Mapping and Scheduling Streaming Applications using SMT Solvers

Pranav Tendulkar

Supervisors:

Dr. Oded Maler

Dr. Peter Poplavko

Verimag, FRANCE

13 October 2014

Multi-core Processors Everywhere



Multi-core Processors Everywhere





Multi-core Processors Everywhere

























Multi-core systems



How To:



Multi-core systems



How To:

• Deploy the application to the platform



Multi-core systems



How To:

- Deploy the application to the platform
- Decide number of processors to use?



Multi-core systems



How To:

- Deploy the application to the platform
- Decide number of processors to use?
- Allocate tasks to processors and schedule them

























Application Model

Task Graph





Application Model

Task Graph



• Tasks : Software procedure



Application Model

Task Graph



• Tasks : Software procedure

annotated with execution time



Application Model

Task Graph



- Tasks : Software procedure
- Edges : Precedence relations





- Tasks : Software procedure
- Edges : Precedence relations

Deployment Problem

Task Graph **Deployment Solution** D ^Processors в н P_2 G P_1 F Time С G

- Tasks : Software procedure
- Edges : Precedence relations

D

■ Mapping : Task ⇒ Processor

н



- Tasks : Software procedure
- Edges : Precedence relations

- Mapping : Task ⇒ Processor
- Scheduling : Task \Rightarrow Time













Solution space is large

ıΒ 10 . 6 6 2 proc., 10 tasks \approx 1000+ potential solutions . 6 6



Deployment problem

• The difficultly in the deployment is that the design space is exponential



- The difficultly in the deployment is that the design space is exponential
- One needs to model complex hardware : Processors, Network, DMA



- The difficultly in the deployment is that the design space is exponential
- One needs to model complex hardware : Processors, Network, DMA
- Multiple Evaluation Criteria
 - Latency
 - Memory used
 - Processors used
 - ...



Research Questions

How to:

model the software



Research Questions

How to:

- model the software
- model the hardware (Processors, Network, DMA)



Research Questions

How to:

- model the software
- model the hardware (Processors, Network, DMA)
- Optimize deployment while dealing with design space explosion



Outline





- Opployment using SMT
- Symmetry elimination
- Distributed memory scheduling
- 6 Design Tools

Conclusions


Overview

Motivation

- 2 Application Model
- Oeployment using SMT
- 4 Symmetry elimination
- Distributed memory scheduling
- Design Tools
- Conclusions



Model of Computation

Synchronous Dataflow graphs (SDF)

by Edward Lee and David Messerschmitt in 1987



Model of Computation

Synchronous Dataflow graphs (SDF)

by Edward Lee and David Messerschmitt in 1987

represents Streaming Applications



Model of Computation

Synchronous Dataflow graphs (SDF)

by Edward Lee and David Messerschmitt in 1987

represents Streaming Applications

















Synchronous DataFlow





Actors

- Pre-processing
- Blur
- Post-processing



Synchronous DataFlow





- Actors
 - Pre-processing
 - Blur
 - Post-processing

• Edges

Blur executes only after Pre-processing finishes



Synchronous DataFlow





- Actors
 - Pre-processing
 - Blur
 - Post-processing
- Edges
 - Blur executes only after Pre-processing finishes

Rates

- Pre-processing produces 4 pieces of an image (tokens)
- Each Blur consumes 1 piece



Synchronous DataFlow



Actor Blur is compact representation of data parallel tasks.





- Actor Blur is compact representation of data parallel tasks.
- All Blur tasks have **same properties** such as execution time.



