

Distributed revision control with Mercurial

Bryan O'Sullivan

Copyright © 2006, 2007 Bryan O'Sullivan.

This material may be distributed only subject to the terms and conditions set forth in version 1.0 of the Open Publication License. Please refer to Appendix **D** for the license text.

This book was prepared from [rev 4700dd38384c](#), dated 2007-12-14 11:08 -0800, using [rev fe2e81229819](#) of Mercurial.

Contents

Contents	i
Preface	2
0.1 This book is a work in progress	2
0.2 About the examples in this book	2
0.3 Colophon—this book is Free	2
1 Introduction	3
1.1 About revision control	3
1.1.1 Why use revision control?	3
1.1.2 The many names of revision control	4
1.2 A short history of revision control	4
1.3 Trends in revision control	5
1.4 A few of the advantages of distributed revision control	5
1.4.1 Advantages for open source projects	6
1.4.2 Advantages for commercial projects	6
1.5 Why choose Mercurial?	7
1.6 Mercurial compared with other tools	7
1.6.1 Subversion	7
1.6.2 Git	8
1.6.3 CVS	8
1.6.4 Commercial tools	9
1.6.5 Choosing a revision control tool	9
1.7 Switching from another tool to Mercurial	9
2 A tour of Mercurial: the basics	10
2.1 Installing Mercurial on your system	10
2.1.1 Linux	10
2.1.2 Solaris	10
2.1.3 Mac OS X	11
2.1.4 Windows	11
2.2 Getting started	11
2.2.1 Built-in help	11
2.3 Working with a repository	12
2.3.1 Making a local copy of a repository	12
2.3.2 What's in a repository?	12
2.4 A tour through history	13
2.4.1 Changesets, revisions, and talking to other people	14
2.4.2 Viewing specific revisions	14
2.4.3 More detailed information	16

2.5	All about command options	16
2.6	Making and reviewing changes	17
2.7	Recording changes in a new changeset	18
2.7.1	Setting up a username	18
2.7.2	Writing a commit message	19
2.7.3	Writing a good commit message	19
2.7.4	Aborting a commit	20
2.7.5	Admiring our new handiwork	20
2.8	Sharing changes	20
2.8.1	Pulling changes from another repository	20
2.8.2	Updating the working directory	21
2.8.3	Pushing changes to another repository	23
2.8.4	Sharing changes over a network	23
3	A tour of Mercurial: merging work	25
3.1	Merging streams of work	25
3.1.1	Head changesets	27
3.1.2	Performing the merge	27
3.1.3	Committing the results of the merge	29
3.2	Merging conflicting changes	29
3.2.1	Using a graphical merge tool	30
3.2.2	A worked example	30
3.3	Simplifying the pull-merge-commit sequence	32
4	Behind the scenes	35
4.1	Mercurial's historical record	35
4.1.1	Tracking the history of a single file	35
4.1.2	Managing tracked files	35
4.1.3	Recording changeset information	35
4.1.4	Relationships between revisions	35
4.2	Safe, efficient storage	36
4.2.1	Efficient storage	37
4.2.2	Safe operation	37
4.2.3	Fast retrieval	37
4.2.4	Identification and strong integrity	38
4.3	Revision history, branching, and merging	38
4.4	The working directory	38
4.4.1	What happens when you commit	39
4.4.2	Creating a new head	39
4.4.3	Merging heads	39
4.5	Other interesting design features	40
4.5.1	Clever compression	40
4.5.2	Read/write ordering and atomicity	41
4.5.3	Concurrent access	41
4.5.4	Avoiding seeks	41
4.5.5	Other contents of the dirstate	42
5	Mercurial in daily use	47
5.1	Telling Mercurial which files to track	47
5.1.1	Explicit versus implicit file naming	47
5.1.2	Aside: Mercurial tracks files, not directories	48
5.2	How to stop tracking a file	48

5.2.1	Removing a file does not affect its history	49
5.2.2	Missing files	49
5.2.3	Aside: why tell Mercurial explicitly to remove a file?	50
5.2.4	Useful shorthand—adding and removing files in one step	50
5.3	Copying files	50
5.3.1	The results of copying during a merge	50
5.3.2	Why should changes follow copies?	52
5.3.3	How to make changes <i>not</i> follow a copy	52
5.3.4	Behaviour of the “hg copy” command	52
5.4	Renaming files	53
5.4.1	Renaming files and merging changes	54
5.4.2	Divergent renames and merging	54
5.4.3	Convergent renames and merging	55
5.4.4	Other name-related corner cases	55
5.5	Recovering from mistakes	55
6	Collaborating with other people	57
6.1	Mercurial’s web interface	57
6.2	Collaboration models	57
6.2.1	Factors to keep in mind	57
6.2.2	Informal anarchy	58
6.2.3	A single central repository	58
6.2.4	Working with multiple branches	58
6.2.5	Feature branches	61
6.2.6	The release train	61
6.2.7	The Linux kernel model	61
6.2.8	Pull-only versus shared-push collaboration	62
6.2.9	Where collaboration meets branch management	62
6.3	The technical side of sharing	62
6.4	Informal sharing with “hg serve”	62
6.4.1	A few things to keep in mind	63
6.5	Using the Secure Shell (ssh) protocol	63
6.5.1	How to read and write ssh URLs	63
6.5.2	Finding an ssh client for your system	64
6.5.3	Generating a key pair	64
6.5.4	Using an authentication agent	64
6.5.5	Configuring the server side properly	65
6.5.6	Using compression with ssh	66
6.6	Serving over HTTP using CGI	67
6.6.1	Web server configuration checklist	67
6.6.2	Basic CGI configuration	67
6.6.3	Sharing multiple repositories with one CGI script	69
6.6.4	Downloading source archives	70
6.6.5	Web configuration options	70
7	File names and pattern matching	73
7.1	Simple file naming	73
7.2	Running commands without any file names	73
7.3	Telling you what’s going on	74
7.4	Using patterns to identify files	75
7.4.1	Shell-style glob patterns	75
7.4.2	Regular expression matching with re patterns	76

7.5	Filtering files	76
7.6	Ignoring unwanted files and directories	77
7.7	Case sensitivity	77
7.7.1	Safe, portable repository storage	77
7.7.2	Detecting case conflicts	78
7.7.3	Fixing a case conflict	78
8	Managing releases and branchy development	79
8.1	Giving a persistent name to a revision	79
8.1.1	Handling tag conflicts during a merge	81
8.1.2	Tags and cloning	81
8.1.3	When permanent tags are too much	82
8.2	The flow of changes—big picture vs. little	82
8.3	Managing big-picture branches in repositories	82
8.4	Don't repeat yourself: merging across branches	83
8.5	Naming branches within one repository	84
8.6	Dealing with multiple named branches in a repository	86
8.7	Branch names and merging	87
8.8	Branch naming is generally useful	88
9	Finding and fixing your mistakes	89
9.1	Erasing local history	89
9.1.1	The accidental commit	89
9.1.2	Rolling back a transaction	89
9.1.3	The erroneous pull	90
9.1.4	Rolling back is useless once you've pushed	90
9.1.5	You can only roll back once	91
9.2	Reverting the mistaken change	91
9.2.1	File management errors	92
9.3	Dealing with committed changes	93
9.3.1	Backing out a changeset	93
9.3.2	Backing out the tip changeset	94
9.3.3	Backing out a non-tip change	94
9.3.4	Gaining more control of the backout process	96
9.3.5	Why “hg backout” works as it does	98
9.4	Changes that should never have been	99
9.4.1	Protect yourself from “escaped” changes	100
9.5	Finding the source of a bug	100
9.5.1	Using the bisect extension	101
9.5.2	Cleaning up after your search	104
9.6	Tips for finding bugs effectively	104
9.6.1	Give consistent input	104
9.6.2	Automate as much as possible	105
9.6.3	Check your results	105
9.6.4	Beware interference between bugs	105
9.6.5	Bracket your search lazily	105
10	Handling repository events with hooks	106
10.1	An overview of hooks in Mercurial	106
10.2	Hooks and security	107
10.2.1	Hooks are run with your privileges	107
10.2.2	Hooks do not propagate	107

10.2.3	Hooks can be overridden	107
10.2.4	Ensuring that critical hooks are run	107
10.3	Care with <code>pretxn</code> hooks in a shared-access repository	108
10.3.1	The problem illustrated	108
10.4	A short tutorial on using hooks	109
10.4.1	Performing multiple actions per event	109
10.4.2	Controlling whether an activity can proceed	110
10.5	Writing your own hooks	110
10.5.1	Choosing how your hook should run	110
10.5.2	Hook parameters	111
10.5.3	Hook return values and activity control	111
10.5.4	Writing an external hook	111
10.5.5	Telling Mercurial to use an in-process hook	111
10.5.6	Writing an in-process hook	111
10.6	Some hook examples	112
10.6.1	Writing meaningful commit messages	112
10.6.2	Checking for trailing whitespace	112
10.7	Bundled hooks	113
10.7.1	<code>acl</code> —access control for parts of a repository	113
10.7.2	<code>bugzilla</code> —integration with Bugzilla	115
10.7.3	<code>notify</code> —send email notifications	118
10.8	Information for writers of hooks	119
10.8.1	In-process hook execution	119
10.8.2	External hook execution	120
10.8.3	Finding out where changesets come from	120
10.9	Hook reference	121
10.9.1	<code>changegroup</code> —after remote changesets added	121
10.9.2	<code>commit</code> —after a new changeset is created	121
10.9.3	<code>incoming</code> —after one remote changeset is added	121
10.9.4	<code>outgoing</code> —after changesets are propagated	122
10.9.5	<code>prechangegroup</code> —before starting to add remote changesets	122
10.9.6	<code>precommit</code> —before starting to commit a changeset	122
10.9.7	<code>preoutgoing</code> —before starting to propagate changesets	123
10.9.8	<code>pretag</code> —before tagging a changeset	123
10.9.9	<code>pretxnchangegroup</code> —before completing addition of remote changesets	123
10.9.10	<code>pretxncommit</code> —before completing commit of new changeset	124
10.9.11	<code>preupdate</code> —before updating or merging working directory	124
10.9.12	<code>tag</code> —after tagging a changeset	124
10.9.13	<code>update</code> —after updating or merging working directory	124
11	Customising the output of Mercurial	126
11.1	Using precanned output styles	126
11.1.1	Setting a default style	127
11.2	Commands that support styles and templates	127
11.3	The basics of templating	127
11.4	Common template keywords	128
11.5	Escape sequences	129
11.6	Filtering keywords to change their results	130
11.6.1	Combining filters	131
11.7	From templates to styles	131
11.7.1	The simplest of style files	133
11.7.2	Style file syntax	133

