

Efficient Handling of Arrays in Dataflow Process Networks

Daniel Baudisch, Jens Brandt, Klaus Schneider

Embedded Systems Group
Department of Computer Science
University of Kaiserslautern, Germany

Outline

- 1 Introduction
- 2 Reduce Communication Costs of Arrays
- 3 Results
- 4 Conclusion

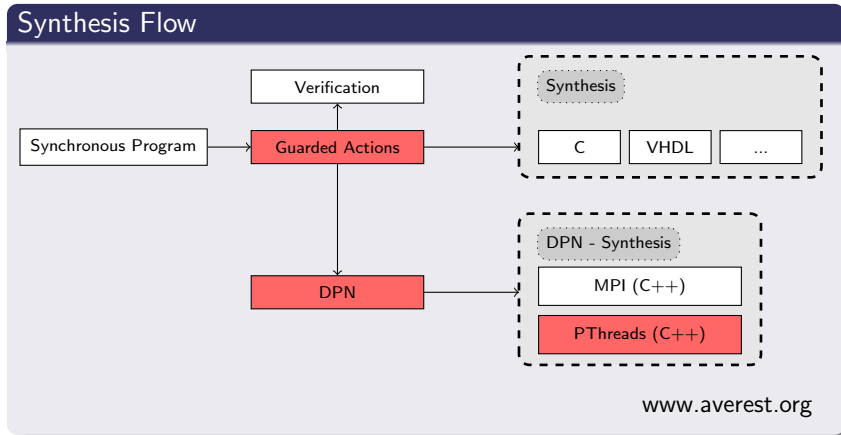
Outline

- 1 Introduction
- 2 Reduce Communication Costs of Arrays
- 3 Results
- 4 Conclusion

Motivation

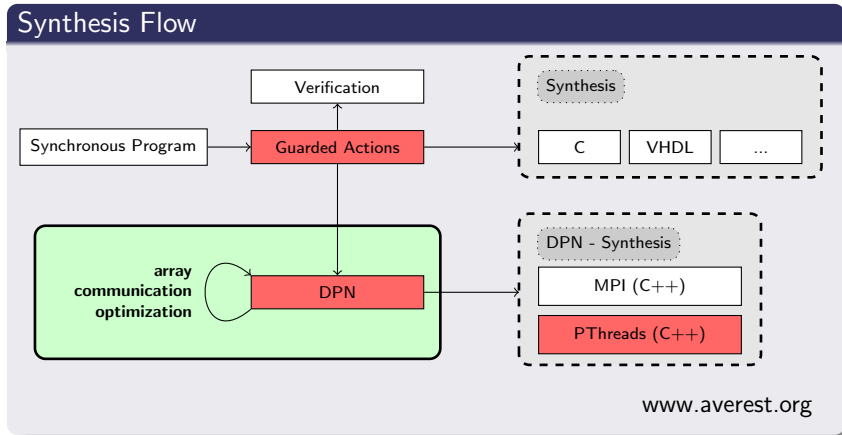
- model-based development of applications
in particular: embedded systems
- synchronous languages, e. g. Esterel, Lustre, Quartz
 - can be used for embedded systems
 - hide communication latencies
 - execution of instructions in perfect synchrony
- synthesis more difficult
(especially for heterogenous/distributed systems)

Synthesis Flow



- here: from DPN to DPN (array communication optimization)

Synthesis Flow



- here: from DPN to DPN (array communication optimization)

Guarded Actions (GA)

System (Example)

Interface:

Inputs: i, d, c

Output: o

Locals: $a[N], j0, j, p$

Guarded Actions:

$p \Rightarrow \text{next}(p) = \text{True}$

$c \wedge p \Rightarrow a[j] = i$

$p \Rightarrow \text{next}(j) = (j + 1) \% N$

$p \Rightarrow j0 = (j - d) \% N$

$p \Rightarrow o = a[j0]$

Guarded Actions (GA)

System (Example)

Interface:

Inputs: i, d, c

Output: o

Locals: $a[N], j0, j, p$

Guarded Actions:

$p \Rightarrow \text{next}(p) = \text{True}$

$c \wedge p \Rightarrow a[j] = i$

$p \Rightarrow \text{next}(j) = (j + 1) \% N$

$p \Rightarrow j0 = (j - d) \% N$

$p \Rightarrow o = a[j0]$

Generic execution

```
while(True)
```

```
    read system inputs
```

```
    execute GA in data-flow order
```

```
        update state of system
```

