

THE NATIONAL UNIVERSITY
of SINGAPORE



School of Computing
Lower Kent Ridge Road, Singapore 119260

TRC8/04

*Processor Frequency Selection in Energy-Aware
SoC Platform Design for Multimedia Application*

*Yahong LIU, Alexander MAXIAGUINE,
Samarjit CHAKRABORTY and Wei Tsang OOI*

August 2004

Technical Report

Foreword

This technical report contains a research paper, development or tutorial article, which has been submitted for publication in a journal or for consideration by the commissioning organization. The report represents the ideas of its author, and should not be taken as the official views of the School or the University. Any discussion of the content of the report should be sent to the author, at the address shown on the cover.

JAFFAR, Joxan
Dean of School

Processor Frequency Selection in Energy-Aware SoC Platform Design for Multimedia Applications

Yanhong Liu* Alexander Maxiaguine[†] Samarjit Chakraborty* Wei Tsang Ooi*

Abstract

Of late, there has been a considerable interest in generic and configurable System-on-Chip platforms specifically targeted towards implementing multimedia applications. A number of such platforms offer the possibility of including processor soft cores which are highly customizable. For voltage/frequency scaled processors, such customization includes the selection of appropriate voltage/frequency operating points which are tuned to the application set to be mapped onto the platform. In this context, we present an analytical framework that can guide a system designer in identifying the frequency ranges that should be supported by the different processors of a platform architecture. This framework can also be used to identify how such frequency ranges depend on the different parameters of the architecture (such as on-chip buffer sizes), and the performance impacts associated with selecting a particular frequency range. In the case of multimedia streaming applications, identifying such performance impacts and tradeoffs involved in customizing a platform architecture is especially difficult due to the bursty nature of on-chip traffic arising out of multimedia processing and the high variability in their execution requirements. The framework presented here is designed to precisely capture such characteristics and can be used in the design-space exploration of energy-aware platform architectures for multimedia processing.

1 Introduction

Today, most portable devices such as PDAs and mobile phones offer support for streaming multimedia applications. Because of flexibility, time-to-market advantages and low design costs, very often such devices are now designed using generic, configurable System-on-Chip (SoC) platforms. Examples of such generic platforms are PrimeXsys from ARM [17] and AcurX from Plamchip [2]. These platforms are targeted towards a wide range applications starting from DVD players and set-top boxes, to network routers and network security processors. However, the advantages of generic platform-based solutions come at the cost of a large performance disparity between such designs and fully customized ASIC-based solutions. Although generic platforms can be configured for specific applications, for example by integrating custom cores into them, they can rarely match the performance and power utilization characteristics of ASICs.

¹Department of Computer Science, National University of Singapore
E-mail: {liuyanho, samarjit, ooiwt}@comp.nus.edu.sg

²Computer Engineering and Networks Laboratory, Swiss Federal Institute of Technology (ETH) Zürich
E-mail: maxiagui@tik.ee.ethz.ch

Application-specific platforms narrow this gap to some extent, and are suited for a “product family”-oriented design approach. Such platforms are customized for a particular application domain, but still support sufficient flexibility to allow them to be configured for specific products belonging to that domain. An example of such a platform is OMAP from Texas Instruments [16], which allow multimedia capabilities to be included in 2.5G and 3G wireless handsets and PDAs. The Eclipse architecture template [18] and the Viper SoC architecture [10], from Philips, are also examples of such application-specific platforms which target advanced set-top box and DTV markets.

However, the design of such application-specific platforms involve delicate tradeoffs between flexibility, cost, performance and power consumption characteristics. For example, increasing flexibility is usually associated with increased cost and low performance. Because of the high computational demands, real-time constraints and low power consumption requirements of portable multimedia devices, *designing* SoC platforms for such devices require a disciplined design methodology. Although, recently there has been some work on *platform management* techniques [21, 22], the issue of *platform design* has not been sufficiently addressed so far. The only significant effort in this direction, that we are aware of, is from Philips in designing the Eclipse architecture template for media processing SoCs [20, 19, 24]. In this paper, we follow this line of work and propose a framework for designing energy-aware SoC platforms specifically targeted towards media processing in portable devices. In the application domain targeted by Eclipse (digital televisions and set-top boxes), power consumption is not a significant issue.

Problem Statement: The problems we are interested in addressing are of the following form. Suppose that we are given a multiprocessor SoC platform architecture “template” and a number of multimedia applications, all of which are required to be supported by this platform. Our job is to derive a (concrete) platform architecture from this template, by choosing appropriate processors, sizes of on-chip buffers and possibly other parameters such as bus widths and cache configurations. The processors to be chosen for this platform support software-controlled voltage and frequency scaling to allow different degrees of power consumption at run time. Therefore, we are also required to choose the frequency/voltage ranges that each processor should support. In this paper we specifically focus on this last issue and identify how this range depends on the other parameters of the platform architecture, such as on-chip buffer sizes.

The results presented in this paper also provide insights into questions such as: if a processor supports only a fixed number of *operating points*, where each such point is characterized by a voltage and a frequency value, then how many such operating points should a processor ideally support and how should these values be chosen? A processor which allows the voltage and frequency values to be changed continuously would typically be more expensive than one which allows these values to be changed in discrete steps and supports only a fixed number of these values or operating points. Today, processors of both these types are available—Intel’s XScale processor is of the former type and Transmeta’s Crusoe processor is of the latter type. Therefore, it is pertinent to ask questions like what kind of performance impacts would choosing a processor of the latter type have, over a more expensive processor which supports a continuous range of frequency values? Further, a platform designer would also be interested in identifying how the frequency range, that needs to be supported by a processor, varies with the available on-chip buffer size. Since on-chip buffers are available only

at a premium because of their high area requirements [25], such information would help in choosing an appropriate tradeoff.

Answers to these questions are becoming increasingly relevant because a number of embedded processor cores today offer a high degree of customization potential, such as instruction set tailoring and register file sizing. For processors which support dynamic frequency/voltage scaling, choosing efficient operating points is therefore becoming a part of this customization procedure. In the context of multimedia applications, this is especially critical because of the complex and bursty nature of on-chip traffic and the high variability in the execution times of multimedia processing tasks—both of these resulting in a highly variable demand on the computational resources available on the chip [25]. Hence, being able to control the processor frequency accurately to counter this variability is important.

Our Results and Relation to Previous Work: The main contribution of this paper is framework which can guide a system designer in identifying the operating frequency range that different processors on a SoC platform architecture should support in order to run a given multimedia application or a class of applications. Identifying such a range accurately is not straightforward because of the reasons mentioned above, i.e. the complex nature of on-chip traffic arising out of multimedia processing and the variability in the execution times of tasks. Moreover, since different applications and input classes might have very different computational demands choosing an appropriate processor frequency range involves several tradeoffs between processor cost, flexibility and on-chip buffer requirements. Our framework can help a system designer in identifying these tradeoffs.

Although there has been a significant amount of work in developing voltage and frequency scheduling algorithms in the context of multimedia applications (see, for example, [1, 6, 9, 12, 15, 26], and also [3] and the references therein) the problem of processor design and processor frequency selection from an energy-aware perspective has received considerably little attention so far. Our work is partly motivated by a recent paper [5], which proposes a linear programming based technique to optimally select operating voltage/frequency points in an embedded processor core. The task to be executed on such a processor is specified as a task graph whose vertices are annotated with execution requirements and deadlines. However, in contrast to this work, the framework proposed here is targeted specifically towards streaming multimedia applications and explicitly models the burstiness in multimedia streams and the variability in the execution requirements of multimedia processing tasks. Other representative work in this direction is [11], which addresses the selection of the processor core and instruction and data cache configuration in the design of variable voltage processors.

The framework presented in this paper is based on the theory of *Network Calculus* [4] which was developed and is still largely used in the context of analyzing communication networks. Very recently, it was extended to analyze SoC architectures in the context of network processors [8, 23]. This work was further extended in [7, 14] to the domain of general SoC platform architectures. Our work in this paper follows this line of development, but extends the underlying theory in two ways. On an abstract level, the analysis schemes presented in the previous papers relied on concrete input instances and could not provide any theoretical guarantees on the performance of an architecture for a *class* of inputs. In contrast to this, the framework presented here can be used to analyze a *class* of input streams, for which a more elaborate theory is necessary—this is explained in detail in Sections 3.1 - 3.3. Secondly, none of the previous results provided means for computing the range of processor

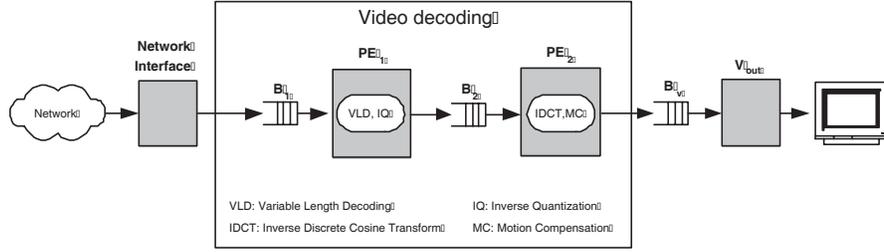


Figure 1: A multiprocessor SoC platform onto which an MPEG-2 decoder application is partitioned and mapped.

frequencies from an input specification. This extension is presented in Section 4.

