

Design and Validation of Application-specific Multiprocessor Platforms based on Networks-on-Chip

Leandro Soares Indrusiak

lsi@cs.york.ac.uk http://www-users.cs.york.ac.uk/lsi

CREDES Workshop – Tallinn - September 2010



About the speaker

- Leandro Soares Indrusiak
 - born in Santa Maria, RS, Brazil in 1974
 - Electrical Engineer, UFSM, Santa Maria, 1995
 - MSc Computer Science, UFRGS, Porto Alegre, 1998
 - Dr.-Ing Computer Science, jointly issued by UFRGS and TU Darmstadt (Germany), 2003
 - 1998 2000 Lecturer in Computer Science
 PUCRS, Uruguaiana, Brazil
 - 2003 2008 Research Fellow in Microelectronics Design TU Darmstadt, Germany
 - 2008 Lecturer in Real Time Embedded Systems University of York, UK



Application-specific Multiprocessor Platforms based on NoCs | L. S. Indrusiak

York - Location and history

- York is 200 miles north of London
 - London: 2 hours

THE UNIVERSITY of York

- Manchester airport: 1.5 hours
- Ancient city historic capital of north with over 2,000 years of history and heritage, and outstanding quality of life









Heslington West





Heslington East – Phase 1





York's position

THE UNIVERSITY of york

UK top 10 World top 100 UK University of the Year 2003 Queen's Anniversary Prize 1996, 2006 & 2008, 2010 Computer Science, Plant Biology, Health Economics, and Social Policy

"One of Britain's academic success stories, forging a reputation to rival Oxford and Cambridge in the space of 40 years" Sunday Times

"The Cambridge of the North"

Financial Times







Computer Science at York

In numbers

THE UNIVERSITY of york

- 38 full-time academics (lecturers, senior lecturers, readers, professors)
- 400 undergraduates (BEng, MEng)
- 130 taught postgraduates (MSc)
- 100 research students (PhD)
- Consistently ranked among the top CS departments in the UK
- Teaching judged "Excellent" by HEFCE
- Majority of the department's research judged "world leading" or "internationally excellent"







Computer Science at York - Teaching

Undergraduates

Programmes
BSc/BEng degree programmes (3 year)
Computer Science
Computer Science with Embedded Systems
Computer Science/Mathematics (Equal)
BSc/BEng degree programmes (4 year sandwich)
Computer Science
Computer Science with Embedded Systems
Computer Science/Mathematics (Equal)
MEng/MMath degree programmes (4 year)
Computer Science with Embedded Systems
Computer Science and Software Engineering
Mathematics / Computer Science (Equal)
MEng/MMath degree programmes (5 year sandwich)
Computer Science with Embedded Systems
Computer Science and Software Engineering

Mathematics / Computer Science (Equal)







Computer Science at York - Teaching

Graduates

MRes in Computational Biology

MSc/Diploma/Certificate in Gas Turbine Control

MSc in Information Technology

MSc in Natural Computation

MSc/Diploma/Certificate in Safety Critical Systems Engineering

MSc in Software Engineering

Certificate in System Safety Engineering

MSc in Computer Science (by research)

MPhil and PhD







THE UNIVERSITY of Jork Application-specific Multiprocessor Platforms based on NoCs | L. S. Indrusiak

Computer Science at York – Research







Design and Validation of Application-specific Multiprocessor Platforms based on Networks-on-Chip

Leandro Soares Indrusiak

lsi@cs.york.ac.uk http://www-users.cs.york.ac.uk/lsi

CREDES Workshop – Tallinn - September 2010



Outline

THE UNIVERSITY of york

- On-Chip Multiprocessing
 - motivation, review
- Application-specific Multiprocessor Platforms
 - validation, concurrency issues, application-platform mapping
- Platform models
 - evaluation, accuracy
- Application models
 - concurrency
- Joint execution of application and platform models
 - case study
- Background theory and tool support
 - actor orientation, Ptolemy II, custom made extensions
 - demo



Outline

- On-Chip Multiprocessing
 - motivation, review
- Application-specific Multiprocessor Platforms
 - validation, concurrency issues, application-platform mapping
- Platform models
 - evaluation, accuracy
- Application models
 - concurrency
- Joint execution of application and platform models
 - case study
- Background theory and tool support
 - actor orientation, Ptolemy II, custom made extensions
 - demo



The problem with CMOS

- CMOS limits can be felt in many ways
 - number of transistors
 - clock frequency
 - power dissipation
- Expectations of performance increase must be met somehow
 - on-chip multiprocessing seems to be the best shot





THE UNIVERSITY of fork Application-specific Multiprocessor Platforms based on NoCs | L. S. Indrusiak
On-chip Multiprocessing
PE



