Design of Parallel & High Performance Computing

Reasoning about Performance III

... Scheduling Proofs

Greedy scheduling: Analysis

Reminder: $T_1 = W$, $T_\infty = D$ $\quad T_p \geq T_1/p, T_\infty$ $\quad T_p \leq D + W/p = T_\infty + T_1/p$

Greedy scheduler:

$\#$ complete steps $\leq T_1/p$

$\#$ incomplete steps $\leq T_\infty$ (every incomplete step reduces critical path by 1)

$\Rightarrow T_p \leq T_1/p + T_\infty$

How far from optimal?

Assume T_p^* is optimal:

$$T_p^* \geq \max\{T_1/p, T_\infty\}$$

$$T_p \leq T_1/p + T_\infty \leq 2\max\{T_1/p, T_\infty\} \leq 2T_p^*$$ so at most factor 2 away

enough parallelism: $W/D = T_1/T_\infty \gg p \iff T_\infty \ll T_1/p$

$$\Rightarrow T_p \leq T_1/p + T_\infty \approx T_1/p$$

Work stealing scheduling: Analysis

Theorem: The scheduler achieves $T_p \leq T_1/p + O(T_\infty)$

Proof sketch: Every processor is either working or stealing

Total time working: T_1

Total time stealing: Every steal has $1/p$ chance to reduce the critical path; hence $O(p\,T_\infty)$

$\Rightarrow$ total time: $T_1 + O(p\,T_\infty)$

$\Rightarrow$ time per processor: total time$/p$ = $T_1/p + O(T_\infty)$

Note: space requirement $S_p \leq p\,S_1$.