Behavioural skeletons

Concern$_1$ ... Concern$_n$

$\{experts\}_1$ ... $\{experts\}_n$

$\{\ \}_i \cap \{\ \}_j \cong \emptyset$

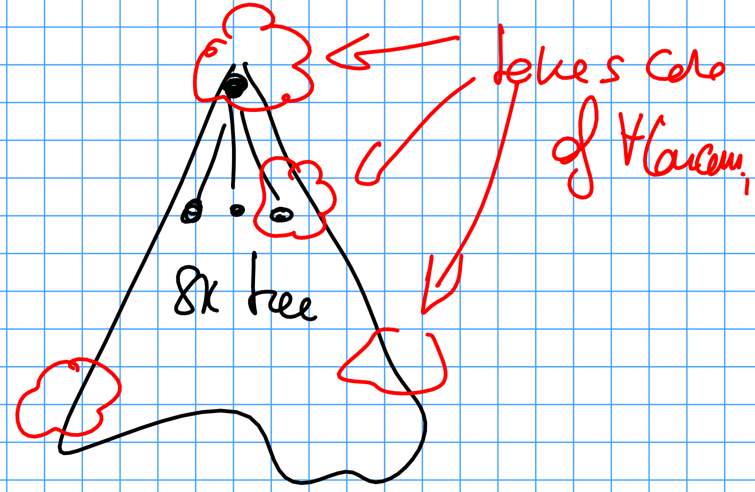

linear combination of all the measures relative to different concerns

Power(#PE, consumption$_i$, ...) $\in [1 \div 1000W]$

Performance(#PE, skeleton tree, comm bandwidth, processing power of the nodes, ...) $\in [1 \div 1000$ GFLOPS$]$ $\in [\mu sec \div sec]$

$\forall$ concern c_i $\quad f_{c_i}(p_1 \dots p_n)$ $\quad p_1 \dots p_n$ parameters hw, sw

$\min \sum f_{c_i}(\quad)$ $\qquad \min \sum_i a_i f_{c_i}(\dots)$

What about "boolean" concerns?

e.g. security

Framework BS

→ write structured parallel programs

→ write rules as "manager programs"

expert may use:

$\{actuators_1 \dots actuators_n\}$

$\{sensors_1 \dots sensors_m\}$

different ("disjoint") set of experts

⇓

each one gives good policies for their own managed concern

⇓

need coordination among these policies

how can I decide that

Decision$_i$ made by MGR_k is in contrast or fits Decision$_j$ made MGR_h ?

which is the common ground I can use to evaluate the decisions ?

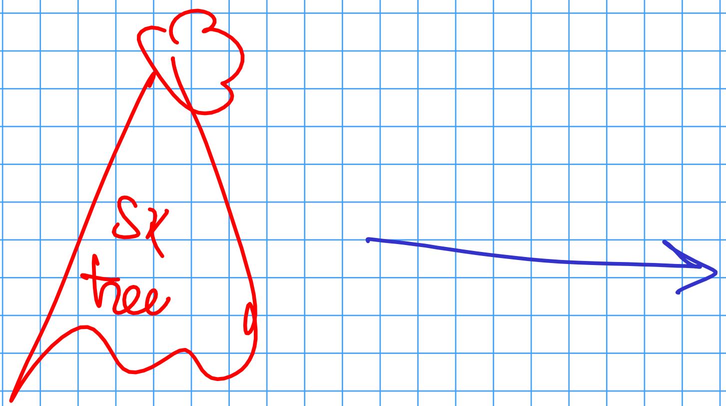

HW/SW configs

graph of the "machines" used (core, GPU, PE) with the comm. channels

decision "MW++" taken by MGR_{perf} :

