

# A Model-Driven Framework for Hardware-Software Co-design of Dataflow Applications

Waheed Ahmad, Bugra M. Yildiz, Arend Rensink, and Mariëlle Stoelinga

University of Twente, The Netherlands

{w.ahmad, b.m.yildiz, arend.rensink, m.i.a.stoelinga}@utwente.nl

**Abstract.** Hardware-software (HW-SW) co-design allows to meet system-level objectives by exploiting the synergy of hardware and software. Current tools and approaches for HW-SW co-design face difficulties coping with the increasing complexity of modern-day application due to, e.g., concurrency and energy constraints. Therefore, an automated modeling approach is needed which satisfies modularity, extensibility and interoperability requirements. Model-Driven Engineering (MDE) is a prominent paradigm that, by treating models as first-class citizens, helps to fulfill these requirements. This paper presents a state-of-the-art MDE-based framework for HW-SW co-design of dataflow applications, based on synchronous dataflow (SDF) graph formalism. In the framework, we introduce a reusable set of three coherent metamodels for creating HW-SW co-design models concerning SDF graphs, hardware platforms and allocation of SDF tasks to hardware. The framework also contains model transformations that cast these models into priced timed-automata models, the input language of the well-known model checker UPPAAL CORA. We demonstrate how our framework satisfies the requirements of modularity, extensibility and interoperability in an industrial case study.

## 1 Introduction

Hardware-software (HW-SW) co-design is an engineering practice that allows to meet system-level objectives by exploiting the synergy of hardware and software through their simultaneous design. For instance, HW-SW co-design allows exploring design alternatives, and helps to improve the development cost and time-to-market. However, current tools and approaches for HW-SW co-design have difficulties coping with the concurrency and increasing complexity of modern-day systems. As a result, the time and effort needed for modeling and validating such designs are negatively affected. In fact, it has been widely recognized that a HW-SW co-design approach must have the following features [4, 10, 12, 15] all of which are not satisfied by the current HW-SW co-designing approaches:

- *Modularity* [4, 10]: The modeling approach should separate different aspects — such as hardware, software and their mappings — to keep their various concerns modular. This allows convenient exploration of design alternatives

concerning hardware and software. Modules targeting different concerns are better maintainable and reusable.

- *Extensibility* [10]: The HW-SW co-designing approach should have convenient extension mechanisms allowing rapid implementation of possible future requirements. This reduces not only the development cost of new products, but also their time-to-market.
- *Interoperability* [4,15]: HW-SW co-designing often involves tools serving different purposes, such as model designing, simulating, integrating etc. The HW-SW modeling approach should support interoperability between these tools, enabling system designers to explore design alternatives rapidly.

Model-Driven Engineering (MDE) is an approach that helps to fulfill the aforementioned requirements [24]. In MDE, the important concepts of the target domain are formally captured in a so-called, *metamodel*. Separate metamodels for the domains of interest help to keep the design modular. All models are instances of a metamodel, or possibly an integrated set of metamodels. Moreover, models can be transformed to the other via *model transformations*, defined at the metamodel level.

In this paper, we model software applications as Synchronous Dataflow (SDF) graphs [16] which are partitioned into tasks, with inter-task dependencies. SDF graphs are well-known computational models for real-time streaming and dataflow applications. This paper presents a novel HW-SW co-design framework based on the principles of MDE. Our framework allows model-driven HW-SW co-designing of SDF applications mapped on multiprocessor hardware platforms, and generate energy-optimal schedules for these SDF applications. To generate energy-optimal schedules, our framework transforms the co-designed SDF and hardware models to priced timed-automata models. The problem of finding energy-optimal schedules (while satisfying minimal throughput requirements) is encoded as an optimization problem, defined as a reachability property over priced timed-automata models. The property is then checked by the model checker UPPAAL CORA [5] that generates an energy-optimal schedule.

