

Refinement of a Perfectly Synchronous Communication Model onto *Nostrum* NoC Best-Effort Communication Service

Zhonghai Lu, Ingo Sander, and Axel Jantsch
Laboratory of Electronics and Computer Systems
Department of Microelectronics and Information Technology
Royal Institute of Technology, Sweden
Isafjordsgatan 39
SE-164 40 Stockholm, Sweden
{zhonghai,ingo,axel}@imit.kth.se

Abstract

We present a novel approach to refine a system model specified with perfectly synchronous communication onto a Network-on-Chip (NoC) best-effort communication service. We propose a top-down procedure with three steps, namely, *channel refinement*, *process refinement*, and *communication mapping*. In channel refinement, synchronous channels are replaced with stochastic channels abstracting the best-effort service. In process refinement, processes are refined in terms of interfaces and synchronization properties. Particularly, we use *synchronizers* to achieve *synchronization consistency*. Within communication mapping, the refined processes and channels are mapped to a NoC architecture. Adopting the *Nostrum* NoC platform as target architecture, we use a digital equalizer as a tutorial example to illustrate the feasibility of our concepts.

1 Introduction

For system design, a synchronous design style is attractive since it allows one to separate timing from function. The designer can focus on the design of the system functionality without being distracted by unnecessary low-level communication details. This also facilitates the verification task, which is a key activity at the system level. Later, the implementation details and design constraints can be gradually filled in by refinement.

Network-on-Chip (NoC) is emerging as a new SoC paradigm to cope with the scalability problem of buses in order to connect tens or perhaps even hundreds of microprocessor-sized heterogeneous resources, such as processor cores, DSPs, FPGAs/ASICs, and memories, enabled on a single chip due to the steady technology scaling. *Nostrum* [MNTJ04, NMÓJ03, TMJ03] is our NoC architecture that provides a packet-switched communication platform. To satisfy different performance/cost requirements, *Nostrum* provides two classes of unicast communication services, namely, Best Effort (BE) and Guaranteed Bandwidth (GB) services. The BE service is connection-less where packets are routed without resource reservation. The GB service is connection-oriented where packets are delivered after enough bandwidth is reserved.

There is a huge gap between an abstract system model and a complex implementation platform like NoC. In order to bridge the gap, we propose a NoC design flow shown in Fig. 1 where we concentrate on the communication problem. A system specified as a synchronous process model that has to be mapped on a NoC. There are three communication-related tasks: *clustering & resource allocation*, *communication refinement*, and *synthesis*. The clustering flattens the hierarchy in the model and groups processes into new processes with perhaps coarser granularity. With resource allocation, the new processes are allocated to HW or SW execution resources. Communication refinement bridges the gap between the

communication model in the specification and the NoC communication implementation via adapters. With synthesis, these processes and adapters are synthesized into HW and/or SW.

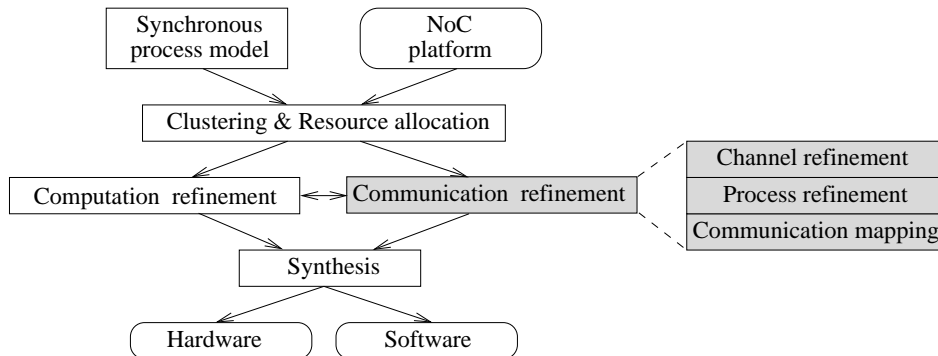


Figure 1: A NoC Design Flow

In this paper, we address the communication refinement that starts from a synchronous communication model and ends with the Nostrum NoC best-effort communication service. Our contributions are (1) a novel approach to realize this communication refinement; (2) a classification of process synchronization properties as *strict*, *nonstrict*, *strong*, and *weak* synchronization in order to formally analyze processes’ local synchronization requirement(s) (Section 5.2); (3) using *synchronizers* (synchronization adapters) to maintain synchronization consistency during refinement (Section 5.3). In a synchronous system model, communication is perfectly synchronous with a global logical clock and cleanly separated from computation. With the NoC communication service, communication introduces variable delays and crosses multiple clock domains connected by a packet-switched network. Clearly the communication in the implementation domain is not synchronous, thus not consistent with that in the specification domain. We will focus on this synchronization issue while keeping the process computation untouched. Note that, this synchronization issue is a process communication property at the system modeling level (signal level), not at the lower implementation level such as shared memory synchronization using locks or semaphores, as well as message passing synchronization using blocking or nonblocking semantics. We assume a clustering is done in a way that the resulting processes remain locally in a synchronous domain. Besides, we consider that a resource maintains a local synchronous region. Consequently a process (after clustering) is to be mapped to one resource and one resource hosts exactly one process.