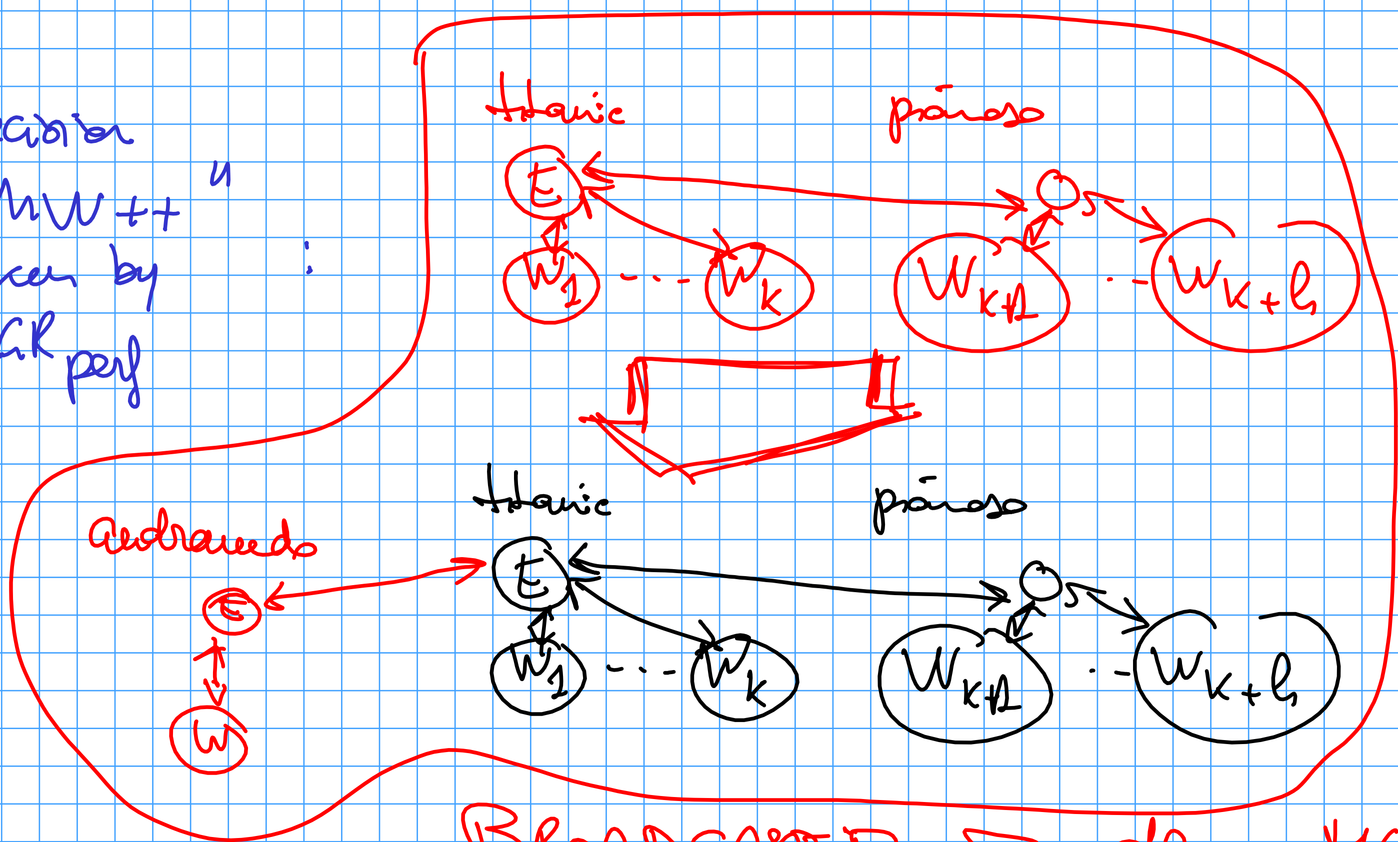

BROADCASTED TO other MGR_i

MGR_{power} < OK / NOK / OK provided ... properties

OK (MW++) but the new PE running the worker should consume less than k WATTS

MGR Perf

R_i :: precondition (true) → $a_1 ; a_2 ; a_3 \| a_4$ (plan)

↑ deploy (E, W code on andromeda)

↑ start code on andromeda

plan $a_1' ; a_2' ; a_3'$

different set of actions guaranteeing property required by the other manager

If this is possible: OK provided ... ⇒ OK with ≠ plan

If not possible: OK provided ... ⇒ NOK