Work Stealing

Guillaume Pallez
Inria

M2 CISD, Enseirb-Matmeca,
Automne 2020

WS: Context

- A parallel platform with p processors

- A task-graph G tp be executed

- Non-clairvoyant setting : the structure of G and/or the execution times of its constitutive tasks are discovered online

Sharing v Stealing

Batch scheduling

Centralized scheduling

- A single list stores all ready tasks
- All processors retrieve work from that list

Advantage(s)

- Global view and knownledge

Drawback(s)

- Does not scale (contentions, etc.)

Sharing v Stealing

Batch scheduling
Centralized scheduling

- A single list stores all ready tasks
- All processors retrieve work from that list

Advantage(s)

- Global view and knowledge

Drawback(s)

- Does not scale (contentions, etc.)

Work stealing
Distributed scheduling

- Each processor owns a list of “its” ready tasks

Advantage(s)

- No contention problem
- Scalable solution

Drawback(s)

- Processors with empty lists do not know where to retrieve work from.

Stealing policies 1/2

Global round-robin

- A global variable holds the identity of the next processor to steal from
- Variable incremented after each steal (successful or not)
- Advantage : eventual progress
- Drawback : centralized solution...

Stealing policies 1/2

Global round-robin

- A global variable holds the identity of the next processor to steal from
- Variable incremented after each steal (successful or not)
- Advantage : eventual progress
- Drawback : centralized solution...

Local round-robin

- Each processor has its own variable indicating the next processor it should try to steal from
- Variable incremented after each steal (successful or not)
- Advantage : eventual progress ; solution is scalable
- Drawback : all stealing processors may attempt to steal from the same processor ; arbitrary notion of “distance” between processors

Stealing policies 2/2

Random stealing (Blumofe and Leiserson)

- The processor to steal from is randomly and uniformly chosen
- Advantage : decentralized ; scalable ; no notion of “distance” ; low probability of simultaneous steal from same processor
- Drawback : performance ? ? ?
  TODO@home : think about what can be a good model for performance of work stealing