomes a distributed system, thereby posing a serious challenge to the traditional CAD flows that are based on the synchronous design paradigm [3]. Furthermore, since wire delays are hard to predict at early stages of the design process, an increasing number of design exceptions in terms of post-layout timing violations forces costly design re-iterations (*timing-closure problem*).

*Latency-insensitive design (LID)* [4], [5], has been proposed as a “correct-by-construction” design methodology to handle the increasing impact of global communication latency in nanometer integrated circuit design without forcing major departures from traditional and well-established design flows. Given a synchronous system specification, e.g. a register-transfer level (RTL) netlist of logic blocks specified and validated using a hardware-description language, a functionally-equivalent latency-insensitive system can be automatically derived by encapsulating each sequential logic block (referred as a *pearl* or *core*) within an automatically generated interface process (a *shell*). The advantage of this transformation is that any communication channel connecting two core/shell pairs can now present a varying latency in terms of number of clock cycles without affecting the functional correctness of the original design. In practice the latency of a channel is changed through the insertion of *relay stations*, that are clocked buffers with twofold storage capacity and simple

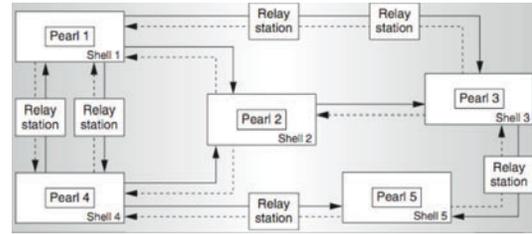


Fig. 1. Shell encapsulation, relay station insertion, and channel back-pressure.

flow-control logic. Hence, LID provides a sound way to address the problem of interconnect delay in nanometer design by simplifying the application of wire pipelining for global communication channels at any stage of the design process and without requiring any re-design of the cores. Furthermore, it simplifies the assembly and reuse of pre-designed cores for building complex SOCs because these can be arbitrarily complex sequential logic blocks as long as they are *stallable*: this is the only prerequisite for LID and it can be easily implemented with *clock gating* mechanisms [4], [5].

In practice, the LID methodology calls for three steps: (1) a strictly synchronous (or *strict*) system is originally designed and validated as a netlist of stallable cores; (2) a *patient* system is automatically derived from the strict system by encapsulating each core within a shell; (3) any number of relay stations can be inserted on any channel between any pair of shells. Fig. 1 shows a latency-insensitive system with five core-pearl pairs connected by point-to-point, unidirectional channels. The shell logic and relay stations implement a *latency-insensitive protocol*, which is designed to accommodate any variation of channels’ latency while guaranteeing that the functional behavior of the original strict system is preserved (*semantics preservation*).

A formal definition of the properties of relay stations and shells is given in a denotational framework as part of the theory of LID [5]. At the core of LID lies the notion of latency-equivalence: two signals are latency equivalent if they present the same ordered streams of data items but possibly with different timing. In a synchronous model of computation the existence of a clock guarantees a common time reference among signals and, therefore, a signal must present an event at each clock cycle [6], [7]. LID distinguishes between the occurrence of an informative event (a valid data item or *valid* token) and a stalling event (*void* token). Any class of latency-equivalent signals contains a single reference signal that does not present stalling events (a *strict* signal) while all the other members of the equivalence class (stalling signals) contain the same sequence of informative events interleaved by one or more stalling events. Following the tagged-signal model [7],

		1	2	3	4	5
LID-2ss	data	A	B	C	C	...
	void	0	0	0	0	...
	stop	0	1	1	1	...
	receiver stalled	0	1	1	1	...
	sender stalled	0	0	0	1	...
<hr/>						
LID-1ss	data	A	B	B	...	...
	void	0	0	0	...	...
	stop	0	1	0	...	...
	receiver stalled	0	1	1	...	...
	sender stalled	0	0	1	...	...

Fig. 2. Simulations of the two latency-insensitive protocols with different back-pressure mechanisms.

the notions of latency-equivalence signals, strict signals, and stalling signals are extended to sets of signals (behaviors) and sets of behaviors (processes) [5].

In a nutshell, LID allows to derive from the original reference strict system specification, which contains only strict processes, any possible latency-equivalent implementation, which contains only patient processes. Each strict process abstracts the core in the original specification while the corresponding latency-equivalent patient process is obtained by composing the core with a shell. While the original cores are not designed to process void tokens, a shell-core pair is a patient process, i.e. it can tolerate the arrival of a void token at any of its I/O channel ports at any given clock cycle and be able to eventually continue with its correct operations.

In a practical implementation, void tokens are used to capture latency variations on communication channels and are processed by the shells in a way that makes them transparent to the cores. In particular, relay stations, which are not present in the original strict design, are initialized with void tokens when introduced in the patient design to pipeline a given channel. Void tokens are then processed by the shells while remaining transparent to the cores. Informally, any shell acts according to an *AND-firing policy*, thereby it stalls its core whenever at least a valid token is missing on one of its input channels. As a shell stalls its core, potential valid tokens that may be present on other input channels are stored locally in input queues within the shell for future processing by the core. In this way each shell dynamically absorbs the latency variations across the channels by realigning the valid tokens before presenting them to the core. Whenever it is not stalled, the core processes valid tokens on its inputs as it does in the original strict system.

Since in practice a queue can only have a finite size, a downlink shell must be able to inform an uplink shell that is necessary to postpone the production of valid token for some cycles (*backpressure*). In the denotational framework of theory of LID, a backpressure event at a given clock cycle is also abstracted as the occurrence of a void token on the channel between the two shells [5]. While the theory of LID defines the general properties that any latency-insensitive protocol must obey, many possible protocol specifications and supporting interface-circuit implementations are conceivable in practice. A protocol that relies on just two control bits, a *void* bit to identify invalid data and *stop* bit to implement backpressure, was first presented in [4] and discussed in more detail together with the supporting interface circuits in [3], [8].

**Contribution.** The latency-insensitive protocol that is dis-

cussed in [3], [4], [8] stipulates that a shell or relay station is stalled whenever the *stop* bit is kept high for *two consecutive* clock cycles. In this paper we refer to this protocol as LID-2ss, which stands for *two-stop-to-stall*. The top of Fig. 2 reports a simulation trace of a channel according to LID-2ss where the receiver is being stalled at cycle 2. Because the receiver is stalled, valid token *A* is not processed and thus is buffered by the receiver’s shell. To avoid buffer overflow and possible loss of the data, the receiver stalls the sender by asserting the *stop* bit both at cycle 2 and 3. Notice that the sender only stalls at cycle 4 holding the valid token *C* on its output port after receiving two *stop* signals. This means token *B* needs to be buffered by a queue in the receiving shell together with token *A*. In fact, both the shell queues and the relay stations have storage capacity equal to two according to the library of interface circuits that were proposed to support LID-2ss.

In this paper we describe a simpler latency-insensitive protocol labeled as LID-1ss, which stands for *one-stop-to-stall*, that is based on a different back-pressure convention. In the new protocol, a shell or a relay station stalls whenever it receives a single *stop* signal, as reported by the simulation trace in the bottom part of Fig. 2: here, the receiver asserts the stop bit only at cycle 2, and the sender begins to stall immediately at cycle 3. In our design a queue of capacity equal to one in the receiver’s shell is sufficient since only data token *A* must be buffered there during stalling while *B* is preserved uplink in the channel for future processing. Notice that our new protocol LID-1ss does not allow us to reduce the storage capacity of a relay station to one because this would reduce the performance of a latency-insensitive system by half as explained in the theory of LID [5]. However, it does allow us to reduce the storage capacity of a shell input queue to one with respect to the original protocol LID-2ss because we can take advantage of the storage capacity within the core <sup>1</sup>.