The rest of the paper is organized as follows. The next section describes a system model of a platform architecture and formally states our problem. Given a specification of the application to be implemented on this architecture and the class of input streams to be processed, in Section 3 we compute bounds on the *service* that needs to be provided by each processor of this architecture. In Section 4 we show how such service bounds can be used to derive the operating frequency range of each processor. Finally, in Section 5 we present a case study involving an MPEG-2 decoder application to illustrate an application of the proposed framework, and also validate the results obtained using detailed simulations.

2 Problem Formulation

In this paper, we consider the following system-level view of multimedia stream processing on a SoC platform. A platform architecture, such as the one shown in Figure 1, consists of multiple processing elements (PEs) onto which the different parts of an application are partitioned and mapped. An input multimedia stream enters a PE, gets processed by the task(s) implemented on this PE, and the processed stream enters another PE for further processing. A model of the architecture in Figure 1 is shown in Figure 2. Each PE has an internal buffer, which is a FIFO channel of fixed capacity, and is used to store the incoming stream to be processed. Finally, the fully processed stream is written into a *playout buffer* which is read by some *real-time client* (RTC) such as an audio or a video output device.

For the sake of generality, we consider any multimedia stream that is seen by a PE belonging to the platform, to be made up of a potentially infinite sequence of *stream objects*. A stream object might be a bit belonging to a compressed bitstream representing a coded video clip, or a macroblock, or a video frame, or an audio sample—depending on where in the architecture the stream exists. For example, in the architecture shown in Figure 1, stream objects entering PE_1 are single bits. But stream objects entering PE_2 are coded macroblocks and the stream objects being written into the playout buffer B_i are decoded macroblocks.

In Figure 2, let $x_i(t)$ denote the number of stream objects that arrive at a PE PE_i during the time interval $[0, t]$. Let $y_i(t)$ (equal to $x_{i+1}(t)$) denote the number of processed stream objects at the output of PE_i (or the input of PE_{i+1}) during the time interval $[0, t]$. The real-time client *RTC* consumes stream objects from the playout buffer at a rate $C(t)$, which again denotes the number of stream objects consumed within the time interval $[0, t]$.

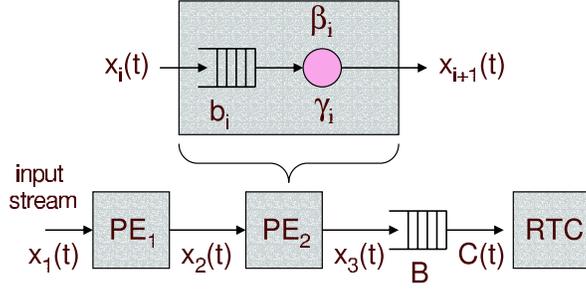


Figure 2: System-level view of multimedia processing on a multiprocessor SoC platform.

The stream entering the processing element PE_i receives a *service* β_i , which is specified by a tuple (β_i^l, β_i^u) . Within *any* time interval of length Δ , it is guaranteed that PE_i will process *at least* $\beta_i^l(\Delta)$ number of stream objects and it will be able to process *at most* $\beta_i^u(\Delta)$ number of stream objects. The functions β_i^l and β_i^u therefore represent lower and upper bounds on the service provided by PE_i and is determined by the time required to process each stream object, the scheduling policy implemented on this PE (in case multiple streams are being processed by it), and also by the voltage/frequency scheduling policy implemented on it. Lastly, each PE PE_i is also associated with a tuple $\gamma_i = (\gamma_i^l, \gamma_i^u)$, where $\gamma_i^l(k)$ and $\gamma_i^u(k)$ denote the minimum and the maximum number of *processor cycles* respectively, that may be required to process *any* k consecutive stream objects belonging to the input stream. γ_i is therefore used to capture the variability in the execution requirements of the different stream objects.

Now recall from Section 1 that for each PE belonging to the platform, we would like to determine the operating frequency range that should be supported by it. If the processor supports only a fixed number of discrete frequency levels, then we would like to determine how should these frequencies be chosen and what kind of performance impacts will this decision have. Note that the platform should be designed to support a *class* (or several classes) of multimedia streams. For example, a portable multimedia device might have a wireless interface through which MPEG-2 coded video streams of two different classes come in—high-quality video clips with 8 Mbps input bit rate and low-quality clips with 4 Mbps input bit rate. The computational demands associated with these two input classes might vary widely, which translates to different operating frequency requirements for any PE on the platform. The input to a PE, when specified using the function $x_i(t)$, however, represents a concrete instance of a stream rather than a *class* of streams. Therefore, to specify the arrival pattern of a class or family of streams, we use an abstraction called *arrival curve* which is similar to the concept of *service* β_i described above. The arrival curve α_{x_i} representing the class of streams that might arrive at the input of PE_i is also specified by a tuple $(\alpha_{x_i}^l(\Delta), \alpha_{x_i}^u(\Delta))$, where the first and the second terms represent the minimum and the maximum number of stream objects that might arrive within any time interval of length Δ . In other words, $\alpha_{x_i}^l(\Delta) \leq x_i(t + \Delta) - x(t) \leq \alpha_{x_i}^u(\Delta)$, $\forall t, \Delta \geq 0$. Therefore, any concrete arrival pattern $x_i(t)$ is lower and upper bounded by the functions $\alpha_{x_i}^l$ and $\alpha_{x_i}^u$ respectively. Similarly, we use $\alpha_{y_i}^l$ and $\alpha_{y_i}^u$ to denote lower and upper bounds on the arrival pattern of the processed stream at the output of PE_i .

Now, let us consider the last PE in the path of a stream, i.e. the PE whose output is written into the playout buffer (see Figure 2). Henceforth, for simplifying the notation, we drop the subscript i

representing the PE identifier. Therefore, as described above, any input instance to this PE is specified by the function $x(t)$ and the class of all input instances is bounded by the arrival curve α_x . Any output arrival pattern from this PE is represented by the function $y(t)$ and the sizes of the internal and the playout buffers are b and B respectively. The consumption pattern of stream objects from the playout buffer is specified by the function $C(t)$ described above. Now, given $\alpha_x(\Delta)$, $\gamma(k)$, $C(t)$ and the buffer sizes b and B , the problem is to compute the set of all possible processor frequencies at which this PE might be run, such that the following constraints are satisfied: (i) the playout buffer never overflows, (ii) it never underflows, and (iii) the internal buffer never overflows.

The constraint on the playout buffer underflow is to ascertain that stream objects can be read out by the output audio/video devices (the RTC in Figure 2) at the rate specified by $C(t)$, and hence the output quality is guaranteed. The constraints on buffer overflow are motivated by the fact that typically on-chip PEs use static voltage and task scheduling policies. This is because using blocking write/read mechanisms efficiently to prevent buffer overflows/underflows either require a multithreaded processor architecture or substantial run-time operating system support for context switching.

Our solution to the above problem consists of two parts. In Section 3 we compute lower and upper bounds on the service (β^l and β^u) that needs to be provided by the PE in order to satisfy the above mentioned buffer constraints. In Section 4, we then show how to compute the frequency range that needs to be supported by the PE in order to realize these service bounds. The extension of these results to any other PE in the path of a stream (i.e. one whose output is not written into the playout buffer, but instead into another PE) is fairly simple, and is also explained in Section 3.4.

Throughout this paper we assume the following processor model: a PE can either support a continuous range of clock frequencies, or a fixed number of discrete frequencies, where this number can also be equal to one—i.e. the PE runs at a fixed frequency and does not support dynamic frequency scaling. Any clock frequency is associated with a minimum operating voltage that needs to be supplied to run the processor at this frequency. We assume that this is the voltage at which the processor is run for any frequency—i.e. voltage and frequency are tightly coupled and determining the frequency results in the voltage also being determined. Hence, we will only be concerned with determining the frequency range or the discrete frequency values for any given PE.

3 Computing Bounds on Service Requirements

Given $\alpha_x(\Delta)$, $C(t)$ and the buffer sizes b and B for the last PE in the path of a stream, in this section we compute the lower and the upper bounds $\beta^l(\Delta)$ and $\beta^u(\Delta)$ on the service that needs to be provided to this stream to satisfy the buffer overflow and underflow constraints described in Section 3. Within any interval of length Δ , if the service provided is less than our computed $\beta^l(\Delta)$, then either the internal buffer might overflow or the playout buffer might underflow. Similarly, if the service provided is greater than the computed $\beta^u(\Delta)$, then the playout buffer might overflow.

Following the notation introduced in Section 3, we use $x(t)$ to denote any arrival pattern of stream objects at the input of the PE and $y(t)$ to denote the arrival pattern at the output of the PE. Recall that the functions x , y and C always denote *cumulative values* over the time interval $[0, t]$, whereas the functions α_x and β take *time interval lengths* as the input parameter.

Notation: For any two functions f and g , the *min-plus convolution* of f and g is given by: $(f \otimes g)(t) = \inf_{s:0 \leq s \leq t} \{f(t-s) + g(s)\}$. The *min-plus deconvolution* of f and g is given by: $(f \oslash g)(t) = \sup_{u:u \geq 0} \{f(t+u) - g(u)\}$. We use $f \wedge g$ to denote the infimum of f and g , or the minimum if it exists, and $f \vee g$ to denote the supremum of f and g , or the maximum if it exists.

