

Parallel Computing Notebook
Jackson Pausback
April 2017

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 is a lot more time efficient to have multiple people doing it at once.

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: Parallel workers flipped 4, the Serial worker flipped 2.

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: The Parallel workers flipped 8, the Serial worker flipped 2.

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: The Parallel workers flipped 1, the Serial worker flipped 2.

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 more work done in the same amount of time.

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: The Parallel workers flipped 4, the Serial worker flipped 2.

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: Parallel workers flipped 16, the Serial Worker flipped 2.

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: Even if you can only flip a few coins and mark down the data for a few coins, it is still more efficient with parallel workers.

Human Parallel Computer - Data Parallelism through Forest Fire Simulations

My number: 5

Total number of students: 22

My probability: 0.227

Percentages: 4.84, 0.34, 0.34, 2.07, 1.38

Iteration counts: 9, 1, 1, 5, 3

Average percentage: 1.8

Average # of iterations: 3.8

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

A: We divided up numbers, which Aaron and the workers did, we gathered our data which the workers did, and we returned our data to Aaron which everyone did.

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

A: Probability, Percentage of burnt, and Iteration.

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

A: When we got out numbers, and when we gave back our averages.

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

A: Less discussion, already have set out ways to give out papers, etc. and to have more computers.

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

A: Split the room in half, and have one instructor teaching each side, or if we wanted more data we could get twice as many computers and use the second instructor for the new wave.


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

A: We had many computers doing the same task with different inputs to gather data.

Domain Decomposition

[Model Website](#)

Click or drag:



Number of dependencies: 16

Workload sizes:



Image: ^, I chose this Image because it is the least possible amount of dependencies possible.

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: It is called dependencies because they depend on the trees around them to tell them when there is a fire. Workload sizes is how much each color has 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: They depend on each other to get the data and communicate it correctly, and the amount of work is called workload sizes because that is how much work each computer has to do.

Q: Why would we want to minimize the dependencies?

A: The more dependencies the more communication which means the longer the task will take to finish.

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

A: Maybe that computer is the most powerful or has the most space.

Parallel Recipes

My serial recipe: How to cook an oven pizza

Materials: Oven, Pizza, Oven dish, Plate, pizza cutter, oven mitt

Instructions: Shrek, Donkey, Fionna, Farquad, Gingerbread man

1. Get your Pizza
2. Unwrap the Pizza
3. Put it on a dish that can go in the oven
4. Go to your oven
5. Press Bake, 400 degrees

6. Wait until the oven has been heated to 400 degrees
7. Pick up your pizza
8. Open the stove
9. Put the pizza in the oven
10. Close the oven door
11. Put on a timer for 19 minutes
12. Wait 19 minutes
13. Get an oven mitt from the place you keep your oven mitt
14. Open the stove
15. Take the pizza out with the oven mitt
16. Set it down somewhere where it won't burn anything
17. Get a circular pizza cutter from the place you keep your pizza cutter
18. Cut the pizza
19. Transfer the slices from the oven dish to a plate
20. Enjoy

Dependencies:

- You have the materials necessary for the process. You can eat all the materials in the Pizza if you want to not get sick/death.

My parallel recipe

Materials: Oven, Pizza, Oven dish, Plate, pizza cutter, oven mitt, shrek, gingerbread, fionna, donkey

Instructions:

1. Gingerbread man stand by the oven
2. Shrek get the pizza
3. Shrek unwrap the pizza
4. Gingerbread man press bake on the oven
5. Gingerbread man press 400 degrees and start
6. Fionna get the oven dish
7. Once the oven is 400 degrees, donkey get the pizza and give it to fionna
8. Fionna put the pizza on the oven dish
9. Fionna put the pizza in the oven
10. Gingerbread man set a timer for 19 mins
11. Wait 19 mins
12. Shrek get an oven mitt
13. Fionna put the oven mitt on
14. Fionna take the pizza out
15. Put the oven dish somewhere where it won't burn anything
16. Donkey get the plate
17. Shrek transfer the slices from the oven dish to the plate
18. Donkey get the pizza cutter

19. Fionna cut the pizza
20. Shrek get the plate and bring it to the table.

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

A: It had a lot more people working, some people had specific jobs (gingerbread man worked the oven).

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

A: Basically the exact same steps just different people doing it, same result too.

