Buffer-space Efficient and Deadlock-free Scheduling of Stream Applications on Multi-core Architectures

Jongsoo Park and William J. Dally
Concurrent VLSI Architecture Group, Stanford University

Background – Stream Compilers

Model of Computation – Synchronous Data Flow (SDF) [1]

- Actors: computation
- Streams: communication
- Many embedded apps.
- Deterministic parallelization
- Efficient communication through DMAs

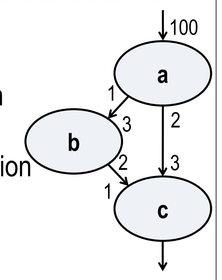

Two important phases

- Partitioning: e.g., map A and B to core 0, C to core1
  - Use METIS or integer linear programming
- Scheduling

Scheduling – Previous Work

Stream graph modulo scheduling (SGMS)[2]

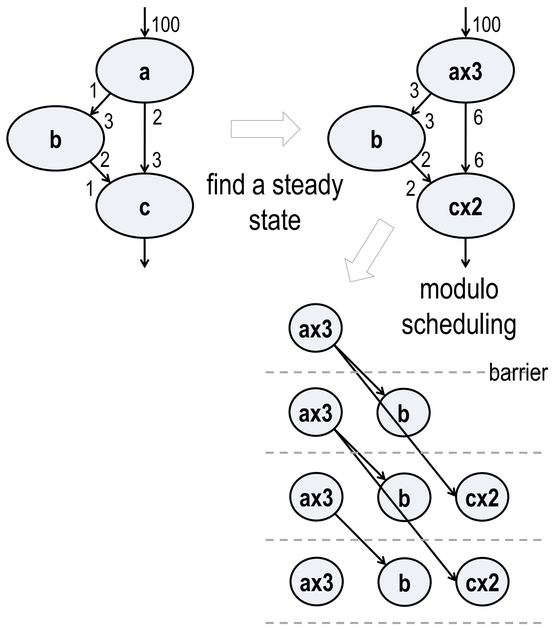

Previous Work – Drawbacks

Buffer-space Explosion: e.g., W-CDMA

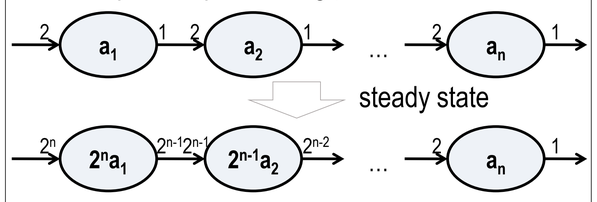

Feedback Loops: e.g., GSM

- Contract feedback loops to single actors, and loose parallelism, or
- Suffer deadlocks.

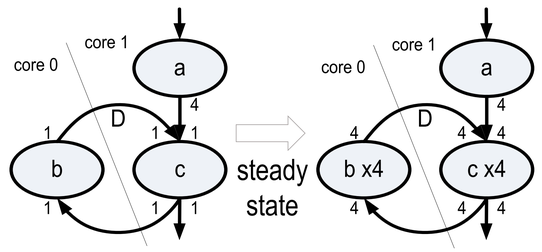

Team Scheduling[3]