We contribute a new set of interface circuits (i.e. shells and relay stations) that support the LID-1ss protocol and offer substantial improvements with respect to previous works in the literature. In particular,

- they offer shorter logic delay and have smaller area overhead than the circuits supporting the original latency-insensitive protocol LID-2ss discussed in [3], [4], [8];
- they offer shorter logic delay and, for many systems, enable higher processing throughput than the interface circuits for *synchronous elastic architectures* that were recently proposed in [10].

We also report on our work to validate both our design and the original design: using the NuSMV model checker we formally verified that the RTL synthesizable implementations of the key LID building blocks (relay stations and shells) is a correct refinement of the corresponding abstract specifications according to the theory of LID [5].

The paper is organized as follows. In Sec. II we briefly overview the related work on latency-insensitive design. The RTL logic of the interface circuits supporting our LID-1ss

<sup>1</sup>To discuss how the performance of a latency-insensitive system can be optimized through relay-station insertion and the sizing of shell input queues goes beyond the scope of this paper and we refer to [9].

protocol is described in detail in Sec. III. We then discuss the formal verification of these circuits in Sec. IV. Finally, in Sec. V. we present a comprehensive set of experimental results that provide a comparative analysis of LID-1ss, LID-2ss, and SEA in terms of logic delay, effect on system’s processing throughput, and area overhead.

## II. RELATED WORK

The LID methodology has recently raised some interests and several extensions and related approaches have been proposed [10]–[15]. Indeed, while it specifies the fundamental properties of any latency-insensitive protocol, the denotational framework used to develop the theory of LID [5] leaves open the possibility of developing various protocol specifications that in turn may lead to practical implementations with different characteristics.

The simpler protocol that we discuss in this paper was already assumed in [13], [14]. Chelcea and Nowick presented a mixed-timing relay station that stalls for one clock cycle if a stop signal is received [13]. As they focus on describing a complete class of low-latency FIFO interfaces for mixed-timing systems, they do not discuss the design of shell blocks to support LID. Lu and Koh use max-plus algebra to analyze the performance of a latency-insensitive system with back-pressure [14]. The model of the protocol that they adopt assumes that a sender is stalled when one or more of its receivers asserts the stop bit. However, neither the design of the shell nor the design of a relay station is provided. Conversely, in this paper we contribute the complete interface logic for a single-clock synchronous system at the RTL level.

Cortadella et al. recently proposed synchronous elastic architectures (SEAs) [10] that are based on the synchronous elastic flow (SELF) protocol; SELF is a new approach to LID that “combines the modularity of asynchronous design with the efficiency of synchronous implementations” [10]. Like the LID-2ss protocol that was originally proposed in [4] and the LID-1ss one that we discuss in the present paper, SELF also relies on *valid* and *stop* bits. Further, SEAs rely on sequential buffers, called *elastic buffers (EB)*, to pipeline long channel wires, as LID relies on relay stations. On the other hand, SEAs do not use the idea of shell interfaces with input queues that store valid tokens during stalling. Instead, in a SEA it is possible to have elastic buffers with multiple input/output channels thanks to special elastic *fork* and *join* control structures [10]: when stalling occurs, each valid but unused token is held by its immediate sender. Robustness with respect to latency variations is achieved in SEAs by combining elastic buffers, fork and join structures while performing an *elasticization* transformation on the original circuit. This step consists essentially of replacing each flip-flop in the core with two transparent latches of different polarity, similar to a master-slave structure, but with independent enable signals for the two latches so that “a mechanism for double-pumping in one cycle” [10] can be realized. By properly setting the enable signals the elasticized core can either operate as usual, or be stalled, or store two output data in the two back-to-back latches. However, using enable signals to control

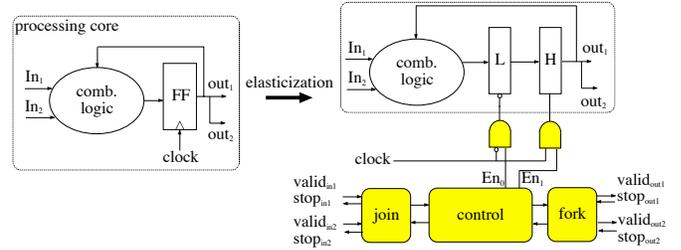


Fig. 3. Elasticizing a core with on SEA interface circuits and clock gating.

clocked latches may incur significant area overhead because additional steering logic is needed [16]. In this paper, a slight modification is made in the SEA interface circuits: the latches are driven by *gated* clock signals to avoid extra steering logic for stalling the core and storing two unconsumed data tokens. This technique was first proposed by Jacobson et al. for their *synchronous interlocked pipelines* [16]. The elasticization of a processing core is illustrated in Fig. 3, where the shaded boxes represent the logic implementing the SEA interface circuits and stalling mechanism. In particular the join control structures differ subtly from LID-1ss interface circuits with respect to the timing of sending a *stop* bit to a sender. In a LID-1ss interface this is sent whenever a queue is full. Instead the join control structure of a processing core with multiple input channels requests all valid tokens to be resent (by asserting the corresponding *stop* bits) whenever at least one invalid tokens arrives at the *same* clock cycle. This may have negative impacts on the performance of a SEA because: (a) it degrades the overall system throughput and (b) it limits the maximum clock frequency at which the final circuit can run due to long combinational paths spanning two interconnect channels. In Section V we present a detailed discussion of these issues in the context of a comparative analysis of the interface circuits for the two approaches.

Suhaib et al. [17] propose a framework for validating families of latency-insensitive protocols by taking a system, transforming it into a latency-insensitive system and then comparing the output behavior of the original system with the one of the transformed system on a subset of possible inputs. This technique is good for the development and debugging phase of new latency-insensitive protocols because it can uncover many bugs quickly without requiring an exhaustive verification. As described in Sec. IV, our approach is more applicable to a later phase in the design of the circuit implementation of a latency-insensitive protocol. In particular, we formally verify the RTL implementation of relay station and shell in a *modular* fashion so that a previously verified synchronous system does not need to be re-verified after it has been transformed into a latency-insensitive system. This approach has several advantages. New systems can be verified independently of the architecture they will operate on. In addition, formally verifying the shell is quite demanding in terms of computational memory: to verify an entire system implementation with numerous cores, each encapsulated in its own shell would be prohibitively expensive at the same level of rigor.

### III. A SIMPLIFIED LATENCY-INSENSITIVE PROTOCOL AND ITS IMPLEMENTATIONS

In this section we discuss in detail the implementation of the simplified latency-insensitive protocol LID-1ss that we introduced in Section I. Briefly, the new protocol differs from the original LID-2ss protocol discussed in [4] in the back-pressure mechanism: the LID-1ss protocol uses a single *stop* bit to stall a sender. For both the shell and the relay station, we first present sample simulations of their I/O behaviors and then explain the details of the RTL designs.

**Shell.** Fig. 4 shows a sample simulation trace of a two-input-two-output shell and its core with the assumption that both input queues have a capacity of two. A block diagram of the shell and its stallable core module is illustrated in Fig. 5(a). The core implements a function  $f: (C_{t+1}, D_{t+1}) = f(A_t, B_t)$ , where  $A_t$  and  $B_t$  are data tokens arriving on input channel  $In_1$  and  $In_2$  while  $C_t$  and  $D_t$  are the tokens produced by the core on output channel  $Out_1$  and  $Out_2$  at time  $t$ , respectively.

