

EPIPURE¹ : A Partitioning and CoDesign Framework For Reconfigurable Computing

Diguet J-Ph.¹, Le Moullec Y.², Bilavarn S.³, Ben Chehida K.⁴, Auguin M.⁴, Gamrat C.⁵

Gogniat G.¹, Philippe J-L.¹, Fornari X.⁶, Fouillart A-M.⁷, Kajfasz P.⁷

(1) LESTER CNRS FRE 2734, Lorient, FR; (2) CISS, Univ. Aalborg, DK; (3) EPFL/ITS/LTS2, Lausanne, CH; (4) I3S, CNRS UMR 6070, Sophia-Antipolis, FR ; (5) CEA/LIST, Saclay, FR; (6) Esterel Technologies, Sophia-Antipolis, FR; (7) Thales LJS, Colombes, FR

ABSTRACT

This paper presents a new global design methodology capable to bridge the gap between an abstract specification level and a heterogeneous reconfigurable architecture level. The Epicure contribution is the result of a joint study on abstraction/refinement methods and a smart reconfigurable architecture within the formal Esterel design tools suite. The original points of this work are : i) a generic HW/SW interface model, ii) a specification methodology that handles the control, includes efficient verification and HW/SW synthesis capabilities, iii) a method for parallelism exploration based on abstract resources/performance estimation expressed in terms of area/delay tradeoffs, iv) a HW/SW partitioning approach that refines the specification into explicit HW configurations and the associated SW control. The Epicure framework shows how a cooperation of complementary methodologies and CAD tools associated with a relevant architecture can significantly improve the designer productivity, especially in the context of reconfigurable architectures.

Keywords

Smart HW/SW interface, Formal programming model, Design space exploration, Reconfigurable computing, Hardware/Software partitioning, Parallelism exhibition and functional estimation, design productivity.

1. INTRODUCTION

Complex applications combine control and data processing that are generally tightly coupled. The design of an efficient architecture that is a correct implementation of the application is a more and more tedious work. This design complexity is due for example to the Hw/Sw separation realized very soon in the design cycle, generally at the specification step, and which lasts over numerous refinement steps, until cosimulation. Approaches based on formal models can be an attractive way to deal with specification validation but there is still a lack of methodologies able to derive straightly from a formal specification a Hw/Sw architecture that realizes the specification. One major difficulty results from the architecture itself that combines different components, each one controlled in a specific manner that is not related with the control semantic in the formal model.

System-level design-space exploration and platform-based design have been addressed by a certain number of research projects. In [9] definitions and a methodology for platform-based design are presented. The main goal of this methodology is to improve the design re-usability and regularity in order to cope with the increasing development costs. In SPADE [10] a modeling and simulation environment provides a design-space exploration at several abstraction levels. The methodology consists in mapping applications described with Kahn Process Networks onto an architecture model which is roughly and manually defined by the designer in a first approach. Both the application and architecture models are simulated. The performance of the system is provided by the simulation results and is then manually analyzed by the designer to perform HW/SW partitioning and iteratively refine the architecture. Platune [11], is a tool for exploring the configuration space of parameterized SoCs. The method enables all Pareto-optimal configurations in terms of execution time and power to be found. This is done

¹ This project was funded in part by a grant from the French ministry of research under the RNTL framework.

by considering the dependencies between the SoC parameters such as cache size and associativity, bus width, buffer sizes etc.

In our project (EPICURE) we investigated the way to reduce the gap between a formal specification model and the architectural specification level. To do this, we consider a common view of the software and the hardware in the architecture in order to lessen the difficulty to perform refinements on the formal specification while the semantic of the model is preserved. The common view of hardware and software is achieved in our approach by an intelligent interface (ICURE) between the software unit and a dynamically reconfigurable unit. Based on this abstract homogeneous architectural model of the heterogeneous architecture we have developed a design methodology based on the visual synchronous SSM model integrated in the Esterel Studio™ framework. The methodology consists of an abstraction tool that estimates implementation characteristics from a C-level description of a task and a partitioning refinement tool that realizes the allocation and the scheduling of the tasks according to the available resources in the architecture.

In the following we will first introduce the reconfigurable hardware/software architecture under assumption. In a second part we will describe the general design flow and methodology. Finally, a third part will illustrate the EPICURE framework on two application examples to show how the EPICURE methodology can be used efficiently.

2. "ICURE"-CENTRIC ARCHITECTURE

Reconfigurable computing has become now a very active and promising field of research investigation. Since the first attempt triggered by the introduction of SRAM based FPGAs in the late 80's, computer architects have made numerous advances for the conception and effective use of reconfigurable computers. Some of those first attempts focused on the system level and used off-the-shelf devices such as advanced FPGAs packed with microprocessors or complex function cores [1,2]. Other works strictly addressed the design aspect and have led to the definition of new architectures and families that could be realized with today CMOS or future nano technologies perspectives. All of those contributions would certainly result into well engineered high performance devices but could fall short when faced with real world applications, due to the lack of easy programming tools.

Among the various possibilities of reconfigurable hardware (RHW) architectures, we chose to focus on the interface that is needed to link a standard CPU and a generic reconfigurable processing unit (RPU). Although this might seem a little conservative from an architectural point of view, it is obviously the right place to introduce the necessary circuitry in order to abstract the CPU/RPU communication channel. The high level of abstraction in the CPU/RPU interface will ease the standardization of communications and development of protocols needed for the HW/SW partitioning tools. This way, the methodology and tools developed in the EPICURE project is also made independent from any specific reconfigurable technology.

The interface has been called ICURE, an acronym for "Intelligent Control Unit for Reconfigurable Element". It acts as a hardware abstraction layer in between the CPU and the RPU. Practically, ICURE main functions are to implement the link between the CPU/RPU and to manage the reconfiguration process in a smart way. The CPU could be any of RISC, VLIW, DSP, micro-controller, etc. and is connected to ICURE via a memory mapped register set through a standard address/data bus. The RPU is a generic reconfigurable "black box" with a number of bi-directional I/O ports and a reconfiguration link. Such a generic interface allows for the connection of any type of existing (FPGAs) or future reconfigurable device.

Reconfiguration of the RPU is driven by a dedicated management unit inside ICURE. Reconfiguration data are stored in a data structure called a context which encapsulates all the necessary bitstream, routing information and object code needed for proper configuration and execution of the RPU logic resources. Along with the configuration data, a context also contains the necessary information about how the CPU should use the RHW. In particular, the context indicates the number, types and sequences of parameters required for a successful

hardware function call. The reconfiguration bitstream and object code are synthesized and compiled using the appropriate back-end tools for the RPU and CPU targets.




