Parallel Computing Workbook
Rohan Kashyap
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: 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: Parallel workers can do the same amount of work in less time.
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 and one serial worker can only 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 and one serial worker can only 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: 8 parallel workers can flip 16 coins and one serial worker can only 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: The parallel workers can do 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: 2 parallel workers can flip 4 coins and 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 and one serial worker can only 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: Parallel workers are always able to keep track of more work.

Human Parallel Computer - Data Parallelism through Forest Fire Simulations
My number: 12
Total number of students: 17
My probability: 0.706
Percentages: 98.26, 99.65, 97.92, 95.84, 98.96
Iteration counts: 18, 20, 19, 26, 19
Average percentage: 98.126
Average # of iterations: 20.4
Q: What were some of the **tasks** we did in this exercise? What were they, and who did them?
A: We ran a model and recorded it, took averages, listened to instructions, etc.
Q: What kinds of **data** did we work with in this exercise?
A:The # of iterations and the percentage of burning
Q: In which steps was there **communication** or **message passing** during this exercise (mark these steps)?
A: Telling Aaron our averages, listening to instructions
Q: In what ways could this exercise have been **optimized** so it could take less time?
A: We could have input our averages directly into the graph
Q: How could we have run this exercise using two instructors instead of one?
A: One instructor graphed the iterations, other graphed percentages
Q: In what ways did we simulate a **parallel computer** in this exercise?
A: We were the parallel computer, because we worked together to create the final graph by each handling a different part of the problem

Domain Decomposition
Model Website

Image:

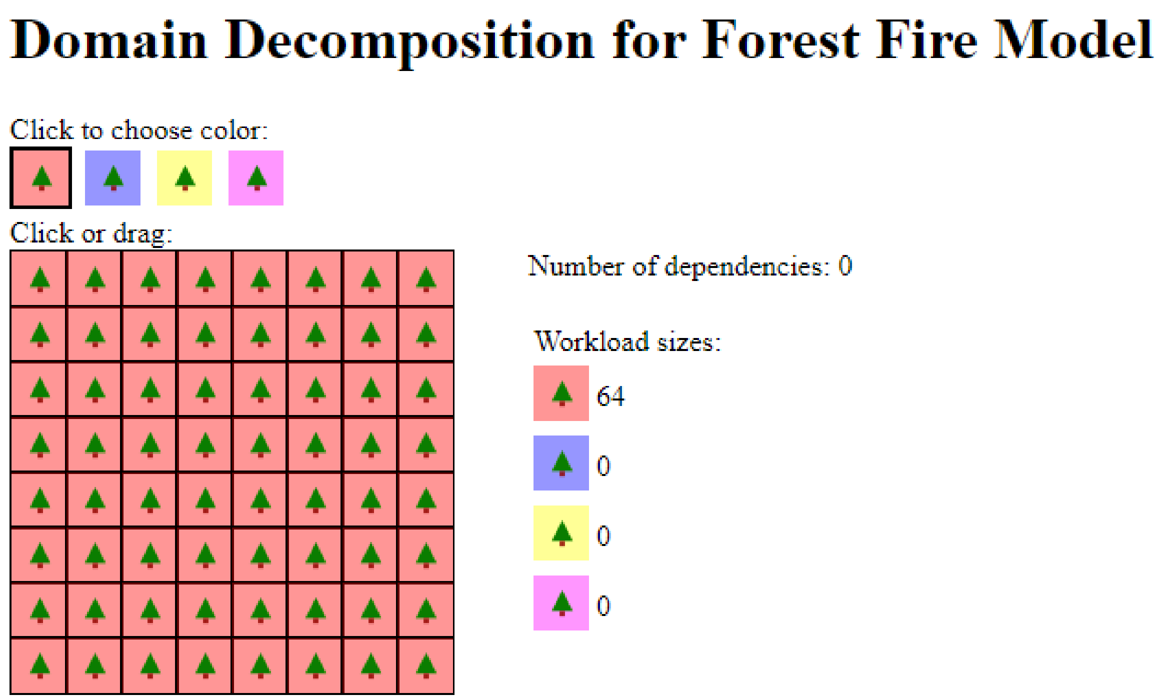

This is the most basic model possible. Here, there are no colors other than red, so there are no dependencies. While this may seem like it is fast and efficient, since one person/computer is doing this, it will be extremely slow. This is one of the two ways to make the most inefficient model.

