Parallel Computing Workbook
Hudson
April 2020

**Flipping Coins in Parallel**
[Model Website]

Q: If each worker can flip one coin per time step, how many time steps does it take the serial worker to flip two coins?
A: 2

Q: How many time steps does it take two parallel workers to flip the two coins?
A: 1

Q: Set the work size to 16. How many time steps does it take the serial worker to flip the coins?
A:16

Q: How many time steps does it take the parallel workers to flip the coins?
A: 8

Q: Set the number of parallel workers to 4. How many time steps does it take the parallel workers to flip the coins now?
A: 4

Q: Set the number of parallel workers to 8. How many time steps does it take the parallel workers to flip the coins now?
A: 2

Q: From what you've seen, what is one reason why it would be a good idea to use parallel workers instead of a serial worker?
A: it takes fewer steps to accomplish the same amount of work

Q: Decrease the number of parallel workers to 2. Decrease the max time to 2. In 2 time steps, how many coins can be flipped by 2 parallel workers compared to one serial worker?
A: 2 parallel workers can flip 4 coins combined; one serial worker can flip 2 coins

Q: Increase the number of parallel workers to 4. In 2 time steps, how many coins can be flipped by 4 parallel workers compared to one serial worker?
A: 4 parallel workers can flip 8 coins, while one serial worker can flip 2 coins

Q: Increase the number of parallel workers to 8. In 2 time steps, how many coins can be flipped by 8 parallel workers compared to one serial worker?

A: 16 coins can be flipped by 8 parallel workers, while only 2 coins can be flipped by serial workers

Q: From what you've seen, what is another reason why it would be a good idea to use parallel workers instead of a serial worker?
A: you can get considerably more work done with parallel workers than you can with serial workers.

Q: Increase the max time to 16 time steps. Decrease the number of parallel workers to 2. Decrease the max worker memory to 2 coins. If each worker can only hold 2 coins in memory, what is the maximum number of coins that can be flipped by 2 parallel workers compared to 1 serial worker?
A: 2 parallel workers flip 4 coins, while one serial worker can only flip 2 coins.

Q: Increase the number of parallel workers to 8. If each worker can only hold 2 coins in memory, what is the maximum number of coins that can be flipped by 8 parallel workers compared to 1 serial worker?
A:8 parallel workers can flip 16 coins, while one serial worker can flip 2 coins.

Q: From what you've seen, what is another reason why it would be a good idea to use parallel workers instead of a serial worker?
A: if memory is a constraint, parallel workers are more efficient.

**Human Parallel Computer - Data Parallelism through Forest Fire Simulations**
Model Website

My number: 13
Total number of students:16
My probability:13/17
Percentages:99.3 97.92 99.65 98.61 98.96
Iteration counts:17 18 19 17 17
Average percentage:92.888%
Average # of iterations:17.6
Q: What were some of the **tasks** we did in this exercise? What were they, and who did them?
A: each person had a different fraction which they plugged into a forest fire simulation. We then shared the data we gathered and graphed it.

Q: What kinds of **data** did we work with in this exercise?
A: We got the percent of trees burned, and the number of time steps it took to complete each fire. Our ID, the average %burned, and the number of time steps, our graphs.

Q: In which steps was there **communication** or **message passing** during this exercise (mark these steps)?

A: When we decided our ID numbers and shared out results with each other and graphed them.

Q: In what ways could this exercise have been **optimized** so it could take less time?
A: if we didn't have to communicate with eatch other.

Q: How could we have run this exercise using two instructors instead of one?
A: one instructor oversaw the collection of data, the other organized it in a meaningful way

Q: In what ways did we simulate a **parallel computer** in this exercise?
A: we simulated a parallel computer by having many people doing a task that all further a single cause

**Domain Decomposition**
[Model Website]

How to take a screenshot. In most cases you do not want/need to download additional software.
[Mac OS]
[Windows]

Image:

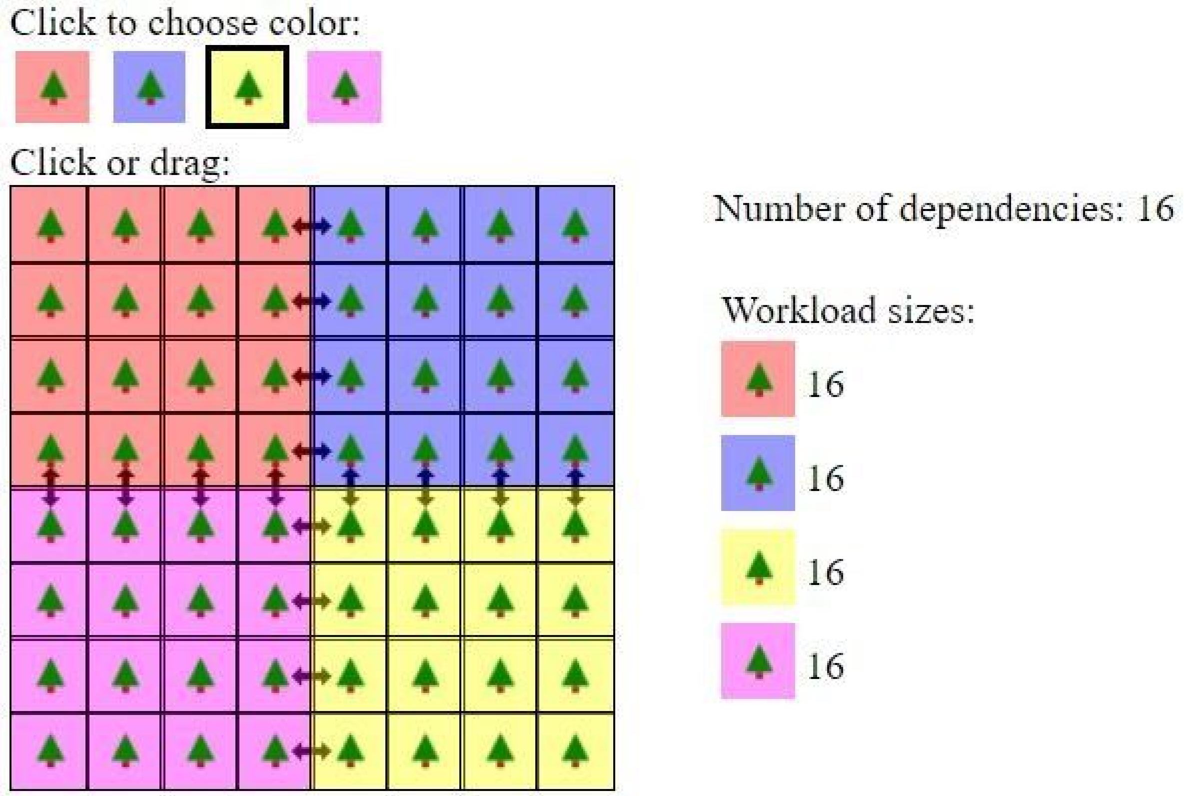


Evenly spread and divied out with low dependencies

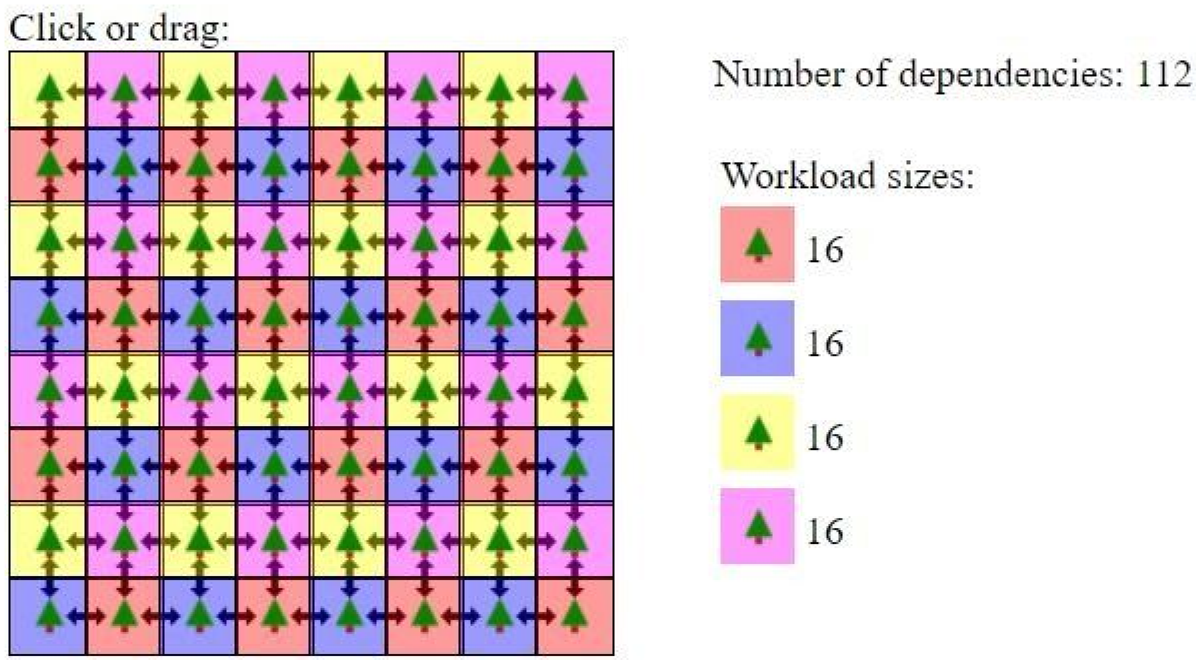


Evenly spread out but with very high dependencies

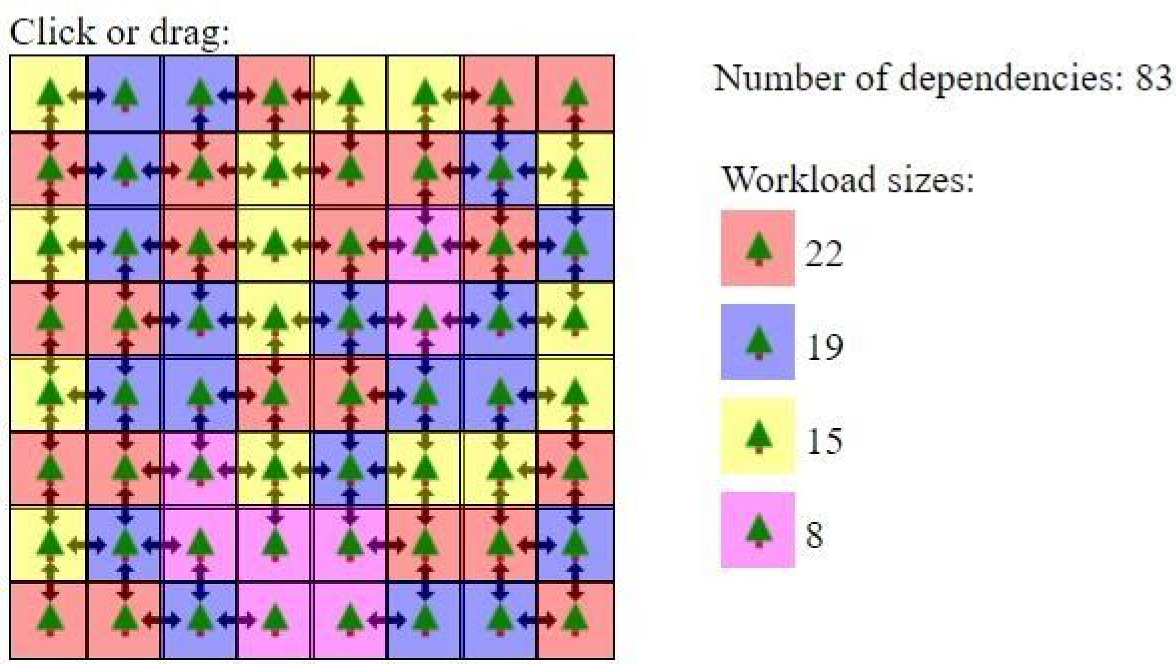


Random, no order, not even

Q: If we assumed each color is assigned to a researcher in a real forest, and each researcher is studying how a fire spreads through the forest, why do we call them **dependencies**? Why do we call them **workload sizes**?

A: we call them dependencies because that is the point where other people depend on another person. We call it workload sizes because that is the amount of work each person needs to do.

Q: If we assumed each color is assigned to a computer running a simulation for that part of the forest, why do we call them **dependencies**? Why do we call them **workload sizes**?
A: we call them dependencies because that is the point where other computers depend on another computer. We call it workload sizes because that is the amount of work each computer needs to do

Q: Why would we want to minimize the dependencies?
A: you would want to minimize dependencies because it will be faster for colors/researchers/computers just to work, not to talk

Q: What are some reasons we might want to give more work to one of the colors/researchers/computers?
A: some colors/researchers/computers may be able to do work faster than others