

Skeletons from grids to multicores

Introduction Scenario Urgencies

Multicores Scenario FastFlow Experimental results Offloading

Non functiona concerns

Behavioural skeletons Hierarchical composition

Multiple concern management

Conclusions

Skeletons from grids to multicores

Marco Danelutto

Dept. of Computer Science, University of Pisa, Italy

October 2010, Amsterdam





from grids to multicores				
M. Danelutto				
Introduction				
Scenario				
Urgencies				
Multicores				
Scenario				
FastFlow				
Experimental				
results				
Offloading				
Non functional				
concerns				
Scenario				
Behavioural				
skeletons				
Hierarchical composition				
Multiple concern				
management				



Contents

Skeletons from grids to multicores

M. Danelutto

Introduction

Scenario Urgencies

Multicores

Scenario FastFlow Experiment

Offloading

Non functiona concerns

Scenario Behavioural skeletons Hierarchical composition

Multiple concern management

Conclusions

1 Introduction

- Scenario
- Urgencies

2 Multicores

- Scenario
- FastFlow
- Experimental results
- Offloading

3 Non functional concerns

- Scenario
- Behavioural skeletons
- Hierarchical composition
- Multiple concern management



Structured parallel programming

Skeletons from grids to multicores
M. Danelutto
Introduction Scenario Urgencies



Structured parallel programming

Skeletons from grids to multicores

M. Danelutto

Introductior Scenario

Multicores Scenario FastFlow Experimental results Offloading

Non functiona concerns

Behavioural skeletons Hierarchical

composition Multiple concer

Conclusions

Algorithmic skeletons

- \blacksquare Cole 1988 \rightarrow common, parametric, reusable parallelism exploitation pattern
- languages & libraries since '90 (P3L, Skil, eSkel, ASSIST, Muesli, SkeTo, Mallba, Muskel, Skipper, BS, ...)
- high level parallel abstractions (parallel programming community)
 - hiding most of the technicalities related to parallelism exploitation
 - directly exposed to application programmes



Structured parallel programming

Skeletons from grids to multicores

M. Danelutto

Introductior Scenario Urgencies

Multicores Scenario FastFlow Experimental results Offloading

Non functiona concerns

Scenario Behavioural skeletons Hierarchical composition Multiple conce

Conclusions

Algorithmic skeletons

- \blacksquare Cole 1988 \rightarrow common, parametric, reusable parallelism exploitation pattern
- languages & libraries since '90 (P3L, Skil, eSkel, ASSIST, Muesli, SkeTo, Mallba, Muskel, Skipper, BS, ...)
- high level parallel abstractions (parallel programming community)
 - hiding most of the technicalities related to parallelism exploitation
 - directly exposed to application programmes

Parallel design patterns

- Massingill, Mattson, Sanders 2000 → "Patterns for parallel programming" book (2006) (software engineering community)
- design patterns à la Gamma book
 - name, problem, solution, use cases, etc.
- concurrency, algorithms, implementation, mechanisms



Concept evolution

Parallelism

Skeletons from grids to multicores

M. Danelutto

Introduction Scenario

Multicores Scenario FastFlow Experimental results Offloading

Von functiona concerns

Behavioural skeletons Hierarchical

Multiple concern management

- parallelism exploitation patterns shared among applicationsseparation of concerns:
 - \blacksquare system programmers \rightarrow efficient implementation of parallel patterns
 - \blacksquare application programmers \rightarrow application specific details



Concept evolution

Skeletons from grids to multicores

M. Danelutto

Introduction Scenario Urgencies

Multicores Scenario FastFlow Experimental results Offloading

Non functiona concerns

Scenario Behavioural skeletons Hierarchical composition Multiple conc

Conclusions

Parallelism

- parallelism exploitation patterns shared among applicationsseparation of concerns:
 - \blacksquare system programmers \rightarrow efficient implementation of parallel patterns
 - \blacksquare application programmers \rightarrow application specific details