Domain Decomposition for Forest Fire Model

Click to choose color:

Click or drag:

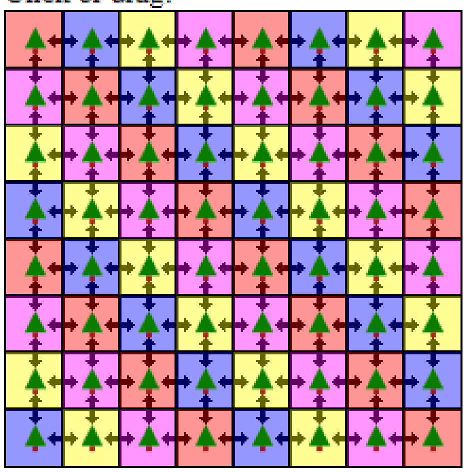

Number of dependencies: 112

Workload sizes:

16

16

16

16

This is the other most inefficient way to do the work. All the computers/people are doing the same amount of work, but there are so many dependencies, 112, that everything is very slow. This would take even longer than the previous image!

Domain Decomposition for Forest Fire Model

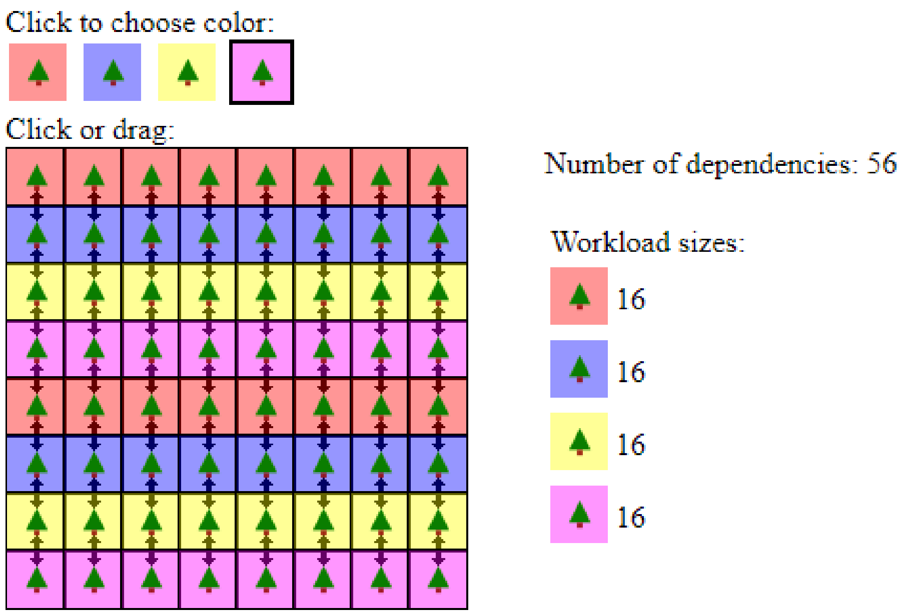

This model is much better than the other two models. The computers/people have the same amount of work, and the dependencies are low enough that the communication is manageable. It could still be improved by putting like colors together, but this would still be relatively efficient.