In the sequel, we outline related work in Section 2. Section 3 gives an overview of our refinement technique and introduces the digital equalizer. Section 4 and 5 present the channel refinement and process refinement in detail, respectively. In Section 6, we describe the communication mapping. Finally we draw conclusions and point out future directions in Section 7.

2 Related Work

Based on the isolation of communication from computation, a large body of work on communication refinement exists in the literature. Through the Virtual Component Interfaces (VCI) of the VSI Alliance [LSdJ⁺00], the COSY-VCC design flow [BKK⁺00] supports communication refinement from specification, to performance estimation and to implementation. IPSIM [CCGM03] developed on top of SystemC 3.0 supports an object-oriented methodology and establishes two inter-module communication layers. The message box layer concerns generic and system-specific communication, while the driver layer implements higher level application-dependent communications. The SpecC methodology defines four levels of abstraction, namely at the specification, architecture, communication and implementation level, and the refinement transformations between them [DGG02]. Jerraya et al. achieved communication refinement via a generic wrapper concept [YNL⁺01]. In the course of communication refinement, methods to allow architecture exploration and communication protocol selection can be found in [LSvdWD01]

and [KM99], respectively. These works do not assume a synchronous specification, thus are not applicable to our context.

With synchronous communication, latency insensitive theory [CMSV01] targets synchronized HW design where synchronization can still be achieved even if interconnecting synchronous IP blocks experiences indefinite wire latencies; Desynchronization for SW design was addressed in [BCG00]. Furthermore, some mathematical frameworks were developed to support refinement-based design methods. Benveniste et al. present a theoretical framework for modeling heterogeneous systems, and derive sufficient conditions to maintain semantic-preserving transformations when deploying a synchronous specification onto GALs and the loosely time-triggered architectures [BCCSV03]. Another framework is proposed in [GTL03] concerning the refinement of a polysynchronous specification, which allows multiple clocks instead of a single clock. All these works are complementary to our work but none of them provides a detailed refinement approach targeting a NoC platform.

3 Refinement Overview

In this section, we first introduce the functional specification with perfect synchrony and the digital equalizer. Then we describe the Nostrum communication services. Finally we outline the refinement procedure.

3.1 Functional Specification with Perfect Synchrony

The synchronous modeling paradigm is based on an elegant and simple mathematical model, which has been shown successful and is the ground of synchronous languages such as Esterel, Signal, Argos and Lustre. The basis is the perfect synchrony hypothesis, i.e., both computation and communication take no observable time. A system is modeled as a set of concurrent communicating processes via signals. Processes use ideal data types and assume infinite buffers. Signals are ordered sequences of events. Each event has a time slot as a slot to convey data. If the data contains useful information, the event is *present* and called a *token*; otherwise, the event is *absent* and modeled as a \perp representing a clock tick. Each signal can be related to the time slots of another signal in an unambiguous way. The output events of a process occur in the same time slot as the corresponding input events. Moreover, they are instantaneously distributed in the entire system and are available to all other processes in the same slot. Receiving processes in turn consume the events and emit output events again in the same time slot. A signal can thus be viewed as an ideal communication channel which has no delay for any event data types (unlimited bandwidth)¹. A process specified in the synchronous paradigm is a synchronous process. For feedback loops, the perfect synchrony creates cyclic dependency between output and input, and thus leads to deadlock, which is resolved with an initial event in the specification. A synchronous model is deterministic, i.e., given the same input streams, it generates the same output streams.




Figure 2: The Digital Equalizer

As a tutorial example, Fig. 2 shows the functional model of an equalizer. It adjusts the bass and treble volume of the audio stream according to button control levels. In addition it prevents the bass level

¹For convenience, we use the term *signal* to express either a sequence of events or the ideal communication medium.

from exceeding a predefined threshold to avoid damaging the speakers. Its function can be described by the following set of equations, where the initial value '1' is used to resolve the feedback loops. This model is specified in functional language Haskell and executable.

$$\begin{aligned}
\text{AudioOut} &= \text{Equalizer}(\text{Buttons}, \text{AudioIn}) \\
&\text{where} \\
\text{AudioOut} &= \text{Sum}(\text{AudioBass}, \text{AudioTreble}) \\
(\text{Bass}, \text{Treble}) &= \text{LevelControl}(\text{Buttons}, \text{AudioOut}) \\
\text{AudioBass} &= \text{BassFilter}(\text{AudioIn}, \text{init} : \text{Bass}) \\
\text{AudioTreble} &= \text{TrebleFilter}(\text{AudioIn}, \text{init} : \text{Treble}) \\
\text{init} &= 1
\end{aligned}$$

