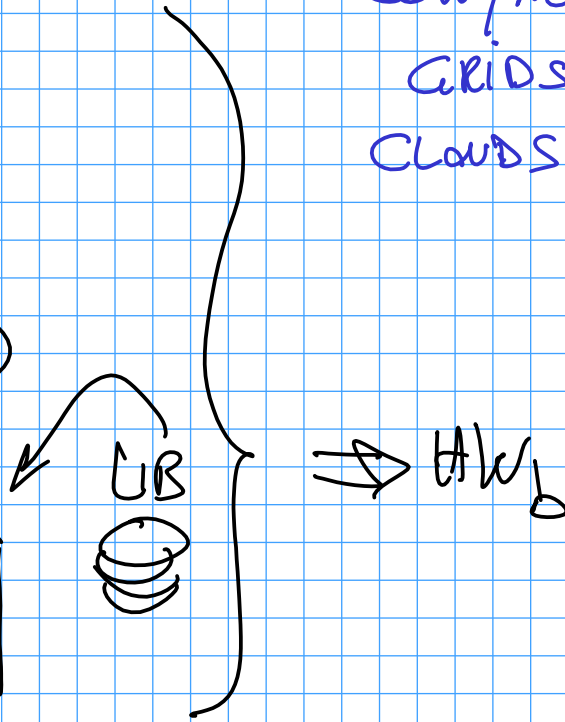




PORTABILITY




functional portability
performance portability

PARALLEL ARCHITECTURES



Example:




need to "adapt" my program

Functional

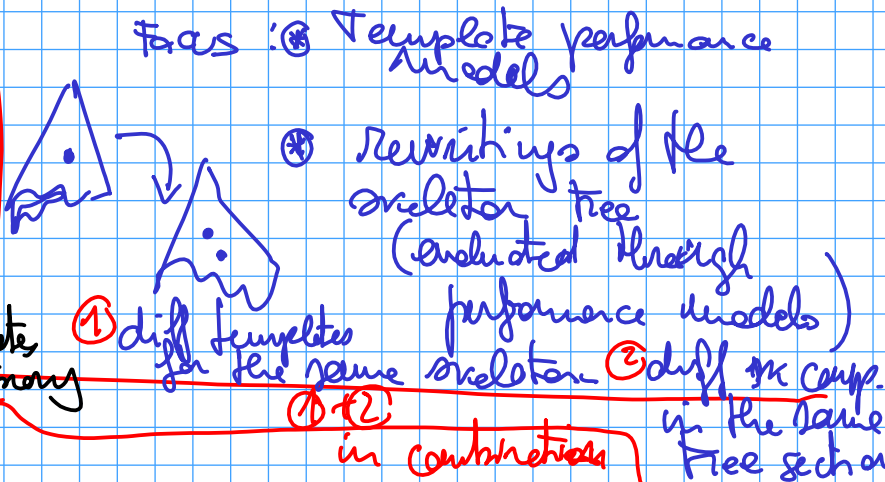
Performance

PORTABILITY


FUNCTIONALITY



HWA → HWB: choosing of different templates from the library



① ② in combination in the same Free section



PAR MDF INTERPRETERS

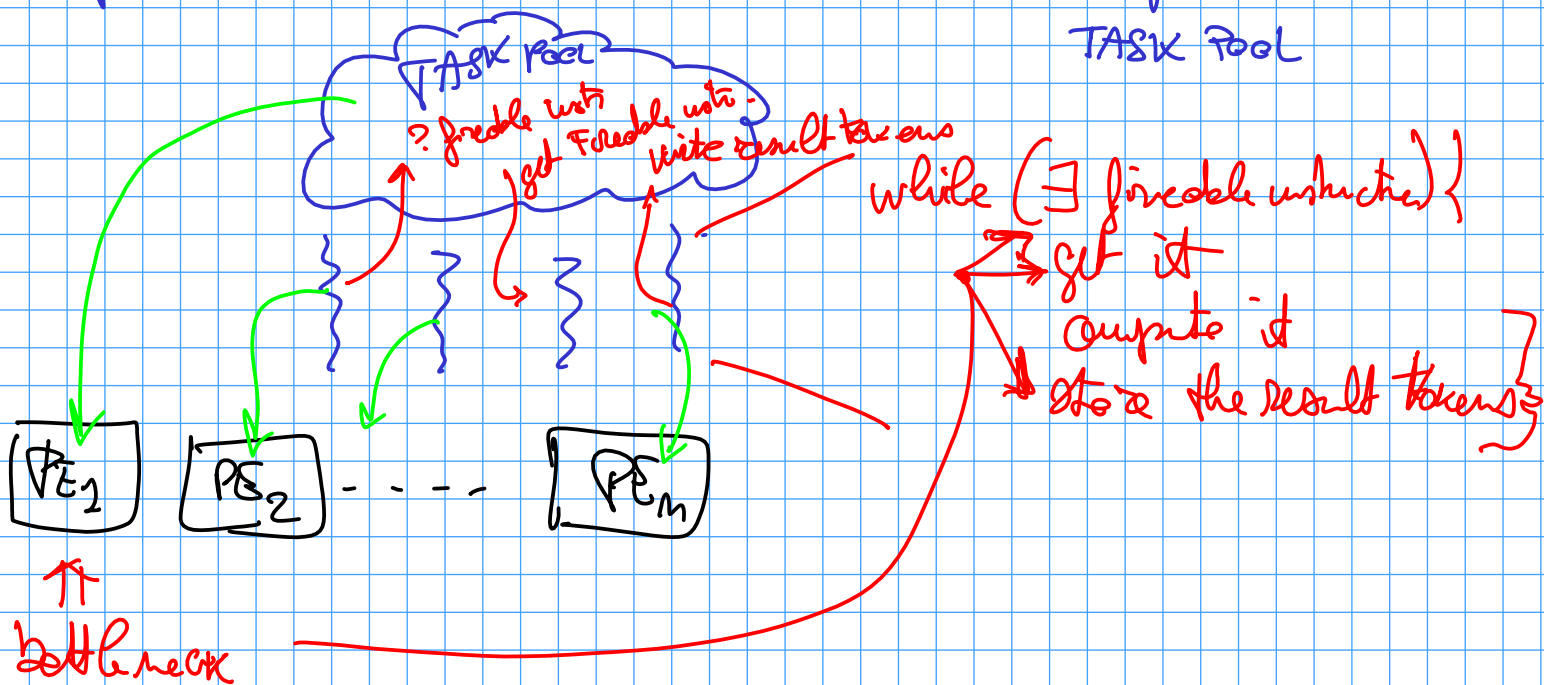
HWA → HWB

port PAR MDF INTERPRETERS HWB

- ① New interpreters as efficient as possible
 - ↳ implement as efficiently as possible the "core" ops of MDF interpreter
 - access to repository
 - fireable instruction "list"
 - executing fireable instructions


MDF → Cow/now

problem: implementation of the MDF instruction repository



"Software" RING connecting a number of "local" TASK Pools such that they look like a single global pool.

INIT: merge MDF graphs over $\forall TP_i$



TP manager::

if (# findable instruction < LOW.W. MARK)



ask findable instructions from neighbour TP_i

if (# findable instruction > HIGH.W. MARK)


start answering findable instructions requests from other TP_i

answer local ops if possible


+ 2 RINGS



cow (16 nodes, single core, with TCP/IP)




WORK STEALING



if TP_i is requested to deliver a fixable instruction with NO fixable instruction in TP_i :

then try to steal a fixable instruction from another TP_j


(random victim selection)



Portability through virtual machines


Skeleton program

⇒ compile ⇒ VM object code



JVM

executes bytecode



what happens to performances?

