Deriving an Efficient, Application-Specific, FPGA-Based Pipelined Processor

- C code is transformed to hardware architecture
- Template is provided to guide application-specific architecture derivation (e.g. Simulated Annealing)
- Template provides high-level temporal parallelism
- Temporal and spatial parallelism are exploited within pipelined stages
  - Port-widening for memory-intensive stages
  - Multiple data paths for compute-intensive stages

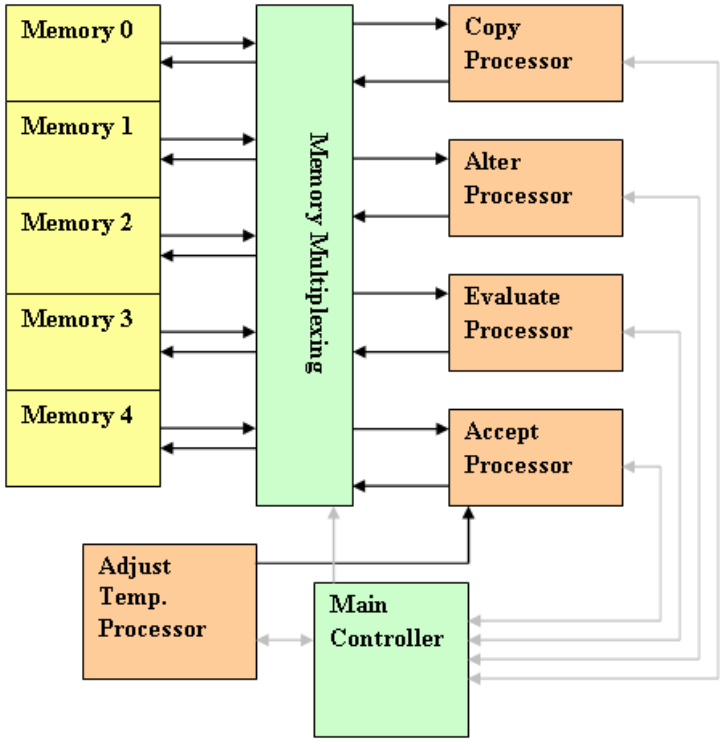

Results

	ASP1	ASP2	ASP3	ASP4
Copy Latency	101 cyc.	101 cyc.	51 cyc.	51 cyc.
Alter Latency	24 cyc.	24 cyc.	24 cyc.	24 cyc.
Evaluate Latency	409 cyc.	78 cyc.	78 cyc.	45 cyc.
Accept Latency	54 cyc.	54 cyc.	54 cyc.	54 cyc.
Adjust Temp. Latency	12 cyc.	12 cyc.	12 cyc.	12 cyc.
Circuit Size (LUT/FF/BRAM)	4,552 2,668 23	4,856 2,972 28	5,341 2,972 49	5,729 3,324 97