11.8	Style files by example	133
11.8.1	Identifying mistakes in style files	133
11.8.2	Uniquely identifying a repository	134
11.8.3	Mimicking Subversion's output	134
12	Managing change with Mercurial Queues	136
12.1	The patch management problem	136
12.2	The prehistory of Mercurial Queues	136
12.2.1	A patchwork quilt	137
12.2.2	From patchwork quilt to Mercurial Queues	137
12.3	The huge advantage of MQ	137
12.4	Understanding patches	138
12.5	Getting started with Mercurial Queues	138
12.5.1	Creating a new patch	139
12.5.2	Refreshing a patch	140
12.5.3	Stacking and tracking patches	140
12.5.4	Manipulating the patch stack	141
12.5.5	Pushing and popping many patches	141
12.5.6	Safety checks, and overriding them	142
12.5.7	Working on several patches at once	142
12.6	More about patches	142
12.6.1	The strip count	142
12.6.2	Strategies for applying a patch	144
12.6.3	Some quirks of patch representation	144
12.6.4	Beware the fuzz	145
12.6.5	Handling rejection	146
12.7	Getting the best performance out of MQ	146
12.8	Updating your patches when the underlying code changes	147
12.9	Identifying patches	148
12.10	Useful things to know about	148
12.11	Managing patches in a repository	148
12.11.1	MQ support for patch repositories	149
12.11.2	A few things to watch out for	150
12.12	Third party tools for working with patches	150
12.13	Good ways to work with patches	150
12.14	MQ cookbook	151
12.14.1	Manage "trivial" patches	151
12.14.2	Combining entire patches	153
12.14.3	Merging part of one patch into another	153
12.15	Differences between quilt and MQ	154
13	Advanced uses of Mercurial Queues	155
13.1	The problem of many targets	155
13.1.1	Tempting approaches that don't work well	155
13.2	Conditionally applying patches with guards	156
13.3	Controlling the guards on a patch	156
13.4	Selecting the guards to use	157
13.5	MQ's rules for applying patches	158
13.6	Trimming the work environment	158
13.7	Dividing up the series file	158
13.8	Maintaining the patch series	159
13.8.1	The art of writing backport patches	160

13.9	Useful tips for developing with MQ	160
13.9.1	Organising patches in directories	160
13.9.2	Viewing the history of a patch	160
14	Adding functionality with extensions	163
14.1	Improve performance with the <code>inotify</code> extension	163
14.2	Flexible diff support with the <code>extdiff</code> extension	165
14.2.1	Defining command aliases	167
14.3	Cherrypicking changes with the <code>transplant</code> extension	167
14.4	Send changes via email with the <code>patchbomb</code> extension	167
14.4.1	Changing the behaviour of patchbombs	168
A	Command reference	169
A.1	“ <code>hg add</code> ”—add files at the next commit	169
A.2	“ <code>hg diff</code> ”—print changes in history or working directory	169
A.2.1	Options	169
A.3	“ <code>hg version</code> ”—print version and copyright information	171
A.3.1	Tips and tricks	171
B	Mercurial Queues reference	173
B.1	MQ command reference	173
B.1.1	“ <code>hg qapplied</code> ”—print applied patches	173
B.1.2	“ <code>hg qcommit</code> ”—commit changes in the queue repository	173
B.1.3	“ <code>hg qdelete</code> ”—delete a patch from the <code>series</code> file	173
B.1.4	“ <code>hg qdiff</code> ”—print a diff of the topmost applied patch	173
B.1.5	“ <code>hg qfold</code> ”—merge (“fold”) several patches into one	173
B.1.6	“ <code>hg qheader</code> ”—display the header/description of a patch	174
B.1.7	“ <code>hg qimport</code> ”—import a third-party patch into the queue	174
B.1.8	“ <code>hg qinit</code> ”—prepare a repository to work with MQ	174
B.1.9	“ <code>hg qnew</code> ”—create a new patch	174
B.1.10	“ <code>hg qnext</code> ”—print the name of the next patch	174
B.1.11	“ <code>hg qpop</code> ”—pop patches off the stack	175
B.1.12	“ <code>hg qprev</code> ”—print the name of the previous patch	175
B.1.13	“ <code>hg qpush</code> ”—push patches onto the stack	175
B.1.14	“ <code>hg qrefresh</code> ”—update the topmost applied patch	176
B.1.15	“ <code>hg qrename</code> ”—rename a patch	176
B.1.16	“ <code>hg qrestore</code> ”—restore saved queue state	176
B.1.17	“ <code>hg qsave</code> ”—save current queue state	176
B.1.18	“ <code>hg qseries</code> ”—print the entire patch series	176
B.1.19	“ <code>hg qtop</code> ”—print the name of the current patch	176
B.1.20	“ <code>hg qunapplied</code> ”—print patches not yet applied	176
B.1.21	“ <code>hg qversion</code> ”	177
B.1.22	“ <code>hg strip</code> ”—remove a revision and descendants	177
B.2	MQ file reference	177
B.2.1	The <code>series</code> file	177
B.2.2	The <code>status</code> file	177
C	Installing Mercurial from source	178
C.1	On a Unix-like system	178
C.2	On Windows	178

D Open Publication License	179
D.1 Requirements on both unmodified and modified versions	179
D.2 Copyright	179
D.3 Scope of license	179
D.4 Requirements on modified works	180
D.5 Good-practice recommendations	180
D.6 License options	180
Bibliography	181
Index	181

List of Figures

2.1	Graphical history of the <code>hello</code> repository	14
3.1	Divergent recent histories of the <code>my-hello</code> and <code>my-new-hello</code> repositories	26
3.2	Repository contents after pulling from <code>my-hello</code> into <code>my-new-hello</code>	27
3.3	Working directory and repository during merge, and following commit	28
3.4	Conflicting changes to a document	30
3.5	Using <code>kdiff3</code> to merge versions of a file	34
4.1	Relationships between files in working directory and filelogs in repository	36
4.2	Metadata relationships	36
4.3	Snapshot of a revlog, with incremental deltas	37
4.4		43
4.5	The working directory can have two parents	44
4.6	The working directory gains new parents after a commit	44
4.7	The working directory, updated to an older changeset	45
4.8	After a commit made while synced to an older changeset	45
4.9	Merging two heads	46
5.1	Simulating an empty directory using a hidden file	48
6.1	Feature branches	61
9.1	Backing out a change using the “ <code>hg backout</code> ” command	95
9.2	Automated backout of a non-tip change using the “ <code>hg backout</code> ” command	96
9.3	Backing out a change using the “ <code>hg backout</code> ” command	97
9.4	Manually merging a backout change	99
10.1	A simple hook that runs when a changeset is committed	109
10.2	Defining a second <code>commit</code> hook	109
10.3	Using the <code>pretxncommit</code> hook to control commits	110
10.4	A hook that forbids overly short commit messages	112
10.5	A simple hook that checks for trailing whitespace	113
10.6	A better trailing whitespace hook	114
11.1	Template keywords in use	129
11.2	Template filters in action	132
12.1	Simple uses of the <code>diff</code> and <code>patch</code> commands	138
12.2	Contents to add to <code>~/.hgrc</code> to enable the MQ extension	139
12.3	How to verify that MQ is enabled	139
12.4	Preparing a repository for use with MQ	139
12.5	Creating a new patch	140

12.6 Refreshing a patch	141
12.7 Refresh a patch many times to accumulate changes	142
12.8 Stacking a second patch on top of the first	143
12.9 Understanding the patch stack with “hg qseries” and “hg qapplied”	144
12.10 Applied and unapplied patches in the MQ patch stack	144
12.11 Modifying the stack of applied patches	145
12.12 Pushing all unapplied patches	145
12.13 Forcibly creating a patch	146
12.14 Using MQ’s tag features to work with patches	149
12.15 The diffstat, filterdiff, and lsdiff commands	151

Preface

Distributed revision control is a relatively new territory, and has thus far grown due to people’s willingness to strike out into ill-charted territory.

I am writing a book about distributed revision control because I believe that it is an important subject that deserves a field guide. I chose to write about Mercurial because it is the easiest tool to learn the terrain with, and yet it scales to the demands of real, challenging environments where many other revision control tools fail.

0.1 This book is a work in progress

I am releasing this book while I am still writing it, in the hope that it will prove useful to others. I also hope that readers will contribute as they see fit.

0.2 About the examples in this book

This book takes an unusual approach to code samples. Every example is “live”—each one is actually the result of a shell script that executes the Mercurial commands you see. Every time an image of the book is built from its sources, all the example scripts are automatically run, and their current results compared against their expected results.

The advantage of this approach is that the examples are always accurate; they describe *exactly* the behaviour of the version of Mercurial that’s mentioned at the front of the book. If I update the version of Mercurial that I’m documenting, and the output of some command changes, the build fails.

There is a small disadvantage to this approach, which is that the dates and times you’ll see in examples tend to be “squashed” together in a way that they wouldn’t be if the same commands were being typed by a human. Where a human can issue no more than one command every few seconds, with any resulting timestamps correspondingly spread out, my automated example scripts run many commands in one second.

As an instance of this, several consecutive commits in an example can show up as having occurred during the same second. You can see this occur in the `bisect` example in section 9.5, for instance.

So when you’re reading examples, don’t place too much weight on the dates or times you see in the output of commands. But *do* be confident that the behaviour you’re seeing is consistent and reproducible.

0.3 Colophon—this book is Free

This book is licensed under the Open Publication License, and is produced entirely using Free Software tools. It is typeset with L^AT_EX; illustrations are drawn and rendered with **Inkscape**.

The complete source code for this book is published as a Mercurial repository, at <http://hg.serpentine.com/mercurial/book>.

Chapter 1

Introduction

1.1 About revision control

Revision control is the process of managing multiple versions of a piece of information. In its simplest form, this is something that many people do by hand: every time you modify a file, save it under a new name that contains a number, each one higher than the number of the preceding version.

Manually managing multiple versions of even a single file is an error-prone task, though, so software tools to help automate this process have long been available. The earliest automated revision control tools were intended to help a single user to manage revisions of a single file. Over the past few decades, the scope of revision control tools has expanded greatly; they now manage multiple files, and help multiple people to work together. The best modern revision control tools have no problem coping with thousands of people working together on projects that consist of hundreds of thousands of files.

1.1.1 Why use revision control?

There are a number of reasons why you or your team might want to use an automated revision control tool for a project.