For reasons which we explain later, let us assume that at time $t = 0$, the buffer fill-levels of the playout buffer and the internal buffer of the last PE in the path of the stream are B_0 and b_0 respectively. Then the constraint on the playout buffer underflow can be stated as (see Figure 2):

$$y(t) \geq C(t) - B_0, \quad \forall t \geq 0 \quad (1)$$

Similarly, the constraint on the playout buffer overflow can be stated as:

$$y(t) \leq C(t) + B - B_0, \quad \forall t \geq 0 \quad (2)$$

Finally, the constraint that the internal buffer in the PE should not overflow, is given by:

$$y(t) \geq x(t) - (b - b_0), \quad \forall t \geq 0 \quad (3)$$

The constraints (1) and (3) can be combined and stated as: $y(t) \geq (C(t) - B_0) \vee (x(t) - (b - b_0))$, $\forall t \geq 0$. Now, if $\beta^l(\Delta)$ is the minimum number of stream objects that is guaranteed to be processed by this PE within any time interval of length Δ , then it can be shown that $y(t) \geq (x \otimes \beta^l)(t)$, $\forall t \geq 0$ [4]. Hence, $(x \otimes \beta^l)(t)$ is the minimum value of $y(t)$ for any t , and therefore the above constraint on $y(t)$ can be reformulated as:

$$(x \otimes \beta^l)(t) \geq (C(t) - B_0) \vee (x(t) - (b - b_0)), \quad \forall t \geq 0 \quad (4)$$

It can be shown [4] that for any functions f, g and h , $g \otimes h \geq f$ if and only if $h \geq f \oslash g$. Using this result, inequality (4) can be reformulated as:

$$\beta^l(t) \geq ((C(t) - B_0) \vee (x(t) - (b - b_0))) \oslash x(t), \quad \forall t \geq 0 \quad (5)$$

Inequality (5) therefore gives a lower bound on the service that needs to be provided by the PE in order to satisfy the playout buffer underflow and the internal buffer overflow constraints.

If $\beta^u(\Delta)$ is the maximum number of stream objects that can be processed by the PE within any time interval of length Δ , then it can be shown that $y(t) \leq (x \otimes \beta^u)(t)$, $\forall t \geq 0$ [4]. Hence, following the same reasoning as above, $(x \otimes \beta^u)(t)$ is the maximum value of $y(t)$ for any t , and by using this, the constraint (2) can be reformulated as:

$$(x \otimes \beta^u)(t) \leq C(t) + B - B_0, \quad \forall t \geq 0 \quad (6)$$

Given the functions $x(t)$, $C(t)$ and the values of B and B_0 , we would like to determine maximum value of $\beta^u(t)$ for any t (or the largest possible function β^u), for which inequality (6) is satisfied. Such a β^u would give an upper bound on the service that can be provided by the PE, which will satisfy the playout buffer overflow constraint.

It may be noted here that in some cases $\beta^u(t)$ can be infinitely large. For example, consider the case where the arrival pattern of stream objects (given by $x(t)$) is exactly the same as the consumption

pattern by the real-time client from the playout buffer, i.e. $x(t) = C(t)$ for all $t \geq 0$. In this case $\beta^u(t)$ can be infinitely large, since no matter how much service is provided by the PE, the playout buffer can never overflow. In fact, $\beta^u(t)$ can be infinitely large if $x(t) \leq C(t) + B - B_0$ for all $t \geq 0$ (this also follows directly from the definition of the min-plus convolution operator).

3.1 Computing Service Bounds for a *Class* of Streams

The above bounds on β^l and β^u are based on a specific instance of the arrival pattern of a stream, i.e. $x(t)$. Hence, these bounds can only guarantee the buffer overflow and underflow constraints for this specific arrival pattern. However, we would like to derive the service bounds for a class of arrival patterns—i.e. all arrival patterns which are bounded by the arrival curve α_x .

3.1.1 Computing the Bound on β^l

For a concrete arrival pattern of stream objects given by $x(t)$, the bound on β^l is given by (from inequality (5):

$$\begin{aligned}
\beta^l(t) &\geq ((C(t) - B_0) \vee (x(t) - (b - b_0))) \otimes x(t) \\
&= \sup_{u \geq 0} \{((C(t+u) - B_0) \vee (x(t+u) - (b - b_0))) - x(u)\} \\
&= \sup_{u \geq 0} \{(C(t+u) - B_0 - x(u)) \vee (x(t+u) - x(u) - (b - b_0))\} \\
&= \max\{\sup_{u \geq 0} (C(t+u) - x(u)) - B_0, \sup_{u \geq 0} (x(t+u) - x(u) - (b - b_0))\} \\
&= \max\{(C \otimes x)(t) - B_0, \alpha_x^u(t) - (b - b_0)\}, \quad \forall t \geq 0
\end{aligned} \tag{7}$$

Since $x(t) \geq \alpha_x^l(t)$ for all $t \geq 0$, for any function f , $(f \otimes x)(t) \leq (f \otimes \alpha_x^l)(t)$ for all $t \geq 0$. Hence, the constraint specified by inequality (7) can be reformulated as:

$$\beta^l(t) \geq ((C \otimes \alpha_x^l)(t) - B_0) \vee (\alpha_x^u(t) - (b - b_0)), \quad \forall t \geq 0$$

Further, let us assume that the consumption pattern of stream objects from the playout buffer, as specified by the function $C(t)$ is lower and upper bounded by the arrival curve α_C , i.e. $\alpha_C^l(\Delta) \leq C(t + \Delta) - C(t) \leq \alpha_C^u(\Delta)$, $\forall t, \Delta \geq 0$. Then the above constraint on β^l can be stated as:

$$\beta^l(t) \geq ((\alpha_C^u \otimes \alpha_x^l)(t) - B_0) \vee (\alpha_x^u(t) - (b - b_0)), \quad \forall t \geq 0 \tag{8}$$

Inequality (8) therefore provides a lower bound on the minimum service that needs to be provided by the PE, in order to satisfy the playout buffer underflow and the internal buffer overflow constraints, where all arrival patterns at the PE are bounded by α_x and all consumption patterns from the playout buffer are bounded by α_C . Once again, recall that α_x and α_C represent a *class* or a *family* of arrival and consumption patterns of stream objects, and not any specific instance of an arrival or a consumption pattern.

3.1.2 Computing the Bound on β^u

From inequality (6), we obtain that the upper bound on the maximum service that can be provided by the PE is given by the largest possible function β^u which satisfies $(x \otimes \beta^u)(t) \leq C(t) + B - B_0$, $\forall t \geq 0$. We know that any instance of an arrival pattern $x(t)$ at the input of the PE is upper bounded by α_x^u , and the lower bound on the consumption pattern of stream objects from the playout buffer is given by α_C^l . Hence, the above constraint on β^u can be reformulated as:

$$(\alpha_x^u \otimes \beta^u)(t) \leq \alpha_C^l(t) + B - B_0, \quad \forall t \geq 0 \quad (9)$$

Again, given the functions α_x^u , α_C^l and B , B_0 , we would like to determine the largest function β^u which satisfies inequality (9). Towards this, note that if $\alpha_x^u(t) \leq \alpha_C^l(t) + B - B_0$, $\forall t \geq 0$, then $\beta^u(t)$ can be infinitely large for all $t > 0$. This corresponds to the case where the rate at which stream objects arrive at the PE is not high enough to overflow the playout buffer, irrespective of how fast they are processed.

Now, let us consider the case where, $\alpha_x^u(t_i) > \alpha_C^l(t_i) + B - B_0$, for $i = 1, \dots, n$, and for all other values of t , $\alpha_x^u(t) \leq \alpha_C^l(t) + B - B_0$. Assuming that $t_i < t_{i+1}$ for $i = 1, \dots, n-1$, in this case the largest function β^u which satisfies inequality (9) is given as follows:

$$\beta^u(t) \leq \begin{cases} 0 & \text{if } t = 0, \\ \alpha_C^l(t_1) + B - B_0 & \forall t \leq t_1 \\ \alpha_C^l(t_{i+1}) + B - B_0 & \forall t_i < t \leq t_{i+1}, i = 1, \dots, n-1 \\ \infty & \forall t > t_n \end{cases} \quad (10)$$

Proof sketch: Suppose $t < t_1$. Then $(\alpha_x^u \otimes \beta^u)(t)$ is equal to $\inf_{0 \leq s \leq t} \{\alpha_x^u(s) + \beta^u(t-s)\}$, which is less than or equal to $(\alpha_x^u(t) + \beta^u(0))$. Since, $\alpha_x^u(t)$ is less than or equal to $\alpha_C^l(t) + B - B_0$ for all $t < t_1$ and $\beta^u(0)$ is equal to 0, it follows that $(\alpha_x^u \otimes \beta^u)(t)$ is also less than or equal to $\alpha_C^l(t) + B - B_0$ for all $t < t_1$. For $t = t_1$, $(\alpha_x^u \otimes \beta^u)(t) \leq \alpha_x^u(0) + \beta^u(t)$. Since, $\alpha_x^u(0) = 0$ and $\beta^u(t_1) = \alpha_C^l(t_1) + B - B_0$, it follows that $(\alpha_x^u \otimes \beta^u)(t_1)$ is less than or equal to $\alpha_C^l(t_1) + B - B_0$. Similar arguments hold for all other values of t . \square

The above upper bound on β^u therefore guarantees that the playout buffer never overflows when the arrival pattern of stream objects at the PE is bounded by α_x and the consumption pattern of stream objects from the playout buffer is bounded by α_C .