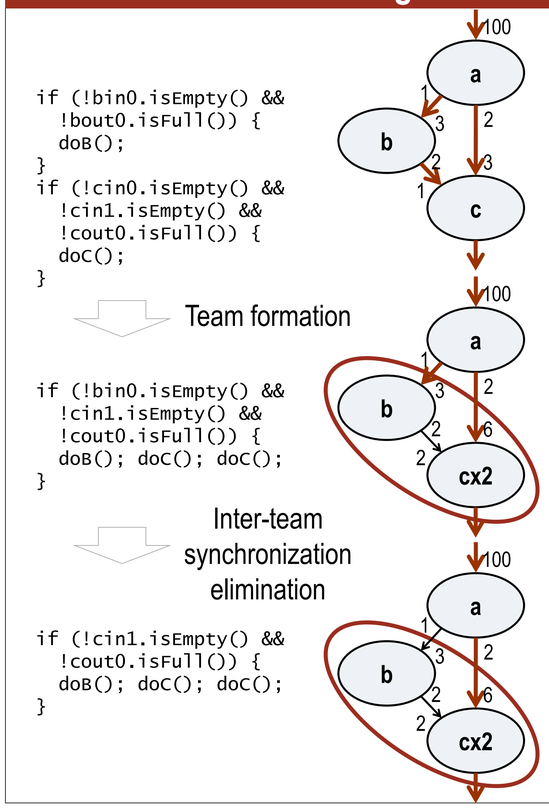

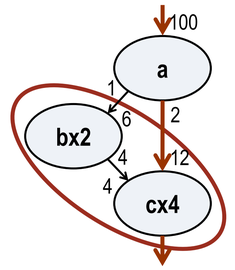

Buffer Sizing

Producer-consumer Pairs

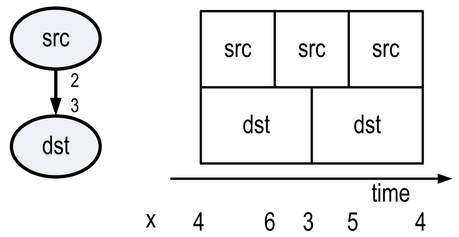

- 2(p + c – gcd(p, c)), where p = # of output tokens per producer firing and c = # of input tokens per consumer firing

General Case

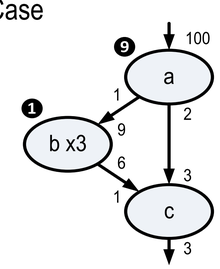

- len(a, c) = 128 → deadlock
- len(a, c) = 180 → serialization

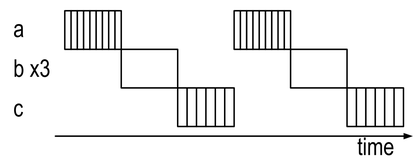

- min len(a,c) to avoid serialization = 400 (see [3]).

Results

- Evaluation Setup
  - StreamIt benchmarks, GSM, and W-CDMA
  - 16-core ELM [4]
  - ELK stream programming language [5]
- Buffer-space Limited Experiment

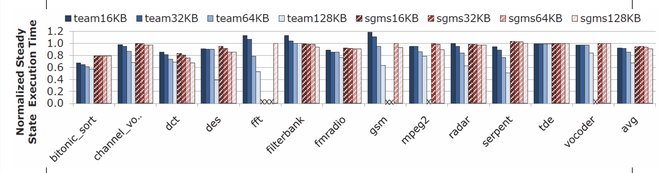

- SGMS cannot satisfy the buffer limit for W-CDMA (16-128KB), fft (16-64KB), gsm (16-32KB), mpeg (16KB), and vocoder (16KB)
- Average speedup of team128KB is 11x
- 37% higher throughput

- Amortization Factor Limited Experiment

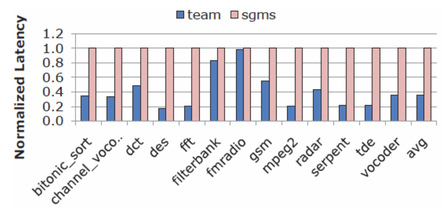

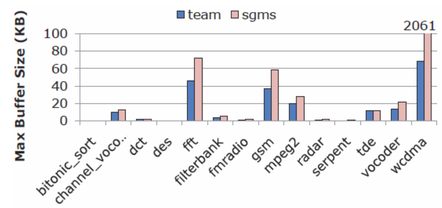

References

[1] Edward A. Lee and D. G. Messerschmitt, Static Scheduling of Synchronous Data Flow Programs for Digital Signal Processing, IEEE Transactions on Computers, 1987
[2] Manjunath Kudlur and Scott Mahlke, Orchestrating the Execution of Stream Programs on Multicore Platforms, PLDI, 2008
[3] Jongsoo Park and William J. Dally, Buffer-space Efficient and Deadlock-free Scheduling of Stream Applications on Multi-core Architectures, SPAA, 2010
[4] J. Balfour, W. Dally, D. Black-Schaffer, V. Parikh, and J. Park, An Energy-efficient Processor Architecture for Embedded Systems, Computer Architecture Letters, 2008
[5] ELK webpage. CVA Group, Stanford University. http://cva.stanford.edu/projects/elk