Retargeting of Compiled Simulators for Digital Signal Processors Using a Machine Description Language

Stefan Pees, Andreas Hoffmann, Heinrich Meyr
Integrated Signal Processing Systems, RWTH Aachen
pees[hoffmann,meyr]@ert.rwth-aachen.de

Abstract

This paper presents a methodology to retarget the technique of compiled simulation for Digital Signal Processors (DSPs) using the modeling language LISA. In the past, the principle of compiled simulation as means for speeding up simulators has only been implemented for specific DSP architectures. The new approach presented here discusses methods of integrating compiled simulation techniques to retargetable simulation tools. The principle and the implementation are discussed in this paper and results for the TI TMS320C6201 DSP are presented.

1 Introduction

Integrating complete systems consisting of hardware and software components on a single chip raises new challenges in the area of verification. Because target hardware is typically available only late in the design cycle, the complete system must be verified by means of *cycle-accurate* simulation. At the same time, *simulation speed* is critical for the verification of such systems and thus an important issue in simulator design [1, 2].

The principle of compiled simulation is to take advantage of a priori knowledge and move frequent operations from simulation run-time to compile-time with the goal of providing the highest possible simulation speed. Compiled simulation of programmable DSP architectures was introduced to speed up the instruction set simulation of programmable DSP architectures [3] and was extended to cycle-accurate models of pipelined processors [4]. So far, the approaches addressing the particular requirements of compiled simulation of DSPs are targeted to a specific processor architecture using a handwritten simulation compiler. However, the task of building a custom simulator for new architectures is extremely error-prone and tedious. It is a very lengthy process of matching the simulator to an abstract model of the processor architecture. These efforts can be reduced significantly by using a retargetable simulator which is generated from machine descriptions [5, 6].

This paper explores the general principles of compiled simulation that can be applied when using a language-based approach. Furthermore, an implementation based on the machine description language LISA [7, 8] is presented. From this description, efficient simulation tools are generated.

2 Related Work

Hardware description languages (HDLs) like VHDL or Verilog are widely used to model and simulate processors, but mainly with the goal of developing hardware. Using these models for instruction-level processor simulation has a number of disadvantages. They cover hardware implementation details which are not needed for performance evaluation and software verification. Moreover, the description of detailed hardware structures has a significant impact on simulation speed [2].

The machine description language nML was developed at TU Berlin [9] and adopted in several projects [1]. While retargetable assemblers and disassemblers can be generated for some DSP processors, it is not possible to produce cycle-accurate simulators for pipelined processor architectures. The main reason is the simple underlying instruction sequencer which does not support pipeline operations like e.g. flushes. Processors with more complex execution schemes like the Texas Instruments TMS320C6x cannot be described, even at the instruction-set level, because of the numerous combinations of parallel and sequential instructions within a fetch packet. These restrictions also apply to the approach of ISDL [10] which is very similar to nML. However, cycle-accurate models of pipelined processor architectures require a pipeline-accurate behavioral description beyond pure semantics. The approach based on the language EXPRESSION [11] incorporates particular mechanisms for the description of memory hierarchies.

However, no results are published that indicate the applicability for cycle-accurate simulation purposes.

The language RADL [12] is derived from earlier work on LISA [7] and extended to support multiple pipelines. But no results are provided on realized simulators based on this language.

To summarize the review, none of the approaches above does support cycle-accurate simulation or fast processor simulators that are based on compiled techniques [4]. Our interest in supporting this technique and the issue of realizing cycle-accurate processor models motivated the introduction of the language LISA which is used in our approach [7, 8].

3 Compiled Simulation

The objective of compiled simulation is to reduce the simulation time. In general, efficient run-time reduction is achieved by accelerating frequent operations. Here, the technique for accelerating operations is to use a priori knowledge during the translation of target program code into simulation code for the host.

The principle of compiled simulation for DSPs corresponds to the ideas that are already successfully implemented in the simulation of synchronous VLSI circuits [13]. Such compiled simulators for DSPs have been realized for specific processor architectures [4]. Re-using the efforts for the implementation of the compiled techniques is extremely difficult since the techniques are implemented in the so-called *simulation compiler* which is highly architecture dependent.

- The step of **instruction decoding** determines the instructions, operands and modes from the respective instruction word. The pipeline structures found in modern DSPs make it obvious that the simulation of these operations consumes a significant amount of simulation time. If we take for example the Texas Instruments TMS320C62x DSP, most instructions actually execute within only one pipeline stage (or cycle), whereas fetching, dispatching, and decoding require six pipeline stages (or cycles).