3.2 Nostrum Communication Services

In Nostrum, each resource R_i ($i = 1, 2, \dots, n$) is equipped with a Resource-Network-Interface (RNI) in order to access the network, as shown in the lower part of Fig. 3. The RNI and the network belong to the Nostrum protocol stack. Nostrum provides a message passing platform with two unicast communication services, i.e., best-effort and guaranteed bandwidth. The BE service [NMÖJ03] is implemented by routing packets. It has no guarantee on timely delivery, but has an upper bound on delivery time. To this end, we assume a network admission protocol that prevents the network from saturation and guarantees bounds on delay. It is connectionless and does not reserve network resources such as storage and link bandwidth, thus has a lower cost. The GB service is implemented by using looped containers and temporally disjoint networks [MNTJ04]. It guarantees bandwidth, which is negotiated during the connection establishment phase. It is connection-oriented and reserves the network resources before transmission and thus has a higher cost. The RNIs hide the service implementation details and make the services *transparently* accessible to applications. The access methods as a standard interface are communication primitives.

Within Nostrum, we define a set of communication primitives for message passing as follows:

- *int open(int src, int dst, int service, struct bandwidth)*: it opens a simplex channel between a source *src* process and a destination *dst* process. The *service* denotes the channel service class, 0 for the BE service, 1 for the GB service. The *bandwidth* is a user-defined record with three fields $\{int\ min_bw, avg_bw, int\ max_bw\}$ which specifies the minimum, average and maximum bandwidth (*Bytes/second*) requirement of the channel. The method returns a unique channel identity number (*cid*) upon successfully opening the channel; otherwise, it returns various reasons of failure, such as a destination invalid, or performance not satisfied.
- *bool write(int cid, void msg)*: it writes *msg* to the specified channel *cid*. The size of messages is bounded. It returns the status of the write.
- *bool read(int cid, void *msg)*: it reads channel *cid* and writes the received data to the address starting at *msg*. It returns the status of the read.

We have implemented these primitives with the BE service using SystemC in our layered NoC simulator *Semla* [TMJ03]. Currently the write() and read() are implemented with nonblocking semantics. *Semla* is programmable as to network topology, process-to-resource allocation, routing algorithm, resource/network clock frequency, and traffic pattern. The current implementation opens channels statically during compile time and the opened channels are never closed.

3.3 The Refinement Procedure

Given a synchronous system specification, our objective is to refine the synchronous communication model onto the Nostrum best-effort (BE) service. To this end, we propose a three-step procedure: *channel*

refinement, process refinement, and communication mapping. We illustrate the procedure via a pair of producer-consumer processes in Fig. 3. The three steps are marked by a circle with a step number inside it.

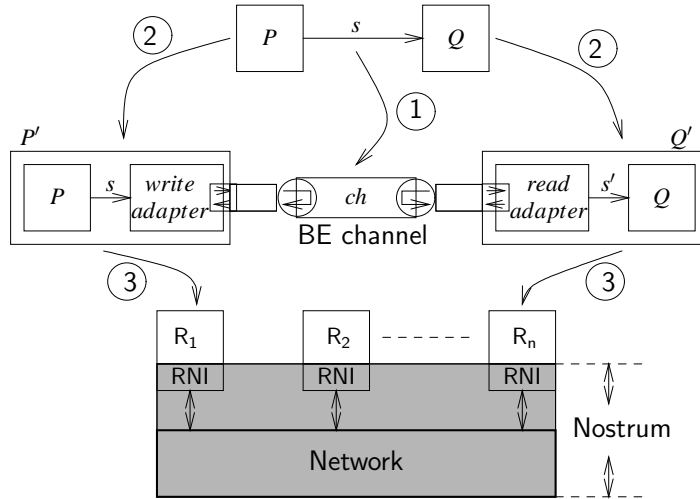


Figure 3: Communication Refinement Overview

Step 1: With **channel refinement**, we first abstract the behavior of Nostrum best-effort service as that of stochastic service channels which are then used to replace the ideal communication channels, i.e., the signals. In Fig. 3, the signal s between the producer P and the consumer Q is refined to a BE service channel ch . As a consequence, signal s becomes signal s' , which is a derived version of s , after being delivered via the service channel. Furthermore, s and s' are no longer synchronous.

Step 2: With **process refinement**, we deal with how a process can be connected to the service interfaces as well as how its synchronization property can be satisfied by *adapters*. Particularly, to guarantee the correctness of the refinement, the process synchronization property from the specification model to the refined model must be consistent. Maintaining this synchronization property is the focus of this paper. Moreover, we consider feedback loops where the process synchronization may be relaxed since a synchronous specification may over specify the system. In Fig. 3, P and Q are adapted with a write and read adapter, respectively. Note that the adapters contain components to interface with the service channels and components for synchronization whenever necessary.