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

A: The Parallel would take less time because of the more people. It would be less efficient in this task because at sometimes 4 people is too many and 2 or 3 would be fine and less crowded.

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

A: No

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

A: You can use serial or parallel methods in both tasks, and in both the parallel is more efficient.

Going Shopping

Solutions:

1. Everyone gets in the Van and they go to the store. They all stay together while in the store. They get a big lunch they can share, like a loaf of bread, a pack of cheese, lettuce, and turkey.
2. Everyone gets in the Van and they go to the store. Half of the kids go with 1 adult, and the other goes with the other half. Everyone gets their own lunchables.
3. Call and get good delivered.
4. Everyone go to the store in the van, kids go in partners and get their food and they meet up in check-out.

What if's:

1. Money is low.
2. Someone is allergic to something.
3. What if the car gets in an accident.
4. What if the kid falls out of the car.
5. What if the store is closed.
6. What if the store doesn't have food.

Q: Where is there inherent **parallelism** in your solutions?

A: In the fourth solution each 2 person team does the same thing to speed up the process.

Q: Where are there **dependencies** between tasks in your solution?

A: The kids are relying on Aaron and Jim to get them food , and Aaron and Jim are relying on the food to be at the store.

Q: Where is there **communication** in your solution?

A: The kids tell Jim and Aaron they don't have lunch. Aaron tells everyone to get in the van, the kids say what they want and they go and get it.

Parallelism in Nature

1. Model link: <http://www.shodor.org/interactivate/activities/RabbitsAndWolves/>

Data: The location and movement of the Rabbits and Wolves. The grass/food and the speed of regrowth. Also initial values of agents.

Tasks: Some of the tasks would be the eating, the wolves eating rabbits and the rabbits eating the grass. The grass regrowing is another task.

Parallelizable data: One computer could control all of an agent's initial values.

Parallelizable tasks: One computer could control the reproduction of the agents, and another on the function of eating. Each computer could control a singular agent's movement.

Q: What patterns do you notice in the types of data and tasks that can be parallel?

A: They are usually the most necessary functions to the model.

Q: What patterns do you notice in the types of data and tasks that cannot be parallel?

A: The types of data and tasks that can't be Parallelized are things that are not applied a lot in the model such as a counter.

Careers in High Performance Computing

Career: High Performance Computing Software Engineer

How HPC can be used in that career:

Sources:

- To be a HPC Software Engineer you have to have high level computing expertise, many types of jobs can use HPC Software Engineer such as colleges, an example I found was at Princeton University they were looking for a HPC Software Engineer for their Astrophysics department. This shows how many possible jobs there are using HPC.

The World's "Fastest" Supercomputers

Q: When was the most recent Top500 list published?

A: November, 2016

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: 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

Q: How many of the Top500 sites in the top 10 are located in the United States?

A: 5

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: The people who made it chose for it not to be.

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

Q: What are **cores**?

A: Cores are a unit used for processing instructions, they receive instructions and then calculate or do actions based on those instructions,

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

A: teraflop

Q: What does **Linpack** measure?

A: the floating-point rate of a computer

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

A: Cost-efficiency or looks.

LittleFe

Where the name comes from: Little Iron

Components:

- Aluminum frame kit (custom item from LittleFe/BCCD team)1Low-voltage wiring harness (custom item from LittleFe/BCCD team)1[ASRock IMB-151D Bay Trail Thin Mini-ITX Motherboard with J1900 Celeron](#)6[Transcend SO-DIMM DDR3L Low Voltage 1600 Memory 4GB](#)6[Solidtek KB-5010BU 13.31" Mini Keyboard with Trackball Mouse](#)1[New A Pair of Laptop Notebook Wifi Bluetooth WWAN Broadband Internal Antenna](#)1[Intel 6200 Dell Intel 6200 IEEE 2.4GHZ-5GHZ 802.11n Wi-Fi Adapter - Half Size Mini PCI Express Card - 300Mbps](#)1[Anker® Unibody Aluminum USB 3.0 to RJ45 Gigabit Ethernet Adapter Supporting 10/100/1000 Mbps Ethernet \[RTL8153 Chipset\]](#)1[Seagate Laptop Thin 500GB 7200RPM SATA 6 GB/s 32 MB Cache 2.5 Inch Hard Disk Drive \(ST500LM021\)](#)1[TopLink 8-Port Gigabit Ethernet Desktop Switch - 10/100/1000 Mbps](#)1[Pelican 1610](#)1[Meanwell PB-360P-12](#)1[1.5ft Cat6 Ethernet Cable - 550MHz UTP Bare Copper, Red](#)5[1.5ft Cat6 Ethernet Cable - 550MHz UTP Bare Copper, Gray](#)1[TeckNet® Basic 3-Button 3D optical USB Wired Mouse](#)1[C2G / Cables to Go 03134 18 AWG Universal Power Cord for NEMA 5-15P to IEC320C13, Black \(10 Feet/3.04 Meters\)](#)1[10inch SATA 6Gbps Cable w/Locking Latch - Neon Blue](#)1Female 15-pin SATA power to female 4-pin fdd power connector1

