


FAST, FLEXIBLE AND FUN: REVISION CONTROL WITH MERCURIAL

Martin Geisler
<mg@aragost.com>

IBM Research – Zurich
September 16th, 2010



ABOUT THE SPEAKER

Martin Geisler:

- ▶ core Mercurial developer:
 - ▶ reviews patches from the community
 - ▶ helps users in our IRC channel
- ▶ PhD in Computer Science from Aarhus University, DK
 - ▶ exchange student at ETH Zürich in 2005
 - ▶ visited ZRL for three months in 2008
- ▶ now working at aragost Trifork, Zürich



OUTLINE

INTRODUCTION

USING MERCURIAL

- Workflows

- Branching

- The Underlying Model

- Using History

COOL EXTENSIONS

- Changing History

- Talking to Other Systems

- Third-Party Tools

WRAPPING UP



OUTLINE

INTRODUCTION

USING MERCURIAL

Workflows

Branching

The Underlying Model

Using History

COOL EXTENSIONS

Changing History

Talking to Other Systems

Third-Party Tools

WRAPPING UP



WHAT IS MERCURIAL?

Main features:

- ▶ fast, **distributed** revision control system
- ▶ robust support for branching **and** merging
- ▶ very flexible and extensible



WHO IS USING IT?

Mercurial is used by:

- ▶ Oracle for Java, OpenSolaris, NetBeans, OpenOffice, ...
- ▶ Mozilla for Firefox, Thunderbird, ...
- ▶ Google
- ▶ many more...

OpenJDK



OpenOffice.org




NetBeans



TESTIMONIALS

Martin Fowler, 2010:

<http://martinfowler.com/>



TESTIMONIALS

Joel Spolsky, 2010: <http://www.jelonsoftware.com/>

Mercurial is better than Subversion.


It is a better way of working on source code with a team. It is a better way of working on source code by yourself.

It is just better.




CENTRALIZED REVISION CONTROL

Single repository, multiple working copies:




CENTRALIZED REVISION CONTROL

Single repository, multiple working copies:




DISTRIBUTED REVISION CONTROL

Mercurial duplicates the history on many servers:



Alice

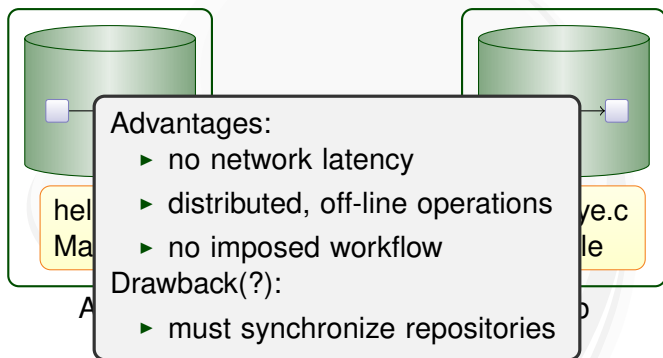


Bob



DISTRIBUTED REVISION CONTROL

Mercurial duplicates the history on many servers:



MOVING CHANGESETS AROUND

Pull and merge:

Alice



Bob



MOVING CHANGESETS AROUND

Pull and merge:

Alice




Bob



MOVING CHANGESETS AROUND

Pull and merge:

Alice



Bob



MOVING CHANGESETS AROUND

Pull and merge:

Alice




Bob




MOVING CHANGESETS AROUND

Pull and merge:



MOVING CHANGESETS AROUND


Pull and merge:



MOVING CHANGESETS AROUND

Pull and merge:

Alice



Bob

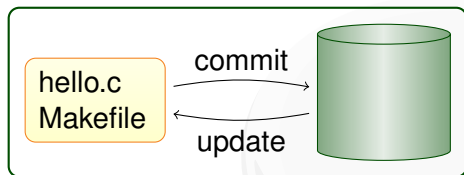


Merging:

- ▶ find common ancestor of A_2 and B_1 : I
- ▶ compute differences between I and B_1
- ▶ apply them to A_2 , taking renames into account



KEY MERCURIAL COMMANDS




Alice

Local commands:

- ▶ **hg commit**: save a snapshot into the current repository
- ▶ **hg update**: checkout revision into working directory
- ▶ **hg merge**: join different lines of history



KEY MERCURIAL COMMANDS



Local commands:

- ▶ **hg commit**: save a snapshot into the current repository
- ▶ **hg update**: checkout revision into working directory
- ▶ **hg merge**: join different lines of history

Network commands:

- ▶ **hg pull**: retrieve changesets from another repository
- ▶ **hg push**: send your changesets to another repository



OUTLINE

INTRODUCTION

USING MERCURIAL

Workflows

Branching

The Underlying Model

Using History

COOL EXTENSIONS

Changing History

Talking to Other Systems

Third-Party Tools

WRAPPING UP



OUTLINE

INTRODUCTION

USING MERCURIAL

Workflows

Branching

The Underlying Model

Using History

COOL EXTENSIONS

Changing History

Talking to Other Systems


Third-Party Tools

WRAPPING UP




WORKFLOW IN A TEAM

Mercurial scales from a single team... :




WORKFLOW BETWEEN COMPANY DIVISIONS

... to enterprise-wide development... :



WORKFLOW BETWEEN TWO COMPUTERS

... to working with yourself:



OUTLINE

INTRODUCTION

USING MERCURIAL

Workflows

Branching

The Underlying Model

Using History

COOL EXTENSIONS

Changing History

Talking to Other Systems

Third-Party Tools


WRAPPING UP



RELEASE BRANCHES




RELEASE BRANCHES




RELEASE BRANCHES




RELEASE BRANCHES




RELEASE BRANCHES




RELEASE BRANCHES




RELEASE BRANCHES



RELEASE BRANCHES



RELEASE BRANCHES



BRANCHES IN SVN

Subversion knows nothing about branches!

- ▶ but SVN has a cheap copy mechanism
- ▶ used for tags and branches



BRANCHES IN SVN

Subversion knows nothing about branches!

- ▶ but SVN has a cheap copy mechanism
- ▶ used for tags and branches

r10


```
trunk/  
  hello.c  
  Makefile  
branches/  
tags/
```



BRANCHES IN SVN

Subversion knows nothing about branches!


- ▶ but SVN has a cheap copy mechanism
- ▶ used for tags and branches



BRANCHES IN SVN

Subversion knows nothing about branches!

- ▶ but SVN has a cheap copy mechanism
- ▶ used for tags and branches



MERGING BRANCHES IN SVN

The support is incomplete and fragile:

- ▶ renamed files are not merged correctly
- ▶ old clients will not update the merge info



MERGING BRANCHES IN SVN

The support is incomplete and fragile:

- ▶ renamed files are not merged correctly
- ▶ old clients will not update the merge info

From the SVN Book:

The bottom line is that Subversion's merge-tracking feature has an **extremely complex** internal implementation, and the `svn:mergeinfo` property is the only window the user has into the machinery. Because the feature is **relatively new**, a numbers of edge cases and possible unexpected behaviors may pop up.

—*Version Control with Subversion*

(Mercurial has robust built-in support for merging branches.)



OUTLINE

INTRODUCTION

USING MERCURIAL

Workflows

Branching

The Underlying Model

Using History

COOL EXTENSIONS

Changing History

Talking to Other Systems

Third-Party Tools


WRAPPING UP



THE UNDERLYING MODEL

A Mercurial changeset conceptually consist of:

- ▶ 0–2 parent changeset IDs:
 - ▶ root changeset has no parents
 - ▶ normal changesets have one parent
 - ▶ merge changesets have two parents
- ▶ date, username, commit message
- ▶ difference from first parent changeset
- ▶ changeset ID is computed as SHA-1 hash of the above
- ▶ makes it impossible to inject **malicious code** on server



OUTLINE

INTRODUCTION

USING MERCURIAL

Workflows

Branching

The Underlying Model

Using History

COOL EXTENSIONS

Changing History

Talking to Other Systems

Third-Party Tools

WRAPPING UP



BROWSING THE HISTORY OF A FILE

The `hg annotate` command is invaluable:

- ▶ you see when each line was introduced
- ▶ you can quickly jump back to earlier versions

History of Mercurial's README file:

```
3942: Basic install:
 445:
3942: $ make                # see install targets
3942: $ make install        # do a system-wide install
3942: $ hg debuginstall      # sanity-check setup
3942: $ hg                    # see help
  0:
# ...
```

Better interface in `hg serve`



SEARCHING FILE CONTENT

Ever wondered when a function was introduced?