- It will track the history and evolution of your project, so you don't have to. For every change, you'll have a log of *who* made it; *why* they made it; *when* they made it; and *what* the change was.
- When you're working with other people, revision control software makes it easier for you to collaborate. For example, when people more or less simultaneously make potentially incompatible changes, the software will help you to identify and resolve those conflicts.
- It can help you to recover from mistakes. If you make a change that later turns out to be in error, you can revert to an earlier version of one or more files. In fact, a *really* good revision control tool will even help you to efficiently figure out exactly when a problem was introduced (see section 9.5 for details).
- It will help you to work simultaneously on, and manage the drift between, multiple versions of your project.

Most of these reasons are equally valid—at least in theory—whether you're working on a project by yourself, or with a hundred other people.

A key question about the practicality of revision control at these two different scales (“lone hacker” and “huge team”) is how its *benefits* compare to its *costs*. A revision control tool that's difficult to understand or use is going to impose a high cost.

A five-hundred-person project is likely to collapse under its own weight almost immediately without a revision control tool and process. In this case, the cost of using revision control might hardly seem worth considering, since *without* it, failure is almost guaranteed.

On the other hand, a one-person “quick hack” might seem like a poor place to use a revision control tool, because surely the cost of using one must be close to the overall cost of the project. Right?

Mercurial uniquely supports *both* of these scales of development. You can learn the basics in just a few minutes, and due to its low overhead, you can apply revision control to the smallest of projects with ease. Its simplicity means you won't have a lot of abstruse concepts or command sequences competing for mental space with whatever you're *really* trying to do. At the same time, Mercurial's high performance and peer-to-peer nature let you scale painlessly to handle large projects.

No revision control tool can rescue a poorly run project, but a good choice of tools can make a huge difference to the fluidity with which you can work on a project.

1.1.2 The many names of revision control

Revision control is a diverse field, so much so that it doesn't actually have a single name or acronym. Here are a few of the more common names and acronyms you'll encounter:

- Revision control (RCS)
- Software configuration management (SCM), or configuration management
- Source code management
- Source code control, or source control
- Version control (VCS)

Some people claim that these terms actually have different meanings, but in practice they overlap so much that there's no agreed or even useful way to tease them apart.

1.2 A short history of revision control

The best known of the old-time revision control tools is SCCS (Source Code Control System), which Marc Rochkind wrote at Bell Labs, in the early 1970s. SCCS operated on individual files, and required every person working on a project to have access to a shared workspace on a single system. Only one person could modify a file at any time; arbitration for access to files was via locks. It was common for people to lock files, and later forget to unlock them, preventing anyone else from modifying those files without the help of an administrator.

Walter Tichy developed a free alternative to SCCS in the early 1980s; he called his program RCS (Revision Control System). Like SCCS, RCS required developers to work in a single shared workspace, and to lock files to prevent multiple people from modifying them simultaneously.

Later in the 1980s, Dick Grune used RCS as a building block for a set of shell scripts he initially called cmt, but then renamed to CVS (Concurrent Versions System). The big innovation of CVS was that it let developers work simultaneously and somewhat independently in their own personal workspaces. The personal workspaces prevented developers from stepping on each other's toes all the time, as was common with SCCS and RCS. Each developer had a copy of every project file, and could modify their copies independently. They had to merge their edits prior to committing changes to the central repository.

Brian Berliner took Grune's original scripts and rewrote them in C, releasing in 1989 the code that has since developed into the modern version of CVS. CVS subsequently acquired the ability to operate over a network connection, giving it a client/server architecture. CVS's architecture is centralised; only the server has a copy of the history of the project. Client workspaces just contain copies of recent versions of the project's files, and a little metadata to tell them where the server is. CVS has been enormously successful; it is probably the world's most widely used revision control system.

In the early 1990s, Sun Microsystems developed an early distributed revision control system, called TeamWare. A TeamWare workspace contains a complete copy of the project's history. TeamWare has no notion of a central repository. (CVS relied upon RCS for its history storage; TeamWare used SCCS.)

As the 1990s progressed, awareness grew of a number of problems with CVS. It records simultaneous changes to multiple files individually, instead of grouping them together as a single logically atomic operation. It does not manage

its file hierarchy well; it is easy to make a mess of a repository by renaming files and directories. Worse, its source code is difficult to read and maintain, which made the “pain level” of fixing these architectural problems prohibitive.

In 2001, Jim Blandy and Karl Fogel, two developers who had worked on CVS, started a project to replace it with a tool that would have a better architecture and cleaner code. The result, Subversion, does not stray from CVS’s centralised client/server model, but it adds multi-file atomic commits, better namespace management, and a number of other features that make it a generally better tool than CVS. Since its initial release, it has rapidly grown in popularity.

More or less simultaneously, Graydon Hoare began working on an ambitious distributed revision control system that he named Monotone. While Monotone addresses many of CVS’s design flaws and has a peer-to-peer architecture, it goes beyond earlier (and subsequent) revision control tools in a number of innovative ways. It uses cryptographic hashes as identifiers, and has an integral notion of “trust” for code from different sources.

Mercurial began life in 2005. While a few aspects of its design are influenced by Monotone, Mercurial focuses on ease of use, high performance, and scalability to very large projects.

1.3 Trends in revision control

There has been an unmistakable trend in the development and use of revision control tools over the past four decades, as people have become familiar with the capabilities of their tools and constrained by their limitations.

The first generation began by managing single files on individual computers. Although these tools represented a huge advance over ad-hoc manual revision control, their locking model and reliance on a single computer limited them to small, tightly-knit teams.

The second generation loosened these constraints by moving to network-centered architectures, and managing entire projects at a time. As projects grew larger, they ran into new problems. With clients needing to talk to servers very frequently, server scaling became an issue for large projects. An unreliable network connection could prevent remote users from being able to talk to the server at all. As open source projects started making read-only access available anonymously to anyone, people without commit privileges found that they could not use the tools to interact with a project in a natural way, as they could not record their changes.

The current generation of revision control tools is peer-to-peer in nature. All of these systems have dropped the dependency on a single central server, and allow people to distribute their revision control data to where it’s actually needed. Collaboration over the Internet has moved from constrained by technology to a matter of choice and consensus. Modern tools can operate offline indefinitely and autonomously, with a network connection only needed when syncing changes with another repository.

1.4 A few of the advantages of distributed revision control

Even though distributed revision control tools have for several years been as robust and usable as their previous-generation counterparts, people using older tools have not yet necessarily woken up to their advantages. There are a number of ways in which distributed tools shine relative to centralised ones.

For an individual developer, distributed tools are almost always much faster than centralised tools. This is for a simple reason: a centralised tool needs to talk over the network for many common operations, because most metadata is stored in a single copy on the central server. A distributed tool stores all of its metadata locally. All else being equal, talking over the network adds overhead to a centralised tool. Don’t underestimate the value of a snappy, responsive tool: you’re going to spend a lot of time interacting with your revision control software.

Distributed tools are indifferent to the vagaries of your server infrastructure, again because they replicate metadata to so many locations. If you use a centralised system and your server catches fire, you’d better hope that your backup media are reliable, and that your last backup was recent and actually worked. With a distributed tool, you have many backups available on every contributor’s computer.

The reliability of your network will affect distributed tools far less than it will centralised tools. You can’t even use a centralised tool without a network connection, except for a few highly constrained commands. With a distributed tool, if your network connection goes down while you’re working, you may not even notice. The only thing you won’t be able to do is talk to repositories on other computers, something that is relatively rare compared with local operations. If you have a far-flung team of collaborators, this may be significant.

1.4.1 Advantages for open source projects

If you take a shine to an open source project and decide that you would like to start hacking on it, and that project uses a distributed revision control tool, you are at once a peer with the people who consider themselves the “core” of that project. If they publish their repositories, you can immediately copy their project history, start making changes, and record your work, using the same tools in the same ways as insiders. By contrast, with a centralised tool, you must use the software in a “read only” mode unless someone grants you permission to commit changes to their central server. Until then, you won’t be able to record changes, and your local modifications will be at risk of corruption any time you try to update your client’s view of the repository.

The forking non-problem

It has been suggested that distributed revision control tools pose some sort of risk to open source projects because they make it easy to “fork” the development of a project. A fork happens when there are differences in opinion or attitude between groups of developers that cause them to decide that they can’t work together any longer. Each side takes a more or less complete copy of the project’s source code, and goes off in its own direction.

Sometimes the camps in a fork decide to reconcile their differences. With a centralised revision control system, the *technical* process of reconciliation is painful, and has to be performed largely by hand. You have to decide whose revision history is going to “win”, and graft the other team’s changes into the tree somehow. This usually loses some or all of one side’s revision history.

What distributed tools do with respect to forking is they make forking the *only* way to develop a project. Every single change that you make is potentially a fork point. The great strength of this approach is that a distributed revision control tool has to be really good at *merging* forks, because forks are absolutely fundamental: they happen all the time.

If every piece of work that everybody does, all the time, is framed in terms of forking and merging, then what the open source world refers to as a “fork” becomes *purely* a social issue. If anything, distributed tools *lower* the likelihood of a fork:

- They eliminate the social distinction that centralised tools impose: that between insiders (people with commit access) and outsiders (people without).
- They make it easier to reconcile after a social fork, because all that’s involved from the perspective of the revision control software is just another merge.

Some people resist distributed tools because they want to retain tight control over their projects, and they believe that centralised tools give them this control. However, if you’re of this belief, and you publish your CVS or Subversion repositories publically, there are plenty of tools available that can pull out your entire project’s history (albeit slowly) and recreate it somewhere that you don’t control. So while your control in this case is illusory, you are forgoing the ability to fluidly collaborate with whatever people feel compelled to mirror and fork your history.

1.4.2 Advantages for commercial projects

Many commercial projects are undertaken by teams that are scattered across the globe. Contributors who are far from a central server will see slower command execution and perhaps less reliability. Commercial revision control systems attempt to ameliorate these problems with remote-site replication add-ons that are typically expensive to buy and cantankerous to administer. A distributed system doesn’t suffer from these problems in the first place. Better yet, you can easily set up multiple authoritative servers, say one per site, so that there’s no redundant communication between repositories over expensive long-haul network links.

Centralised revision control systems tend to have relatively low scalability. It’s not unusual for an expensive centralised system to fall over under the combined load of just a few dozen concurrent users. Once again, the typical response tends to be an expensive and clunky replication facility. Since the load on a central server—if you have one at all—is many times lower with a distributed tool (because all of the data is replicated everywhere), a single cheap server can handle the needs of a much larger team, and replication to balance load becomes a simple matter of scripting.

If you have an employee in the field, troubleshooting a problem at a customer’s site, they’ll benefit from distributed revision control. The tool will let them generate custom builds, try different fixes in isolation from each other, and

search efficiently through history for the sources of bugs and regressions in the customer's environment, all without needing to connect to your company's network.

1.5 Why choose Mercurial?

Mercurial has a unique set of properties that make it a particularly good choice as a revision control system.