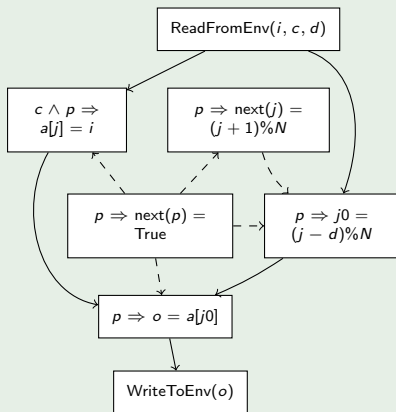
```
        generate outputs of system
```

```
    write system outputs
```

```
endwhile
```


Guarded Actions (GA)

Dependency Graph



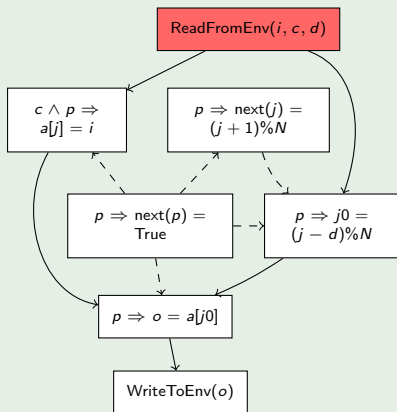
Generic execution

```

while(True)
  read system inputs
  execute GA in data-flow order
  update state of system
  generate outputs of system
  write system outputs
endwhile
  
```

Guarded Actions (GA)

Dependency Graph



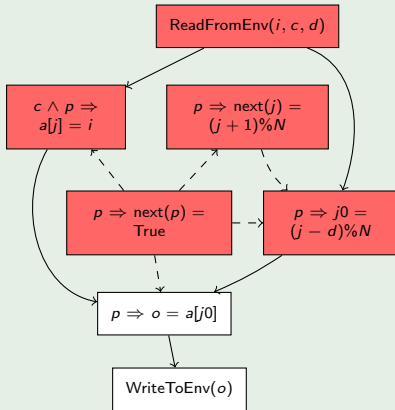
Generic execution

```

while(True)
  read system inputs
  execute GA in data-flow order
  update state of system
  generate outputs of system
  write system outputs
endwhile
  
```

Guarded Actions (GA)

Dependency Graph

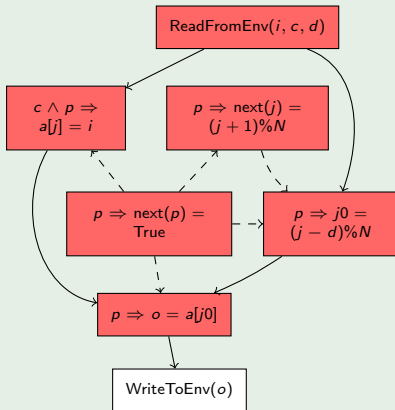


Generic execution

```
while(True)
  read system inputs
  execute GA in data-flow order
  update state of system
  generate outputs of system
  write system outputs
endwhile
```

Guarded Actions (GA)

Dependency Graph




Generic execution

```

while(True)
  read system inputs
  execute GA in data-flow order
  update state of system
  generate outputs of system
  write system outputs
endwhile
  
```

Guarded Actions (GA)

Dependency Graph




Generic execution

```

while(True)
  read system inputs
  execute GA in data-flow order
  update state of system
  generate outputs of system
  write system outputs
endwhile
  
```

Guarded Actions (GA)

DPN




Notes

- DPN is gained by replacing dependencies by FIFO-buffers \Rightarrow “Theory of Latency Insensitive Design” (McMillan et. al)
- writer must be uniquely determined: group nodes writing to the same variable
- grouping of nodes / partitioning increase computation effort per node

Guarded Actions (GA)

DPN




Synthesis (PThreads)

- create thread for each node
- each thread:


```
while(True)
  read input buffers
  execute GA of node
  write output buffers
endwhile
```

Guarded Actions (GA)

DPN



Issue

- transport of arrays expensive
- here:
only few elements are written in each step
(recognized by variable array indices)
 \Rightarrow only changes should be committed

Outline

- 1 Introduction
- 2 Reduce Communication Costs of Arrays
- 3 Results
- 4 Conclusion

Related Work

- Gupta et. al.: analysis of structured programs with data parallelism
CFG to generate 'Section Communication Descriptors' and 'Availability Section Descriptors'
- Arvind: demand-driven optimization
Compute data only if it is required. Reduces computations, keeps communication.
- CAL (OpenDF): arrays only as local states
Communication is left to programmer.
- *more related work in paper*