3.2 Computing Service Bounds in Terms of Number of Processor Cycles

The lower and the upper bounds on the service that needs to be guaranteed by a PE, as given by inequalities (8) and (10), are specified in terms of the minimum and the maximum number of stream objects that need to be processed within any given time interval. However, due to the data-dependent variability in the execution times of multimedia tasks, the number of processor cycles required to completely process any stream object might be highly variable. As explained in Section 2, this variability can be captured by the function γ , which we refer to as the *workload curve*. $\gamma^l(k)$ denotes the minimum number of processor cycles required to process any k consecutive stream objects and $\gamma^u(k)$ denotes the maximum number of processor cycles that may be required to process any k consecutive

stream objects. Therefore, it follows from the last subsection that $\gamma^u(\beta^l(\Delta))$ is the minimum number of processor cycles that must be provided to a stream within any time interval of length Δ to guarantee that the playout buffer never underflows and the internal buffer at the PE never overflows. Similarly, $\gamma^l(\beta^u(\Delta))$ is the maximum number of processor cycles that may be provided to a stream within any time interval of length Δ to guarantee that the playout buffer never overflows.

Here we would like to point out that from our definition of the function β , it follows that $\beta^l(t) \geq \beta^l(s) + \beta^l(t-s)$ for all $t \geq 0$ and $0 \leq s \leq t$. Similarly, $\beta^u(t) \leq \beta^u(s) + \beta^u(t-s)$ for all $t \geq 0$ and $0 \leq s \leq t$. However, the bounds given by inequalities (8) and (10) need not satisfy these properties. Let us assume that inequality (8) is of the form $\beta^l(t) \geq f(t)$, $\forall t \geq 0$ and inequality (10) is of the form $\beta^u(t) \leq g(t)$, $\forall t \geq 0$, i.e. $f(t)$ is the right hand side term of inequality (8) and $g(t)$ is the right hand side term of inequality (10). Now let us define two functions σ^l and σ^u as follows:

$$\sigma^l(\Delta) = \begin{cases} 0 & \text{if } \Delta = 0 \\ \gamma^u(f(\Delta)) & \text{if } \Delta = 1 \\ \max\{\gamma^u(f(\Delta)), (\sigma^l \otimes \sigma^l)(\Delta)\} & \text{if } \Delta > 1 \end{cases} \quad (11)$$

$$\sigma^u(\Delta) = \begin{cases} 0 & \text{if } \Delta = 0 \\ \gamma^l(g(\Delta)) & \text{if } \Delta = 1 \\ \min\{\gamma^l(g(\Delta)), (\sigma^u \otimes \sigma^u)(\Delta)\} & \text{if } \Delta > 1 \end{cases} \quad (12)$$

The functions $\sigma^l(\Delta)$ and $\sigma^u(\Delta)$ are therefore defined over $\Delta = 0, 1, 2, \dots$, and denote the minimum and the maximum number of processor cycles that should be provided to a stream within *any* time interval of length Δ for all the buffer overflow and underflow constraints to be satisfied. Moreover, it can be shown that these two functions satisfy the properties that any function which bounds the service provided by a PE should satisfy, i.e. $\sigma^l(t) \geq \sigma^l(s) + \sigma^l(t-s)$ for all $t \geq 0$ and $0 \leq s \leq t$ and $\sigma^u(t) \leq \sigma^u(s) + \sigma^u(t-s)$ for all $t \geq 0$ and $0 \leq s \leq t$.

3.3 Bounding the Analysis Interval

So far, our computation of the service bounds σ^l and σ^u were based on the fact that the arrival curves α_x and α_C and the workload curve γ are known for all possible time interval lengths $\Delta \geq 0$. These curves would usually be derived by simulating the processing or execution of several representative audio/video samples on a template platform architecture, as explained in Section 1. The traces collected from such a simulation—from the different parts of the platform architecture, such as the arrival pattern of stream objects in front of PE_2 in Figure 1—are then analyzed to derive the different arrival and workload curves. However, since these representative audio/video samples would always be of finite length, the curves or bounds derived from the resulting traces would also be of finite length. But the platform designed on the basis of these finite length traces might later be used to process larger audio/video samples. Hence, we would like to guarantee the buffer overflow and underflow constraints on input streams of any length (provided they satisfy the bounds dictated by the arrival and the workload curves), although the analysis and the design of the platform is based on only finite length representative inputs.

We would however like to point out here that in practice the above issue will not be of major concern to any system designer. He would use sufficiently long (but finite length) representative

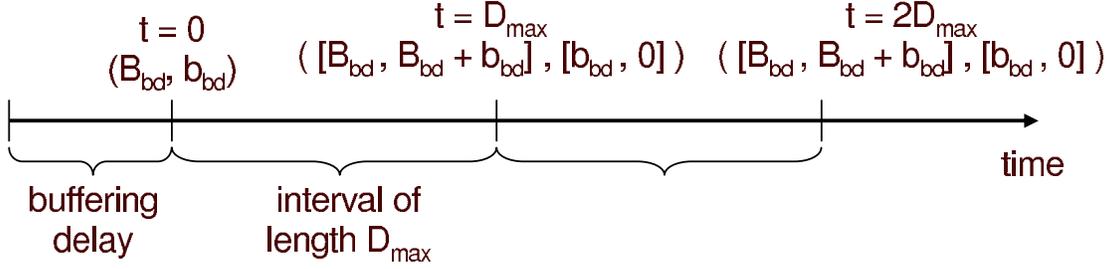


Figure 3: Fill levels of the playout buffer and the internal buffer as the system evolves over time.

audio/video samples in the initial simulation phase to derive the bounds (i.e. arrival and workload curves) that any input belonging to the class represented by these audio/video samples is expected to satisfy. Based on these bounds, the platform architecture in question would be designed. When such an architecture processes input streams which are longer in duration than the samples used for designing the architecture, it is assumed that the variability of the entire stream is bounded by the variability existing in the sample inputs. Such assumptions are not specific to our framework and are common whenever a system is designed based on *representative inputs* (for example, see [13]).

However, it is also possible to formally address this issue within our framework. Towards this, note that the motivation behind using the arrival curves and the workload curves is (i) to capture the burstiness in the on-chip traffic associated with multimedia processing on a multiprocessor SoC architecture [25], and (ii) to capture the variability in the execution time of multimedia processing tasks (which also adds to the burstiness in the on-chip traffic). However, such burstiness and variability is restricted to only short time intervals, and in the long-term both the traffic and the execution requirements exhibit average-case behaviour. To capture such a behaviour, we assume that for all multimedia streams belonging to a particular *class*, the total execution requirement from a PE over a (sufficiently long) prespecified time interval of length Δ_{max} is fixed. However, there is variability over time intervals of length smaller than Δ_{max} and such variability is captured by the workload curves. Similarly, the real-time client at the output also consumes a fixed number of stream objects over any such time interval of length Δ_{max} and exhibits variability over time intervals of smaller lengths. This implies that if we look at a PE over any time interval of length Δ_{max} , then within this time interval the processor requires a fixed number of processor cycles to process as many stream objects as those are consumed by the real-time client within this time interval.

The evolution of a system exhibiting such a behaviour is shown in Figure 3. The first phase consists of a *buffering delay*—during this time the output real-time client does not consume any stream objects and the playout buffer is filled up to a certain level. Our analysis begins at the end of this phase, i.e. at time $t = 0$, and starting from this time the real-time client starts consuming stream objects from the playout buffer. Let the playout buffer fill level at time $t = 0$ (i.e. after the buffering delay) be B_{bd} and the internal buffer fill level at this time be equal to b_{bd} . Over the time interval $t = 0$ to $t = \Delta_{max}$ the service provided by the PE is bounded by the functions σ^l and σ^u described in Section 3.2. However, in contrast to the derivation shown in Section 3.2, these functions are not computed over *all* possible time interval lengths $\Delta > 0$, but are only computed for all $\Delta \leq \Delta_{max}$. Further, in the computation of β^l and β^u (from which σ^l and σ^u are derived), the value of B_0 is set to 0 and b_0 to b_{bd} while calculating β^l , and the value of B_0 is set to $B_{bd} + b_{bd}$ while calculating

β^u . This is because the evolution of the system, as shown in Figure 3, can be seen as a sequence of consecutive *episodes*, each of which is of length Δ_{\max} . As explained above, within each such episode, a fixed number of processor cycles are required to process a sequence of stream objects, the number of which is equal to the number of stream objects consumed by the real-time client within this episode. To guarantee this, we assume that the minimum service provided by the PE is exactly sufficient to completely process these number of stream objects. This service must however be greater than $\mathcal{I}(\Delta)$ for all $\Delta \leq \Delta_{\max}$ to guarantee the buffer overflow and underflow constraints within this time interval. It is easy to note that this minimum service is exactly equal to the long-term average rate at which processor cycles must be provided to process the input stream. If the service provided is exactly equal to this average rate, then at time $t = \Delta_{\max}$ the fill levels of the playout buffer and the internal buffer are again equal to their fill levels at time $t = 0$, i.e. B_{bd} and b_{bd} respectively. However, if the total number of processor cycles provided is greater than what is dictated by the average rate, then at time $t = \Delta_{\max}$ the buffer fill levels of the playout buffer and the internal buffer can be at most $B_{bd} + b_{bd}$ and 0 respectively. The same argument holds for all the subsequent episodes. Therefore, at the start of any episode, the maximum fill level of the playout buffer is equal to $B_{bd} + b_{bd}$ and that of the internal buffer is equal to b_{bd} . Notice that when the initial buffer backlog B_0 is non-zero, the PE can provide less service than the long-term average rate since the number of stream objects required to be generated by the PE within this episode can be reduced by B_0 . So we set B_0 to 0 while calculating β^l . Therefore, setting B_0 and b_0 to be equal to the above values while calculating σ^l and σ^u guarantees that all the buffer constraints are satisfied within any episode, and hence also at all times.