Split-Join Graphs

we use split-join graphs : restriction of SDF

still covering perhaps 90% of use cases in the literature



Split-Join Graphs

we use split-join graphs : restriction of SDF

still covering perhaps 90% of use cases in the literature

a simple example:



- α : spawn and split
- $1/\alpha :$ wait and join



Split-Join Graphs

we use split-join graphs : restriction of SDF

still covering perhaps 90% of use cases in the literature

a simple example:



- $\alpha:$ spawn and split
- $1/\alpha :$ wait and join





Restrictions compared to general SDF





Restrictions compared to general SDF

Split-join does not support:

Stateful actors





Restrictions compared to general SDF

Split-join does not support:

- Stateful actors
- Non-proportional rates





Restrictions compared to general SDF

Split-join does not support:

- Stateful actors
- Non-proportional rates
- Initial tokens and cyclic paths





Overview

Motivation

- 2 Application Model
- Opployment using SMT
- 4 Symmetry elimination
- Distributed memory scheduling
- Design Tools
- Conclusions









- Boolean variables
 - in₀, in₁, in₂ ...
 - out_0 , out_1 , out_2 ...





- Boolean variables
 - in₀, in₁, in₂ ...
 - out_0 , out_1 , out_2 ...
- Constraints
 - $\operatorname{out}_0 = \operatorname{in}_0 \lor \operatorname{in}_1 \oplus \operatorname{in}_2 \ldots$





- Boolean variables
 - in₀, in₁, in₂ ...
 - out_0 , out_1 , out_2 ...
- Constraints
 - $\operatorname{out}_0 = \operatorname{in}_0 \lor \operatorname{in}_1 \oplus \operatorname{in}_2 \ldots$





- Boolean variables
 - in_0 , in_1 , in_2 ...
 - out_0 , out_1 , out_2 ...
- Constraints
 - $\operatorname{out}_0 = \operatorname{in}_0 \lor \operatorname{in}_1 \oplus \operatorname{in}_2 \dots$







- Boolean variables
 - in_0 , in_1 , in_2 ...
 - out_0 , out_1 , out_2 ...
- Constraints
 - $\operatorname{out}_0 = \operatorname{in}_0 \lor \operatorname{in}_1 \oplus \operatorname{in}_2 \dots$







- Boolean variables
 - in_0 , in_1 , in_2 ...
 - out_0 , out_1 , out_2 ...
- Constraints
 - $\operatorname{out}_0 = \operatorname{in}_0 \lor \operatorname{in}_1 \oplus \operatorname{in}_2 \dots$







- Boolean variables
 - in_0 , in_1 , in_2 ...
 - out_0 , out_1 , out_2 ...
- Constraints
 - $\operatorname{out}_0 = \operatorname{in}_0 \lor \operatorname{in}_1 \oplus \operatorname{in}_2 \dots$







- Boolean variables
 - in_0 , in_1 , in_2 ...
 - out_0 , out_1 , out_2 ...
- Constraints
 - $\operatorname{out}_0 = \operatorname{in}_0 \lor \operatorname{in}_1 \oplus \operatorname{in}_2 \dots$







- Boolean variables
 - in_0 , in_1 , in_2 ...
 - out_0 , out_1 , out_2 ...
- Constraints
 - $\operatorname{out}_0 = \operatorname{in}_0 \lor \operatorname{in}_1 \oplus \operatorname{in}_2 \dots$









Actor	Α	В				С
Tasks	A ₀	B ₀	B_1	B_2	B ₃	C ₀
Description	Variables					





Actor	Α	AB				С
Tasks	A ₀	B_0	B_1	B_2	B ₃	C ₀
Description	Variables					
Start time	xA_0	xB ₀	xB_1	xB_2	xB_3	xC_0





Actor	A B				С	
Tasks	A ₀	B_0	B_1	B_2	B_3	C ₀
Description	Variables					
Start time	xA ₀	xB ₀	xB_1	xB_2	xB_3	xC_0
Allocated proc.	pA ₀	pB ₀	pB_1	pB_2	pB_3	pC ₀





Actor	A B				С	
Tasks	A ₀	B_0	B_1	B_2	B ₃	C ₀
Description	Variables					
Start time	xA ₀	xB_0	xB_1	xB_2	xB_3	xC_0
Allocated proc.	pA ₀	pB_0	pB_1	pB_2	pB ₃	pC ₀
Duration	dA	dB				dC