Basic Idea

- create local copy a^R of array a in each node R reading a
- determine tuple(s) containing update information in writer:
 - guard $valid_a$ (condition whether update has to be done)
 - index idx_a (address of element that is overwritten)
 - value rhs_a
- communicate update information instead of modified array to readers
- *detailed pseudo-algorithm in the paper*

Example

Origin

$$c \wedge p \Rightarrow a[j] = i$$

$$p \Rightarrow o = a[j_0]$$

Modified

$$\begin{aligned} c \wedge p &\Rightarrow \text{valid}_a = \text{true} \\ \text{valid}_a &\Rightarrow \text{idx}_a = j \\ \text{valid}_a &\Rightarrow \text{rhs}_a = i \\ \text{valid}_a &\Rightarrow a[\text{idx}_a] = \text{rhs}_a \end{aligned}$$

$$\begin{aligned} p &\Rightarrow o = a^R[j_0] \\ \text{valid}_a &\Rightarrow a^R[\text{idx}_a] = \text{rhs}_a \end{aligned}$$

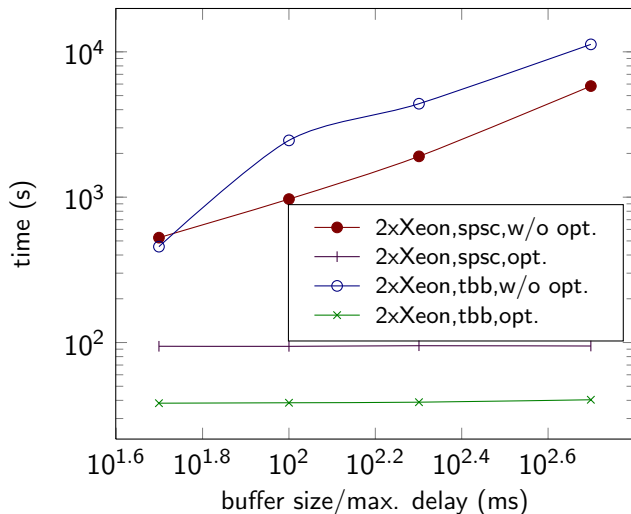
Some Details

- take care of timing (immediate / delayed)
- index tuple for multi-dimensional array access
⇒ neglect constants
- when to use communication reduction
(size of update) * (number of potential update) < size of array
- *re-use update tuple* for assignments with same index
- *writer of memory cell is uniquely determined*
allows to parallelize operations on multi-dimensional arrays

Outline

- 1 Introduction
- 2 Reduce Communication Costs of Arrays
- 3 Results**
- 4 Conclusion

Benchmarks - Results



Benchmarks - Results

Benchmark systems: 2x Xeon X5450 (= 8x 3.00GHz)
i5-750 (4x 2.66GHz)

Benchmark Name	Speedup on i5-750		Speedup on 2xXeon	
	custom queue	Intel TBB queue	custom queue	Intel TBB queue
Delay (max. 50ms)	53.29	19.02	5.59	11.94
Delay (max. 100ms)	98.19	39.90	10.31	63.84
Delay (max. 200ms)	189.43	66.30	20.04	113.10
Delay (max. 500ms)	343.85	208.37	61.44	278.69
MatrixMult (16x16)	1.14	1.35	1.18	1.33
MatrixMult (32x32)	1.88	1.47	1.30	3.03
MatrixMult (48x48)	1.56	1.53	1.33	1.56
Pitchshift	809.00	289.00	145.64	340.08
Landscape Rend. (H=400)	8.35	3.96	4.20	5.31
Landscape Rend. (H=800)	34.73	4.66	7.38	9.81
Landscape Rend. (H=1600)	36.41	20.26	17.97	20.05

Outline

- 1 Introduction
- 2 Reduce Communication Costs of Arrays
- 3 Results
- 4 Conclusion

Conclusion + Future Work

- synchronous MoC hides communication latencies
⇒ nice for parallel programming
- in practice: communication costs time
- done: reduction of communication by sending update of single array elements instead of complete array
 - never slower
 - may be faster
- further issues:
 - generic reduction of communication
⇒ endo- and isochronous systems to reduce overhead
big idea: consider when a variable is changed and when it is needed

The End

Thank you for your attention!
Questions? Suggestions? Ideas?