Several scenarios are illustrated in this trace. In cycle 1 both channels  $In_1$  and  $In_2$  present valid data tokens, and, therefore, the core can be fired to produce valid output tokens ( $C_2$  and  $D_2$ ) at cycle 2. At cycle 2 the void input token of channel  $In_1$  (void bit is high) causes the shell to stall the core at cycle 3. Therefore, both the output tokens at cycle 3 are marked as void with their *voidOut* bits being asserted by the shell.

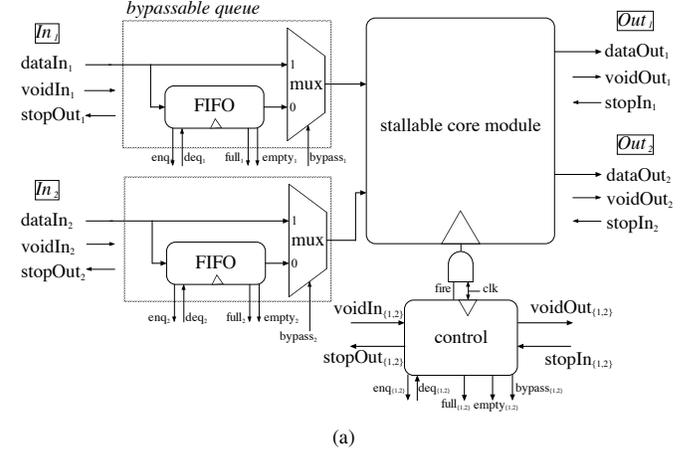
The scenario in which the shell receives back-pressure happens at cycle 5, when the downlink receiver of channel  $Out_2$  asserts the *stopIn<sub>2</sub>* bit. Thus the output token  $D_4$  is regarded as void at cycle 5. The core is stalled at cycle 6, and both  $C_4$  and  $D_4$  are repeated at cycle 6. However, since the downlink receiver of channel  $Out_1$  has already sampled  $C_4$ , the void bit is set for the repeated  $C_4$  so the same token will not be sampled twice on channel  $Out_1$ . The accompanying *void* bit of  $D_4$ , on the other hand, is not set because token  $D_4$  on channel  $Out_2$  has not been sampled yet. In this case  $D_4$  is sampled at the end of cycle 6 (when the clock edges arrives to start cycle 7).

What follows from cycle 6 shows the case when an input queue is full. The stop request from the downlink of channel  $Out_2$  causes the input queue of channel  $In_2$  to be filled up at cycle 6 (two valid tokens are stored in channel  $In_2$ 's queue at the end of cycle 5, due to the stalls at cycle 3 and 6), thus a stop request is raised to the uplink sender of channel  $In_2$ . Note that at cycle 6 the shell is not able to store token  $B_6$ . The same token is thus resent on channel  $In_2$  and is sampled by the shell at cycle 7.

Next we present the details of the shell RTL logic design. Fig. 5(a) reports a block diagram of a two-input-two-output shell, and the logic functions of the controller is listed in Fig. 5(b). The control logic is general and can be easily scaled to handle an arbitrary number of input and output channels. All the logic functions are quite simple and can be implemented with few logic gates. The clock gating signal *fire* decides whether the core module is fired or stalled. It is asserted when each channel presents a valid token either directly from the channel input or from its input queue, and no stop request has arrived on any output channel. The second condition can be

		1	2	3	4	5	6	7	8	9	10	11
$In_1$	<i>dataIn<sub>1</sub></i>	$A_1$	$A_1$	$A_2$	$A_3$	$A_4$	$A_5$	$A_6$	$A_6$	$A_8$	$A_8$	$A_9$
	<i>voidIn<sub>1</sub></i>	0	1	0	0	0	0	0	1	1	0	0
	<i>stopOut<sub>1</sub></i>	0	0	0	0	0	0	0	0	0	0	0
$In_2$	<i>dataIn<sub>2</sub></i>	$B_1$	$B_2$	$B_3$	$B_4$	$B_5$	$B_6$	$B_6$	$B_6$	$B_6$	$B_8$	$B_9$
	<i>voidIn<sub>2</sub></i>	0	0	0	0	0	0	0	1	1	0	0
	<i>stopOut<sub>2</sub></i>	0	0	0	0	0	1	0	0	0	0	0
$Out_1$	<i>dataOut<sub>1</sub></i>	$C_1$	$C_2$	$C_2$	$C_3$	$C_4$	$C_4$	$C_5$	$C_6$	$C_7$	$C_7$	$C_8$
	<i>voidOut<sub>1</sub></i>	0	0	1	0	0	1	0	0	0	0	1
	<i>stopIn<sub>1</sub></i>	0	0	0	0	0	0	0	0	0	0	1
$Out_2$	<i>dataOut<sub>2</sub></i>	$D_1$	$D_2$	$D_2$	$D_3$	$D_4$	$D_4$	$D_5$	$D_6$	$D_7$	$D_7$	$D_8$
	<i>voidOut<sub>2</sub></i>	0	0	1	0	0	0	0	0	0	1	0
	<i>stopIn<sub>2</sub></i>	0	0	0	0	1	0	0	0	0	0	0

Fig. 4. Sample I/O behavior of the new shell. Shaded data tokens are bubbles.



$$\begin{aligned}
 \text{fire} &= \bigwedge_{i \in \mathcal{I}} (\text{voidIn}_i + \overline{\text{empty}}_i) \cdot \bigvee_{j \in \mathcal{O}} (\text{stopIn}_j \cdot \overline{\text{voidOut}}_j) \\
 \forall j \in \mathcal{O} \text{ voidOut}_j &= \begin{cases} 0 & \text{if } \text{stopIn}_j \cdot \overline{\text{voidOut}}_j \text{ is true} \\ \text{fire} & \text{otherwise} \end{cases} \\
 \forall i \in \mathcal{I} \text{ stopOut}_i &= \text{full}_i \\
 \forall i \in \mathcal{I} \text{ enq}_i &= \text{voidIn}_i \cdot (\text{fire} + \overline{\text{empty}}_i) \cdot \overline{\text{full}}_i \\
 \forall i \in \mathcal{I} \text{ deq}_i &= \overline{\text{empty}}_i \cdot \text{fire} \\
 \forall i \in \mathcal{I} \text{ bypass}_i &= \overline{\text{empty}}_i
 \end{aligned}
 \tag{b}$$

Fig. 5. (a) A block diagram of a two-input-two-output shell and a stallable core module. (b) Logic functions of the shell controller.

detected by checking the current *stopIn* and *voidOut* bits for each output channel. If the *voidOut<sub>j</sub>* bit is high for some output channel  $j$ , the downlink receiver of channel  $j$  has received the latest valid token. In this case the core module can proceed even if the receiver requests to stop.

The *voidOut<sub>j</sub>* bit informs to the downlink module on output channel  $j$  whether the current token is a valid token or not. It is a sequential signal buffered by an edge-triggered flip-flop. The condition  $\text{stopIn}_j \cdot \overline{\text{voidOut}}_j = \text{true}$  means that the downlink module on channel  $j$  is not able to process the current (also the latest) valid data token. In this case the core module will be stalled, the current token will be repeated, and *voidOut<sub>j</sub>* will be set low. In all other cases the value of the *voidOut<sub>j</sub>* bit depends on whether the core module will be fired.

