

CS61A Lecture 30

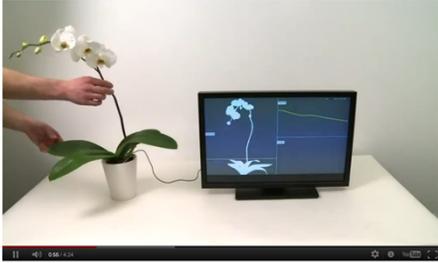
MapReduce

Jom Magrotker
UC Berkeley EECS
August 8, 2012



COMPUTER SCIENCE IN THE NEWS

Touch Your Philodendron and Control Your Computer: Technology Turns Any Plant Into an Interactive Device



<http://www.sciencedaily.com/releases/2012/08/120808094053.htm>
http://www.dorseyresearch.com/research/projects/ht_botanica_dsp.htm



TODAY

- MapReduce
 - Description
 - Examples
- Social implications of Computing
 - Therac-25 case study



RECAP: REDUCE

```
reduce(lambda so_far, elt: so_far * elt,
        [1, 3, 8, 7], 2)

((((2 * 1) * 3) * 8) * 7)
```



MAP AND REDUCE

Many problems can be solved by mapping a function over a sequence, and then reducing the new items.

To find the sum of squares of items in a list:

```
from functools import reduce
from operator import add
def sum_squares(list):
    return
    reduce(add, ← 2. ... and add the squared items together.
           map(lambda num: num ** 2, ← 1. Square each item in a list...
               list),
           0)
```



MAP AND REDUCE: EXAMPLE

Count the number of letters in the words of a list

```
def count_letters(wordlist):
    return \
    reduce(add, ← 2. ... and add the lengths together.
           map(lambda word: len(word), ← 1. Find the length of each word in a list...
               wordlist),
           0)
```



MAP AND REDUCE: EXAMPLE

Count the number of words in a list

```
def count_words(wordlist):
    return \
        reduce(___,
               map(_____,
                   wordlist),
               0)
```



7 Cal

MAP AND REDUCE: EXAMPLE

Count the number of words in a list

```
def count_words(wordlist):
    return \
        reduce(add,
               map(lambda word: 1,
                   wordlist),
               0)
```



8 Cal

MAP AND REDUCE: EXAMPLE

Count the number of words in a list of even length

```
def count_words(wordlist):
    return \
        reduce(___,
               map(_____,
                   wordlist),
               0)
```



9 Cal

MAP AND REDUCE: EXAMPLE

Count the number of words in a list of even length

```
def count_words(wordlist):
    return \
        reduce(add,
               map(lambda word: 1 if len(word)%2 == 0 else 0,
                   wordlist),
               0)
```



10 Cal

MAP AND REDUCE: EXAMPLE

Count the number of items in a list that satisfy a given predicate

```
def count_pred(pred, list):
    return \
        reduce(___,
               map(_____,
                   list),
               0)
```



11 Cal

MAP AND REDUCE: EXAMPLE

Count the number of items in a list that satisfy a given predicate

```
def count_pred(pred, list):
    return \
        reduce(add,
               map(lambda item: 1 if pred(item) else 0,
                   list),
               0)
```



12 Cal

MAP AND REDUCE: EXAMPLE

Filter out the items in a list that satisfy a given predicate

```
def filter(pred, list):
    return \
        reduce(_____,
              map(lambda item: _____ if pred(item) \
                  else __,
                  list),
              __)
```



MAP AND REDUCE: EXAMPLE

Filter out the items in a list that satisfy a given predicate

```
def filter(pred, list):
    return \
        reduce(lambda so_far, sublist: so_far + sublist,
              map(lambda item: [item] if pred(item) \
                  else [],
                  list),
              [])
```



ANNOUNCEMENTS

- You are now done with all projects and homework assignments! Congratulations! 😊
- Final lecture: You choose!
Piazza poll will be sent out today asking for suggestions for tomorrow's lecture.

