

Parallel Computing Workbook
Seokyoung Hong
April 2018

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: **Two time steps**

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

A: **One time step**

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

A: **Sixteen time steps**

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

A: **Eight time steps**

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: **Four time steps**

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: **Two time steps**

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 divides up the work so it takes less time steps to complete a task with parallel workers.**

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: **Four coins compared to two 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: **Eight coins compared to two 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: **Sixteen coins compared to two 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: **More work can be done with parallel workers in a fixed time frame.**

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: **Four coins can be held in memory compared to two 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: **Sixteen coins can be held in memory compared to two 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: **Parallel workers have a higher maximum memory to store data with.**

Human Parallel Computer - Data Parallelism through Forest Fire Simulations

My number: **9**

Total number of students: **17**

My probability: **52.9%**

Percentages: **41.86%, 70.93%, 0.34%, 0.34%, 61.93%**

Iteration counts: **21, 28, 1, 1, 20**

Average percentage: **35.08%**

Average # of iterations: **14.2**

Q: What were some of the **tasks** we did in this exercise? What were they, and who did them?

A: **Listening to the instructor, running a simulation based on your percentage, writing down the data of the percentage tree burned and number of iterations, calculating the averages of data collected. Students did them.**

Q: What kinds of **data** did we work with in this exercise?

A: **Percentages and number of iterations, average number of iterations, average percentage, ID number.**

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

A: **Handing the paper that we wrote on to the leader.**

Q: In what ways could this exercise have been **optimized** so it could take less time?

A: **Having more than one instructor, creating a shared doc where all the students could plug their data into.**

Q: How could we have run this exercise using two instructors instead of one?

A: **Set one instructor for even numbers, one instructor for odd numbers to graph the data. Or, one instructor makes one graph, and another instructor makes another graph.**


Q: In what ways did we simulate a **parallel computer** in this exercise?

A: **We had multiple workers working on the different tasks, so in the same amount of time we had more work done, and since that work was used to create a more accurate graph, our result was better. Also we learned how important communication is in parallel computing, but the downside is that the more communication there is the more processing power it takes and stuff.**


Domain Decomposition

[Model Website](#)


Image:




I call this the Reverse German Engineering, a model made to maximize the inefficiency of computing some work. The reason is that it contains so many dependencies with other computers that it will drive up the amount of communication needed and thus be less efficient.



Behold the Master of Efficiency, the model that requires the least amount of dependencies and thus being most efficient. Since each area contains an equal amount of workload while being clustered with the same type of computers to make minimal contact with others, efficiency is good.



Here is an example of a model that has very few dependencies, but not an even spread of the workload. This model will not be very efficient despite having fewer dependencies because only one type of computer has the majority of the tasks to complete. However, it might be! If one computer is faster than the others this will work out (okay, not *exactly* this kind of model, but something similar)



Okay I made this one because it's very symbolic and deep. The yellow square is ingrained in the negative balance of the photo, and symbolizes going against the pattern of oppressive society, and we as the viewers are to contemplate if the price of freedom is worth going against the grain and ruining the previous work of our ancestors. I'm selling it on Amazon for 20 grand.

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 they depend other researches to send and receive information in order to complete their work. We call them workload sizes because that's how much data each researcher has to collect over the area of the forest.

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 they depend other squares (or other computers) around them to compute work. Since they don't have the same memory or RAM, they have to send that information using a network cable or the ethernet. We call them workload sizes because that's how much work each computer has to complete.

Q: Why would we want to minimize the dependencies?

A: It lowers the amount of communication required to do work and I imagine that makes computing the work more efficient. And, if one computer type gets a computation wrong, more other types of computers will get that wrong information (better risk management).

Q: What are some reasons we might want to give more work to one of the colors/researchers/computers?

A: They might be better at doing a kind of specific work compared to other computers. Also, it will decrease the amount of dependencies needed. Finally, it could be that some computers are specialized - one type can work in a line, one can work in blocks, and it would be more efficient to lay out the work in a pattern.

