Optimizing for Data-Parallelism in Kahn Process Networks

[Extended Abstract]

Andrés Goens (student)
Chair for Compiler Construction
TU Dresden, Germany
andres.goens@tu-dresden.de

Jeronimo Castrillon (advisor)
Chair for Compiler Construction
TU Dresden, Germany
jeronimo.castrillon@tu-dresden.de

ABSTRACT

Dataflow Models of Computation (MoC) are extensively used to leverage compile-time information and enable efficient scheduling of applications. With the multicore era, interest in these MoCs has reflourished, since they enable an automated extraction of parallelism and using heretogeneous resources. This is in part due to the formal nature of the models, which guarantees a firm grip over the flow of data and a deterministic execution. This formalized approach permits a formulation of the problem of mapping and scheduling dataflow applications to hardware architectures as an optimization problem, such that (near-) optimal solutions can be found via design-space exploration (DSE). While dataflow models excel at extracting task-level parallelism and some instances of pipeline parallelism, they traditionally struggle when data-level parallelism (DLP) is involved; the static nature of the graphs behind the models makes it difficult.

For models with static data-flow patterns, like Synchronous Data Flow (SDF), extensions have been put forward to enable efficient DLP [1]. However, with the growing complexity and data-dependecies of modern algorithms, the static nature of SDF falls short on many problems. For example, decoders for multimedia formats which exploit the nature of the data for compression, like H.265 (HEVC), or in the Long-Term Evolution (LTE) telecommunication standard, which requires base-stations to dynamically adapt to the number of users to efficiently use the system's resources.

In this work we extend a dynamic dataflow model, Kahn Process Networks (KPNs), to enable DLP. We do so in a parametric fashion, by adding parametric processes with a formal foundation that is transparently compatible with the semantics of the original KPNs. By formally adding a parameter, a DSE step can optimize for the right amount of DLP to a use-case scenario, as an additional parameter to the mapping and scheduling of processes.. Furthermore, it preserves deterministic properties that make KPNs valuable, which allows switching between different parameters at runtime, in a fashion similar to [2].

CGO '17 Austin, Texas USA
 ISBN 123-4567-24-567/08/06. . . $15.00
DOI: 10.475/123_4

Figure 1: Optimizing Data-Level Parallelism for the N-Body Problem as Parametrized Kahn Process Network.

Using a pre-processor extension in a modern commercial KPN-based compiler for MPSoCs, we evaluated our parametrized KPNs for applications with data-parallelism. Figure 1 shows the results of an N-body simulation, which is dynamic in nature as well, since the number of bodies is not known at compile-time. In it, n is a parameter introduced by our model extensions, and the experiments were executed on machine with two Intel Xeon E5-2630v2 processors running at 2.6 GHz. It shows the effectiveness of the extension, with almost linear scaling thanks to the data-parallism, which peaks after $n = 32$, however, and for larger numbers the process overhead overcomes the parallism advantages. This optimization target would not be possible with the existing methods. Current work is involved in also testing with an LTE application and speaker-recognition software, and future work will concentrate in leveraging the formal properties of this extension to map to other execution models, e.g. OpenCL [3].

1. REFERENCES

[1] B. Bhattacharya and S. S. Bhattacharyya, "Parameterized dataflow modeling for dsp systems," *IEEE Transactions on Signal Processing*, vol. 49, no. 10, pp. 2408–2421, 2001.

[2] L. Schor, I. Bacivarov, H. Yang, and L. Thiele, "Adapnet: Adapting process networks in response to resource variations," in *Compilers, Architecture and Synthesis for Embedded Systems (CASES), 2014 International Conference on.* IEEE, 2014, pp. 1–10.

[3] W. Lund, S. Kanur, J. Ersfolk, L. Tsiopoulos, J. Lilius, J. Haldin, and U. Falk, "Execution of dataflow process networks on opencl platforms," in *2015 23rd Euromicro International Conference on Parallel, Distributed, and Network-Based Processing.* IEEE, 2015, pp. 618–625.