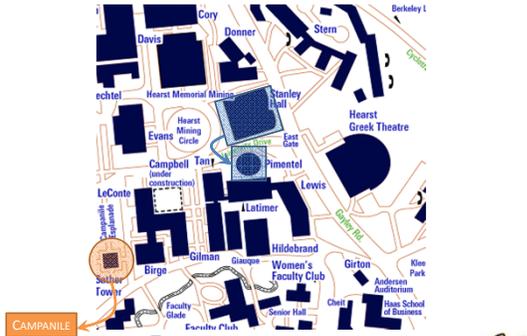


ANNOUNCEMENTS: FINAL

- Final is **Thursday, August 9**.
 - Where? 1 Pimentel.
 - When? 6PM to 9PM.
 - How much? All of the material in the course, from June 18 to August 8, will be tested.
- Closed book and closed electronic devices.
- One 8.5" x 11" 'cheat sheet' allowed.
- No group portion.
- Post-final laser tag?



ANNOUNCEMENTS: FINAL



http://berkeley.edu/map/may/0306.html



HOW MUCH DATA?

As of 2007, Google processes **20 petabytes** of data per day. Now, they probably process thousands of petabytes per day.

How much is a petabyte?

Source: <http://moss.com/blog/wp-content/uploads/2008/07/whatsapetabyte.pdf>



HOW MUCH DATA?

Assume that every 0 or 1 in a file is as large as a pixel on a screen (approx. 100 μm^2).

Then, a homework file, of size (approximately) 5 **kilobytes**, is about as big as

■



HOW MUCH DATA?

A music file has (approximately) 6 **megabytes**.
(1 megabyte = 1024 kilobytes)

■




HOW MUCH DATA?

One DVD-R has (approximately) 5 **gigabytes**.
(1 gigabyte = 1024 megabytes)

■



This is about as big as the total area of all pages of the *paper* copy of a homework.



HOW MUCH DATA?

One terabyte is 1024 **gigabytes**.
One petabyte is 1024 **terabytes**.

This is as much information as would fit in:
6 billion photos on Facebook, or
around 13 years of HD video.



HOW MUCH DATA?

One petabyte is an **extremely large amount of data**.
We *cannot* process this on one computer alone.

My computer would take around **7 years** to process all of 20 petabytes *exactly once*, and that assumes that all it *ever* does is process data and *never fails*.

Companies like Google do not have this luxury: they need data to be processed as quickly as possible.



PARALLELISM

Idea: Have a *lot* of computers working at once on different chunks of the same data, *independently* of other chunks and computers.



MAPREDUCE

All data is available at the outset; results are not consumed until processing completes

↓

Framework for batch processing of Big Data

System used by programmers to build applications

Very large and complex data sets

25

MAPREDUCE: HISTORY

- Developed at Google in 2003.
- Programmers had to consider a lot of details as part of their code: what to do when a computer fails, and how to divide the work up among different computers.
- However, many operations fit neatly into a *map* step, followed by a *reduce* step.

26

MAPREDUCE

The MapReduce framework is an *abstraction*: it allows programmers to code, while taking care of the underlying messy details, such as dealing with failure and distributing work across computers.

All the programmers had to do was create:

- a **mapper** (the function to map), and
- a **reducer** (the function to combine mapped results).

27

MAPREDUCE: EXECUTION MODEL

The diagram illustrates the execution model. It starts with an **Input** box at the top. Below it, a **MAPPER** block contains six circles labeled 'M'. Arrows point from the input to these mappers. Below the mappers is an **Intermediate** stage with six boxes containing key-value pairs: k1:v, k1:v, k2:v, k1:v, k3:v, k4:v, k4:v, k5:v, k4:v, k1:v, k3:v. A **SORTER** block (green) points to a **Group by Key** operation, which results in a **Grouped** stage with five boxes: k1:v, v, v, v, k2:v, k3:v, v, k4:v, v, v, k5:v. A **REDUCER** block (orange) points to five circles labeled 'R'. Arrows point from the grouped data to these reducers. Finally, an **Output** box is at the bottom. A purple callout bubble points to the intermediate stage with the text 'Key-value pairs!'.