3.4 Extending the Analysis to Other PEs

The framework presented so far is based on the assumption that the PE being analyzed is the last one in the path of a stream i.e. its output is directly written into the playout buffer. Now let us consider a PE, whose output is fed into another PE i.e. the next PE in the path of the stream. An example of such a PE is PE_1 in Figures 1 and 2. To derive the service bounds for this PE, let us denote the arrival curve corresponding to the arrival pattern of stream objects at the internal buffer of PE_2 as α_{x_2} . Similarly, let the arrival pattern of stream objects at the internal buffer of PE_1 be bounded by α_{x_1} , and let the size of this internal buffer be b_1 . Then bounds on β^l and β^u (such as those given by inequalities (8) and (10)) for PE_1 can be calculated from α_{x_1} , α_{x_2} and b_1 . However, the only constraint that needs to be satisfied in this case is that the internal buffer of PE_1 should not overflow. The resulting bounds on the service are therefore much simpler than the ones derived above, and hence we omit them here. This same scheme can be applied to other PEs in the path of the stream which are away from the playout buffer. If all of the PEs provide a service in accordance with the bounds computed for them, then it is guaranteed that none of the internal buffers in the architecture will overflow, and the playout buffer will neither overflow and nor underflow.

4 Computing Processor Frequency Ranges

Given the service bounds σ^l and σ^u for a PE, in this section we compute the discrete frequency levels or the frequency range that must be supported by the PE in order to realize these service bounds. For any given multimedia application and a class of input streams to be processed, accurately determining

the appropriate processor frequency range is a non-trivial problem. A straightforward approach would be to choose a processor frequency that is sufficient to process stream objects at a rate equal to the long-term average rate at which stream objects are consumed by the output real-time client. However, since the input stream is bursty in nature and there is data-dependent variability in the execution requirement of stream objects, processing stream objects at this average rate might lead to buffer overflows and underflows. Therefore, the commonly followed practice is to choose a processor frequency which is slightly higher than what is required to process stream objects at the average consumption rate. However, there is no well known formal guideline on how much higher this chosen frequency should be. If the chosen frequency is much higher than what is required to counter the effects of bursts, then complicated blocking reads and writes must be used, which requires substantial operating system support.

The situation is much more complicated when the PE in question has to process multiple classes of input streams or multiple applications. In such cases, the different input classes might have different computational demands and hence require different processor frequencies. Here it is important to determine the range of processor frequencies that must be supported for each class. If these ranges overlap, then the processor might support some frequency belonging to this overlapping range. But if these ranges do not overlap, then multiple frequency levels need to be supported. Further, the processor frequency range to be supported by a PE is heavily dependent on the size of the on-chip buffers. A platform designer would therefore be interested in obtaining insights into this dependency. Our results presented below would help in obtaining such insights. These results can be summarized as follows:

- For any application and a class of input streams, we can statically generate frequency schedules for a PE which satisfy all the buffer constraints. Such schedules specify the frequency with which the PE should be run at any time.
- We derive a frequency range (f_{\min}, f_{\max}) such that *all* feasible frequency scheduling algorithms will only use frequencies within this range. Therefore, it would be sufficient if the PE supports frequencies belonging to this range only.
- We also derive a frequency range $(\bar{f}_{\min}, \bar{f}_{\max})$ such that at any time, choosing any frequency within this range will lead to all the buffer constraints being satisfied. However, such frequency schedules will be a subset of *all* possible feasible frequency schedules. Further, $(\bar{f}_{\min}, \bar{f}_{\max}) \subseteq (f_{\min}, f_{\max})$. The main difference between the above two frequency ranges is that if a PE supports only the latter range, then the set of possible frequency schedules is restricted but the frequency scheduling algorithm is simpler because at any time, any frequency within this range can be selected. We call the former range the *History-Dependent Frequency (HDF) Range* since the frequency that can be chosen at any time depends on the frequencies chosen at previous time instants. The latter range is called the *History-Independent Frequency (HIF) Range*. The main motivation behind identifying the HDF range is that it allows more flexibility in choosing a frequency schedule: based on the HDF range, a designer can choose a processor which supports only a fixed number of frequency levels and runs this processor using a static frequency schedule which contains only these frequency levels. However, these frequency levels might not belong to the HIF range.

- Finally, our framework can also be used to identify how the bounds (f_{\min}, f_{\max}) and $(\bar{f}_{\min}, \bar{f}_{\max})$ change by changing the on-chip buffer sizes (both the playout and the internal buffers).

For simplicity, we assume that the processor frequency can be changed at each time unit. Then during any run of the processor over a time interval of length n (where n may be equal to Δ_{\max} defined in Section 3.3), let its frequency values be f_1, \dots, f_n , i.e. f_i is the frequency at which the processor is run during the time interval $t = (i - 1, i]$. If the service being offered to a stream as a result of this schedule has to be bounded by the service curves σ^l and σ^u , then f_1, \dots, f_n is required to satisfy the following inequalities:

$$\begin{aligned}
\sigma^l(1) &\leq f_i && \leq \sigma^u(1) \quad \forall i = 1, 2, \dots, n \\
\sigma^l(2) &\leq f_i + f_{i+1} && \leq \sigma^u(2) \quad \forall i = 1, 2, \dots, n - 1 \\
\sigma^l(3) &\leq f_i + f_{i+1} + f_{i+2} && \leq \sigma^u(3) \quad \forall i = 1, 2, \dots, n - 2 \\
&\dots && \\
\sigma^l(n) &\leq f_i + f_{i+1} + f_{i+2} + \dots + f_{i+n-1} && \leq \sigma^u(n) \quad \forall i = 1
\end{aligned}$$

The above constraints may be summarized as follows. For all $\Delta = 1, 2, \dots, n$ and $i = 1, 2, \dots, n - \Delta + 1$,

$$\sigma^l(\Delta) \leq \sum_{j=0}^{\Delta-1} f_{i+j} \leq \sigma^u(\Delta) \quad (13)$$

From the constraints given by inequality (13), it may be seen that any frequency value f_i is dependent on all the previously assigned frequencies. To be more clear about how this dependency is, from inequality (13) we identify all those inequalities that only include f_i and the values in f_1, \dots, f_{i-1} , which is shown in the following,

$$\begin{aligned}
\sigma^l(1) &\leq f_i && \leq \sigma^u(1) \\
\sigma^l(2) &\leq f_{i-1} + f_i && \leq \sigma^u(2) \\
\sigma^l(3) &\leq f_{i-2} + f_{i-1} + f_i && \leq \sigma^u(3) \\
&\dots && \\
\sigma^l(i) &\leq f_1 + f_2 + \dots + f_{i-1} + f_i && \leq \sigma^u(i)
\end{aligned} \quad (14)$$

The above inequality shows that f_i is dependent on the sum of all the previous k assigned frequencies proximate to it, for any $k = 0, \dots, i - 1$.

Randomly Generated Frequency Schedules: Suppose that all the previous frequencies f_1, \dots, f_{i-1} have been assigned, by solving inequality (14), the lower and upper bounds on any f_i is given as follows:

$$\begin{aligned}
f_1^l &= \sigma^l(1) \\
f_1^u &= \sigma^u(1) \\
f_2^l &= \max\{\sigma^l(1), \sigma^l(2) - f_1\} \\
f_2^u &= \min\{\sigma^u(1), \sigma^u(2) - f_1\} \\
f_3^l &= \max\{\sigma^l(1), \sigma^l(2) - f_2, \sigma^l(3) - (f_1 + f_2)\} \\
f_3^u &= \min\{\sigma^u(1), \sigma^u(2) - f_2, \sigma^u(3) - (f_1 + f_2)\} \\
&\dots
\end{aligned}$$

Hence,

$$\begin{aligned} f_i^l &= \max_{1 \leq j \leq i-1} \{ \sigma^l(1), \sigma^l(i-j+1) - \sum_{p=j}^{i-1} f_p \} \\ f_i^u &= \min_{1 \leq j \leq i-1} \{ \sigma^u(1), \sigma^u(i-j+1) - \sum_{p=j}^{i-1} f_p \} \end{aligned} \quad (15)$$

f_i^l and f_i^u therefore give lower and upper bounds on the frequency that can be assigned during the time interval $(i-1, i]$, provided the frequencies at all the previous time intervals are known. To generate a static frequency schedule, we can choose any $f_i \in [f_i^l, f_i^u]$, and the chosen f_i will determine the range $[f_{i+1}^l, f_{i+1}^u]$.

