

## A Reconfigurable Multi-Core Architecture to Support SPMD Applications

John K. Antonio  
School of Computer Science  
University of Oklahoma  
Norman, OK

In a multiprocessor architecture, the number of processors that can be effectively utilized to execute a SPMD (Single Program Stream, Multiple Data Stream) application is generally dictated by two factors: the *granularity of the application* and the *degree of coupling* among the processors of the architecture. Granularity is a relative measure of how frequent inter-processor communication (or synchronization) is required. Coarse granularity means there is generally a long period of time between communication events; fine granularity means there is a relatively short period of time between communication events. The degree of coupling of a multiprocessor architecture is a relative indicator of the communication bandwidth among the processors. A loosely coupled architecture has a relatively low bandwidth interconnection network; a tightly coupled architecture has a relatively high bandwidth interconnection network.

For a given SPMD application with a fixed problem size, increasing the number of processors used to execute the application on a given multiprocessor architecture generally causes the application granularity to become finer, i.e., the time between communication events decreases. When granularity becomes too fine relative to the coupling of a given multiprocessor architecture, then inter-processor communication time overhead becomes significant, resulting in diminishing overall performance. Thus, fine granularity applications require a more tightly coupled architecture in order to achieve improved and scalable performance.

In this paper, a reconfigurable multi-core architecture is introduced. The reconfigurable resources of the architecture are employed to adjust the bandwidth of the interconnection among the cores and/or the processing capacity of co-processing units associated with each core. Devoting more reconfigurable resources to the interconnection among cores has the benefit of making the coupling among the cores tighter, but at the expense of reducing the ability for each core to locally exploit instruction-level parallelism. By devoting more reconfigurable resources to each co-processing unit, the potential performance of each core is improved, but at the expense of loosening the coupling among the cores. Thus, at one extreme, the proposed architecture can be configured to be a tightly coupled collection of simple cores; at the other extreme the architecture can be configured to be a relatively loosely coupled collection of powerful cores. An analytical study is included to demonstrate the efficacy of the approach by illustrating that the configuration of the architecture can be optimally determined to match the granularity of a given SPMD application.