

# ACCEPTOR: a model and a protocol for real-time multi-mode applications on reconfigurable heterogeneous platforms

Joël Goossens  
joel.goossens@ulb.ac.be  
Université libre de Bruxelles  
Brussels, Belgium

Antonio Paolillo  
antonio.paolillo@ulb.ac.be  
Université libre de Bruxelles  
Brussels, Belgium

Xavier Poczekajlo  
xavier.poczekajlo@ulb.ac.be  
Université libre de Bruxelles  
Brussels, Belgium

Paul Rodriguez  
paul.rodriguez.lobera@pm.me  
Université libre de Bruxelles  
Brussels, Belgium

## ABSTRACT

In this work, we consider hard real-time applications scheduled upon heterogeneous multiprocessor platforms. The originality of this study is to consider multi-mode real-time applications (software aspects) *and reconfigurable*-heterogeneous hardware platforms (composed of CPUs, GPUs, FPGAs...). Our approach is based on a multi-mode protocol, for mode-dependent tasks upon reconfigurable hardware. The goal is to handle predictable switches between different task sets *and* different hardware settings. The novelty here is the *dynamic* hardware *and* software reconfigurability. First, we propose a formal model of the applications and reconfigurable hardware platforms. We then propose and prove correct a mode change protocol. We propose in particular a validity test for the verification of the timing constraints of the application –including the time allowed to complete a mode change. We also perform a complete evaluation. We study the theoretical complexity of the protocol, use a simulation to evaluate the efficiency of our solution, and finally propose a competitive analysis of our protocol to prove that it is 2-competitive.

## CCS CONCEPTS

• **Computer systems organization** → **Real-time systems**; *Multicore architectures*.

## KEYWORDS

real-time scheduling, mode change protocol, multiprocessor, reconfigurable hardware

## ACM Reference Format:

Joël Goossens, Xavier Poczekajlo, Antonio Paolillo, and Paul Rodriguez. 2019. ACCEPTOR: a model and a protocol for real-time multi-mode applications on reconfigurable heterogeneous platforms. In *Proceedings of RTNS'19: 27th*

*International Conference on Real-Time Networks and Systems (RTNS'19)*. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

## 1 INTRODUCTION

The performance requirements for embedded real-time applications are rising. Machine vision systems and intelligent vehicles are only a few examples taken from the wide range of those demanding applications. Several techniques are used to design those real-time applications. One of them is to model the different functionalities as a set of different recurrent tasks. Some applications have the particularity to go through several *states* (or *modes*) during their lifespans. Those modes may be for an example event-driven (the detection of a specific event by the system) or time-driven (day/night modes). Each mode may then be designed independently, with its own set of functionalities modelled as several task sets: one for each mode. Only one mode and its task set can be active at a time. Executing a real-time multi-mode application requires to be able to schedule each mode, but also to handle the transition from one mode to another. During a transition, the tasks from the old mode must be completed, and the tasks from the new mode must be enabled. The transition from one mode to another must be bounded by a real-time constraint, defined at design time.

Dividing an application into several distinct modes improves the schedulability analysis. Indeed, the analysis is then closer to what will happen during the life-span of the application which makes it more accurate. By dividing the tasks into several task subsets—one for each mode, the maximal instantaneous workload is reduced, leading to a potential reduction of the hardware requirements. This technique has been well studied on uni-processor and multi-processor systems (see [13–15, 19]).

Modern hardware offers to the designer of real-time applications the possibility to be reconfigured during run-time. Such hardware can be reconfigured in several ways: it can be turned on and off (to adjust the energy consumption to the current workload), or its speed and/or functionalities. Modifying the hardware functionalities at run-time allows to reduce the hardware requirements. Indeed, several specialised processors working at distinct time instants may be replaced by a single reconfigurable processor (any hardware that may be used to execute tasks, and may change its behaviour at run-time). For example, last generation FPGAs (field-programmable gate array) offer all those possibilities. This kind of hardware is included in new embedded processing platforms, which

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*RTNS'19, Nov 06–08, 2019, Toulouse, France*

© 2019 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

makes them heterogeneous in nature. For example, the Xilinx Zynq UltraScale+™ system on chip [22] is a single die integrating several ARM processing cores from different architectures together with other heterogeneous components (64 bits quad-core Cortex® A53 cores, dual-core Cortex® R5 real-time cores, a Mali™-400 MP2 GPU and an FPGA programmable logic). This processor's block diagram is illustrated by Fig. 1.

In this paper, we propose a new approach: combining both software and hardware reconfigurations. In this approach, each mode is defined by a set of functionalities and a specific hardware configuration adapted to those functionalities. The hardware reconfigurations are handled during the transition from one mode to another. To the best of our knowledge, combining both multi-mode software and hardware reconfiguration has never been studied before. Nevertheless, fitting the hardware to the current functionalities of the application still leads to a more efficient execution (in terms of speeds, energy consumption and/or hardware requirements).

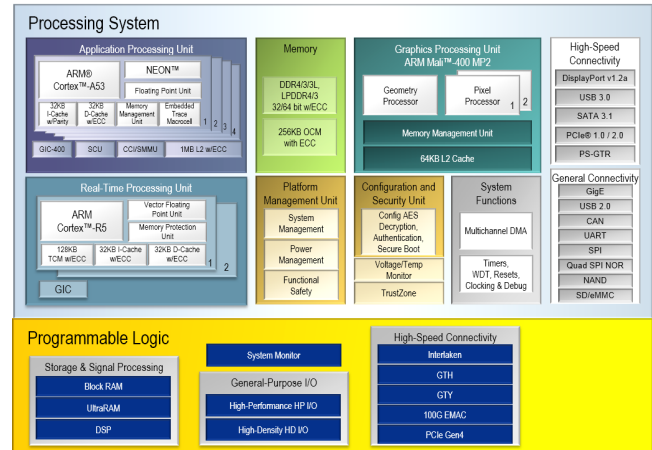
**Contributions.** The key contributions of this paper are the following:

- the—to our knowledge—first model for multi-mode applications on reconfigurable heterogeneous platforms;
- an asynchronous aperiodic mode change protocol to schedule such applications on reconfigurable multi-processor systems;
- a validity test for this protocol based on an upper-bound of the transition duration;
- a complete evaluation through simulations, complexity and competitive analysis of the contributions.

The remaining of the paper is organised as follows. Section 2 presents the related work. Section 3 introduces the model used in the paper. Section 4 describes the mode change protocol proposed as a first solution to handle the transition between the different modes. Section 5 proposes an upper-bound of the transitions duration, when using the introduced protocol. Section 6 provides a validity test and proves its correctness. Section 7 evaluates the protocol through a complete evaluation using simulations for the upper-bound efficiency, and both a complexity and competitive analysis of the protocol. It also discusses some of the limitations of the model. Section 8 concludes the paper.

## 2 RELATED WORK

**Multi-mode.** A survey [19] proposes various solutions and defines the main vocabulary and concepts for multi-mode applications on *uni*-processor systems. Concerning *multi*-processors, the literature reports several multi-mode protocols which handle the transition from one mode to another. Regarding *global* scheduling, V. Nelis' works include several protocols for *homogeneous* or *heterogeneous related* multi-processors (see for instance [13–15]). Those protocols combine different paradigms: periodic/aperiodic—a protocol is periodic iff it has mode-independent tasks, synchronous/asynchronous—a protocol is synchronous iff dependent tasks of two different modes can never be active simultaneously—task systems. More recently, Shih et al. [21] provided a schedulability analysis for global scheduling of mode change for the imprecise computation model upon identical multi-processors.



**Figure 1: The Zynq UltraScale+™ EG processor block diagram.** EG devices feature a quad-core ARM® Cortex-A53 platform running up to 1.5GHz, combined with dual-core Cortex-R5 real-time processors, a Mali-400 MP2 graphics processing unit, and a 16nm FinFET+ programmable logic [22].

Concerning *partitioned* scheduling a short contribution by Marinho et al. [12] formalises the scheduling problem and shows two counter-intuitive phenomena. Emberson et al. [6] proposed heuristics to handle the mode change. Lastly, Goossens et al. [9] consider the partitioned scheduling problem of multi-mode real-time systems upon *identical* multi-processors. The authors propose two methods for handling mode changes in partitioned scheduling.

In this work we extend our previous contributions where we introduced the study of multi-mode protocols on multi-processor platforms, with reconfigurable hardware [8].

**Scheduling upon heterogeneous platforms.** Cluster-based scheduling for heterogeneous systems has been well-studied, and Raravi et al. [18] propose, to the best of our knowledge, the most efficient approach to this problem.