Actor	A B				С	
Tasks	A ₀	B_0	B_1	B_2	B ₃	C ₀
Description	Variables					
Start time	xA ₀	xB_0	xB_1	xB_2	xB_3	xC_0
Allocated proc.	pA ₀	pB_0	pB_1	pB_2	pB ₃	pC ₀
Duration	dA	dB				dC

- Precedence Constraints
 - $xB_0 \ge (xA_0 + dA)$





Encoding deployment with constraints



В С Actor Α Tasks C_0 A_0 B_0 B_1 B₂ B_3 Description Variables Start time xA_0 xB₀ xB₁ xB₂ xB_3 xC_0 Allocated proc. $\mathbf{p}\mathbf{B}_1$ pB_2 pC_0 pA_0 pB_0 pB_3 Duration dA dB dC


Encoding deployment with constraints



Actor	A B					С
Tasks	A ₀	B_0	B_1	B_2	B_3	C ₀
Description	Variables					
Start time	xA ₀	xB_0	xB_1	xB_2	xB_3	xC ₀
Allocated proc.	pA ₀	pB_0	pB_1	pB_2	pB ₃	pC ₀
Duration	dA	dB				dC

- Precedence Constraints
 - $xB_0 \ge (xA_0 + dA)$
- Mutual Exclusion Constraints

• if
$$(pB_1 = pB_2)$$
 then
 $xB_1 \ge (xB_2 + dB) \lor xB_2 \ge (xB_1 + dE)$

Latency Cost

• Latency = $(\mathbf{x}\mathbf{C}_0 + \mathbf{d}\mathbf{C})$





Multi-criteria Problem



Latency





Multi-criteria Problem





Multi-criteria Problem



Tendulkar

Multi-criteria Problem



Multi-criteria Problem



















Problem Monotonicity

 $\begin{array}{l} \mbox{Latency} \leq 4 \\ \mbox{\#Proc} \leq 2 \\ \mbox{Not Possible} \end{array}$





Problem Monotonicity

 $\begin{array}{l} \text{Latency} \leq 4 \\ \#\text{Proc} \leq 2 \\ \textbf{Not Possible} \end{array}$

Latency = 2 #Proc = 1 Also Not Possible





Problem Monotonicity

 $\begin{array}{l} \text{Latency} \leq 4 \\ \#\text{Proc} \leq 2 \\ \textbf{Not Possible} \end{array}$

Latency = 2 #Proc = 1 Also Not Possible







Design Space Exploration

Split-join Graph



























Design Space Exploration



Timeout: Cannot decide SAT / UNSAT in a given TIME-BUDGET.







Divide cost space using grids



• sat points • unsat points • not yet explored points



- Divide cost space using grids
- One SMT query per point on the grid



• sat points • unsat points • not yet explored points



- Divide cost space using grids
- One SMT query per point on the grid
- Finer grid after every iteration



sat points
 unsat points
 not yet explored points



- Divide cost space using grids
- One SMT query per point on the grid
- Finer grid after every iteration
- Don't query in known area



sat points
 unsat points
 ont yet explored points



Overview

Motivation

- 2 Application Model
- 3 Deployment using SMT
- 4 Symmetry elimination
- Distributed memory scheduling
- 6 Design Tools
- Conclusions



Task Symmetry





















- all instances of actor C are similar (symmetric)
- No change in latency !