Our framework consists of three metamodels: (1) a metamodel for SDF graphs; (2) a metamodel for Platform Application Models (PAMs), which describe the processor types and their power levels, and the cost of switching between the power levels; and (3) a metamodel for expressing potential allocations of the tasks in an SDF graph to the processor types in a PAM. As mentioned earlier, our framework considers the model checker UPPAAL CORA for generating energy-optimal schedules. Therefore, for supporting the generation of UPPAAL CORA models, we also use an existing UPPAAL metamodel developed at the University of Paderborn [1]. The models conforming to three metamodels explained earlier, are transformed to UPPAAL CORA models automatically via *model transformations* in the framework. We have already described our method of using priced timed-automata for the purpose of energy optimization in [2], and therefore out of scope of this paper. Rather the novelty of this paper is the design prospects of using MDE. Later in Section 5, we demonstrate our framework as

an evidence to show the benefits of MDE namely, modularity, extensibility, and interoperability. The main contributions of this paper are as follows:

- We introduce the insights of state-of-the-art model-driven engineering approach into the embedded systems community, in particular for the domain of HW-SW co-design.
- We propose a reusable set of three coherent, extensible metamodels for HW-SW co-design.<sup>1</sup>
- We define and apply model transformations from the dataflow domain to the model-checking domain, obtaining an automated tool to compute energy-optimal schedules for dataflow applications.
- We demonstrate that our fully automated framework provides modularity, extensibility and interoperability between tools, via an industrial case study.

The rest of the paper is structured as follows: Section 2 provides the related work. Section 3 gives an overview of our framework and Section 4 describes the framework components in detail. Section 5 evaluates our framework using an industrial case study, and Section 6 concludes the paper.

The extended version of this paper can be found in the technical report [3].

## 2 Related Work

There exists a plethora of commercial and academic tools for HW-SW co-designing [4, 6, 9, 10, 12, 15]. Here, we only present the closest studies.

The state-of-the-art toolsets in the realm of HW-SW co-design of dataflow applications, are Octopus [4] and Ptolemy [15]. The Octopus toolset [4], in comparison to our approach, does not consider any metamodels. Furthermore, this toolset uses Java libraries for model transformation. Rather than Java, which is a general-purpose language, we use ETL that is specifically designed as a domain-specific language for model transformations. The lack of metamodels and model transformation language cause challenges in extensibility and maintainability, which are in fact stated as a future directions of the work in [4]. Ptolemy [15] is another well-known toolset for supporting HW-SW co-design of dataflow applications. However, similar to Octopus, it is not based on MDE which poses challenges in reusability and maintainability.

The closest works to ours are presented in [6] and [9]. Both of these papers utilize MDE techniques for HW-SW co-designing of embedded systems. In contrast to our work, these papers consider generic software applications and hardware models. We, on the other hand, analyze real-life software applications and hardware models enabled with power management.

To the best of our knowledge, this paper presents the first model-driven HW-SW co-design framework for dataflow applications that provides modularity, interoperability, and extensibility.

<sup>1</sup> All metamodels, model transformations, and case studies discussed in this paper can be found at <https://github.com/utwente-fmt/COMET>. An instruction manual for replicating the experiments is also given in this repository.

### 3 The Model-Driven Framework

#### 3.1 Model-Driven Engineering

Models are powerful tools to express behavior, structure and other properties in many domains such as mathematics, engineering, and other natural sciences. Model-Driven Engineering (MDE) is a software engineering approach that considers models not only as documentation, but also adopts them as the basic abstraction to be used throughout all engineering disciplines and in any application domain [7]. The models in MDE are closer to some particular domain concepts rather than the computing concepts. These models are considered equal to the code since they are formally defined and have execution semantics.

To define models, we need to specify their language as a model of these models at a more abstract level that is so-called *metamodels*. In their common use, metamodels capture the concepts of some domain or application; and define the permitted structure and behaviour, to which models must adhere. Therefore, metamodels describe the syntax of models [18].

MDE allows interoperability between different domains (and tools in these domains) via *model transformations*. Models transformations satisfy interoperability and furthermore save effort and reduce errors by automating the model derivation and modification process.

#### 3.2 Overview of Model-Driven Framework

Figure 1 shows the detailed overview of our framework. The HW-SW co-design of the application consists of the first four steps:

- In step 1, an SDF model of the software application is created using the SDF<sup>3</sup> tool [21] in an XML format specific to the tool.
- In step 2, the SDF model is transformed to an SDF model that conforms to the metamodel we defined for SDF graphs.
- In step 3, the hardware platform model is created using *PAM Visual Editor* that is a graphical editor for specifying Platform Application Models (PAMs). This model conforms to the PAM metamodel we defined for PAMs.
- In step 4, an allocation model is created for specifying the mapping of the tasks in the SDF model to the processor types in the PAM.