- The step of **operation sequencing** determines the sequence of operations to be executed for each instruction of the application program. This step can be implemented in a compiled simulator by generating a two dimensional table (see figure 1). One dimension of this table represents the instructions of the DSP program, the other contains pointers to functions that contribute to the transition function which drives the simulator into the next control step.

address	simulator function	simulator function	simulator function	simulator function
80561	&sim_func_11	&sim_func_12	&sim_func_13	...
80562	&sim_func_21	&sim_func_22	&sim_func_23	...
80563	&sim_func_31	&sim_func_32	&sim_func_33	...
...	...	...	...	

Figure 1: Simulation Table.

- **Operation instantiation and simulation loop unfolding** unfolds the simulation loop that drives the simulation into the next state and instantiates the respective simulation code for each instruction of the application program. This is implemented in the compiled simulator by generating individual behavioral code for each instruction of the DSP program.

Between the two extremes of fully compiled and fully interpretive simulation, partial implementation of the compiled principle is possible by implementing only some of these steps. Higher levels of compiled simulation can be achieved by investing substantially more design effort and exploiting highly architecture-specific properties. There are two levels of compiled simulation which are of particular interest – the levels which we call *static scheduling* and *dynamic scheduling* of the simulation. In case of the dynamic scheduling, the task of selecting operations from overlapping instructions in the pipeline are scheduled at run-time of the simulation. The static scheduling already schedules the operations at compile-time.

4 Model Requirements of the Simulation Compiler

Beyond the general requirements of retargetable simulators that are generated from machine descriptions, compiled simulation requires specific information on the target processor architecture to perform the above steps of the simulation compiler.

4.1 Decoding

During the decoding step, the instruction type, the operands and execution modes are determined. The operands may come from different sources (registers file, immediate, indirect) and they can have different types (signed, unsigned, fixed-point, floating point) and word lengths. Execution modes and condition codes may further specify the operation. Decoding is performed by extracting this information from the respective instruction word.

DSPs typically feature extensive non-orthogonal instruction set coding which makes decoding complex and rises considerable issues in the formal capture of the decoding mechanisms using a machine description language. Most approaches such as nML avoid this problem by capturing the non-orthogonal coding in the behavioral model. However, this is no representation which hardly allows to distinguish (simulation) *run-time* operations from *compile-time* operations – those operations that already can be performed during simulation compilation. In LISA, the distinction between these two types of operations is made by means of particular IF-ELSE and SWITCH-CASE statements which are discussed later.

4.2 Operation Sequencing

In order to perform operation sequencing, the precedence of operations composing one instruction and the inter-instruction precedence must be determined. The complexity of this task rapidly grows with the depth and mechanisms of the instruction pipeline.

The processor model must provide detailed information on the pipeline structure and its mechanisms in order to enable this step. The LISA language with its detailed pipeline model enables the description of all pipeline structures and the intra-instruction precedence of operations. Figure 2 shows the intra-instruction precedence relations of operations for a simple four-stage pipeline (with the stages IF, ID, EX, WB).

Figure 2: Intra-instruction precedence.

4.3 Operation Instantiation

The inter-instruction precedence of operations can be derived from the overlapping of instructions in the pipeline for the case that no control hazards occur. Figure 3 depicts both, the intra- and inter-instruction precedence relations of operations. The simulation compiler has to compose operations from overlapping instructions to form the transition function that drives the simulation into the next state. Such operations are shown in vertical columns in figure 3.

Due to control hazards such as jumps, branches and exceptions, the program execution may follow different paths which causes multiple possible combinations of operations in the pipeline. For this reason, the simulation compiler has to generate code for all these possible combinations. During run-time of the simulation, the appropriate path is selected.

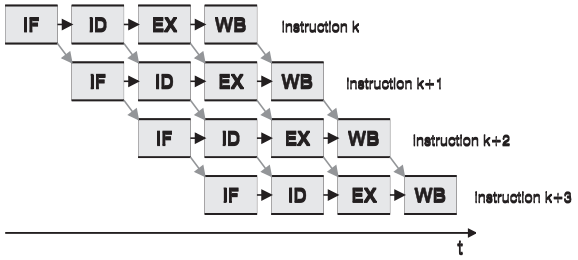

Figure 3: Precedence of instructions in the pipeline.

5 LISA Language

LISA descriptions are composed of *resource declarations* on the one hand and of *operations* on the other hand. The declared resources build the storage objects of the hardware architecture (e.g. registers, memories, pipelines) which capture the state of the system and which can be used to model the limited availability of resources for operation access.

Operations are the basic objects in LISA. They represent the designer's view of the behavior, the structure, and the instruction set of the programmable architecture. Operation definitions collect the description of different properties in several sections, such as the operation behavior (in the BEHAVIOR section), instruction set information, and timing.

Operations are formed by a header line and the operation body. The header line consists of the keyword OPERATION and its identifying name:

```
OPERATION name_of_operation
{
  sections...
}
```

For more details on the LISA language, please refer to [8].