Parallel Recipes

My serial recipe: How to efficiently consume a Big Mac® in 30 easy steps :)

Materials:

1. Big Mac
2. Fries
3. Soda
4. Ketchup

Instructions:

1. Spread materials on a paper tray with your hands
2. Carefully grasp and pick up Big Mac with both hands
3. Bring Big Mac closer to face
4. Open mouth
5. Insert Big Mac into open mouth
6. Bite down with teeth
7. Chew on the morsel until the morsels are small enough to
8. Swallow
9. Lower Big Mac
10. Scan tray for supplementing foods
11. Grab fry from fry container with dominant hand
12. Spread fry in ketchup with a dipping motion
13. Bring fry closer to mouth
14. Open mouth and bite down

15. Chew until morsels are small enough to
16. Swallow
17. Repeat steps 10-16 to your own discretion of hungriness.
18. Grab soda cup with dominant hand
19. Bring soda cup closer to your face
20. Put mouth on straw
21. Make sucking motion with mouth against the straw until an adequate amount of solution is displaced inside your mouth.
22. Swallow
23. Move soda cup a little farther apart with dominant hand
24. Rinse and repeat (quite literally) at your discretion of thirstiness.
25. Repeat steps 1-24 until entire meal is finished.
26. Wipe mouth with napkin
27. Use napkin to wipe grease off fingers
28. Gain weight and depression
29. Become insecure
30. Slowly spiral downhill into a self destructive pattern of overeating and insecurity until the point where nothing in your life can bring happiness as the first Big Mac that you have tasted and you live your life trying to chase that high, until the point where you are checked into a rehabilitative center against your will by your few remaining family members and you live the rest of your life strapped up into nutrient feeding tubes.

Dependencies: (Can't do one step unless you have done another step)

- Cannot eat before grasping Big Mac and bringing it closer to your face and opening your mouth
- Cannot swallow before you chew (I mean, you shouldn't).
- Cannot lower Big mac before you pick it up
- Cannot drink before grasping soda and making a sucking motion with your mouth
- Cannot clean up after eating before eating
- Cannot become insecure and gain weight without eating a ton of Big Macs

My parallel recipe - how to consume an infinite amount of Big Macs with infinite people :D

Materials:

1. Infinite Big Mac
2. Infinite amount of ketchup
3. Infinite amount of fries
4. Infinite amount of soda

Instructions:

1. Each person grabs a tray of a Big Mac meal
2. Line up all the people in a circle
3. Each person grasps Big Mac with both hands
4. Each person brings Big Mac closer to face

5. Each person opens their mouth
6. Each person inserts Big Mac into open mouth
7. Each person bites down with teeth
8. Each person chews on the morsel until the morsels are small enough to
9. Swallow
10. Each person lowers Big Mac
11. Each person hands remaining amount of Big Mac to the person to the right of them with their right hand
12. The person to the right of the person passing them the Big Mac grabs the patty with their left hand
13. **While the passing of the 'Mac is happening**, each person scans tray for supplementing foods
14. Each person grabs fry from fry container with their teeth
15. Each person spreads fry in ketchup with a dipping motion with their teeth
16. Each person bites down
17. Each person chews until morsels are small enough to
18. Swallow
19. Each person lowers passed Big Mac
20. Each person grabs soda cup with their dominant hand
21. Each person brings soda cup closer to their face
22. Each person puts mouth on straw
23. Each person makes a sucking motion with their mouth against the straw until an adequate amount of solution is displaced inside their mouth to
24. Swallow
25. Each person moves soda cup a little farther apart with dominant hand
26. Each person rinses and repeat (quite literally) at their discretion of thirstiness.
27. Each person repeats steps 1-21 until entire meal is finished in the circle, consuming over infinite Big Macs.
28. Each person wipes mouth with napkin
29. Each person uses napkin to wipe grease off fingers
30. Thankfully to the amount of people completing this recipe, they help each other as a support system before they spiral out of control and help watch each other's weight, and they all live fulfilling and happy lives

Q: In what ways was your **parallel** recipe different than your **serial** (non-parallel) recipe?

A: **To maximize efficiency, instead of having the unbearably long time to pick up a Big Mac and lower it to eat it, each person in the parallel recipe passes the big mac in an assembly line when consuming the supplementing foods. Additionally, in some steps, there is more going on due to the passing, the benefit of having multiple workers.**

Q: In what ways was your **parallel** recipe the same as your **serial** (non-parallel) recipe?

A: **Eating the supplementing foods is the same as the serial recipe because it is less efficient to pass around the fries and drinks, instead those steps are the same.**

Q: In what ways was your parallel recipe more efficient? In what ways was it less efficient?

A: **The ways my parallel recipe was more efficient was that the passing of the Big Macs reduced the amount of time between eating supplementing foods and the meal. The way is was less efficient is that an error in passing could set off the rest of the workers in the circle, causing a catastrophe.**

Q: Did anything need to change about the resources/materials/ingredients/tools in your recipe when you went from serial to parallel?

A: **Not really, just we needed more of them.**

Q: In your parallel recipe, what was the most number of workers you could keep busy at once?

