Abstract

The evolutionary development of reconfigurable logic devices has seen a gradual drift from fine-grain to coarser-grain architectures, but this remained firmly grounded on their processing capabilities. These improvements have unfortunately not been accompanied by similar advancements in support for data-flow. Although embedded memory blocks have been added and routing resources improved, data-flow remains the bottleneck in most implementations, particularly for DSP applications.

This research aims at improving an application's data-flow implementation by presenting memory design abstractions for structured data access. These abstractions can then either be mapped to a dedicated address generation unit, or be offered as a library of address generation modules for platform-independent optimised synthesis. Besides requiring fewer routing resources, the dedicated implementation exhibits an additional benefit towards area (up to 18 times) and performance (up to 2.6 times).

The benefit of reconfigurable logic devices over digital signal processors is mostly related to their potential for more parallelism. Since memory blocks in reconfigurable platforms are generally smaller than in DSP processors, the scalability of this memory with localised address generation is studied. Different memory partitioning schemes for improved concurrent accesses are investigated in detail. Design exploration for a 2D filter indicates area and performance variations up to 55 and 2.7 times, respectively.

Some model applications are used to demonstrate the memory architecture design-space. The implementations are investigated with respect to the on-chip buffer size,

the potential parallelism and possibility of appending new data during processing.

The presented work indicates the benefits of dedicated address generation and memory design abstractions to improve the data-flow implementation on reconfigurable devices. The approach is scalable, eases the design process and can be applied to a large set of applications.