- all instances of actor C are similar (symmetric)
- No change in latency !
- Huge number of such symmetric solutions




- all instances of actor C are similar (symmetric)
- No change in latency !
- Huge number of such symmetric solutions
- Add constraints to eliminate all but one









• lexicographic order : $C_{00} \ll C_{01} \ll C_{10} \ll C_{11}$





- lexicographic order : $C_{00} \ll C_{01} \ll C_{10} \ll C_{11}$
- enforce lexicographic order in schedule: s(u) < s(u') for $u \ll u'$





- lexicographic order : $C_{00} \ll C_{01} \ll C_{10} \ll C_{11}$
- enforce lexicographic order in schedule: $s(u) \leq s(u')$ for $u \ll u'$
- $s(\mathbf{C}_{00}) \le s(\mathbf{C}_{01}) \le s(\mathbf{C}_{10}) \le s(\mathbf{C}_{11})$





- lexicographic order : $C_{00} \ll C_{01} \ll C_{10} \ll C_{11}$
- enforce lexicographic order in schedule: $s(u) \leq s(u')$ for $u \ll u'$
- $s(\mathbf{C}_{00}) \le s(\mathbf{C}_{01}) \le s(\mathbf{C}_{10}) \le s(\mathbf{C}_{11})$

















Task Symmetry : Theorem



• Theorem : Every group has a lexicographic schedule





- Theorem : Every group has a lexicographic schedule
- Corollary : No feasible schedule is lost













Pareto Exploration

Exploration : Processors vs Latency $\alpha = 30$



Pareto Exploration



without symmetry breaking

Exploration : Processors vs Latency $\alpha = 30$



Pareto Exploration



without symmetry breaking

with symmetry breaking

Exploration : Processors vs Latency $\alpha = 30$



Pareto Exploration



without symmetry breaking

with symmetry breaking

Exploration : Processors vs Latency $\alpha = 30$

Solver Performance Timeouts reduce ! The gap between SAT and UNSAT points is smaller.

Video Decoder

3D cost space $(\mathbf{C}_L, \mathbf{C}_P, \mathbf{C}_B)$ exploration, \mathbf{C}_B - total buffer size

MPEG video decoder:





Video Decoder

3D cost space $(\mathbf{C}_L, \mathbf{C}_P, \mathbf{C}_B)$ exploration, \mathbf{C}_B - total buffer size

MPEG video decoder:





۰

with symmetry constraints

without symmetry constraints



Video Decoder

3D cost space $(\mathbf{C}_L, \mathbf{C}_P, \mathbf{C}_B)$ exploration, \mathbf{C}_B - total buffer size





Video Decoder

3D cost space $(\mathbf{C}_L, \mathbf{C}_P, \mathbf{C}_B)$ exploration, \mathbf{C}_B - total buffer size

MPEG video decoder:





۰

with symmetry constraints

without symmetry constraints

V

Better Pareto points

Tendulkar

Video Decoder

3D cost space $(\mathbf{C}_L, \mathbf{C}_P, \mathbf{C}_B)$ exploration, \mathbf{C}_B - total buffer size





۰

with symmetry constraints

without symmetry constraints

Better Pareto points in same TIME-Budget !



Distributed memory scheduling



Distributed memory scheduling

• So far we ignored the communication costs



Distributed memory scheduling

- So far we ignored the communication costs
- For distributed memory, communication needs to be modeled



Overview

Motivation

- 2 Application Model
- 3 Deployment using SMT
- Symmetry elimination
- Distributed memory scheduling
 - Design Tools

Conclusions



Kalray MPPA-256

512 KB	USMC		PCle	inter laken		DDR
Quad Core						GPIOs
E.						_
Inter Iaken						Inter
Quad Core						Quad
512 KB						KB 512
DDR						Quad Core
GPIOs		PCle		interlaken		512 KB
L						



Kalray MPPA-256



• 16 compute clusters



Kalray MPPA-256



• 16 compute clusters



Kalray MPPA-256



16 compute clusters
 16 processors



Kalray MPPA-256



- 16 compute clusters
 - 16 processors
 - 2 MB Shared Memory



Kalray MPPA-256



• 16 compute clusters

- I6 processors
- 2 MB Shared Memory
- DMA



Kalray MPPA-256



• 16 compute clusters

- 16 processors
- 2 MB Shared Memory
- DMA
- Toroidal 2D network



Kalray MPPA-256



- 16 compute clusters
 - 16 processors
 - 2 MB Shared Memory
 - DMA
- Toroidal 2D network
Design Flow