- It is easy to learn and use.
- It is lightweight.
- It scales excellently.
- It is easy to customise.

If you are at all familiar with revision control systems, you should be able to get up and running with Mercurial in less than five minutes. Even if not, it will take no more than a few minutes longer. Mercurial's command and feature sets are generally uniform and consistent, so you can keep track of a few general rules instead of a host of exceptions.

On a small project, you can start working with Mercurial in moments. Creating new changes and branches; transferring changes around (whether locally or over a network); and history and status operations are all fast. Mercurial attempts to stay nimble and largely out of your way by combining low cognitive overhead with blazingly fast operations.

The usefulness of Mercurial is not limited to small projects: it is used by projects with hundreds to thousands of contributors, each containing tens of thousands of files and hundreds of megabytes of source code.

If the core functionality of Mercurial is not enough for you, it's easy to build on. Mercurial is well suited to scripting tasks, and its clean internals and implementation in Python make it easy to add features in the form of extensions. There are a number of popular and useful extensions already available, ranging from helping to identify bugs to improving performance.

1.6 Mercurial compared with other tools

Before you read on, please understand that this section necessarily reflects my own experiences, interests, and (dare I say it) biases. I have used every one of the revision control tools listed below, in most cases for several years at a time.

1.6.1 Subversion

Subversion is a popular revision control tool, developed to replace CVS. It has a centralised client/server architecture.

Subversion and Mercurial have similarly named commands for performing the same operations, so if you're familiar with one, it is easy to learn to use the other. Both tools are portable to all popular operating systems.

Subversion lacks a history-aware merge capability, forcing its users to manually track exactly which revisions have been merged between branches. If users fail to do this, or make mistakes, they face the prospect of manually resolving merges with unnecessary conflicts. Subversion also fails to merge changes when files or directories are renamed. Subversion's poor merge support is its single biggest weakness.

Mercurial has a substantial performance advantage over Subversion on every revision control operation I have benchmarked. I have measured its advantage as ranging from a factor of two to a factor of six when compared with Subversion 1.4.3's *ra_local* file store, which is the fastest access method available). In more realistic deployments involving a network-based store, Subversion will be at a substantially larger disadvantage. Because many Subversion commands must talk to the server and Subversion does not have useful replication facilities, server capacity and network bandwidth become bottlenecks for modestly large projects.

Additionally, Subversion incurs substantial storage overhead to avoid network transactions for a few common operations, such as finding modified files (`status`) and displaying modifications against the current revision (`diff`). As a result, a Subversion working copy is often the same size as, or larger than, a Mercurial repository and working directory, even though the Mercurial repository contains a complete history of the project.

Subversion is widely supported by third party tools. Mercurial currently lags considerably in this area. This gap is closing, however, and indeed some of Mercurial's GUI tools now outshine their Subversion equivalents. Like Mercurial, Subversion has an excellent user manual.

Because Subversion doesn't store revision history on the client, it is well suited to managing projects that deal with lots of large, opaque binary files. If you check in fifty revisions to an incompressible 10MB file, Subversion's client-side space usage stays constant. The space used by any distributed SCM will grow rapidly in proportion to the number of revisions, because the differences between each revision are large.

In addition, it's often difficult or, more usually, impossible to merge different versions of a binary file. Subversion's ability to let a user lock a file, so that they temporarily have the exclusive right to commit changes to it, can be a significant advantage to a project where binary files are widely used.

Mercurial can import revision history from a Subversion repository. It can also export revision history to a Subversion repository. This makes it easy to "test the waters" and use Mercurial and Subversion in parallel before deciding to switch. History conversion is incremental, so you can perform an initial conversion, then small additional conversions afterwards to bring in new changes.

1.6.2 Git

Git is a distributed revision control tool that was developed for managing the Linux kernel source tree. Like Mercurial, its early design was somewhat influenced by Monotone.

Git has a very large command set, with version 1.5.0 providing 139 individual commands. It has something of a reputation for being difficult to learn. Compared to Git, Mercurial has a strong focus on simplicity.

In terms of performance, Git is extremely fast. In several cases, it is faster than Mercurial, at least on Linux, while Mercurial performs better on other operations. However, on Windows, the performance and general level of support that Git provides is, at the time of writing, far behind that of Mercurial.

While a Mercurial repository needs no maintenance, a Git repository requires frequent manual "repacks" of its metadata. Without these, performance degrades, while space usage grows rapidly. A server that contains many Git repositories that are not rigorously and frequently repacked will become heavily disk-bound during backups, and there have been instances of daily backups taking far longer than 24 hours as a result. A freshly packed Git repository is slightly smaller than a Mercurial repository, but an unpacked repository is several orders of magnitude larger.

The core of Git is written in C. Many Git commands are implemented as shell or Perl scripts, and the quality of these scripts varies widely. I have encountered several instances where scripts charged along blindly in the presence of errors that should have been fatal.

Mercurial can import revision history from a Git repository.

1.6.3 CVS

CVS is probably the most widely used revision control tool in the world. Due to its age and internal untidiness, it has been only lightly maintained for many years.

It has a centralised client/server architecture. It does not group related file changes into atomic commits, making it easy for people to "break the build": one person can successfully commit part of a change and then be blocked by the need for a merge, causing other people to see only a portion of the work they intended to do. This also affects how you work with project history. If you want to see all of the modifications someone made as part of a task, you will need to manually inspect the descriptions and timestamps of the changes made to each file involved (if you even know what those files were).

CVS has a muddled notion of tags and branches that I will not attempt to even describe. It does not support renaming of files or directories well, making it easy to corrupt a repository. It has almost no internal consistency checking capabilities, so it is usually not even possible to tell whether or how a repository is corrupt. I would not recommend CVS for any project, existing or new.

Mercurial can import CVS revision history. However, there are a few caveats that apply; these are true of every other revision control tool's CVS importer, too. Due to CVS's lack of atomic changes and unversioned filesystem hierarchy, it is not possible to reconstruct CVS history completely accurately; some guesswork is involved, and re-names will usually not show up. Because a lot of advanced CVS administration has to be done by hand and is hence

error-prone, it's common for CVS importers to run into multiple problems with corrupted repositories (completely bogus revision timestamps and files that have remained locked for over a decade are just two of the less interesting problems I can recall from personal experience).

Mercurial can import revision history from a CVS repository.

1.6.4 Commercial tools

Perforce has a centralised client/server architecture, with no client-side caching of any data. Unlike modern revision control tools, Perforce requires that a user run a command to inform the server about every file they intend to edit.

The performance of Perforce is quite good for small teams, but it falls off rapidly as the number of users grows beyond a few dozen. Modestly large Perforce installations require the deployment of proxies to cope with the load their users generate.

1.6.5 Choosing a revision control tool

With the exception of CVS, all of the tools listed above have unique strengths that suit them to particular styles of work. There is no single revision control tool that is best in all situations.

As an example, Subversion is a good choice for working with frequently edited binary files, due to its centralised nature and support for file locking. If you're averse to the command line, it currently has better GUI support than other free revision control tools. However, its poor merging is a substantial liability for busy projects with overlapping development.

I personally find Mercurial's properties of simplicity, performance, and good merge support to be a compelling combination that has served me well for several years.

1.7 Switching from another tool to Mercurial

Mercurial is bundled with an extension named `convert`, which can incrementally import revision history from several other revision control tools. By “incremental”, I mean that you can convert all of a project's history to date in one go, then rerun the conversion later to obtain new changes that happened after the initial conversion.

The revision control tools supported by `convert` are as follows:

- Subversion
- CVS
- Git
- Darcs

In addition, `convert` can export changes from Mercurial to Subversion. This makes it possible to try Subversion and Mercurial in parallel before committing to a switchover, without risking the loss of any work.

The “`hg convert`” command is easy to use. Simply point it at the path or URL of the source repository, optionally give it the name of the destination repository, and it will start working. After the initial conversion, just run the same command again to import new changes.

Chapter 2

A tour of Mercurial: the basics

2.1 Installing Mercurial on your system

Prebuilt binary packages of Mercurial are available for every popular operating system. These make it easy to start using Mercurial on your computer immediately.

2.1.1 Linux

Because each Linux distribution has its own packaging tools, policies, and rate of development, it's difficult to give a comprehensive set of instructions on how to install Mercurial binaries. The version of Mercurial that you will end up with can vary depending on how active the person is who maintains the package for your distribution.

To keep things simple, I will focus on installing Mercurial from the command line under the most popular Linux distributions. Most of these distributions provide graphical package managers that will let you install Mercurial with a single click; the package name to look for is `mercurial`.

Debian `apt-get install mercurial`

Fedora Core `yum install mercurial`

Gentoo `emerge mercurial`

OpenSUSE `yum install mercurial`

Ubuntu Ubuntu's Mercurial package is based on Debian's. To install it, run the following command.

```
1 apt-get install mercurial
```

The Ubuntu package for Mercurial tends to lag behind the Debian version by a considerable time margin (at the time of writing, seven months), which in some cases will mean that on Ubuntu, you may run into problems that have since been fixed in the Debian package.

2.1.2 Solaris

XXX.

2.1.3 Mac OS X

Lee Cantey publishes an installer of Mercurial for Mac OS X at <http://mercurial.berkwood.com>. This package works on both Intel- and Power-based Macs. Before you can use it, you must install a compatible version of Universal MacPython [BI]. This is easy to do; simply follow the instructions on Lee’s site.

2.1.4 Windows

Lee Cantey also publishes an installer of Mercurial for Windows at <http://mercurial.berkwood.com>. This package has no external dependencies; it “just works”.

Note: The Windows version of Mercurial does not automatically convert line endings between Windows and Unix styles. If you want to share work with Unix users, you must do a little additional configuration work. XXX Flesh this out.

2.2 Getting started

To begin, we’ll use the “hg version” command to find out whether Mercurial is actually installed properly. The actual version information that it prints isn’t so important; it’s whether it prints anything at all that we care about.

```
1 $ hg version
2 Mercurial Distributed SCM (version fe2e81229819)
3
4 Copyright (C) 2005–2007 Matt Mackall <mpm@selenic.com> and others
5 This is free software; see the source for copying conditions. There is NO
6 warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.
```

2.2.1 Built-in help

Mercurial provides a built-in help system. This is invaluable for those times when you find yourself stuck trying to remember how to run a command. If you are completely stuck, simply run “hg help”; it will print a brief list of commands, along with a description of what each does. If you ask for help on a specific command (as below), it prints more detailed information.

```
1 $ hg help init
2 hg init [-e CMD] [--remotecmd CMD] [DEST]
3
4 create a new repository in the given directory
5
6     Initialize a new repository in the given directory.  If the given
7     directory does not exist, it is created.
8
9     If no directory is given, the current directory is used.
10
11     It is possible to specify an ssh:// URL as the destination.
12     Look at the help text for the pull command for important details
13     about ssh:// URLs.
14
15 options:
16
17 -e --ssh          specify ssh command to use
18 --remotecmd      specify hg command to run on the remote side
```

```
19
20 use "hg -v help init" to show global options
```

For a more impressive level of detail (which you won't usually need) run “hg help -v”. The -v option is short for --verbose, and tells Mercurial to print more information than it usually would.