Fig. 1 - The "ICURE" interface

The main tasks that are performed by the ICURE interface are:

- Take care of the reconfiguration phase control (Reconfiguration control). This is done through a simple automaton that implements the proper logic for sending reconfiguration data to the RPU. For example, when using standard “off the shelf” FPGA, this controller will take care of sending the bitstream to the FPGA.
- Handle hardware / software data routing and exchange. By programming a crossbar type resource, a convenient routing of signals is performed in between the RPU I/O ports (The hardware port) and the dual ported memory (the software port) for the selected context. In that respect, the ICURE interface allows for software data types (int, float, etc) to hardware data types (signals) seamless conversion. The routing table is changed each time a new context is loaded.
- Handle data synchronization between resources. The ICURE interface includes a dual ported memory that allows sharing small amount of data (32 kbits) in an asynchronous way between the CPU and RPU. For larger data transfers, an external access is provided to the ICURE interface through a DMA port.
- Take care of the context control (ICM Icure Context Management). Contexts for a given application are stored in a dedicated memory usually during an initialization phase. Upon request from the CPU, the context manager (ICM) will configure the required context. The ICM will first check if a context is already configured and not running and will eventually perform the configuration phase. Once a context configuration is initiated, the ICM extracts the three parts of the context data structure (HW data, I/O routing, SW code) and direct them to the corresponding ICURE resources (Reconf control, Routing, Shared CPU memory).

All the ICURE interface features are made available to the controlling CPU through a set of API functions that allow for a seamless use of the reconfigurable resources. Using a hardware function context is basically a two step process : In a first configuration phase, the CPU asks for a given context by calling a `load_context()` function. This configuration phase is entirely and automatically performed by the ICURE interface. Upon its completion, the CPU can enter the execution phase and use the desired hardware function through a simple call mechanism. To some extent, this concept is similar to the Set/Execute model proposed in [3].

The introduction of the ICURE interface in between the CPU and RPU along with its dedicated API, allows for a standardized mechanism for the control and access of the hardware reconfigurable tasks. The EPICURE partitioning framework and tools described hereafter will thus help in defining and switching contexts using the ICURE standardized approach using a two step approach:

1. Help to selecting the appropriate applications functions and tasks for use in a context.
2. Introduce context switching at the right places inside the application control execution graph based on application and/or user constraints.

3. SPECIFICATION/EXECUTION MODEL

3.1 Specifications

The model is specified using the Safe State Machines (SSM) formalism developed by Esterel Technologies [5]. SSM is a visual synchronous model, integrated and Esterel Studio™. SSM is the commercial version of SyncCharts, an academic synchronous graphical model designed in the mid-nineties [15] as a graphical notation for the Esterel language [16,17]. As such, SyncCharts was given mathematical semantics fully compatible with Esterel semantics. Like Esterel, SSMs are devoted to the programming of control-dominated software or hardware systems. Those systems are reactive, that is, they continuously react to the stimuli coming from their environment by sending back some other stimuli. They are purely input-driven: they react at apace imposed by their environment. A reactive application usually performs both data-handling and control-handling. Esterel and SSMs are imperative languages especially designed to deal with control-handling. In particular, they support hierarchical automata, parallelism and preemption mechanisms at any level of nesting.

The key point here is that the underlying semantics guarantees an unambiguous interpretation of an Esterel Studio™ design. Furthermore, thanks to this semantics, one can produce “Correct-by-Construction” code from the specification.

3.2 The Esterel Studio Approach

Esterel Studio is a design environment optimized for complex control-dominated IPs (such as DMAs, protocols, cache controllers, I/O subsystems, etc.). It is dedicated to the capture of formal design specifications using SSM (see. Fig.2), and enables the formal verification of properties very early in the design phase. Its use automates the production of synthesizable RTL code (VHDL and Verilog), both for prototyping and production purposes.

- Rigorous executable specifications, with powerful sequential control and preemption primitives, unlimited hierarchical design and synthesis flow.
- Generated VHDL, Verilog, C, C++, all share the same exact behavior. It is then possible to generate C code allowing interactive simulations with source code highlighting, which is very helpful during design phase. Then the user can produce a final HDL code, with same behavior. Automatic generation of VHDL and C testbenches with cycle accuracy check are also possible for non-regression testing.
- Co-simulation with all HDL simulators.
- Formal verification of properties and assertions: Thanks to its underlying model, it is also possible to apply formal verification on models. The designer defines a property and the Esterel Studio verifier can check if that property holds. The property is itself expressed using the Esterel Studio graphical or textual formalisms: the designer keeps the same environment to develop his/her designs and their properties. If the property fails, the tool produces a scenario that can be reused for validation purposes with the simulator.

3.3 Specifying Models for EPICURE

The EPICURE methodology proposes to describe the control of the application using the SSM, and to let the data processing part as calls to procedures within the design. The key point is to analyze such calls made during a computation cycle, in order to determine their final implementation. This analysis step results into a separation of two kind of procedures: those implemented in hardware, and those remaining in software. Note that the main

controller remains implemented as a software application. Currently there is no automatic way to define a split of the global controller into smaller ones.

To determine which procedure calls are made during a cycle, we use the OC intermediate format of the Esterel V5 compiler. This format describes the model as a flat automaton, where any internal communications in the original design have been resolved, and where parallelism and hierarchy disappear.

The format describes for each state, the transitions to the next states as a direct acyclic graph where leaves represent the next states. Along the paths of the graph, there are actions that are either typical procedure calls, or tests (input model tests, or internal data tests) that generate branchings to other graphs. From this graph, one can determine the calls that are performed, their order, and then analyze the effective cost of a transition. This is held by the partitioning process, see 4.3. The OC format contains all the procedures used by the model, which is used by the estimation phase to determine the procedures that have to be estimated. Figure 3 describes the compilation flow using OC (note this description is schematized: we will not enter into the Esterel compiler details).

The SSM description is first compiled into some internal format, possibly OC. At this stage, one can already make a verification of activities. Applying verification during the design phase is an interesting feature, because it may help to find bugs in the specification at early stages, and brings more confidence to the system under design.

Then, one can go through the code generation process. In particular, it's possible to produce a C code for interactive simulation. If the verification phase exhibits problems (a non verified property), then the verifier generates scenarii leading to counter examples. Those scenarii can be played interactively within the simulator.

Once the final design meets all the requirements, the final code can be generated. In our case, we generate the OC intermediate code in a first step. Since the Esterel compiler can read any of its internal format and starts the compilation process from its corresponding level, we can use an OC code properly modified outside from the compiler.




Figure 2 SSM Example during simulation




Figure 3 Compilation flow using OC

4. DESIGN FLOW

4.1 Overview

The Epicure framework is based on the ICURE and API models and on the three step design flow depicted in Fig.4. The first step is to specify the application. This step relies on an original execution model that enables to handle both control and data-flow parts of the specification in order to exhibit the reconfiguration and

acceleration opportunities. The specification strategy is guided by the ability of the designer to make a decision about the control/data-flow boundaries. A simple way is provided to guide this choice according to the application analysis. We will see in section 5 that these decisions can have a huge impact on the final results.

The specification is performed with Esterel Studio. At the lowest specification level, data-flow links are inserted through C procedure calls within the Esterel specification.