<http://research.google.com/archive/mapreduce-osd04-slides/index-auto-0007.html>
28

MAPREDUCE

Map Phase
Apply a *mapper* over inputs, and emit intermediate *key-value pairs*.

Sort Phase
Group the intermediate key-value pairs by *keys*.

Reduce Phase
For each intermediate *key*, apply a *reducer* to “accumulate” all values associated with that key.

29

MAPREDUCE

Running example

Count all the vowels in Shakespeare’s works.

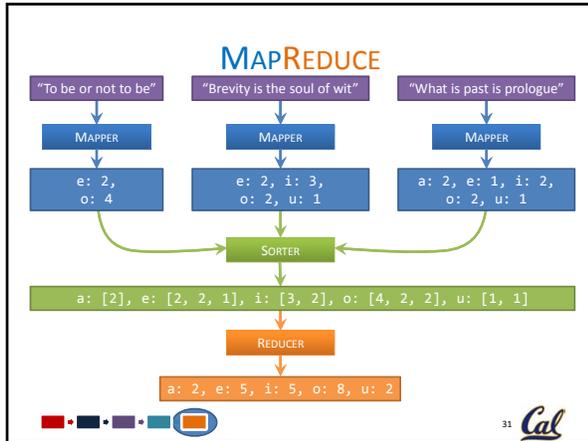
We need a **mapper** and a **reducer**.

Idea:

The **mapper** will count the number of vowels in each line, for each vowel.

The **reducer** will take the counts for each vowel from each line and total all the counts.

30



MAPREDUCE: MAPPER

```
import sys
from ucb import main
from mr import emit

def emit_vowels(line):
    for vowel in 'aeiou':
        count = line.count(vowel)
        if count > 0:
            emit(vowel, count)

@main
def run():
    for line in sys.stdin:
        emit_vowels(line)
```

2. Check if the line has a vowel.

3. If it does, "emit" – or send to the *standard output* – a key-value pair that associates the vowel with its count.

1. For every line in a file, fed through the *standard input* of this program...

32

MAPREDUCE: REDUCER

```
import sys
from ucb import main
from mr import emit, values_by_key

@main
def run():
    for key, value_iterator in values_by_key(sys.stdin):
        emit(key, sum(value_iterator))
```

Input: Lines of text representing key-value pairs, grouped by key
Output: Iterator over the keys and associated values. The "associated values" are also returned using an iterator.

1. Iterate over the keys and associated values, fed through the *standard input*...

2. ... and "emit" each key to the *standard output*, followed by the sum of the associated values.

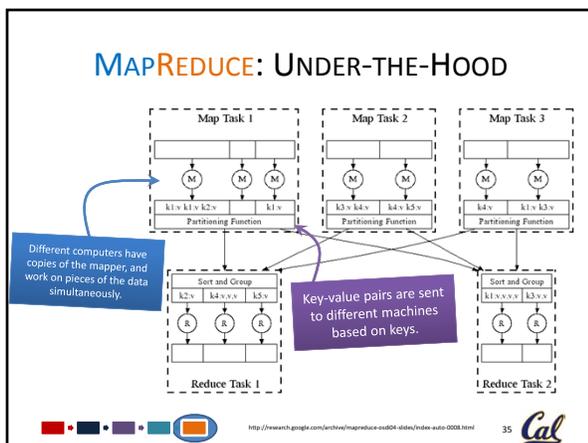
33

STANDARD INPUT AND OUTPUT

The *standard input* is where the program receives its input from external sources, such as the user or a file.

The *standard output* is where the program will send its output, such as the screen or another file.

34



MAPREDUCE: PRACTICE

What is a **mapper** and **reducer** that we can use to count the number of times particular words appear in a file?

Mapper: For every line in a file, emit...

Reducer: For every key,

36

MAPREDUCE: PRACTICE

What is a **mapper** and **reducer** that we can use to count the number of times particular words appear in a file?

Mapper: For every line in a file, emit a key-value pair that associates a “special” word with the number of times it appears in that line.

Reducer: For every key, sum up all the values associated with that key.



MAPREDUCE: PRACTICE

What is a **mapper** and **reducer** that we can use to count the number of lines in a file?