MPI

provides a set of comm. instructions




drivers for different hw

library (200+ entries) (calls)

runtime lib (10 entries)

TCP/IP drivers	SYS V drivers	Posix driver	Mellorix driver
COM/NOV	SMP Multicores		



SMP metacore (x86_64 ~ 60 cores)

MPI program written for COM + TCP/IP
↳ asynch comm

hdl = start-send (. . .)
compute something else

if (check (hdl) == completed) go on
else wait (hdl)

OpenMP

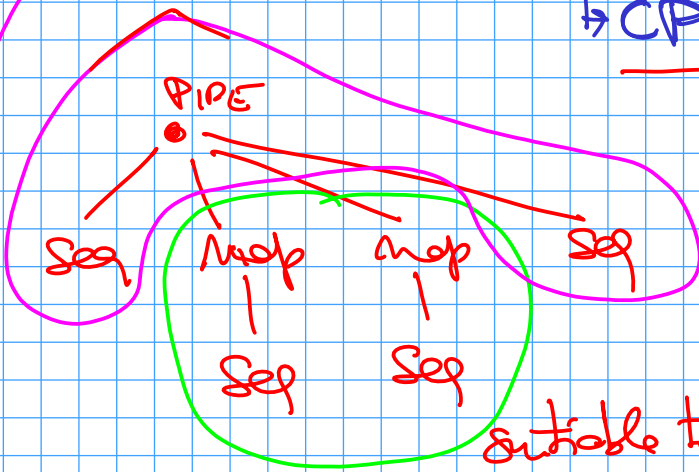
#pragma omp parallel for (. . .)
for (. . .) { . . . }

hw₁ shared L2

hw₂ shared M

PORTABILITY → heterogeneous architectures

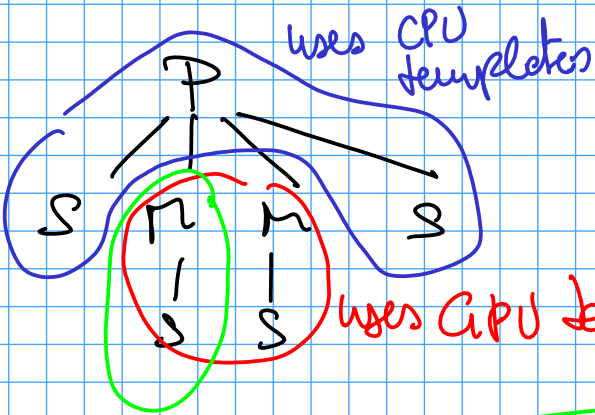
→ CPU + GPU



only suitable to target CPU cores

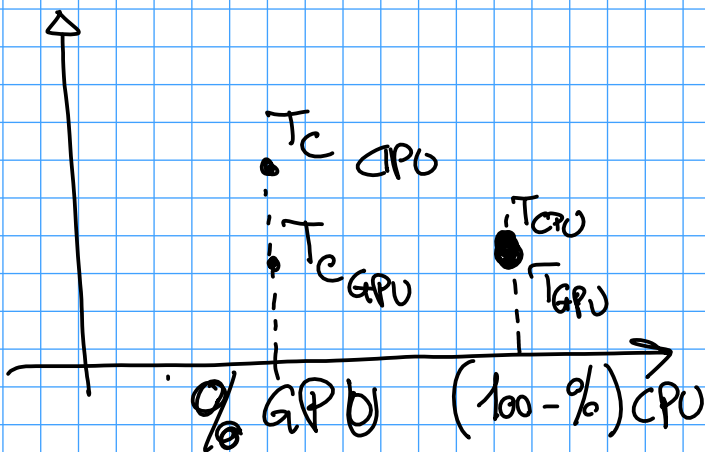

suitable to be executed on GPU cores

targetability : use techniques supporting portability to support usage of different target hw at the same time !



hw: i3 + GPU Intel (OpenCL)

Needs to know how to split the overall task among CPU & GPU



#task

$$T_{CPU}(\#t) = f(\cdot)$$

$$T_{GPU}(\#t) = g(\cdot)$$