Fig.4 Epicure global design Flow

The resulting file is an OC format code. Then, the second step of the Epicure flow is carried out by the Design Trotter framework whose entry point is the list of C procedures from the OC code. The estimation step provides FPGA resources / cycle-budget trade-off curves. Thus the higher level of HW granularity is controlled by the designer by means of function selection and then lead to a set of functions candidates for the estimation flow. Based on the set of HW and SW solutions, the third step consists in selecting a pertinent implementation regarding design constraints such as performance, area and power, but also considering the reconfiguration opportunities. The partitioning solution is then given in the OC format. In case of a HW implementation, the C procedure call is replaced by a HW execution directive that is sent to the ICURE interface to control the FPGA context loading and HW/SW data exchanges. Finally, the OC to C conversion is made by Esterel Studio.

4.2 Estimation

4.2.1 Introduction

This work is based on the following observations: Firstly, the trade-off between energy savings, area use and real-time constraints, which is required for future embedded systems, imposes a very optimized use of the silicon die. Thus the necessary improvements can only be reached through an implementation of massive spatial parallelisms, a use of limited clock frequencies and the design of more or less dedicated hardware modules. This search of parallelism and dedicated hardware opportunities constitutes an important task of the design flow. This task is located before the complete definition of the target architecture and can be described as a system-level architectural exploration. Secondly, if we now observe that CAD tools for co-synthesis and co-simulation have reached a reasonable degree of maturity, system-level exploration tools still remain at a research stage. However, fast estimations based on largely abstract and uncompleted specifications are strongly required at a system level [6]. Most of the current tools provide an exploration based on a fixed library of pre-characterized functions. Such

libraries usually quantify a software implementation and one or two hardware solutions. Thus, the architectural exploration is bounded by the fixed granularity of functions and limited to a hardware/software partitioning. It appears that another stage has to be developed in order to really guide or parameterize some architectural choices. To cope with the architectural exploration process, we have implemented the estimation flow depicted in Fig.5-A. The different steps of the framework are detailed in the following sections : Specification, Characterization, Exploration and Physical Mapping.

4.2.2 Contribution

The problem of resource estimation can be divided in two steps. The first one is the design space exploration problem that focuses on the definition of pertinent architectural solutions, and the second one is the estimation of the mapping results of a given architectural solution on a target FPGA (see [7] for details). The originality of our approach is to combine both aspects within a single framework.

Previous methods and tools depend on a designer experience to define an initial architecture that will be further refined by mean of design-space exploration. The first step of our method consists in exploring and exploiting the parallelism potential of an application. Those exploration results are reported using trade-off curves, where a point corresponds to a potential architecture. This is a guidance tool for the designer that helps him to define the initial architecture. Therefore this approach is complementary with previous ones and helps to choose or build an architecture. The finally selected architecture can then be refined by existing co-design methods or by the next steps of the EPICURE flow. Our contribution focuses on the three following items:

Exhibition and exploitation of parallelism. Parallelism is a key parameter in the design of digital systems. It has a direct impact on several performance parameters such as execution time, energy consumption and area. Therefore we seek to explore the parallelism potential of an application in terms of a) type (data-transfer and data-processing), b) granularity level, and c) type (spatial, temporal). In our work these aspects are addressed thanks to our graph-based internal representation (HCDFG) and via the time constrained core of our estimator (used with several time constraints). This estimator rapidly provides dynamic exploration of an application by means of parallelism vs. delay trade-off curves in which a point corresponds to a potential architecture. The designer has then access to a set of solutions from which he can extract the most promising ones regarding the constraints. These solutions can then be refined and mapped to technology dependent architectural models using the next steps of our flow [7,8] or existing co-design tools;

Target architecture definition. At the system level, we assume the target architecture is not defined yet. At that level the designer considers algorithmic operators (resources executing the operations found in the specification) since the goal is to guide the definition of the architecture. After the function set pruning, the designer can use the exploration tool at the architectural level where algorithmic operators are associated to pre-synthesized arithmetic operators.

Impact of specification choices. Since the core of the estimation framework is based on a fast scheduler, the designer can evaluate very rapidly the impact of algorithmic transformations. This feature is very important since it enables the exploration of several specifications for a given application.

4.2.3 Specification

4.2.3.1 Definitions

In the Epicure project, each C function issued from the system specification is transformed into a Control and Data Flow Graph (HCDFG) at the beginning of the estimation step. There are three types of elementary (i.e., non hierarchical) nodes, the granularity of which depends on the granularity of the architectural model:

- A **processing node** (P) represents an arithmetic or logic operation, the granularity of the node depends on the architectural model : (ALU, MAC, +,-, etc.).
- A **memory node** represents a data-transfer (DT). The main node parameters are the transfer direction (read / write), the data format and the hierarchy level that can be fixed by the designer.
- A **conditional node** represents a test operation (if, case, loops, etc.)

Three types of oriented edges are used to indicate scheduling constraints :

- A **control dependency** indicates an order between operations without data-transfer, for instance a test operation must be scheduled before the mutual exclusive branches. The control dependency edges can also be used to impose a given order between independent operations or graphs with the intention of favoring resource optimization (data-reuse for instance).
- A **scalar data dependency** between two nodes A and B indicates that node A uses a scalar data produced by node B.
- A **multidimensional data dependency** is a data-dependency where the data produced is no more a scalar but an array. For instance such an edge is created between a loop CDFG is reading an array transformed by another loop CDFG.
- A **DFG** is a graph that contains only elementary memory and processing nodes. Namely it represents a sequence of non conditional instructions of the original C code.
- A **CDFG** is a graph that represents a test or a loop pattern with their associated DFGs.
- A **HCDFG** is a graph that contains only elementary conditional nodes, HCDFGs and CDFGs.

4.2.3.2 Graph creation rules

The decomposition principle is quite simple. The graph is traveled with a depth-first search algorithm. A H/CDFG is created when a conditional node is found in the next hierarchy level. When no more conditional nodes are found, a DFG is built. In order to ease the estimation process, CDFG patterns have been defined to identify rapidly "loop, if, etc." specific nodes. Another important point is that the model covers the processing complexity of the complete application. Thus, the computation of array indexes (address computation), conditional tests and loop index evolution are represented with DFGs. A HCDFG example is given in Fig. 6.

4.2.4 Characterization

The aim of metrics computation is to guide the designer in its choices of implementation using the functions characteristics. A first category of metrics can be computed before scheduling. The **MoM** metric indicates the ratio of data-transfer (global data) compared to the processing operations (local or reused data), so it indicates the bandwidth requirements of the graph. The **CoM** metric provides the ratio of test operations over all other kind of operations. A graph presenting a high CoM doesn't appear as a good candidate for hardware acceleration.

The second category of metrics is based on the asap-alap dates computed before DFG scheduling. The γ metric provides the average parallelism over the critical path. A high value indicates thean important speed up potentialspeed for a given graph (thus energy savings capabilities). The **DRM** is the aData Reuse Metric, namely the ratio between data transfers and computations. Compared to the MoM metric, the Data Reuse Metric is more accurate but also more complex. All metrics are first computed for DFGs. Then, they are extended to all granularity levels by mean of combinations, which rules are based on the patterns of Fig.5-C [5]. The HCDFG representation allows an easy implementation of various other metrics.