Domain Decomposition for Forest Fire Model

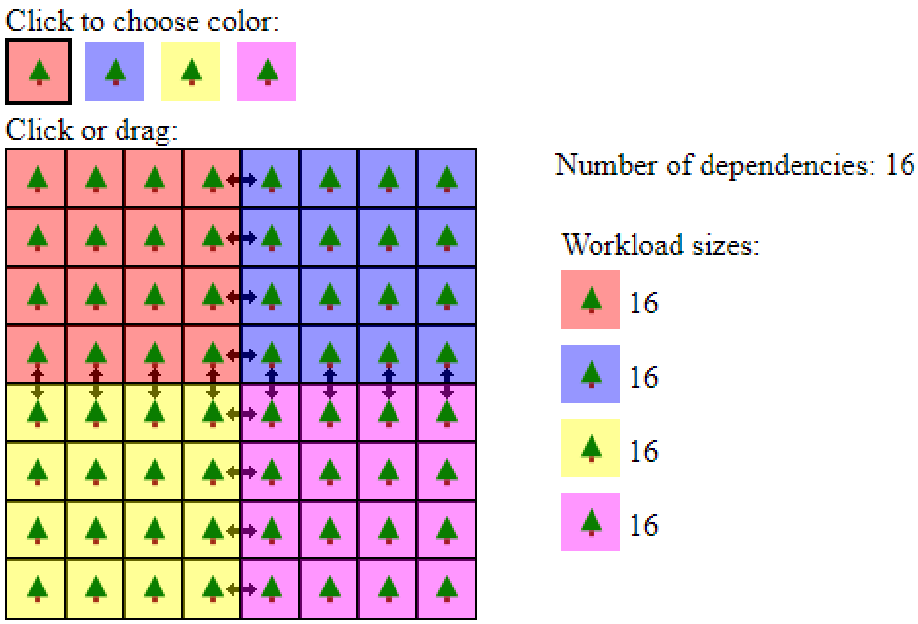

This is the best possible model. Each computer/person has the same amount of work, like the two models before it, but it only has 16 dependencies! Because the like computers/people are put together, they form four giant 'supercomputers/superpeople.' That causes the dependencies to be between the supercomputers/superpeople. This is the most efficient and fastest way to model the trees.

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: They are called dependencies because they depend on the other researchers for information. The workload sizes term refers to the amount or size of work each computer 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: Like the researchers, the computers do not have the ability, or in this case the capacity, to do all the work by themselves, so when they communicate, they are dependant on each other. The amount of work they can do is called the workload size.

Q: Why would we want to minimize the dependencies?

A: The less communication and dependencies, the computers/researchers would work faster and more efficiently.

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

A: They might be more powerful or able to do more work in general. They may also be specialized in that field, so they would be more efficient.

<u>Parallel Recipes</u>

My serial recipe: How to unpack/prepare a violin to be played

Materials:

1. Violin
2. Violin Case
3. Bow
4. Rosin
5. Stand
6. Music
7. Shoulder Rest
8. Violin Tuner

Instructions:

1. Get violin case
2. Open case
3. Remove violin from case
4. Remove shoulder rest from case
5. Turn violin upside down
6. Put shoulder rest on bottom of violin
7. Flip violin to the front
8. Get violin tuner
9. Set violin tuner to ‘A’ string
10. Pluck ‘A’ string (third from left)
11. Turn knob below string
12. Repeat steps 10-11 until violin string sounds like note coming from tuner
13. Set violin tuner to ‘D’ string
14. Pluck ‘D’ string (second from left)
15. Turn knob below string
16. Repeat steps 14-15 until violin string sounds like note coming from tuner
17. Set violin tuner to ‘G’ string
18. Pluck ‘G’ string (first from left)
19. Turn knob below string
20. Repeat steps 18-19 until violin string sounds like note coming from tuner
21. Set violin tuner to ‘E’ string
22. Pluck ‘E’ string (fourth from left)
23. Turn knob below string
24. Repeat steps 22-23 until violin string sounds like note coming from tuner
25. Set down violin carefully
26. Remove bow from case
27. Turn screw at bottom of bow until bow hairs are as taut as desired
28. Set down bow carefully
29. Remove rosin from case
30. Open rosin

31. Pick up bow
32. Put rosin on bow hairs
33. Rub rosin up and down bow hairs until desired level of rosin is reached
34. Close rosin
35. Put rosin in case
36. Set down bow carefully
37. Get stand
38. Unfold stand
39. Raise stand to desired height
40. Set down stand
41. Remove music from case
42. Open music to desired song
43. Place music on stand
44. Pick up violin and bow
45. Put violin on left shoulder, pointing left
46. Hold end of bow in right hand

Dependencies:
- Must pick up and set down violin and bow before doing anything
- Must tune strings, rosin bow, and set up stand/music separately
- Must place, remove, and set down everything in order before doing other things

My parallel recipe: How to unpack/prepare a violin to be played

Materials:
1. Violin
2. Violin Case
3. Bow
4. Rosin
5. Stand
6. Music
7. Shoulder Rest
8. Violin Tuner

Instructions:
1. P1: Get violin case
2. P1: Open case
3. P1: Remove violin from case
4. P1: Remove shoulder rest from case
5. P1: Turn violin upside down
6. P1: Put shoulder rest on bottom of violin
7. P1: Flip violin to the front
8. P1: Get violin tuner
9. P1: Set violin tuner to ‘A’ string
10. P1: Pluck ‘A’ string (third from left)
11. P1: Turn knob below string