Application Graph



Design Flow





Design Flow



Goals

- Load balance the groups
- Minimize data exchange



Design Flow





Design Flow



Goals

Minimize distance between communicating groups

Design Flow





Design Flow





Design Flow



Goals

- Minimize Latency
- Minimize Buffer size





Tasks and Transfers



- Tasks and Transfers
 - Cluster Mapping



- Tasks and Transfers
 - Cluster Mapping
 - Processor and DMA Mapping



- Tasks and Transfers
 - Cluster Mapping
 - Processor and DMA Mapping
 - Start time



- Tasks and Transfers
 - Cluster Mapping
 - Processor and DMA Mapping
 - Start time
- Edges



- Tasks and Transfers
 - Cluster Mapping
 - Processor and DMA Mapping
 - Start time
- Edges
 - Communication buffer size



- Tasks and Transfers
 - Cluster Mapping
 - Processor and DMA Mapping
 - Start time
- Edges
 - Communication buffer size
- Application



- Tasks and Transfers
 - Cluster Mapping
 - Processor and DMA Mapping
 - Start time
- Edges
 - Communication buffer size
- Application
 - Latency















Task	Description	Resources used	Task duration	
I	Initialization	Processor and DMA	Constant	







Task	Description	Resources used	Task duration	
I	Initialization	Processor and DMA	Constant	
G	Network Transfer	Only DMA	Transfer size dependent	



Model Transformation

An example application graph:





Model Transformation

An example application graph:





Model Transformation

An example application graph:



Partition-Aware graph:





Model Transformation

An example application graph:



Partition-Aware graph:







Model Transformation

An example application graph:



Partition-Aware graph:







Model Transformation

An example application graph:



Partition-Aware graph:







Model Transformation

An example application graph:



Partition-Aware graph:







JPEG Decoder Example





JPEG Decoder Example



VLD : Variable Length Decoder



JPEG Decoder Example



VLD : Variable Length Decoder

IQ / IDCT : Inverse Quantization / Inverse Discrete Cosine Transform



JPEG Decoder Example



VLD : Variable Length Decoder

IQ / IDCT : Inverse Quantization / Inverse Discrete Cosine Transform

Color : Color Conversion





- $C_{ au}$: Max. workload per group
- C_{η} : Total communication cost
- C_z : No. of Groups





- $C_{ au}$: Max. workload per group
- C_{η} : Total communication cost
- C_z : No. of Groups

Solution	Allocated group			Exploration Cost		
Solution	vld	iq	color	$\mathrm{C}_{ au}$	C_η	$\mathbf{C}_{\mathbf{z}}$
P_{s0}	0	1	2	424012	12384	3
P_{s1}	0	0	1	758116	2736	2
P_{s2}	0	0	0	934288	0	1
P_{s3}	0	1	1	510276	9648	2





- $C_{ au}$: Max. workload per group
- C_{η} : Total communication cost
- C_z : No. of Groups

Solution	Allocated group			Exploration Cost		
Solution	vld	iq	color	$\mathrm{C}_{ au}$	C_η	$\mathbf{C}_{\mathbf{z}}$
P_{s0}	0	1	2	424012	12384	3
P_{s1}	0	0	1	758116	2736	2
P_{s2}	0	0	0	934288	0	1)
P_{s3}	0	1	1	510276	9648	2





- $C_{ au}$: Max. workload per group
- C_{η} : Total communication cost
- C_z : No. of Groups