Intel 8086 (1978)



(2008)

On-chip Multiprocessing

- Many architectural alternatives
 - homogeneous vs. heterogeneous processing
 - shared memory vs. distributed memory
 - on-chip interconnect
 - point-to-point, crossbar
 - on-chip bus
 - network-on-chip

	· _1)
	一般的问题
	n de ge de g
IL SIL SIL SIL S	周 我看 我看 我看 我
aniarani.	

source: Intel



Cell Processor

 nine processing elements: one Power Processor Element (PPE) and eight Synergistic Processing Elements (SPE)

• PPE has separated I & D L1 cache (32KB each)

• Each SPE can only access its 256KB of local storage (LS) and uses its memory flow controller (MFE) to perform DMA operations to/from LS (non-blocking)

- Bandwidth (3.2 GHz)
 - SPE <-> LS = 2x 25.6 GB/s
 - MFC <-> EIB = 2x 25.6 GB/s
 - MIC <-> EIB = 2x 25.6 GB/s

• L1 <-> L2 = 2x 51.2 GB/s





- ARM Cortex-9 MPCore
 - can have 1-4 Cortex-A9 cores with separated I & D L1 cache (32KB each)
 - data caches of all cores are fully coherent
 - interface to external components be made cache-coherent
 - on-chip interconnect based on AMBA standard





ARM Cortex-9 MPCore – TI OMAP 4 Platform





Tilera TILEPro64





- As the number of processing elements increase, on-chip communication becomes a major issue
 - point-to-point: huge area overhead
 - on-chip bus: doesn't scale (restricted parallelism)
 - in both cases: long wires cause timing issues and excessive power consumption







Networks-on-Chip (NoC)

- multi-hop, packet-based communication
- scalable, high bandwidth
- Iow power (shorter wires)
- distributed arbitration
- regular, reusable
- supports GALS paradigm
- communication latency
- ▶ area
- complex to design





Networks-on-Chip (NoC)





Outline

THE UNIVERSITY of york

- On-Chip Multiprocessing
 - motivation, review
- Application-specific Multiprocessor Platforms
 - validation, concurrency issues, application-platform mapping
- Platform models
 - evaluation, accuracy
- Application models
 - concurrency
- Joint execution of application and platform models
 - case study
- Background theory and tool support
 - actor orientation, Ptolemy II, custom made extensions
 - demo



Application-specific Multiprocessor Platforms

Fundamental questions

- can I meet the constraints of a given application A if I run it over platform P?
 - performance and timing
 - power
- no? isn't there a way to optimise it?
 - mapping and scheduling
 - voltage and frequency scaling
- yes? can I guarantee it?
 - hard/soft real-time



THE UNIVERSITY of York Application-specific Multiprocessor Platforms based on NoCs | L. S. Indrusiak

Application-specific Multiprocessor Platforms

- Many architectural alternatives
 - Iarge design space
 - hard to evaluate which alternative platform is the optimal match for the requirements imposed by a given application



Alberto Sangiovanni-Vincentelli

 Application-platform mapping plays a critical role



Application-specific Multiprocessor Platforms

 Application-Platform mapping isn't really a problem when dealing with sequential software and uniprocessor hardware





THE UNIVERSITY of York Application-specific Multiprocessor Platforms based on NoCs | L. S. Indrusiak

Application-specific Multiprocessor Platforms

207

 It has been studied for decades by researchers in parallel and distributed computing



On the Mapping Problem

SHAHID H. BOKHARI, MEMBER, IEEE

Abstract—In array processors it is important to map problem modules onto processors such that modules that communicate with each other lies, as for as possible, on adjacent processors. This mapping problem is formulated in graph theoretic terms and shown to be equivalent, in its most general form, to the graph isomorphism problem. The problem is also very similar to the bandwidth reduction problem for sparse matrices and to the quadratic assignment problem. It appears unlikely that an efficient exact algorithm for the general mapping problem will ever be found. Research in this area must com-

IEEE TRANSACTIONS ON COMPUTERS, VOL. C-30, NO. 3, MARCH 1981

It appears unlikely that an efficient exact algorithm for the general mapping problem will ever be found. Research in this area must concentrate on efficient heuristics that find good solutions in most cases. A heuristic algorithm that proceeds by sequences of pairwise interchanges alternating with probabilistic jumps is described. This algorithm has been used to solve practical mapping problems on a specific array processor (the finite element machine) with good results. Results for a set of practical problems are tabulated, several of which are illustrated.

Index Terms—Adjacency matrices, array processing, assignment, computer networks, distributed processors, finite element machine, graph isomorphism, heuristic algorithm, mapping problem, pairwise interchange.