12. P1: Repeat steps 10-11 until violin string sounds like note coming from tuner
13. P1: Set violin tuner to ‘D’ string
14. P1: Pluck ‘D’ string (second from left)
15. P1: Turn knob below string
16. P1: Repeat steps 14-15 until violin string sounds like note coming from tuner
17. P1: Set violin tuner to ‘G’ string
18. P1: Pluck ‘G’ string (first from left)
19. P1: Turn knob below string
20. P1: Repeat steps 18-19 until violin string sounds like note coming from tuner
21. P1: Set violin tuner to ‘E’ string
22. P1: Pluck ‘E’ string (fourth from left)
23. P1: Turn knob below string
24. P1: Repeat steps 22-23 until violin string sounds like note coming from tuner
25. P2: At the same time-
26. P2: Remove bow from case
27. P2: Turn screw at bottom of bow until bow hairs are as taut as desired
28. P3: Remove rosin from case
29. P3: Open rosin
30. P3: Give rosin to P2
31. P2: Put rosin on bow hairs
32. P2: Rub rosin up and down bow hairs until desired level of rosin is reached
33. P2: Close rosin
34. P2: Put rosin in case
35. P3: At the same time-
36. P3: Get stand
37. P3: Unfold stand
38. P3: Raise stand to desired height
39. P3: Set down stand
40. P3: Remove music from case
41. P3: Open music to desired song
42. P3: Place music on stand
43. P2: Give bow to P1
44. P1: Put violin on left shoulder, pointing left
45. P1: Hold end of bow in right hand

<u>Q: In what ways was your **parallel** recipe different than your **serial** (non-parallel) recipe?</u>
A:The three people tuned the violin, rosined the bow, and set up the stand/music at the same time.
<u>Q: In what ways was your **parallel** recipe the same as your **serial** (non-parallel) recipe?</u>
A: The people still had to follow the same steps to do all those things.
<u>Q: In what ways was your parallel recipe more efficient? In what ways was it less efficient?</u>
A:It was more efficient because the three people did the three main tasks at once, but there was some dependencies, like P3 giving P2 the rosin.

Q: Did anything need to change about the resources/materials/ingredients/tools in your recipe when you went from serial to parallel?
A: No I did not. Because of the end goal, there was no need to add more of anything.
Q: In your parallel recipe, what was the most number of workers you could keep busy at once?
A: 3 people, as that was how many main tasks there were.
Q: In what ways do you think this activity relates to computing and parallel computing?
A: The people work together to solve the problem faster and more efficiently, than came together at the end, just like in a multicore computer chip.

Careers in High Performance Computing

Career: Medicine
How HPC can be used in that career:

- Image guided radiation therapy
- Drug development
- Simulating trials
- Visualization of processes
- Taking CAT scans, x rays, etc of the human body
- MRI imaging
- Microsurgery robots
- Controlling life support for those in critical condition

Sources:

- http://medicalphysicsweb.org/cws/article/opinion/34995
- https://insidehpc.com/2015/06/hpc-in-medical-applications/

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: China
Q: How many **cores** does it have?
A: 10,649,000
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: 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: Third
Q: Why doesn't Blue Waters appear on that list?
A: It wants to focus on memory and petascale computing rather than overall power. They want to focus on practical and difficult real-world problems.

Source: https://www.hpcwire.com/2012/11/16/blue_waters_opts_out_of_top500/
Q: What are **cores**?
A: A computing unit in a computer that can either work by itself, or work in parallel with other cores.
Q: What does **TFLOP/S** stand for?
A: Teraflop
Q: What does **Linpack** measure?
A: The amount of time needed to solve a dense number of linear equations
Q: What would be some different ways to rank supercomputers?
A: Cores, Rmax, Rpeak, Power Consumption, Use

LittleFe
Where the name comes from: Little Iron (Slang for supercomputers used to be Big Iron)
Components:
- Motherboard
- CPU
- GPGPU
- Heat Sink
- RAM
- ROM
- Storage (HDD)
- Network (Cables & Routers)
- 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: The computer is most likely more powerful and perhaps has more power than a local supercomputer. It also does not cost money (probably)
Q: What are the disadvantages?
A: Not everyone is able to use it, difficult to access. With local supercomputers, you can also configure it however you want for whatever you want to do.

Parallel Computing: Terminology and Examples
Slides