The analysis of the co-design for energy-optimal schedules is conducted using the UPPAAL CORA model checker. This is achieved in the last three steps:

- In step 5, the co-design is transformed to a priced timed-automata model that conforms to the UPPAAL metamodel.
- In step 6, the priced-timed-automata model is transformed to the format accepted by the model checker.
- In step 7, we analyze the resulting model to compute the energy-optimal schedule using the UPPAAL CORA model checker.




Fig. 1: Detailed overview of our framework. The elements with dark background color represent the new contributions.

Although the steps in Figure 1 show a general guideline for a HW-SW co-design of a system from scratch, a different strategy can be adopted according to the requirements of the system design. For example, if a system designer needs to analyze how a software application runs on various hardware platforms, s/he can create an SDF model by follow steps 1 and 2 only once and then create several PAM models by conducting step 3 multiple times.

Detailed explanation of the framework are given in Section 4.

### 3.3 Tooling Choices

To realize the model-driven approach, we have created metamodels using ECore in Eclipse Modeling Framework (EMF) [19]. EMF provides a plethora of plugins

to support various functionalities, such as querying, validation, and transformation of EMF models. For instance, using the EuGENia plugin [14], we have created PAM Visual Editor based on Graphical Editing Framework (GMF).

The model transformations have been implemented using Epsilon Transformation Language (ETL) [13], which is one of the domain-specific languages provided by the Epsilon framework. ETL supports many input-to-many output model transformations; it also allows the users to inherit, import and reuse other Epsilon modules in the transformations.

## 4 Details of the Model-Driven Framework

This section presents our concrete instantiation of the model-driven framework by describing our modeling choices in some detail. We recall the formal (mathematical) definitions of the domain concepts and discuss how we have chosen to translate them to metamodel elements.

### 4.1 SDF Graphs

Typically, real-time streaming applications execute a set of periodic tasks, which consume and produce a fixed amount of data. Such applications are naturally modeled as SDF graphs.

**Definition and Metamodel** An SDF graph is a directed, connected graph in which tasks are represented by *actors*. Actors communicate with each other via streams of data elements, represented by *channels* (the streams) that carry *tokens* (an abstraction of the data elements). Each channel  $(a, b, p, q)$  connects a producer actor  $a$  to a consumer actor  $b$ , and specifies *production* and *consumption* rates  $p$  and  $q$ , respectively, both given by integer values. The execution of an actor  $a$  is known as the *firing* of  $a$ ; as a result,  $q$  tokens are removed from any channel  $(b, a, p, q)$  of which  $a$  is the consumer, and  $p$  tokens are added to any channel  $(a, b, p, q)$  of which  $a$  is the producer. Formally:

**Definition 1.** *An SDF graph is a tuple  $G = (A, D, \text{Tok}_0)$  where  $A$  is a finite set of actors,  $D \subseteq A^2 \times \mathbb{N}^2$  is a finite set of channels, and  $\text{Tok}_0 : D \rightarrow \mathbb{N}$  denotes the initial number of tokens on each channel.*

Some notation: given an SDF graph  $G$  as above, the sets of *input* and *output channels* of an actor  $a \in A$  are defined respectively as  $\text{In}(a) = \{(b, a, p, q) \in D \mid b \in A, p, q \in \mathbb{N}\}$  and  $\text{Out}(a) = \{(a, b, p, q) \in D \mid b \in A, p, q \in \mathbb{N}\}$ .

*Example 1.* Figure 2 shows the SDF graph of the Viola-Jones face detector [23], kindly provided by the company Recore Systems, that is used in their face recognition system. The SDF graph contains seven actors (*im\_read*, *dupl.im*, *integral*, *haar\_det*, *haar\_scal*, *col\_obj* and *grp\_rect*) representing the tasks performed in face detection. For example, *im\_read* captures the scene containing one or more faces, *haar\_det* detects the regularities in the human face called *Haar features*, and *grp\_rect* groups the rectangles having similar Haar features.




Fig. 2: SDF graph of Viola-Jones face detector




Fig. 3: SDF Metamodel

The *SDF Metamodel* capturing the concepts of Definition 1 is shown in Figure 3. Recall that an SDF graph is a tuple  $G = (A, D, \text{Tok}_0)$ .

