structure ++ - - - → policies ++

MM → skeletons

structure ↓ MDF → PO interpreter

• map
|
seq (IP)

A = coll of rows   B = coll of columns

∀ row in A and col in B → IP(row$_i$ x col$_j$) → C_{ij}

map pattern

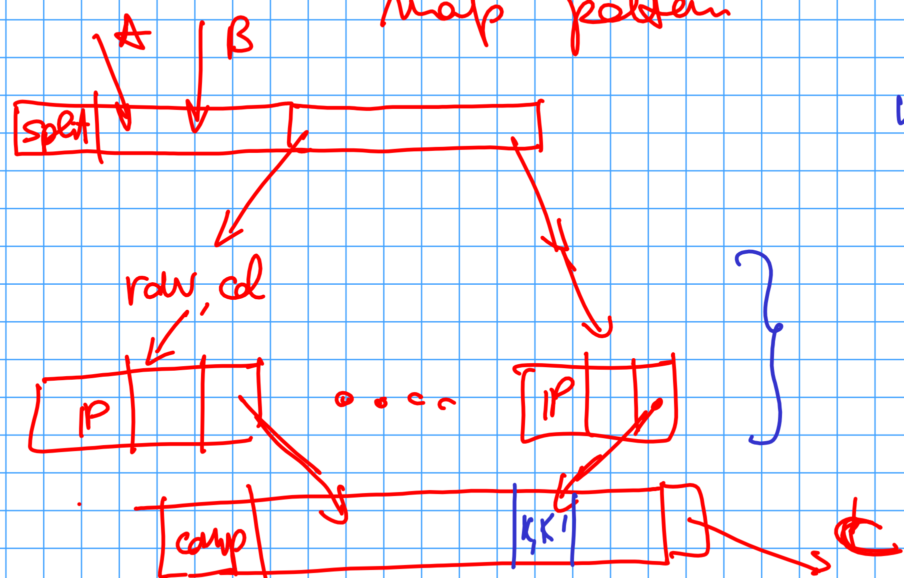

interpreter ⇒ MW = 16 (cow)

A, B, C   128 x 128

PE$_i$

IP R_k, $C_{k'}$

IP | R_k, $C_{k''}$

send ref   value

async tasks

TASK POOL MANAGER

affinity scheduling techniques

unbalanced computations perform better in MDF than

Template for map

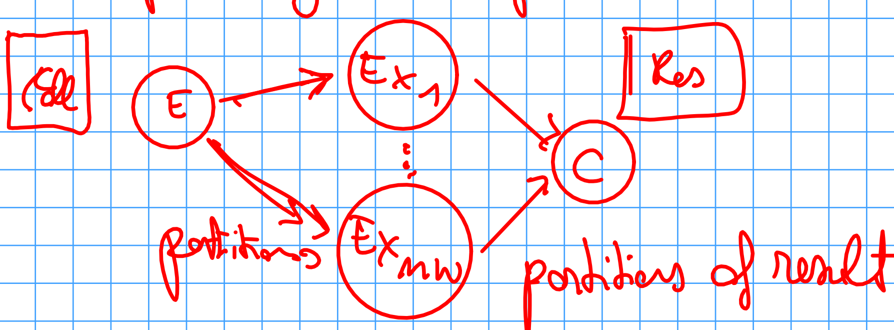

on balanced comp the static template is better

↳ P3L (early '90)

MDF gain (unbal comp) 30% 50%

loose (bal comp) 5%-7%

COMPOSITION

DATA // STREAM //

MAP (PIPE (S_1 , S_2))

PIPE (MAP (S_1))