A: **In my parallel recipe, all my workers, no matter how many there are at once, could be kept busy at once. But since that is based on an arbitrary work that is limited to how many big Macs that we have, that technically means only one worker is kept busy at once.**

Q: In what ways do you think this activity relates to computing and parallel computing?

A: **I tried to demonstrate in this activity the benefits of having infinite workers to increase the amount of work done in a set amount of time. In terms of parallel computing more computers are getting more work done (by having a greater quantity of computers working on the same task).**

***When trying to make a recipe a shorter amount of time by adding more workers, there is a limit to keeping each worker busy because the some tasks require other tasks to be completed before starting them.**

***A serial algorithm is not completely necessary for making a parallel algorithm.**

Careers in High Performance Computing

Career: Oil Industry

How HPC can be used in that career:

- Supercomputers are used to scan for oil and gas reservoirs deep below the surface of the Earth
- Oil industry used to do paper printouts of soundwaves sent underground to find resources, but with the supercomputer they can use exploring drones or workers with sensors attached to their heads to get data that the supercomputer can process. Can analyze area in weeks, in the past it took months and months.
- Another example - manages energy assets upstream production and mining. (Eni)

Sources:

- <https://gineersnow.com/industries/oil-gas/supercomputer-helps-company-search-oil-gas-deep>
- <https://www.bloomberg.com/news/articles/2018-01-18/eni-takes-on-total-and-bp-as-oil-majors-crank-up-computing-power>

The World's "Fastest" Supercomputers

Q: When was the most recent Top500 list published?

A: **November 2017**

Q: What is the name of the fastest supercomputer in the world according to the most recent list?

A: **Sunway Taihulight**

Q: Where is that supercomputer located?

A: **National Supercomputing Center (Wuxi, China)**

Q: How many **cores** does it have?

A: **10,649,600**

Q: How much **peak performance (RPEAK)** does it have?

A: **125,435.9 TFlop/s**

Q: How many of the Top 500 sites in the top 10 are located in the United States?

A: **Four**

Q: If the **Blue Waters** supercomputer was capable of a **peak performance** of **13,000 TFLOP/S** when it came on-line in 2012, where would it be listed in the November 2012 list?

A: **#3**

Q: Why doesn't Blue Waters appear on that list?

A: **Blue Waters dropped out of the competition because they wanted to be computing practical things for science rather than using them to work on trivial problems for a competition.**

Source: https://www.hpcwire.com/2012/11/16/blue_waters_opts_out_of_top500/

Q: What are **cores**?

A: **A processing unit that receives instructions and performs calculations**

Q: What does **TFLOP/S** stand for?

A: **A trillion floating point operations per second**

Q: What does **Linpack** measure?

A: **Linpack measures a system's floating point computational power.**

Q: What would be some different ways to rank supercomputers?

A: **R max (Tflop/s), number of cores, power efficiency, or subjective things like partition in scientific journals, usefulness, etc.**

LittleFe

Where the name comes from: **Play on "big iron," which is slang for supercomputers.**

Components:

- CPU
 - Cores
 - GPGPU
- Heat Sink
- RAM (Random Access Memory)
- ROM (Read Only Memory)
- Hard Drive
- Network
 - Firewire Cable

- Ethernet Cable, Switch, or Router
- Infiniband
- Power Supply

Blue Waters demo

[YouTube video](#)

Q: What are the advantages to using a remote supercomputer as compared to a local supercomputer like LittleFe?

A: **It's a lot more times powerful, obviously. A lot more people can access it to perform computations at the same time from anywhere in the world. If you sign up for a remote supercomputer like Blue Waters, you have a lot more resources to access, like experts or staff that have their careers dedicated to high performance computing - FOR FREE!**

Q: What are the disadvantages?

A: **Requires a lot of space and power, and an intricate cooling system. If you want to divvy up the power with lots of people you have to be willing to share. Additionally, with a local supercomputer, you have the compute whatever your little heart desires- you have the freedom and don't have to fill out lots of proposals and forms. Finally, you don't need an Internet connection to access LittleFe.**

Parallel Computing: Terminology and Examples

[Slides](#)

- **Serial** - instructions are executed one at a time
- **Parallel** - multiple instructions are completed in parallel
- **Core** - entity of CPU that executes directions
- **Shared memory** - multiple cores can share RAM
- **Distributed memory** - cores don't share RAM
- **Node** - a grouping of cores and their shared memory
- **Cluster** - a grouping of nodes and the network that connects them
- **Supercomputer** - a really big fast cluster