- *SynchronousDataFlowRoot* is the root of a model, in which everything else is contained; it corresponds to  $G$ .
- *Actor* corresponds to the set  $A$ ; the associations *incomingChannels* and *outgoingChannels* represent the derived functions *In* and *Out* from  $A$  to sets of channels.
- *Channel* corresponds to the set  $D$ . The 4-tuples  $(a, b, p, q) \in D$  are represented in the metamodel by the *source* and *target* associations (for  $a$  and  $b$ ), respectively the *sourceRate* and *targetRate* attributes (for  $p$  and  $q$ ). *initialTokens* represents the function  $\text{Tok}_0$ ; thus, it has been modeled as an attribute of *Channel*, rather than as a separate function.

With respect to the mathematical definition, there are two differences: (i) whereas a channel  $(a, b, p, q)$  is completely determined by its constituent values, due to the nature of metamodels a *Channel* has its own identity (and so conceivably there could be two *Channels* with the same 4-tuple of values), which can not occur in the mathematical set up in Definition 1; (ii) the function  $\text{Tok}_0$  has been combined with *Channel*. This removes some of the modularity of the mathematical model, at the benefit of simplicity.

**Model Creation** In our framework, SDF models are created in steps 1 and 2 of Figure 1. The starting point is an SDF graph created using the well-known open-source SDF<sup>3</sup> tool [21] (step 1). This tool produces output in the form of an XML document, adhering to its own schema (fixed in an XSD). To bring such documents into our framework, we have defined an SDF<sup>3</sup>-to-SDF *Transformation* which produces models conforming to the SDF metamodel of Figure 3. The transformation definition involves a systematic mapping of the SDF<sup>3</sup> concepts to our SDF metamodel concepts.

## 4.2 Platform Application Models

A Platform Application Model (PAM) models the multi-processor platform to which the application, modeled as an SDF graph, is mapped. Our PAMs support several features, including (1) heterogeneity, i.e., there can be multiple processors with different types, (2) a partitioning of the processors in voltage/frequency islands, (3) frequency levels each processor can run on, (4) power consumed by a processor at a certain frequency, both when in use and when idle, and (5) power overhead required to switch between frequency levels.

### Definition and Metamodel

**Definition 2.** *Given an SDF graph  $G = (A, D, \text{Tok}_0)$  with a set of actors  $A$ , a platform application model (PAM) is a tuple  $\mathcal{P} = (\Pi, \zeta, F, P_{idle}, P_{occ}, P_{tr}, \tau_{act})$  consisting of*

- a finite set of processors  $\Pi = \{\pi_1, \dots, \pi_n\}$ . We assume that  $\Pi$  is partitioned into disjoint blocks of voltage/frequency islands (VFIs) such that  $\bigcup \Pi_i = \Pi$ , and  $\Pi_i \cap \Pi_j = \emptyset$  for  $i \neq j$ ,
- a function  $\zeta : \Pi \rightarrow 2^A$  indicating which processors can handle which actors,
- a finite set  $F = \{f_1, \dots, f_m\}$  of discrete frequencies available to all processors,
- a function  $P_{occ} : \Pi \times F \rightarrow \mathbb{N}$  denoting the power consumption (static plus dynamic) of a processor operating at a certain frequency  $f \in F$  in the operating state,
- a function  $P_{idle} : \Pi \times F \rightarrow \mathbb{N}$  denoting the power consumption (static) of a processor operating at a certain frequency  $f \in F$  in the idle state,
- a partial function  $P_{tr} : \Pi \times F^2 \rightarrow \mathbb{N}$  denoting the transition overhead between frequencies for each processor  $\pi \in \Pi$ , and
- a function  $\tau_{act} : A \times F \rightarrow \mathbb{N}_{\geq 1}$  denoting the actual execution time of each actor (in  $A$ ) mapped to a processor at a certain frequency level (in  $F$ ).

*Example 2.* Exynos 4210 is a state-of-the-art processor used in high-end platforms such as Samsung Galaxy Note, SII, etc. Table 1 shows three frequencies (MHz)  $\{f_1, f_2, f_3\} \in F$  [17] and corresponding experimental power consumption. We assume that our PAM contains four Exynos 4210 processors, i.e.,  $\Pi = \{\pi_1, \pi_2, \pi_3, \pi_4\}$ . The processors are partitioned into 2 VFIs, i.e.,  $\Pi_1 = \{\pi_1, \pi_2\}$  and  $\Pi_2 = \{\pi_3, \pi_4\}$ . We assume that the power overhead (W) of all  $\pi \in \Pi$  is,  $P_{tr}(\pi, f_1, f_2) = P_{tr}(\pi, f_2, f_3) = 0.2$  and  $P_{tr}(\pi, f_3, f_2) = P_{tr}(\pi, f_2, f_1) = 0.1$ .