4.2.5 Exploration

The exploration is based on two sub-steps that are (C)DFG scheduling and (C)DFG combinations.

4.2.5.1 Adaptive DFG and CDFG-loop scheduling

Simple DFG : The estimation and exploration of the resource parallelism starts with the scheduling of DFGs. The objective is to quickly produce graph schedules for various cycle budgets in order to produce the parallelism/cycle budget tradeoffs illustrated in Fig.5-B. The range of delay varies from the minimum cycle budget (Δc) equal to the critical path to the delay required for a complete sequential execution (Δs).path We use a time-constrained list-scheduling heuristic where the number of resources of type i allocated is given by the lower bound : $N_{operations(i)} / N_{bcyles}$. This lower bound is recomputed online during the scheduling execution with the remaining operations and cycles in order to avoid final over-allocations. However, the scheduling policy is a key choice for resource minimization, it depends on the DFG orientation namely the critical resource (DT or P).

So the DRM metric is used to cope with the priority scheduling order: DT scheduling first, P scheduling first or simultaneous DT/P scheduling.

Loop CDFG : We consider loop unfolding when the time constraint Δ_i cannot be reached. The minimum unfolding factor α is computed with the maximum distance of inter-iteration dependency d_{max} (computed from the array index graphs) [12]. The loop core is duplicated (α times) and the asap dates of the duplicated graphs are shifted according to the value of d_{max} . Then the resulting DFG is scheduled.

4.2.5.2 CDFG hierarchical combinations

The extension to all granularity levels is obtained by a combination of the tradeoff curves previously computed for the DFGs and CDFG-loops of the function. The estimation process of a complete function is hierarchical. First, the lowest levels (DFGs) are processed by means of adaptive scheduling, OARC and draw pile. Then, the CDFGs are estimated using the loop and control patterns described previously. Finally, combination rules are applied in order to estimate sequential and parallel executions of graphs (and allow to estimate all hierarchy patterns). During this bottom-up approach, the combination order is guided by the criticality metric (cf. \ref{subsection:carac}) : when three or more elements have to be combined, the two most critical ones are combined first and so on. Moreover, we aim at producing very fast estimates to allow the exploration of large design-spaces. So, considering the great number of points in each trade-off curve, we cannot apply an exhaustive search of the Pareto points. Instead, we use a technique based on the peculiar points (PP) of the CDFGs trade-off curves. A peculiar point corresponds to a solution for which the cost in terms of resources decreases for a given time constraint. Note that the term resource is used for processing units, data transfers and local memories as well. Combinations are based on four algorithms according to the patterns of Fig.4-C. For each case, the combination rules are applied for processing and memory units. Briefly these rules follow the general following ideas :

- Sequential (HC)DFG : cycle budgets are added and maximum values are considered for resource estimation;
- Parallel (HC)DFG : maximum of cycle budgets are considered and resources are added;
- Exclusive (HC)DFG: the designer can select the worst case in terms of cycle budget and resource or select a finer solution based on branch probabilities issued from traces. In this last possibility has been used in the framework of the Epicure project.
- Loop (HC)DFG : Loop scheduling is only considered for the deeper nested loop, higher levels of nesting are executed sequentially.

Thus, the tradeoff curves are finally available from the DFG to the top (HC)DFG level representing the whole function. As a result we obtain the following results for each HCDFG hierarchy level and for all pertinent cycle budget :

- The number of processing operators used (MUL, ALU, MAC) depending on the association operation / operator defined in the architecture model;
- The size of local memories based on the combination of local memory sizes required for each DFG during the scheduling step.
- The total memory size required (MS).
- The number of simultaneous memory accesses (N_{ROM} , N_{RAM_READ} , N_{RAM_WRITE}).
- The unrolling factor
- The total number of control states (N_s).

The background memory size is computed with array first declarations. The framework provides also provide an histogram that represents the ratio of data types (bytes, long, etc..) and data nature (constant, temporary variables, arrays, scalar variable). A snapshot of the Desing trotter IHM is given in Fig. 8.

4.2.6 Physical Mapping

From the complete set of previous estimations (Fig 7-a.), a designer can make a more relevant choice among the proposed SW and HW solutions according to the metric values and tradeoff curves. Then, the abstract scheduling profiles are projected onto an implementation technology under architecture constraints using a technology description (Technology Architecture and Performance Model - TAPM) and architectural rules (Fig. 7-b.). Separate models are used for control, memory and datapath parts in a way to achieve more accuracy and define realistic solutions.




Fig. 5 Estimation flow




Fig. 6 A HCDFG specification example




Fig. 7 A HCDFG specification example

Figure 7 illustrates the estimation of the mapping process. As stated in section 4.2.5, the abstract scheduling profiles are characterized for a given cycle budget by the number of processing operators, the estimation of the memory size, the number of simultaneous memory accesses, the number of control states and the unrolling factor. We defined this way a RTL characterization from which we can derive accurate estimations of resource

occupation vs. execution time for a given target technology (expected results of the physical mapping). Results are then provided in terms of logic cells (e.g. CLB, logic elements) and dedicated cells (e.g. embedded memories, DSP operators) occupation for a given execution time of the algorithm. Information about the target technology such as operators / memories area and delay are thus needed and described in the TAPM file.




Fig. 8 Design Trotter Interface for design Estimation at each HCDFG hierarchy level

For better accuracy, a specific computation process is used for each unit of the architecture:

- Memory

A basic approach is applied to compute the area of the memory unit from the RTL characterization (number of simultaneous memory accesses, total memory size) and the TAPM (area and access times of memories). This computation process is capable of considering different type of implementation cells (logic or dedicated cells), thus to provide even more realistic estimation values. The area of the memory unit is then estimated using the following rule:

$$N_{CELL} = MS * W_{RAM} / N_{BITS/CELL}$$

where MS is the total memory size, W_{RAM} represent the bit width of the data to be stored and $N_{BITS/CELL}$ the number of memory bits implementable within a single cell. The same approach can be extended to ROM memories.

The number of control signals is derived for the evaluation of the control cost: $N_{CS} = N_{RAM} * N_{CS/RAM}$, where N_{RAM} is the number of memories (maximum of simultaneous ram reads and ram writes) and $N_{CS/RAM}$ is the number of control signals for each memory (typically, write enable and address signals left to the control unit in the current version).

- Datapath

The processing unit is estimated by adding the contribution of each processing operator which area is given by the TAPM file. For example 3 eight bit adders requires 12 logic cells on a Virtex FPGA (slices in this case), as each adder requires 4 slices to be correctly mapped on this device. Here again the number of control signals is computed in order to derive the impact on the controller.

- Control

The area of the control unit is derived from the number of control states N_S and control signals N_{CS} . Control logic is supposed to be implemented using a ROM memory. The data bit width depends then on the number of states and control signals:

$$W_{CTL} = \log_2(N_S) + N_{CS}$$

The total area of the control logic is estimated as the area of a $N_S * W_{CTL}$ ROM memory.