2.3 Working with a repository

In Mercurial, everything happens inside a *repository*. The repository for a project contains all of the files that “belong to” that project, along with a historical record of the project's files.

There's nothing particularly magical about a repository; it is simply a directory tree in your filesystem that Mercurial treats as special. You can rename or delete a repository any time you like, using either the command line or your file browser.

2.3.1 Making a local copy of a repository

Copying a repository is just a little bit special. While you could use a normal file copying command to make a copy of a repository, it's best to use a built-in command that Mercurial provides. This command is called “hg clone”, because it creates an identical copy of an existing repository.

```
1 $ hg clone http://hg.serpentine.com/tutorial/hello
2 destination directory: hello
3 requesting all changes
4 adding changesets
5 adding manifests
6 adding file changes
7 added 5 changesets with 5 changes to 2 files
8 2 files updated, 0 files merged, 0 files removed, 0 files unresolved
```

If our clone succeeded, we should now have a local directory called `hello`. This directory will contain some files.

```
1 $ ls -l
2 total 4
3 drwxrwxr-x 3 bos 1000 4096 Dec 10 19:55 hello
4 $ ls hello
5 Makefile hello.c
```

These files have the same contents and history in our repository as they do in the repository we cloned.

Every Mercurial repository is complete, self-contained, and independent. It contains its own private copy of a project's files and history. A cloned repository remembers the location of the repository it was cloned from, but it does not communicate with that repository, or any other, unless you tell it to.

What this means for now is that we're free to experiment with our repository, safe in the knowledge that it's a private “sandbox” that won't affect anyone else.

2.3.2 What's in a repository?

When we take a more detailed look inside a repository, we can see that it contains a directory named `.hg`. This is where Mercurial keeps all of its metadata for the repository.

```
1 $ cd hello
2 $ ls -a
3 .  ..  .hg  Makefile  hello.c
```

The contents of the `.hg` directory and its subdirectories are private to Mercurial. Every other file and directory in the repository is yours to do with as you please.

To introduce a little terminology, the `.hg` directory is the “real” repository, and all of the files and directories that coexist with it are said to live in the *working directory*. An easy way to remember the distinction is that the *repository* contains the *history* of your project, while the *working directory* contains a *snapshot* of your project at a particular point in history.

2.4 A tour through history

One of the first things we might want to do with a new, unfamiliar repository is understand its history. The “`hg log`” command gives us a view of history.

```
1 $ hg log
2 changeset: 4:b57f9a090b62
3 tag:      tip
4 user:     Bryan O'Sullivan <bos@serpentine.com>
5 date:     Tue Sep 06 15:43:07 2005 -0700
6 summary:  Trim comments.
7
8 changeset: 3:ff5d7b70a2a9
9 user:     Bryan O'Sullivan <bos@serpentine.com>
10 date:    Tue Sep 06 13:15:58 2005 -0700
11 summary:  Get make to generate the final binary from a .o file.
12
13 changeset: 2:057d3c2d823c
14 user:     Bryan O'Sullivan <bos@serpentine.com>
15 date:    Tue Sep 06 13:15:43 2005 -0700
16 summary:  Introduce a typo into hello.c.
17
18 changeset: 1:82e55d328c8c
19 user:     mpm@selenic.com
20 date:    Fri Aug 26 01:21:28 2005 -0700
21 summary:  Create a makefile
22
23 changeset: 0:0a04b987be5a
24 user:     mpm@selenic.com
25 date:    Fri Aug 26 01:20:50 2005 -0700
26 summary:  Create a standard "hello, world" program
27
```

By default, this command prints a brief paragraph of output for each change to the project that was recorded. In Mercurial terminology, we call each of these recorded events a *changeset*, because it can contain a record of changes to several files.

The fields in a record of output from “`hg log`” are as follows.

changeset This field has the format of a number, followed by a colon, followed by a hexadecimal string. These are *identifiers* for the changeset. There are two identifiers because the number is shorter and easier to type than the hex string.

user The identity of the person who created the changeset. This is a free-form field, but it most often contains a person’s name and email address.

date The date and time on which the changeset was created, and the timezone in which it was created. (The date and time are local to that timezone; they display what time and date it was for the person who created the changeset.)

summary The first line of the text message that the creator of the changeset entered to describe the changeset.

The default output printed by “hg log” is purely a summary; it is missing a lot of detail.

Figure 2.1 provides a graphical representation of the history of the `hello` repository, to make it a little easier to see which direction history is “flowing” in. We’ll be returning to this figure several times in this chapter and the chapter that follows.




Figure 2.1: Graphical history of the `hello` repository

2.4.1 Changesets, revisions, and talking to other people

As English is a notoriously sloppy language, and computer science has a hallowed history of terminological confusion (why use one term when four will do?), revision control has a variety of words and phrases that mean the same thing. If you are talking about Mercurial history with other people, you will find that the word “changeset” is often compressed to “change” or (when written) “cset”, and sometimes a changeset is referred to as a “revision” or a “rev”.

While it doesn’t matter what *word* you use to refer to the concept of “a changeset”, the *identifier* that you use to refer to “a *specific* changeset” is of great importance. Recall that the `changeset` field in the output from “hg log” identifies a changeset using both a number and a hexadecimal string.

- The revision number is *only valid in that repository*,
- while the hex string is the *permanent, unchanging identifier* that will always identify that exact changeset in *every* copy of the repository.

This distinction is important. If you send someone an email talking about “revision 33”, there’s a high likelihood that their revision 33 will *not be the same* as yours. The reason for this is that a revision number depends on the order in which changes arrived in a repository, and there is no guarantee that the same changes will happen in the same order in different repositories. Three changes *a, b, c* can easily appear in one repository as 0, 1, 2, while in another as 1, 0, 2.

Mercurial uses revision numbers purely as a convenient shorthand. If you need to discuss a changeset with someone, or make a record of a changeset for some other reason (for example, in a bug report), use the hexadecimal identifier.

2.4.2 Viewing specific revisions

To narrow the output of “hg log” down to a single revision, use the `-r` (or `--rev`) option. You can use either a revision number or a long-form changeset identifier, and you can provide as many revisions as you want.

```

1 $ hg log -r 3
2 changeset: 3:ff5d7b70a2a9
3 user:      Bryan O'Sullivan <bos@serpentine.com>
4 date:      Tue Sep 06 13:15:58 2005 -0700
5 summary:    Get make to generate the final binary from a .o file.
6
7 $ hg log -r ff5d7b70a2a9
8 changeset: 3:ff5d7b70a2a9
9 user:      Bryan O'Sullivan <bos@serpentine.com>
10 date:      Tue Sep 06 13:15:58 2005 -0700
11 summary:    Get make to generate the final binary from a .o file.
12
13 $ hg log -r 1 -r 4
14 changeset: 1:82e55d328c8c
15 user:      mpm@selenic.com
16 date:      Fri Aug 26 01:21:28 2005 -0700
17 summary:    Create a makefile
18
19 changeset: 4:b57f9a090b62
20 tag:       tip
21 user:      Bryan O'Sullivan <bos@serpentine.com>
22 date:      Tue Sep 06 15:43:07 2005 -0700
23 summary:    Trim comments.
24

```

If you want to see the history of several revisions without having to list each one, you can use *range notation*; this lets you express the idea “I want all revisions between *a* and *b*, inclusive”.

```

1 $ hg log -r 2:4
2 changeset: 2:057d3c2d823c
3 user:      Bryan O'Sullivan <bos@serpentine.com>
4 date:      Tue Sep 06 13:15:43 2005 -0700
5 summary:    Introduce a typo into hello.c.
6
7 changeset: 3:ff5d7b70a2a9
8 user:      Bryan O'Sullivan <bos@serpentine.com>
9 date:      Tue Sep 06 13:15:58 2005 -0700
10 summary:    Get make to generate the final binary from a .o file.
11
12 changeset: 4:b57f9a090b62
13 tag:       tip
14 user:      Bryan O'Sullivan <bos@serpentine.com>
15 date:      Tue Sep 06 15:43:07 2005 -0700
16 summary:    Trim comments.
17

```

Mercurial also honours the order in which you specify revisions, so “hg log -r 2:4” prints 2,3,4 while “hg log -r 4:2” prints 4,3,2.

2.4.3 More detailed information

While the summary information printed by “hg log” is useful if you already know what you’re looking for, you may need to see a complete description of the change, or a list of the files changed, if you’re trying to decide whether a changeset is the one you’re looking for. The “hg log” command’s `-v` (or `--verbose`) option gives you this extra detail.

```
1 $ hg log -v -r 3
2 changeset: 3:ff5d7b70a2a9
3 user:      Bryan O'Sullivan <bos@serpentine.com>
4 date:      Tue Sep 06 13:15:58 2005 -0700
5 files:     Makefile
6 description:
7 Get make to generate the final binary from a .o file.
8
9
```

If you want to see both the description and content of a change, add the `-p` (or `--patch`) option. This displays the content of a change as a *unified diff* (if you’ve never seen a unified diff before, see section 12.4 for an overview).

```
1 $ hg log -v -p -r 2
2 changeset: 2:057d3c2d823c
3 user:      Bryan O'Sullivan <bos@serpentine.com>
4 date:      Tue Sep 06 13:15:43 2005 -0700
5 files:     hello.c
6 description:
7 Introduce a typo into hello.c.
8
9
10 diff -r 82e55d328c8c -r 057d3c2d823c hello.c
11 --- a/hello.c      Fri Aug 26 01:21:28 2005 -0700
12 +++ b/hello.c      Tue Sep 06 13:15:43 2005 -0700
13 @@ -11,6 +11,6 @@
14
15  int main(int argc, char **argv)
16  {
17  -     printf("hello, world!\n");
18  +     printf("hello, world!");
19     return 0;
20 }
21
```

2.5 All about command options

Let’s take a brief break from exploring Mercurial commands to discuss a pattern in the way that they work; you may find this useful to keep in mind as we continue our tour.

Mercurial has a consistent and straightforward approach to dealing with the options that you can pass to commands. It follows the conventions for options that are common to modern Linux and Unix systems.

- Every option has a long name. For example, as we’ve already seen, the “hg log” command accepts a `--rev` option.