Step 3: Finally, together with a process-to-resource allocation scheme, the **communication mapping** is to implement both the adapters and service channels on a NoC, in this case, the Nostrum simulator Semla. In Fig. 3, the refined processes P' and Q' are mapped to the resources R_1 and R_n , respectively. Accordingly, the service channel ch is implemented via the interfaces provided by the RNIs of the resources R_1 and R_n .

4 Channel Refinement

We first abstract the behavior of the Nostrum best-effort service resorting to a stochastic approach, then analyze the impact of stochastic channels on the system model, particularly, the disruption of the perfect synchrony assumption.

4.1 The Behavior Model of Nostrum Best-Effort Service

The performance of the Nostrum BE service is *nondeterministic* in nature since the message delivery experiences dynamic contention scenarios in the RNIs and network. Nevertheless, the message delivery

time is not completely indeterminate. Given the characteristics of a packet-switched network such as topology, routing algorithm and flow control scheme, the behavior of message delivery is a function of the network traffic (both total traffic amount and traffic patterns). For the GB service, the bandwidth is guaranteed but the delay may be jittery.

To capture the performance characteristics of the best-effort service, we resort to a stochastic approach. Formally, we develop a unicast BE service channel as a point-to-point *stochastic* channel: given an input signal of messages $\{m_1, m_2, \dots, m_n\}$ to a service channel, the output signal is $\{d_1, m_1, d_2, m_2, \dots, d_n, m_n\}$, where d_i denotes the delay of message m_i ($i = 1, 2, \dots, n$) which may be expressed in terms of the number of absent (\perp) values; d_i is subject to a distribution with a minimum $d_{i,min}$ and maximum $d_{i,max}$ value determined by the service implementation, network traffic and the distance of the two ends of the service channel. If $d_i = n$ (n is a natural number), it means that there are n absent values between message m_{i-1} and message m_i . We identify two important properties of the behavior of the service channel: (1) d_i is varying; (2) d_i is bounded. This behavior is purely viewed from the perspective of application processes and its implementation details are hidden. In addition, the stochastic channel model is generic.

4.2 Impact of the Stochastic Channels

Replacing the ideal channel (zero delay and unlimited bandwidth) with a stochastic channel (varying delay and limited bandwidth) leads to the violation of the synchrony assumption. In the specification, a channel is ideal so that we can use a *single* signal s to connect a producer to a consumer process. After replacing it with a service channel, the signal can be seen as being *split* into a pair of signals, the original signal s and its derived signal s' , as shown in Fig. 3. For a process with two synchronous input signals, for example, the *Sum* process of the equalizer (Fig. 2), if both signals s_3 and s_4 are delivered via a service channel, they are split, resulting in two derived signals s'_3 and s'_4 , which are now the input signals to the *Sum* process. Apparently, the two pairs of signals, s_3 and s'_3 , s_4 and s'_4 , and the two derived signals s'_3 and s'_4 are not synchronous. A synchronous system becomes globally asynchronous, leading to possibly nondeterministic behavior which deviates from the specification. It is therefore important to maintain *synchronization consistency* during the refinement for correctness.

5 Process Refinement

We first briefly consider how to interface with the service channels in general, and then discuss the synchronization property of processes and methods to achieve synchronization consistency. The granularity of a process in this context is a synchronous domain resulting from the clustering. At the system level (a composition of processes), we discuss feedback loops.

5.1 Interfacing with the Service Channels

Once an ideal channel is replaced by a service channel, the processes can not be directly connected to the interface of the service channel. They must be *adapted* in terms of data and control because (1) the input/output data type of a service channel is a bounded message while a signal in the specification assumes an ideal data type, whose length is finite but arbitrary, e.g., a 32/64-bit integer, a 64-bit floating point or a user-defined 256-bit record type etc.; (2) the service channel has bounded buffers and limited bandwidth while a signal uses unlimited resources. The sending and receiving of messages use shared resources and thus control functionality has to be added to maintain the message delivery properties such as reliability and causality etc. The control function typically enables to allocate shared resources, schedule multiple threads and achieve thread-level synchronization. These adaptations are achieved by a writer and reader process. Specifically, to interface with the service channels, a producer needs to be wrapped with a *writer*, a consumer with a *reader*.

5.2 Process Synchronization Property

In the system model, all signals of each process are synchronous. In spite of this, whether or not the input signals of a process must be synchronous, i.e., the synchronization property of a process, is subject to the evaluation condition of processes, specifically, the local condition(s) to *evaluate* the input events. Because of the tight synchronization in the model, some processes may be over specified, limiting the implementation alternatives. During the refinement, the designer(s) must closely inspect and determine the synchronization property of the processes.

