

OC4MC

# OCaml for MultiCore

Deuxième réunion du GDR GPL/LTP

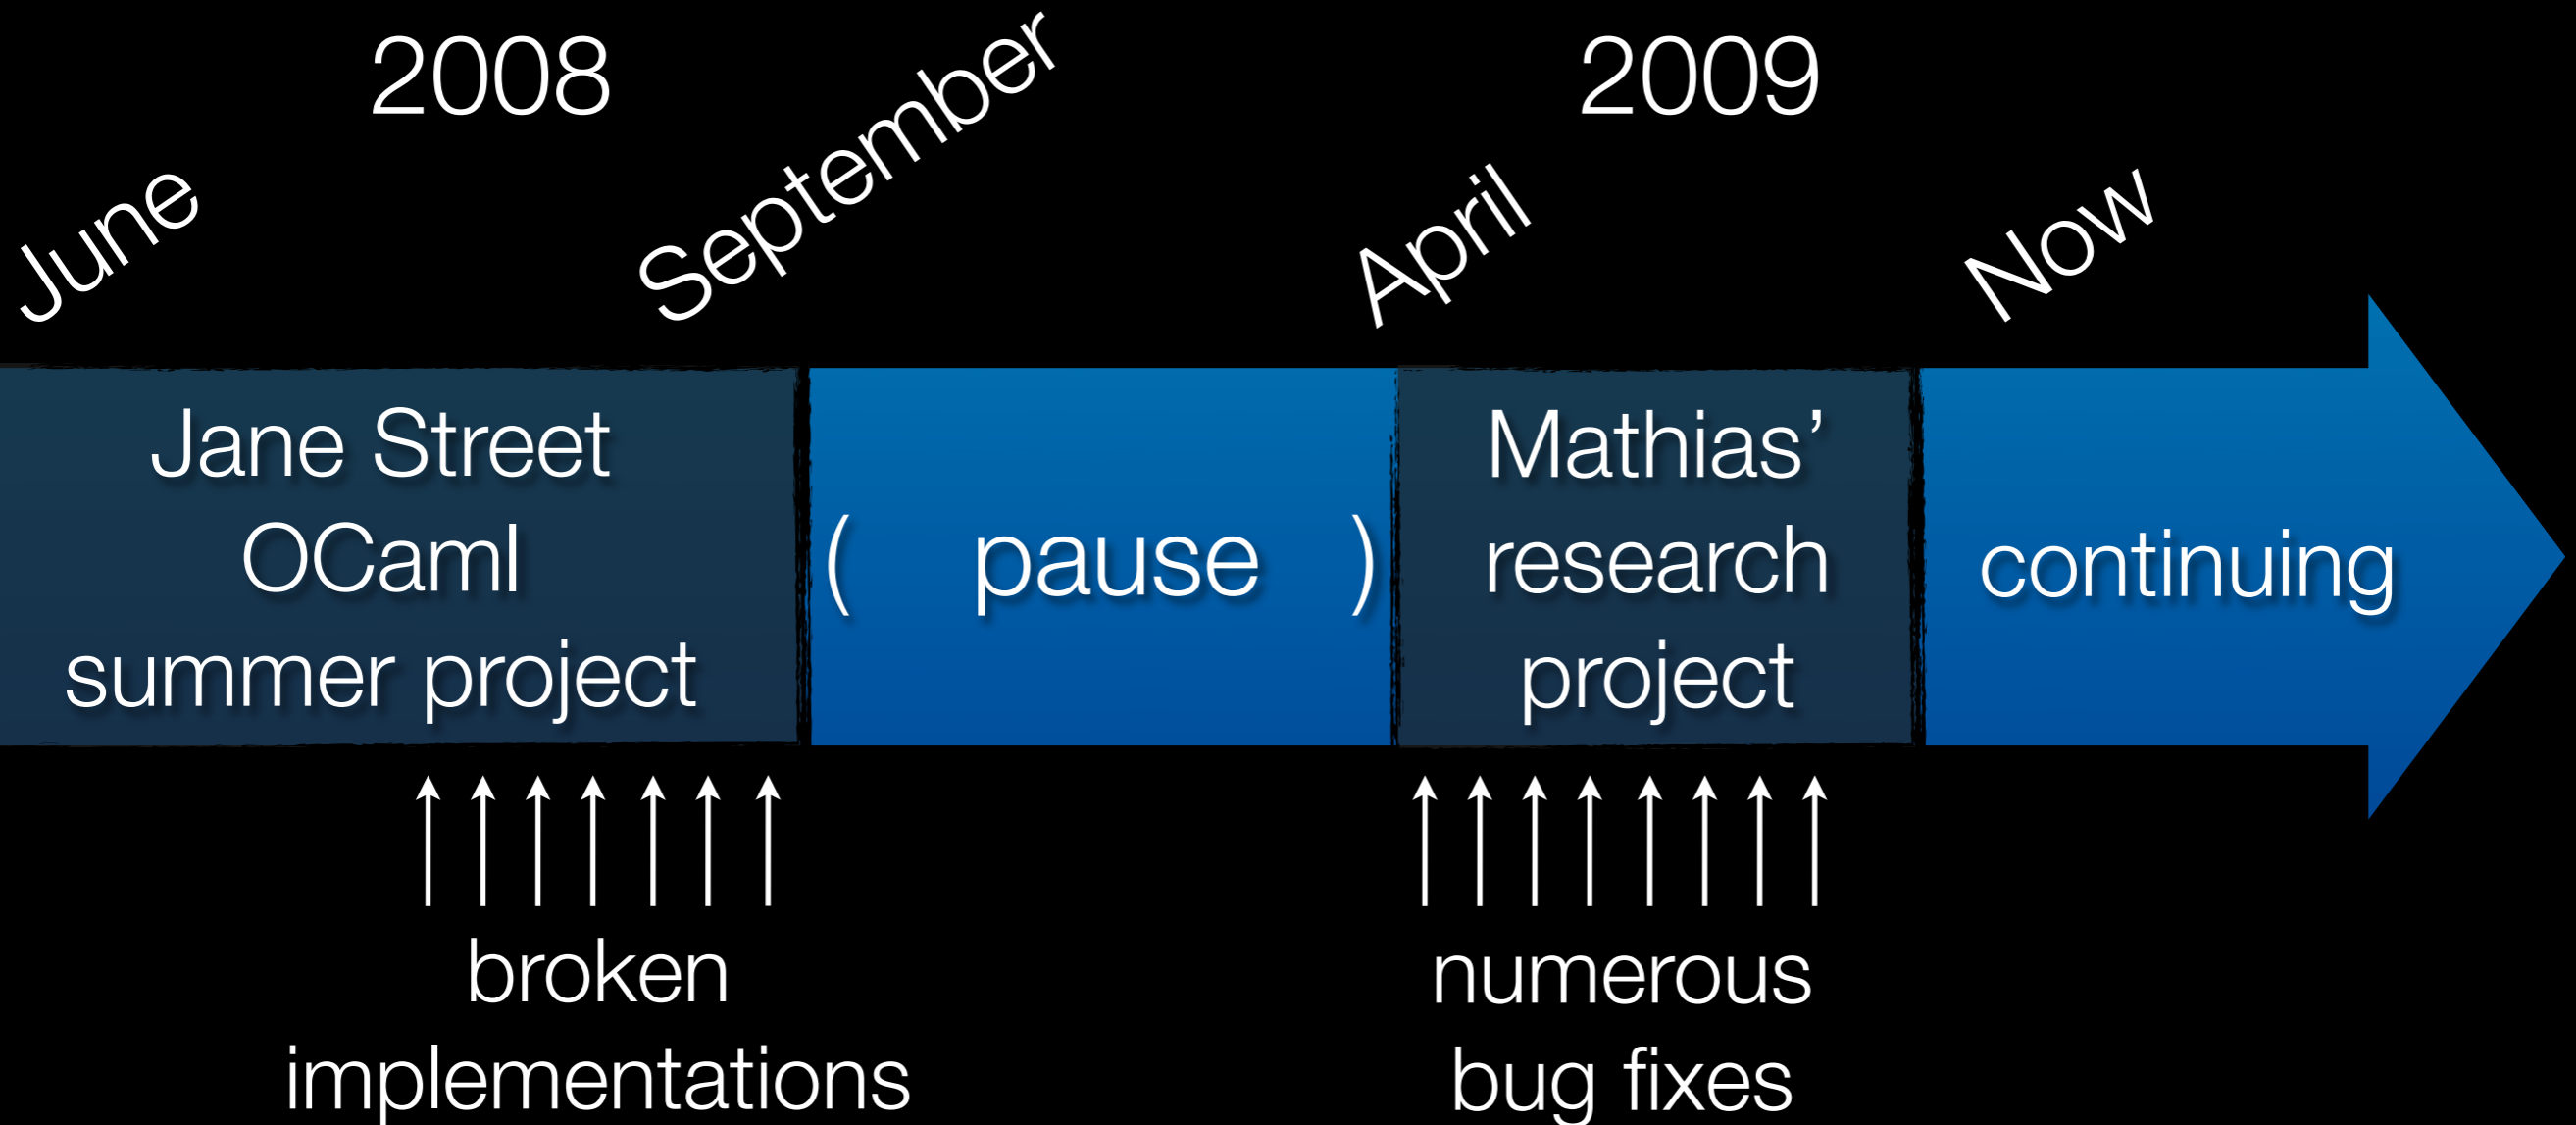
LACL / Université Paris XII

21 Octobre 2009

Mathias Bourgoïn, Adrien Jonquet  
Benjamin Canou, Philippe Wang  
Emmanuel Chailloux