Solution	Allocated group			Exploration Cost		
Solution	vld	iq	color	$\mathrm{C}_{ au}$	C_η	$\mathbf{C}_{\mathbf{z}}$
P_{s0}	0	1	2	424012	12384	3
P_{s1}	0	0	1	758116	2736	2
P_{s2}	0	0	0	934288	0	1
P_{s3}	0	1	1	510276	9648	2


Motivation Application Model Deployment using SMT Symmetry elimination Distributed memory scheduling Design Tools Conclusions

JPEG Decoder Example



Tendulkar

JPEG Decoder Example

JPEG decoder latency on Kalray platform



StreamIt Benchmarks



StreamIt Benchmarks



Overview

Motivation

- 2 Application Model
- 3 Deployment using SMT
- 4 Symmetry elimination
- Distributed memory scheduling

6 Design Tools

7 Conclusions





Runtime













































StreamExplorer

































StreamExplorer

Written in Java



- Written in Java
- 32k+ lines of Code.



StreamExplorer

- Written in Java
- 32k+ lines of Code.

Runtime

Written in C++



StreamExplorer

- Written in Java
- 32k+ lines of Code.

- Written in C++
- 14k+ lines of Code.



Overview

Motivation

- 2 Application Model
- 3 Deployment using SMT
- 4 Symmetry elimination
- Distributed memory scheduling
- Design Tools





Conclusions and Future Work

Conclusions:



Conclusions and Future Work

Conclusions:

• Symmetry elimination finds better solutions



Conclusions and Future Work

Conclusions:

- Symmetry elimination finds better solutions
- Combined Optimization with Communication modeling


Conclusions:

- Symmetry elimination finds better solutions
- Combined Optimization with Communication modeling
- Automated design flow for distributed memory



Conclusions:

- Symmetry elimination finds better solutions
- Combined Optimization with Communication modeling
- Automated design flow for distributed memory



Conclusions:

- Symmetry elimination finds better solutions
- Combined Optimization with Communication modeling
- Automated design flow for distributed memory

Future Work:

Spread actor over multiple clusters



Conclusions:

- Symmetry elimination finds better solutions
- Combined Optimization with Communication modeling
- Automated design flow for distributed memory

- Spread actor over multiple clusters
- Network route selection and scheduling



Conclusions:

- Symmetry elimination finds better solutions
- Combined Optimization with Communication modeling
- Automated design flow for distributed memory

- Spread actor over multiple clusters
- Network route selection and scheduling
- Pipelined scheduling



Conclusions:

- Symmetry elimination finds better solutions
- Combined Optimization with Communication modeling
- Automated design flow for distributed memory

- Spread actor over multiple clusters
- Network route selection and scheduling
- Pipelined scheduling
- Scheduling under uncertainty



Contributions

- P. Tendulkar, P. Poplavko, and O. Maler. "Symmetry Breaking for Multi-criteria Mapping and Scheduling on Multicores". In: FORMATS. 2013
- P. Tendulkar, P. Poplavko, I. Galanommatis, and O. Maler. "Many-Core Scheduling of Data Parallel Applications using SMT Solvers". In: DSD. 2014
- P. Tendulkar, P. Poplavko, and O. Maler. Strictly Periodic Scheduling of Acyclic Synchronous Dataflow Graphs using SMT Solvers. Tech. rep. Verimag Research Report, 2014
- P. Tendulkar and S. Stuijk. "A Case Study into Predictable and Composable MPSoC Reconfiguration". In: IPDPS RAW Workshop. 2013
- S. Saidi, P. Tendulkar, T. Lepley, and O. Maler. "Optimizing Explicit Data Transfers for Data Parallel Applications on the Cell Architecture". In: ACM TACO (2012)
- S. Saidi, P. Tendulkar, T. Lepley, and O. Maler. "Optimizing two-dimensional DMA transfers for scratchpad Based MPSoCs platforms". In: *Microprocessors and Microsystems* (2013)



Motivation Application Model Deployment using SMT Symmetry elimination Distributed memory scheduling Design Tools Conclusions

Thank You



Questions?