The major data-path components in a shell are the bypassable queues that store unused valid tokens from input channels. Its minimum forward latency is zero. The bypassable queue is implemented as a standard FIFO whose output is multiplexed with the incoming data of the channel. If the queue is empty, the controller selects the data token from the input channel and passes it to the core module. The

	1	2	3	4	5	6	7	8	9	10	11
<i>dataIn</i>	A <sub>1</sub>	A <sub>1</sub>	A <sub>2</sub>	A <sub>2</sub>	A <sub>3</sub>	A <sub>4</sub>	A <sub>4</sub>	A <sub>5</sub>	A <sub>6</sub>	A <sub>7</sub>	A <sub>7</sub>
<i>voidIn</i>	0	1	0	1	0	0	1	0	0	0	0
<i>stopOut</i>	0	0	0	0	0	0	0	0	0	1	0
<i>dataOut</i>	*	A <sub>1</sub>	A <sub>1</sub>	A <sub>2</sub>	A <sub>2</sub>	A <sub>3</sub>	A <sub>4</sub>	A <sub>4</sub>	A <sub>5</sub>	A <sub>5</sub>	A <sub>6</sub>
<i>voidOut</i>	0	0	1	0	1	0	0	0	0	0	0
<i>stopIn</i>	0	0	0	0	1	0	1	0	1	0	0

Fig. 6. Sample I/O behavior of the new relay station.

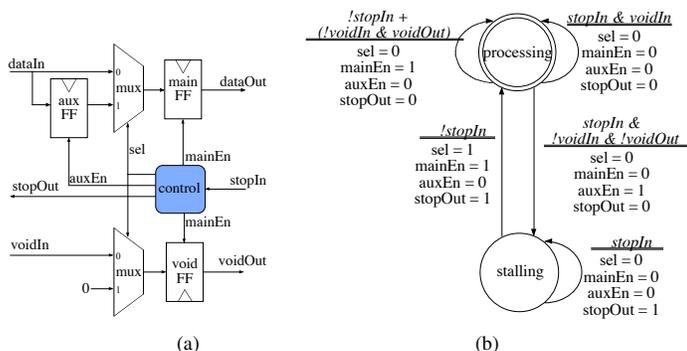


Fig. 7. (a) Block diagram of the new relay station; (b) The state transition diagram of its controller.

internal queue is a *sequential* element: all of the operations (i.e. enqueue and dequeue) and the update of its status (i.e. full or empty) take place at each clock edge. Hence all of the *stopOut* signals, which are the full signals from the queue, are sequential signals.

**Relay station.** Fig. 6 reports sample I/O behaviors of a relay station. From cycle 1 to 4, the relay station simply relays the received data, void or not, from its input channel to its output channel. At cycle 9, the relay station receives a stop request from its downlink receiver. It then stalls (and repeats its output token) for one cycle to avoid overflow its downlink receiver. Meanwhile, the incoming data token at cycle 9 is buffered in the relay station’s internal storage, and the stop request is sent to its uplink sender at next clock cycle.

Sometimes, an optimization can be applied to avoid stalling the relay station when the downlink receiver asserts the *stopIn* bit. This is shown at cycle 5 to 6. At cycle 5 the relay station receives the stop request and emits a void token at the same time. Because the void token will not be sampled by its downlink receiver, the relay station can safely continue to relay data tokens at cycle 6 without being stalled.

Another optimization occurs when the relay station absorbs a stop request instead of relaying it to its uplink sender. For instance, at cycle 7 the relay station receives a void token from its uplink and a stop request from its downlink. It can actually discard the void token received at cycle 7, instead of buffering it, and simply repeat its current output at cycle 8. In this way, it avoids propagating the stop request.

Fig. 7(a) shows an implementation of the relay station for the proposed latency-insensitive protocol; Fig. 7(b) reports the state transition diagram of its controller. The new relay station uses two edge-triggered flip-flops to store incoming data tokens, and one flip-flop to buffer the *voidOut* bit. The two flip-flops storing data tokens provide the necessary twofold storage capacity. The output of the main flip-flop is the

data output of the relay station. The controller decides when to update the three flip-flops and sets *stopOut* and *voidOut* bits according to the protocol. The control logic is discussed next.

The controller is a two-state Mealy finite state machine with three input and four output signals. The initial state is the *processing* state, which enables the main flip-flop and sets the *stopOut* bit low. In the stalling state, instead, the relay station uses both the main and the auxiliary flip-flops to store data tokens, and requests the uplink sender to stop sending more data tokens by asserting its *stopOut* bit. Note that the value of the *stopOut* bit depends only on the current state of the controller, and thus no combinational path exists between *stopIn* and *stopOut*.

The switching from the *processing* state to the *stalling* state is triggered by the condition that the *stopIn* bit is high, and both the *voidIn* and *voidOut* bits are low. The asserted *stopIn* bit indicates that the receiver is not able to process the output data taken of the relay station. Hence the relay station has to maintain its output token by keeping the same data in the main flip-flop. On the other hand, the relay station must save the incoming valid token (indicated by low values of *voidIn* and *stopOut*) in the auxiliary flip-flop, and enter the *stalling* state. Note that the incoming *voidIn* bit is not saved in the void flip-flop, because in this case it is always low (this is part of the condition to switch from the *processing* to the *stalling* state) and thus can be easily recovered.

The relay station goes back from the *stalling* to the *processing* state when its downlink receiver deasserts the *stopIn* bit, indicating that it is ready to receive more valid data tokens. Then, the relay station moves the token saved in the auxiliary flip-flop to the main flip-flop. It also updates the void flip-flop with a constant low value because the accompanying void bit of the data token in the auxiliary flip-flop must be deasserted.

#### IV. FORMAL VERIFICATION OF THE LID PROTOCOL IMPLEMENTATIONS

An important compositional result is proven as part of the theory of latency-insensitive design [5]: if all modules in a strict system are replaced by corresponding latency-equivalent patient modules, then the resulting system is patient and latency equivalent to the original one. Naturally, this theoretical result is not enough to guarantee that a particular implementation of a latency-insensitive system is correct. The theory tells us that we can build a patient system out of patient parts, but we must also verify that the parts (the actual implementations of the shells and relay stations) are patient. On the other hand, we can verify the implementations of shells and relay stations in isolation because according to the compositionality rule for latency equivalence of patient processes, a system composed of shell-core pairs and relay stations is also latency equivalent to the original strict system.

We first translated by hand the synthesizable VERILOG code implementing the logic of the shell and relay station described in Section III into the NuSMV language [18]. Then we used the NuSMV model checker to verify that they are correct refinements of the specifications given in the LID theory.

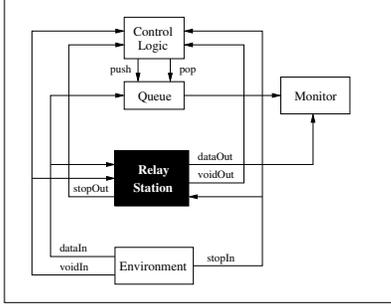


Fig. 8. Verification framework for a relay station.

In particular we verified the design for properties related to latency equivalence, liveness, and storage capacity. For a relay station this is sufficient to prove that it is a patient process. The shell is a little trickier. For the shell, patience also depends on the functionality of the core that the shell encapsulates and the shell implementation varies slightly depending on the number of input and output channels of its core.