5.1 Formal Description of Non-orthogonal Coding Fields

In LISA, non-orthogonal coding is expressed by additional conditional statements that can be used to structure the processor model. The purpose of these new conditional statements is to express the coding dependencies between different operations. Following the syntax of programming languages, they have the form of IF-ELSE and SWITCH-CASE statements.

We will now discuss an example. Figure 4 displays the coding of a simplified instruction word. There are three instructions `add`, `sub`, and `mul` whose execution is also controlled by the coding field `mode` which selects between

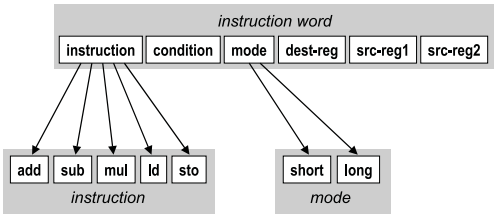

Figure 4: Non-orthogonal coding fields.

`short` and `long` operands and their specific arithmetic. However, the other instructions `ld` and `sto` use the `mode` field for a different purpose. Possible LISA code for the `add` instruction is shown in example 1.

```
OPERATION add
{
  DECLARE { REFERENCE mode; }
  IF (mode == short) {
    BEHAVIOR { dest_lo = src1_lo + src2_lo; }
  }
  ELSE {
    BEHAVIOR {
      dest_lo = src1_lo + src2_lo;
      carry = dest_lo >> 16;
      dest_lo &= 0xFFFF;
      dest_hi = src1_hi + src2_hi + carry;
    }
  }
}
```

Example 1: Formal expression of non-orthogonality.

Here, the IF-THEN-ELSE statement encloses two alternative sections with their respective behavioral description of the operation `add`. This formal representation lets the simulation compiler distinguish these two cases and generate specific simulation code.

6 Implementation Results

In order to evaluate the applicability and efficiency of compiled simulation in a retargetable environment, we implemented the first two steps of section 3 – compile-time decoding and operation sequencing – in our experimental tool suite. As shown in Figure 5, a LISA compiler takes the processor model and translates in into a data base. The information in this data base is accessible for the simulation compiler generator which produces source code in C++ for a processor-specific simulation-compiler. This simulation compiler translates application of the target processor into a simulation table which becomes a part of the final simulator.

We have chosen the Texas Instruments TMS320C6201 DSP as reference processor for our experimental analysis of the obtainable simulation speed. The TMS320C6201 was described in LISA as a cycle-based model. Although

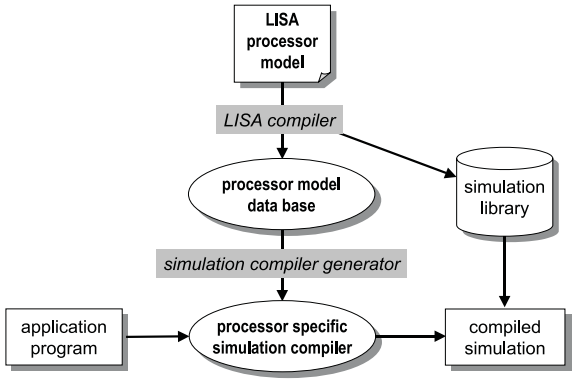

Figure 5: Retargetable, compiled simulation tools.

the architecture of this processor with two pipelines consisting of eleven pipeline stages is very complex, the LISA description including the memory interface was realized by one designer in 6 weeks. The complete translation of this model with the LISA compiler and the simulation compiler generator takes less than 35 seconds on a Sparc Ultra 10 workstation. As a comparison, a custom compiled simulator for the less complex TMS320C54x (six-stage pipeline) the same designer has spent more than 12 months.

6.1 Simulator Benchmarks

In order to evaluate the simulation speed of our generated, compiled simulator of the TI C6201 we used the `sim62x`, version 2.0 which is part of the TI's software development tools for our reference. The benchmarks are based on three typical DSP algorithms, a FIR filter, the ADPCM G.721 codec, and the GSM speech encoder. All measurements were made on a Sparc Ultra 10.

Compilation time of object code into a compiled simulation was measured on three reference applications. The required time and respective application size is shown in figure 6. The compilation speed is calculated by relating the number of compiled instructions to compilation time.

Application	FIR	ADPCM	GSM
code size [kByte]	2112	3072	58592
compilation time [s]	0.97	1.38	27.31
compilation speed [instructions/s]	544	557	536

Figure 6: Simulation Compilation Speed.

For all applications measured, the compilation speed ranges between 530 and 560 instructions/s, even for our GSM coder that nearly requires the whole internal memory space of the DSP.

