



DF implementation



TASK:

Reduce overheads



ONCE & FOR ALL

- setup / termination

PER NODE computed

• fireable nodes

↑ moved to interpreters + code & + data


- synchronize accesses to the repository

↓

- move results (tokens) back to the repository


main overhead introduced (RUN TIME)

interpreters



✓ interpreter ⇒ diff PE

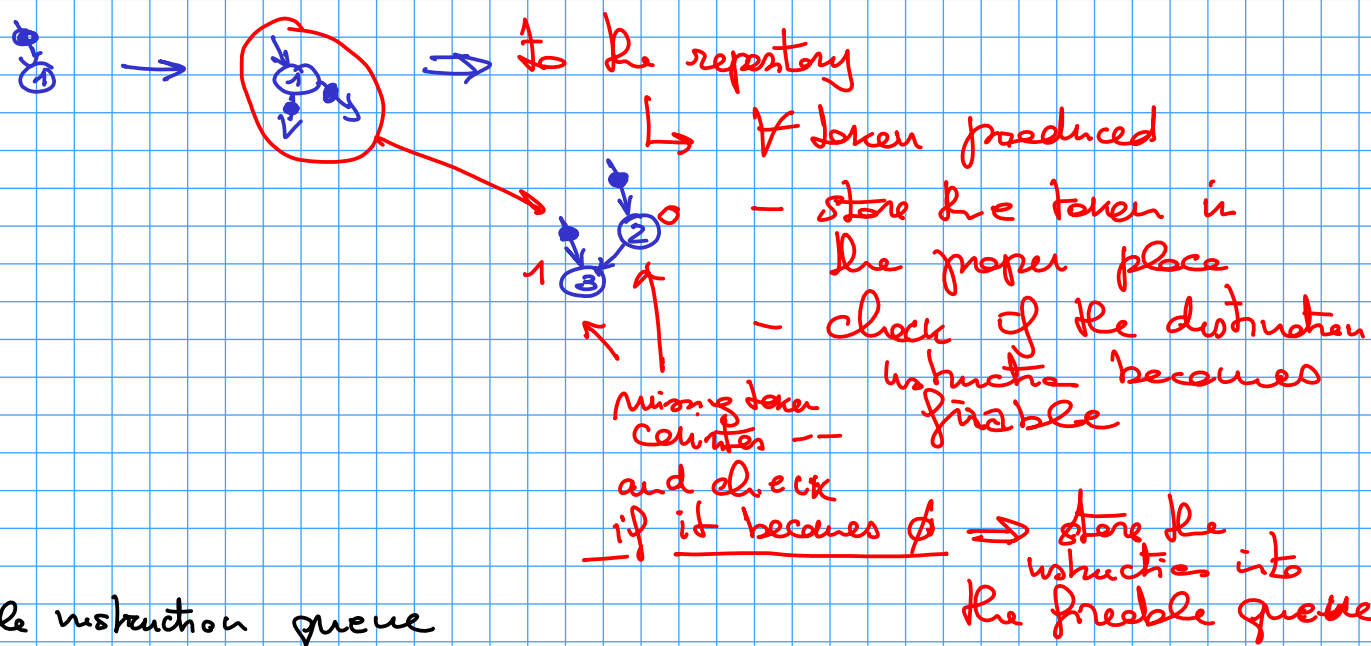
find fireable instructions / node



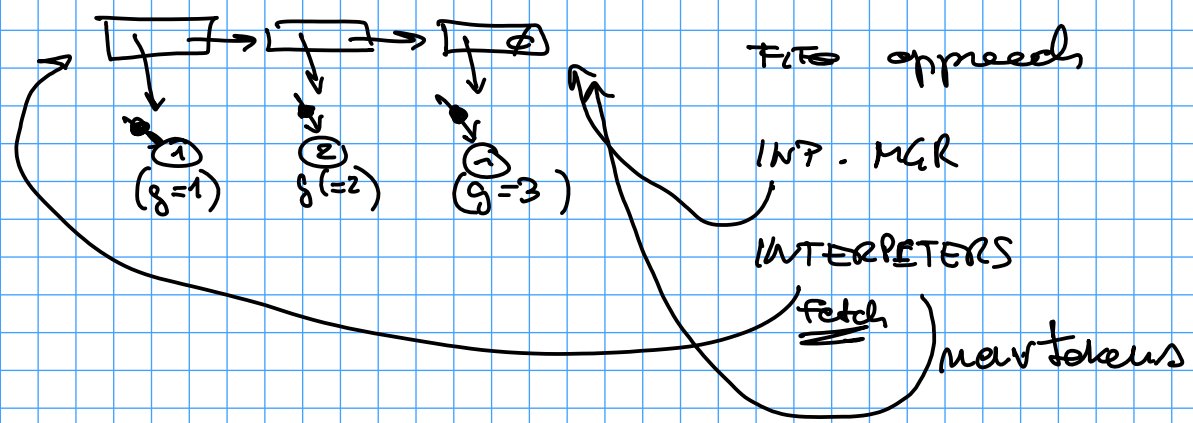
1) input data set available (INPUT MANAGER)