- Most options have short names, too. Instead of `--rev`, we can use `-r`. (The reason that some options don't have short names is that the options in question are rarely used.)
- Long options start with two dashes (e.g. `--rev`), while short options start with one (e.g. `-r`).
- Option naming and usage is consistent across commands. For example, every command that lets you specify a changeset ID or revision number accepts both `-r` and `--rev` arguments.

In the examples throughout this book, I use short options instead of long. This just reflects my own preference, so don't read anything significant into it.

Most commands that print output of some kind will print more output when passed a `-v` (or `--verbose`) option, and less when passed `-q` (or `--quiet`).

2.6 Making and reviewing changes

Now that we have a grasp of viewing history in Mercurial, let's take a look at making some changes and examining them.

The first thing we'll do is isolate our experiment in a repository of its own. We use the `hg clone` command, but we don't need to clone a copy of the remote repository. Since we already have a copy of it locally, we can just clone that instead. This is much faster than cloning over the network, and cloning a local repository uses less disk space in most cases, too.

```
1 $ cd ..
2 $ hg clone hello my-hello
3 2 files updated, 0 files merged, 0 files removed, 0 files unresolved
4 $ cd my-hello
```

As an aside, it's often good practice to keep a "pristine" copy of a remote repository around, which you can then make temporary clones of to create sandboxes for each task you want to work on. This lets you work on multiple tasks in parallel, each isolated from the others until it's complete and you're ready to integrate it back. Because local clones are so cheap, there's almost no overhead to cloning and destroying repositories whenever you want.

In our `my-hello` repository, we have a file `hello.c` that contains the classic "hello, world" program. Let's use the ancient and venerable `sed` command to edit this file so that it prints a second line of output. (I'm only using `sed` to do this because it's easy to write a scripted example this way. Since you're not under the same constraint, you probably won't want to use `sed`; simply use your preferred text editor to do the same thing.)

```
1 $ sed -i '/printf/a\\tprintf("hello again!\\n");' hello.c
```

Mercurial's `hg status` command will tell us what Mercurial knows about the files in the repository.

```
1 $ ls
2 Makefile hello.c
3 $ hg status
4 M hello.c
```

The `hg status` command prints no output for some files, but a line starting with "M" for `hello.c`. Unless you tell it to, `hg status` will not print any output for files that have not been modified.

The "M" indicates that Mercurial has noticed that we modified `hello.c`. We didn't need to *inform* Mercurial that we were going to modify the file before we started, or that we had modified the file after we were done; it was able to figure this out itself.

It's a little bit helpful to know that we've modified `hello.c`, but we might prefer to know exactly *what* changes we've made to it. To do this, we use the `hg diff` command.

```

1 $ hg diff
2 diff -r b57f9a090b62 hello.c
3 --- a/hello.c      Tue Sep 06 15:43:07 2005 -0700
4 +++ b/hello.c      Mon Dec 10 19:55:40 2007 +0000
5 @@ -8,5 +8,6 @@ int main(int argc, char **argv)
6     int main(int argc, char **argv)
7     {
8         printf("hello, world!\n");
9 +     printf("hello again!\n");
10         return 0;
11     }

```

2.7 Recording changes in a new changeset

We can modify files, build and test our changes, and use “hg status” and “hg diff” to review our changes, until we’re satisfied with what we’ve done and arrive at a natural stopping point where we want to record our work in a new changeset.

The “hg commit” command lets us create a new changeset; we’ll usually refer to this as “making a commit” or “committing”.

2.7.1 Setting up a username

When you try to run “hg commit” for the first time, it is not guaranteed to succeed. Mercurial records your name and address with each change that you commit, so that you and others will later be able to tell who made each change. Mercurial tries to automatically figure out a sensible username to commit the change with. It will attempt each of the following methods, in order:

1. If you specify a `-u` option to the “hg commit” command on the command line, followed by a username, this is always given the highest precedence.
2. If you have set the `HGUSER` environment variable, this is checked next.
3. If you create a file in your home directory called `.hgrc`, with a `username` entry, that will be used next. To see what the contents of this file should look like, refer to section 2.7.1 below.
4. If you have set the `EMAIL` environment variable, this will be used next.
5. Mercurial will query your system to find out your local user name and host name, and construct a username from these components. Since this often results in a username that is not very useful, it will print a warning if it has to do this.

If all of these mechanisms fail, Mercurial will fail, printing an error message. In this case, it will not let you commit until you set up a username.

You should think of the `HGUSER` environment variable and the `-u` option to the “hg commit” command as ways to *override* Mercurial’s default selection of username. For normal use, the simplest and most robust way to set a username for yourself is by creating a `.hgrc` file; see below for details.

Creating a Mercurial configuration file

To set a user name, use your favourite editor to create a file called `.hgrc` in your home directory. Mercurial will use this file to look up your personalised configuration settings. The initial contents of your `.hgrc` should look like this.

```
1 # This is a Mercurial configuration file.
2 [ui]
3 username = Firstname Lastname <email.address@domain.net>
```

The “[ui]” line begins a *section* of the config file, so you can read the “username = ...” line as meaning “set the value of the username item in the ui section”. A section continues until a new section begins, or the end of the file. Mercurial ignores empty lines and treats any text from “#” to the end of a line as a comment.

Choosing a user name

You can use any text you like as the value of the username config item, since this information is for reading by other people, but for interpreting by Mercurial. The convention that most people follow is to use their name and email address, as in the example above.

Note: Mercurial’s built-in web server obfuscates email addresses, to make it more difficult for the email harvesting tools that spammers use. This reduces the likelihood that you’ll start receiving more junk email if you publish a Mercurial repository on the web.

2.7.2 Writing a commit message

When we commit a change, Mercurial drops us into a text editor, to enter a message that will describe the modifications we’ve made in this changeset. This is called the *commit message*. It will be a record for readers of what we did and why, and it will be printed by “hg log” after we’ve finished committing.

```
1 $ hg commit
```

The editor that the “hg commit” command drops us into will contain an empty line, followed by a number of lines starting with “HG:”.

```
1 empty line
2 HG: changed hello.c
```

Mercurial ignores the lines that start with “HG:”; it uses them only to tell us which files it’s recording changes to. Modifying or deleting these lines has no effect.

2.7.3 Writing a good commit message

Since “hg log” only prints the first line of a commit message by default, it’s best to write a commit message whose first line stands alone. Here’s a real example of a commit message that *doesn’t* follow this guideline, and hence has a summary that is not readable.

```
1 changeset: 73:584af0e231be
2 user:      Censored Person <censored.person@example.org>
3 date:      Tue Sep 26 21:37:07 2006 -0700
4 summary:   include buildmeister/commondefs.  Add an exports and install
```

As far as the remainder of the contents of the commit message are concerned, there are no hard-and-fast rules. Mercurial itself doesn’t interpret or care about the contents of the commit message, though your project may have policies that dictate a certain kind of formatting.

My personal preference is for short, but informative, commit messages that tell me something that I can’t figure out with a quick glance at the output of “hg log --patch”.

2.7.4 Aborting a commit

If you decide that you don't want to commit while in the middle of editing a commit message, simply exit from your editor without saving the file that it's editing. This will cause nothing to happen to either the repository or the working directory.

If we run the “`hg commit`” command without any arguments, it records all of the changes we've made, as reported by “`hg status`” and “`hg diff`”.

2.7.5 Admiring our new handiwork

Once we've finished the commit, we can use the “`hg tip`” command to display the changeset we just created. This command produces output that is identical to “`hg log`”, but it only displays the newest revision in the repository.

```
1 $ hg tip -vp
2 changeset: 5:fc920a075b40
3 tag:      tip
4 user:     Bryan O'Sullivan <bos@serpentine.com>
5 date:     Mon Dec 10 19:55:40 2007 +0000
6 files:    hello.c
7 description:
8 Added an extra line of output
9
10
11 diff -r b57f9a090b62 -r fc920a075b40 hello.c
12 --- a/hello.c      Tue Sep 06 15:43:07 2005 -0700
13 +++ b/hello.c      Mon Dec 10 19:55:40 2007 +0000
14 @@ -8,5 +8,6 @@ int main(int argc, char **argv)
15     int main(int argc, char **argv)
16     {
17         printf("hello, world!\n");
18 +       printf("hello again!\n");
19         return 0;
20     }
21
```

We refer to the newest revision in the repository as the tip revision, or simply the tip.

2.8 Sharing changes

We mentioned earlier that repositories in Mercurial are self-contained. This means that the changeset we just created exists only in our `my-hello` repository. Let's look at a few ways that we can propagate this change into other repositories.

2.8.1 Pulling changes from another repository

To get started, let's clone our original `hello` repository, which does not contain the change we just committed. We'll call our temporary repository `hello-pull`.

```
1 $ cd ..
2 $ hg clone hello hello-pull
3 2 files updated, 0 files merged, 0 files removed, 0 files unresolved
```

We'll use the “hg pull” command to bring changes from my-hello into hello-pull. However, blindly pulling unknown changes into a repository is a somewhat scary prospect. Mercurial provides the “hg incoming” command to tell us what changes the “hg pull” command *would* pull into the repository, without actually pulling the changes in.

```
1 $ cd hello-pull
2 $ hg incoming ../my-hello
3 comparing with ../my-hello
4 searching for changes
5 changeset: 5:fc920a075b40
6 tag: tip
7 user: Bryan O'Sullivan <bos@serpentine.com>
8 date: Mon Dec 10 19:55:40 2007 +0000
9 summary: Added an extra line of output
10
```

(Of course, someone could cause more changesets to appear in the repository that we ran “hg incoming” in, before we get a chance to “hg pull” the changes, so that we could end up pulling changes that we didn’t expect.)

Bringing changes into a repository is a simple matter of running the “hg pull” command, and telling it which repository to pull from.

```
1 $ hg tip
2 changeset: 4:b57f9a090b62
3 tag: tip
4 user: Bryan O'Sullivan <bos@serpentine.com>
5 date: Tue Sep 06 15:43:07 2005 -0700
6 summary: Trim comments.
7
8 $ hg pull ../my-hello
9 pulling from ../my-hello
10 searching for changes
11 adding changesets
12 adding manifests
13 adding file changes
14 added 1 changesets with 1 changes to 1 files
15 (run 'hg update' to get a working copy)
16 $ hg tip
17 changeset: 5:fc920a075b40
18 tag: tip
19 user: Bryan O'Sullivan <bos@serpentine.com>
20 date: Mon Dec 10 19:55:40 2007 +0000
21 summary: Added an extra line of output
22
```

As you can see from the before-and-after output of “hg tip”, we have successfully pulled changes into our repository. There remains one step before we can see these changes in the working directory.