**Reconfigurable hardware.** All the research concerning reconfigurable platforms combining CPUs and FPGA elements used for real-time systems is relatively new, as the platforms themselves are new to the market [11].

Cornil et al. [4] assessed the research challenges to face with this kind of platforms. Ahmad et al. [20] provided tools to optimise the design of real-time applications running on reconfigurable devices (with regards to different metrics such as performance and energy consumption). Pagani et al. [17] proposed the integration of Dynamic Partial Reconfiguration (DPR, a technique to reconfigure an FPGA at run-time) as part of a provided service of operating systems Biondi et al. [2, 3] provided several timing analyses and run-time framework works that make use of DPR, enabling reconfigurable heterogeneous platforms as target candidates for real-time systems. They also provided an extensive state of the art as part of their research paper [2]. Their approach is based on modelling the dependency between heterogeneous components with self-suspending tasks, waiting for resources to free on remote processing units.




Figure 2: Graph transition

Later, Pagani et al. [16] provided an implementation of their DPR framework for the Linux operating system. Bini [1] presented the Adaptive Fair Scheduler technique, that considers resource allocation and provide guarantees to the application. His approach is general enough to be applied to heterogeneous and reconfigurable computing.

### 3 MODEL

In this section, we present the first contribution of our paper: the model. To the best of our knowledge, this model is the first to combine both multi-mode application model and reconfigurable platform model. A multi-mode application on a reconfigurable platform is composed of different modes  $M^h = \langle \tilde{\Theta}^h, T^h, \Delta^h \rangle$ . Each mode will execute a different task set on the platform configured in a specific way for this mode. The application is composed of  $\mu$  modes  $M \doteq \{M^1, M^2, \dots, M^\mu\}^1$ , executed on a platform containing  $m$  processors.

Two use-cases of the model are shown in Section 3.6, as examples. All the notations introduced in this section are referenced in Table 1.

#### 3.1 Mode transitions

The application executes at any instant a mode  $M^h$ . The active mode may only change during a Mode Change Phase. A Mode Change Phase is triggered when the system receives a *Mode Change Request*. It can only be received whenever the platform is not already performing a Mode Change Phase. When a Mode Change Request  $MCR(M^{dst})$  occurs at  $t_{MCR}$ , the current mode is immediately deactivated, and new mode  $M^{dst}$  must be activated by  $t_{MCR} + \Delta^{dst}$ .  $\Delta^{dst}$  is a real-time constraint specified at design time for each mode  $M^h$ . The Mode Change Phase ends when the new mode  $M^{dst}$  is activated. Do note that this constraint depends only from the destination mode  $M^{dst}$ , independently from  $M^{src}$ .

**3.1.1 Mode Change Graph.** In an application, some transitions will never occur. The possible transitions are represented in the *Mode Change Graph*  $\mathcal{G} \doteq \{V, E \subseteq V^2\}$ . The Mode Change Graph is a directed graph, where  $V$  contains one and only one node for each mode  $M^h \in M$ , and  $E$  represents all the possible transitions from one mode to another. A Mode Change Phase from a mode  $M^{src}$  to a mode  $M^{dst}$  is possible if and only if  $(M^{src}, M^{dst}) \in E$ . A graph example is given in Fig. 2. In this example, the mode after  $M^2$  can only be  $M^3$  but not  $M^1$  because the directed edge from  $M^1$  to  $M^2$  goes only in one way.

<sup>1</sup>The notation " $\doteq$ " means equal by definition.

### 3.2 Task model

**3.2.1 Job.** The platform executes jobs. A job  $J_i$  is defined by three parameters  $(a_i, c_i, d_i)$ . The arrival time in the system is denoted by  $a_i$ . The worst case execution time (WCET)  $c_i$  is computed on an arbitrary platform running at a unitary rate (i.e. one computing unit per time-unit). Finally,  $d_i$  is its absolute deadline. The platform must therefore execute the job  $J_i$  for  $c_i$  computing units between  $a_i$  and  $d_i$  to be completed. A job is said to be active if it has arrived in the system and is not completed yet. It is said to be inactive otherwise.

**3.2.2 Task.** In the application,  $n$  tasks are forming the task set. A task  $\tau_i$  is defined by three parameters  $(C_i, D_i, T_i)$ .  $C_i$  is the WCET,  $D_i$  is the relative deadline and  $T_i$  is the minimum inter-arrival time. In this work, we are considering implicit deadlines only:  $\forall i, D_i = T_i$ .

A task  $\tau_i$  generates an infinite sequence of jobs. A task releases a job if and only if it is enabled. When a task is disabled, its incomplete jobs are called rem-jobs. It is enabled (resp. disabled) when its mode becomes active (resp. inactive). A job  $J_i$  released at  $a_i$  has an absolute deadline of  $d_i \doteq a_i + D_i$  and  $c_i \doteq C_i$ .

### 3.3 Platform model

From simple uniprocessor platforms, modern models are now also targeting complex platforms. Multiprocessors platforms can be either modelled as an *identical*, *uniform* or *unrelated* machine [7]. An *identical machine* is a platform on which all the processors are identical. A *uniform machine* is formed of similar processors with different performances. In this paper, we consider *unrelated machines*. This kind of platform is formed of processors having different instruction sets and/or different performance. More than unrelated processors, our model considers reconfigurable processors. Those processors may change their instruction sets or performances at run-time.

The platform  $P$  is composed of  $m$  processors. Each processor  $p_q$  has a type  $\pi^k$  with  $k \in [1, \dots, \phi]$ , and  $q \in [1, \dots, m]$ . There are  $\phi$  different processor types. The  $q^{\text{th}}$  element of  $\Pi$ —denoted  $\Pi_q$ —contains the type of the processor  $p_q$ . The type defines the possible configurations for the processor.

**Configurations.** A processor of type  $\pi^k$  is configured at any time in configuration  $\theta_c \in \Theta_k$ . A configuration defines several parameters like the instruction set of the processor or the speed. The set of sets  $\Theta \doteq \{\Theta_1, \Theta_2, \dots, \Theta_\phi\}$  contains the possible set of configurations for all types. A configuration is accessible to processors of one and only one type: i.e.  $\forall k, k', k \neq k' \implies \Theta_k \cap \Theta_{k'} = \emptyset$ . There are  $o$  different configurations, with  $o \doteq \sum_{k=1}^{\phi} |\Theta_k|$ . Reconfiguring a processor is not instantaneous. It takes  $\delta_c$  time-units (denoted as the reconfiguration delay) to reconfigure a processor of type  $\pi^k$  to  $\theta_c$  if  $\theta_c \in \Theta_k$ , otherwise it takes  $+\infty$ . This model deliberately abstracts some implementation details. For example, we provide no notation to get the configuration of a specific processor at a given time.

**(Non-)Reconfigurable.** A processor  $p_q$  of type  $\pi^k$  may be *reconfigurable* or *non-reconfigurable*. It is said to be *reconfigurable* if and only if  $|\Theta_k| > 1$ .

**(In-)Active.** A processor may be turned on and off dynamically. This may be used to manage power consumption. A processor

$p_q$  is inactive at time  $t_1$  if and only if it is configured in an idle-configuration. An idle-configuration is a configuration in which a processor cannot execute any job, i.e. its speed is null for any job. It is said active otherwise.

### 3.4 Job progression rate

On heterogeneous machines, the job progression rate (speed at which a job is being completed) depends on both the job and the processor. Specifically, the job progression rate  $r_{i,c}$  on the processor  $p_q$  depends on both the job  $J_i$  and the current configuration  $\theta_c$  of  $p_q$ . Formally, Job  $J_i$  completes  $t \times r_{i,c}$  computing units when executed on a processor configured in  $\theta_c$  for  $t$  time-units. This amount may be null if the task cannot be executed on this configuration.

### 3.5 Mode model

A mode  $M^h = \langle \tilde{\Theta}^h, T^h, \Delta^h \rangle$  is a combination of a specific hardware configuration, a specific software and a real-time constraint. In our model, the software is represented by the task set to execute and the hardware by the configurations of the processors. The real-time constraint bounds the maximum delay to reconfigure the whole platform to the new mode. In the literature, previous works—as Nélis' work in [15]—were only considering software reconfiguration. Reconfiguring both *hardware and software* is the originality of our work.

Each mode has a cluster-based approach. It will consider several independent groups of processors—denoted as clusters, with a global approach inside each cluster. To do so, each mode splits its task set into several task subsets, and schedule each subset upon a distinct cluster.

**3.5.1 Mode specifications.** The mode  $M^h$  is defined by the following specifications:

**Hardware.** The processors must be configured in a specific way. The configurations vector  $\tilde{\Theta}^h \doteq \langle m_1^h, m_2^h, \dots, m_o^h \rangle$  contains for a specific mode  $M^h$  the required number  $m_c^h$  of processors configured in  $\theta_c$ .