**Verification approach.** Fig. 8 and Fig. 9 illustrate our verification approach for the relay station and the shell respectively. The verification framework consists of the component-under-verification (CUV) together with the environment, queue, and monitor modules. The environment generates data items, the valid bits, and the stop bits in an unconstrained manner: at each clock cycle, the environment may non-deterministically choose a value for *dataIn*, and non-deterministically set *voidIn* and *stopIn* to either *true* or *false* values. This enables verification under all possible input sequences; if any possible input sequence fails, a counterexample is generated. The monitor checks the correctness of the property to be verified by comparing the stream(s) of valid data produced by the CUV versus the stream(s) of data that passed through the queue. The correct functioning of a latency-insensitive component is checked under the assumption that its environment obeys the latency-insensitive protocol i.e. the environment holds a data token until it is sampled by the component. We do not impose this assumption on the environment and instead track the sampling of data tokens according to the latency-insensitive protocol.

The queue is a FIFO used to store the valid data tokens sampled by the monitor until they are matched with the output tokens. It has standard push and pop operations for adding new valid tokens to the tail of the queue and popping valid tokens off the head of the queue. A valid data token is pushed in the queue whenever the CUV latches in the token. Similarly a valid data token is popped off the queue whenever the CUV outputs a data token. These decisions are made by the queue control logic based on the values of the *stop* and *void* bits. The queue's pop signal is forwarded to the monitor, and when a pop occurs the monitor compares the queue's output to the CUV's output.

For the verification of the relay station a simple FIFO is sufficient because the relay station itself has simple store-and-forward behavior. For the verification of the shell, we also need a core module to perform computation on the given

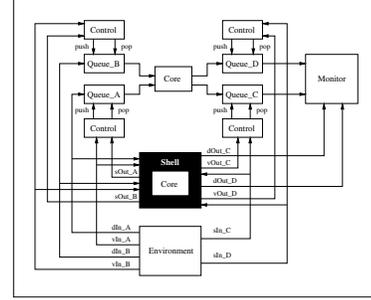


Fig. 9. Verification framework for a shell.

inputs and produce output data. We chose a 2-input, 2-output core that computes in parallel the two-input NAND and NOR logic operations and stores the results in two internal flip-flops. Separate queues are maintained for each incoming channel, and a second core module is instantiated outside the shell. When both input queues have valid data tokens, these are passed to the core and the results are stored in an output queue. The monitor compares the output of the shell with the data in the output queue.

**Formal Properties.** We checked the properties of latency equivalence, liveness, and storage capacity. The latency equivalence property expresses that there is no loss, duplication or reordering of valid tokens in a data stream. To test latency equivalence of the relay station, we checked that the relay station's outgoing data stream is latency equivalent to its incoming data stream. To verify latency equivalence of the two-input two-output shell, we compared the data tokens produced by the core alone and those produced by the core/shell pair.

The liveness property expresses progress in the system. A component is live if it produces meaningful data provided the environment allows it. We imposed a fairness constraint on the environment for the *void* and *stop* bits so that the environment generates valid data items infinitely often and enables the downlink stream infinitely often. The liveness property states that the component generates valid data tokens infinitely often and enables the uplink stream infinitely often.

The storage capacity property checks that the number of data items in the monitor queue never exceeds the storage capacity of the component. The relay station capacity is equal to two. The storage capacity of the shell depends on the size of its internal queue, which is at least equal to one.

The above properties were verified individually for the shell and relay station Verilog implementations. All of the properties passed verification. The latency equivalence property was also tested on known erroneous implementations of both the shell and relay station. The verification failed and generated counterexamples as expected. The verification was performed on a machine with 2 AMD Opteron™ processors and 3.5 GB memory over Redhat Linux with the Fedora Core 6, and NuSMV version 2.4.1. Time and memory usage from the verification experiments are summarized in Table I.

Property	Module name	Time	Memory
Latency	Relay station	0.2 sec	7.2 MB
Equivalence	Shell	15.5 min	2.4 GB
Liveness	Relay station	5.5 sec	14.3 MB
	Shell	1.4 hours	2.4 GB

TABLE I  
MEMORY AND TIME STATISTICS FOR THE VERIFICATION TASKS.

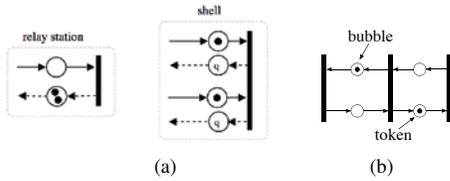


Fig. 10. Marked graph models of (a) LID-1ss and (b) SEA interface circuits.

## V. COMPARISONS OF LID INTERFACE CIRCUITS AND SYNCHRONOUS ELASTIC ARCHITECTURES

In this section we present a comparative analysis of the new class of interface circuits implementing the proposed latency-insensitive protocol LID-1ss versus the interface circuit implementation of the original LID-2ss protocol and the interface circuits for synchronous elastic architectures (SEAs) proposed in [10]. We completed the semicustom design of the three classes of circuits with a  $90nm$  industrial standard-cell library in order to compare them in terms of system throughput, logic delay, as well as area overhead.

In Section II we provided a brief overview of SEAs and clarified that they do not use the concept of shell interfaces but rely instead on elastic fork and join structures. In the sequel, however, whenever it is convenient we will use the term “shell” also to refer to the SEA interface logic for a processing core and, in particular, to the composition of the control logic of the substitute elastic buffer with the fork and join control structures.

**System Throughput.** To make a system robust with respect to communication latency through the application of either LID or elasticization may have a negative impact on its performance measured as processing throughput. This is defined as the ratio of the number of valid tokens over the number of valid tokens plus void tokens that the system processes over time. Since both a relay station (RS) and an elastic buffer (EB) are initialized with a void token and since void tokens may create more void tokens whenever they stall a computation, the placement of RSs or EBs on channels that belong to feedback loops and/or re-convergent paths may induce permanent degradation of the system throughput. The system throughput can be computed exactly by using either *marked graph* models [9], [10], [19], or equivalently max-plus algebra [14]. Fig. 10 shows the marked graph models for the interface circuits of LID-1ss and SEA [10]. Note that in the shell model the sizes of the shell queues are represented by a variable  $q$  whose value may be set statically (at design time) to optimize performance [9]. These models are compositional as they inherit their topological structure from the modeled system. Fig. 11 reports the LID-1ss model and the SEA model for the system shown in Fig. 12(a). Note that in the LID-1ss

model each transition takes a single time unit to fire. Instead in the SEA model a transition takes half a time unit to fire because it is a latch-based design.

The maximum sustainable processing throughput of a LID or SEA system is equal to the reciprocal of the *cycle time* of its corresponding marked graph model: the cycle time is equal to the largest cycle metric across all its cycles; the cycle metric is equal to the sum of each transition’s firing time divided by the number of tokens along the cycle. (an invariant number in a marked graph) [20].<sup>2</sup> For both models in Fig. 11 we highlighted the critical cycles, i.e. cycles having the highest cycle metric. The LID-1ss-based implementation has a throughput of  $3/4 = 0.75$ , assuming all input queues in a shell have a capacity of one [9], [14]. The throughput of the SEA version, on the other hand, is lower:  $2/3 = 0.67$ . In this particular example, the ideal system throughput, equal to 1, can still be achieved for both implementations. For the LID-1ss version it is necessary either to insert an additional relay station between cores  $B$  and  $C$  (or  $A$  and  $B$ ) or to raise to two the size of the input queue in the  $C$  shell for the channel  $B \rightarrow C$ . The second approach is called *optimal channel queue sizing* [9], [14]. Since the SEA join structures do not use queues, the only solution to improve the throughput is to insert an additional elastic buffer between cores  $B$  and  $C$  (or  $A$  and  $B$ ).