New architectures

- Heterogeneous in Hw & Sw
- Multicore NUMA, cache coherent architectures



Concept evolution

Skeletons from grids to multicores

M. Danelutto

ntroduction Scenario Urgencies

Multicores Scenario FastFlow Experimental results Offloading

Non functiona concerns Scenario

Behavioural skeletons Hierarchical composition Multiple conc

Conclusions

Parallelism

- parallelism exploitation patterns shared among applicationsseparation of concerns:
 - \blacksquare system programmers \rightarrow efficient implementation of parallel patterns
 - \blacksquare application programmers \rightarrow application specific details

New architectures

- Heterogeneous in Hw & Sw
- Multicore NUMA, cache coherent architectures

Further non functional concerns

security, fault tolerance, power management, ...



Urgencies

- Skeletons from grids to multicores
- M. Danelutto
- Introductior Scenario Urgencies
- Multicores Scenario FastFlow Experimental results Offloading
- Non functiona concerns
- Behavioural skeletons Hierarchical composition Multiple conce
- Conclusions

- \rightarrow Targeting multi/many cores
 - different implementation issued and solutions
 - completely different computational grains to be addressed
- \rightarrow Targeting non functional concerns
 - autonomic management of independent non functional concerns
 - co-management of different non functional concerns



Urgencies

- Skeletons from grids to multicores
- M. Danelutto
- Introduction Scenario Urgencies
- Multicores Scenario FastFlow Experimental results Offloading
- Non functiona concerns
- Scenario Behavioural skeletons Hierarchical composition Multiple conce

Conclusions

- \rightarrow Targeting multi/many cores
 - different implementation issued and solutions
 - completely different computational grains to be addressed
- \rightarrow Targeting non functional concerns
 - autonomic management of independent non functional concerns
 - co-management of different non functional concerns

Structured parallel programming models

- can be exploited to address both issues
- synergies among the solutions may be exploited as well



Skeletons from grids to multicores

M. Danelutto

Introduction Scenario Urgencies

Multicores

Scenario FastFlow Experimental results Offloading

Non functiona concerns

Behavioural skeletons Hierarchical composition

Multiple concern management

Conclusions

Targeting multicores



Targeting multi/many cores

Skeletons from grids to multicores

Introduction Scenario Urgencies

Multicores Scenario

FastFlow Experimental results Offloading

Non function concerns

Behavioural skeletons Hierarchical composition Multiple conc

Conclusions

Structured parallel programming on $\ensuremath{\mathsf{COW}}/\ensuremath{\mathsf{NOW}}$

Implementation template based

• P3L, eSkel, Muesli, SkeTo \rightarrow collection of

 $\langle Architecture, ProcessNetwork, Model \rangle$

Macro Data Flow based

■ Lithium/Muskel, Skipper, Calcium (Skandium) → compile skeletons to MacroDataFlow graphs + Dstributed MDF interpreter



Targeting multi/many cores

Skeletons from grids to multicores

I<mark>ntroduction</mark> Scenario Urgencies

Scenario

Experimental results Offloading

Non function concerns

Scenario Behavioural skeletons Hierarchical composition Multiple conc

Conclusions

Structured parallel programming on $\ensuremath{\mathsf{COW}}/\ensuremath{\mathsf{NOW}}$

Implementation template based

• P3L, eSkel, Muesli, SkeTo \rightarrow collection of

 $\langle Architecture, ProcessNetwork, Model \rangle$

Macro Data Flow based

■ Lithium/Muskel, Skipper, Calcium (Skandium) → compile skeletons to MacroDataFlow graphs + Dstributed MDF interpreter

Emphasis

- communication latency hiding
- avoid unnecessary data copies



Multi/many core features

Skeletons from grids to multicores

M. Danelutto

Introduction Scenario Urgencies

Multicores

Scenario FastFlow Experimenta results Offloading

Non functiona concerns

Behavioural skeletons Hierarchical composition

Multiple concern management

Conclusions

Shared memory hierarchy

- NUMA (C vs. M accesses, non uniform intercore interconnection strctures)
- chache coherence (snoopy vs. directory based)
- Control abstractions
 - threads (user vs. system space, completion vs. time sharing)
 - processes