2.8.2 Updating the working directory

We have so far glossed over the relationship between a repository and its working directory. The “hg pull” command that we ran in section 2.8.1 brought changes into the repository, but if we check, there’s no sign of those changes in the working directory. This is because “hg pull” does not (by default) touch the working directory. Instead, we use the “hg update” command to do this.

```

1 $ grep printf hello.c
2     printf("hello, world!\n");
3 $ hg update tip
4 1 files updated, 0 files merged, 0 files removed, 0 files unresolved
5 $ grep printf hello.c
6     printf("hello, world!\n");
7     printf("hello again!\n");

```

It might seem a bit strange that “hg pull” doesn’t update the working directory automatically. There’s actually a good reason for this: you can use “hg update” to update the working directory to the state it was in at *any revision* in the history of the repository. If you had the working directory updated to an old revision—to hunt down the origin of a bug, say—and ran a “hg pull” which automatically updated the working directory to a new revision, you might not be terribly happy.

However, since pull-then-update is such a common thing to do, Mercurial lets you combine the two by passing the `-u` option to “hg pull”.

```

1 hg pull -u

```

If you look back at the output of “hg pull” in section 2.8.1 when we ran it without `-u`, you can see that it printed a helpful reminder that we’d have to take an explicit step to update the working directory:

```

1 (run 'hg update' to get a working copy)

```

To find out what revision the working directory is at, use the “hg parents” command.

```

1 $ hg parents
2 changeset: 5:fc920a075b40
3 tag:      tip
4 user:     Bryan O'Sullivan <bos@serpentine.com>
5 date:     Mon Dec 10 19:55:40 2007 +0000
6 summary:  Added an extra line of output
7

```

If you look back at figure 2.1, you’ll see arrows connecting each changeset. The node that the arrow leads *from* in each case is a parent, and the node that the arrow leads *to* is its child. The working directory has a parent in just the same way; this is the changeset that the working directory currently contains.

To update the working directory to a particular revision, give a revision number or changeset ID to the “hg update” command.

```

1 $ hg update 2
2 2 files updated, 0 files merged, 0 files removed, 0 files unresolved
3 $ hg parents
4 changeset: 2:057d3c2d823c
5 user:     Bryan O'Sullivan <bos@serpentine.com>
6 date:     Tue Sep 06 13:15:43 2005 -0700
7 summary:  Introduce a typo into hello.c.
8
9 $ hg update
10 2 files updated, 0 files merged, 0 files removed, 0 files unresolved

```

If you omit an explicit revision, “hg update” will update to the tip revision, as shown by the second call to “hg update” in the example above.

2.8.3 Pushing changes to another repository

Mercurial lets us push changes to another repository, from the repository we're currently visiting. As with the example of "hg pull" above, we'll create a temporary repository to push our changes into.

```
1 $ cd ..
2 $ hg clone hello hello-push
3 2 files updated, 0 files merged, 0 files removed, 0 files unresolved
```

The "hg outgoing" command tells us what changes would be pushed into another repository.

```
1 $ cd my-hello
2 $ hg outgoing ../hello-push
3 comparing with ../hello-push
4 searching for changes
5 changeset: 5:fc920a075b40
6 tag: tip
7 user: Bryan O'Sullivan <bos@serpentine.com>
8 date: Mon Dec 10 19:55:40 2007 +0000
9 summary: Added an extra line of output
10
```

And the "hg push" command does the actual push.

```
1 $ hg push ../hello-push
2 pushing to ../hello-push
3 searching for changes
4 adding changesets
5 adding manifests
6 adding file changes
7 added 1 changesets with 1 changes to 1 files
```

As with "hg pull", the "hg push" command does not update the working directory in the repository that it's pushing changes into. (Unlike "hg pull", "hg push" does not provide a -u option that updates the other repository's working directory.)

What happens if we try to pull or push changes and the receiving repository already has those changes? Nothing too exciting.

```
1 $ hg push ../hello-push
2 pushing to ../hello-push
3 searching for changes
4 no changes found
```

2.8.4 Sharing changes over a network

The commands we have covered in the previous few sections are not limited to working with local repositories. Each works in exactly the same fashion over a network connection; simply pass in a URL instead of a local path.

```
1 $ hg outgoing http://hg.serpentine.com/tutorial/hello
2 comparing with http://hg.serpentine.com/tutorial/hello
3 searching for changes
4 changeset: 5:fc920a075b40
```

```
5 tag:          tip
6 user:         Bryan O'Sullivan <bos@serpentine.com>
7 date:         Mon Dec 10 19:55:40 2007 +0000
8 summary:      Added an extra line of output
9
```

In this example, we can see what changes we could push to the remote repository, but the repository is understandably not set up to let anonymous users push to it.

```
1 $ hg push http://hg.serpentine.com/tutorial/hello
2 pushing to http://hg.serpentine.com/tutorial/hello
3 searching for changes
4 ssl required
```

Chapter 3

A tour of Mercurial: merging work

We've now covered cloning a repository, making changes in a repository, and pulling or pushing changes from one repository into another. Our next step is *merging* changes from separate repositories.

3.1 Merging streams of work

Merging is a fundamental part of working with a distributed revision control tool.

- Alice and Bob each have a personal copy of a repository for a project they're collaborating on. Alice fixes a bug in her repository; Bob adds a new feature in his. They want the shared repository to contain both the bug fix and the new feature.
- I frequently work on several different tasks for a single project at once, each safely isolated in its own repository. Working this way means that I often need to merge one piece of my own work with another.

Because merging is such a common thing to need to do, Mercurial makes it easy. Let's walk through the process. We'll begin by cloning yet another repository (see how often they spring up?) and making a change in it.

```
1 $ cd ..
2 $ hg clone hello my-new-hello
3 2 files updated, 0 files merged, 0 files removed, 0 files unresolved
4 $ cd my-new-hello
5 $ sed -i '/printf/i\\tprintf("once more, hello.\\n");' hello.c
6 $ hg commit -m 'A new hello for a new day.'
```

We should now have two copies of `hello.c` with different contents. The histories of the two repositories have also diverged, as illustrated in figure 3.1.

```
1 $ cat hello.c
2 /*
3  * Placed in the public domain by Bryan O'Sullivan. This program is
4  * not covered by patents in the United States or other countries.
5  */
6
7 #include <stdio.h>
8
9 int main(int argc, char **argv)
10 {
11     printf("once more, hello.\\n");
```

```

12     printf("hello, world!\n");
13     return 0;
14 }
15 $ cat ../my-hello/hello.c
16 /*
17  * Placed in the public domain by Bryan O'Sullivan. This program is
18  * not covered by patents in the United States or other countries.
19  */
20
21 #include <stdio.h>
22
23 int main(int argc, char **argv)
24 {
25     printf("hello, world!\n");
26     printf("hello again!\n");
27     return 0;
28 }

```




Figure 3.1: Divergent recent histories of the my-hello and my-new-hello repositories

We already know that pulling changes from our my-hello repository will have no effect on the working directory.

```

1 $ hg pull ../my-hello
2 pulling from ../my-hello
3 searching for changes
4 adding changesets
5 adding manifests
6 adding file changes
7 added 1 changesets with 1 changes to 1 files (+1 heads)
8 (run 'hg heads' to see heads, 'hg merge' to merge)

```

However, the “hg pull” command says something about “heads”.

3.1.1 Head changesets

A head is a change that has no descendants, or children, as they're also known. The tip revision is thus a head, because the newest revision in a repository doesn't have any children, but a repository can contain more than one head.




Figure 3.2: Repository contents after pulling from my-hello into my-new-hello

In figure 3.2, you can see the effect of the pull from my-hello into my-new-hello. The history that was already present in my-new-hello is untouched, but a new revision has been added. By referring to figure 3.1, we can see that the *changeset ID* remains the same in the new repository, but the *revision number* has changed. (This, incidentally, is a fine example of why it's not safe to use revision numbers when discussing changesets.) We can view the heads in a repository using the “hg heads” command.

```
1 $ hg heads
2 changeset: 6:fc920a075b40
3 tag:      tip
4 parent:   4:b57f9a090b62
5 user:     Bryan O'Sullivan <bos@serpentine.com>
6 date:     Mon Dec 10 19:55:40 2007 +0000
7 summary:   Added an extra line of output
8
9 changeset: 5:6b25c84e540c
10 user:     Bryan O'Sullivan <bos@serpentine.com>
11 date:     Mon Dec 10 19:55:42 2007 +0000
12 summary:   A new hello for a new day.
13
```

3.1.2 Performing the merge

What happens if we try to use the normal “hg update” command to update to the new tip?

```

1 $ hg update
2 abort: update spans branches, use 'hg merge' or 'hg update -C' to lose changes

```

Mercurial is telling us that the “hg update” command won’t do a merge; it won’t update the working directory when it thinks we might be wanting to do a merge, unless we force it to do so. Instead, we use the “hg merge” command to merge the two heads.

```

1 $ hg merge
2 merging hello.c
3 0 files updated, 1 files merged, 0 files removed, 0 files unresolved
4 (branch merge, don't forget to commit)

```




Figure 3.3: Working directory and repository during merge, and following commit

This updates the working directory so that it contains changes from *both* heads, which is reflected in both the output of “hg parents” and the contents of `hello.c`.

```

1 $ hg parents
2 changeset: 5:6b25c84e540c
3 user:      Bryan O'Sullivan <bos@serpentine.com>
4 date:      Mon Dec 10 19:55:42 2007 +0000
5 summary:   A new hello for a new day.
6
7 changeset: 6:fc920a075b40
8 tag:       tip
9 parent:    4:b57f9a090b62
10 user:     Bryan O'Sullivan <bos@serpentine.com>
11 date:     Mon Dec 10 19:55:40 2007 +0000
12 summary:  Added an extra line of output
13
14 $ cat hello.c
15 /*
16  * Placed in the public domain by Bryan O'Sullivan. This program is
17  * not covered by patents in the United States or other countries.

```

```

18  */
19
20  #include <stdio.h>
21
22  int main(int argc, char **argv)
23  {
24      printf("once more, hello.\n");
25      printf("hello, world!\n");
26      printf("hello again!\n");
27      return 0;
28  }

```

3.1.3 Committing the results of the merge

Whenever we’ve done a merge, “hg parents” will display two parents until we “hg commit” the results of the merge.

```

1  $ hg commit -m 'Merged changes'

```

We now have a new tip revision; notice that it has *both* of our former heads as its parents. These are the same revisions that were previously displayed by “hg parents”.

```

1  $ hg tip
2  changeset: 7:53b340e61e98
3  tag:      tip
4  parent:   5:6b25c84e540c
5  parent:   6:fc920a075b40
6  user:     Bryan O'Sullivan <bos@serpentine.com>
7  date:     Mon Dec 10 19:55:42 2007 +0000
8  summary:  Merged changes
9

```

In figure 3.3, you can see a representation of what happens to the working directory during the merge, and how this affects the repository when the commit happens. During the merge, the working directory has two parent changesets, and these become the parents of the new changeset.