- ▶ `hg grep` can help you!


Example: When was `hg forget` introduced?

```
% hg grep --all 'def forget' commands.py
commands.py:8902:+:def forget(ui, repo, *pats, **opts):
commands.py:3522:-:def forget(ui, repo, *pats, **opts):
commands.py:814:-:def forget(ui, repo, file1, *files):
commands.py:814:+:def forget(ui, repo, *pats, **opts):
# ...
```




REVISION GRAPH BISECTION

You've found a bug! When was it first introduced?
Use `hg bisect` to mark good and bad revisions:




REVISION GRAPH BISECTION

You've found a bug! When was it first introduced?
Use `hg bisect` to mark good and bad revisions:




REVISION GRAPH BISECTION

You've found a bug! When was it first introduced?
Use `hg bisect` to mark good and bad revisions:




REVISION GRAPH BISECTION

You've found a bug! When was it first introduced?
Use `hg bisect` to mark good and bad revisions:




REVISION GRAPH BISECTION

You've found a bug! When was it first introduced?
Use `hg bisect` to mark good and bad revisions:




REVISION GRAPH BISECTION

You've found a bug! When was it first introduced?
Use `hg bisect` to mark good and bad revisions:




REVISION GRAPH BISECTION

You've found a bug! When was it first introduced?
Use `hg bisect` to mark good and bad revisions:




REVISION GRAPH BISECTION

You've found a bug! When was it first introduced?
Use `hg bisect` to mark good and bad revisions:



REVISION GRAPH BISECTION

You've found a bug! When was it first introduced?
Use `hg bisect` to mark good and bad revisions:



OUTLINE

INTRODUCTION

USING MERCURIAL

Workflows

Branching

The Underlying Model

Using History

COOL EXTENSIONS

Changing History

Talking to Other Systems

Third-Party Tools

WRAPPING UP



MERCURIAL IS EXTENSIBLE

You can add new functionality to Mercurial:

- ▶ ships with 30+ extensions
- ▶ wiki lists 75+ extensions
- ▶ extensions can change basically everything
- ▶ helps to keep the core small and focused



OUTLINE

INTRODUCTION

USING MERCURIAL

Workflows

Branching

The Underlying Model

Using History

COOL EXTENSIONS

Changing History

Talking to Other Systems

Third-Party Tools

WRAPPING UP



MOVING CHANGESETS AROUND

Tired of all those merges? Use the **rebase** extension!

- ▶ Revision graph:



MOVING CHANGESETS AROUND

Tired of all those merges? Use the **rebase** extension!


- ▶ Revision graph:



MOVING CHANGESETS AROUND

Tired of all those merges? Use the **rebase** extension!


- ▶ Revision graph:




MOVING CHANGESETS AROUND

Tired of all those merges? Use the **rebase** extension!

- ▶ Revision graph:




- ▶ Merge:




MOVING CHANGESETS AROUND

Tired of all those merges? Use the **rebase** extension!


- ▶ Revision graph:



- ▶ Merge:




- ▶ Rebase:




MOVING CHANGESETS AROUND

Tired of all those merges? Use the **rebase** extension!


- ▶ Revision graph:



- ▶ Merge:



- ▶ Rebase:



- ▶ Beware: public changes should never be rebased.