I. INTRODUCTION

M OST arrays of processors are incompletely connected, that is, a direct link does not connect each pair of processors. The reasons for this include: 1) the fact that the total number of links in completely connected systems increases as the square of the number of processors—a growth rate that is unacceptable in most cases, and 2) the number of a heurstic algorithm that has been developed to solve this problem for a specific array processor. We start by giving a mathematical formulation of the problem in Section II. In Section III we show that in its most general form the mapping problem is equivalent to the graph isomorphism problem, one of the classical unsolved combinatorial problems. We point out the similarities between the mapping problem and the bandwidth reduction and quadratic assignment problems. Exact solutions for neither of these problems exist and they are solved approximately using heuristic algorithms.

In Section IV we describe how the mapping problem arises when solving structural problems on the finite element machine (FEM), an array of processors currently under development at NASA Langley Research Center. In Section V we describe a simple heuristic algorithm that has been implemented and used to find mappings for the finite element machine with very encouraging results. Results for a number of test cases are tabulated, several of which are illustrated.

II. MATHEMATICAL FORMULATION

Let the graph of the problem to be mapped onto the array be denoted $G_p = \langle V_p, E_p \rangle$, where the nodes or vertices V_p correspond to the set of modules and each edge $(x, y) \in E_p$

RTSyork 28

Application-specific Multiprocessor Platforms

 Application-Platform mapping is a critical issue when it comes to multiprocessor platforms

REALLY?



RTS Jork 29

Application-specific Multiprocessor Platforms

THE UNIVERSITY of York

Sources: Gordon ASPLOS'06, Kudlur PLDI'08





application

Application-specific Multiprocessor Platforms

 Application-Platform mapping is a critical issue when it comes to multiprocessor platforms



must be able to explore concurrency at the application level



platform

application

Application-specific Multiprocessor Platforms

 Application-Platform mapping is a critical issue when it comes to multiprocessor platforms



can be done dynamically to improve performance or to increase dependability



platform

Application-Platform Mapping

- The simplest formulation of the mapping problem resembles a graph isomorphism problem
 - Papplication is a graph G = G(A, C), where a_i∈ A is an application task and c_{i,i} ∈ C represents the communication from a_i to a_i
 - platform is a graph Gⁱ = G(B, D), where b_i ∈ B is a processor and d_{i,j} ∈ D represents the channel from b_i to b_i
 - objective is to map tasks onto processors such that task that communicate with each other lie on adjacent processors
- There is no known polynomial time solution for the graph isomorphism problem – NP-hard



Application-Platform Mapping

 More sophisticated problem formulations may include additional information to the application and platform graphs, so that different objectives can be met

```
platform:
processing power of b<sub>i</sub>
power consumption and latency of d<sub>i,i</sub>
```

application: computational cost and deadline of a_i volume, max latency and required bandwidth of c_{i,i}

 objectives: reduce bandwidth requirements on the platform, reduce power consumption, balance thermal dissipation, etc.



Application-Platform Mapping

- Platform models over-simplify the complex design space of on-chip multiprocessor platforms
- Application and Platform models disregard the temporal dimension
 - must include further details on time and concurrency



Outline

THE UNIVERSITY of york

- On-Chip Multiprocessing
 - motivation, review
- Application-specific Multiprocessor Platforms
 - validation, concurrency issues, application-platform mapping
- Platform models
 - evaluation, accuracy
- Application models
 - concurrency
- Joint execution of application and platform models
 - case study
- Background theory and tool support
 - actor orientation, Ptolemy II, custom made extensions
 - demo


- How to evaluate the quality of platform alternatives?
- HDL-based approach
 - cycle-accurate model of the platform
 - hard to create, debug and modify
 - synthesisable
 - traffic generators
 - DE simulator