# project timeline



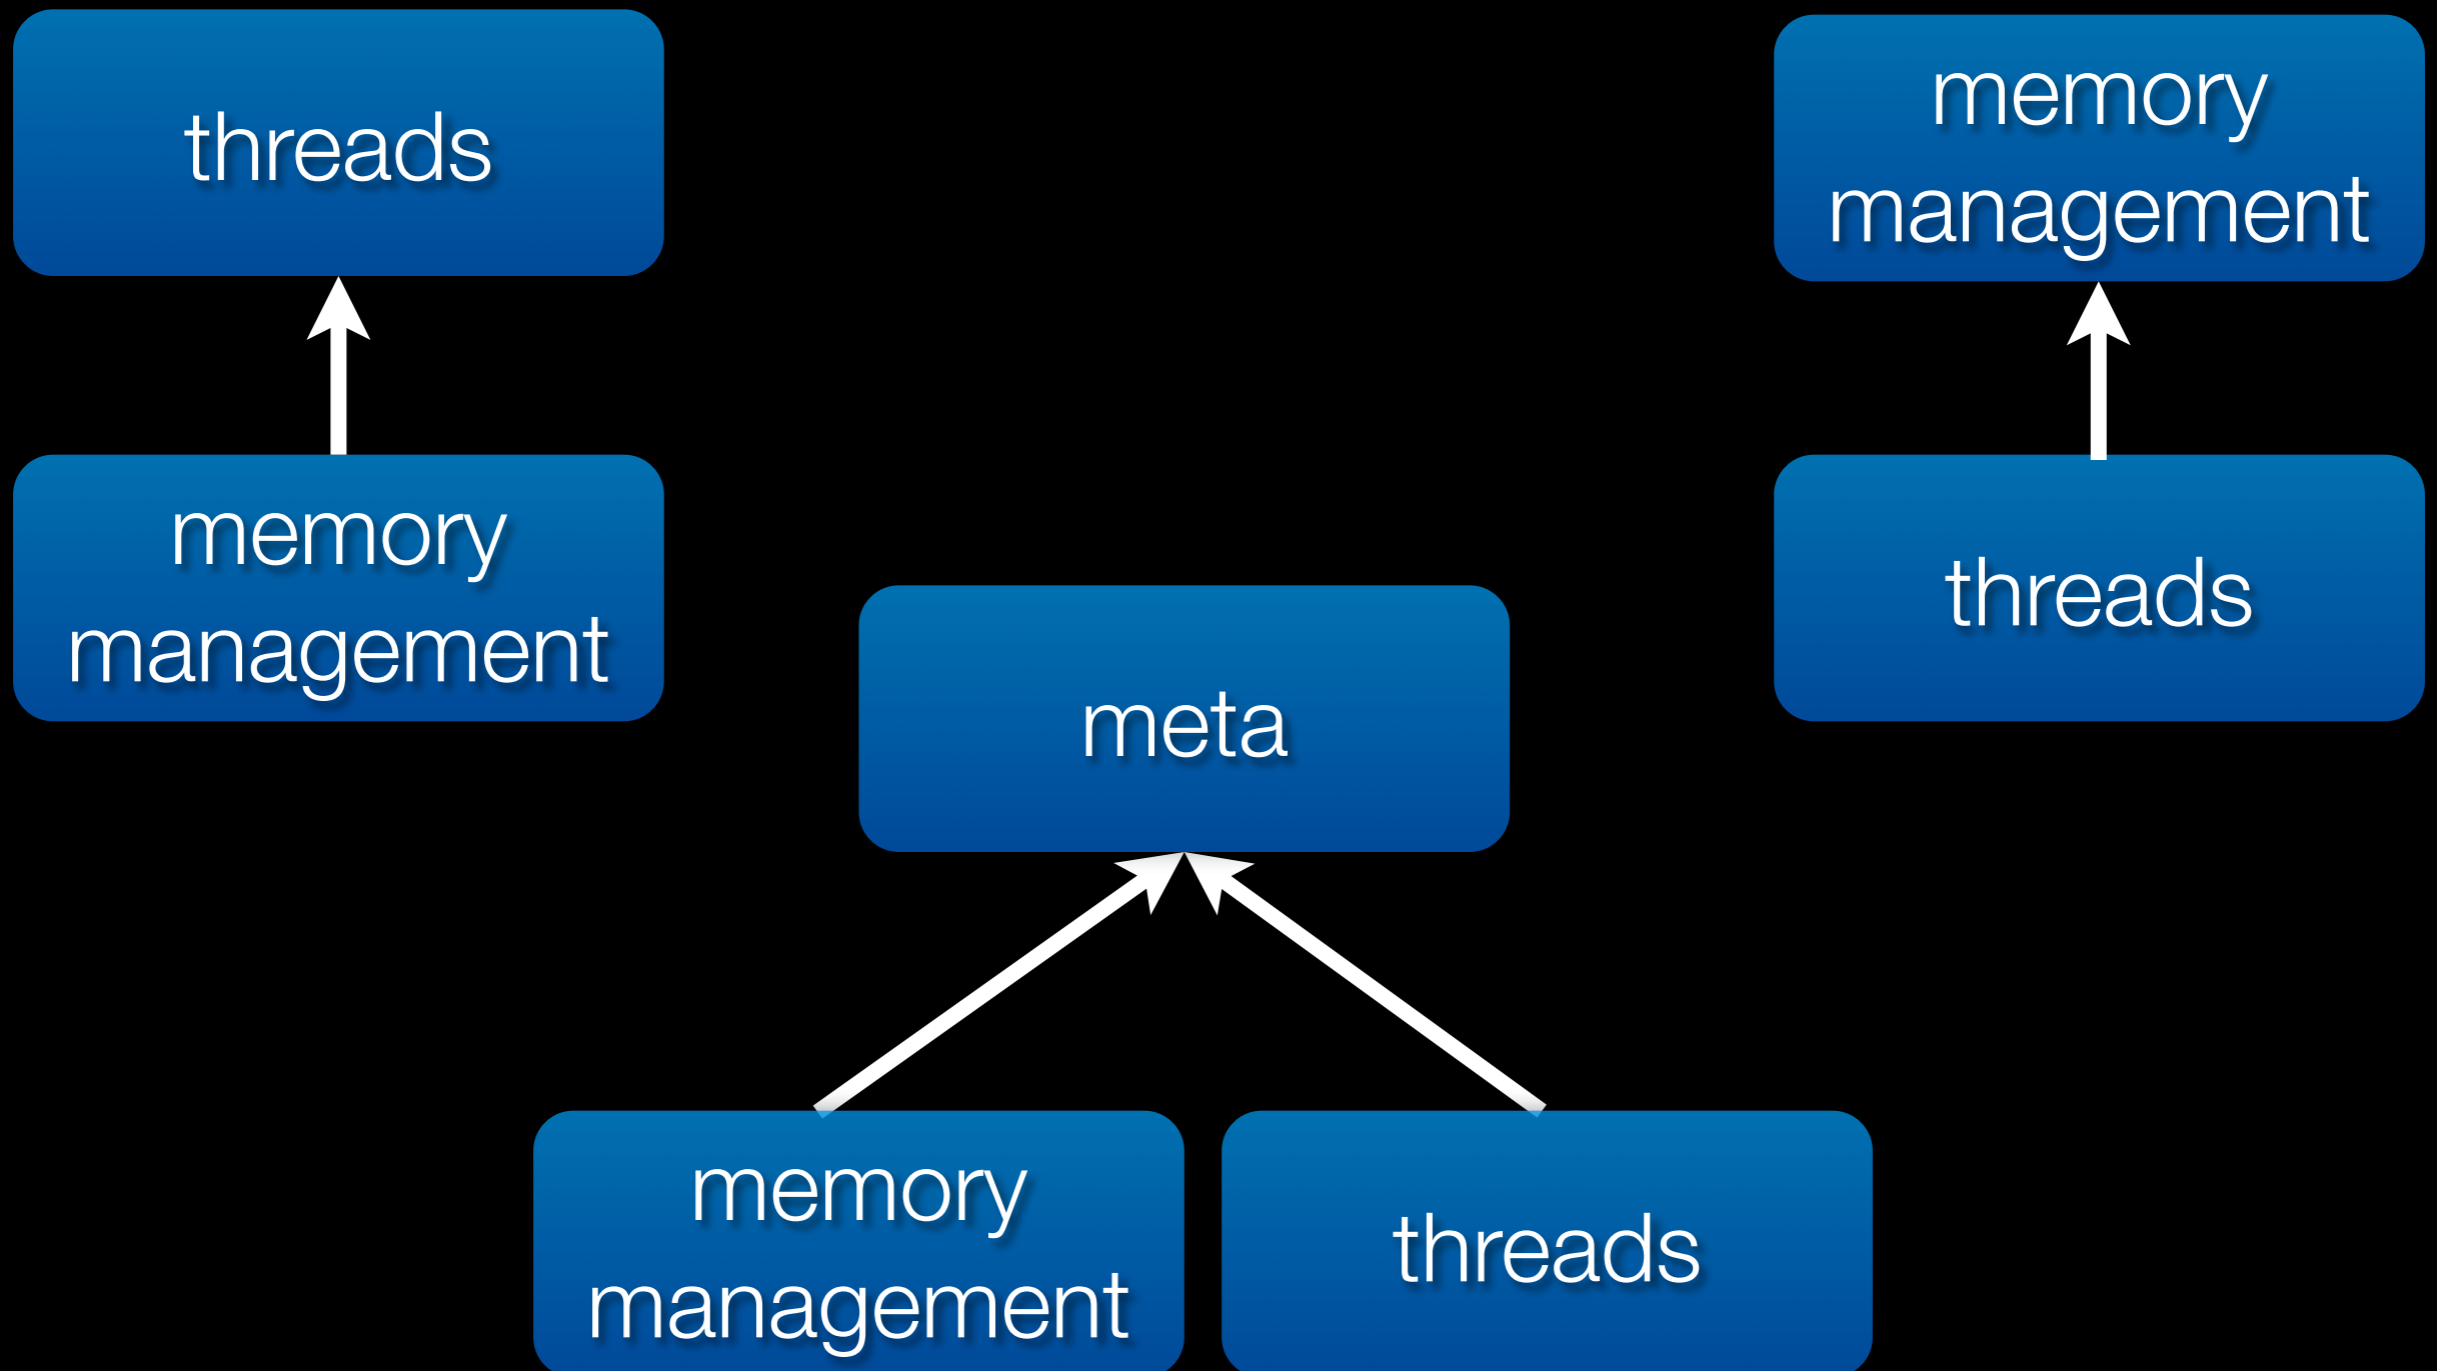
# multithreading based multicore profit with OCaml

- ✦ (super light) preliminaries
- ✦ chapter 1 / OCaml & Parallelism
- ✦ chapter 2 / OCaml for MultiCore
- ✦ chapter 3 / Performance
- ✦ chapter 4 / Conclusion & Future Work

preliminaries

**dependencies mess**

# dependencies



# interdependencies



if there is a garbage collector and there are threads, then their implementations are probably interdependent

chapter 1

# OCaml & Parallelism

(what a mess around)