HDF Range: To compute the HDF range, that was mentioned above, we first compute the lower and upper bounds on f_i , by solving inequality (13). From equality (15), we know that the bounds on f_i is dependent on all values of $\sum_{p=j}^{i-1} f_p$. Like f_i , $\sum_{p=j}^{i-1} f_p$ is also dependent on all the previous frequencies f_1, \dots, f_{j-1} , the constraints for which can be derived in a similar way to that for f_i (refer to inequality (14)). We define two functions $(\sum_{p=j}^{i-1} f_p)^l$ and $(\sum_{p=j}^{i-1} f_p)^u$. The first provides a lower bound, and the second an upper bound on the sum $f_j + \dots + f_{i-1}$.

$$\begin{aligned} \left(\sum_{p=j}^{i-1} f_p \right)^l &= \max_{1 \leq q \leq j-1} \left\{ \sigma^l(i-j), \sigma^l(i-q) - \left(\sum_{p=q}^{j-1} f_p \right)^u \right\} \\ \left(\sum_{p=j}^{i-1} f_p \right)^u &= \min_{1 \leq q \leq j-1} \left\{ \sigma^u(i-j), \sigma^u(i-q) - \left(\sum_{p=q}^{j-1} f_p \right)^l \right\} \end{aligned}$$

Using the above two functions, we can derive the bounds on f_i as

$$\begin{aligned} f_i^{\min} &= \max_{1 \leq j \leq i-1} \left\{ \sigma^l(1), \sigma^l(i-j+1) - \left(\sum_{p=j}^{i-1} f_p \right)^u \right\} \\ f_i^{\max} &= \min_{1 \leq j \leq i-1} \left\{ \sigma^u(1), \sigma^u(i-j+1) - \left(\sum_{p=j}^{i-1} f_p \right)^l \right\} \end{aligned}$$

f_i^{\min} is the smallest possible processor frequency that can be assigned during the time interval $(i-1, i]$, and f_i^{\max} is the largest possible processor frequency that can be assigned during this time interval. Then the HDF range is defined as: $f_{\min} = \min_{i=1, \dots, n} \{ f_i^{\min} \}$, and $f_{\max} = \max_{i=1, \dots, n} \{ f_i^{\max} \}$. A dynamic programming method is also proposed to compute the HDF range, as shown in algorithm 1.

HIF Range: From the service curves, the HIF range is defined as: $\bar{f}_{\min} = \max_{\Delta=1, \dots, n} \{ \sigma^l(\Delta) / \Delta \}$, and $\bar{f}_{\max} = \min_{\Delta=1, \dots, n} \{ \sigma^u(\Delta) / \Delta \}$. As mentioned above, frequency schedules restricted to this range can choose any frequency value within this range at any time instant.

Proof sketch: When we replace the service curve $(\sigma^l(\Delta), \sigma^u(\Delta))$ by $(\bar{f}_{\min} \cdot \Delta, \bar{f}_{\max} \cdot \Delta)$ in inequality (13), it is easy to see that any such frequency schedule satisfies the new inequality. Since for any $\Delta = 1, \dots, n$, $\sigma^l(\Delta) \leq \bar{f}_{\min} \cdot \Delta$ and $\bar{f}_{\max} \cdot \Delta \leq \sigma^u(\Delta)$, it follows that any such frequency schedule satisfies inequality (13), and thus satisfies the buffer overflow and underflow constraints.

5 Case Study: MPEG-2 Decoder Application

In this section we present a case study to illustrate an application of the framework presented in the last two sections. Towards this, we map a MPEG-2 decoder application onto a platform architecture consisting of two PEs. The goal is to compute the processor frequency range that needs to be supported by one of the PEs and also identify how this range changes by changing the on-chip buffer size.

As shown in Figure 1, the MPEG-2 decoder application is partitioned into a set of tasks executing in parallel on two PEs of the system architecture. PE_1 executes the variable length decoding (VLD)

Algorithm 1 Algorithm of Computing HDF Range

Input: service curves $\sigma^l(\Delta)$, $\sigma^u(\Delta)$ and time interval length n ;

Definition: arrays $F^l[n, n]$ and $F^u[n, n]$;

Initialization: $F^l(i, k) \leftarrow \sigma^l(k)$, $F^u(i, k) \leftarrow \sigma^u(k)$ for all $1 \leq i, k \leq n$;

for $i \leftarrow 2$ to n **do**

for $j \leftarrow i$ to 2 **do**

$F^l(j, i - j + 1) \leftarrow \max_{1 \leq q \leq j-1} \{ \sigma^l(i - j + 1), \sigma^l(i + 1 - q) - F^u(q, j - q) \}$
 /* $F^l(j, i - j + 1)$ stores $(\sum_{p=j}^i f_p)^l$ */
 $F^u(j, i - j + 1) \leftarrow \min_{1 \leq q \leq j-1} \{ \sigma^u(i - j + 1), \sigma^u(i + 1 - q) - F^l(q, j - q) \}$
 /* $F^u(j, i - j + 1)$ stores $(\sum_{p=j}^i f_p)^u$ */

end for

end for

for $i \leftarrow 1$ to n **do**

$f_i^{\min} \leftarrow F^l(i, 1)$, $f_i^{\max} \leftarrow F^u(i, 1)$;

end for

$f_{\min} = \min_{i=1, \dots, n} \{ f_i^{\min} \}$, $f_{\max} = \max_{i=1, \dots, n} \{ f_i^{\max} \}$;

and the inverse quantization (IQ) tasks, while PE_2 executes the inverse discrete cosine transform (IDCT) and the motion compensation (MC) tasks. A compressed video bit stream arrives from the network interface into the input buffer of PE_1 . After processing on PE_1 the partially decoded stream consisting of stream objects called *macroblocks* enters the buffer B_2 in front of PE_2 . PE_2 reads the buffer one macroblock at a time and computes for each macroblock the IDCT and MC functions. After that the video stream emerges out of PE_2 as a fully decoded stream of macroblocks. This stream is written into the playout buffer B_v , which is read at a constant rate by the video output port V_{out} . The video output port represents the real-time client (RTC) in this setup. The rate with which it reads the playout buffer B_v is determined by the resolution and the frame rate of the decoded MPEG-2 video sequence. As we mentioned above, the service provided to the video stream on the PEs of such an architecture is dependent on the buffer constraints. In the above setup, at any point in time none of the buffers B_1 , B_2 and B_v are allowed to overflow, and the playout buffer B_v should not underflow.

Determining the service that must be offered by the PEs to the stream, and thereby identifying their feasible clock frequency ranges under the given buffer constraints is not an easy task. The main complexity of the problem stems from the highly variable load imposed on the PEs of the architecture by the video stream. For example, let us consider the load imposed by the stream on PE_2 . Firstly, the execution demand of IDCT and MC tasks performed by PE_2 varies for different types of macroblocks (e.g. because of the various kinds of motion compensation methods that have to be applied to the compressed macroblocks). Secondly, the arrival pattern of macroblocks into the buffer B_2 has a high degree of burstiness, which is caused by the variability in the execution demand of the tasks executing on PE_1 . The overall result is a very complex and variable nature of the processing load imposed on PE_2 . This further increases the burstiness of the macroblock stream emerging at its output and entering the playout buffer B_v .

Now, using the above example of PE_2 we will demonstrate how the proposed methodology can be applied to compute the required service bounds σ and the associated feasible clock frequency ranges

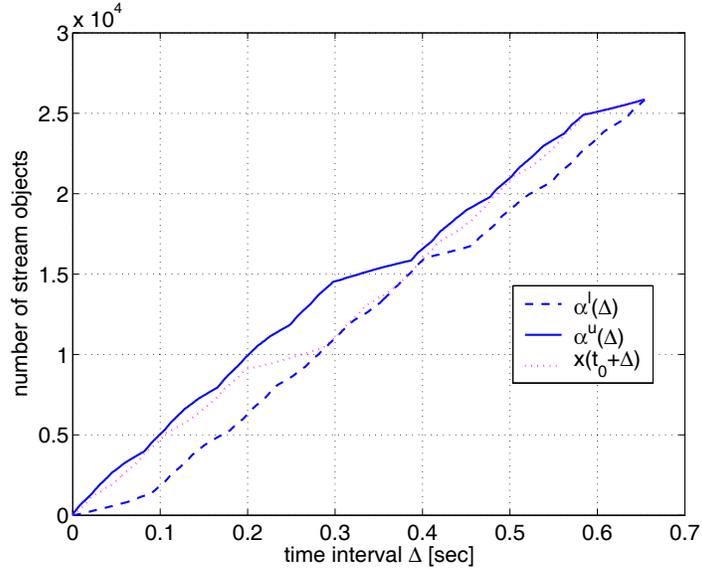


Figure 4: Arrival curves (α_x^l, α_x^u) of the macroblock stream on the output of PE_1 for the video sequence $video_1$. A fragment of the function $x(t)$ for $video_1$ is shown in this figure. Note that it is bounded by the corresponding arrival curves.

of PE_2 for the given MPEG-2 decoder application.