2) interpreter x produces new tokens after computing some instruction


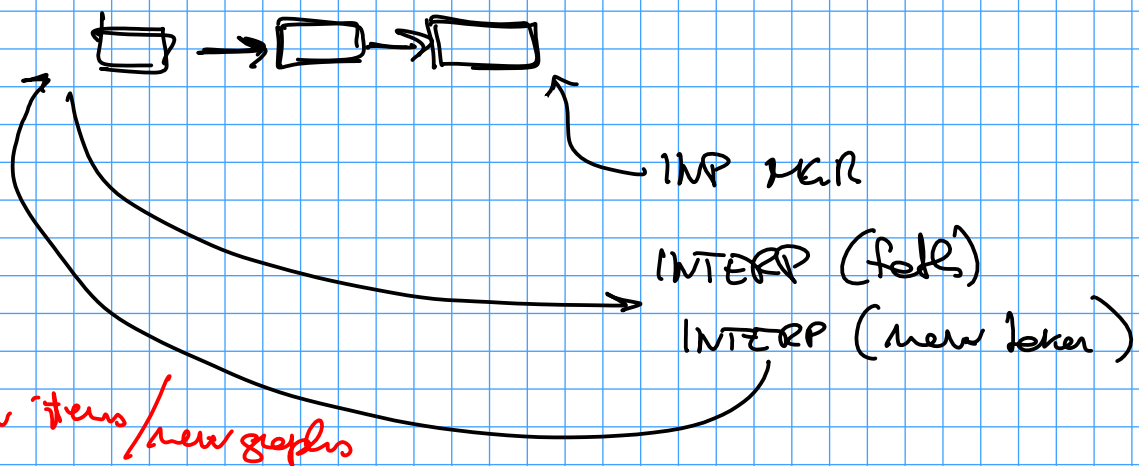



fireable instruction queue

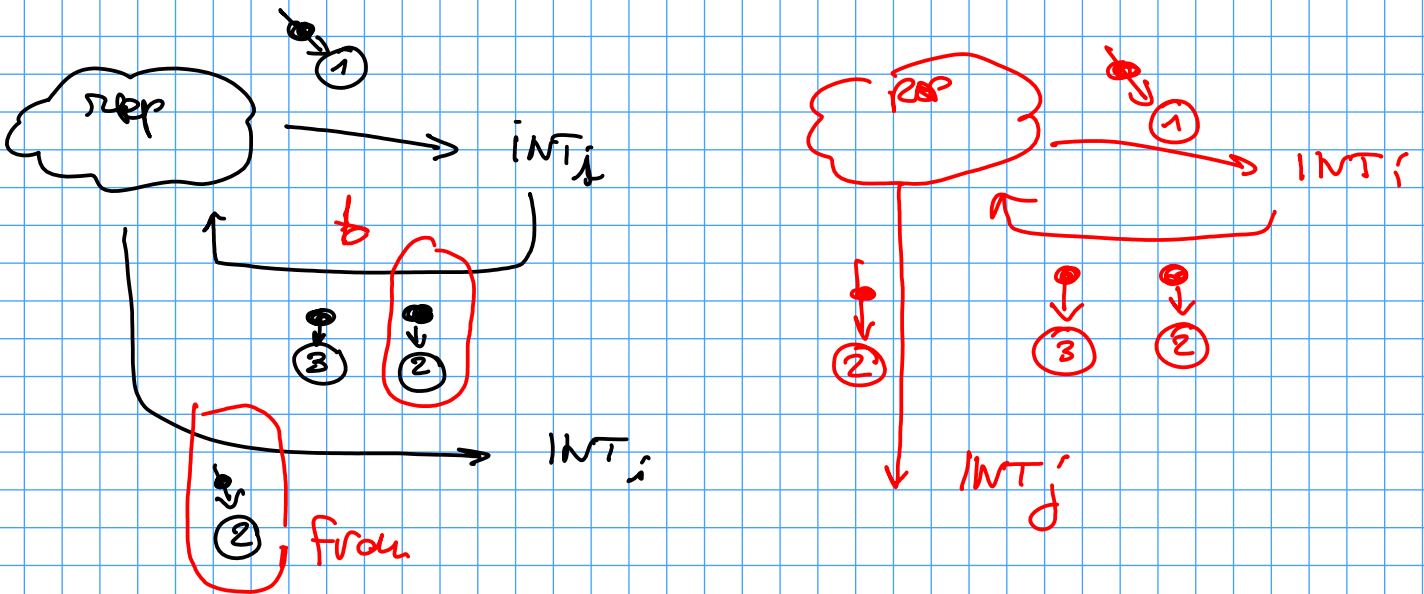





FIFO policy + priority on the graph id



source overhead: moving data & code to/from interpreters



Reduce overhead \Rightarrow avoid communications


INT_i sends source data to REP
& keeps a copy in a local cache

REP sends ② (ref to the locally cached data)
instead of sending copy of the whole data

Further optimize:


REP sends ① \rightarrow ②
to INT_i

INT_i computes ①
sends ③ to REP
& the computes ②




different kind of graphs with different needs as for as caching is concerned

similarities in directory based cache coherence protocols
 & managers in the repository
 + interpreters talking each other




⇒ schedule the black graph to INT₁
 & the blue one to INT₂



⇒ no data parallelism
 stream parallelism only

⇓
 good throughput
 bad latency