Blue Waters demo

[YouTube video](#)

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

A: The remote supercomputers can be much larger because you don't have to be in the same room such as LittleFe. Also like BlueWater, you can access it from across the country if you have the passcode.

Q: What are the disadvantages?

A: If you are accessing it across the country, it has to communicate information back and forth which takes time. Also they are usually much more expensive than something like LittleFe.

Parallel Computing: Terminology and Examples

[Slides](#)

Shopping for Your Own Supercomputer

- Part: [Corsair Carbide Series CC-9011070-WW Black Steel Micro ATX / Mini-ITX Air 240 High Airflow MicroATX and Mini-ITX PC Case ...](#)
Cost: 85 x 1
Link: <https://www.newegg.com/Product/Product.aspx?Item=N82E16811139044&ignorebr=1>
- Part: [WD Black 5TB Performance Desktop Hard Disk Drive - 7200 RPM SATA 6Gb/s 128MB Cache 3.5 Inch - WD5001FZWX](#)
Cost: 220 x 1
Link: <https://www.newegg.com/Product/Product.aspx?Item=N82E16822236971&ignorebr=1>
- Part: [NETGEAR ProSAFE 5-Port Gigabit Ethernet Switch \(GS205\)](#)
Cost: 15 x 1
Link: <https://www.newegg.com/Product/Product.aspx?Item=N82E16833122608&ignorebr=1>
- Part: [EVGA SuperNOVA 1000 G2 120-G2-1000-XR 80+ GOLD 1000W Fully Modular Includes FREE Power On Self Tester Power Supply](#)
Cost: 140 x 1
Link: <https://secure.newegg.com/Shopping/ShoppingCart.aspx?Submit=view>
- Part: [G.SKILL 6GB \(3 x 2GB\) 240-Pin DDR3 SDRAM DDR3 1600 \(PC3 12800\) Triple Channel Kit Desktop Memory Model ...](#)
Cost: 50 x 5
Link: <https://www.newegg.com/Product/Product.aspx?Item=N82E16820231225&ignorebr=1>
- Part: [MSI H110M Gaming LGA 1151 Intel H110 HDMI SATA 6Gb/s USB 3.1 Micro ATX Intel Motherboard](#)
Cost: 67 x 5
Link: <https://www.newegg.com/Product/Product.aspx?Item=N82E16813130918&ignorebr=1>
- Part: [AMD FX-9590 Vishera 8-Core 4.7 GHz Socket AM3+ 220W FD9590FHHKWO Desktop Processor - Black Edition](#)

Cost:160 x 5

Link:<https://www.newegg.com/Product/Product.aspx?Item=N82E16819113347&ignorebr=1>

- Part:[Cooler Master Hyper T2 - Compact CPU Cooler with Dual Looped Direct Contact Heatpipes](#)

Cost:17 x 5

Link:<https://www.newegg.com/Product/Product.aspx?Item=N82E16835103107&ignorebr=1>

- Part:[BYTECC C6EB-3B 3 ft. Cat 6 Blue Enhanced 550MHz Patch Cables](#)

Cost:3 x 5

Link:<https://www.newegg.com/Product/Product.aspx?Item=N82E16812270013&ignorebr=1>

Q: What are the most important parts of a supercomputer?

A: Motherboard, CPU, and CPU Cooler in my opinion.

Q: What assumptions did you make when you did your shopping? Which of these assumptions were false?

A: Cheaper parts are the best way to go, but I realized a lot of the great deals are probably broken parts or not even the same part.