The required configurations for all the modes are specified as a vector of configuration vectors:  $\tilde{\Theta} \doteq \langle \tilde{\Theta}^1, \tilde{\Theta}^2, \dots, \tilde{\Theta}^\mu \rangle$ .

**Software.** The mode  $M^h$  has to execute a specific task set  $T^h$ , containing  $\nu^h$  tasks. To partition the task set among several clusters,  $T^h$  itself is divided into task subsets: one per configuration present in the mode  $M^h$ . The processors configured in  $\theta_c$  must run the task subset  $T^{h,c}$ . The set  $T \doteq \{T^1, T^2, \dots, T^\mu\}$  contains the task subsets for all the modes. The tasks are said to be *mode-dependent*: each task may appear in at most one mode, i.e.  $\forall h, h', h \neq h' \implies (U T^{h,c}) \cap (U T^{h',c}) = \emptyset$ . This assumption will be discussed in Section 7.4.

**Delay.** Switching from one mode to another is not instantaneous. Rem-jobs cannot be stopped before completion and reconfiguring a processor takes a delay. However, this delay must be bounded to take into account the real-time constraints of the application.  $\Delta^h$  represents the maximum possible delay for reconfiguring the system after a Mode Change Request to  $M^h$ . The set  $\Delta \doteq \{\Delta^1, \Delta^2, \dots, \Delta^\mu\}$  contains the real-time constraint of each mode.

**Table 1: Summary of the notations**

A mode	$M^h \in M$
Mode change instant	$t_{\text{MCR}}$
The mode change graph	$\mathcal{G}$
A job	$J_i$
A task	$\tau_i$
A processor	$p_q \in P$
A type	$\pi^k$
Type of processor $p_q$	$\Pi_q$
A configuration of type $\pi^k$	$\theta_c \in \Theta_k$
Reconfiguration delay to $\theta_c$	$\delta_c$
Progression rate for $J_i$ on $\theta_c$	$r_{i,c}$
Configurations requirements of Mode h	$\tilde{\Theta}^h \in \tilde{\Theta}$
Task subset of a mode $M^h$ for proc. in $\theta_c$	$T^{h,c}$
Task set of a mode $M^h$	$T^h \in T$
Mode $M^h$ 's real time constraint	$\Delta^h \in \Delta$

### 3.6 Model example

In this section, we provide instances of our model to match existing real-world platforms.

**Example 1.** As we explained, the platform is organised in several unrelated clusters of identical machines. When using DPR (Dynamic Partial Reconfiguration), the Zynq UltraScale+™, may be configured such that it is composed of the following identical platforms:

- a four cores identical platform meant to execute general purpose application software (the 4 Cortex-A53 cores);
- a dual core identical platform that may host highly critical and predictable software (the dual Cortex-R5 cores);
- a GPU dedicated to display information (the Mali™-400 GPU). We model it as a single processor;
- 4 independent hardware accelerators that are each dedicated to one specific task. These accelerators are hosted in the Programmable Logic (FPGA) which may arbitrarily be divided in 4 independent slots by the system designer (and managed by the DPR engine). Each one of these accelerators/slots corresponds to a single processor and the whole Programmable Logic is then modelled as an identical platform with 4 processors.

In this example, we assumed that the system designer used Dynamic Partial Reconfiguration (DPR) capabilities offered by the Zynq chip family [23] (notice that earlier research allow the use of DPR in real-time systems as described in Section 2). Using DPR, we assume that the designer defined 4 processors that are reconfigurable as defined in Section 3.3.

In this example, the platform  $P$  is composed of the following  $m = 11$  processors:

- $p_1, p_2, p_3, p_4$ : the application cores, forming the Application Processing Unit (APU). All these cores have the same architecture (A53) and belong to type  $\pi^1$ :  $\Pi_1 = \Pi_2 = \Pi_3 = \Pi_4 = \pi^1$ . As the APU is a static CPU, it has only one configuration, meaning  $\Theta_1 = \{\theta_1\}$ , where  $\theta_1$  denotes the only possible configuration of the APU.
- $p_5, p_6$ : the real-time cores, forming the Real-Time Processing Unit (RPU). These cores have the same architecture (R5)

and then belong to type  $\pi^2$ . Similarly to the APU, we have  $\Theta_2 = \{\theta_2\}$ , where  $\theta_2$  denotes the only possible configuration of the RPU.

- $p_7$ : it represents the GPU, which has its own type  $\pi^3$  and is also a static processor:  $\Theta_3 = \{\theta_3\}$  where  $\theta_3$  denotes the only possible configuration of the GPU.
- $p_8, p_9, p_{10}, p_{11}$ : the different dynamically partially reconfigurable processors of the chip Programmable Logic (from the FPGA). These 4 processors are independent. We assume that the processors are specialised for some processing: we have two types of processors  $\pi^4$  and  $\pi^5$ . We assume that in this example,  $p_8$  and  $p_9$  are image processing kernels (allowing to process an input image and apply a filter on it), and that  $p_{10}, p_{11}$  implement cryptographic accelerators in order to implement several block cipher algorithms. The image processing kernels can be in three configurations: sepia, sobel and grayscale. The cryptographic accelerators can be configured either to run AES or 3DES block ciphers. Therefore:

$$\begin{cases} \Theta_4 = \{\theta_{\text{sepia}}, \theta_{\text{sobel}}, \theta_{\text{gray}}\} \\ \Theta_5 = \{\theta_{\text{aes}}, \theta_{\text{des}}\} \end{cases}$$

The rates of task  $\tau_{\text{aes}}$  depend on the configuration of its processor. This task can be performed only on RPU, GPU or a specialised FPGA. Therefore, we could have :  $r_{\text{aes},\text{aes}} \gg r_{\text{aes},3} > r_{\text{aes},2} > r_{\text{aes},\text{des}} = 0$ . Obviously, this task is completed faster on a specialised FPGA, and cannot be executed on an FPGA wrongly configured (hence  $r_{\text{aes},\text{des}} = 0$ ). This platform example is depicted in Fig. 3(a). The platform is represented with several divisions into *clusters*. Typical reconfiguration times of the Programmable Logic are of the order of a few milliseconds [17]. These timings could be used to set the values of  $\delta_{\text{sepia}}, \delta_{\text{sobel}}, \delta_{\text{gray}}, \delta_{\text{aes}}$  and  $\delta_{\text{des}}$ .

**Example 2.** The ARM big.LITTLE architecture is a technique allowing to switch between high-performance cores and low-power cores. A typical design of this architecture is composed of 4 A57 cores (the high-performance cluster) and 4 A53 cores (the low-power cluster) used alternately. These kind of platforms suits our model as the multi-mode reconfiguration is implemented in the hardware. In this case, we model the platform  $P$  as a set of  $m = 8$  processors:

- $p_1, p_2, p_3, p_4$ : the high-performance cluster. All cores are of type  $\pi^1$  and  $\Theta_1 = \{\theta_{A57}, \theta_{\text{off } 57}\}$ , these configurations meaning that the cores are either active or inactive.
- $p_5, p_6, p_7, p_8$ : the low-power cluster. They are of type  $\pi^2$  and, similarly to the other type, we have  $\Theta_2 = \{\theta_{A53}, \theta_{\text{off } 53}\}$

In some platform designs such as Samsung Exynos 5 Octa, only one kind of core can run simultaneously—the operating system must explicitly switch the whole platform from one kind of core to the other. This constraint can be enforced by defining the clusters and modes accordingly and by using the inactive configurations  $\theta_{\text{off } 53}$  and  $\theta_{\text{off } 57}$ . For example, if we define two modes  $M^{53}$  and  $M^{57}$ ,  $\Theta^{53} = \langle m_{A53}^{53}, m_{A57}^{53} \rangle = \langle 4, 0 \rangle$  and  $\Theta^{57} = \langle m_{A53}^{57}, m_{A57}^{57} \rangle = \langle 0, 4 \rangle$ . This means that during mode  $M^{53}$ , only the processors of type  $\pi^2$  will be active.

These two examples illustrate how generic our model is and how it is able to capture the essence of real-world platforms.

## 4 PROTOCOL

In this section, we introduce the second contribution of this paper: the protocol to handle mode transitions on a reconfigurable platform. It is called ACCEPTOR, for : AsynChronous ClustEr-based ProTOcol for multi-mode applications on Reconfigurable platforms. ACCEPTOR conducts the reconfiguration phase: it schedules the rem-jobs, performs the requested reconfigurations and enables the new mode tasks. It also respects the real-time constraints such as the delay for new mode's activation and the rem-jobs hard deadlines.