# multitask programming

join calculus  
concurrent dataflow  
cooperative threads  
preemptive threads  
cluster  
grid  
CUDA  
Ada  
(heavy) process  
p2p  
Io  
OpenCL  
ASM  
C  
SML  
Oz  
Java  
OCaml  
Concurrent ML  
internet computing  
JOCaml  
Alice  
Haskell  
Join Java  
MultiLisp  
Scala  
Concurrent Haskell  
Clojure  
PPE+SPEs  
shared memory  
message passing  
assembly tweaking  
CELL  
SIMD GPU  
CCS  
NVIDIA GeForce GT130  
MIMD CPU  
Intel Core Quad  
 $\pi$ -calculus  
pict



(still the same mess)

# multitask programming

Concurrent Haskell      JOCaml      Concurrent ML

Alice      Join Java       $\pi$ -calculus      lo  
MultiLisp      Ada      pict      CCS      Oz

OC4MC

OCaml Haskell SML Scala      join calculus  
Clojure      concurrent dataflow

Java

OpenCL  
CUDA

cooperative threads      internet computing

preemptive threads

C      (heavy) process

ASM assembly tweaking

message passing  
shared memory

p2p  
cluster  
grid

SIMD GPU

PPE+SPEs

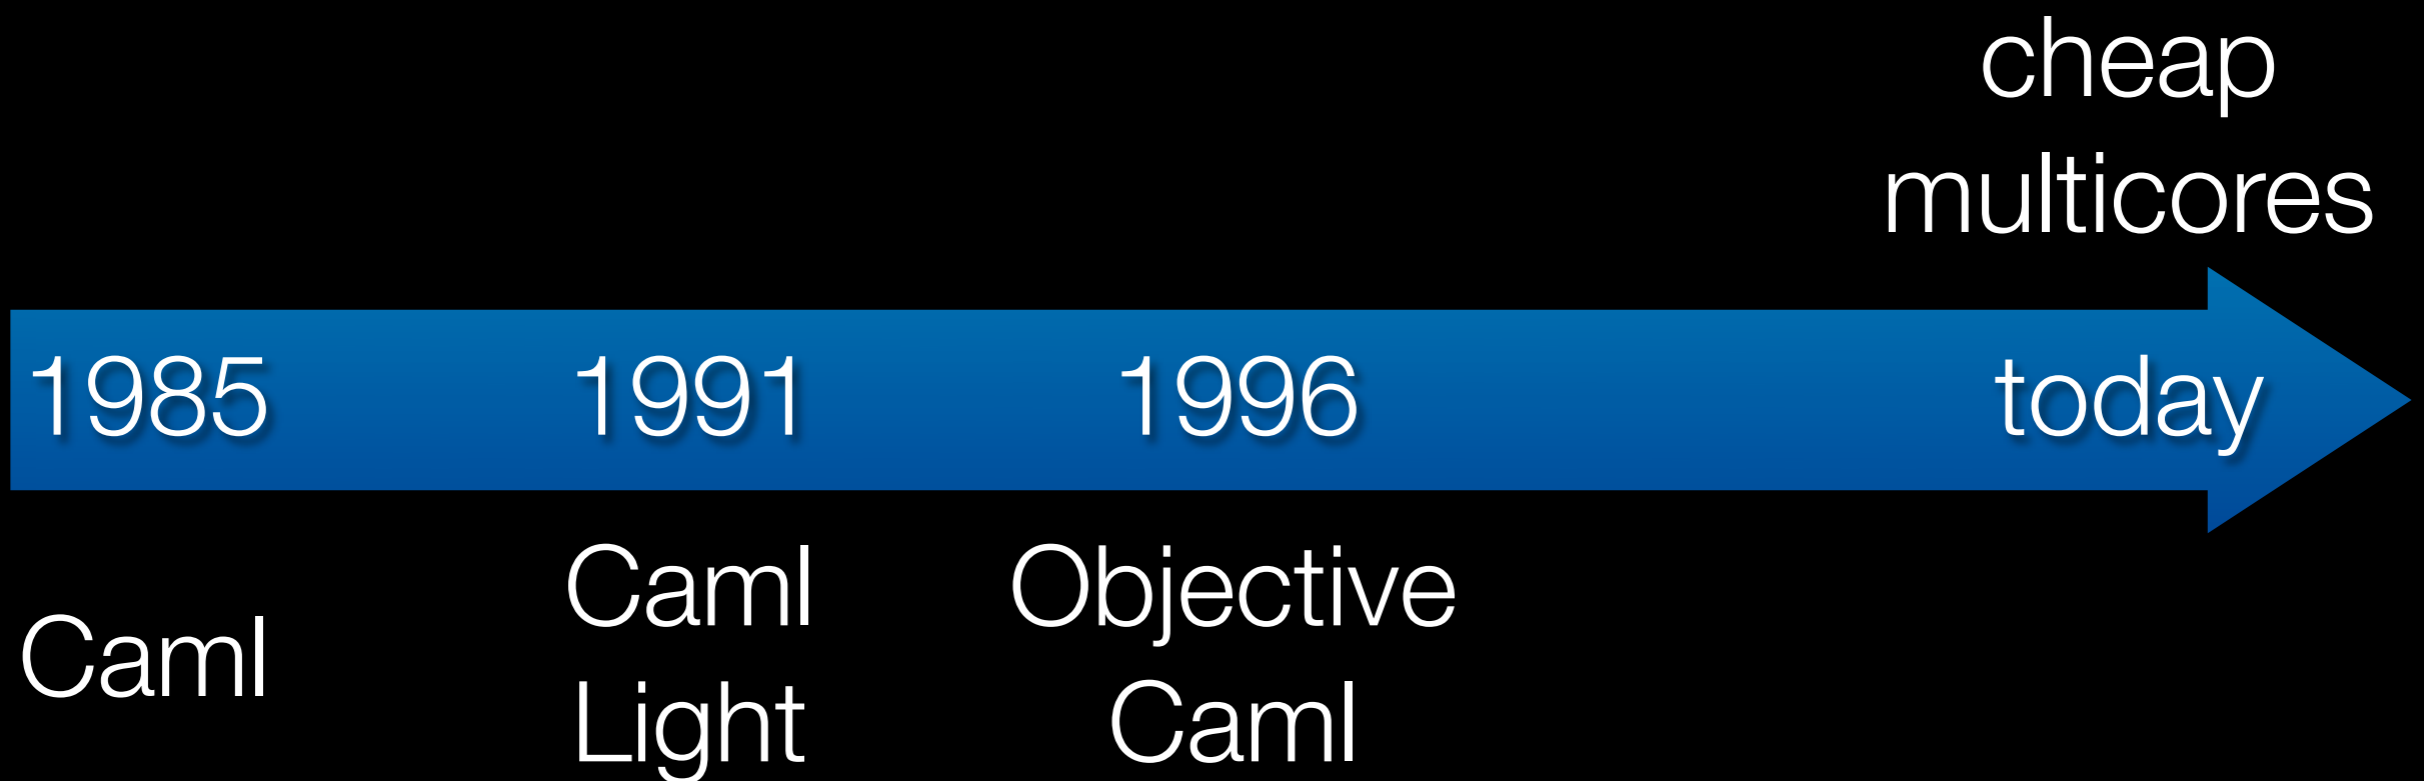
MIMD CPU

NVIDIA GeForce GT130

CELL

Intel Core Quad<sup>s</sup>

# some Caml history



# Objective Caml

- ✦ Functional-based multiparadigm language
- ✦ Distribution by INRIA
  - ✦ known for its efficiency
  - ✦ non-parallel concurrent threads

