

# Reduce

Mark Greenstreet and Ian M. Mitchell

CpSc 418 – January 10, 2018

- Finish Erlang Processes
- [Example Problem](#)
- [Table of Contents](#)



Unless otherwise noted or cited, these slides are copyright 2017 by Mark Greenstreet and are made available under the terms of the Creative Commons Attribution 4.0 International license <http://creativecommons.org/licenses/by/4.0/>

# Objectives

- Understand how reduce combines values using a tree.
- Describe the performance issues for reduce: trade-offs of time for computation and time for communication
- Describe 2 or 3 examples of reduce.

# CPSC 418 Poetry Competition

- The competition:
  - ▶ Everyone writes a poem.
  - ▶ Everyone submits to poem to Mark (the contest judge).
  - ▶ Mark reads all of the poems, compares them, selects the best poem.
  - ▶ The winner receives an original manuscript of the complete poems of [Lu Bai](#), signed by the author.

# Sequential Time for the Poetry Competition

- $N$  students in the class.
- $t_r$  to read and rank two poems.
- Total time  $(N - 1)t_r$ .
- Works fine until  $N$  becomes so large that we can't judge all the poems in a reasonable amount of time.

## Parallel version

- The procedure
  - ▶ Clone  $P$  copies of Mark.
  - ▶ Each Mark-clone reads and ranks  $N/P$  poems and sends the best poem to the original Mark.
  - ▶ The original Mark receives  $P$  candidates for the best poem, and selects the best one.
  - ▶ The winner receives the prize.
- Parallel Time
  - ▶ Each Mark clone takes time \_\_\_\_\_, but they all judge their poems in
  - ▶ The original Mark takes time \_\_\_\_\_.
  - ▶ The total time is \_\_\_\_\_.
  - ▶ Simplify this to get \_\_\_\_\_.
  - ▶  $SpeedUp = \frac{T_{seq}}{T_{par}} =$  \_\_\_\_\_.

# Bureaucratic Overhead

- To satisfy UBC privacy policies, the messages between the Mark-clones and the original mark must be sent in special envelopes.
- There's lots of special procedure for handling these envelopes, takes time  $\lambda$  to send or receive a message.
- The original Mark receives  $P$  messages from the  $P$  clones. This takes time  $\lambda P$ .
- The total time is now: \_\_\_\_\_.
- $SpeedUp = \frac{T_{seq}}{T_{par}} =$  \_\_\_\_\_.

Can we do better?

# Revenge of the Clones

- While Mark is working through the pile of envelopes, some of the clones realize that they could pair up and combine their results.
  - ▶ This costs  $\lambda$  time, the clones have to follow the rules as well.
  - ▶ The original Mark ends up with half as many envelopes to handle.
- What is the total time? \_\_\_\_\_.
- What is the speed-up? \_\_\_\_\_.

# Up a Tree with the Poetry

- The optimization on the previous slide worked great.
  - ▶ We're computer scientists, let's apply the optimization recursively.
  - ▶ Viewed from another angle, this is an example of divide-and-conquer.
- Combine the results in a tree.
  - ▶ How many levels in the tree?
  - ▶ How much time at each level?
- What is the total time? \_\_\_\_\_.
- What is the speed-up? \_\_\_\_\_.



# Is there more to life than poetry?

- Find the largest element in a list.
- Find the sum of the elements in a list.
- Count the number of 3s in a list.
- What to these all have in common?
- We'll look at more examples on Friday

# The Reduce Pattern

- We have a problem that takes  $T_{\text{seq}}(N)$  sequential time, where  $N$  is the “size” of the problem instance.
- We can divide this into  $P$  tasks with “perfect” speed-up:
  - ▶ Each task takes time  $T_{\text{seq}}(N)/P$  time.
  - ▶ Combining the results takes  $\lceil \log_2(P) \rceil \lambda$  time.
- $$\text{SpeedUp} = \frac{T_{\text{seq}}(N)}{T_{\text{seq}}(N)/P + \lceil \log_2(P) \rceil \lambda}$$
- What happens to *SpeedUp* as  $P$  goes large (for fixed  $N$ )?
- What happens to *SpeedUp* as  $N$  goes large (for fixed  $P$ )?