The description of the protocol will be preceded by an introduction about clusters.

### 4.1 Clustered platform

A clustered platform is composed of several groups of processors called *clusters*. Formally, a cluster is formed by several processors, configured *identically*. At a given instant  $t$ , a processor of type  $\pi^k$  configured in  $\theta_c$  belongs to one and only one cluster. This cluster is composed of all processors of type  $\pi^k$  configured in  $\theta_c$ .

In this paper we will use the following popular concepts regarding multi-mode problem:

**Work-conserving** A scheduler is *work-conserving* iff the processors can be idle only when there are no jobs waiting to execute.

**The cluster's makespan** is the required time to complete all the rem-jobs generated by the task set of the cluster. [15] proposed an upper-bound  $\overline{MS}$  that will be used in the protocol.

**Idle** A cluster is said to be *idle* iff all its processor are idle. A processor is idle if it is neither executing tasks nor being reconfigured.

A cluster( $\pi^k$ ) is composed of processors of type  $\pi^k$ , configured identically.

A cluster( $\theta_c$ ) is a cluster composed of processors configured in  $\theta_c$ .


Because the processors may be reconfigured, the clusters may (only) be modified during a Mode Change Phase. Clustering is based on the configuration of the system, and may thus change from one mode to another. Fig. 3 represents the same platform with two different clustering. For example,  $p_8$  is on a cluster with  $p_9$  in the first mode (see (a)) but is then alone in the second mode (see (b)).

### 4.2 ACCEPTOR


The ACCEPTOR protocol is a periodic asynchronous protocol. It can be used with sporadic task sets, scheduled by a clustered work-conserving scheduler.

The approach is to schedule the rem-jobs, and reconfigure each processor once it becomes idle. This protocol is composed of an offline phase and a run-time phase. The offline phase determines for each cluster the required reconfigurations for each possible Mode Change Phase (see Section 4.2.1). The run-time phase is dedicated to schedule the rem-jobs and reconfigure the processors (see Section 4.2.2). This run-time phase focuses on each cluster individually. An illustration is given in Fig. 4. Finally, an example of the protocol is given in Section 4.2.3.

**4.2.1 Offline: computing the required reconfigurations.** The offline phase simply computes which reconfigurations must be done for each Mode Change Phase from  $M^{\text{src}}$  to  $M^{\text{dst}}$ . This will be done



**Figure 3: Two possible divisions of the example platform (based on the Zynq UltraScale+™) into several clusters.**



**Figure 4: Protocol illustration**

iteratively for each type  $\pi^k$ . To do so, the protocol first computes the differences between the required reconfigurations of both modes.

Then, it computes the makespan of each  $\pi^k$ -cluster, using the bound provided by [15].

And finally it assigns the longest (resp. shortest) required reconfigurations (based on the reconfiguration delays) to processors belonging to clusters( $\pi^k$ ) having the shortest (resp. longest) makespans (based on the makespan upper-bound from [15]). The assignment is denoted  $\{\theta_c, \theta_{c'}\}$  where  $\theta_c$  is the configuration of the processor in  $M^{\text{src}}$  and  $\theta_{c'}$  is the configuration of the processor in  $M^{\text{dst}}$ .

Formally,

- Compute for each clusters( $\theta_c$ ) in each mode  $M^h$  the makespan upper-bound  $\overline{MS}$  for the task subset  $T^{h,c}$  on  $m_c^h$  processors;

Then,  $\forall (M^{\text{src}}, M^{\text{dst}}) \in \mathcal{G}, k \in [1, \dots, \phi]$ :

- (1) The missing configurations for clusters( $\pi^k$ ) are all configurations  $\theta_c \in \Theta_k$  where  $m_c^+ > 0$ , with  $m_c^+ \doteq \max(0, m_c^{\text{dst}} - m_c^{\text{src}})$ . Symmetrically, the excess configurations are all  $\theta_c \in \Theta_k$  where  $m_c^- > 0$ , with  $m_c^- \doteq \max(0, m_c^{\text{src}} - m_c^{\text{dst}})$ .
- (2) To perform the reconfiguration, we use a vector containing missing configurations. This vector will be used at run-time to track the required configurations to perform. Store  $\tilde{\theta}^+ \doteq \langle \theta_c, \theta_{c'}, \dots \rangle$  a vector of missing configurations.  $\tilde{\theta}^+$  contains  $m_c^+$  times the configuration  $\theta_c$ . Order  $\tilde{\theta}^+$  by reconfiguration delay, in decreasing order.

- (3) Symmetrically, store  $\tilde{\theta}^- \doteq \langle \theta_c, \theta_{c'}, \dots \rangle$  a vector of excess configurations.  $\tilde{\theta}^-$  contains  $m_c^-$  times the configuration  $\theta_c$ . Order  $\tilde{\theta}^-$  by the cluster's makespan of the processors configured in  $\theta_{c''} \in \tilde{\theta}^-$ , in increasing order.
- (4) Bind the  $x^{\text{th}}$  longest reconfiguration to the  $x^{\text{th}}$  idle processor. The  $x^{\text{th}}$  longest reconfiguration  $\theta_c$  is determined using its reconfiguration delay  $\delta_c$ . The  $x^{\text{th}}$  idle processor is the  $x^{\text{th}}$  processor to become idle on this cluster during this Mode Change Phase. To do so, define the assignment  $\{\theta_c, \theta_{c'}\}$ , where  $\theta_c$  is the  $x^{\text{th}}$  element of  $\tilde{\theta}^-$  and  $\theta_{c'}$  is the  $x^{\text{th}}$  of  $\tilde{\theta}^+$ .
- (5) Store the assignment  $\{\theta_c, \theta_{c'}\}$  in the multi-set  $RT^{\text{src}, \text{dst}}$ .

The results will be used at run-time. Because a cluster is formed by processors of the same type and configured identically the required reconfigurations for the cluster( $\theta_c$ ) are contained in the multiset  $\{\{\theta_{c'}, \theta_{c''}\} \mid c = c' \wedge \{\theta_{c'}, \theta_{c''}\} \in RT^{\text{src}, \text{dst}}\}$ .

**4.2.2 Run-time: scheduling and reconfiguring.** At run-time, the protocol must successfully schedule the rem-jobs, and reconfigure the processors as computed in the offline phase. This phase is straightforward: it will schedule the tasks using the same scheduler as before. The only difference concerns the reconfigurations: when a processor becomes idle, it may be reconfigured. For each cluster( $\theta_c$ ) it takes as input the list of remaining rem-jobs  $J$  and the required reconfigurations to make  $\{\{\theta_{c'}, \theta_{c''}\} \mid c = c' \wedge \{\theta_{c'}, \theta_{c''}\} \in RT^{\text{src}, \text{dst}}\}$ .

Formally,

- Schedule  $J$  using the same scheduler as before;
- When a processor becomes idle, reconfigure it to the longest required reconfiguration  $\theta_{\text{dst}}$  that hasn't been done yet, where  $\theta_{\text{dst}} \in \{\{\theta_{c'}, \theta_{\text{dst}}\} \mid c = c' \wedge \{\theta_{c'}, \theta_{\text{dst}}\} \in RT^{\text{src}, \text{dst}}\}$ .

Whenever  $m_c^{\text{dst}}$  processors are reconfigured to  $\theta_c$ , the cluster( $\theta_c$ ) of the new mode  $M^{\text{dst}}$  is formed and its task set is activated. This phase ends once every cluster is idle and reconfigured. The new mode  $M^{\text{dst}}$  is then enabled.

**4.2.3 Example.** Fig. 4 depicts an example of the protocol's execution. The platform is composed of 4 processors ( $m = 4$ ). Processors  $p_1$  and  $p_2$  share the same type (i.e.  $\Pi_1 = \Pi_2$ ), as  $p_3$  and  $p_4$  (i.e.  $\Pi_3 = \Pi_4$ ). Two processors must be configured in  $\theta_1$  for mode  $M^{\text{src}}$ :  $\tilde{\theta}^{\text{src}} = \{2, 2, 0, 0, 0\}$ . For example,  $m_1^{\text{src}} = 2$ . For mode  $M^{\text{dst}}$ :  $\tilde{\theta}^{\text{dst}} = \{0, 0, 2, 1, 1\}$ . This means that there will be three clusters in this mode, because there are three different types of required configurations. Here are some of the different reconfigurations delay:  $\delta_3 = 1, \delta_4 = 3, \delta_5 = 2$ . The real-time constraint is  $\Delta^{\text{dst}} = 10$ .