For certain systems, however, an SEA-based implementation cannot achieve the same system throughput of an implementation based on either LID-1ss or LID-2ss. This is due to the particular structure of these systems that may present particular combinations of reconvergent paths and/or feedback loops. For example, for the system shown Fig. 12(b) an implementation based on LID-1ss or LID-2ss can achieve higher system throughput than a SEA-based implementation. Note that the system has a similar reconvergent path from  $A$  to  $C$  as the example in Fig. 12(a), but it has two additional *cycles*:  $(A, B, E, A)$  and  $(B, C, D, B)$ . In a LID-1ss implementation, to achieve the ideal throughput equal to 1 it is necessary to increase the input queue size of channel  $B \rightarrow C$  in  $C$ ’s shell to 2. In this case, however, it is impossible for a corresponding SEA to achieve such an ideal throughput because, at best, one can insert an additional elastic buffer between  $B$  and  $C$  (or  $A$  and  $B$ ), which brings the throughput up to  $3/4$  (the cycle with the inserted EB becomes the new critical cycle).

The two examples in Fig. 12(b) show the impact on system throughput that input queues at a join point have. Insufficient queue size at a join point, like in the LID-1ss shell with queues of size one or in the SEA join structures that lack queues, degrades the system throughput.<sup>3</sup> The reason is the following: whenever an input queue is full at a join point, the uplink sender, informed by the stop signal (back-pressure), must re-send the same data token until the queue has room to accept it. The more such re-sending happens, as in a SEA join structure, the more throughput degradation may occur.

<sup>2</sup>This can be computed by solving the *maximum cycle mean problem* for which a number of efficient algorithms have been proposed [21], [22].

<sup>3</sup>It should be possible to derive an implementation of interface circuits for the SELF protocol that instead of being based on SEA join structures uses input queues like in a LID shell block.

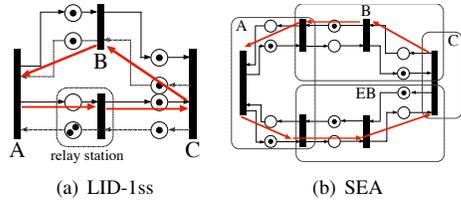


Fig. 11. Marked graph models of the example in Fig. 12(a).

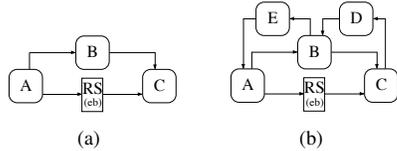


Fig. 12. Examples of systems with unbalanced reconvergent paths.

**Interface Logic Delay.** The delay of LID and SEA interface logic affects the overall system performance in two ways. First, the longest combinational logic path within an interface or across two communicating interfaces might become the new critical path of the system, and thus determine the maximum clock frequency at which the system can run. Second, when pipelining a wire using repeaters, either relay stations (RS) or elastic buffers (EB), the smaller the *cross-interface* logic delay between two communicating interfaces is, the further the two interfaces can be stretched away without inserting repeaters in-between. Thus the deployment of interfaces with smaller cross-interface logic delay can result in less number of RSs/EBs used for wire pipelining. Because each inserted RS/EB introduces an additional void token into the system and may potentially reduce system throughput, it is desirable to design interfaces with minimal cross-interface logic delay.

In order to analyze the logic delays of the various interface circuits we synthesized their RTL Verilog implementations<sup>4</sup> with a 90nm industrial standard cell library using SYNOPSIS DESIGN COMPILER. As shown in Fig. 13, the interface logic is assumed to drive optimally buffered wires [1], [23]. The critical logic delays within each individual interface and across the logic of communicating interfaces are then extracted using DESIGN COMPILER static timing analyzer.

For the LID-2ss and LID-1ss designs, which are based on edge-triggered flip-flops (FFs), the slack is derived by subtracting the maximum logic delay between two flip-flops and the flip-flop setup time from the clock period. For the SEA design, which is based on level-sensitive latches, the slack is calculated by subtracting the maximum logic delay between two active-high (or active-low) latches and latch setup time from the clock period.<sup>5</sup> When calculating cross-interface slacks, as shown in Fig. 13 (LID-2ss and LID-1ss) and Fig. 14

<sup>4</sup>We derived the LID-2ss and LID-1ss implementations, and obtained a gate-level circuit implementation from the authors of SEA. We slightly changed the latter to avoid excess area overheads, as discussed in Section II.

<sup>5</sup>Although a latch-based design allows *time borrowing*, the total delays over a path spanning a chain of active-high and -low latches must stay within a fixed number of clock periods determined by the number of high-low latch pairs. To simplify the analysis without sacrificing accuracy, we assumed that the path between two active-high (or -low) latches must be within one clock period.

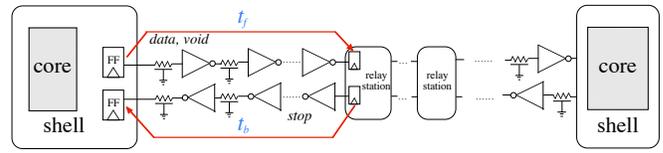


Fig. 13. Long wires are optimally buffered by repeaters.

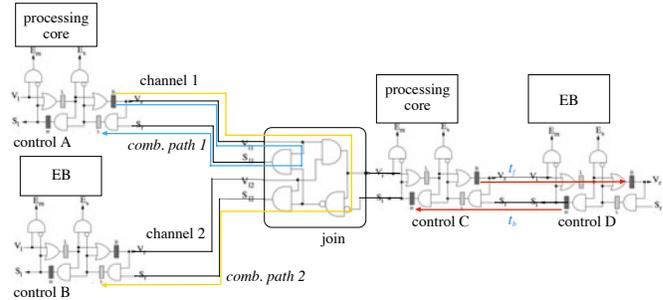


Fig. 14. Combinational paths due to the join structure (left) and SEA slack computation (right).

(SEA), the delays of forward paths (data and *void/valid*)  $t_f$  and of backward paths (*stop*)  $t_b$  are both considered (without counting delays of buffered wires across the channel).

Fig. 15(a) and Fig. 15(b)-15(e) summarize the results of our analysis of the impacts of logic delay on system performance in terms of the minimum slacks and the maximum physical lengths of interconnects as allowed by the three sets of interface logic respectively. Fig. 15(a) reports the minimum slacks left in each interface logic and the four possible combinations of communicating interface logic when running at 500 MHz clock rate, while ignoring the delays of buffered interconnects. The channel width is assumed to be 64-bit wide, and each core has two input channels. The more slack an interface logic has, the faster clock rate can be applied. LID-1ss has more slack in all but one scenarios, and thus enjoys faster clock rates than LID-2ss and SEA. Conversely, the slack of the shell-shell pair in SEA is significantly low. This may either limit the system clock frequency, or require the insertion of an additional elastic buffer between the two shells to increase available slack. But inserting an elastic buffer introduces a void token and, therefore, it may lower the system throughput.