STREAM // DATA //

STATE fwd

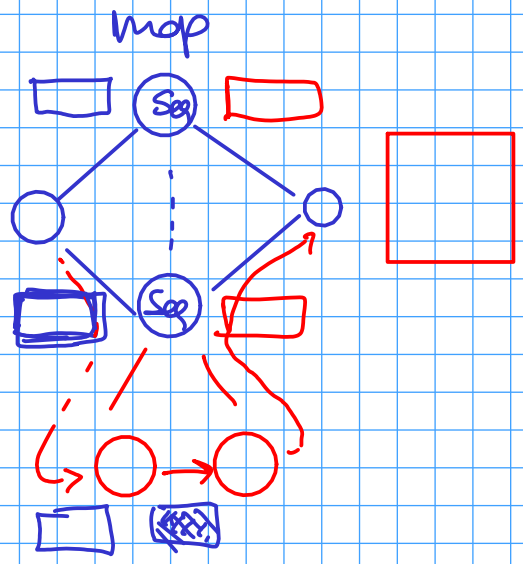

TWO TIER MODEL

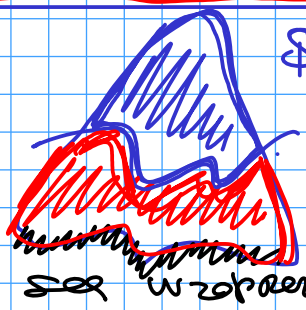

Stream parallelism

Data parallelism

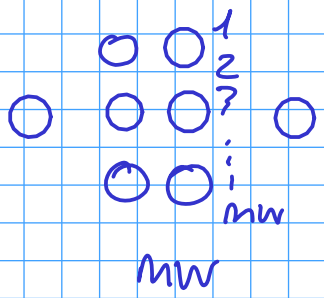

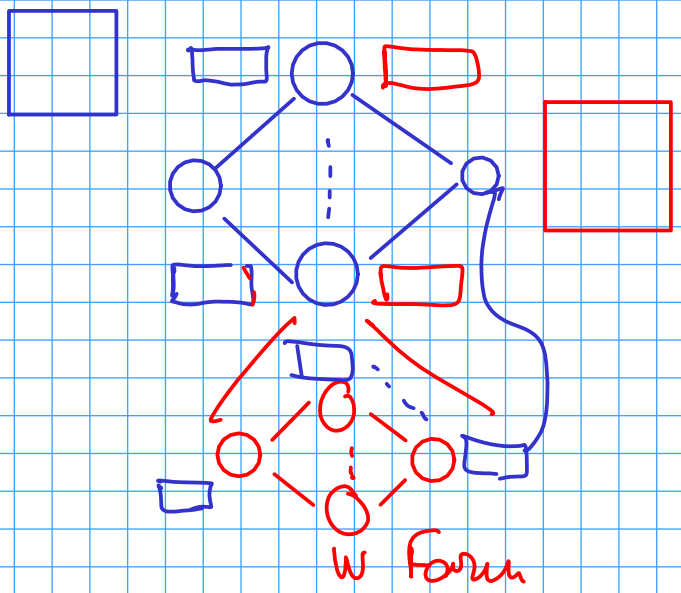

STATE

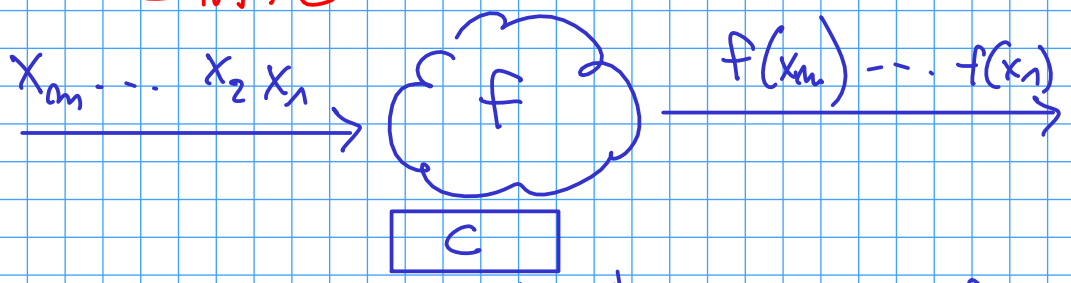

counts the items computed so far

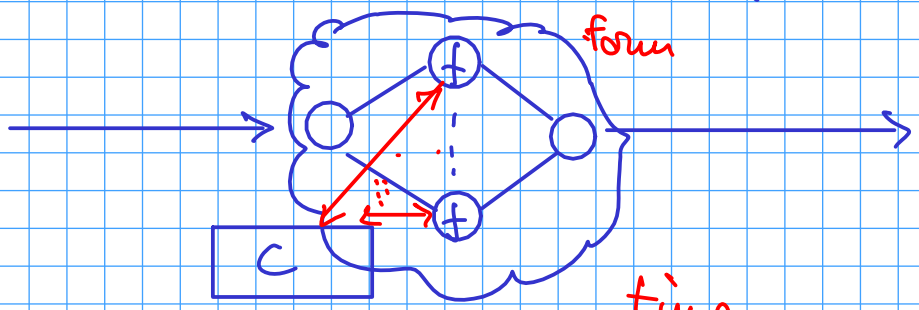

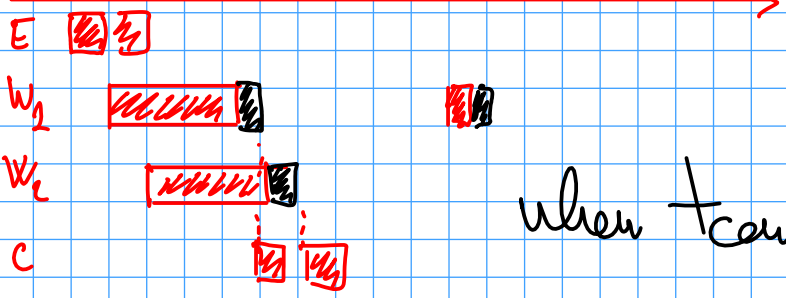

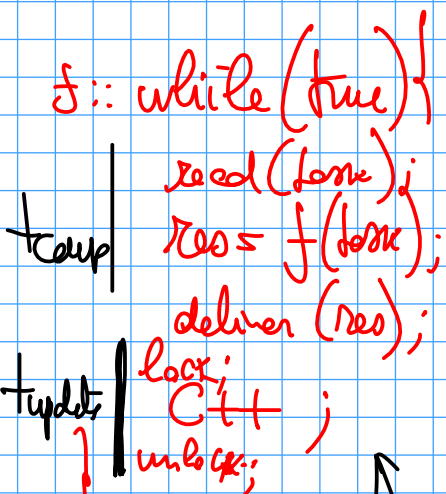

when t_{comp} is executable t_{update}

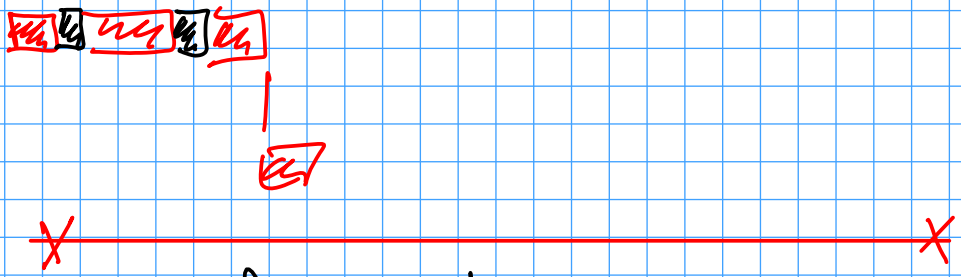

Patterns for state

1) Read only data structures

↳ read them as many times we need with limit of the access bandwidth

2) "accumulator" variable x ⇒ parallel reduce

$x = f(x, c)$

read

associative and commutative