The vectors  $\tilde{\theta}^- = \{2, 2, 0, 0, 0\}$  and  $\tilde{\theta}^+ = \{0, 0, 2, 1, 1\}$  can be trivially computed. With this, we obtain the required reconfiguration to be performed. Using the makespan of cluster 1 and cluster 2, we can assign those reconfigurations. At the end of the computation, the multi-set  $RT^{\text{src}, \text{dst}} = \{(\theta_1, \theta_3), (\theta_1, \theta_3), (\theta_2, \theta_4), (\theta_2, \theta_5)\}$

At  $t = 0$ ,  $p_1$  and  $p_2$  form the cluster( $\theta_1$ ) denoted as cluster 1.  $p_3$  and  $p_4$  form the cluster( $\theta_2$ ) denoted as cluster 2. On cluster 1, we can see that the task  $\tau_4$  finishes its job at  $t_{\text{MCR}} = 4$ . Thus, it has no job in the system when the Mode Change Phase starts and generates no rem-job. The same occurs for  $\tau_6$ . The others tasks have all a job active, which becomes a rem-job at  $t_{\text{MCR}}$ . For example, the execution of  $\tau_5$  starts at  $t = 3$  and continues afterwards.



During the Mode Change Phase, the same task pattern occurs, because the scheduler is the same as before.  $p_4$  becomes idle at  $t = 5$ . There are no other jobs to execute which leads the protocol to reconfigure it. The cluster( $\theta_2$ ) has the following reconfigurations to make:  $(\theta_2, \theta_4), (\theta_2, \theta_5)$ . The reconfigurations to  $\theta_5$  are made first because  $\delta_5 > \delta_4$ . These reconfigurations end at  $t = 8$ . Because  $m_5^{\text{dst}} = 1$ , a cluster of the new mode is formed and its task set  $\{\tau_{12}, \tau_{13}\}$  is activated. On the other hand, even if  $p_2$  has been reconfigured at  $t = 8$ , because  $m_3^{\text{dst}} = 2$ : the cluster is not formed yet. At  $t = 9$ , all the new mode's clusters have been formed. The mode  $M^{\text{dst}}$  can then be enabled. The phase last for  $9 - 4 < \Delta^{\text{dst}}$  units of time, and all the rem-jobs were successfully scheduled so the Mode Change Phase succeeds.

## 5 MAKESPAN UPPER-BOUND

To assert the feasibility of a given system using the ACCEPTOR protocol, we now introduce an upper-bound of the instant in which the cluster's rem-jobs are completed, and the required reconfigurations are done. This instant is denoted as  $\overline{\text{reconfigured}}$ , and can be immediately derived to obtain an upper-bound on any Mode Change Phase. Then, the validity of the protocol will be demonstrated in Section 6.

The new upper-bound denoted as  $\overline{\text{reconfigured}}$  will be used to verify the feasibility of an application running on a given platform. Before that, we will introduce a result from [13] from which we will build on: the upper-bound on the  $\text{Idle}_k$  instant denoted as  $\overline{\text{Idle}_k}$ . An  $\text{Idle}_k$  instant is the earliest instant such as at least  $k$  processors are idle.

LEMMA 5.1 (LEMMA 2.11 IN [13]). *Suppose that  $J$  is ordered by non-decreasing job processing times, i.e.,  $c_1 \leq c_2 \leq \dots \leq c_n$ , all starting at  $t = 0$ . Then, whatever the job priority assignment, an upper-bound on the  $\text{Idle}_k$  instant,  $1 \leq k \leq m$ , when scheduled by a FJP work-conserving global scheduler upon  $m$  identical processors, with no reconfiguration is given by:*

$$\overline{\text{Idle}_k} = \begin{cases} c_k & \text{if } n = m \\ \frac{\sum_{i=1}^n c_i + (k-1) \times c_{n-m+k}}{m} & \text{otherwise} \end{cases} \quad (1)$$

The new upper-bound  $\overline{\text{reconfigured}}$  is based on  $\overline{\text{Idle}_k}$ . It simply searches for the processor with the maximum sum of idle time + reconfiguration delay.

COROLLARY 5.2 (UPPER BOUND ON THE CLUSTER'S MAKESPAN). *Suppose that  $J$  is ordered by non-decreasing job processing times, i.e.,  $c_1 \leq c_2 \leq \dots \leq c_n$ , all starting at  $t = 0$ .*

*Suppose that the required reconfigurations' delays are in the vector  $\Delta$ , ordered by reconfiguration time in decreasing order, and that  $\delta_q$  is the  $q^{\text{th}}$  element of the vector.*

*Then, whatever the job priority assignment, an upper-bound on the cluster's makespan  $\overline{\text{reconfigured}}$ , when scheduled by a FJP work-conserving clustered scheduler is*

$$\overline{\text{reconfigured}} = \max_{q=1}^m \{\overline{\text{Idle}_q} + \delta_q\} \quad (2)$$

PROOF. The proof is trivially obtained by construction.  $\square$

## 6 VALIDITY

A validity test is used to determine whether the Mode Change Phases of an application will be successfully managed by a multi-mode protocol, on a specific system. To ensure this, we need to verify that every deadline will be met during the transition phase, and that the new mode will always be enabled on time.

We now introduce the sustainability. This notion will be useful to prove that scheduling rem-jobs is at most as difficult as scheduling the recurrent tasks.

**Sustainability.** A scheduler  $S$  is sustainable iff for any set of jobs  $J$  schedulable by  $S$ , any jobset  $J'$  derived from  $J$  where all jobs are identical or smaller (equal or smaller C) will be feasibly and non-strictly faster scheduled by  $S$ .

Lemma 6.1 from [10] will be used to show that any scheduler used in the protocol will be predictable, by hypothesis. This lemma applies to any uniform multiprocessor platform. It also applies to identical multiprocessor platform, because identical is a special case of uniform multiprocessor platform.

LEMMA 6.1 ([5]). *Any work-conserving and Fixed Job Priority (FJP) algorithm is predictable on uniform multiprocessor platforms.*

The schedulability of the rem-jobs during Mode Change Phase is ensured by the Lemma 6.1.

LEMMA 6.2 (REM-JOB'S SCHEDULABILITY). *Every rem-job's deadline of every cluster will be respected during every possible Mode Change Phase, iff (i) the scheduler is a clustered, FJP, preemptive, work-conserving scheduler and (ii)  $T^{\text{src},c}$  is schedulable by the scheduler upon  $m_c^{\text{src}}$  processors configured in  $\theta_c$ .*

PROOF. Because of (i) and Lemma 6.1, we know that the scheduler is predictable. Hence, the scheduler is predictable and (ii): the rem-jobs of the taskset  $T^{\text{src},c}$  are schedulable on their cluster. Thus, the schedulability is ensured.  $\square$

Lemma 6.3 ensures that the real-time constraint of each mode will be respected during every possible transition.

LEMMA 6.3 (TRANSITION'S TIME CONSTRAINT). *The protocol will respect the time constraint  $\Delta_{\text{dst}}$  iff  $\forall \text{src}, \text{dst}, (M^{\text{src}}, M^{\text{dst}}) \in \mathcal{G}, \forall \text{cluster}(\theta_c)$ , with  $\Delta_c^{\text{src},\text{dst}}$  being the required reconfigurations's delays ordered in decreasing order:*

$$\Delta_{\text{dst}} \geq \max_{\text{cluster}(\theta_c)} \overline{\text{reconfigured}}(T^{\text{src},c}, m_c^{\text{src}}, \Delta_c^{\text{src},\text{dst}})$$

PROOF. The right part of the inequation computes the maximal upper-bound  $\overline{\text{reconfigured}}$ . By construction, if the hypothesis is respected, every cluster of each possible transition will be reconfigured when required.  $\square$

Theorem 6.4 proves the validity of our protocol, under its hypothesis.

THEOREM 6.4 (VALIDITY OF ACCEPTOR). *The protocol ACCEPTOR is valid iff (i) the scheduler is a clustered, FJP, preemptive, work-conserving scheduler, (ii)  $T^{\text{src},c}$  is schedulable by the scheduler upon  $m_c^{\text{src}}$  processors configured in  $\theta_c$  and (iii)  $\forall \text{src}, \text{dst}, (M^{\text{src}}, M^{\text{dst}}) \in$*

$\mathcal{G}, \forall \text{cluster}(\theta_c)$ , with  $\Delta_c^{\text{src}, \text{dst}}$  being the required reconfigurations's delays ordered in decreasing order:

$$\Delta_{\text{dst}} \geq \max_{\text{cluster}(\theta_c)} \overline{\text{reconfigured}}(T^{\text{src}, c}, m_c^{\text{src}}, \Delta_c^{\text{src}, \text{dst}})$$

PROOF. To prove the validity of the protocol, we must prove that every rem-job's deadline will be respected during the Mode Change Phase and that the real-time constraint  $\Delta_{\text{dst}}$  will be respected during each possible transition from every mode  $M^{\text{src}}$  to mode  $M^{\text{dst}}$ .