Inspired by [LP95], we use *firing rules* to discuss the synchronization property of *synchronous processes*. For a synchronous process with n input signals, PI is a set of N input patterns, $PI = \{I_1, I_2, \dots, I_N\}$. The input patterns of a synchronous process describe its firing rules, which give the conditions of evaluating input events at each event cycle. I_i ($i \in [1, N]$) constitutes a set of event patterns, one for each of n input signals, $I_i = \{I_{i,1}, I_{i,2}, \dots, I_{i,n}\}$. A pattern $I_{i,j}$ contains only one element that can be either a token wildcard $*$ or an absent value \perp , where $*$ does not include \perp . Based on the definition of firing rules, we propose four levels of process synchronization properties as follows:

- *Strict synchronization.* All the input events of a process must be present before the process evaluates and consumes them. The only rule that the process can fire is $PI = \{I_1\}$ where $I_1 = \{[*], [*], \dots, [*]\}$.
- *Nonstrict synchronization.* Not all the input events of a process are absent before the process fires. The process can *not* fire with the pattern $I = \{[\perp], [\perp], \dots, [\perp]\}$.
- *Strong synchronization.* All the input events of a process must be either present or absent in order to fire the process. The process has only two firing rules $PI = \{I_1, I_2\}$, where $I_1 = \{[*], [*], \dots, [*]\}$ and $I_2 = \{[\perp], [\perp], \dots, [\perp]\}$.
- *Weak synchronization.* The process can fire with any possible input patterns. For a 2-input process, its firing rules are $PI = \{I_1, I_2, I_3, I_4\}$ where $I_1 = \{[*], [*]\}$, $I_2 = \{[\perp], [\perp]\}$, $I_3 = \{[*], [\perp]\}$ and $I_4 = \{[\perp], [*]\}$.

We can identify processes with a *strict*, *strong*, and *weak* synchronization property in the equalizer (Fig. 2). The *BassFilter* (s_0 and s_1) and *TrebleFilter* (s_0 and s_2) have a strict synchronization. Both filters are composed of a FIR filter and an amplifier. The FIR filter is specified as an FSM, whose state transition is sensitive to time, thus a \perp value in an audio stream can change the values of its output sequence. Meanwhile, the amplifier must have an amplification level, thus a \perp value makes the amplifier undefined. The *Sum* process (s_3 and s_4) has a strong synchronization. It is a combinational process and thus tolerable to events with a \perp value. However, the two events of s_3 and s_4 must be synchronized before being processed since they represent the low and high frequency parts of the same audio sample. The *LevelControl* (s_b and s_5) process has a weak synchronization. It can fire even when either or both of the events of s_b and s_5 are absent since pressing buttons happens irregularly and the bass level surpassing the threshold occurs only aperiodically.

5.3 Achieving Synchronization Consistency

Apparently, for processes with a strict or strong synchronization, their synchronization properties can not be satisfied if any of the input signals passes through a service channel since the delays via the channel are stochastic. Although globally asynchronous, the processes can be locally synchronized by using adapters to satisfy their synchronization properties. To achieve strong synchronization, we use a synchronizer process *sync*; to achieve strict synchronization, we use three processes, *sync*, *deSync* and *addSync*. We use a two-input process to illustrate these processes in Fig. 4. A synchronizer process *sync* aligns the tokens of its input events, as shown in Fig. 4a. It does not change the time structure of the input signals. A desynchronizer *deSync* removes the absent values, as shown in Fig. 4b. All its input

signals must have the same token pattern, resembling the output signals of the *sync* process. Removing absent values implies that the process is *stalled*. The desynchronizer changes the timing structure of the input signals, which must be recovered in order to prevent from incurring unexpected behavior of other processes that use the timing information. An add-synchronizer *addSync* adds the absent values to recover the timing structure, as shown in Fig. 4c. It must be used in relation to a *deSync* process. If the input events of the *deSync* is a token, the *addSync* reads one event from its internal buffers for each output signal; otherwise, it outputs a \perp event. As can be seen, the two processes *deSync* and *addSync* are used as a pair to assist processes to fulfill strictness.