Overview

- SDF and Split Join graphs
- Symmetry Breaking
- 10 Design Flow Details
- DMA transfer granularity
- 12 Run-time Management

Split-Join Graphs

Hypothesis supported by StreamIt:1

- Total 763 actors analyzed in various applications
 - 94% are stateless
 - 6% are stateful
 - 45% have states due to algorithm
 - 55% have avoidable states
- Odd rates exists but are rare
 CD-DAT benchmark used as an example

Converts CD audio (44.1 kHz) to digital audio tape (48 kHz)

[1] W. Thies and S. Amarasinghe. "An Empirical Characterization of Stream Programs and Its Implications for Language and Compiler

Design". In: PACT. 2010

Overview



- Symmetry Breaking
- 10 Design Flow Details
- 10 DMA transfer granularity
- 12 Run-time Management



Proof Sketch



modify a feasible schedule such that: $s(v_0) \le s(v_1) \le s(v_2) \le ...$

prove that precedence constraints are satisfied

 \checkmark here: for neutral channels (α = 1), unfolded to (v_h, v'_h)





start-time compatible



new hier. index; , new precedence relation



Proof Sketch



modify a feasible schedule such that: $s(v_0) \le s(v_1) \le s(v_2) \le ...$

prove that precedence constraints are satisfied

 \checkmark here: for neutral channels (α = 1), unfolded to (v_h, v'_h)



lexicographic order



start-time compatible



new hier. index; new precedence relation









Proof Sketch





take successor [j]



Proof Sketch





take successor [j]



Proof Sketch



take successor [j] by definition there exist j + 1 same or earlier successors



Proof Sketch



take successor [j] by definition there exist j + 1 same or earlier successors



Proof Sketch



take successor [j] by definition there exist j + 1 same or earlier successors their original predecessors finish before successor [j]:



Proof Sketch



take successor [j] by definition there exist j + 1 same or earlier successors their original predecessors finish before successor [j]:



Proof Sketch



take successor [j] by definition there exist j + 1 same or earlier successors their original predecessors finish before successor [j]:



Proof Sketch



take successor [j] by definition there exist j + 1 same or earlier successors their original predecessors finish before successor [j]: j + 1 predecessors finish before, hence the earliest j + 1 ones as well



Proof Sketch



take successor [*j*] by definition there exist j + 1 same or earlier successors their original predecessors finish before successor [*j*]: j + 1 predecessors finish before, hence the earliest j + 1 ones as well predecessor [*j*] finishes before successor [*j*]



Overview

- 8 SDF and Split Join graphs
- Symmetry Breaking
- 0 Design Flow Details
- DMA transfer granularity
- 12 Run-time Management



Design Flow





Overview

- 8 SDF and Split Join graphs
- Symmetry Breaking
- 10 Design Flow Details
- 1 DMA transfer granularity
- 12 Run-time Management

Buffering Algorithm

For architectures with DMA and limited local memory





Data transfer granularity



Data transfer granularity

Additional complexity with multiple processors





DMA transfer granularity optimization





Experimental Evaluation

Characterization of DMA of IBM Cell B.E.:





Experimental Evaluation

Synthetic Application Benchmark:





Overview



12 Run-time Management



The context



- Multiple configurations for each application
- Applications start / stop dynamically
- How to:
 - select a configuration for each application?
 - re-configure the applications?



CompSoC platform



Features:

- CompOSe real-time operating system
- Predictable Æthereal network-on-chip
- TDM application scheduling for composability
- **composable**: The changes in an application don't affect other running applications



Resource Manager conceptual view



Resource manager Design:

- System RM : takes re-configuration decisions
- Application RM : implements re-configuration decisions



Resource Manager on the platform



Resource manager Implementation:

- System RM: is a separate application
- Application RM:
 - organized in master-slave(s) configuration
 - is a part of user application


Experiment with JPEG Decoder