- Because of the hypotheses (i-ii), Lemma 6.2 can be applied and thus prove that every job's deadline will be respected during the Mode Change Phase.
- Because of the hypotheses (i-iii), Lemma 6.3 can be applied. Every real-time constraint will be therefore respected.

Because the job's deadlines and the real-time constraint will be respected, the protocol is valid when the hypotheses (i-iii) are respected.  $\square$

## 7 EVALUATION

In this section, we present a complete evaluation of the contribution. We first evaluate in Section 7.1 the theoretical worst case time complexity. Then, we use simulations in Section 7.2 to evaluate the pessimism of the bound introduced in Section 5. Section 7.3 performs a competitive analysis of the protocol itself, and finally Section 7.4 discuss a limitation of the model.

### 7.1 Complexity

We now evaluate the theoretical worst case time complexity of both the offline and run-time phases of the protocol.

Concerning the offline phase: the computation of  $\overline{\text{Idle}}_k$  has a worst case time complexity of  $O(n)$ , as MS. First, the protocol computes  $C$  makespan upper-bounds, where  $C \leq m$  is the number of clusters and  $n$  the largest number of tasks that a cluster has to schedule, hence a complexity of  $O(C \times n) = O(m \times n)$ . Steps 2 and 3 performed on clusters contain a sort on at most  $m$  elements, and Step 3 also computes  $\overline{\text{Idle}}_k$  for at most  $m$  processors; Thus, complexity of Step 2 is  $O(m \times \log(m))$  and complexity Step 3 is  $O(m \times \log(m) + m) = O(m \times \log(m))$ . The total complexity is thus  $O(m \times \log(m) + m \times n)$ .

Concerning the run-time phase, it has the same complexity as the scheduler. Choosing the reconfiguration to perform can be made in  $O(1)$ .

Complexities for both offline and run-time phases are very low, and make the protocol scalable.

### 7.2 Empirical pessimism of $\overline{\text{reconfigured}}$

To the best of our knowledge, the upper-bound  $\overline{\text{Idle}}_k$  pessimism has not been evaluated, neither theoretically nor empirically. We provided here an empirical evaluation of  $\overline{\text{reconfigured}}$ , which is based on  $\overline{\text{Idle}}_k$ . We measure the pessimism ratio  $\frac{\text{upper-bound}}{\text{mode change time}}$ , where the *mode change time* is a measurement of the simulated Mode Change Phase, and where the *upper-bound* corresponds to the computation of  $\overline{\text{reconfigured}}$  for a given cluster.

The experiment has been conducted for clusters with  $m = 2, 4, 8, 16, 32, 64, 128$  processors and  $n = m \times 7$  being the number of tasks




Figure 5: Pessimism of  $\overline{\text{reconfigured}}$

scheduled on this cluster. The reconfiguration times are uniformly chosen in the range  $[0, 10]$ , where 0 indicates that no reconfiguration is required. We generate approximately 1 000 feasible task sets with utilisations  $\in (p - 0.1, p]$ , where we increase  $p$  from 0.6 to 1.0 in steps of 0.1. The deadlines are uniformly chosen in the range  $[1, 20]$ . As a scheduler, we use Global-RM and remove any non-schedulable taskset from our experiments. For those experiments, time is discrete.

The results are displayed in Fig. 5 and can be read this way: the average pessimism ratio is 1.14 for a cluster with  $m = 4$  and a taskset utilisation of 0.9. This means that the upper-bound is (on average) 14% larger than the actual duration in this configuration. One can observe two major trends: first, the more processors and tasks there are, the more pessimistic the bound gets. Secondly, a higher utilisation leads to a higher (but bounded) pessimism until  $u \in (0.9, 1.0]$  for most cases. At  $u = 1.0$ , the pessimism ratio drops to 1.0. Concerning the first assertion: it is worth mentioning that this phenomena is (empirically) bounded: the size of the clusters are doubled when the pessimism per cluster size grows less and less. Even more: the pessimism is near equal for clusters of size  $m = 32, 64$  and  $128$ . Concerning the second assertion: the pessimism drops to 1.0 when the utilisation is superior to 0.91 because a lot of tasks with low inter-arrival time. This leads to very low  $\overline{\text{Idle}}_k$ , being negligible in comparison with the highest reconfiguration time. Thus, the upper-bound is exact on those cases. Also, this upper-bound is never higher than 1.20 which is very accurate.

Those experimental results shows that the upper-bound  $\overline{\text{reconfigured}}$  is accurate and can be used for clusters having a large amount of processors or a high utilisation.

### 7.3 Competitive analysis of ACCEPTOR

After the empirical evaluation of the upper-bound  $\overline{\text{reconfigured}}$ 's pessimism, we now evaluate the compete analysis of the protocol itself. A protocol is said to be  $\lambda$ -competitive if it takes at most  $\lambda \times C$  time to complete a Mode Change Phase, where  $C$  is the optimal Mode Change Phase duration. To perform the competitive analysis



of the protocol, we search for a worst case, i.e. the system for which the relative difference between our protocol's performances and the optimal one is the highest. We first define what a *squeezable* system is, and then prove several properties on those systems, which leads to prove that they correspond to the worst case. Once we have this worst case, we compute the competitiveness of the protocol. An example of a *squeezable* system is depicted in Fig. 6. The top graph shows a specific Mode Change Phase handled by ACCEPTOR and below is the same Mode Change Phase handled by an optimal one.

Lemmas 7.1–7.3 prove several properties on the *squeezable* systems. Lemmas 7.4–7.5 introduce two properties that stand for any system, which leads to a global property in Corollary 7.6. Lemmas 7.7–7.8 use both the properties on *squeezable* systems and the global property to prove that the system  $S$  is a worst case in Corollary 7.9. Using the worst case, Theorem 7.10 proves that ACCEPTOR is  $\frac{2m-1}{m}$ -competitive on a cluster composed of  $m$  processors. Then, Corollary 7.11 proves that ACCEPTOR is 2-competitive.

As defined in Section 5, an  $\text{Idle}_k$  instant is the earliest instant such as at least  $k$  processors are idle. Formally,  $\text{Idle}_k(J, m)$  is the earliest instant such as at least  $k$  processors may be idle when the job set  $J$  is scheduled upon  $m$  processors, for any work-conserving scheduler. Unlike  $\text{Idle}_k$ , this value is the exact instant. We also introduce the following new notations:  $\text{Idle}_{\max}(S)$  is the maximal length of the period during which a processor is idle when a Mode Change Phase of a system  $S$  is handled by ACCEPTOR. The required time to handle the Mode Change Phase of  $S$  with an optimal protocol (resp. ACCEPTOR) is denoted  $|\text{OPT}(S)|$  (resp. denoted  $|\text{US}(S)|$ ). The optimal protocol doesn't have the same limitation as ACCEPTOR. For example, it doesn't have to schedule the rem-jobs using a work-conserving scheduler. Finally, we define the notion of *squeezable* system, in the context of a Mode Change Phase.

**Squeezable** A system  $S$  is said to be *squeezable* if and only if  $S$  has only one required reconfiguration  $\theta_c$ —with a reconfiguration delay of  $\delta_c$ —to perform during the Mode Change Phase and a set of rem-jobs  $J$  on a cluster having  $m$  processors such that:

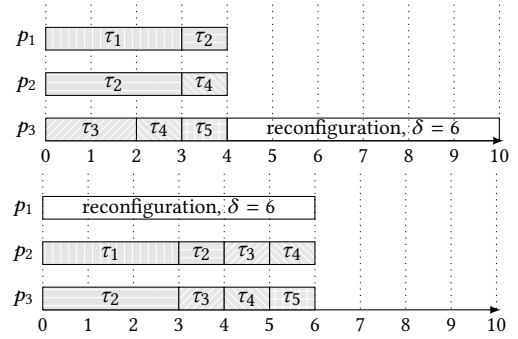
- i  $\sum C_i = \delta_c \times (m - 1)$ ,
- ii  $\text{Idle}_1(J, m) = \text{Idle}_2(J, m) = \dots = \text{Idle}_m(J, m) = \delta_c \times \frac{m-1}{m}$ ,
- iii  $\text{Idle}_1(J, m-1) = \text{Idle}_2(J, m-1) = \dots = \text{Idle}_{m-1}(J, m-1) = \delta_c$ ,
- iv  $\forall i, D_i \geq \delta_c$ .