No.	Frequency(MHz)	$P_{idle}$ (W)	$P_{occ}$ (W)
1	1400	0.4	4.6
2	1222	0.3	3.2
3	1033	0.1	1.8

Table 1: Example platform description




Fig. 4: PAM Metamodel

The *PAM Metamodel* capturing most of the concepts of Definition 2 is shown in Figure 4. A brief explanation can be given as follows:

- *PlatformApplicationModelRoot* stands for the PAM as a whole.
- *ProcessorType* collects the characteristics of a set of processors. In the meta-model, the power and frequency characteristics of a processor are associated with its *type*, creating a reusable layer of indirection with respect to the mathematical model.
- *Processor* stands for the elements of  $\Pi$ . Each *Processor* has a *type* association to the corresponding *ProcessorType*.
- *VoltageFrequencyIsland* stands for the clusters  $\Pi_i$  in the VFI partitioning of  $\Pi$ . The element-of relationship between a processor and its VFI is captured by the (opposite) *island* and *processors* associations.
- *ProcessorState* associates the working/idle state of a processor (type) (the boolean *isWorking* attribute), combined with a *frequency* level, to a *powerConsumption* value. This encodes the  $P_{occ}$  and  $P_{idle}$  functions of the mathematical definition.
- *ProcessorStateChange* encodes the  $P_{tr}$  function of the definition: each instance associates a *powerCost* with a certain pair of *source* and *target* *ProcessorStates*.

In a major change with respect to the mathematical definition, we have chosen not to include the  $\zeta$  and  $\tau_{act}$  functions in the PAM, but to isolate them in a separate allocation model. This enhances the modularity of the modeling




Fig. 5: Allocation Metamodel. *Capability* refers to *Actor* in the SDF metamodel and *ProcessorState* in the PAM metamodel.

framework. Apart from this change, all elements of Definition 2 are clearly recognizable in the metamodel, though sometimes encoded in a different manner. In particular, we have introduced the processor types as an intermediate level to enhance modularity;  $P_{occ}$  and  $P_{idle}$  are combined in *ProcessorState*; and  $P_{tr}$  is encoded as *ProcessorStateChange*.

**Model Creation** The creation of PAMs corresponds to step 3 in Figure 1. Although EMF provides a default tree-based model editor, we have built PAM Visual Editor, a domain-specific visual editor for PAMs, by benefiting from state-of-the-art MDE techniques. To build PAM Visual Editor, we have used EuGENia, which can automatically generate a visual editor from an annotated ECore metamodel. We show an example PAM created using this visual editor in Section 5.

### 4.3 Allocation Models

In a heterogeneous system, the freedom of assigning actors  $a \in A$  to processors  $\pi \in \Pi$  is constrained by which processors can be utilized to execute a particular actor. Thus, in order to run an SDF model on a PAM, we need to know (1) which SDF actors can be run on which processors of the PAM and (2) what their execution times are at given frequencies. This information is encoded in an allocation model, which relates both the SDF and PAM models. Allocation models conform to *Allocation Metamodel* that we define to represent this concern.

The information related to allocation concern is a part of Definition 2, but we have chosen to define it as a separate *Allocation Metamodel* for the sake of modularity to make the PAM metamodel independent of the SDF metamodel.

**Definition and Metamodel** The information to be represented in the Allocation metamodel consists of the  $\zeta$  and  $\tau_{act}$  functions of Definition 2. The *Allocation Metamodel* is shown in Figure 5. It contains:

- *AllocationRoot*, which stands for the combined allocation functions  $\zeta$  and  $\tau_{act}$  of Definition 2.
- *Capability*, following  $\tau_{act} : A \times F \rightarrow \mathbb{N}_{>1}$  in Definition 2, refers to *Actor* in the SDF metamodel, and *ProcessorState* (defining the frequency of the processor)

in the PAM metamodel, and yields the time needed to execute the actor at the processor state. At the same time, *ProcessorState* also encodes which processor type an actor can be executed on.