From the knowing of the memory, control and processing costs, we can derive a global area estimate by adding the contribution of each unit of the architecture. Then, the algorithm execution time is derived from the cycle budget and clock period defined during the scheduling process, and this complete the physical mapping estimation process of the of the hardware functions.

Concerning software estimations, by now the processor models are too coarse to obtain accurate results. Instead, we consider the solutions that present a parallelism compatible with the target processor. Accurate estimations are obtained after profiling.

Results provided by means of CLB and dedicated cells provided to the I3S partitioning tool.

4.3 Partitioning

4.3.1 Introduction

The goal of partitioning in our methodology is to allocate and schedule the application tasks such that the whole execution time is minimized, with a constraint on the RPU resources. This partitioning step takes into account the dynamic reconfiguration capabilities of the RPU. Figure 9 represents a magnification of the partitioning step in the Epicure global design flow illustrated in Fig. 4. This step requires the following inputs:

- The estimations in terms of time/resources tradeoffs for each task of the application, resulting from the estimation step (*.estim*).
- A set of architectural parameters (*.param*) describing :
 - ICURE : memory size, access time and width, ICURE API descriptions.
 - the target RPU : reconfiguration time/LC, maximum resource size (in terms of LCs and DCs).
 - the target buses : width, access time and throughput.

The OC code generated by the ESTEREL compiler (*.oc*). This code is an intermediate format common to LUSTRE and ESTEREL that allows efficient C code generation. The structure of an OC code consists typically in several lists of objects referenced by the automaton (lists of user defined types, procedures, actions...) and the explicit automaton itself. This automaton is described in the Mealy's formalism where actions (like conditional structures and external C procedure calls) are labeled on edges. For example, in figure 10, the action list 0 1 2 3 4 5 between the states <1> and <2> consists of a series of 6 external C procedure calls.

We first start by analyzing the OC code to split the control part from the data flow part using our *OcDataExtractor* tool depicted in figure 9. Then we build the Data Flow Graphs (DFGs) considering transitions or paths in the automaton. In addition to the estimations (*.estim*) and the architectural parameters (*.param*), these DFGs represent the inputs of our partitioning method based on a genetic algorithm (GA). After partitioning all the DFGs, we get a set of schedules including computations on the CPU and the RPU, data communications and reconfiguration operations. These information are integrated in the initial OC control structure to rebuild a semantically correct OC code. The GA also delivers the set of contexts intended to be run sequentially on the RPU.

We developed the *OcDataExtractor* tool to translate the OC code into an internal representation. This is made by adding test nodes when meeting conditional structures (IF statements, tests on ESTEREL signal presence) and then building a set of DFGs over transitions or paths of the automaton focusing on the external C procedure calls.

The control part consists of the state sequencing of the automaton considering the inserted test nodes, and the data flow part consists of a list of external C procedure calls over each automaton transition or path.

4.3.1. DFG Construction

In order to guarantee a coherence in the data locality between DFGs, we assume that all the data (variables) are stored in the double port internal memory of ICURE. The tasks of each DFG, respectively, read (and write back) data from (to) this memory in the beginning (at the end) of the execution. These explicit data transfers greatly limit the hazard and simplify the data coherency management. In figure 10, two memory nodes are added to our internal representation.




Fig. 9 : the partitioning flow




Fig. 10 : Data dependency check and DFG construction

4.3.1.1 DFG construction over transitions

In this case, the automaton is split into transitions⁴, and a DFG is built over each one.

A simple list of call actions does not exhibit the parallelism potential that can take advantage of the architecture capabilities. Therefore, a data dependence analysis is applied to the input and output parameters of the procedures to derive a DFG representation from this list of call actions. These actions operate on variables according to the following syntax :

Action_number : call Procedure_number (variables passed by reference) (variables passed by value).

A procedure call (denoted task in the sequel) can start its execution when all its parent tasks and incoming edges⁵ have completed their own executions. Edges between tasks are annotated with the size (in Bytes) of the transmitted variables : procedure P₄, in figure 10, passes by value the variable 1 (of type 0 : u_int = 2 Bytes) to procedure P₆ and procedure P₅ passes by reference the variable 2 (of type 1 : user_type = 64 Bytes) to P₆.

4.3.1.2 DFG construction over paths :

Partitioning can be applied to each DFG to get a set of configurations of the hardware and software units according to each list of procedure calls. However, according to the granularity level used to develop the

² In the sequel, we denote by Transition the connection between two state nodes or a state node and a test node.

³ In the sequel, we denote by Edge the connection between two tasks or a task and the memory.

⁴ In the sequel, we denote by Transition the connection between two state nodes or a state node and a test node.

⁵ In the sequel, we denote by Edge the connection between two tasks or a task and the memory.

ESTEREL description of the application, the volume of computation provided on each transition may be too weak to exploit the architecture efficiently. Therefore, we consider the ability to operate on paths of the automaton rather than single transitions.

Using a profiling of the application, we set the probabilities of transition on each exiting branch of a test-node and we speculate just once considering the most likely branch (90% of probability for the example of the transition between test node $test_2$ and state node 6 in figure 11). Considering more than one test-node in a path is feasible only if the composition of probabilities remains high enough. DFGs encountered in the path are concatenated into a single DFG (figure 11) by extending the data dependence analysis to all the global DFG tasks.

In the case where there is no data dependence between the tasks of two contiguous DFGs of the same path, we add control dependence edges (dashed edges without data exchange) assuring the temporal dependencies of actions belonging to transitions between successive states (example of the control dependence edge between tasks P_6 and P_7 in figure 11). In the contrary case, we save a write/read operation compared to a DFG construction over a transition. Task P_6 in figure 11, no longer have to store its computation result in memory but send it directly to P_8 .

Partitioning DFGs over a path including test-nodes allows the reconfiguration unit to be pre-configured according to the most likely branch.

Once the DFGs is built, our goal is to partition the different DFGs associated with transitions (or paths) of the automaton, minimizing the global execution time under a maximum RPU resource constraint. Our approach,




Fig. 11: DFG Construction over a path

based on a Genetic Algorithm detailed in [18], performs an offline, non pre-emptive and static scheduling of the DFG tasks.

4.3.2. HW/SW partitioning using a Genetic Algorithm

To take into account the dynamic reconfiguration capabilities of this kind of architectures, the partitioning problem can be split in three different steps :

- Spatial partitioning : the classical assignment of the DFG tasks to the architecture components.
- Temporal partitioning : an architecture dependent step that consists in the definition of the different configurations that will run sequentially on the RPU. Configuration sequencing virtually extends the

application space. Complex algorithms that requires more processing resources than available on a given RPU are split into separate configurations and executed sequentially.

- Scheduling : a scheduling that handles the reconfiguration times (the times for configuration switching) in addition to the execution and communication times.




Fig. 12 : Implementation point selection

4.3.2.1 The Spatial partitioning :

This partitioning step is performed by the solution encoding step of the GA. This allows the GA to process a design space exploration by generating different mappings of the tasks on the CPU and the RPU. Note that this mapping scheme is not a binary binding of each task to the CPU or to the RPU (as the one in [14]), but a binding to an implementation point (from the estimation phase).