Reduce overhead ⇒ do apply better scheduling




⇓
 AFFINITY
 SCHEDULING

AFFINITY SCHEDULING

INT: get token_i ; computes produces token_j
and keeps token_i & token_j in
a local "cache" managed with LRU
(sw cache) policy

REPO: try to schedule the computation
graph node ($\text{param}_1 \dots \text{param}_n$)
 token_i token_j
where he knows token_x needed by the
computation have been produced

knowing that INT_k produced token_j

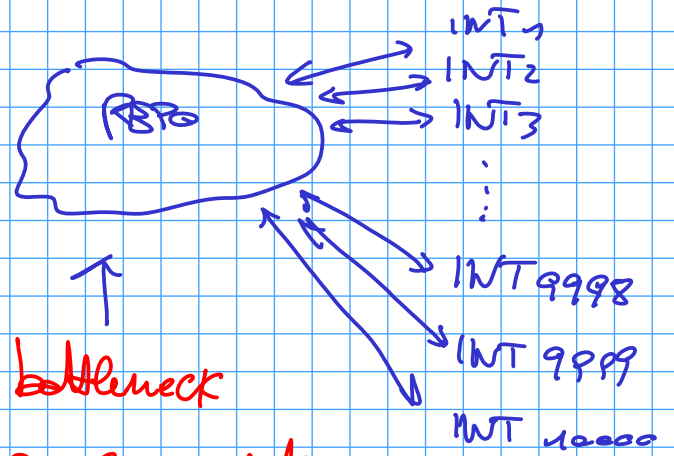



may experience a token
cache "fault"
then? → see (*)
below

- ⇒ when INT_k produces token_j :
- ① stores token_j in the local cache
 - ② deploys back token_j to the REPO

then? (*) raise a fault

- assign the token_j to the REPO
- wait
- when it comes: compute

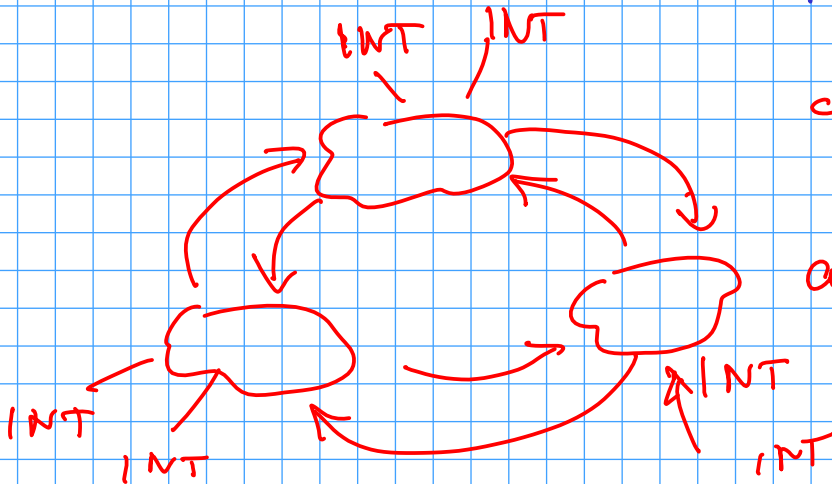



single point of failure

Reduce overhead?

How?

exploit locality \Rightarrow parallel repository



clockwise ring to push tokens

anticlockwise to fetch freeable instructions