The metamodel is in fact more expressive than the mathematical definition: for instance, the execution time of an actor is not constrained to be always the same for a given frequency level; instead, it may also depend on the processor type.

**Model Creation** The creation of Allocation models corresponds to step 4 in Figure 1. It is supported out-of-the-box via the default tree-based model editor provided by EMF.

#### 4.4 Common Metamodel

In addition to those discussed above, Figure 1 also shows an element called *Common Metamodel*. This demonstrates an MDE technique for reuse: this metamodel defines the general concept of *Identifiable*, which has a string-valued *identifier* attribute; *Actors* and *ProcessorTypes* are subtypes of *Identifiable* and thereby inherit this feature. Whenever (during extension of the framework) additional reusable concepts are introduced, these can be added to the common metamodel.

#### 4.5 Priced Timed-Automata Models

As mentioned earlier, we use priced timed-automata for energy optimization. Once the SDF, PAM and allocation models are available, one can generate the priced timed-automata model using the *Co-Design-to-UPPAAL Transformation* and successive model-to-text transformation. These correspond to steps 5 and 6 in Figure 1. In Step 7 in Figure 1, the energy-optimal schedule of an SDF graph can be generated. The ideas in [2] related to how the entities in SDF, PAM and allocation models can be mapped to priced timed-automata models and how the energy-optimal schedule can be calculated are reused in steps 5 and 7.

## 5 Case Study and Evaluation

In this section, we show the effectiveness of our framework for HW-SW co-design by applying it on a case study. We also demonstrate how our framework satisfies the features stated in Section 1, namely: (1) modularity, (2) extensibility, and (3) interoperability. We also evaluate the timing performance of our framework with the help of some other case studies.

### 5.1 Case Study

As a case study, we consider the dataflow application of the Viola-Jones face detector in Example 1 mapped to a platform with 4 processors of the type Exynos 4210 in Example 2.




Fig. 6: Energy-optimal schedule on four Exynos 4210 processors. The white, red, and green blocks denote the frequencies  $f_1$ ,  $f_2$ , and  $f_3$  respectively.

Following step 1 in Figure 1, we have created the SDF graph of the Viola-Jones face detection using SDF<sup>3</sup>. This SDF graph was already given in Figure 2. In step 2, we apply SDF<sup>3</sup>-to-SDF transformation to generate the SDF model conforming to our metamodel for the SDF graph.

In step 3, we create the PAM using the visual editor, as described in Section 4.2. The details of this PAM can be found in our technical report [3].

After we have the PAM and SDF models, we create the allocation model that assigns the actors to the processor states with their execution times in step 4.

Once we have the SDF, PAM, and allocation models, we apply the Co-Design-to-UPPAAL transformation in step 5 and the model-to-text transformation in step 6 to generate the priced-timed-automata model that is compatible with UPPAAL CORA. In step 7, we follow the approach presented in [2] to compute the energy-optimal schedule for some given throughput requirements. Figure 6 shows the energy-optimal schedule, for the time per graph iteration constraint of 650 ms for our example. The schedule shows the execution order of the actors at the specific frequency and processors.

## 5.2 Evaluation

*a) Modularity:* To show the modularity of our framework with a concrete example, let us consider the following scenario: We want to analyze the energy consumption of the face detection application on a different hardware platform, viz., Intel Core2 Duo E6850.

Table 2 shows the frequencies and corresponding power consumption of this new processor [17]. For this scenario, we only change the processor type while keeping the number of processors and VFI distributions the same. Now, all we have to do is to develop a new PAM corresponding to this new platform specification and

generate the corresponding priced timed-automata model. We reuse the existing SDF model of the application without making any modifications. Using the framework, we derive the energy-optimal schedule shown in Figure 7 on the new hardware platform, for time per graph iteration constraint of 650 ms.

No.	Freq.	$P_{idle}$	$P_{occ}$
1	3006	0.4	55
2	2338	0.3	34
3	1776	0.1	22

Table 2: Platform description

$\pi 1$	idle					haar_scal	idle		
$\pi 2$	idle					haar_scal	idle		
$\pi 3$	idle					haar_scal	idle		
$\pi 4$	idle	im_read	dupl_im	integral	haar_det	haar_scal x 22	haar_scal	col_obj	grp_rect

time

Fig. 7: Energy-optimal schedule on four Intel Core 2 Duo E6850 processors

*b) Extensibility:* One can extend the framework using the following mechanisms: introducing new models with related metamodels and new transformations; extending existing transformations or metamodels.