EDITING HISTORY

Inspired by `git rebase -i`, `histedit` lets you

- ▶ reorder changesets:



EDITING HISTORY

Inspired by `git rebase -i`, `histedit` lets you

- ▶ reorder changesets:



- ▶ fold changesets:



EDITING HISTORY

Inspired by `git rebase -i`, `histedit` lets you

- ▶ reorder changesets:



- ▶ fold changesets:



- ▶ drop changesets:



EDITING HISTORY

Inspired by `git rebase -i`, `histedit` lets you

- ▶ reorder changesets:



- ▶ fold changesets:



- ▶ drop changesets:



- ▶ edit changesets:



OUTLINE

INTRODUCTION

USING MERCURIAL

Workflows

Branching

The Underlying Model

Using History

COOL EXTENSIONS

Changing History

Talking to Other Systems

Third-Party Tools

WRAPPING UP



MIGRATING HISTORY

The **convert** extension can import history:

- ▶ CVS, SVN, Git, Bazaar, Darcs, ...
- ▶ incremental conversion
- ▶ many options for fiddling with branches, authors, ...



MIGRATING HISTORY

The **convert** extension can import history:

- ▶ CVS, SVN, Git, Bazaar, Darcs, ...
- ▶ incremental conversion
- ▶ many options for fiddling with branches, authors, ...

Interestingly, **convert** can import from Mercurial:

- ▶ **--filemap** lets you exclude and rename files
- ▶ **--branchmap** lets you rename branches



INTERFACING WITH SUBVERSION

The `hgsubversion` extension let's you:

- ▶ use `hg clone` on a SVN URL
- ▶ use `hg pull` to convert new SVN revisions
- ▶ use `hg push` to commit changesets to SVN server

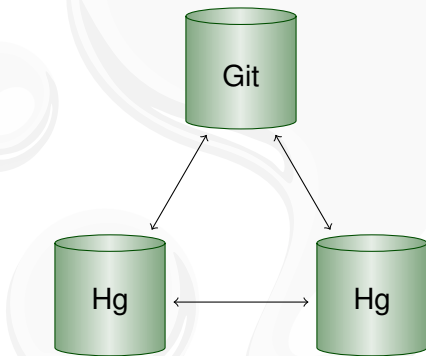
Goal: make `hg` a better Subversion client than `svn`!



INTERFACING WITH GIT

Need to work on a Git repository? Try **hg-git**!

- ▶ Mercurial extension: you get the nice **hg** command line
- ▶ round-tripping: changeset hashes are preserved



OUTLINE

INTRODUCTION

USING MERCURIAL

Workflows

Branching

The Underlying Model

Using History

COOL EXTENSIONS

Changing History

Talking to Other Systems

Third-Party Tools

WRAPPING UP



THIRD-PARTY TOOLS

Tools with Mercurial support:

- ▶ Shell integration: TortoiseHg (Windows, Mac, Linux)
- ▶ IDEs: Eclipse, NetBeans, IntelliJ, Visual Studio, Emacs...
- ▶ Project Support: Trac, JIRA, Maven, Hudson, BuildBot...



OUTLINE

INTRODUCTION

USING MERCURIAL

- Workflows

- Branching

- The Underlying Model

- Using History

COOL EXTENSIONS

- Changing History

- Talking to Other Systems

- Third-Party Tools

WRAPPING UP



MERCURIAL IN A NUTSHELL

Mercurial changes the way you develop:

- ▶ simple yet strong model for **both** branching and merging
- ▶ power tool instead of necessary evil
- ▶ light-weight and snappy



MORE INFORMATION

- ▶ Mercurial homepage:
<http://mercurial.selenic.com/>
- ▶ *Mercurial: The Definitive Guide*:
<http://hgbook.red-bean.com/>
- ▶ Getting Started:
<http://mercurial.aragost.com/kick-start/>
<http://mercurial.ch/>
<http://hginit.com/>
- ▶ Some free Mercurial hosting sites:
<http://bitbucket.org/>
<http://code.google.com/>
<http://sourceforge.net/>
<http://www.codeplex.com/>



CONTACT

Please get in touch if you have more questions or have considered using Mercurial in your organization:

- ▶ Email: mg@aragost.com
- ▶ IRC: [mg](#) in [#mercurial](#) on irc.freenode.net



MERCURIAL CONTRIBUTORS

From <http://ohloh.net/p/mercurial/map:>



Showing 50 of [325 contributors](#).

MERCURIAL CONTRIBUTORS

From <http://ohloh.net/p/mercurial/map>:

