


## Class 37: How to Find Aliens (and Factors)



CS150: Computer Science  
University of Virginia  
Computer Science

David Evans  
<http://www.cs.virginia.edu/~evans>

CS150 Fall 2008: Lecture 27: Finding Aliens

2 Computer Science

## Finding Aliens



Arecibo Antenna, Puerto Rico  
From [http://www.acl.gatech.edu/research/07/finding\\_aliens.html](http://www.acl.gatech.edu/research/07/finding_aliens.html)

CS150 Fall 2008: Lecture 27: Finding Aliens

2 Computer Science

## Finding the Aliens

```
for signal in signals:
    power = findPowerSpectrum (signal)
    if (isAlien (power)):
        print "Found an alien!" + signal
```

CS150 Fall 2008: Lecture 27: Finding Aliens

3 Computer Science

## Processing Signals

- Power spectrum
- Find patterns in signal
- Eliminate natural and human-made signals
- Today:
  - BlueGene, 280Tflops/s
  - ~\$200 per Gigaflop
- No success finding aliens



Cray T3E-1200E (~1998)  
1 Teraflop/s  
• Trillion floating point operations/sec

CS150 Fall 2008: Lecture 27: Finding Aliens

4 Computer Science

## Finding the Aliens Cheaper

```
parfor signal in signals:
    power = findPowerSpectrum (signal)
    if (isAlien (power)):
        print "Found an alien!" + signal
```

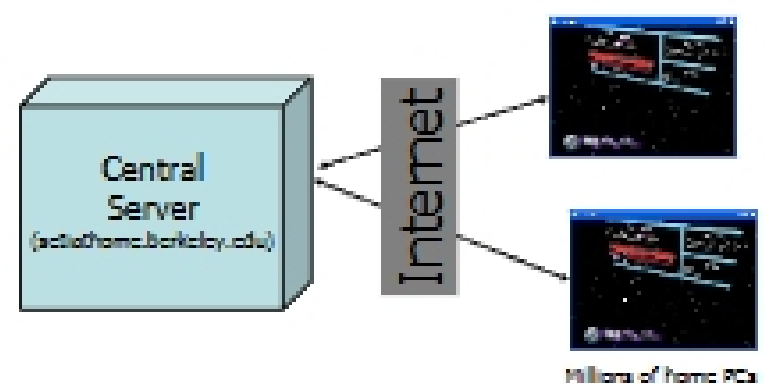
Parallel for: Instead of doing each element sequentially in order, we can do each element in parallel on a different machine.

Note: python does not actually have parfor, but other languages do.

CS150 Fall 2008: Lecture 27: Finding Aliens

5 Computer Science


## Public Distributed Computing



CS150 Fall 2008: Lecture 27: Finding Aliens


6 Computer Science

## BOINC (Seti@Home) Client



CS180 Fall 2008, Lecture 27: Finding Aliens

## BOINC Manager



CS180 Fall 2008, Lecture 27: Finding Aliens

## Incentives are Necessary

Rank	Name	Worked on (hours)	CPU hours	Worked on (days)
1	...	200,000	10,000,000	...
2	...	150,000	7,500,000	...
3	...	100,000	5,000,000	...
4	...	75,000	3,750,000	...
5	...	50,000	2,500,000	...
6	...	25,000	1,250,000	...
7	...	10,000	500,000	...
8	...	5,000	250,000	...

CS180 Fall 2008, Lecture 27: Finding Aliens

## Incentives are Dangerous

- People will cheat\*
- How to cheat?
  - Respond with the “there are no aliens” message without actually doing all the work
- Chances of getting caught?
  - 0 (Assumes all jobs have no aliens. So far this is true.)

\* Only applies in real world, not of UVA.

CS180 Fall 2008, Lecture 27: Finding Aliens

## Preventing Cheaters

- Send the same job to multiple workers
  - Wastes computing
  - What if the cheater controls many machines?
- Instead of response being “no aliens” make clients send a response that proves they did computation
- Sometimes send out fake jobs that have aliens in them
  - Clients must find the fake aliens
  - Need to make sure the fake jobs look just like real jobs
  - (Airport security scanners work like this also)

Doug Scajda and colleagues at University of Richmond work on these problems (see link on notes)

CS180 Fall 2008, Lecture 27: Finding Aliens

## Why is finding aliens so “easy”?

Note: we haven't yet found any aliens, but it's easy to set up the computation...

- Can be broken into many tasks
- Each task can be done completely independently
  - No shared memory between the tasks
- The data to describe a task and response is small compared to the computing
  - SETI@home jobs are 350KB data download, 1KB upload, 3.9 Trillion operations (several hours on PC)

CS180 Fall 2008, Lecture 27: Finding Aliens

## "Harder" Task


- Dividing your PS8 project work among your team

PS8 Update: you have 1 week left!

If you do not have basic functionality working by tomorrow, arrange to meet with me

## Scheduling Meetings

Alice wants to schedule a meeting with Bob and Colleen

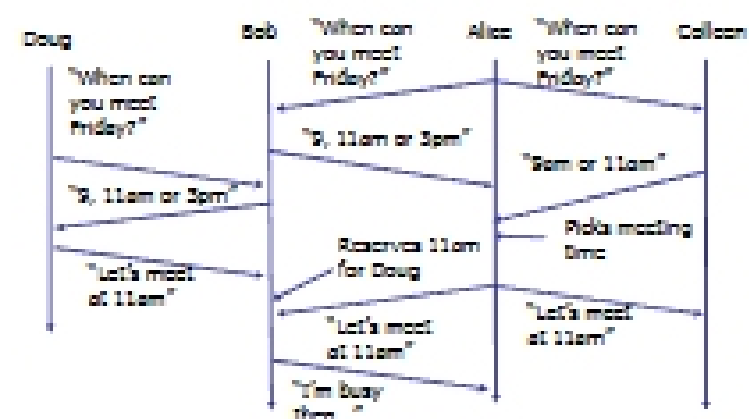


## Partial Ordering of Events

- Sequential programs give us a *total ordering* of events: everything happens in a determined order
- Concurrency gives us a *partial ordering* of events: we know some things happen before other things, but not total order

Alice asks to schedule meeting before Bob replies  
 Alice asks to schedule meeting before Colleen replies  
 Bob and Colleen both reply before Alice picks meeting time  
 Alice picks meeting time before Bob reserves time on calendar

## Race Condition



## Preventing Race Conditions

- Use locks to impose ordering constraints
- After responding to Alice, Bob reserves all the times in his response until he hears back (and then frees the other times)

## Locking