Mapper: For every line in a file, emit...

Reducer: For every key,



MAPREDUCE: PRACTICE

What is a **mapper** and **reducer** that we can use to count the number of lines in a file?

Mapper: For every line in a file, emit the key-value pair (“line”, 1).

Reducer: For every key, sum up all the values associated with that key.



MAPREDUCE

What does the framework provide us?

- **Fault Tolerance:** A machine or hard drive might crash.
The framework will automatically re-run failed tasks.
- **Speed:** Some machine might be slow because it is overloaded.
The framework can run multiple copies and keep the result of the one that finishes first.



MAPREDUCE

What does the framework provide us?

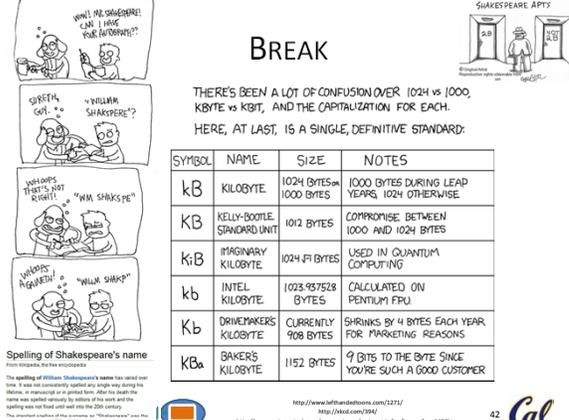
- **Network locality:** Data transfer is expensive.
The framework tries to schedule map tasks on the machines that hold the data to be processed.
- **Monitoring:** When will my job finish?
The framework provides a web-based interface describing jobs.



BREAK

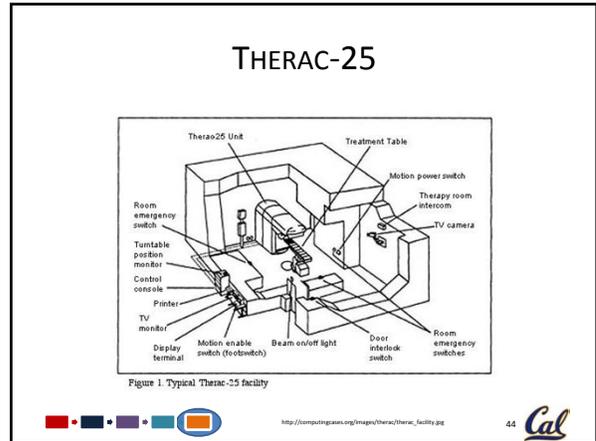
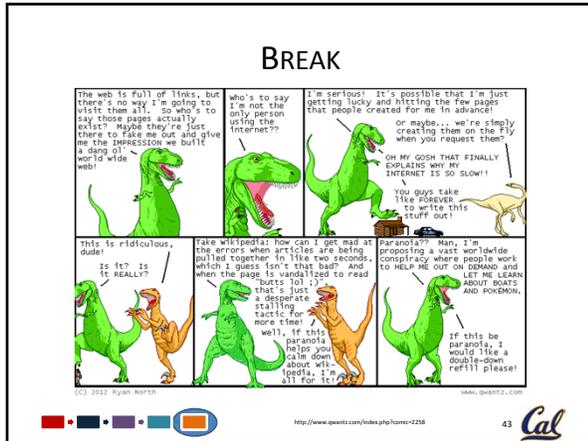
THERE'S BEEN A LOT OF CONFUSION OVER 1024 vs 1000, KBYTE vs KBIT, AND THE CAPITALIZATION FOR EACH. HERE, AT LAST, IS A SINGLE, DEFINITIVE STANDARD:

| SYMBOL | NAME | SIZE | NOTES |
|--------|----------------------------|---------------------|--|
| kB | KILOBYTE | 1024 BYTES | 1000 BYTES DURING LEAD YEARS, 1024 OTHERWISE |
| KB | KELLY-BOOTLE STANDARD UNIT | 1012 BYTES | COMPROMISE BETWEEN 1000 AND 1024 BYTES |
| KiB | IMAGINARY KILOBYTE | 1024 FI BYTES | USED IN QUANTUM COMPUTING |
| kb | INTEL KILOBYTE | 1023-957528 BYTES | CALCULATED ON PENTIUM FPU |
| Kb | DRIVEMAKERS KILOBYTE | CURRENTLY 908 BYTES | SHRINKS BY 4 BYTES EACH YEAR FOR MARKETING REASONS |
| KBa | BAKER'S KILOBYTE | 1152 BYTES | 9 BITS TO THE BYTE SINCE YOU'RE SUCH A GOOD CUSTOMER |



Spelling of Shakespeare's name
The spelling of William Shakespeare's name has varied over time. It was not consistently spelled any single way during his lifetime, as mentioned in a printed form. After his death the name was spelled variously by editors of his work and the spelling was not fixed until after the 20th century. The original spelling of the printing of "Shakespeare" was the





THERAC-25

- Built by AECL (Atomic Energy of Canada Limited)
- Dual-mode machine:
 - Treat using low-energy electron beams.
 - Treat using high-energy X-ray beams.
- Turntable with items for each mode placed on it: rotates to the correct position before a beam is started up.

http://computingwaves.org/ieee_materials/therac/supporting_docs/therac_case_narr/Machine_Design.html
<http://en.wikipedia.org/wiki/Therac-25>
45

THERAC-25

What happened?

- High-power beam activated instead of the low-power beam.
- High-power beam had approximately 100 times the intended dose of radiation.
- Six accidents between 1985 and 1987.
- Three deaths from radiation poisoning.

http://computingwaves.org/ieee_materials/therac/supporting_docs/therac_case_narr/Machine_Design.html
<http://en.wikipedia.org/wiki/Therac-25>
46

THERAC-25

Why? Software bugs and user interface issues.

- A one-byte counter frequently “overflowed”, or got too big for the memory allocated to it. If an operator provided manual input at the same time as the overflow, the interlock would fail.
- Failure only occurred during a nonstandard sequence of keystrokes: this went unnoticed for a long time.

http://computingwaves.org/ieee_materials/therac/supporting_docs/therac_case_narr/Machine_Design.html
<http://en.wikipedia.org/wiki/Therac-25>
47

THERAC-25

Why? Software bugs and user interface issues.

- Therac-25 did *not* have the hardware interlocks that its predecessors had, to prevent the wrong beam from activating.
- The developers relied on software interlocks, which were incorrectly coded.
- The code reused software from older models, where the hardware interlocks could not report that they were triggered.

http://computingwaves.org/ieee_materials/therac/supporting_docs/therac_case_narr/Machine_Design.html
<http://en.wikipedia.org/wiki/Therac-25>
48

THERAC-25

Why? Software bugs and user interface issues.

- The combination of hardware and software was never tested until at the hospital.
- If the system noticed that something was wrong, it would halt the X-ray beam and display the word "MALFUNCTION", followed by an unhelpful error code.

The operator could override the warning and proceed anyway.

- Complaints were not believed initially.



Source: http://computingcases.org/case_materials/thrac25/supporting_docs/thrac_case_study/Machine_Design.html
<http://en.wikipedia.org/wiki/Therac-25>



49

THERAC-25



http://bitstrangeness.com/2006



50

THERAC-25

What can we learn?

- Computers are *everywhere*: this means that the code you write can potentially be used in a life-or-death situation.
- In a complex codebase, software bugs can be *really hard* to catch and fix.
- This can be ameliorated significantly with *proper design and frequent testing*.



51

THERAC-25

What can we learn?

- Test not just the software, but how it interacts with hardware *and even the user*.
- The user is *not your enemy*.
 - A significantly large project will account for many possible erroneous inputs from the user.
 - The user should be provided useful output to understand the cause for error.



52

CONCLUSION

- MapReduce is a framework that allows programmers to process large amounts of data in parallel, without having to concern themselves with details such as computer failure and task assignment.
- Proper design and test of software is becoming increasingly essential, especially in a world where computing is ubiquitous.



53