# some parallel threads issues

- ✦ runtime library support
  - ✦ reentrance (beware of shared static variables)
  - ✦ memory allocation/collection

# addressing some parallel threads issues

- ✦ runtime library support
  - ✦ reentrance (beware of shared static variables)
    - ✦ use (POSIX) `__thread` facility
    - ✦ transform to function parameters
    - ✦ refactoring
  - ✦ memory allocation/collection

# addressing some parallel threads issues

- ✦ runtime library support
  - ✦ reentrance (beware of shared static variables)
  - ✦ **memory allocation/collection**
    - ✦ look at the guts, be scared, run or make a choice
      - ✦ learn, understand, adapt
      - ✦ learn, remove parts & rewrite from scratch

# INRIA OCAML

# memory management

- ✦ **allocation**

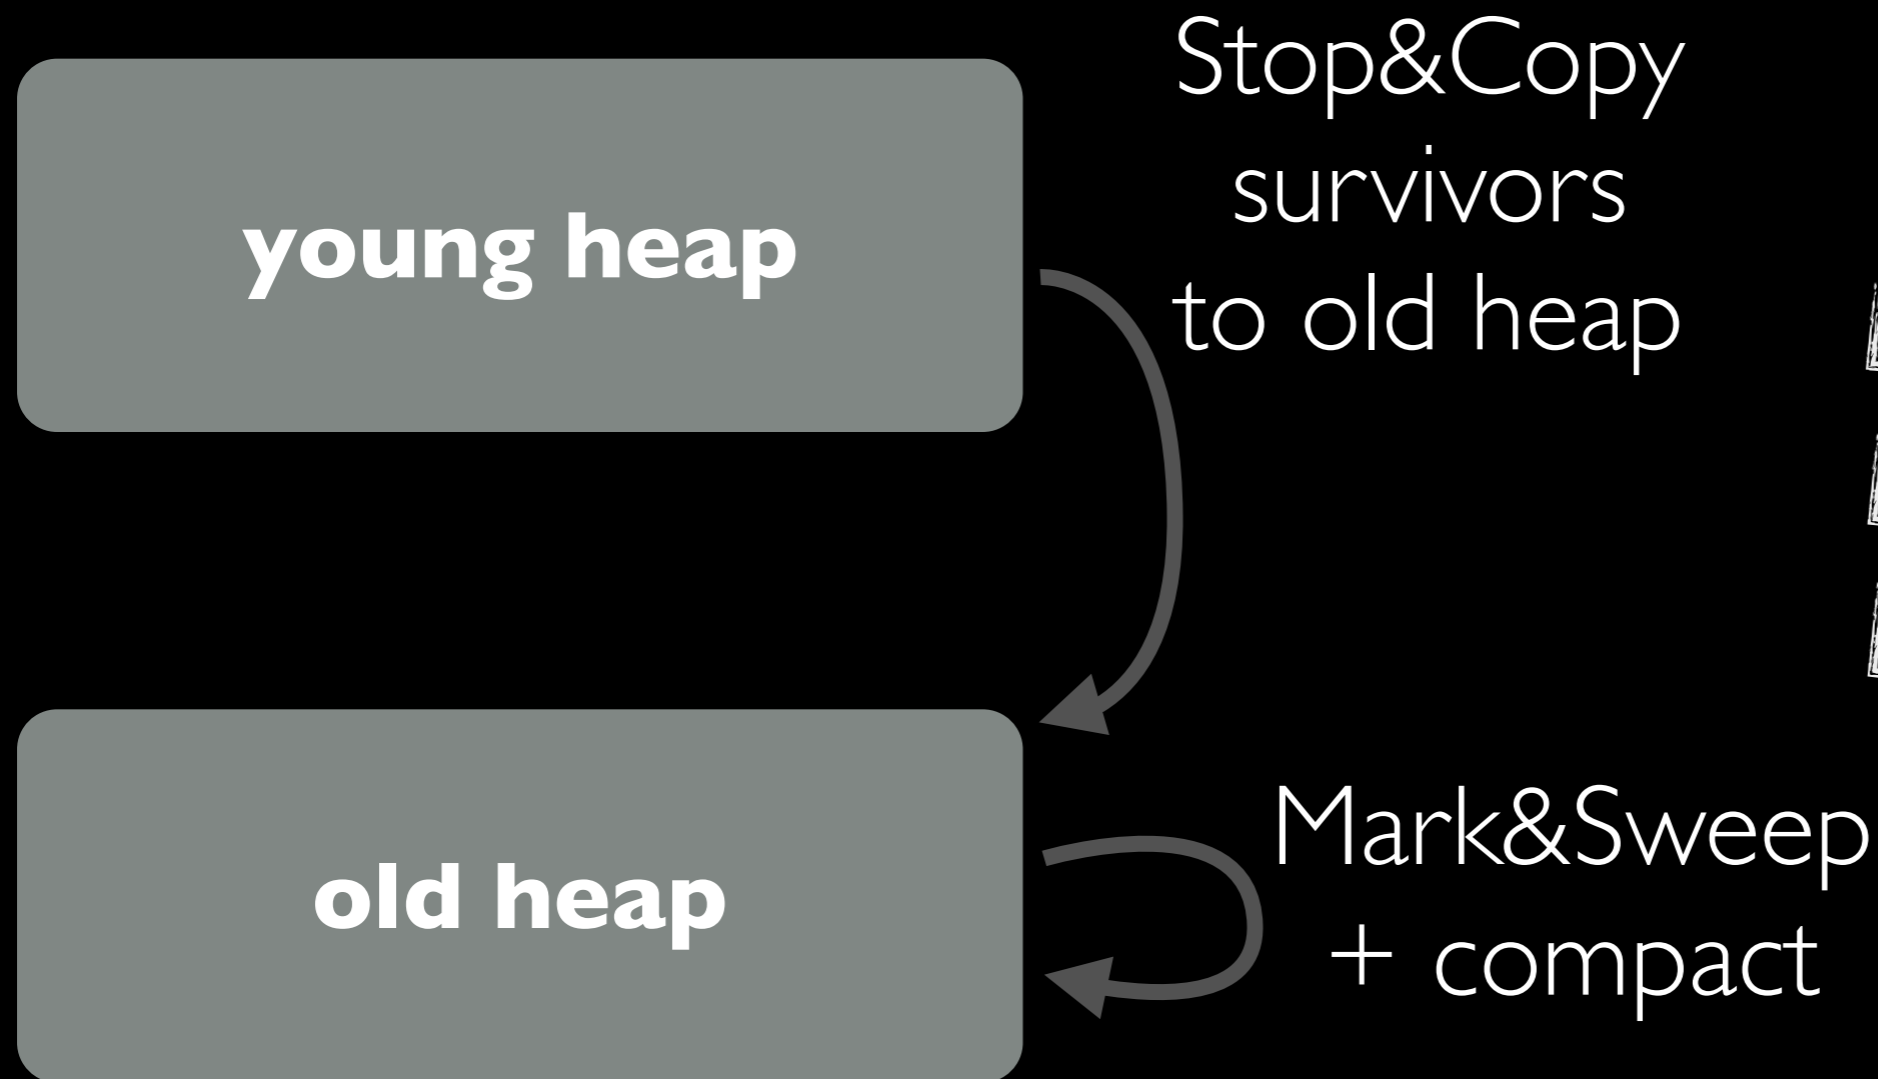
- ✦ `heap_ptr = heap_ptr - size(val)`

- ✦ **two-generation collection**

- ✦ minor collection: stop&copy

- ✦ major collection: incremental mark&sweep&compact

# INRIA OCAML garbage collection



**very efficient**  
**generational**  
**incremental**

***(runtime lib implementation in C code + asm)***



chapter 2

# OCaml for MultiCore

(influence of the)



# past experience

**Pagano et al.** (JFLA 2007, PADL 2008, ICFP 2009)

- ✦ alternative runtime lib implementation for software certification required by civil avionics norms
  - ✦ memory management too hard to explain thus impossible to certify
  - ✦ remove concurrency, marshalling, weak pointers, ...
  - ✦ result: from 16 000 lines of C code to 4 500

# addressing some parallel threads issues

- ✦ runtime library support
  - ✦ reentrance (beware of shared static variables)
  - ✦ memory allocation/collection
    - ✦ look at the guts, be scared, run or make a choice
      - ✦ learn, understand, adapt
      - ✦ learn, remove parts & rewrite from scratch

and what about scheduling?

# interdependencies



if the memory management implies moving values,  
then it can stop a thread from accessing values...

# OCaml for MultiCore

- ✦ **new memory management**
  - ✦ **stop** the world & two-generation **copy**
  - ✦ “inherited” mechanisms
    - ✦ stop at allocation, blocking operations
- ✦ **parallel threads**