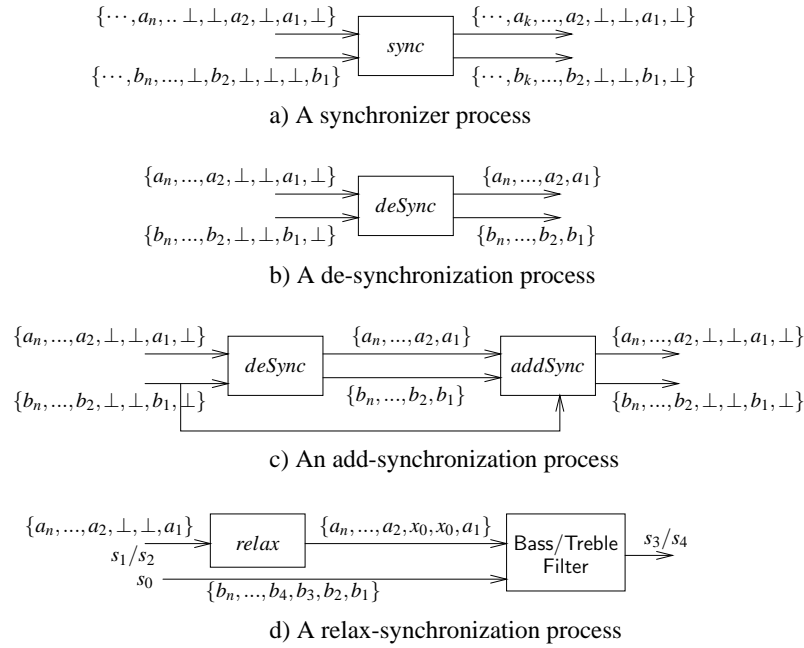


Figure 4: Processes for Synchronization




Figure 5: Read/Write Adapters for A Process with Strong Synchronization




Figure 6: Read/Write Adapters for A Process with Strict Synchronization

We can now use these synchronizers in connection with the *reader* and *writer* processes to wrap the original processes to interface with the service channels and maintain the synchronization consistency from the specification model to the refined model. For instance, as shown in Figure 5, we use a *sync* process together with a pair of *reader/writer* processes to wrap the *sum* process in the equalizer model to maintain its strong synchronization. We use the three processes, *sync*, *deSync* and *addSync*, together with a pair of *reader/writer* processes to wrap the *Bass/Treble Filter* process (Fig. 2) to maintain their strict synchronization.

The refinement of processes with a nonstrict synchronization should be individually investigated according to their firing rules.

5.4 Feedback Loops

In the specification, feedback loops are resolved by using initial events. If the feedback signals pass through a service channel, the delays are nondeterministic. If following the initial event approach in the refinement procedure, we encounter a problem since we are not certain how many initial events are required to resolve the deadlock. Consider the *Bass/Treble Filter*, if the tokens of s_1/s_2 are not available, it can not fire. This implies it may not be able to process enough audio samples in time, leading to violate the system's performance constraint. However, if the amplification level signals, s_1 (*Bass*) and s_2 (*Treble*), are delayed and thus not available, the amplifiers should continue functioning by, for example, using the previous amplification level or simply using a constant level like 1. In this case, the effect of pressing buttons may be delayed several cycles. This is tolerable since the human sensing of the changes in the audio volume takes some time.

By this observation, we can in fact *relax* the strict synchronization of the processes *Bass/Treble Filter*, using a relax-synchronization process *relax* illustrated in Fig. 4d. If the input event is a token, it outputs the token; otherwise, a token x_0 is emitted. The exact value of x_0 is application dependent. Relaxing synchronization is a design decision leading to behavior discrepancy between the specification and the refined model. It must be used carefully to ensure that it does not cause to violate the system requirements.

6 Communication mapping

The inputs to this task are the refined model as well as a process-to-resource allocation scheme; the output is a communication implementation on Semla.

6.1 Channel Mapping

With a resource allocation scheme, all processes are allocated to resources in a one-to-one manner. Note that this is not a limitation but due to the assumption on the clustering and resources (refer to Section 1). With such a clustering, inter-process signals, which represent inter-resource communications, are mapped to service channels. Since the processes may be hierarchical, we need to flatten the hierarchy to the level that each signal mapped to a service channel can be uniquely identified with a pair of a producer and a consumer process with *finer* granularity. For simplicity, we do not consider mapping multiple service channels to one implementation channel. Mapping channels is thus straightforward. Each pair of processes communicating via a service channel in the refined model results in its dedicated unicast implementation channel, which is mapped to the open channel primitive *open()*. For example, with the producer-consumer case, a BE channel setup is fulfilled by a single line of code: `int ch[1] = open(P,Q,BE_SERVICE,NULL)`.

6.2 Communication Process Mapping

After the process refinement, the refined processes consist of the original computational process, the writer and reader, and perhaps the synchronizer(s) to satisfy their synchronization properties. Our refinement keeps the original processes intact. Therefore, the tasks of communication process mapping are to implement the adapters for writing (*writer*), reading (*reader*), and the synchronizers such as *sync*, *deSync*, *addSync* and *relax*, and to coordinate the writing and reading operations if needed.