- MASTER																														
-🔶 data_out	001A		1		1						F			0	• •		00)1A												
- SLAVE1																														
-🔶 data_in	0000	8 1	0000		00) () ()		L		0)		0)0)(0000		00)00												
-🔶 page	0	0	1	()(1	<u>))1)</u>	<u>r m</u>	2000))1	<u>XI</u>)1)	(1)	10	0		0	Ш	01	0	110		0)	0	10)				
🔶 intr_signal	0	1		1	1 1	11	T II	1	1	1					11		П													
- SLAVE2																														
-🔶 data_in	0000	0000	j (0000			XX	0		I X				0		100	000	0.)0.)0.			. 0)					0 <mark>I</mark>
📣 page	1	0) (0)	0 0	0	0	0)	0 0)0				10		m	1)))	T))	XX	001	χŋ	XΊ)))1))1	01	αx	m
🔶 intr_signal	0		11		1	111	1.11	11	11	TT	11 1	ΤĪ	Π		11		Ш				T		1	1	1	1				
- SLAVE3																														
-🔶 data_in	0000	0000	0000			0	000	1	Т	0	Υ			0000	1 I	Т	0.	T	11	ľ	Ĩ	1	1	1	1	X	X.	i I	I	
-🔶 page	0	10					I110	10	0)	0 10))0	10)O	10		0)	0 Y Y	WY.))0	10	10	10	10	10	10	0	0)	0)	0 1	0 10
📥 intr signal	0				11				Τ	1	1	1	1		11	<u> </u>	III		I II		Ī		11		11	ĪĪ	ĪĪ		11	
- SLAVE4																														
-🔶 data_in	0000	(0000	(0000		<u>)</u> 0000										000)		$\Box \mathfrak{I}$	X	I)	I)	I)	н	I)I	00	00	
-🔶 page	0	0))0	0)0)0)0)0)0)0	10	0	0	0)	0)
Intr_signal	0					- 11								1			111		I I	Π		L II		11	П	Ш	П	TT	II	
- SLAVE5																														
-🔶 data_in	0000	0000	0000		0000	1	0000								••		ŧŒ	000	I.)())())0)0)	T	T
-🔶 page	0	10																110	0 XC) (0	10	10	10	10	0	0	0	0 1	0)	0)
🔶 intr_signal	0				1										11		H									1				



- How to evaluate the quality of platform alternatives?
- Abstract NoC models
 - cycle/flit-level accuracy
 - PAT (payload abstraction)
 - static analysis
 - hybrid
- Modelling trade-offs
 - accuracy
 - observability
 - simulation speed













THE UNIVERSITY of york

Simplified models

- ► PAT
 - packet-level accuracy, 1-position buffers
 - complete packet is abstracted by header and trailler
 - header's way through the network is fully simulated, trailler latencies are calculated
 - less than 5% error for average latency (compared with cycle-accurate HDL)
- real-time analysis
 - priority-based virtual channels
 - static analysis of worst-case interference between network flows
- hybrid model transaction-level (TLM)
 - priority-based virtual channels
 - analysis of interference between network flows
 - less than 10% error for average and worst case latency (compared with cycleaccurate model)



THE UNIVERSITY of York

Simplified models



application	simulation time (s)					
flows	cycle-accurate	TLM				
20	17105.42	5.47				
40	31490.13	19.17				
60	55561.87	45.78				
80	70231.76	83.06				
100	86626.07	105.28				



application	simulation time (s)					
time (s)	cycle-accurate	TLM				
1	1465.18	7.67				
2	2895.69	7.89				
4	5978.32	7.98				
8	11806.91	8.28				
20	29927.24	9.77				
200		26.55				
800		86.29				
1400		169.06				







But how useful are those traffic generators?



critical point if we are targeting application-specific platforms



Can we have models of actual processors, OS and application software?



> yes, but it will be slow to simulate and hard to debug/modify



We use an application model instead





Course Outline

THE UNIVERSITY of york

- On-Chip Multiprocessing
 - motivation, review
- Application-specific Multiprocessor Platforms
 - validation, concurrency issues, application-platform mapping
- Platform models
 - evaluation, accuracy
- Application models
 - concurrency
- Joint execution of application and platform models
 - case study
- Background theory and tool support
 - actor orientation, Ptolemy II, custom made extensions
 - demo



Main focus on

THE UNIVERSITY of York

- concurrency model
- inter-process communication behaviour
- Executable





THE UNIVERSITY of Jork Application-specific Multiprocessor Platforms based on NoCs | L. S. Indrusiak

Application models

- Many concurrent programming models available
 - threads
 - concurrent processes with message-passing
 - actors
 - streams/dataflow
 - CSP
 - timed automata
 - <add your favorite here>





 The choice of which concurrent programming model should be adopted depends on

- application domain
- familiarity by the designer/developer
- availability of stable flows and tools
- predictability, and ability to match underlying HW/SW

model	tools
threads	libraries in Java and C++
message passing	OpenMP
actors	Simulink, Ptolemy
streams / dataflow	Streamlt, CUDA
timed automata	UPAAL



THE UNIVERSITY of York





THE UNIVERSITY of fork



- captures the concurrent behaviour of the application
- completely independent from the platform model
- for a given mapping, it can be jointly executed with the platform model
 - joint execution allows back-annotation of performance and power consumption figures



Course Outline

THE UNIVERSITY of york

- On-Chip Multiprocessing
 - motivation, review
- Application-specific Multiprocessor Platforms
 - validation, concurrency issues, application-platform mapping
- Platform models
 - evaluation, accuracy
- Application models
 - concurrency
- Joint execution of application and platform models
 - case study
- Background theory and tool support
 - actor orientation, Ptolemy II, custom made extensions
 - demo