As an example, suppose we want to extend our platform models with the concept of “battery”. The current version of the platform assumes energy source to be ideal such that the system never runs out of energy. However, we want to include batteries in our HW-SW co-design as resources. This extension can be achieved through the following steps:

1. *Adding a metamodel for batteries:* An example battery metamodel is shown in Figure 8. This metamodel defines the number of batteries in the system with their initial capacities (Coulomb).

2. *Extending the Co-Design-to-UPPAAL model transformation:* In order to include the battery model in the scheduling analysis, we have to transform the concepts in the battery metamodel to the concepts in the priced timed-automata domain. This is achieved through extending the Co-Design-to-UPPAAL model transformation in step 5. The extension to the transformation will create the dependency of the processors on the batteries, in such a way that the processors consume charge from these batteries. The extension will further generate separate templates for battery and battery scheduler. The template of the battery keeps track of the current charge. When the battery gets empty, it informs the battery scheduler via synchronization. In that case, the battery scheduler template activates the next available battery. When all batteries are out-of-charge, the processors cannot run anymore. We can extend the Co-Design-to-UPPAAL model transformation without modifying it since ETL allows to extend and reuse transformation modules. Please note that the model-to-text transformation in step 6 stay unaffected by this extension.

*c) Interoperability:* In our framework, we utilize SDF<sup>3</sup> for creating SDF graphs and UPPAAL CORA for deriving energy-optimal schedule. To automatically generate UPPAAL CORA models from SDF<sup>3</sup> models, we have implemented model transformations in our framework, thus providing interoperability.

### 5.3 Timing Performance

To determine the timing performance of our framework, we consider five real-life case studies namely, a Viola-Jones face detector in Figure 2, a MPEG-4




Fig. 8: Battery Metamodel

Decoder [22], an Audio Echo Canceller [11], an MP3 Decoder [20], and an MP3 playback application [25]. We also used an artificial bipartite SDF graph [8] with 4 actors. We assume that these case studies are mapped on Exynos 4210 processors having two frequencies.

We examine the timing performance of our framework in two parts: the first part is the timing performance of our framework, i.e., cumulative computation time of steps 2 (SDF<sup>3</sup>-to-SDF transformation), 5 and 6 (Co-Design-to-UPPAAL and model-to-text transformations). The second part is the timing performance of obtaining the optimal schedule via UPPAAL CORA model checker, i.e., step 7.

The case studies show that the time step 2, 5 and 6 take in total increase insignificantly as the number of available processor increases. This is due to the slight increase in the model size with the addition of processors. For step 7, the time required to complete increases exponentially as the number of processor increases, which is because of the fact that the size of the state-space created by the model checker increases exponentially with the size of the model itself.

The details of these case studies can be found in our technical report [3].

## 6 Conclusions and Future Work

In this paper, we have presented a model-driven framework for HW-SW co-design of dataflow applications. In our framework, we have proposed a reusable set of three coherent metamodels for HW-SW co-design domain. To provide interoperability among domains, we have defined a reusable set of extensible model transformations. We have demonstrated that our framework satisfies the modularity, extensibility and interoperability requirements with a case study.

As future direction of our work, we plan to extend our framework with other analysis techniques such as simulation and automated HW-SW partitioning. We also plan to add code generation functionality to our framework.

**Acknowledgements.** This research is funded by the EU FP7 project SENSATION (318490) and NWO project BEATS (612.001.303). The authors are grateful to Kim Sunesen from Recore Systems B.V. for providing the case study.

## References

1. Software Engineering Group, University of Paderborn. <https://www.hni.uni-paderborn.de/en/software-engineering/>. Accessed: 2016-01-14.
2. W. Ahmad, P.K.F. Hölzenspies, M.I.A. Stoelinga, and J. van de Pol. Green computing: Power optimisation of VFI-based real-time multiprocessor dataflow applications. In *DSD'15*, pages 271–275, Aug 2015.
3. W. Ahmad, B. M. Yildiz, A. Rensink, and M. Stoelinga. A model-driven framework for hardware-software co-design of dataflow applications. Technical Report TR-CTIT-16-09, University of Twente, 2016.
4. T. Basten, R. Hamberg, F. Reckers, and J. Verriet. *Model-Based Design of Adaptive Embedded Systems*. Springer Publishing Company, 2013.
5. G. Behrmann, K. G. Larsen, and J. I. Rasmussen. Optimal scheduling using priced timed automata. *SIGMETRICS Perform. Eval. Rev.*, 32(4):34–40, Mar. 2005.