This notions will be heavily used in the following lemmas, corollaries and theorems.

**LEMMA 7.1.** *Any mode change of any squeezable system handled by the protocol ACCEPTOR will have the reconfiguration entirely done alone, i.e. while all the processors except one are idle.*

**PROOF.** ACCEPTOR first schedules the tasks and launches a reconfiguration when a processor is idle. All the processors are idle at the same time because of the property ii) of *squeezable* systems, thus the reconfiguration will be done while all the processors except one are idle. This property is illustrated in Fig. 6.  $\square$

**LEMMA 7.2.** *During any Mode Change Phase of any squeezable system  $S$  handled by ACCEPTOR, if the reconfiguration is done while all the processors except one are idle then at least a processor will be idle for  $\delta_c$  time-units. Formally,  $\text{Idle}_{\max}(S) = \delta_c$ .*



**Figure 6: A squeezable system with  $m = 3$  and  $n = 5$**

**PROOF.** The reconfiguration is done while all the other processors are idle. So, those processors will be idle for  $\delta_c$  time-units. The lemma follows.  $\square$

**LEMMA 7.3.** *Any mode change of any squeezable system  $S$  may be handled by an optimal protocol in  $\delta_c$  time-units. Formally:  $|\text{OPT}(S)| = \delta_c$ .*

**PROOF.** It is obvious to see that it is possible to schedule all the rem-jobs on  $m-1$  processors, and to reconfigure a processor at  $t = 0$ . The rem-jobs will be complete at  $t = \text{Idle}_{m-1}(J, m-1) = \delta_c$  and the reconfiguration will be done at  $t = \delta_c$ . Because of iv), the deadlines will be respected. The lemma follows. This lemma is illustrated in Fig. 6.  $\square$

**LEMMA 7.4.** *Any mode change of any system  $S$  having the required reconfigurations  $\tilde{\Theta}$  and a set of rem-jobs  $J$  on a cluster having  $m$  processors takes more than  $C$  time-units when handled by an optimal protocol Optimal, where  $C$  is the maximum of the longest rem-jobs and the longest reconfiguration delay. Formally:  $|\text{OPT}(S)| \geq C$ , with  $C \doteq \max\{\max_{i=1}^n \{C_i\}, \max_{\theta_c \in \tilde{\Theta}} \{\delta_c\}\}$ .*

**PROOF.** No intra-parallelism is allowed, so the system  $S$  may not last less than the longest element to schedule, whether it is a job or a reconfiguration.  $\square$

**LEMMA 7.5.** *In any mode change handled by ACCEPTOR of any system  $S$  having the required reconfigurations  $\tilde{\Theta}$  and a set of rem-jobs  $J$  on a cluster having  $m$  processors, a processor cannot be idle more than  $C$  time-units where  $C \doteq \max\{\max_{i=1}^n \{C_i\}, \max_{\theta_c \in \tilde{\Theta}} \{\delta_c\}\}$ . Formally,  $\text{Idle}_{\max}(S) \leq C$ .*

**PROOF.** ACCEPTOR uses a work-conserving scheduler. Thus, no processor may be idle if a job or reconfiguration is waiting to execute. It immediately follows that a processor is idle only if no job or reconfiguration are waiting to execute. Thus, a processor will be idle until all the elements being executed are complete whether they are jobs and/or reconfigurations. Those elements may not be longest than  $C$  time-units. The lemma follows.  $\square$

**COROLLARY 7.6.** *For any mode change handled by ACCEPTOR of any system  $S$  having the required reconfigurations  $\tilde{\Theta}$  and a set of rem-jobs  $J$  on a cluster having  $m$  processors, the ratio  $\frac{\text{Idle}_{\max}(S)}{|\text{OPT}(S)|} \leq 1$ .*

PROOF. We have the two following inequations:  $|OPT(S)| \geq C$  (see Lemma 7.4) and  $Idle_{\max}(S) \leq C$  (see Lemma 7.5). It follows that the maximal value of  $\frac{Idle_{\max}(S)}{|OPT(S)|}$  is reached when  $Idle_{\max}(S) = |OPT(S)| = C$ , with  $C \doteq \max\{\max_{i=1}^n \{C_i\}, \max_{\theta_c \in \tilde{\Theta}} \{\delta_c\}\}$ . The maximal ratio is hence 1.  $\square$

LEMMA 7.7. *Any mode change of any squeezable system  $S$  leads to a ratio  $\frac{Idle_{\max}(S)}{|OPT(S)|} = 1$ , which is the maximum value.*

PROOF. The property directly follows from Lemmas 7.2–7.3 and Corollary 7.6.  $\square$

LEMMA 7.8. *When handling any mode change of any squeezable system  $S$ , the ratio  $\frac{|US(S)|}{|OPT(S)|} = \frac{Idle_{\max}(S)}{|OPT(S)|} \times \frac{m-1}{m} + 1$ .*

PROOF. To prove this lemma, we first show that

$$|US(S)| = |OPT(S)| + Idle_{\max}(S) \times \frac{m-1}{m}$$

From Lemma 7.2–7.3, we know that  $|OPT(S)| = \delta_c$  and  $Idle_{\max}(S) = \delta_c$ . We also know that ACCEPTOR will schedule first the rem-jobs for  $\delta_c \times \frac{m-1}{m}$  time-units (ii), and then do the reconfiguration for  $\delta_c$  time-units. Thus,  $|US(S)| = |OPT(S)| + Idle_{\max}(S) \times \frac{m-1}{m} = \delta_c + \delta_c \times \frac{m-1}{m}$ . This implies that

$$\frac{|US(S)|}{|OPT(S)|} = \frac{|OPT(S)| + Idle_{\max}(S) \times \frac{m-1}{m}}{|OPT(S)|} \quad (3)$$

$$= 1 + \frac{m-1}{m} \times \frac{Idle_{\max}(S)}{|OPT(S)|} \quad (4)$$

The lemma follows.  $\square$

COROLLARY 7.9. *When handling any mode change of any squeezable system  $S$  the ratio  $\frac{|US(S)|}{|OPT(S)|}$  is maximised. Thus, any squeezable system is a worst-case of ACCEPTOR.*

PROOF. Lemma 7.7 states that this corollary's hypotheses leads to a maximal ratio  $\frac{Idle_{\max}(S)}{|OPT(S)|}$  and Lemma 7.8 shows that maximising the ratio  $\frac{Idle_{\max}(S)}{|OPT(S)|}$  is equivalent to maximising the ratio  $\frac{|US(S)|}{|OPT(S)|}$ . Thus, this corollary follows.  $\square$

THEOREM 7.10. *For any system composed of  $m$  processors, the protocol ACCEPTOR is  $\frac{2m-1}{m}$ -competitive.*

PROOF. To prove that ACCEPTOR is  $\frac{2m-1}{m}$ -competitive, we have to prove that for any system  $S$ , the ratio  $\frac{|US(S)|}{|OPT(S)|} \leq \frac{2m-1}{m}$ . The Corollary 7.9 states that the maximal ratio will be obtained on any squeezable system  $S$ . On such a system, the ratio

$$\frac{|US(S)|}{|OPT(S)|} = \frac{Idle_{\max}(S)}{|OPT(S)|} \times \frac{m-1}{m} + 1 \quad (\text{Lemma 7.8}). \quad (5)$$

$$= \frac{m-1}{m} + 1 \quad (\text{Lemma 7.7}). \quad (6)$$

Thus, ACCEPTOR takes  $\frac{2m-1}{m}$  of the required time than the optimal protocol. The theorem follows.

Fig. 6 shows an example of such a system with  $m = 3$ ,  $n = 5$  and  $\delta_c = 6$ .  $\square$

COROLLARY 7.11. *For any system composed of  $m$  processors, the protocol ACCEPTOR is 2-competitive. This upper-bound is tight.*

PROOF. Theorem 7.10 states that ACCEPTOR is  $\frac{2m-1}{m}$ -competitive. It stands that

$$\forall m > 0, \frac{2m-1}{m} < 2$$

Thus, ACCEPTOR is 2-competitive. This upper-bound is tight, because of the following limit:

$$\lim_{m \rightarrow \infty} \frac{2m-1}{m} = 2 \quad \square$$

## 7.4 Handling mode independent tasks

In our model, all tasks are mode dependent. Real-time applications may have mode independent tasks: i.e. tasks that run during the whole lifespan of the application such as the OS. Our model 'as is' does not permit such tasks to run. However, a trivial extension with no effect on the validity nor the complexity would be to restrain those tasks to a specific set of processors with no reconfiguration allowed. Still, mode dependent tasks could run on such a cluster and because of Lemma 6.1, their rem-jobs would be correctly scheduled.