SW and HW runtimes of each task are estimated as explained in section 4.2. A zero- LCs implementation point of a task denotes a SW only implementation. The number of implementation points can differ for each task depending on the exploitation of the available parallelism within the task.

The encoding of any solution corresponds to the binding of each task to an implementation point. The area is evaluated as a set of 'Cells' that can be grouped in two main categories: the Logic Cells (LCs) (functional elements for constructing logic) and the Dedicated Cells (DC) (e.g. Block-RAMs, dedicated Multipliers). As a consequence, the Area/ Time trade off curves are multi-dimensional. The GA encoding method codes a chromosome C with an array of genes of length N where N is the number of tasks. Each gene $C(i)$ is an integer representing a percentage. The maximum 100% value is associated with the most LCs -based expensive implementation of task τ_i . The implementation point selected is the nearest point to $C(i)$ on the LC axis. If there is more than one implementation point presenting the same LC number, we proceed with a projection of the implementation points on the area plan (figure 12) and compare the DCs on the next dimension picking also the nearest point to $C(i)$ on that axis, and so on. All the solutions delivered by this encoding method are viable. This permit us to code the multidimensional coordinates with a single integer, and greatly facilitates the conception of mutation and crossover generation operators.

The RPU affected tasks must be gathered in contexts (or clusters) to finally achieve the solution evaluation. A solution is evaluated by its overall execution time. This requires to define a schedule of the tasks including reconfigurations and data transfers between tasks.

4.3.2.2 The Temporal partitioning and scheduling :

The temporal *partitioning* and scheduling steps are considered jointly because of the important interaction existing between those two problems. These are solved by the individual evaluation step of our AG.

Then the goal is to gather the HW allocated tasks, for each DFG, into n different contexts with the following condition that :

- the size of each context does not exceed the maximum RPU size.
- there is a precedence relation between the n contexts.

- the data to be passed to the next context are buffered in the ICURE memory.
 The GA evaluation step is based on a clustering heuristic inspired by the COSYN method [4].




Fig. 13 : Our clustering heuristic

Let's assume that an encoding solution allocates only tasks P_3 and P_5 of the DFG example of figure 11 to a SW implementation (the grey tasks in figure 13) and all the other tasks to HW. Given the architectural parameters (.param) and the size of the transferred variables (labels of the edges), preliminary communication times can easily be computed [18].

We first assign priority levels to the tasks, starting from the DFG's leafs. The priority level of a task is the longest path from the task to a leaf evaluated as computation and communication costs (the task priority is in parentheses Figure 13-a). To reduce the schedule length, we need to decrease the length of the longest path by clustering tasks (grouping tasks in contexts) along it in order to reduce the communication costs.

We pick the unclustered task $\tau_i(T_{exec}(i), S_{LC}(i), S_{DC}(i))$ with the highest priority, where $T_{exec}(i)$ is the execution time and $S_{LC}(i)$ (respectively $S_{DC}(i)$) the number of LCs (respectively DCs) defined by the corresponding implementation point (for simplicity reasons, we set the execution time and resource utilization of each task to a constant given by the table of figure 13). This context building is iterated with the next unclustered task presenting the highest priority task until reaching the available maximum resources of the RPU is reached (in terms of LCs and DCs). Then, a new cluster is created and the process is repeated until all the HW mapped tasks are assigned to clusters.

For a maximum RPU size of 400 LCs and 100 DCs of the example of figure 13, we build two RPU contexts of four tasks each.

With partial reconfiguration capabilities, the reconfiguration time depends on the quantity N_k of LCs needed for mapping the context k on the RPU. We evaluate the reconfiguration time of the context k by :

$$T_{reconf}(k) = N_k \cdot T_{reconf/LC}$$

Where $T_{reconf/LC}$ is the reconfiguration time per LC. When the RPU only supports complete reconfiguration, $T_{reconf}(k)$ is set to the total reconfiguration time. Partitioning approaches like [13] only consider total reconfiguration of the RPU.

Partial and complete reconfigurations are supported in this methodology, but overlaps between computations and reconfigurations are not considered. This last feature transforms the temporal partitioning problem to a constrained placement one that is difficult to solve if we do not consider hypotheses on the size and shape of the configurations [14].

The reconfiguration time for the two contexts of figure 13 is :

$$T_{reconf} = 4 * S_{LC} * T_{reconf/LC} = 10 \text{ time units.}$$

Once the contexts are defined, the algorithm updates the *intra-Context* and *inter-Contexts* communication times. *Intra-Context* communication times are set to zero and *inter-Contexts* communication times are computed considering the data writing (reading) to (from) the ICURE memory and the reconfiguration time of the next context.

Hence, after updating the communication times, the global execution time is computed with the ASAP execution time of each task. The maximum ASAP execution time among all the leaves is the cost of the solution. The cost of the solution proposed in figure 13 is the global execution time = 108 time units.

4.3.2.3 The other steps of the GA :

The fitness (execution time) of every chromosome (solution) delivered by the GA is evaluated by allowing its ranking onto the current population. Afterwards, the GA selects (selection step) the parents that will reproduce in the next generation (generation step). Finally, a renewal step is necessary to integrate the new offspring into the current population.

- The selection step : The selection of solutions by the GA is performed by a *Tournament* technique. A number ($N_{parents}$) of tournaments are performed, each one opposes a given number of individuals randomly chosen in the current population to finally select the fittest to be one of the parents allowed to reproduce.
- The generation step : Mutation and crossover operators are used on the $N_{parents}$ selected individuals to generate the $N_{children}$ solutions representing the new offspring.
- The renewal step : After generation of the new offspring, the renewal of the population is performed according to the *elitism* principle. *Clones* are not allowed in our renewal procedure. When a number of generations N_{gen} has passed without improvements of the best individual, the GA halts and displays the best encountered solution.

4.3.3 OC updating after partitioning :

After partitioning all the DFGs deduced from the external procedure calls in the explicit automaton generated from ESTEREL, we get a set of schedules including computations on the processor and on the reconfigurable unit, data communications and reconfiguration operations. These information are integrated in the initial OC control structure by the *OcRebuilder* tool (figure 9). The GA also delivers the set of contexts, that will run sequentially on the RPU, for the *HW* synthesis tools.

WeIn our approach, we considerin that the control structure (automaton states sequencing and test nodes executions) is executed by the processor. The processor must also executealso the tasks allocated to itself and must control the reconfigurable part through the ICURE interface. This is achieved by a back-annotation of the OC code considering an interrupt-based execution scheme, a generic scheduling function on the processor and the use of APIs implemented in ICURE for context management.

5. RESULTS

Two benchmarks are considered to illustrate the EPICURE paradigm: a video supervision application (ICAM) and a part of JPEG 2000 decoder. The former illustrates that estimation and partitioning can be applied efficiently on dataflow dominated applications and the latter shows a way to cope with complexity in the case of control dominated applications.