5.1 Computing the Service Bounds and the Frequency Range for PE_2

Before we can compute the service bounds σ and the corresponding HDF and HIF ranges for PE_2 , we need to obtain the arrival curves α_x and α_C and the workload curves γ characterizing the stream processed by PE_2 , and the real-time client V_{out} . In general, the construction of such curves can be performed in many different ways. In some cases it is possible to derive the curves analytically from a formal specification of the system and its environment, whereas in other cases a simulation- and trace-based analysis approach may be necessary and indeed sufficient for the problem at hand. In our case, we adopt the latter method to obtain the α_x and the γ curves. Towards this we collect execution traces from a simulation of an *abstract model* of the platform architecture and then analyze them to derive the required curves. The abstract model of the platform is based on an instruction set simulator, which we use to obtain the traces of execution demands of the MPEG-2 decoder tasks. In our experimental setup we use a customized version of the SimpleScalar instruction set simulator for this purpose.

The arrival curves (α_x^l, α_x^u) of the stream at the output of PE_1 can be obtained by measuring the execution demands of the VLD and IQ tasks for each macroblock in the video sequence and by taking into account (i) the constant arrival rate of the compressed bit stream at the input of PE_1 , and (ii) the amount of bits allocated to encode each macroblock in the stream. In this procedure we first obtain the cumulative function $x(t)$ describing the arrival of macroblocks in the buffer B_2 shown in Figure 1. We then analyze it using time intervals of different lengths Δ to obtain the arrival curves. Figure 4 shows the arrival curves (α_x^l, α_x^u), which we have obtained by applying the above method for a representative 4 Mbps video sequence $video_1$.

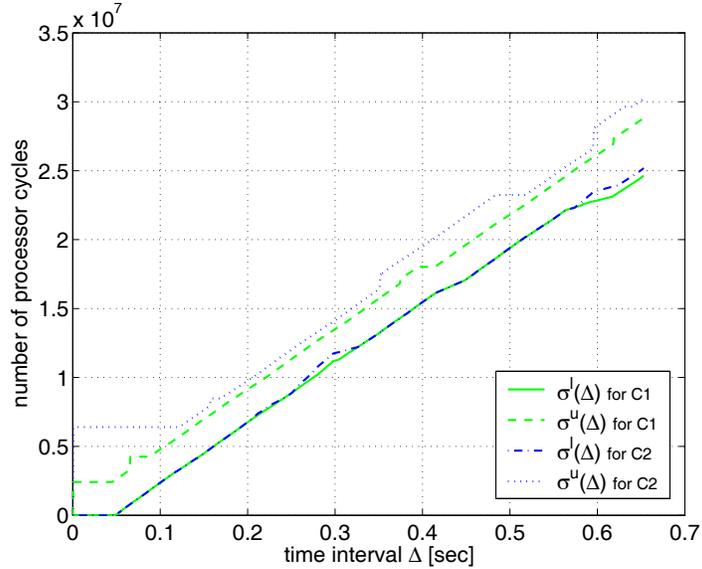


Figure 5: Service bounds (σ^l, σ^u) for *video*₁ for two different system configurations *C1* and *C2*, where $C1 = \{B_2 = 4500, B_v = 5000\}$ and $C2 = \{B_2 = 4000, B_v = 5600\}$.

In similar way, we can derive the arrival curves (α_C^l, α_C^u) . However, in this case since we precisely know the characteristics of the real-time client we do not need to rely on the simulation. α_C^l and α_C^u can be constructed analytically by using the fact that the real-time client reads macroblocks from the playout buffer at a predefined constant rate.

To obtain the workload curves (γ^l, γ^u) , we first collect a trace of execution demands for the pair of tasks IDCT and MC executing on PE_2 . Using a method similar to the one used for obtaining the arrival curves α_x , we then analyze this trace to identify the maximum and the minimum processing demand imposed by any sequence of k consecutive macroblocks within the video sequence.

Now we apply the results presented in Section 3 to compute the cycle-based service bounds (σ^l, σ^u) corresponding to the service that must be offered by PE_2 , to any video stream belonging to the class of streams bounded by the curves α_x , α_C and γ . The bounds (σ^l, σ^u) corresponding to the example video sequence *video*₁, for two different *system configurations* are shown in Figure 5. The two system configurations differ only in the sizes of the buffers B_v and B_2 . By examining the plots on Figure 5 we can see that even a relatively small change in the available buffer space may have a considerable impact on the service bounds. Furthermore, the distribution of the total on-chip buffer space among the different buffers may also have an impact on the service bounds..

Video sequences belonging to different classes of streams may have very different on-chip buffer requirements. Therefore, the service bounds for these sequences, and hence their feasible clock frequency ranges might also be very different. This information about how different these ranges might be for different classes of video sequences can be efficiently obtained from the service bounds σ , as described in Section 4. In our example, using the proposed framework we have computed the HIF ranges for two classes of video sequences characterized by different input bit stream rates, i.e. for 4 Mbps and 8 Mbps MPEG-2 streams. The resulting HIF ranges are shown in Figure 6.

Figure 6 shows the dependency of the HIF range on the playout buffer size for 4 Mbps and a

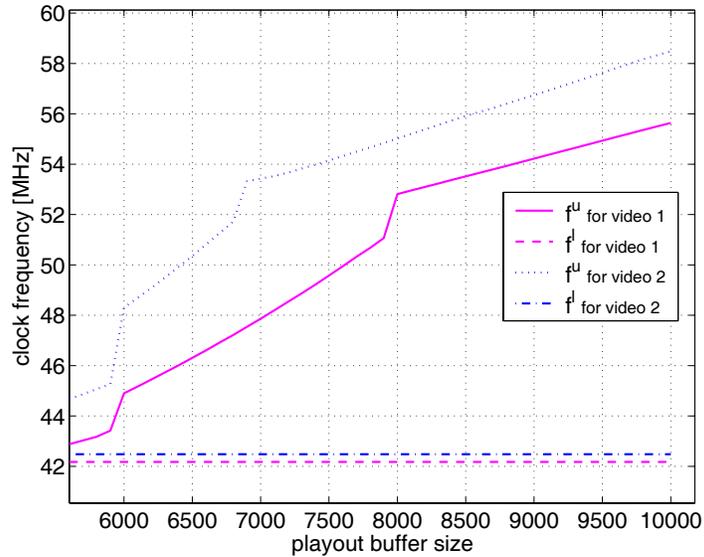


Figure 6: Dependency of HIF ranges on the playout buffer size for two different classes of the MPEG-2 video streams: 4 Mbps ($video_1$) and 8 Mbps ($video_2$). The size of buffer B_2 is fixed to 4000 macroblocks.

8 Mbps MPEG-2 video streams. In this figure it can be seen that the playout buffer size has a considerable impact on the upper frequency bound f^u . By increasing the buffer size, the maximum frequency with which PE_2 can run, also increases. This corresponds to the intuitive understanding that the larger the playout buffer size, the more bursty the incoming stream can be.

Figure 6 shows an overlap in the HIF ranges of the two classes of video streams. This implies that for any *feasible* playout buffer size and a fixed size of B_2 (set to 4000 macroblocks), we can always find a clock frequency with which PE_2 can be run for video sequences belonging to both the input classes (i.e. 4 Mbps and 8 Mbps input rates). It may be noted here playout buffers only beyond a certain size are feasible—meaning that, for them feasible service bounds σ exists. It may also be noted that in general such a common clock frequency for any two input classes might not exist for a single system configuration. In such cases it will be necessary to support multiple frequency ranges/values, where the frequency level at which the processor is run depends on the class to which the input belongs. Alternatively, the configuration of the system can be changed (for example by increasing buffer sizes), till the frequency ranges of two input classes overlap. When this happens, once again it would be sufficient for the processor to support a single frequency level belonging to this overlapping range. Our methodology can be used to efficiently identify such design tradeoffs in the case of configurable platform architectures.

5.2 Validation of the Analytical Bounds

To validate our framework we simulated the platform architecture using static schedules generated using the framework. Towards this, we used a detailed simulator of the system shown in Figure 1. The simulator consisted of a transaction level model of the system architecture written in SystemC, and

video sequence	buffer sizes		schedule	measured backlogs	
	B_2	B_v		B_2	B_v
$video_1$	4000	5600	f^u	3335	4763
			f^l	3719	4621
			rand	3614	4844
			rand	3678	4980
	4500	6000	f^u	2680	5080
			f^l	4078	4615
			rand	3472	3472
			rand	3614	4844
$video_2$	4000	5600	f^u	2479	4525
			f^l	3971	4132
			rand	3368	4754
			rand	3502	4500
	4500	5000	f^u	2786	4330
			f^l	4004	4132
			rand	3227	4401
			rand	2818	4402

Table 1: The maximum buffer fill levels obtained by simulating a static frequency schedule for PE_2 that was derived using the proposed framework. $video_1$ and $video_2$ are 4 Mbps and 8 Mbps MPEG-2 video streams respectively.

the models of PEs were based on a customized version of the SimpleScalar instruction set simulator. Using this simulation setup, we measured the maximum backlogs and recorded any buffer underflows that occurred as a result of running the system with a static frequency schedule for PE_2 (which was generated using our framework). The results of this simulation are summarized in Table 1.