Case study

THE UNIVERSITY of York

- Application model
 - autonomous vehicle
 - model based on actors and message sequence charts



Platform models

- analytical
- hybrid
- cycle-accurate



Case study

- Application model
 - autonomous vehicle
 - model based on actors and message sequence charts

- Platform models
 - analytical
 - hybrid
 - cycle-accurate





THE UNIVERSITY of York Application-specific Multiprocessor Platforms based on NoCs | L. S. Indrusiak

Case study





THE UNIVERSITY of fork Application-specific Multiprocessor Platforms based on NoCs | L. S. Indrusiak

Case study



Mapping 1 3x3	Mapping 2 3x3	Mapping 3 3x3
---------------	---------------	---------------

Interaction	Communication Latency (ms)							
Interaction	Abstract	Hybrid	Cycle-accurate					
OR	0.04	0.09	0.12					
TPA	0.07	0.09	0.14					
DA	0.24	0.27	0.39					
SR	0.68	0.94	1.35					
Р	0.71	0.99	1.41					



Course Outline

THE UNIVERSITY of york

- On-Chip Multiprocessing
 - motivation, review
- Application-specific Multiprocessor Platforms
 - validation, concurrency issues, application-platform mapping
- Platform models
 - evaluation, accuracy
- Application models
 - concurrency
- Joint execution of application and platform models
 - case study
- Background theory and tool support
 - actor orientation, Ptolemy II, custom made extensions
 - demo



- Modeling and simulation framework developed at UC Berkeley
- Focus on modeling concurrent components

http://ptolemy.eecs.berkeley.edu





THE UNIVERSITY of fork Application-specific Multiprocessor Platforms based on NoCs | L. S. Indrusiak

Ptolemy II





THE UNIVERSITY of York Application-specific Multiprocessor Platforms based on NoCs | L. S. Indrusiak

Ptolemy II





THE UNIVERSITY of York





Data Types







Functional composition





Concurrency in Ptolemy II

- Concurrency modeling is a primary goal of Ptolemy II
- Based on the concept of actor-orientation
 - explicit definition of the model of time and concurrency
 - support heterogeneous models
 - support domain polymorphism





Concurrency in Ptolemy II

THE UNIVERSITY of York




- According to Atkinson, Hewitt (MIT, 1977)
 - Conceptually an actor is an object which has both procedural and data aspects
 - A process is a totally ordered sequence of computational events where each event consists of sending a message to an actor
 - The actor model of computation has encouraged the development of a paradigm based on a society of experts communicating by passing messages

- According to Gul Agha (MIT, 1986)
 - Actors are computational agents which may function in parallel (have their own thread of control) and communicate through asynchronous message passing
 - The idea behind actor languages is to provide the syntactic constructs that can free the programmer from having to worry about the details of a concurrent execution



- How to free the programmer from having to worry about the details of a concurrent execution?
- One way is to remove unnecessary data dependencies created by the assignment operation (von Neumann bottleneck)





THE UNIVERSITY of York Application-specific Multiprocessor Platforms based on NoCs | L. S. Indrusiak

- Functional programming allows for the possibility of concurrent evaluation of expressions in a program
- (sum (4 5))
- 9



 Functional programming allows for the possibility of concurrent evaluation of expressions in a program

```
(sum ( (mult (
(sum (4 5)) (sum (4 5))
)
(sum (4 5))
)
```



- Functional programming is unable to address the problem of history-sensitive behaviors (functions don't keep state)
- Actors avoid assignment operations but allow the modeling of history-sensitive behaviors
 - Gul Agha proposed the concept of replacement, which states that an actor must specify/create the actor which will handle the next message/communication



THE UNIVERSITY of York





- According to Lee (Berkeley, 2004)
 - In actor-oriented design, components are concurrent objects that communicate via messaging, as opposed to abstract data structures that interact via procedure calls
 - Actors are conceptually concurrent, but unlike Agha's actors, they need not have their own thread of control
 - Although communication is still through some form of message passing, it need not be strictly asynchronous



- Actors execute and communicate with other actors in a model
- A have a well defined component interface, which abstracts the internal state and behavior of an actor, and restricts how an actor interacts with its environment
- The interface includes ports that represent points of communication for an actor, and parameters that are used to configure the operation of an actor



THE UNIVERSITY of York





- Actors communicate through channels that pass data from one port to another according to some messaging scheme
- In object-oriented design, components interact primarily by transferring control through method calls, while in actor-oriented design they interact by sending messages through channels
- Actors interact only with the channels that they are connected to and not directly with other actors



Method-call based:





- The syntactic structure of an actor-oriented design model says little about the semantics
- The semantics is determined by a model of computation (MoC), which define the operational rules for executing a model
- Such rules determine when actors perform internal computation, update their internal state, and perform external communication
- The model of computation also defines the nature of communication between components
- According to Lee, the MoC is expressed as a Director in an actor-oriented model



THE UNIVERSITY of York





- MoCs formalize a platform of well defined execution semantics that can help mapping applications onto implementation-related platforms
- Heterogeneous models can be created through the hierarchical composition of different MoCs



Source: E. A. Lee



Concurrent Models of Computation

- Rules that define how actors behave and communicate concurrently in a given model
- According to the rule definitions, models can have specific properties, such as being statically schedulable or time safe
- Examples: Discrete Events, Communicating Sequential Processes, Continuous Time, Process Networks, Synchronous Dataflow, Giotto, etc.



THE UNIVERSITY of fork Application-specific Multiprocessor Platforms based on NoCs | L. S. Indrusiak

Concurrent Models of Computation

Models can have different representations of time

- absolute time
 - T is a totally ordered closed connected set
 - $T = \Re$
- discrete time
 - T is a totally ordered discrete set
- untimed
 - precedences
 - T is a partially ordered discrete set





THE UNIVERSITY of fork Application-specific Multiprocessor Platforms based on NoCs | L. S. Indrusiak

Concurrent Models of Computation

Models can have different representations of time

- absolute time
 - continuous time (CT)

- discrete time
 - discrete events (DE)
- untimed
 - dataflow models
 - Kahn process networks (KPN)
 - synchronous dataflow (SDF)



THE UNIVERSITY of York

- In the DE MoC, actors communicate by sending events, which comprehends a token and a tag
 - token is a data value
 - tag includes a timestamp and a microstep (used to order events with the same timestamp)
- Events are processed chronologically according to their timestamps
- Actors whose available input events are the oldest (having the earliest time stamp of all pending events) are activated



- Model time is global: all actors share the same timestamp and microstep
- When actors send a token through an output port, this token is encapsulated as an event and stored in a global event queue
- Unlike specified otherwise by the actor, the event is tagged with the current model time and microstep (actor processes tokens without delay)



- Actors can be activated at arbitrary times by sending *pure events* to themselves with a future timestamp (no need for ports)
- In the global event queue, events are sorted based on their tags, including time stamps and microsteps
- An event is removed from the global event queue when the model time reaches its time stamp, and if it has a data token, then that token is put into the destination input port



THE UNIVERSITY of York



source: Lee 2004



- Simultaneous events
 - the execution of events with the same timestamp must be ordered (in a deterministic way)
 - achieve a dataflow behavior for events with the same timestamp



source: Lee 2005



Typical uses:

THE UNIVERSITY of York

- digital systems
- queue systems
- networks
- transportation systems
- commerce
- Similar to:
 - VHDL, Verilog, SES Workbench



- Actors represent components that interact via continuous time signals
- They typically specify algebraic or differential relations between inputs and outputs, and can be modeled by ordinary differential equations (ODEs)
- Such models are "solved" rather than simulated



THE UNIVERSITY of York





- Solution is a continuous function of time (waveform)
- The precise solution of a set of ODEs is usually impossible to be found using digital computers
- Numerical solutions are approximations of the precise solution



- Numerical solution approximate the state trajectory of a differential equation by estimating its value at discrete points
- The choice of the points depends on the function
- Using numerical solvers, the solution is a discrete-event signal







Typical uses:

THE UNIVERSITY of York

- analog circuits
- plant dynamics in control systems
- mechanical systems
- heat flows
- Communication channels
- Similar to:
 - Spice, Agilent ADS, Simulink



Dataflow Model of Computation

- Dataflow models represent a given system
 - as a directed graph
 - where each node performs computation (actors)
 - and edges represent channels through which nodes can exchange data (tokens)



- In Von Neumann imperative style, program counter rules the execution
- Dataflow models don't define the sequence on which actors are executed (fired)
- Actors can fire when inputs are available (actors without inputs can fire anytime)
- Exposes potential concurrency: data movement rules the execution



Dataflow Model of Computation

- Actors execute concurrently and communicate by exchanging tokens through FIFO-buffered channels
 - computation performed by a given actor can be described by imperative language



- Buffers usually treated as unbounded for flexibility
 - buffer sizes on final implementation are a design decision
- Read operations are destructive: once read, the token is removed from the buffer



THE UNIVERSITY of fork Application-specific Multiprocessor Platforms based on NoCs | L. S. Indrusiak

Dataflow Model of Computation

- Applications of dataflow models
 - signal processing systems
 - systems dealing with continuous streams of data





- Easier to extract parallelism
- Buffered channels are used in such systems anyway



Kahn Process Networks (KPN)