Multi/many core features

Skeletons from grids to multicores

M. Danelutto

Introduction Scenario Urgencies

Multicores

Scenario FastFlow Experimenta results Offloading

Non functiona concerns

Behavioural skeletons Hierarchical composition Multiple conc

Conclusions

Shared memory hierarchy

- NUMA (C vs. M accesses, non uniform intercore interconnection strctures)
- chache coherence (snoopy vs. directory based)

Control abstractions

- threads (user vs. system space, completion vs. time sharing)
- processes

Focus

- synchronization overheads
- data access patterns
- thread-core binding, affinity scheduling



FastFlow

Skeletons from grids to multicores

- M. Danelutto
- Introduction Scenario Urgencies
- Multicores Scenario FastElow
- Experimental results Offloading
- Non function concerns
- Scenario Behavioural skeletons Hierarchical composition Multiple conc
- Conclusions

Advanced programming framework

- targeting multicores
- minimizing synchronization latencies
- streaming support through skeletons
- expandable
- open source





FastFlow: simple streaming networks

Skeletons from grids to multicores

- Introduction Scenario Urgencies
- Multicores Scenario FastFlow
- Experimental results Offloading
- Non function concerns
- Scenario Behavioural skeletons Hierarchical composition Multiple cond
- management
- Conclusions

Single Producer Single Consumer (SPSC) queue

- uses results from the '80
- lock-free, wait-free
- no memory barriers for Total Store Order processor (e.g. Intel, AMD)
- single memory barrier for weaker memory consistency models (e.g. PowerPC)
- ightarrow very low latency in communications





FastFlow: simple streaming networks

Skeletons from grids to multicores

- M. Danelutto
- Introduction Scenario Urgencies
- Multicores Scenario FastElow
- Experimental results Offloading
- Non functiona concerns Scenario
- Behavioural skeletons Hierarchical composition Multiple concer management
- Conclusions

Other queues: SPMC MPSC MPCP

- one-to-many, many-to-one and many-to-many sychronization and data flow
- use a explicit arbiter thread
- providing lock-free and wait-free arbitrary data-flow graphs
- cyclic graphs (provably deadlock-free)







FastFlow: high level programming abstractions



M. Danelutto

Introduction Scenario Urgencies

Multicores Scenario FastElow

Experimenta results Offloading

Non functiona concerns Scenario Behavioural skeletons Hierarchical

composition Multiple concern management











Shows the advantage of lock free comms



Skeletons from grids to multicores

Matrix multiplication

parallel for i only

	matmu	l_ff_v1	matm	ul_ff_v2	matmul.	OpenMP
n. workers	Time S	Speedup	Time	Speedup	Time	Speedup
4	2982.24	3.6	2691.74	4.0	2767.48	3.9
8	1495.87	7.2	1435.41	7.5	1455.73	7.4
12	1496.01	7.2	1431.34	7.5	1492.87	7.2
16	1495.74	7.2	1426.42	7.6	1433.31	7.5

Conclusions

Experimental results



Skeletons from grids to multicores

Experimental results

Matrix multiplication

parallel for i only

	matm	ul_ff_v1	matm	ul_ff_v2	matmul	OpenMP
n. workers	Time	Speedup	Time	Speedup	Time	Speedup
4	2982.24	3.6	2691.74	4.0	2767.48	3.9
8	1495.87	7.2	1435.41	7.5	1455.73	7.4
12	1496.01	7.2	1431.34	7.5	1492.87	7.2
16	1495.74	7.2	1426.42	7.6	1433.31	7.5

FastFlow vs. OpenMP

Comparable at quite coarse grain



FastFlow: different appls

Skeletons from grids to multicores

M. Danelutto

Multicores Scenario FastFlow Experimental results Offloading

Non functiona concerns Scenario Behavioural skeletons Hierarchical composition Multiple concerr