Fig. 15(b)-15(e) report maximum allowable wire lengths between four different pairs of communicating interface circuits at various clock frequencies. LID-1ss allows the maximum interconnect lengths in all four possible scenarios. The “X” marks indicate that at the given clock frequency the timing constraint is not met in the corresponding pair of communicating interfaces, so additional RS/EB must be inserted between them or the pair must be physically close to avoid long interconnect wires. The former solution might decrease system throughput; the latter might constrain physical design tools.

The maximum physical lengths of interconnects allowed between the RS-shell or shell-shell pairs in SEA are shorter than what the corresponding slacks imply. This is because the join structure used in the two-input “shell” in SEA creates multiple combinational paths running across a single channel

twice or spanning across two channels, as indicated in Fig. 14. Therefore the slack available between the two-input shell and its uplink counterparts are shared among the interface logic and the corresponding forward path and backward path between them. As a result, the join structure allows a much shorter physical length for the interconnects, and physical design tools must be used to carefully “balance” the lengths of the “joined” wires to avoid timing violations. These combinational paths are introduced by the interface logic with multiple input channels (here the two-input shell), regardless of whether the senders are elastic buffers or other processing cores.

Notice that the combinational paths created by the SEA join structure are unavoidable. In fact, the lack of input queues at the receiver’s end forces the buffering of unused valid data tokens at the immediate sender’s end. Hence a multi-input core receiving an invalid token must request the re-transmission of all the valid tokens received at the *same* clock cycle as they arrive. Consequently, combinational paths between the communicating interface logic are required.

The above analysis of logic delay shows that the proposed LID-1ss interface logic can support higher system clock rate and throughput than LID-2ss and SEA counterparts. The reason is that the interface logic of LID-1ss has more slack, and requires a smaller number of wire pipelining elements (relay stations) because it allows longer interconnect between its interface logic. Latch-based SEA design does provide additional flexibility to the physical design tools because time borrowing allows an elastic buffer to tolerate varying wire delays and thus to be placed in a wider range of area.

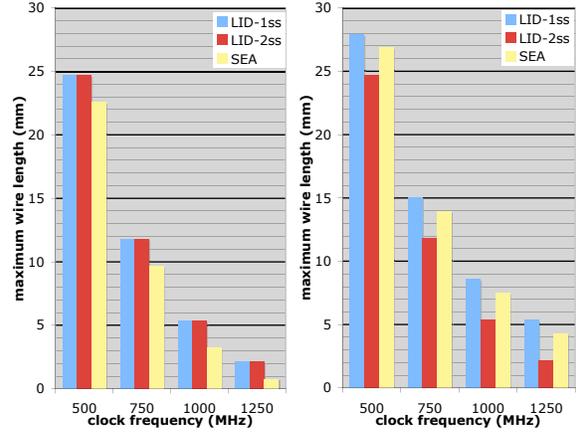
**Area Overhead Comparisons.** Shell interfaces, relay stations and elastic buffers do occupy active silicon area and therefore represent a necessary area overhead of any latency-insensitive design approach. We analyzed and compared area overhead figures for the three approaches discussed in this paper after performing logic synthesis and technology mapping.

Fig. 16(a) reports the area overhead of the shell designs in LID-2ss and LID-1ss (for both queue of size one and two) over a range of different channel widths; Fig. 16(b) shows the corresponding overhead incurred in elasticization of processing cores with different number of flip-flops. The area overhead of the LID-1ss shell with queue of size two is roughly the same as the one of the LID-2ss shell while the LID-1ss shell with queue of size one is smaller. In fact, the area of a shell is dominated by the area of its queues, which depends on the widths of the input channels. For a SEA, the area overhead of elasticizing a processing core grows with the number of flip-flops contained in the core. This is because the substitute latches require a little more area than the replaced edge-triggered flip-flops.

Fig. 16(c) compares the area overhead of the three LID shells and their SEA counterpart when they are used to encapsulate different instances of a  $32 \times 32$  pipelined multiplier synthesized from the SYNOPSIS DESIGNWARE IP core library. For a number of pipeline stages varying from 2 to 6 the bar diagram reports the absolute area of the synthesized multipliers as well as the area of the corresponding shells. The overhead ratios between each shell’s area and the multiplier’s area is labeled on top of each corresponding bar. As expected,

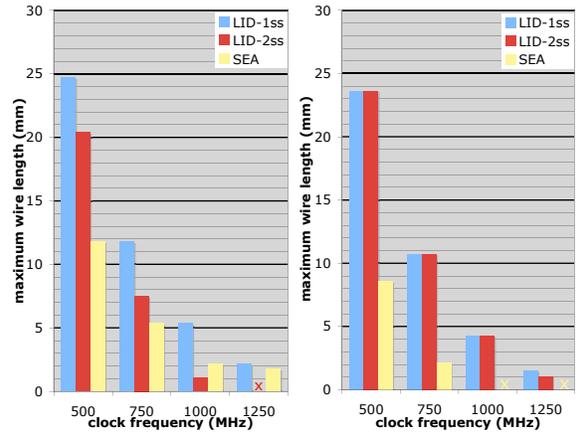
	shell	RS	shell-RS	RS-RS	RS-shell	shell-shell
LID-1ss	1.23	1.28	1.32	1.5	1.33	1.24
LID-2ss	1.14	1.23	1.32	1.32	1.1	1.27
SEA	1.24	1.00	1.21	1.44	1.31	0.92

(a) Slacks (in nanoseconds) of interface logic at 500 MHz clock rate.



(b) 2-in-2-out shell  $\rightarrow$  RS

(c) relay station  $\rightarrow$  relay station



(d) RS  $\rightarrow$  2-in-2-out shell

(e) 2-in-2-out shell  $\rightarrow$  2-in-2-out shell

Fig. 15. Minimum slacks and maximum physical lengths of interconnects allowed by interface logic. The input queue size of LID-1ss shell is two.

the absolute area of the shells in LID-2ss and LID-1ss are constant regardless the number of pipeline stages, but the area overhead ratio of the LID-1ss shell’s area drop from 16% to 13% (in the case of input queue size  $q = 1$ ) as the multiplier’s logic grows (the same trend applies to LID-2ss). In contrast, the area of the SEA “shell” grows slightly with the number of pipeline stages, and its area overhead ratio grows from 5% to 10%. In this example the area overhead of LID-1ss and LID-2ss is significant but this is greatly reduced for IP cores that are more complex than a pipelined multiplier.

Fig. 16(d) reports the area of relay stations and elastic buffers over a range of different channel widths. The area overhead of the latch-based SEA elastic buffers is 2/3 of their LID-2ss and LID-1ss counterparts thanks to the clever use of two latches to provide the needed twofold capacity. Due to the more complex steering logic between its flip-flops LID-1ss relay stations are slightly larger than the LID-2ss ones.