6. L. Bondé, C. Dumoulin, and J.-L. Dekeyser. *Advances in Design and Specification Languages for SoCs*, chapter Metamodels and MDA Transformations for Embedded Systems, pages 89–105. 2005.
7. A. R. da Silva. Model-driven engineering: A survey supported by the unified conceptual model. *Computer Languages, Systems & Structures*, 43:139–155, 2015.
8. M. Geilen, T. Basten, and E. Stuijk. Minimising buffer requirements of synchronous dataflow graphs with model checking. In *DAC '05*, pages 819–824, 2005.
9. I. Gray, N. Matragkas, N. Audsley, L. Indrusiak, D. Kolovos, and R. Paige. Model-based hardware generation and programming - the MADES approach. In *ISORCW'11*, pages 88–96, March 2011.
10. K. Grüttner, P. A. Hartmann, K. Hylla, S. Rosinger, W. Nebel, F. Herrera, E. Villar, C. Brandolese, W. Fornaciari, G. Palermo, C. Ykman-Couvreur, D. Quaglia, F. Ferrero, and R. Valencia. The COMPLEX reference framework for HW/SW co-design and power management supporting platform-based design-space exploration. *Microprocessors and Microsystems - Embedded Hardware Design*, 37:966–980, 2013.
11. J. P. Hausmans, S. J. Geuns, M. H. Wiggers, and M. J. Bekooij. Compositional temporal analysis model for incremental hard real-time system design. In *EMSOFT '12*, pages 185–194, 2012.
12. P. Herber and S. Glesner. A HW/SW co-verification framework for SystemC. *ACM TECS*, 12(1s):61:1–61:23, Mar. 2013.
13. D. S. Kolovos, R. F. Paige, and F. A. Polack. The Epsilon transformation language. In *ICMT '08*, pages 46–60, 2008.
14. D. S. Kolovos, L. M. Rose, R. F. Paige, and F. A. C. Polack. Raising the level of abstraction in the development of GMF-based graphical model editors. In *MISE '09*, pages 13–19, 2009.
15. E. A. Lee. Embedded software. *Advances in Computers*, 56:56–97.
16. E. A. Lee and D. G. Messerschmitt. Synchronous data flow: Describing signal processing algorithm for parallel computation. In *COMPCON'87*, pages 310–315, 1987.
17. S. Park, J. Park, D. Shin, Y. Wang, Q. Xie, M. Pedram, and N. Chang. Accurate modeling of the delay and energy overhead of dynamic voltage and frequency scaling in modern microprocessors. *TCAD*, May 2013.
18. J. Sprinkle, B. Rumpe, H. Vangheluwe, and G. Karsai. Metamodelling. In *Model-Based Engineering of Embedded Real-Time Systems*, pages 57–76. Springer, 2010.
19. D. Steinberg, F. Budinsky, E. Merks, and M. Paternostro. *EMF: Eclipse Modeling Framework*. Pearson Education, 2008.
20. S. Stuijk. *Predictable Mapping of Streaming Applications on Multiprocessors*. PhD thesis, 2007.
21. S. Stuijk, M. Geilen, and T. Basten. SDF<sup>3</sup>: SDF For Free. In *ACSD'06*, pages 276–278, June 2006.
22. B. Theelen, M. C. W. Geilen, T. Basten, J. P. M. Voeten, S. V. Gheorghita, and S. Stuijk. A scenario-aware data flow model for combined long-run average and worst-case performance analysis. In *MEMOCODE'06*, pages 185–194, 2006.
23. P. Viola and M. Jones. Rapid object detection using a boosted cascade of simple features. In *CVPR'01*, pages I–511–I–518 vol.1, 2001.
24. M. Völter, T. Stahl, J. Bettin, A. Haase, and S. Helsen. *Model-driven software development: technology, engineering, management*. 2013.
25. M. H. Wiggers. *Aperiodic multiprocessor scheduling for real-time stream processing applications*. PhD thesis, 2009.