In SystemC, processes are implemented as modules. The readers/writers may be implemented as separate modules or in the same modules as processes. We implement a process and its adapter(s) in

a single module. For implementation, execution control in the module must be considered. Suppose the module has a single thread of control, we need to find a Periodic Admissible Sequential Sequence (PASS) for process executions [LSJ02]. For the process in Fig. 6, a PASS could be $PASS = \{reader, sync, desync, compute, addsync, writer\}$. Besides, a control signal $write_rdy$ must be asserted by the *writer* to the *reader* to enable the reading from the channel(s) for the next-round execution of the PASS, as shown in Fig. 6. This leads to a local feedback loop, and we adopt the initial event approach to deal with. In this case, $write_rdy$ is initially asserted. Using the communication primitives defined in Section 3.2, the SystemC module for Fig. 6 is sketched as follows, with each component explained in commentary:

```

process_class::Process(){
//initially write_rdy=1;
//read_ch0_rdy=0; read_ch1_rdy=0
//sync_rdy=0; compute_done=0;
if (write_rdy==1){
//(1) reader: nonblocking read ch1 and ch2
if (read_ch0_rdy==0)
if ((read(ch[0],&r_msg1))==true)
read_ch0_rdy=1;
if (read_ch1_rdy==0)
if ((read(ch[1],&r_msg2))==true)
read_ch1_rdy=1;
//(2) sync: synchronize the two events
if (read_ch0_rdy==1 && read_ch1_rdy==1)
sync_rdy=1;
else sync_rdy=0;
//(3) deSync: desynchronization by guard
if (sync_rdy==1 && compute_done==0){
//process computation
//return w_msg and set compute_done to 1
w_msg=compute(r_msg1 ,r_msg2);
write_rdy=0; compute_done=1;}
}
//(4) addSync: fill synchronization
if (sync_rdy==1 && compute_done==1) {
//(5) writer: nonblocking write ch3
if (write_rdy==0)
if (write(ch[3],w_msg)==true){
write_rdy=1;
sync_rdy=0; compute_done=0;
read_ch0_rdy=0; read_ch1_rdy=0;}
}
}
}

```

In the implementation domain, whether to emit and pass \perp as a message via a service channel or not can be a design decision that must be handled carefully. To preserve the semantics, \perp must be emitted and passed. However, it incurs too much overhead on computation and communication, and may not be useful since its value is useless. Therefore it is usually neglected. Only in cases where the timing information carried by \perp is used by other processes, it must be emitted and passed as a special value. In the equalizer case, \perp is neglected since its timing information is not used by any of the four processes, therefore it does not affect the system behavior.

We have implemented the equalizer in Semla. The purpose is to validate the concepts of our refinement approach. Fig. 7a illustrates the mapped equalizer in a 4x4 mesh NoC. All the five inter-resource signals s_1, s_2, \dots, s_5 (Fig. 2) use the BE service. To simplify the discussion on performance, the resources and the network use the same clock frequency. The network switches operate in a synchronous manner with the switching per hop taking one cycle. The message streams on s_3 and s_4 are injected into the network conservatively so that a new audio sample will not be processed by the filters until the previous sample has been handled by the *Sum* process. This implies that the audio samples are not processed in

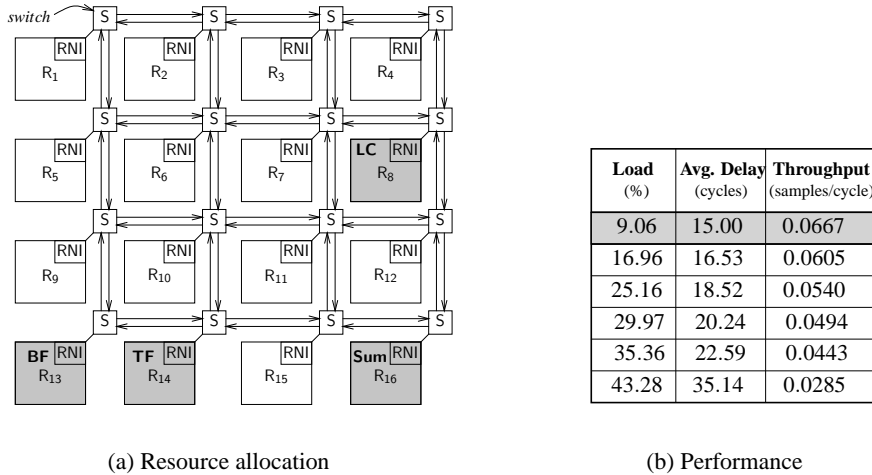


Figure 7: The Equalizer Mapped on A NoC

a pipeline fashion in the network. In addition, we inject background traffic with uniformly distributed random destinations in the network. The motivation is to load the network with reasonable amount of traffic since the equalizer example can only make use of a small fraction of the network capacity. Fig. 7b shows the equalizer performance, where the network load is the average percentage of active links per cycle. The process computations are function calls and complete instantly. We observe the average delay that is the time (in cycles) to process one sample. Since the audio processing is not pipelined, the throughput (samples/cycle) is simply the inverse of the average delay. In Fig. 7b, the first row shows the case where there is no background traffic. As expected, when the network is increasingly loaded, the average delay is increased and the throughput decreased. The average delay can be seen as the time to respond to a button press or to activate bass control. We noted that the audio output sequences are different from those observed from the specification due to relaxing the synchronization for the feedback loops. We conducted other experiments in which we removed the feedback loops, and could validate that the output sequences agree with each other in all traffic setting cases.