- Proposed by Gilles Kahn in 1974
- Processes (actors) communicate asynchronously through unbounded buffers
 - non-blocking write
 - blocking read
- A process that reads from an empty channel will stall and can only continue when the channel contains sufficient tokens
- Processes are not allowed to test an input channel for existence of tokens without consuming them



THE UNIVERSITY of fork Application-specific Multiprocessor Platforms based on NoCs | L. S. Indrusiak

Kahn Process Networks (KPN)

- At any given point, a process is either doing some computation (active) or it is blocked waiting for data (read blocked) on exactly one of its input channels
 - it cannot wait for data from more than one channel simultaneously





Kahn Process Networks (KPN)

- Important property: DETERMINISM
 - the sequence of values communicated through the channels is completely determined by the model
 - a KPN can be executed using a completely parallel schedule, a completely sequential schedule, or any schedule in between, always yielding the same output results for a given input sequence
 - which process to activate at a certain moment has to be decided during execution time, based on the current situation



THE UNIVERSITY of fork Application-specific Multiprocessor Platforms based on NoCs | L. S. Indrusiak

Kahn Process Networks (KPN)

Example



```
process f (InFIFO<int> u, OutFIFO<int> v, OutFIFO<int> w)
{
    int k;
    while (true){
        k = u.wait();
        if (k % 2 = 0) v.send (k);
        else w.send(k);
    }
    send() writes a data
    value on an output FIFO
```


THE UNIVERSITY of york Application-specific Multiprocessor Platforms based on NoCs | L. S. Indrusiak

Kahn Process Networks (KPN)



```
process g (InFIFO<int> u, OutFIFO<int> v)
{
    int k;
    while (true){
        k = u.wait();
        v.send (k);
    }
}
```



THE UNIVERSITY of fork Application-specific Multiprocessor Platforms based on NoCs | L. S. Indrusiak

Kahn Process Networks (KPN)



```
process h (InFIFO<int> t, InFIFO<int> u, OutFIFO<int> v)
{
    int k; boolean s=true;
    while (true){
        if (s) then k = t.wait();
        else k = u.wait();
        v.send (k);
        s=!s;}
```



THE UNIVERSITY of York Application-specific Multiprocessor Platforms based on NoCs | L. S. Indrusiak

Kahn Process Networks (KPN)





































































































































Example



Schedule

• f f f f g_u h h g_u g_d g_d g_u h h h


Kahn Process Networks (KPN)

- A given process is only affected by the sequence of tokens on its inputs
 - ▶ it can't tell whether they arrive early, late, or in what order
 - it will behave the same in any case
 - the sequence of tokens it puts on its outputs is the same regardless of the timing of the tokens on the inputs



Kahn Process Networks (KPN)

- Schedule must be determined at runtime
 - it doesn't affect functional behavior
 - challenge is to avoid accumulation of tokens
 - not all systems can be scheduled without token accumulation



- whether a KPN can execute in bounded memory is undecidable
- a number of algorithms can find a schedule to execute a KPN in bounded memory (if such schedule exists)
 - usually require run-time deadlock detection



- Proposed by Lee and Messerschmitt in 1987
- Restricted type of KPN useful for modeling simple dataflow systems
- Constraint: each actor in a model reads and writes a fixed number of tokens every time it is fired
- Schedule can be statically determined
 - simpler and faster implementation of the model execution engine, bounded memory usage
 - deadlocks are avoided



- Adequate for signal processing systems
 - number of consumed and produced tokens per firing is independent of the data and known beforehand





- Adequate for signal processing systems
 - number of consumed and produced tokens per firing is independent of the data and known beforehand



homogeneous SDF graph

Delay D – nth token consumed by its successor is the n-1th token produced by its predecessor

initialized with d "zero" tokens



- Scheduling SDF
 - must respect precedence graph
 - no process fired unless all tokens it consumes are available



- ▶ valid schedules: A B C , B A C
- invalid schedule: CAB



- Periodic Admissable Sequential Schedule (PASS)
 - assumes infinite stream of data
 - periodic schedule can be applied repetitively on input stream without accumulating tokens in the buffers



PASS:



Synchronous Dataflow (SDF)

- Periodic Admissable Sequential Schedule (PASS)
 - formalism to determine PASS
 - 1. build the topology matrix Γ



token consumption is negative



- Periodic Admissable Sequential Schedule (PASS)
 - formalism to determine PASS
 - 2. determine the relative firing frequency of each node by finding the smallest positive integer vector q such as Γ q = 0





- Periodic Admissable Sequential Schedule (PASS)
 - formalism to determine PASS
 - 3. if rates can be established, any scheduling algorithm that avoids buffer underflow will produce a correct schedule if it exists