Conclusions

Microbenchmarks

Benchmark	Parameters	Skeleton used	Speedup / #cores
Matrix Mult.	1024×1024	farm no collector	7.6 / 8
Quicksort	50M integers	D&C	6.8 / 8
Fibonacci	Fib(50)	D&C	9.21 / 8



FastFlow: different appls

Skeletons from grids to multicores

Experimental results

${\it Microbenchmarks}$

Benchmark	Parameters	Skeleton used	Speedup / #cores
Matrix Mult.	1024×1024	farm no collector	7.6 / 8
Quicksort	50M integers	し&C	6.8 / 8
Fibonacci	Fib(50)	し&C	9.21 / 8

Real use cases

Application	Skeleton used	Measure	Value
YaDT-FF (C4.5 data mining) StochKit-FF (Gillespie)	D&C farm	Speedup Scalabitily	4.5-7.5 10-11
SWPS3-FF (Gene matching)	farm no collector	GCUPS	12.5-34.5



FastFlow: offloading

Skeletons from grids to multicores

Offloading

General purpose methodology

use FastFlow as an efficient, fine grain sw accelerator

1 //<includes> 2 #define N 1024 3 int A[N][N].B[N][N].C[N][N]; 1 5 int main() { // < init A.B.C> for(int i=0:i < N:++i)2 for(int j=0; j<N; ++j) { 10 12 int _C=0: for(int k=0:k<N:++k)13 _C += A[i][k]∗B[k][j]; ❸ 15 C[i][j]=_C; 16 Ξ 17 ė 18 ٦ 6 19 20 a) Original code (sequential)





FastFlow pros and cons

Skeletons from grids to multicores

M. Danelutto

Introductior Scenario Urgencies

Multicores Scenario FastFlow Experimental results Offloading

Non function: concerns Scenario Behavioural skeletons Hierarchical composition

Multiple concern management

Conclusions

Pros:

- very low overhead introduced
- high level abstractions available to application programmers
- streaming fully supported

Cons:

- parallelization of code requires "more code" w.r.t. classical approaches
- "structured" approach to parallelism required



Skeletons from grids to multicores

Introduction Scenario Urgencies

Multicores Scenario FastFlow Experimental results Offloading

Non functional concerns

Scenario Behavioural skeletons Hierarchical composition Multiple concer management

Conclusions

Targeting non functional concerns



The scenario

Skeletons from grids to multicores

M. Danelutto

Introductior Scenario Urgencies

Multicores Scenario FastFlow Experimental results Offloading

Non functiona concerns

Scenario

Behavioural skeletons Hierarchical composition Multiple concerr management

Conclusions

Several important non functional features to be considered:

- \blacksquare performance \rightarrow throughput, latency, load balancing
- security \rightarrow data, code
- \blacksquare fault tolerance \rightarrow checkpointing, recovery strategies
- power management → power/speed tradeoff
- Non functional
 - does not contribute to function computed
 - policies & strategies most likely in the background of system programmers

Autonomic management

 separation of concerns: system programmers (NF) vs. appl programmers (FUN)



Behavioural skeletons

Skeletons from grids to multicores

M. Danelutto

Introduction Scenario Urgencies

Multicores Scenario FastFlow Experimental results Offloading

Non functiona concerns

Behavioural skeletons

Hierarchical composition Multiple concern management

Conclusions

Def: Behavioural skeleton

Co-design of a component including:

- parallelism exploitation pattern
- autonomic management of some non functional concern

Co design

- improves efficiency
- exploits knowledge relative to structure of the computation
- in the design and implementation of suitable management policies



BS: user view





Sample Behavioural skeleton

Skeletons from grids to multicores

M. Danelutto

Introduction Scenario Urgencies

Multicores Scenario FastFlow Experimental results Offloading

Non functiona concerns

Scenario Rebaviour:

Behavioural skeletons

Hierarchical composition Multiple concerr management

Conclusions

Functional replication BS

Parallel pattern

- Master-worker with variable number of workers.
- Master schedules tasks to available workers.

Performance manager

- Interarrival time faster than service time → increase parallelism degree, unless communication bandwidth is saturated.
- \blacksquare Interarrival time slower that service time \rightarrow decrease the parallelism degree.
- \blacksquare Recent change \rightarrow do not apply any action for a while.



Implementation: behavioural skeleton





Implementation: manager

Skeletons from grids to multicores

M. Danelutto

ntroduction Scenario Urgencies

Multicores Scenario FastFlow Experimental results Offloading

Non function concerns

Scenario Behavioural

skeletons Hierarchical composition Multiple concern management

Conclusions

Monitor

perceive pattern status

Analyse

figure out policies

Plan

devise strategy

Execute

implement decisions on pattern





Implementation: manager

Skeletons from grids to multicores

M. Danelutto

Introduction Scenario Urgencies

Multicores Scenario FastFlow Experimental results Offloading

Non functiona concerns

Behavioural skeletons

Hierarchical composition Multiple concerr management

Conclusions

Monitor

perceive pattern status

Analyse

figure out policies

Plan

devise strategy

Execute

implement decisions on pattern



Cyclic execution of a JBoss business rule engine

- \blacksquare RULE :: Priority, Trigger, Condition \rightarrow Action
- Rule set :: derived from user supplied *contract*





Hierarchical composition

Skeletons from grids to multicores

M. Danelutto

Introduction Scenario Urgencies

Multicores Scenario FastFlow Experimental results Offloading

Non functiona concerns Scenario

Behavioural skeletons

Hierarchical composition

Multiple concern management

Conclusions

Program → Composition of BS pipe(seq, farm(seq), seq)

Hierarchy of managers

- user contract directed
 - to top level manager
- propagated (possibly modified) through manager tree

Management "exceptions" manager not able to ensure contract

- \rightarrow raises violation to upper manager in hierarchy
- \rightarrow enters *passive* mode waiting for a new contract



Hierarchical composition

Skeletons from grids to multicores

M. Danelutto

Introduction Scenario Urgencies

Uulticores Scenario FastFlow Experimental results Offloading

Non functiona concerns Scenario

Behavioural skeletons

Hierarchical composition

Multiple concern management

Conclusions

Program → Composition of BS pipe(seq, farm(seq), seq)

Hierarchy of managers

- user contract directed to top level manager
- propagated (possibly modified) through manager tree

Management "exceptions" manager not able to ensure contract

- \rightarrow raises violation to upper manager in hierarchy
- \rightarrow enters *passive* mode waiting for a new contract





Hierarchical composition

Skeletons from grids to multicores

M. Danelutto

Introduction Scenario Urgencies

Multicores Scenario FastFlow Experimental results Offloading

Non functiona concerns Scenario

Behavioural skeletons

Hierarchical composition

Multiple concern management

Conclusions

Program → Composition of BS pipe(seq, farm(seq), seq)

Hierarchy of managers

- user contract directed to top level manager
- propagated (possibly modified) through manager tree

Management "exceptions" manager not able to ensure contract

- \rightarrow raises violation to upper manager in hierarchy
- → enters *passive* mode waiting for a new contract







BS: sample run

Skeletons from grids to multicores

- Introduction Scenario Urgencies
- Multicores Scenario FastFlow Experimental results Offloading
- Non functiona concerns
- Scenario Behavioural skeletons
- Hierarchical composition
- Multiple concern management

Conclusions

Medical image processing

- pipe(seq(getImage), farm(seq(renderImage)), seq(displayImage))
- \blacksquare contract \rightarrow 0.3 to 0.7 images per second
- initial condition
 - enough processing resources
 - image provider stage too slow



BS: sample run

Skeletons from grids to multicores

- M. Danelutto
- Introduction Scenario Urgencies
- Multicores Scenario FastFlow Experimental results Offloading
- Non functiona concerns
- Scenario
- Behavioural skeletons
- Hierarchical composition
- Multiple concern management
- Conclusions