Simulation speed was quantified by running an application on the respective simulator and relating the simulation time to the processed number of cycles. The reference simulator from TI achieved between 2k and 9k cycles/s whereas our generated simulator runs at speeds between 288k and 403k cycles/s at the same accuracy level. This corresponds to factors of 47x to 170x faster simulation as shown in figure 7.

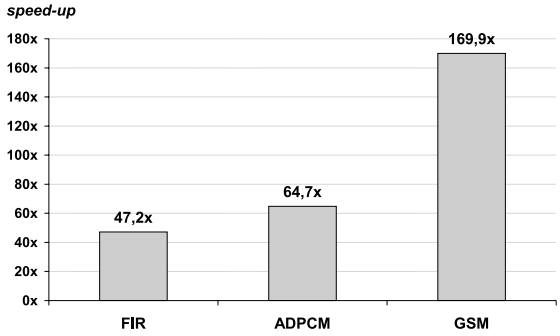

Figure 7: Speed-up: LISA simulator vs. TI sim62x.

7 Conclusion and Future Work

In this paper, we presented the new approach of applying compiled simulation to retargetable processor simulation environments. The compiled technique is a technology that enables fast simulation of programmable DSP architectures. Up to now, compiled simulation has only been implemented for specific processor architectures. Retargetable, compiled simulation based on a machine description language puts specific requirements on the instruction set model. The complete instruction coding must be described formally to enable a high degree of compiled simulation. The processor description language LISA is able to provide such models. In a case study, compiled simulation techniques are implemented for a model of the Texas Instruments TMS320C6201 DSP. Our generated, compiled simulator based on the LISA description runs at 47-170 times higher simulation speed than the commercially available instruction set simulator from TI without any loss in accuracy.

Our future work will focus on modeling further real-life processor architectures and retargetable compiled simulators that provide the third step of compilation – operation instantiation. Another issue is the integration of software simulators into HW/SW co-simulation environments. Furthermore, the goal of the ongoing language design is to address retargetable compiler back-ends as well.

References

[1] M. Hartoog, J. Rowson, *et al.*, "Generation of software tools from processor descriptions for hardware/software codesign," in *Proc. of the Design Automation Conference (DAC)*, Jun. 1997.

[2] J. Rowson, "Hardware/Software co-simulation," in *Proc. of the ACM/IEEE Design Automation Conference (DAC)*, 1994.

[3] V. Živojnović, S. Tjiang, and H. Meyr, "Compiled simulation of programmable DSP architectures," in *Proc. of IEEE Workshop on VLSI Signal Processing*, (Sakai, Osaka), pp. 187–196, Oct. 1995.

[4] S. Pees, V. Živojnović, A. Ropers, and H. Meyr, "Fast Simulation of the TI TMS 320C54x DSP," in *Proc. Int. Conf. on Signal Processing Application and Technology (ICSPAT)*, (San Diego), pp. 995–999, Sep. 1997.

[5] M. Barbacci, "Instruction set processor specifications (ISPS): The notation and its application," *IEEE Transactions on Computers*, vol. C-30, pp. 24–40, Jan. 1981.

[6] A. Fauth and A. Knoll, "Automatic generation of DSP program development tools using a machine description formalism," in *Proc. of the ICASSP - Minneapolis, Minn.*, 1993.

[7] V. Živojnović, S. Pees, and H. Meyr, "LISA – machine description language and generic machine model for HW/SW co-design," in *Proceedings of the IEEE Workshop on VLSI Signal Processing*, (San Francisco), Oct. 1996.

[8] S. Pees, A. Hoffmann, V. Zivojnovic, and H. Meyr, "LISA – Machine Description Language for Cycle-Accurate Models of Programmable DSP Architectures," in *Proceedings of the Design Automation Conference (DAC)*, (New Orleans), June 1999.

[9] A. Fauth, M. Freericks, and A. Knoll, "Generation of hardware machine models from instruction set descriptions," in *Proc. of the IEEE Workshop on VLSI Signal Processing*, 1993.

[10] G. Hadjiyiannis, S. Hanono, and S. Devadas, "ISDL: An instruction set description language for retargetability," in *Proc. of the Design Automation Conference (DAC)*, Jun. 1997.

[11] A. Halambi, P. Grun, et *al.*, "EXPRESSION: A language for architecture exploration through compiler/simulator retargetability," in *Proc. of the Conference on Design, Automation & Test in Europe (DATE)*, Mar. 1999.

[12] C. Siska, "A processor description language supporting retargetable multi-pipeline DSP program development tools," in *Proc. of the Int. Symposium on System Synthesis (ISSS)*, Dec. 1998.

[13] Z. Barzilai, *et al.*, "HSS - A high speed simulator," *IEEE Trans. on CAD*, vol. CAD-6, pp. 601–616, July 1987. 1987.