# OC4MC Partial Collection


**Stop&Copy**

pages


shared heap

thread 1    thread 2    thread 3


Before



**Partial  
GC**

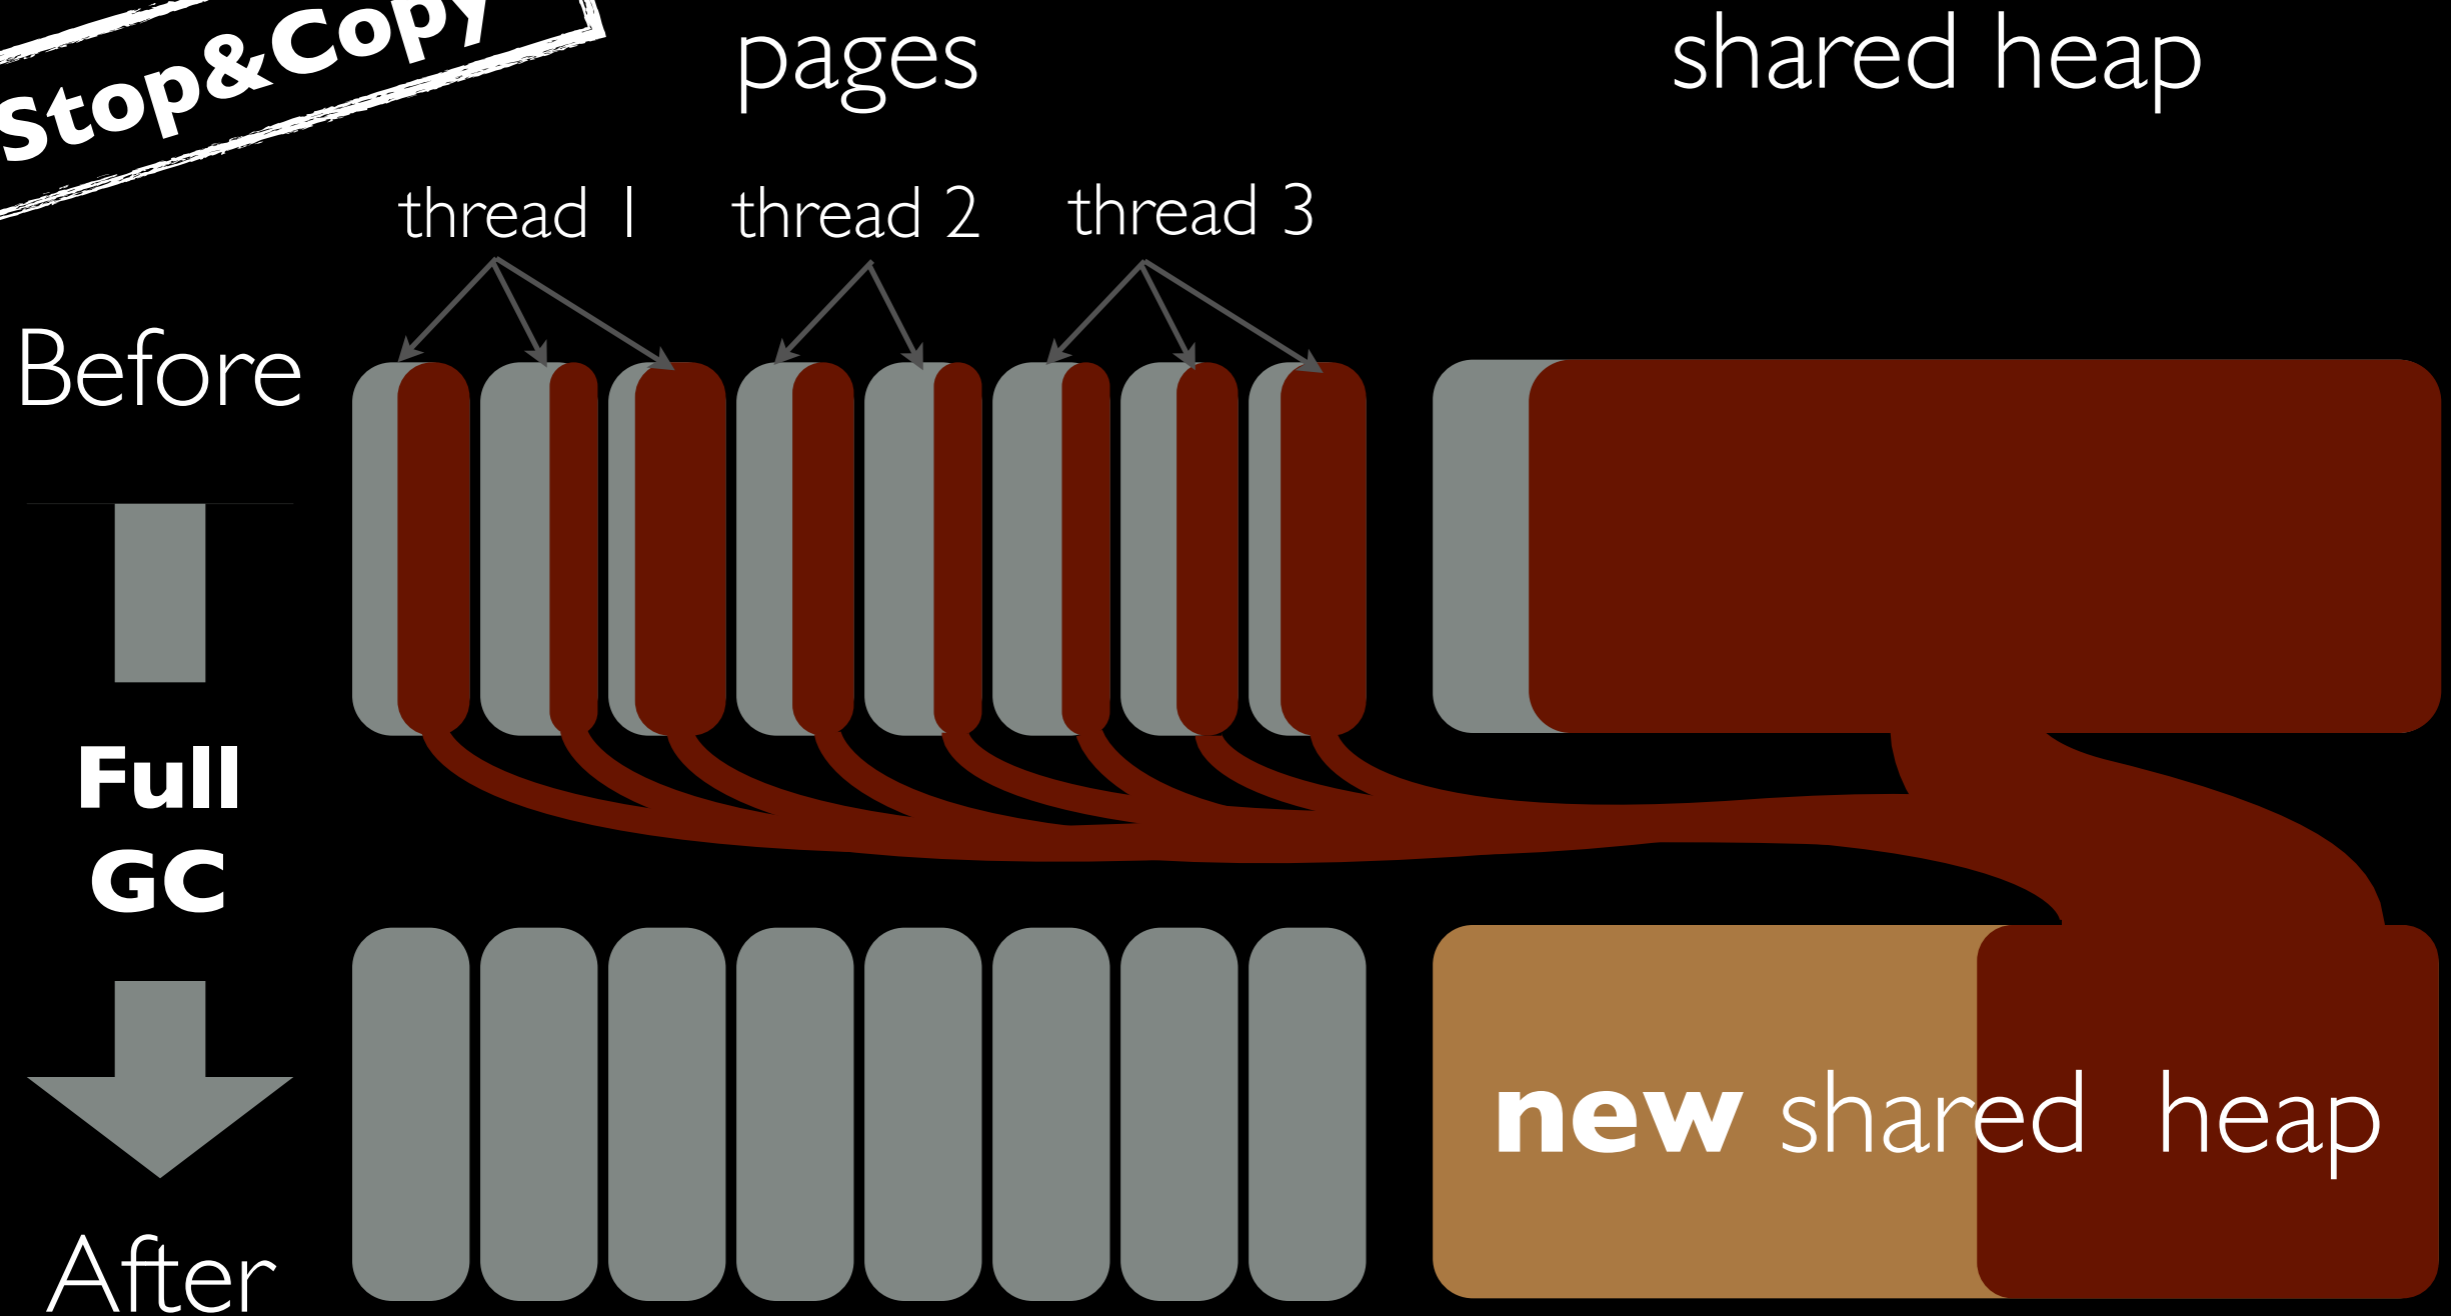


After

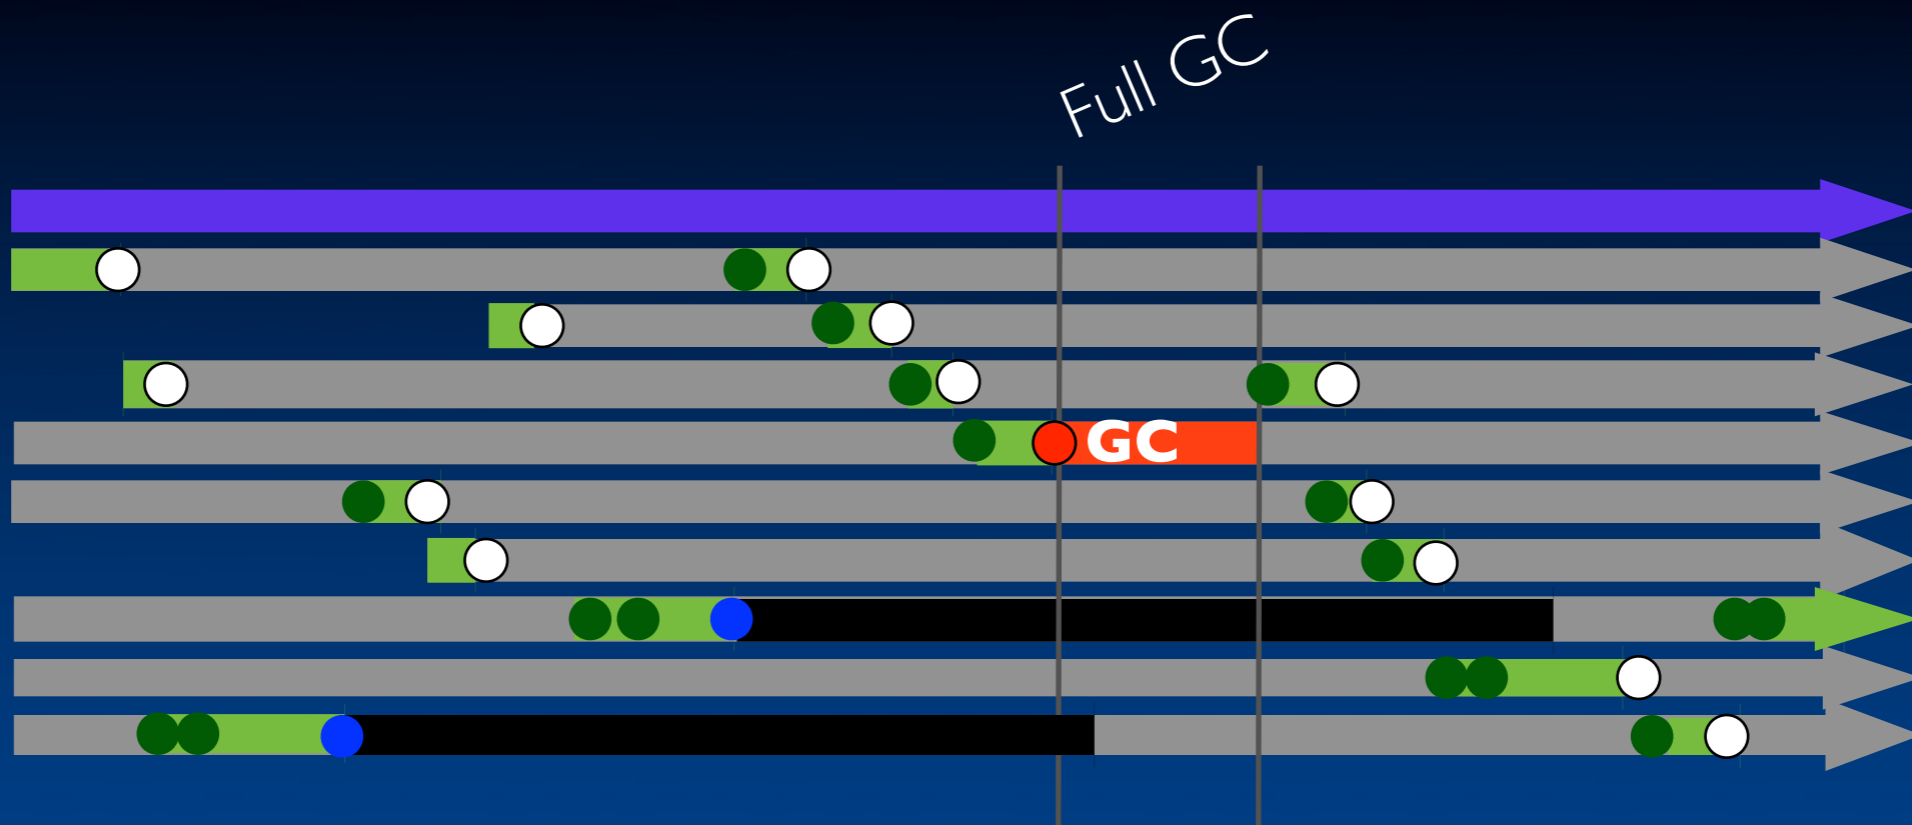


# OC4MC Full Collection

**Stop&Copy**




# Execution Sample



- allocation failure
  - suspended allocation
  - blocking operation
  - allocation ok
- blocking operation
  - thread is running normally
  - thread is sleeping
  - tick thread



## Execution Sample



chapter 3

# performance

# Speedups

|                           | <b>O'CamI</b>        | <b>OC4MC</b>      |                      |                      |                      |                      |
|---------------------------|----------------------|-------------------|----------------------|----------------------|----------------------|----------------------|
| # of threads              | 1TH                  | 1TH               | 2TH                  | 4TH                  | 8TH                  | 16TH                 |