PASS: A B C C



1 2

- Periodic Admissable Sequential Schedule (PASS)
 - not always possible to find a PASS





- Periodic Admissable Sequential Schedule (PASS)
 - how to test?
 - rank(Γ) = s 1 (s stands for the number of actors on the graph)
 - necessary condition for the existence of a PASS





- Periodic Admissable Sequential Schedule (PASS)
 - how to test?
 - rank(Γ) = s 1 (s stands for the number of actors on the graph)
 - necessary condition for the existence of a PASS





rank(r) = 3



- Periodic Admissable Sequential Schedule (PASS)
 - how to test?
 - ▶ rank(Γ) = s 1 (s stands for the number of actors on the graph)
 - necessary condition for the existence of a PASS





- Periodic Admissable Sequential Schedule (PASS)
 - slightly more complex example





- Periodic Admissable Sequential Schedule (PASS)
 - slightly more complex example





Synchronous Dataflow (SDF)

- Periodic Admissable Sequential Schedule (PASS)
 - slightly more complex example



Possible schedules: BBBCDDDDAA BDBDBCADDA BBDDBDDCAA ... many more Which one should I choose?



- Periodic Admissable Sequential Schedule (PASS)
 - how to choose a schedule?





Synchronous Dataflow (SDF)

- Periodic Admissable Sequential Schedule (PASS)
 - how to choose a schedule?



useful for a reconfigurable hardware implementation

Possible schedules: BBBCDDDDAA BDBDBCADDA BBDDBDDCAA ... many more single appearance schedule



Synchronous Dataflow

- Typical uses:
 - signal processing



Synchronous Dataflow

THE UNIVERSITY of york

Another example: SDF scheduling in two steps

- Establish relative execution rates by solving a system of linear equations
- Determine periodic schedule by simulating system for a single round







- FIR.firingCount = 3
- FIR2.firingCount = 1
- input.tokenConsumptionRate = 6
- output.tokenProductionRate = 4

(c) Balance Equation Solution



source: Neuendorffer 2004

Application-specific Multiprocessor Platforms based on NoCs | L. S. Indrusiak

Synchronous Dataflow

THE UNIVERSITY of York





Synchronous Dataflow

- Exercise:
 - Consider the actor oriented model below following a synchronous dataflow model of computation. Calculate the balance equations and find a valid schedule for it in such a way that no actor requires buffers larger than 3 on its input ports.





Synchronous Dataflow

- Exercise:
 - ▶ 2A E = 0
 - ► A C = 0
 - ▶ 2B 2C = 0
 - ▶ 2B D = 0
 - ▶ 2C D = 0
 - ▶ 3D 3 E = 0



▶ A = B = C = 1

▶ D = E = 2

- Valid schedules:
 - ► A B C D D E E (requires 6 position buffer at D-E channel)
 - ► A B C D E D E (all buffers have 3 positions or less)



Course Outline

THE UNIVERSITY of york

- On-Chip Multiprocessing
 - motivation, review
- Application-specific Multiprocessor Platforms
 - validation, concurrency issues, application-platform mapping
- Platform models
 - evaluation, accuracy
- Application models
 - concurrency
- Joint execution of application and platform models
 - case study
- Background theory and tool support
 - actor orientation, Ptolemy II, custom made extensions
 - demo



Application Modelling

Extending actor-orientation with types and explicit ordering

- UML suitable visual representation for the definition of polymorphic type systems (class diagrams) and ordering relations (sequence diagram)
- but UML is not an executable specification language





UML sequence diagrams within actors

Application-specific Multiprocessor Platforms based on NoCs | L. S. Indrusiak

- Recall the formal definition of MSC

THE UNIVERSITY of York

- tuple ⟨*P*, *E*, *C*, *I*, *m*, < ⟩ where</p>
 - ► *P* is a finite set of processes
 - *E* is a finite set of events
 - *C* is a finite set of names for messages
 - ▶ *I*: $E \to T = \{ p!q(a), p?q(a), p(a) \mid p \neq q \in P, a \in C \}$
 - *m*: $S \rightarrow R$
 - $< \subseteq E \times E$ is a acyclic relation between events consisting of:
 - a total order on E_p for every $p \in P$, and
 - *s* < *r*, *whenever m*(*s*)=*r*



UML sequence diagrams within actors

- Recall the formal definition of MSC
 - order of events (message occurrence) within a process (lifeline) is a total order
 - the reflexive-transitive closure of < (denoted as <*) is a partial order on the complete set E
 - enough for the definition of an untimed model of computation



 different possibilities were explored and integrated as a library of directors on an extended version of Ptolemy II



Application modelling

 Application modeling based on behavioral patterns and polymorphic type systems





Platform modelling

THE UNIVERSITY of York