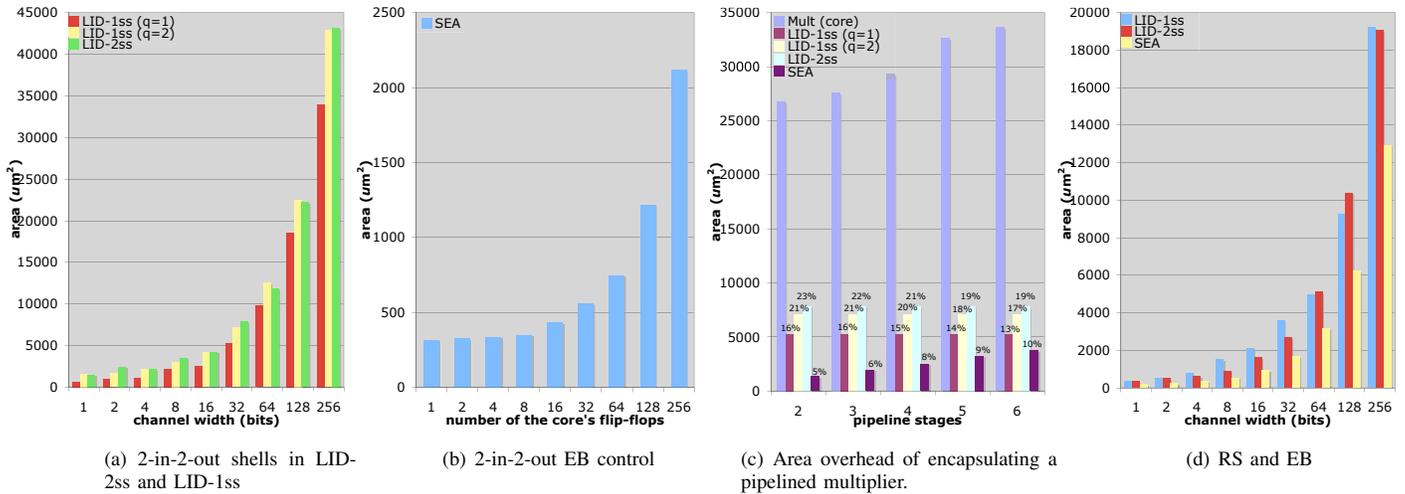


Fig. 16. Area of synthesized interface circuits.

## VI. CONCLUDING REMARKS

We proposed a new class of interface circuits to support latency-insensitive design based on LID-1ss, a simpler latency-insensitive protocol. We presented a detailed experimental analysis comparing the LID-1ss interface circuits to those supporting the original protocol discussed in [4], [8], that we called LID-2ss, as well as to the interface circuits for synchronous elastic architectures that were proposed in [10]. We showed that LID-1ss offers clear improvements in terms of area overhead and logic delay with respect to LID-2ss. With respect to the interface circuits for synchronous elastic architectures the LID-1ss interface circuits have smaller logic delay and, for many systems, enable higher processing throughput.

## VII. ACKNOWLEDGEMENTS

The authors would like to thank Jordi Cortadella for providing the SEA interface circuits and Michael Theobald and Franjo Ivančić for helpful discussions. This research is partially based upon work supported by the NSF under Grant No. 0541278, an NDSEG fellowship, and the GSRC.

## REFERENCES

- [1] R. Ho, K. W. Mai, and M. A. Horowitz, "The future of wires," *IEEE Proc.*, vol. 89, no. 4, pp. 490–504, Apr. 2001.
- [2] D. Matzke, "Will physical scalability sabotage performance gains?" *IEEE Computer*, vol. 30, pp. 37–39, Sep. 1997.
- [3] L. P. Carloni and A. L. Sangiovanni-Vincentelli, "Coping with latency in SOC design," *IEEE Micro*, vol. 22, no. 5, pp. 24–35, Sep-Oct 2002.
- [4] L. P. Carloni, K. L. McMillan, A. Saldanha, and A. L. Sangiovanni-Vincentelli, "A methodology for "correct-by-construction" latency insensitive design," in *Proc. of the Intl. Conf. on Computer-Aided Design*. San Jose, CA: IEEE, Nov. 1999, pp. 309–315.
- [5] L. P. Carloni, K. L. McMillan, and A. L. Sangiovanni-Vincentelli, "Theory of latency-insensitive design," *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, vol. 20, no. 9, pp. 1059–1076, Sep. 2001.
- [6] A. Benveniste, P. Caspi, S. Edwards, N. Halbwachs, P. L. Guernic, and R. de Simone, "The synchronous language twelve years later," *Proc. of the IEEE*, vol. 91, no. 1, pp. 64–83, Jan. 2003.
- [7] E. A. Lee and A. Sangiovanni-Vincentelli, "A Framework for Comparing Models of Computation," *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, vol. 17, no. 12, pp. 1217–1229, Dec. 1998.
- [8] L. P. Carloni, "The role of back-pressure in implementing latency-insensitive systems," *Electr. Notes Theor. Comput. Sci.*, vol. 146, no. 2, pp. 61–80, 2006.
- [9] R. Collins and L. Carloni, "Topology-based optimization of maximal sustainable throughput in a latency-insensitive system," in *To appear in the Proc. of Design Automation Conf. (DAC)*, Jun. 2007.
- [10] J. Cortadella, M. Kishinevsky, and B. Grundmann, "Synthesis of synchronous elastic architectures," in *Proc. of the Design Automation Conf.*, 2006, pp. 657–662.
- [11] A. Agiwal and M. Singh, "An architecture and a wrapper synthesis approach for multi-clock latency-insensitive systems," in *Proc. of the Intl. Conf. on Computer-Aided Design*, 2005, pp. 1006–1013.
- [12] M. R. Casu and L. Macchiarulo, "A new approach to latency insensitive design," in *Proc. of the Design Automation Conf.*, 2004, pp. 576–581.
- [13] T. Chelcea and S. M. Nowick, "Robust interfaces for mixed-timing systems," *IEEE Trans. on Very Large Scale Integrated Systems.*, vol. 12, no. 8, pp. 857–873, 2004.
- [14] R. Lu and C.-K. Koh, "Performance analysis of latency-insensitive systems," *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, vol. 25, no. 3, pp. 469–483, Mar. 2006.
- [15] M. Singh and M. Theobald, "Generalized latency-insensitive systems for single-clock and multi-clock architectures," in *Proc. of the Conf. on Design, Automation and Test in Europe*, 2004, pp. 1008–1013.
- [16] H. Jacobson, P. Kudva, P. Bose, P. Cook, S. Schuster, E. Mercer, and C. Myers, "Synchronous interlocked pipelines," in *Proc. of the Intl. Symp. on Asynchronous Circuits and Systems*, Apr. 2002, pp. 3–12.
- [17] S. Suhaib, D. Mathaikutty, D. Berner, and S. Shukla, "Validating families of latency insensitive protocols," *IEEE Trans. on Computers*, vol. 55, no. 11, pp. 1391–1401, 2006.
- [18] A. Cimatti, E. Clarke, F. Giunchiglia, and M. Roveri, "NUSMV: a new Symbolic Model Verifier," in *Proc. of the Intl. Conf. on Computer-Aided Verification*, July 1999, pp. 495–499.
- [19] L. P. Carloni and A. L. Sangiovanni-Vincentelli, "Performance analysis and optimization of latency insensitive systems," in *Proc. of the Design Automation Conf.*, Jun. 2000, pp. 361–367.
- [20] C. V. Ramamoorthy and G. S. Ho, "Performance evaluation of asynchronous concurrent systems using Petri nets," *IEEE Tran. on Software Engineering*, vol. 6, no. 5, pp. 440–449, Sep. 1980.
- [21] R. M. Karp, "A characterization of the minimum cycle mean in a digraph," *Discrete Mathematics*, vol. 23, pp. 309–311, 1978.
- [22] A. Dasdan and R. Gupta, "Faster maximum and minimum mean cycle algorithms for system-performance analysis," *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, vol. 17, pp. 889–899, Oct. 1998.
- [23] J. M. Rabaey, A. Chandrakasan, and B. Nikolić, *Digital integrated circuits: a design perspective*. Prentice-Hall, Inc. Upper Saddle River, NJ, USA, 2002.