We evaluated several frequency schedules for PE_2 that are bounded by the computed HDF and HIF ranges obtained using the proposed framework. These ranges correspond to different system configurations and classes of the video streams. In Table 1 these schedules are indicated as f^u , f^l , and $rand$. f^u and f^l denote cases where PE_2 was run with a single clock frequency determined by the upper \bar{f}_{max} and the lower \bar{f}_{min} HIF bounds respectively (see Figure 6 for a reference). $rand$ indicates randomly generated static schedules from the HDF bounds (f_{min}, f_{max}) , as explained in Section 4. Two such randomly generated schedules are shown in Figure 7.

In all the simulations we performed, the maximum backlogs measured in the buffers never exceeded the buffer sizes. Furthermore, our simulation results also showed that playout buffer underflows never occurred for any of the simulated frequency schedules. These simulations therefore validate the proposed framework and suggest its practicality. Finally, we would once again like to point out that obtaining equivalent results using purely simulation based approaches is extremely time consuming and such approaches usually fail to provide any formal performance guarantees.

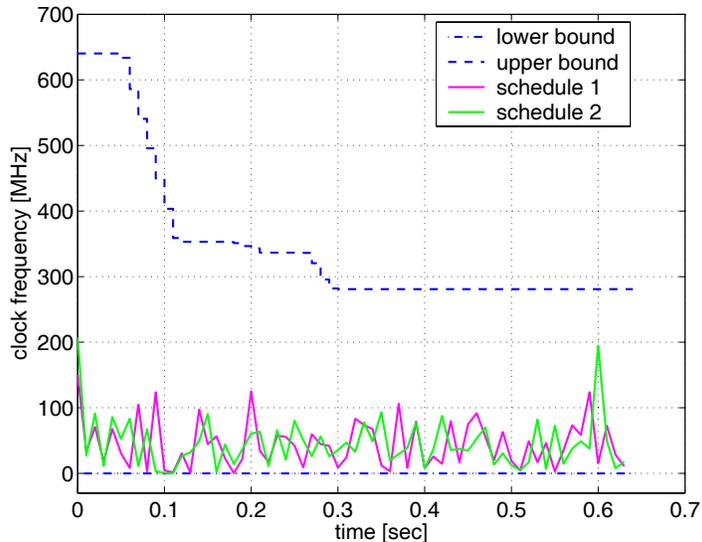


Figure 7: The HDF range and two randomly generated schedules obtained from the service bounds σ .

6 Concluding Remarks

In this paper we presented an analytical framework that can be used for the design space exploration of parameters or configurations of SoC platform architectures for multimedia processing, that contain processor cores which support dynamic voltage/frequency scaling. In contrast to simulation based approaches, which usually follow a trial-and-error approach and involve very high simulation times, the proposed framework can provide useful insights into the design space and can aid a system designer in systematically tuning a platform architecture for a class of applications.

Throughout this paper we have assumed that a PE processes a single input stream. However, in general a PE might process multiple input streams. For example, in the architecture shown in Figure 1, assume that on PE_2 , in addition to the IDCT and the MC tasks, a MP3 decoder task is also implemented. In such a case, this PE processes a audio stream in addition to the video stream shown in the figure. Further, a task scheduler based on some scheduling policy chooses the stream to be processed at any time instant. Given a specification of the audio stream, our framework can be extended to identify the frequency range to be supported by the PE in this case. In addition to the theory presented in this paper, this requires a modelling of the task scheduling discipline and how the *total* service offered by the PE is divided among the two streams. Some of ideas pertaining to such a scheme may be found in [7].

References

- [1] A. Acquaviva, L. Benini, and B. Ricc . An adaptive algorithm for low-power streaming multimedia processing. In *Conference on Design, Automation and Test in Europe (DATE)*, 2001.
- [2] PALM-DP-2000 AcurX configurable SoC platform.
<http://www.palmchip.com/>.

- [3] H. V. Antwerpen, N. Dutt, R. Gupta, S. Mohapatra, C. Pereira, N. Venkatasubramanian, and R. von Vignau. Energy-aware system design for wireless multimedia. In *IEEE Design, Automation and Test in Europe (DATE)*, 2004.
- [4] J.-Y. Le Boudec and P. Thiran. *Network Calculus - A Theory of Deterministic Queuing Systems for the Internet*. LNCS 2050, 2001.
- [5] M. Buss, T. Givargis, and N. Dutt. Exploring efficient operating points for voltage scaled embedded processor cores. In *24th IEEE Real-Time Systems Symposium (RTSS)*, 2003.
- [6] H. Kim C. Im and S. Ha. Dynamic voltage scheduling technique for low-power multimedia applications using buffers. In *International Symposium on Low Power Electronics and Design (ISLPED)*, California, USA, 2001.
- [7] S. Chakraborty, S. Künzli, and L. Thiele. A general framework for analysing system properties in platform-based embedded system designs. In *6th Design, Automation and Test in Europe (DATE)*, Munich, Germany, February 2003.
- [8] S. Chakraborty, S. Künzli, L. Thiele, A. Herkersdorf, and P. Sagmeister. Performance evaluation of network processor architectures: Combining simulation with analytical estimation. *Computer Networks*, 41(5), 2003.
- [9] K. Choi, K. Dantu, W.-C. Cheng, and M. Pedram. Frame-based dynamic voltage and frequency scaling for a MPEG decoder. In *ICCAD*, 2002.
- [10] S. Dutta, R. Jensen, and A. Rieckmann. Viper: A multiprocessor SOC for advanced set-top box and digital TV systems. *IEEE Design & Test of Computers*, 18(5):21–31, September-October 2001.
- [11] I. Hong, D. Kirovski, G. Qu, M. Potkonjak, and M.B. Srivastava. Power optimization of variable-voltage core-based systems. *IEEE Trans. on Computer Aided-Design of Integrated Circuits and Systems*, 18(12), 1999.
- [12] C.J. Huges, J. Srinivasan, and S.V. Adve. Saving energy with architectural and frequency adaptations for multimedia applications. In *34th Annual International Symposium on Microarchitecture (MICRO)*, 2001.
- [13] K. Lahiri, A. Raghunathan, and S. Dey. System level performance analysis for designing on-chip communication architectures. *IEEE Trans. on Computer Aided-Design of Integrated Circuits and Systems*, 20(6):768–783, 2001.
- [14] A. Maxiaguine, S. Künzli, S. Chakraborty, and L. Thiele. Rate analysis for streaming applications with on-chip buffer constraints. In *ASP-DAC*, 2004.
- [15] S. Mohapatra, R. Cornea, N. Dutt, A. Nicolau, and N. Venkatasubramanian. Integrated power management for video streaming to mobile handheld devices. In *ACM Multimedia (MM)*, 2003.
- [16] OMAP for 2.5G and 3G: Overview, Texas Instruments. <http://www.ti.com/sc/omap/>.

- [17] PrimeXsys Platforms Overview, ARM.
<http://www.arm.com/products/solutions/PrimeXsysPlatforms.html>.
- [18] M.J. Rutten, J.T.J. van Eijndhoven, E.G.T. Jaspers, P. van der Wolf, O.P. Gangwal, and A. Timmer. A heterogeneous multiprocessor architecture for flexible media processing. *IEEE Design & Test of Computers*, 19(4):39–50, July-August 2002.
- [19] M.J. Rutten, J.T.J. van Eijndhoven, and E.-J.D. Pol. Design of multi-tasking coprocessor control for eclipse. In *10th International Workshop on Hardware/Software Codesign (CODES)*, Colorado, USA, May 2002.
- [20] M.J. Rutten, J.T.J. van Eijndhoven, and E.-J.D. Pol. Robust media processing in a flexible and cost-effective network of multi-tasking coprocessors. In *14th Euromicro Conference on Real-Time Systems (ECRTS)*, 2002.
- [21] K. Sekar, K. Lahiri, and S. Dey. Dynamic platform management for configurable platform-based system-on-chips. In *International Conference on Computer Aided Design (ICCAD)*, San Jose, California, November 2003.
- [22] K. Sekar, K. Lahiri, and S. Dey. Configurable platforms with dynamic platform management: An efficient alternative to application-specific system-on-chips. In *17th International Conference on VLSI Design (VLSI)*, Mumbai, India, January 2004.
- [23] L. Thiele, S. Chakraborty, M. Gries, and S. Künzli. A framework for evaluating design tradeoffs in packet processing architectures. In *39th Design Automation Conference (DAC)*, 2002.
- [24] P. van der Wolf, W.M. Kruijtzter, and J.T.J. van Eijndhoven. System-level design of embedded media systems. In *Tutorial at the 15th International Conference on VLSI Design (VLSI) and Asia and South Pacific Design Automation Conference (ASP-DAC) (joint conference)*, Bangalore, India, January 2002.
- [25] G. Varatkar and R. Marculescu. On-chip traffic modeling and synthesis for MPEG-2 video applications. *IEEE Transactions on VLSI*, 12(1), January 2004.
- [26] W. Yuan and K. Nahrstedt. Energy-efficient soft real-time cpu scheduling for mobile multimedia systems. In *19th ACM Symposium on Operating Systems Principles (SOSP)*, 2003.