5.1 ICAM

5.1.1 ICAM description

The first benchmark used to evaluate the result quality of our approach is a video supervision application. This application is a part of the ICAM⁶ project involving the LCEI lab of the CEA in collaboration with Thales, Atmel, Philips, WV, DC, Siemens, Alstom... that aims to develop an intelligent, cheap, and miniaturized camera with a high embedded computation power.

As depicted in figure 14, the video supervision application performs the motion detection and the object labelling given a fix background image. This application has a real time constraint of 40 ms per image.




Fig. 14 : The ICAM video supervision application

5.1.2 ICAM Results

Starting from a C ANSI description non optimized for a particular platform, we notice that this application is a set of sequential and iterative data flow treatments limiting de facto the efficiency of a parallel coarse grain architecture. So we introduced a level of parallelism by pipelining a part of the application (Background subtraction and filtering). The whole benchmark is a SSM/C specification. Considering a coarse grain decomposition of the application, we set the number of external C procedures to 16.

The Esterel compiler generates an *oc* file with a 16 state automaton. The *OCDataExtractor* tool applied to this automaton builds a set of 14 DFGs considering transitions and 10 considering paths. These DFGs contain a maximum of 5 procedure calls per transition and 16 per path (we notice here that the path with the total 16 procedure calls of the application corresponds to the nominal path).

The first step of the estimation has consisted in profiling the code with a set of video sequences. In the second step a selection of interesting functions (i.e., those with exploration potential) has been performed. This selection has been possible thanks to the characterization step, of which the results are shown in Tab.1. The third step was

⁶ ICAM : Intelligent CAMera

the intra-function estimation. This step explore the potential parallelism of the functions and generates resource vs. time constraints (abstract cycles) trade-off curves. Then a new selection step has been performed in order to select the most promising solutions. To do so, we have used a speed-up vs. resource extra-cost criterion. Finally the selected solutions have been processed by the last step of the estimation flow: the physical mapping. In this case we have used the Xilinx V400EPQ2. Some HW projection results are presented in Tab.2 Finally, the estimations results have been transmitted to the partitioning step.

Given the execution time estimations on the ARM922 processor and the Xilinx Virtex-E family FPGAs (with a reconfiguration time per LC of 5 μ s), we partition the DFGs on each transition and each path. The execution time of the Genetic algorithm over all the possible paths is about 10 minutes with an initial population size of 300 individuals.

	Gamma	MOM [0,1]	COM [0,1]
lc_testGravity	43,88	0,78	0,22
lc_labeling	10,31	0,74	0,07
lc_changeBackground	5,62	0,76	0,03
lc_reconstDilat	4,75	0,65	0,32
lc_dilatBin	4,69	0,70	0,02
lc_histoThreshold	4,00	0,64	0,29
lc_envelop	3,91	0,66	0,13
lc_absolute	2,60	0,71	0,08
lc_thresholdAdapt	2,20	0,75	0,08
lc_convolveTabHisto	1,27	0,70	0,03
lc_div	1,25	0,73	0,00
lc_getHistogram	1,22	0,75	0,00
lc_setValue	1,14	0,78	0,00
lc_add	1,11	0,75	0,00
lc_sub	1,11	0,75	0,00
lc_erodBin	1,10	0,73	0,01

Tab. 1 : ICAM characterization step

A partitioning result example over a transition is illustrated in figure 16. This gphant chart shows the scheduling on the RPU, the CPU and also the ICURE memory occupation highlighting the tasks execution time, communication time and reconfiguration time.

By summing the execution times of the transitions within the nominal path, we reach a total execution time of 36.1 ms as shown in figure 15 (*OC_Transition*) and 32.4 ms considering the nominal dataflow path (*OC_Path*). We compare these results with the partitioning of the entire application as a *Single_DFG* (without adding control by pipelining the application) and the execution time of the *SW_only* implementation. Considering the nominal path in the partitioning leads to bigger DFGs than those extracted on single transitions and therefore larger configuration contexts. So we find two contexts for *OC_Path* on a Virtex XCV2000E (43200 LCs and 150 RAM Blocks) and 10 contexts for *OC_Transition* on a smaller Virtex XCV1000E (27648 LCs and 96 RAM Blocks). As mentioned in paragraph 4.3.1.2, the DFG construction over the automaton paths allows to save a significant data exchange time between the processing elements and the ICURE memory. That can be seen on figure 15, considering a one DFG construction over the nominal path (*OC_Path*) gives an improvement of 4 ms compared to the concatenation of 12 DFGs (in *OC_Transition*).

The ICAM benchmark has proven the ability of the EPICURE paradigm to efficiently explore coarse grain applications. Now, let us see what is the behavior of our methodology when we consider more complex, fine grain, control dominated applications.

We considered the JPEG 2000 Entropic decoder as a second benchmark to establish the efficiency of our methodology.

N°	Time (ns)	States #	LC #	DC #
25	20191042	4384	868	357
26	46196075	4415	428	191
27	82309179	857	340	56
28	153223136	614	296	36
31	194375278	618	296	32
36	266391525	625	296	28

Tab. 2 : convolveTabHisto hardware




Fig. 15: Execution Time/Image and resource utilization




Fig.16 : A partitioning solution over a transition

JPEG 2000 :

5.1.3 JPEG 2000 description

JPEG2000 is a new image coding standard for different types of still images (bi-level, gray-level, colour, multi-component) with different characteristics (natural images, scientific, medical, remote sensing imagery, text, rendered graphics, etc.) allowing different imaging models (client/server, real-time transmission, image library archival, limited buffer and bandwidth resources, etc.) preferably within a unified system.

In addition to the basic compression functionality, numerous other features are provided, including:

- progressive recovery of an image by fidelity or resolution
- region of interest coding, whereby different parts of an image can be coded with different fidelity
- random access to particular regions of an image without needing to decode the entire code stream
- a flexible file format with provisions for specifying opacity information and image sequences
- good error resilience.

Taking also into account wireless constraints, JPEG2000 is well adapted to real-time embedded systems such as smart phones (MMS) or digital video systems. It provides low bit-rate operation with rate-distortion optimisation and subjective image quality performance superior to existing standards, without sacrificing performance at other points in the rate-distortion spectrum.

In the study, focused was done on the decoder part. The figure 17 gives a description of the decoder process.

Bytes are extracted from the JPEG 2000 codestream, then the parser read the data and, as soon as an entity (marker, packet header data, code block data) is ready to be decoded, the decoder process it.

Two types of data exist:

- Marker: to initialize image information (size of the image, size of the code block, type of wavelet transform, rate...)
- Packet & Packet header : input of the entropic decoder

When the parser is processed, the code block are entropic decoded. Two processing stages are executed:

- Coefficient bit modelling which re-order and reconstruct the data
- Arithmetic decoder which take bytes and contexts and generate bits to reconstruct wavelet coefficients

Finally, the process ends with, successively, the inverse quantization, the inverse wavelet transform and the inverse colour transform in order to build the image.




