Cellflow: a Parallel Application Development Environment with Run-Time Support for the Cell BE Processor

Martino Ruggiero, Michele Lombardi,
Michela Milano, Luca Benini
DEIS - University of Bologna
Italy

INTRODUCTION

- Providing support for multimedia applications on MPSoC platforms remains a significant research challenge
- New tools for efficient mapping of applications onto hardware platforms
- We propose a novel mapping framework:
  - Programming
  - Automatic HW resource management
  - Allocation & Scheduling
    - Computation-efficient w.r.t. state-of-the-art commercial solvers

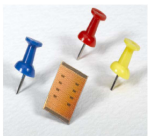

Cell BE STI:

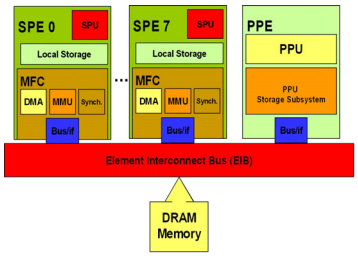

- Heterogeneous system:
  - 1 PPE
  - 8 SPEs
- Element Interconnect Bus:
  - DMA-based
- Limited Local Storage
  - Only 256KB
- **Explicit Resource management**

A multi-core system architecture

It addresses:

- **Server applications:**
  - Next generation IBM Blade Servers.
- **High-performance embedded applications:**
  - Gaming (Sony PS3).
  - Aerospace and defence.
  - Medical imaging.

Cellflow:

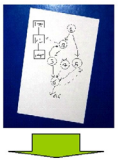

- A software development **toolkit** to help programmers in software implementation
- Starting from a high level task and data flow graph, software developers can easily and quickly build their **application infrastructure**.
- Programmers can intuitively translate high level representation into C-code using our facilities and libraries
- The main goals are:
  - guarantees on **high performance** and **constraint satisfaction**;
  - **predictable** application execution after the optimization step.

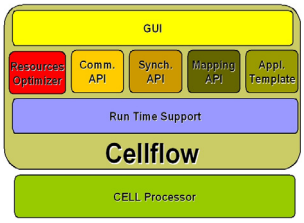

FRAMEWORK

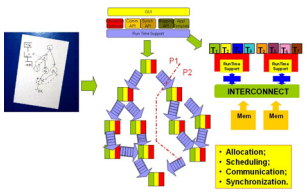

LOGIC-BASED BENDER DECOMPOSITION

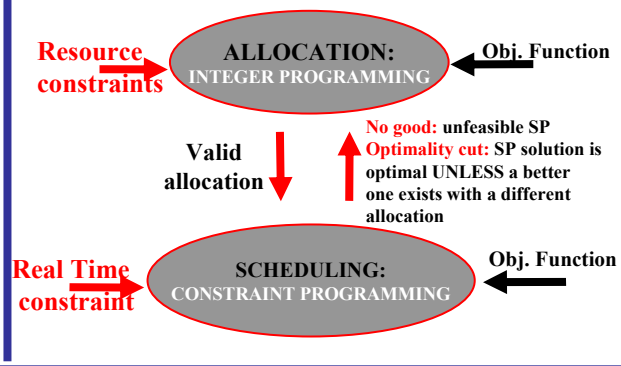

MULTI-STAGE DECOMPOSITION

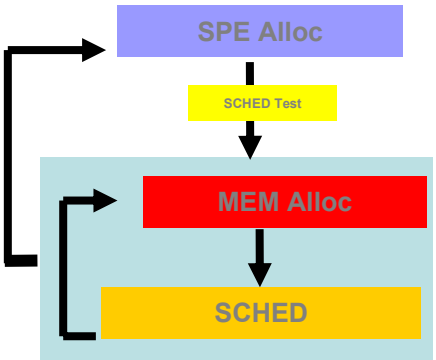

EXPERIMENTAL RESULTS

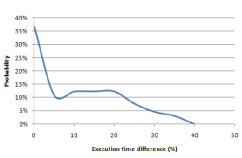

Exact vs. Heuristic Scheduler:

- Heuristic:
  - RR resource allocation + List scheduling
- Up to 40% makespan difference
- 15% in average

TD vs pure-CP

		CP		TD	
Number of tasks	Number of arcs	time (sec.)	> TL	time (sec.)	> TL
15	9-13	0.01	0	0.31	0
15	14-26	0.02	0	0.62	0
25	30-55	0.10	0	369.66	2
25	56-65	0.05	0	530.96	2
30	47-71	1.25	2	620.13	11
30	73-82	0.12	0	834.45	8

task durations are not dependant by allocation decisions

		CP		TD	
Number of tasks	Number of arcs	time (sec.)	> TL	time (sec.)	> TL
10-11	4-11	16.70	0	3.67	0
12-13	8-14	116.92	2	11.19	0
14-15	8-15	81.50	8	10.25	0
16-17	11-17	34.66	11	29.53	0
18-19	13-19	66.47	15	72.56	1
20-21	16-22	400.41	16	248.00	2
22-23	19-26	30.78	18	355.15	3
24-25	20-29	—	20	200.00	9
26-27	23-29	—	20	425.00	6
28-29	25-35	—	20	742.73	9

task durations are dependant by allocation decisions

TD vs BD

	TD			BD		Timed out		
ntasks	SPE it.	MEM it.	time	PM it.	time	TD ∧ BD	¬ TD ∧ BD	TD ∧ ¬ BD
10-11	12	13	3.95	12	71.10	0	0	0
12-13	17	21	11.59	13	151.38	0	1	0
14-15	19	28	14.78	14	145.19	0	0	0
16-17	29	38	42.61	18	388.89	0	2	0
18-19	46	70	245.17	28	863.00	1	5	0
20-21	70	90	665.35	23	1291.90	4	8	0
22-23	33	69	1304.92	19	1686.00	12	6	1
24-25	29	42	1486.15	8	1623.00	11	4	3
26-27	18	41	1523.50	4	1701.67	12	4	3
28-29	13	19	1800.00	3	1721.00	19	0	1

Table 1. Performance tests

Validation of optimizer solutions:

- Comparison between optimizer prediction and real one
- Average difference: 4.8%
- Standard Deviation: 2.41

Demonstrators:

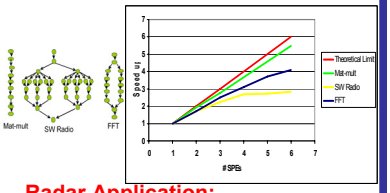

Radar Application:

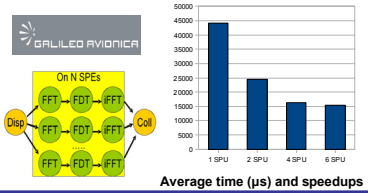

Average time (μs) and speedups