Design Tools Session

Model-based multicore architecture exploration with CoFluent Studio

J. Lemaitre, V. Perrier
CoFluent Design, Nantes, France, {jerome.lemaitre; vincent.perrier}@cofluentdesign.com

Abstract

To match multicore architectures that are massively parallel, applications should be specified using a parallel language. With CoFluent models, this parallelism is captured graphically in an unambiguous manner. In a graphical representation of a functional behavior, functions run independently of each other and exchange data in a network. Inside each function, the behavior can be refined with sequential C code. These models are automatically converted to transaction-level SystemC code, which is executed to validate the behavior of the system in time independently of any architecture constraints. These constraints are introduced in platform models, which consist of abstract software processing units (CPUs or DSPs), hardware processing units (ASICs or FPGAs), storage units and communication units (busses or routing networks).

Platform components are parameterized in terms of speed, number of cores (typically ranging from 1 to 100), scheduling policies, power consumption and cost, etc. No dedicated hardware models or instruction set simulators are required for easier and faster architecture definition. When mapping application models onto platform models, a new SystemC model is generated that takes the multiprocessor/multicore environment properties into account to execute the functions and communications.

This approach allows designers to rapidly explore the impact of high-level parameters on the behavior and performance of the complete system. In particular, the impact of the number of cores on the power consumption and resource load of a system such as a MPEG2 encoder can efficiently be analyzed as illustrated in this poster.