Fig. 17: Decoder synopsis

The yellow box from figure 17 is described in figure 18:




Fig. 18: Entropic decoder

5.1.4 Results on the JPEG 2000 Entropic Decoder

We tried to model the entropic decoder of the JPEG 2000 standard with no a priori choices on granularity level. So we considered a very fine grain specification of the entropic decoder of the JPEG 2000 standard in terms of external C procedures (we used 27 C procedures at the granularity of : logical shifts, clear, comparison ...). The resultant specification is illustrated in figure 19. The Esterel compiler generated an automaton with 17158 states.

The *OcDataExtractor* tool applied to the *oc* generated file builds 44924 DFGs considering transitions and 44158 DFGs considering paths with a maximum of 1 procedure call per transition and 3 per path. If we apply it to a subpart of the Entropic decoder with 311 states, we get 4037 DFGs considering transitions and 3183 considering paths. This high number of DFGs greatly limits the possibility to apply partitioning on each of them.

By scanning methodically the different generated DFGs, we have found that there is a large redundancy in their structure and content. So, we added a redundancy check function to the *OC2Automaton* tool to allow it creating just the non-redundant DFGs for the partitioning step. This procedure is based on a comparison of a

sequence of actions operating on a list of variables. With this kind of functions, we are sure to compare the same list of actions (C procedure calls, test nodes) operating on the same variables.

The results are given by the figure 20. We can notice that the number of DFGs is drastically reduced: 99.7% less, from 44924 DFGs to 107 for DFG construction on the transitions of the total Entropic decoder automaton. The genetic algorithm can be easily applied to these DFGs providing results for the whole application. However, due to the very fine grain specification used to describe the entropic decoder, these results have a limited interest in practice. The aim of this example was only an illustration of the ability of the methodology to deal with large complexity.

These results show also that, in the specification phase, we can not really make abstraction from the underlying estimation and partitioning tools when making choices on the treatments granularity. Despite the complexity reduction capability of our methodology, the partitioning results led us to reconsider the SSM/C specification and the external C procedures granularity.





Fig. 19: Mixed SSM/C specification of the Entropic Decoder

Fig. 20: DFG redundancy reduction

6. CONCLUSION

As a first result it is important to stress that the previous data is the result of a fully automatic process from a set of tools that includes graph extraction from a mixed SSM/C specification, system and architectural estimation and hardware/software partitioning. Such a methodology is made possible by the 'ICURE' hardware abstraction layer in between the CPU and RPU. As an additional benefit this generic interface allows for the EPICURE methodology to become applicable regardless of the hardware architecture of either the CPU or the RPU.

The EPICURE framework is a proof of concept regarding a complete design flow that alleviates the tedious system level design task for run time reconfiguration computing by allowing a fast exploration of numerous potential solutions. Therefore the overall design time is tremendously reduced.

The EPICURE experience also opens some complementary researches issues regarding the specification step. On the one hand, the estimation process can be performed at different levels of granularity, however the top level is fixed by the designer by means of C functions calls. On the other hand, we observe that the partitioning tool can reduce the DFG redundancy, which usually shows up that a coarser grain specification level could have been chosen. Thus, it appears that the designer choices still influence the final results. But this effect could be

drastically reduced with an additional methodology dedicated the application specification issue. EPICURE already brings some solutions. Actually, the characterization step of the estimation flow provides metric to measure the control influence and can guide the designer in his choice. Moreover a blurred separation between Control and Data-flow can be considered by adding some control extraction capabilities in the EPICURE flow. From a general point of view, the combination of specification / design space exploration is opens an interesting research field with real productivity gain perspectives.

The EPICURE experience could not be entirely condensed in a single paper. Thus, a hardware implementation of the "ICURE" interface is currently under progress, new architectural models have been added in the estimation flow to extend the potential system targeted. And power consumption has been considered so as to promote energy efficiency through dynamic HW/SW partitioning.

7. REFERENCES

- [1] <http://www.altera.com/products/devices/excalibur>,
- [2] <http://www.xilinx.com/products>
- [3] SIMA M., VASSILIADIS S., COTOFANA S., VAN EIJNHOFEN J., VISSERS K. A Taxonomy of Custom Computing Machines, 1st *PROGRESS workshop*, Utrecht, The Netherlands, October 13, 2000.
- [4] DAVE B.P., LAKSHMINARAYANA G., JHA. N. COSYN: HW/SW co-synthesis of embedded systems, DAC, 1997.
- [5] <http://www.esterel-technologies.com>
- [6] PLANTIN P., and STOY E., Aspects on system level design, 7th Int. Symp. on HW/SW Codesign, 1999, Rome, Italy
- [7] BILAVARN S., GOGNIAT G. PHILIPPE J-L., Fast Prototyping of Reconfigurable Architectures From a C Program, IEEE ISCAS, Bangkok, 2003.
- [8] LE MOULLEC Y., DIGUET J-Ph., and PHILIPPE J-L, A Power Aware System-Level Design Space Exploration Framework, IEEE DDECS, Brno, Czech Republic, Apr. 2002.
- [9] SANGIOVANNI-VINCENTELLI A., Defining Platform-based Design, EEDesign of EETimes, Feb. 2002.
- [10] LIEVERSE P., STEFANOV T., VAN DER WOLF T. and DEPRETTERE ED F., System Level Design with Spade: an M-JPEG Case Study, ICCAD, Nov. 2001.
- [11] GIVARGIS T., VAHID F. and HENKEL J., System-level Exploration for Pareto-optimal Configurations in Parameterized Systems-on-a-chip, ICCAD, Nov., 2001.
- [12] LE MOULLEC Y., DIGUET J-Ph., and PHILIPPE J-L, Design-Trotter: a Multimedia Embedded Systems Design Space Exploration Tools, IEEE Work. on Multimedia Signal Processing, USA, dec., 2002.
- [13] CHATHA K.S. and VEMURI R., HW/SW Codesign for Dynamically Reconfigurable Architectures, 9th Int. Work. on Field Programmable Logic and Applications (FPL'99), Glasgow, Springer Verlag, Sep., 1999.
- [14] MEI B., SCHAUMONT P. and VERNALDE S., HW/SW Partitioning and Scheduling Algorithm for Dynamically Reconfigurable Embedded Systems, 11th ProRISC workshop on Circuits, Systems and Signal Processing, Veldhoven, Netherlands, Nov. 2000.
- [15] ANDRÉ C., Representation and Analysis of Reactive Behaviors: A Synchronous Approach, Computational Engineering in Systems Applications (CESA), Lille (F), July 1996. Publisher: IEEE-SMC, pp 19-29.
- [16] BOUSSINOT F., De SIMONE R., The ESTEREL Language. Another Look at Real Time Programming, Proceedings of the IEEE, 79:1293-1304, 1991.
- [17] BERRY G., The foundations of Esterel. In G. Plotkin, C. Stirling, and M. Tofte, editors, Proof, Language, and Interaction: Essays in Honour of Robin Milner, MIT Press, 2000.
- [18] BEN CHEHIDA K. AND AUGUIN M., Partitioning reactive Data Flow Applications On Dynamically reconfigurable, Systems, IFIP Conf. on Very Large Scale Integration of System-on-Chip VLSI-SOC'03, Germany 1 - 3 December 2003