Sometimes you need to think big thoughts



Not all performance problems will be solved with an incremental approach

- "Do we have to do it this way?"
- "Is there a better way to do this?"

• "Do I have to do this at all?"

1



"Through the hoop, Bob! Through the hoop!"

Traditional example: Sorting a new deck of cards

Method 1: Pattern recognition

- There are a finite number of possible arrangements
- Find which one you have, and then reorder
- $52! = 4x10^{66}$ so will need about $2*4x10^{66}/2$ comparisons

Method 2: Bubble sort

- Scan through, finding the smallest number
- Then repeat, scanning through the N-1 that's left
- Cost is $O(N^2)$ "sum of numbers from 1 to N" = $52*(52+1)/2 = 1.4x10^3$

Method 3: Better sorts - Shell sort, syncsort, split sort, ...

- Even for arbitrary data, better sort algorithms exist
- $O(N \log_2 N) = k * 52 * 5.7 = k * 300$, where "k" is time per operation
- For N large, important gain regardless of k
- As ideas improve, k has come down from 5 to about $1.2 \implies 360$

Method 4: Bin sort ("Solitaire sort")

- Use knowledge that there are 52 items with unique known labels
- Throw each card into the right bin with 52 calculations: O(N)

Method 5: New decks are already sorted (No operations!)

Bob Jacobsen, UC Berkeley





Tools and Techniques Lecture 2 Telling pions from kaons via Cherenkov light



Pions & Kaons have similar interactions in matter, differ in mass

Particles moving faster than light in a medium (glass, water) emit light

- Angle is related to velocity
- Light forms a cone







*>c

Radius of the reconstructed circle give particle type:



generic B Bbar events



How to make this fit?







CERN

School of Computing

Bad news: Rings get messy due to ambiguities in bouncing

Bob Jacobsen, UC Berkeley

Simple event with five charged particles:





Why is this hard?

Brute-force circle-finding is an O(N⁴) problem

• Basic algorithm: Are these four points consistent with a 'circle'?

Important to understand how cost grows with input size:







Realistic solution for DIRC?





Use what you know:

- Have track trajectories, know position and angle in DIRC bars
- All photons from a single track will have the same angle w.r.t. track No reason to expect that for photons from other tracks

For each track, plot angle between track and every photon - O(N)

- Don't do pattern recognition with individual photons
- Instead, look for overall pattern you already know is present



Not perfect, but optimal?

Bob Jacobsen, UC Berkeley

"But each operation is so much slower..."



How do I compare a "fast" O(N⁴) algorithm with a slow O(N)?



Many realistic problems deal with lots of data items

• Sharp coding is unlikely to save you a factor of 50² per calculation

Where else do we see this pattern?



What do we do when we can't figure out the exact answer?









Big things are different from small things

Bob Jacobsen, UC Berkeley

The life time of HEP software

Software is a long-term commitment

Many releases of the software are needed over its lifetime to fix bugs, add new features, support new platforms etc

How do we cope?

We try to find a way of working that leads to success

We create a "process" for building systems

We devise methods of communicating and record keeping: "models"

We use the best tools & methods we can lay our hands on

And we use a lot of optimism!

Can't technology save us?

More

We've built a series of ever-larger tools to handle large code projects:

Git for controlling and versioning code Tools for building "releases" of systems Tools for "configuration management"

But we struggle against three forces:

- •We're always building bigger & more difficult systems
- •We're always building bigger & more difficult collaborations
- •And we're the same old people

Net effect: We're always pushing the boundary of what we can do

Stupidity got us into this mess; why can't it get us out? - Will Rogers

How we got here:

First, you just wrote a big program But soon it was so big you wanted help So you broke it into pieces/files/modules But how do you share work on those?

CERN

School of Computing

Bob Jacobsen, UC Berkeley

Tools and Techniques Lecture 2 Version Control Systems (Hg, SVN, Git)

As systems & collaborations grow, efficiency goes down "Version" idea: Track changes from one version to next

Anybody can get a specific set of source

Big advantage: checkout is not exclusive

- More than one developer can have the same file checked out
- Developers can control their own use of the code for read, write
- Changes can come from multiple sources
- Tool handles (most) of the conflict resolution

More

Scaling is still an issue

Everybody is sharing a single repository

Every commit is immediately visible to everybody else

Development stands on shifting sand

Detailed records, but little understanding

Workarounds!

External record keeping tools

Package Coordinators

Issue with this arise at large & small level

At the level of developers and contributions, needed way to manage this

- Both tools and procedures
 - We'll be discussing & exercising git as typical tool
 - Individual collaborations have their own ways of sharing info

At the collaboration leveled, need procedures to ensure it all works

• "Nightly builds"

Now common in HEP - Gives early feedback on consistency problems

- "Continuous Integration", including automated testing Only works when people actually integrate early and often
- Reduces problems, but integration is still a lot of work

Copyright ③ 1995 United Feature Syndicate, Inc. Redistribution in whole or in part prohibited More

When Boeing wanted to design the 747, they had two choices:

- 1. Hire "SuperEngineer", who could do it alone
- 2. Hire 7,200 engineers and organize them to cooperate

Which did they choose?

Why?

What can we learn from this?

At first, Git looks like a simple file system...