7 Conclusions and Future Work

Communication refinement is a crucial step in a NoC design flow. We have presented a refinement approach that enables us to map a perfectly synchronous communication model onto the NoC best-effort service accessible through communication primitives. Particularly we classify the synchronization properties of processes and describe methods to achieve synchronization consistency during the refinement upon the violation of the perfect synchrony hypothesis. For feedback loops, we relax the synchronization with the tolerance of system requirements. In this paper we use Nostrum as our target, but with few adjustments, the approach should be applicable for other NoC platforms as well.

In future work, we plan to realize automatically analyzing the synchronization properties of processes, and then during refinement, we take either automatic analysis that yields correct synchronization and system behavior, or manual analysis with manual design decisions on the synchronization refinement combined with a systematic verification of the resulting implementation. For the refinement of feedback loops, we intend to use Nostrum GB service to achieve a systematic solution. Moreover, we will consider optimization of the communication refinement for performance enhancement.

ACKNOWLEDGEMENTS

The work reported in this paper was supported by the Swedish government within the SOCWARE program.

References

- [BCCSV03] A. Benveniste, L. Carloni, P. Caspi, and A. Sangiovanni-Vincentelli. Heterogeneous reactive systems modeling and correct-by-construction deployment. In *Proceedings of the Third International Conference on Embedded Software*, 2003.
- [BCG00] A. Benveniste, B. Caillaud, and P. Le Guernic. Compositionality in dataflow synchronous languages: specification and distributed code generation. *Information and Computation*, 163:125–171, 2000.
- [BKK⁺00] J.-Y. Brunel, W.M. Kruijtzter, H.J.H.N. Kenter, F. Petrot, L. Pasquier, E.A. de Kock, and W.J.M. Smits. COSY communication IP's. In *Proceedings of the 37th Design Automation Conference*, Los Angeles, California, June 2000.
- [CCGM03] Marcello Coppola, Stephane Curaba, Miltos Grammatikakis, and Giuseppe Maruccia. IPSIM: SystemC 3.0 enhancements for communication refinement. In *Proceedings of Design Automation and Test in Europe*, 2003.
- [CMSV01] Luca P. Carloni, Kenneth L. McMillan, and Alberto L. Sangiovanni-Vincentelli. Theory of latency-insensitive design. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 20(9):18, September 2001.
- [DGG02] Rainer Dömer, Daniel D. Gajski, and Andreas Gerstlauer. SpecC methodology for high-level modeling. In *Proceedings of the Ninth IEEE/DATC Electronic Design Processes Workshop*, April 2002.
- [GTL03] Paul Le Guernic, Jean-Pierre Talpin, and Jean-Christophe Le Lann. Polychrony for system design. *Journal of Circuits, Systems and Computers*, 12(3):261–303, December 2003.
- [KM99] P.V. Knudsen and J. Madsen. Integrating communication protocol selection with hardware/software codesign. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 18(8):1077 – 1095, 1999.
- [LP95] Edward A. Lee and Thomas M. Parks. Dataflow process networks. *Proceedings of the IEEE*, 1995.
- [LSdJ⁺00] C.K. Lennard, P. Schaumont, G. de Jong, A. Haverinen, and P. Hardee. Standards for system-level design: practical reality or solution in search of a question? In *Proceedings of Design Automation and Test in Europe*, March 2000.
- [LSJ02] Zhonghai Lu, Ingo Sander, and Axel Jantsch. A case study of hardware and software synthesis in ForSyDe. In *Proceedings of the 15th International Symposium on System Synthesis*, October 2002.
- [LSvdWD01] Paul Lieverse, Todor Stefanov, Pieter van der Wolf, and Ed Depretter. System level design with SPADE: an M-JPEG case study. In *Proceedings of the IEEE/ACM International Conference on Computer-Aided Design*, 2001.
- [MNTJ04] Mikael Millberg, Erland Nilsson, Rikard Thid, and Axel Jantsch. Guaranteed bandwidth using looped containers in temporally disjoint networks within the Nostrum network on chip. In *Proceedings of the Design Automation and Test Europe Conference (DATE)*, 2004.
- [NMÖJ03] Erland Nilsson, Mikael Millberg, Johnny Öberg, and Axel Jantsch. Load distribution with the proximity congestion awareness in a network on chip. In *Proceedings of the Design Automation and Test Europe (DATE)*, pages 1126–1127, 2003.
- [TMJ03] Rikard Thid, Mikael Millberg, and Axel Jantsch. Evaluating NoC communication backbones with simulation. In *Proceedings of the IEEE NorChip Conference*, 2003.
- [YNL⁺01] Sungjoo Yoo, Gabriela Nicolescu, Damien Lyonnard, Amer Baghdadi, and Ahmed A. Jerraya. A generic wrapper architecture for multi-processor SoC cosimulation and design. In *Proceedings of the Ninth International Symposium on Hardware/Software Codesign*, Copenhagen, Denmark, 2001.