Multiple concern management

Skeletons from grids to multicores

M. Danelutto

Introduction Scenario Urgencies

Vulticores Scenario FastFlow Experimental results Offloading

Non functional

Scenario Behavioural skolotops

Hierarchical

Multiple concern management

Conclusions

Indipendent manager hierarchies

- take care of indipendent concerns (e.g. Performance and Security)
- must coordinate independently taken decisions
 - to avoid instability





BS: Coordination protocol

Skeletons from grids to multicores

Introduction Scenario Urgencies

Multicores Scenario FastFlow Experimental results Offloading

Non functiona concerns

Scenario Behavioural skeletons

Hierarchical composition

Multiple concern management

Conclusions

Fireable rule (Manager X) $R_z :: Trig_i, Cond_j, Prio_k \rightarrow a_1, \ldots, a_n$

Step 0 broadcast changes in the computation graph eventually induced by R_z

Step 1 gather answers from all other manager (hierarchies)

Step 2 analyze answers:

- all $OK \rightarrow perform a_1, \ldots, a_n$
- at least one NOK \rightarrow ABORT R_z & lower $Prio_k$
- require(Property_m) \rightarrow perform $a'_1, \ldots, a'_{n'}$ such that Property_m is ensured



BS: coordination protocol

Skeletons from grids to multicores

M. Danelutto

Introduction Scenario Urgencies

Vulticores Scenario FastFlow Experimental results Offloading

Non functiona concerns

Scenario

skeletons

Hierarchical composition

Multiple concern management

Conclusions

 $\mathsf{program} \to \mathsf{farm}(\mathsf{seq})$

contract :: parDegree=8





BS pros and cons

Skeletons from grids to multicores

M. Danelutto

ntroduction Scenario Urgencies

Vulticores Scenario FastFlow Experimental results Offloading

Non functiona concerns

Scenario Rehavioura

skeletons

Hierarchical

Multiple concern management

Conclusions

Pros:

- full decoupling of system and application programmer concerns
- reconfigurable autonomic management through JBoss rule engine
- two prototypes available: GCM BS (ProActive/GCM based) and LIBERO (pure Java/RMI, configurable action/sensor interface)

Cons:

- further investigation needed on multiple concern management
- compiler *contract* \rightarrow *JBoss rules* needed



Skeletons

from grids to multicores

Conclusions

Skeletons from grids to multicores

Grids

- Introduction Scenario Urgencies
- Multicores Scenario FastFlow Experimental results Offloading
- Non functional concerns
- Scenario Behavioural skeletons Hierarchical composition
- Multiple concern management

Conclusions

 BS suitable to manage typical features of grids: heterogeneity, distribution, security, ...

Multicores

efficient mechanisms supporting typical grains

Next step

- adopt BS like management to improve (at run time) the performance of FastFlow
- e.g. depending on current system behaviour
 - experiment, evaluate, adopt new skeleton composition e.g. farm(pipe(seq,seq)) → pipe(farm(seq), farm(seq))



References

Skeletons from grids to multicores

M. Danelutto

Introduction Scenario Urgencies

Vulticores Scenario FastFlow Experimental results Offloading

Non function: concerns

Scenario Behavioural skeletons Hierarchical composition Multiple conc

Conclusions

These slides available at http://www.di.unipi.it/~marcod follow "Papers" then "Talks" tabs

FastFlow home page at http://calvados.di.unipi.it/dokuwiki/doku.php? id=ffnamespace:about source code (GPL) at

http://sourceforge.net/projects/mc-fastflow/

- ProActive/GCM BS home page at http://gridcomp.ercim.eu/content/view/26/34/
- LIBERO prototype: ask author(s) at marcod@di.unipi.it



Skeletons from grids to multicores

Introduction Scenario Urgencies

Multicores Scenario FastFlow Experimental results Offloading

Non functiona concerns

Behavioural skeletons Hierarchical composition

Multiple concern management

Conclusions

Any questions?

marcod@di.unipi.it
marco.danelutto@danelutto.org