| Sieve<br><i>speedup</i>   | 60s<br><b>1.06</b>   | 64s<br><b>1</b>   | 32s<br><b>2</b>      | 16s<br><b>4</b>      | 10s<br><b>6.4</b>    | 9s<br><b>7.11</b>    |
| Matmult<br><i>speedup</i> | 15.5s<br><b>1.17</b> | 18.2s<br><b>1</b> | 9.5s<br><b>1.92</b>  | 4.7s<br><b>3.88</b>  | 2.5s<br><b>7.28</b>  | 2.5s<br><b>7.28</b>  |
| Life<br><i>speedup</i>    | 24.3s<br><b>1.02</b> | 24.7s<br><b>1</b> | 16.6s<br><b>1.49</b> | 13.7s<br><b>1.80</b> | 15.1s<br><b>1.64</b> | 15.2s<br><b>1.62</b> |

little number of threads

|                 | <b>O'CamI</b> | <b>OC4MC</b> |
|-----------------|---------------|--------------|
| Sieve CML-style | 89s           | 59s          |
| <i>speedup</i>  | <b>1</b>      | <b>1.50</b>  |

great number of threads

# Slowdowns

- ✦ GC algorithm kept as simple as possible
  - ✦ heap growth can slow down the program
  - ✦ very functional style (many short life objects) combined with some long life objects shows stop&copy weakness
- ✦ Predictable weakness

chapter 4

# **conclusion & future work**

# Conclusion

it's working!

<http://www.algo-prog.info/ocmc/>

conclusion

## open source distribution

- ✦ Linux x86 64-bit
- ✦ Alternative multicore capable runtime library
  - ✦ memory manager replacement
  - ✦ thread library replacement
- ✦ Available as a patch for OCaml 3.10.2
  - ✦ <http://www.algo-prog.info/ocmc/>

# Conclusion

- ✦ working multicore capable threads for OCaml
- ✦ proof of feasibility
- ✦ potential good performance

<http://www.algo-prog.info/ocmc/>



# Related & Future Work

- ✦ OCaml concurrent/parallel extensions
  - ✦ what if they are used with OC4MC
- ✦ Use a limited number of threads by implementing abstractions on system threads
  - ✦ e.g. Parallel Concurrent ML
- ✦ Alternative GC algorithm

<http://www.algo-prog.info/ocmc/>