You bring out a copy, work on it, and commit Git repository contains all that history

More

Committing to the Main Branch

Committing on a Branch and Merging to Main

Committing on a Branch and Merging to Main

WorkBranch

WorkBranch

Committing on a Branch and Merging to Main

Merging

Because Git focuses on commits, not on file versions, powerful merging

133 commits loaded

Bob Jacobsen, UC Berkeley

Multiple repositories with easy transfer of commits between

More

Bob Jacobsen, UC Berkeley

More than just mirroring

More

More than just mirroring

More than just mirroring

More

Branches are key

- **Develop on a separate branch** ٠
- Future Big Feature on branch
- And another one for || work
- Pays off for bug fix!
- Git merge to get fix across ٠
- Feature done, merges in
- New branch holds release
- and its inevitable fixes
- until <u>merge</u> and release main.
- Meanwhile, work proceeds ٠
- And the process repeats ٠

Gives understandable story?

Bob Jacobsen, UC Berkeley

Using all that history:

My feature broke between 0.1 and 1.0

Which commit broke it?

"git bisect" works through the graph

Was it in 0.2? No?

. . . .

Was it in merge before the release branch? Yes

I found a bug in a specific commit SHA

Which releases does it affect?

What's not affected?

"git diff tag1.0...SHA" to see if included

"git log" and "git revlog" explore history

Graphical representations can help a lot gitk, gitg tools

Complex! Linear history in repository would resolve these much easier

Git Rebase: An Editor for the Story

Git Rebase: An Editor for the Story

AS A PROJECT DRAGS ON, MY GIT COMMIT MESSAGES GET LESS AND LESS INFORMATIVE.

CERN

School of Computing

You want me to trust how many people?

How do you give 6,000 people access to a central repository? Use a distributed repository and "pull requests"

Git-based developers have a full local repository Commits have full context

"Push" moves all that to target

A "pull request" <u>sends</u> all that to somebody at the target, who can accept or not

When accepted, the merge is completed & both repositories in sync (Pull requests rarely rejected outright - usually it's "fix these things and resend")

Strong tools exist to make pull requests easy: CI test results, etc automated

More

Life Cycle of a Pull Request

Bob is working on his laptop, and commits another change locally:

```
% git commit -m"Cover rest of classes" help/en/html/tools
[ctc-tools 79c28b4c93] Cover rest of classes
1 file changed, 14 insertions(+)
```

Life Cycle of a Pull Request

Bob is working on his laptop, and commits another change locally:

```
% git commit -m"Cover rest of classes" help/en/html/tools
[ctc-tools 79c28b4c93] Cover rest of classes
1 file changed, 14 insertions(+)
```

He's ready for that work to be reviewed, and wants to move it to a repository that's always online:

```
% git push
Counting objects: 8, done.
Delta compression using up to 4 threads.
Compressing objects: 100% (7/7), done.
Writing objects: 100% (8/8), 1.07 KiB | 0 bytes/s, done.
Total 8 (delta 6), reused 0 (delta 0)
remote: Resolving deltas: 100% (6/6), completed with 6 local objects.
To https://github.com/bobjacobsen/JMRI.git
3d35322e43..79c28b4c93 ctc-tools -> ctc-tools
```

Life Cycle of a Pull Request

Computing

Open a pull request

Create a new pull request by comparing changes across two branches. If you need to, you can also compare across forks.

Life Cycle of a Pull Request

Once created:

Continuous integration tests are run

All checks have passed 4 successful checks		Hide all checks	
~	Service VersionEye — All software dependencies are fine. You are awesome!	Details	
~	o continuous-integration/appveyor/pr — AppVeyor build succeeded	Details	
~	continuous-integration/travis-ci/pr — The Travis CI build passed	Details	
~	coverage/coveralls — Coverage increased (+0.02%) to 33.589%	Details	

Reviews happen

Merge checks are done

V

Merge pull request

This branch has no conflicts with the base branch Merging can be performed automatically.

And finally, somebody with authorization can click this:

✓ Create a merge commit

All commits from this branch will be added to the base branch via a merge commit.

Squash and merge

The 18 commits from this branch will be combined into one commit in the base branch.

Rebase and merge

The 18 commits from this branch will be rebased and added to the base branch.

to complete the merge onto the desired branch in the main repository.

You can also open this in GitHub Desktop or view command line instructions.

How do you use this all?

Individually:

Use it to work independently Both of others, and of yourself!
Collaborate on intermediate results Clean branches easy to share: "Try bobj/FixIssue10343"
Shape your work result to make it understandable Comments, squashing, comments, rebasing as tools
Integrate early and often! Pull "main" and make sure work is still OK

For a collaboration project:

Help people work at the scales they need to Individually, in small groups, large groups, ...
Control how code is added/updated Shaping contents of common development, releases
Make the contents understandable Tags, known branching / linear history

Series summary

Software engineering is the art of building complex computer systems

It's ideas and techniques spring from our need to handle size & complexity

As you do your own work & develop your own skills, consider:

- How your effort effects or contributes to things 10X, 100X, 1000X larger
- How you'll do things different/better when it's your problem

Exercises on Tuesday

Test Frameworks Performance Profiling Memory Issues Code Management

2:00 PM	Study time / daily sports	2:00 PM	Study time / daily sports
		3:30	Student presentations
1:45	Coffee break	РМ	
1:00	Tools and Techniques - exercises - Bob Jacobsen		
РМ	(UC Berkeley)		
		4:30	Coffee break
		PM	Data Salanga - avaraigan - Rob Jacoboon (UC
		5:00	Data Science - exercises - Bob Jacobsen (UC
		PM	Berkeley)

Instructions to get started on Indigo (Tools & Techniques E1) <u>https://indico.cern.ch/event/1125271/contributions/4723248/</u> More

You'll work in pairs. Try to find somebody with complementary skills!

Learn about each topic, spend more time on the ones that interest you. Speed is not the issue: no reward for first done, no complaint about last.

Think about what you're doing: There are larger lessons to be found!