Doing so removes the limitation in a very easy way and makes the model usable for real-world applications.

## 8 CONCLUSIONS AND FUTURE WORK

In this research we propose a new model combining *dynamic hardware with software reconfigurability*. This is the *first* model for multi-mode real-time applications on reconfigurable heterogeneous platforms. We also propose a protocol to handle the Mode Change Phase of such application and prove its validity. An upper-bound on the Mode Change Phase is also introduced. Then, a complete evaluation of the contributions is done by evaluating the usability of the protocol (by evaluating its worst-case time complexity), its performance (by evaluating the competitiveness of the protocol, and the empirical pessimism of the bound) and discussing the constraint of the model on mode-independent tasks and showing that this is not a limitation.

Several extensions to this work can be done. First, the protocol could be improved in terms of competitiveness. We could improve the model to have a dynamic number of processors at run-time, which is allowed by FPGA reconfiguration. We could also try to introduce in a more advance way mode-independent tasks. Finally, other kind of possible future work is to put this model and protocol (and future extensions) into practice by implementing those in a Real-Time Operating System supporting a heterogeneous reconfigurable platform such as the Zynq UltraScale+™. Such experiments would validate the approach of this theoretical work in a close-to-industry, real-world scenario.

## ACKNOWLEDGEMENT

This work is supported in part by the H2020-ICT-2015 Innovation Action 688403 TULIPP and by the *Fédération Wallonie-Bruxelles*, Concerted Research Action named SOFIST. Lastly, the authors would like to thank Dragomir Mилоjević for his helpful collaboration and Juan M. Rivas for his careful review of our manuscript and his insightful comments.

## REFERENCES

- [1] Enrico Bini. 2016. Adaptive Fair Scheduler: Fairness in Presence of Disturbances. In *Proceedings of the 24th International Conference on Real-Time Networks and Systems, RTNS 2016, Brest, France, October 19-21, 2016*, Alain Plantec, Frank Singhoff, Sébastien Faucou, and Luís Miguel Pinho (Eds.). ACM, 129–138. <https://doi.org/10.1145/2997465.2997468>
- [2] Alessandro Biondi, Alessio Balsini, Marco Pagani, Enrico Rossi, Mauro Marinoni, and Giorgio C. Buttazzo. 2016. A Framework for Supporting Real-Time Applications on Dynamic Reconfigurable FPGAs. In *2016 IEEE Real-Time Systems Symposium, RTSS 2016, Porto, Portugal, November 29 - December 2, 2016*. IEEE Computer Society, 1–12. <https://doi.org/10.1109/RTSS.2016.010>
- [3] A. Biondi and G. Buttazzo. 2017. Timing-aware FPGA partitioning for real-time applications under dynamic partial reconfiguration. In *2017 NASA/ESA Conference on Adaptive Hardware and Systems (AHS)*. 172–179. <https://doi.org/10.1109/AHS.2017.8046375>
- [4] Martin Cornil, Antonio Paolillo, Joël Goossens, and Ben Rodriguez. 2017. Research and implementation challenges of RTOS support for heterogeneous computing platforms. In *Heterogeneous Architectures and Real-Time Systems Seminar, Brussels*.
- [5] Liliana Cucu-Grosjean and Joël Goossens. 2010. Predictability of Fixed-Job Priority schedulers on heterogeneous multiprocessor real-time systems. *Inf. Process. Lett.* 110, 10 (2010), 399–402. <https://doi.org/10.1016/j.ipl.2010.03.009>
- [6] Paul Emberson and Iain Bate. 2007. Minimising task migration and priority changes in mode transitions. In *13th IEEE Real Time and Embedded Technology and Applications Symposium*. IEEE, 158–167.
- [7] S. Funk, J. Goossens, and S. Baruah. 2001. On-line scheduling on uniform multiprocessors. In *Proceedings 22nd IEEE Real-Time Systems Symposium (RTSS 2001) (Cat. No.01PR1420)*. 183–192. <https://doi.org/10.1109/REAL.2001.990609>
- [8] Joël Goossens and Xavier Poczekajlo. 2017. Multimode application on a reconfigurable platform: Introducing a new model and a first protocol. In *11th Junior Researcher Workshop on Real-Time Computing*. 1–4.
- [9] Joël Goossens and Pascal Richard. 2013. Partitioned scheduling of multimode multiprocessor real-time systems with temporal isolation. In *Proceedings of the 21st International Conference on Real-Time Networks and Systems*. ACM, 297–305.
- [10] Rhan Ha and Jane W.-S. Liu. 1994. Validating Timing Constraints in Multiprocessor and Distributed Real-Time Systems. In *Proceedings of the 14th International Conference on Distributed Computing Systems, Poznan, Poland, June 21-24, 1994*. IEEE Computer Society, 162–171. <https://doi.org/10.1109/ICDCS.1994.302407>
- [11] Laura Hopperton. 2011. embedded world: Xilinx introduces 'industry's first' extensible processing platform. (2011). <https://goo.gl/EuNhir>
- [12] José Marinho, Gurulingesh Raravi, Vincent Nélis, and Stefan M Petters. 2011. Partitioned Scheduling of Multimode Systems on Multiprocessor Platforms: when to do the Mode Transition? *RTSOPS* (2011).
- [13] Vincent Nélis. 2010. *Energy-Aware Real-Time Scheduling in Embedded Multiprocessor Systems*. Ph.D. Dissertation. Université libre de Bruxelles.
- [14] Vincent Nélis, Björn Andersson, José Marinho, and Stefan M. Petters. 2011. Global-EDF Scheduling of Multimode Real-Time Systems Considering Mode Independent Tasks. In *23rd Euromicro Conference on Real-Time Systems, ECRTS 2011, Porto, Portugal, 5-8 July, 2011*. 205–214. <https://doi.org/10.1109/ECRTS.2011.27>
- [15] Vincent Nélis, Joël Goossens, and Björn Andersson. 2009. Two Protocols for Scheduling Multi-mode Real-Time Systems upon Identical Multiprocessor Platforms. In *Euromicro Conference on Real-Time Systems*. 151–160.
- [16] M. Pagani, A. Balsini, A. Biondi, M. Marinoni, and G. Buttazzo. 2017. A Linux-based support for developing real-time applications on heterogeneous platforms with dynamic FPGA reconfiguration. In *2017 30th IEEE International System-on-Chip Conference (SOCC)*. 96–101. <https://doi.org/10.1109/SOCC.2017.8226015>
- [17] Marco Pagani, Mauro Marinoni, Alessandro Biondi, Alessio Balsini, and Giorgio Buttazzo. 2016. Towards real-time operating systems for heterogeneous reconfigurable platforms. *OSPERS 2016* (2016), 49.
- [18] Gurulingesh Raravi, Björn Andersson, Vincent Nélis, and Konstantinos Bletsas. 2014. Task assignment algorithms for two-type heterogeneous multiprocessors. *Real-Time Systems* 50, 1 (2014), 87–141. <https://doi.org/10.1007/s11241-013-9191-3>
- [19] Jorge Real and Alfons Crespo. 2004. Mode Change Protocols for Real-Time Systems: A Survey and a New Proposal. *Real-Time Systems* 26, 2 (2004), 161–197. <https://doi.org/10.1023/B:TIME.0000016129.97430.c6>
- [20] Ahmad Sadek, Ananya Muddukrishna, Lester Kalms, Asbjørn Djupdal, Ariel Podlubne, Antonio Paolillo, Diana Goehring, and Magnus Jahre. 2018. Supporting Utilities for Heterogeneous Embedded Image Processing Platforms (STEM): An Overview. In *Applied Reconfigurable Computing, Architectures, Tools, and Applications*, Nikolaos Voros, Michael Huebner, Georgios Keramidas, Diana Goehring, Christos Antonopoulos, and Pedro C. Diniz (Eds.). Springer International Publishing, Cham, 737–749.
- [21] Chi-Sheng Shih and Chang-Min Yang. 2017. Schedulability Analysis of Mode Change for Imprecise Computation on Multi-Core Platforms. In *Proceedings of the International Conference on Research in Adaptive and Convergent Systems, RACS 2017, Krakow, Poland, September 20-23, 2017*. ACM, 261–268. <https://doi.org/10.1145/3129676.3129720>
- [22] Xilinx. 2018. Zynq UltraScale+ MPSoC. (2018). <https://www.xilinx.com/products/silicon-devices/soc/zynq-ultrascale-mpsoc.html>
- [23] Christian Kohn (Xilinx). 2013. Partial Reconfiguration of a Hardware Accelerator on Zynq-7000 All Programmable SoC Devices (XAPP1159). (2013).