3.2 Merging conflicting changes

Most merges are simple affairs, but sometimes you’ll find yourself merging changes where each modifies the same portions of the same files. Unless both modifications are identical, this results in a *conflict*, where you have to decide how to reconcile the different changes into something coherent.

Figure 3.4 illustrates an instance of two conflicting changes to a document. We started with a single version of the file; then we made some changes; while someone else made different changes to the same text. Our task in resolving the conflicting changes is to decide what the file should look like.

Mercurial doesn’t have a built-in facility for handling conflicts. Instead, it runs an external program called `hgmerge`. This is a shell script that is bundled with Mercurial; you can change it to behave however you please. What it does by default is try to find one of several different merging tools that are likely to be installed on your system. It first tries a few fully automatic merging tools; if these don’t succeed (because the resolution process requires human guidance) or aren’t present, the script tries a few different graphical merging tools.

It’s also possible to get Mercurial to run another program or script instead of `hgmerge`, by setting the `HGMERGE` environment variable to the name of your preferred program.




Figure 3.4: Conflicting changes to a document

3.2.1 Using a graphical merge tool

My preferred graphical merge tool is `kdiff3`, which I'll use to describe the features that are common to graphical file merging tools. You can see a screenshot of `kdiff3` in action in figure 3.5. The kind of merge it is performing is called a *three-way merge*, because there are three different versions of the file of interest to us. The tool thus splits the upper portion of the window into three panes:

- At the left is the *base* version of the file, i.e. the most recent version from which the two versions we're trying to merge are descended.
- In the middle is "our" version of the file, with the contents that we modified.
- On the right is "their" version of the file, the one that from the changeset that we're trying to merge with.

In the pane below these is the current *result* of the merge. Our task is to replace all of the red text, which indicates unresolved conflicts, with some sensible merger of the "ours" and "theirs" versions of the file.

All four of these panes are *locked together*; if we scroll vertically or horizontally in any of them, the others are updated to display the corresponding sections of their respective files.

For each conflicting portion of the file, we can choose to resolve the conflict using some combination of text from the base version, ours, or theirs. We can also manually edit the merged file at any time, in case we need to make further modifications.

There are *many* file merging tools available, too many to cover here. They vary in which platforms they are available for, and in their particular strengths and weaknesses. Most are tuned for merging files containing plain text, while a few are aimed at specialised file formats (generally XML).

3.2.2 A worked example

In this example, we will reproduce the file modification history of figure 3.4 above. Let's begin by creating a repository with a base version of our document.

```
1 $ cat > letter.txt <<EOF
2 > Greetings!
3 > I am Mariam Abacha, the wife of former
4 > Nigerian dictator Sani Abacha.
```

```

5 > EOF
6 $ hg add letter.txt
7 $ hg commit -m '419 scam, first draft'

```

We'll clone the repository and make a change to the file.

```

1 $ cd ..
2 $ hg clone scam scam-cousin
3 1 files updated, 0 files merged, 0 files removed, 0 files unresolved
4 $ cd scam-cousin
5 $ cat > letter.txt <<EOF
6 > Greetings!
7 > I am Shehu Musa Abacha, cousin to the former
8 > Nigerian dictator Sani Abacha.
9 > EOF
10 $ hg commit -m '419 scam, with cousin'

```

And another clone, to simulate someone else making a change to the file. (This hints at the idea that it's not all that unusual to merge with yourself when you isolate tasks in separate repositories, and indeed to find and resolve conflicts while doing so.)

```

1 $ cd ..
2 $ hg clone scam scam-son
3 1 files updated, 0 files merged, 0 files removed, 0 files unresolved
4 $ cd scam-son
5 $ cat > letter.txt <<EOF
6 > Greetings!
7 > I am Alhaji Abba Abacha, son of the former
8 > Nigerian dictator Sani Abacha.
9 > EOF
10 $ hg commit -m '419 scam, with son'

```

Having created two different versions of the file, we'll set up an environment suitable for running our merge.

```

1 $ cd ..
2 $ hg clone scam-cousin scam-merge
3 1 files updated, 0 files merged, 0 files removed, 0 files unresolved
4 $ cd scam-merge
5 $ hg pull -u ../scam-son
6 pulling from ../scam-son
7 searching for changes
8 adding changesets
9 adding manifests
10 adding file changes
11 added 1 changesets with 1 changes to 1 files (+1 heads)
12 not updating, since new heads added
13 (run 'hg heads' to see heads, 'hg merge' to merge)

```

In this example, I won't use Mercurial's normal `hgmerge` program to do the merge, because it would drop my nice automated example-running tool into a graphical user interface. Instead, I'll set `HGMERGE` to tell Mercurial to use the non-interactive `merge` command. This is bundled with many Unix-like systems. If you're following this example on your computer, don't bother setting `HGMERGE`.

```

1 $ export HGMERGE=merge
2 $ hg merge
3 merging letter.txt
4 merge: warning: conflicts during merge
5 merging letter.txt failed!
6 0 files updated, 0 files merged, 0 files removed, 1 files unresolved
7 There are unresolved merges, you can redo the full merge using:
8     hg update -C 1
9     hg merge 2
10 $ cat letter.txt
11 Greetings!
12 <<<<<< /tmp/tour-merge-conflictFb6vGk/scam-merge/letter.txt
13 I am Shehu Musa Abacha, cousin to the former
14 =====
15 I am Alhaji Abba Abacha, son of the former
16 >>>>>> /tmp/letter.txt~other.dmqz8N
17 Nigerian dictator Sani Abacha.

```

Because `merge` can't resolve the conflicting changes, it leaves *merge markers* inside the file that has conflicts, indicating which lines have conflicts, and whether they came from our version of the file or theirs.

Mercurial can tell from the way `merge` exits that it wasn't able to merge successfully, so it tells us what commands we'll need to run if we want to redo the merging operation. This could be useful if, for example, we were running a graphical merge tool and quit because we were confused or realised we had made a mistake.

If automatic or manual merges fail, there's nothing to prevent us from "fixing up" the affected files ourselves, and committing the results of our merge:

```

1 $ cat > letter.txt <<EOF
2 > Greetings!
3 > I am Bryan O'Sullivan, no relation of the former
4 > Nigerian dictator Sani Abacha.
5 > EOF
6 $ hg commit -m 'Send me your money'
7 $ hg tip
8 changeset: 3:55085a603229
9 tag: tip
10 parent: 1:4f3647dfc1a6
11 parent: 2:162de8ba5e8f
12 user: Bryan O'Sullivan <bos@serpentine.com>
13 date: Mon Dec 10 19:55:43 2007 +0000
14 summary: Send me your money
15

```

3.3 Simplifying the pull-merge-commit sequence

The process of merging changes as outlined above is straightforward, but requires running three commands in sequence.

```

1 hg pull
2 hg merge
3 hg commit -m 'Merged remote changes'

```

In the case of the final commit, you also need to enter a commit message, which is almost always going to be a piece of uninteresting “boilerplate” text.

It would be nice to reduce the number of steps needed, if this were possible. Indeed, Mercurial is distributed with an extension called `fetch` that does just this.

Mercurial provides a flexible extension mechanism that lets people extend its functionality, while keeping the core of Mercurial small and easy to deal with. Some extensions add new commands that you can use from the command line, while others work “behind the scenes,” for example adding capabilities to the server.

The `fetch` extension adds a new command called, not surprisingly, “`hg fetch`”. This extension acts as a combination of “`hg pull`”, “`hg update`” and “`hg merge`”. It begins by pulling changes from another repository into the current repository. If it finds that the changes added a new head to the repository, it begins a merge, then commits the result of the merge with an automatically-generated commit message. If no new heads were added, it updates the working directory to the new tip changeset.

Enabling the `fetch` extension is easy. Edit your `.hgrc`, and either go to the `[extensions]` section or create an `[extensions]` section. Then add a line that simply reads “`fetch` ”.

```
1 [extensions]
2 fetch =
```

(Normally, on the right-hand side of the “=” would appear the location of the extension, but since the `fetch` extension is in the standard distribution, Mercurial knows where to search for it.)

File Edit Directory Movement Diffview Merge Window Settings Help



A (Base): /tmp/letter.txt~base.C8-rvc ... Top line 1 B: /home/bos/scam/letter.txt

Greetings!

I am Mariam Abacha, the_wife_of former
Nigerian dictator Sani Abacha.

Greetings!

I am Shehu_Musa
Nigerian dictator

Output : /home/bos/scam/letter.txt

Greetings!

? <Merge Conflict>

Nigerian dictator Sani Abacha.

Chapter 4

Behind the scenes

Unlike many revision control systems, the concepts upon which Mercurial is built are simple enough that it's easy to understand how the software really works. Knowing this certainly isn't necessary, but I find it useful to have a "mental model" of what's going on.

This understanding gives me confidence that Mercurial has been carefully designed to be both *safe* and *efficient*. And just as importantly, if it's easy for me to retain a good idea of what the software is doing when I perform a revision control task, I'm less likely to be surprised by its behaviour.

In this chapter, we'll initially cover the core concepts behind Mercurial's design, then continue to discuss some of the interesting details of its implementation.

4.1 Mercurial's historical record

4.1.1 Tracking the history of a single file

When Mercurial tracks modifications to a file, it stores the history of that file in a metadata object called a *filelog*. Each entry in the filelog contains enough information to reconstruct one revision of the file that is being tracked. Filelogs are stored as files in the `.hg/store/data` directory. A filelog contains two kinds of information: revision data, and an index to help Mercurial to find a revision efficiently.

A file that is large, or has a lot of history, has its filelog stored in separate data (".d" suffix) and index (".i" suffix) files. For small files without much history, the revision data and index are combined in a single ".i" file. The correspondence between a file in the working directory and the filelog that tracks its history in the repository is illustrated in figure 4.1.

4.1.2 Managing tracked files

Mercurial uses a structure called a *manifest* to collect together information about the files that it tracks. Each entry in the manifest contains information about the files present in a single changeset. An entry records which files are present in the changeset, the revision of each file, and a few other pieces of file metadata.

4.1.3 Recording changeset information

The *changelog* contains information about each changeset. Each revision records who committed a change, the changeset comment, other pieces of changeset-related information, and the revision of the manifest to use.

4.1.4 Relationships between revisions

Within a changelog, a manifest, or a filelog, each revision stores a pointer to its immediate parent (or to its two parents, if it's a merge revision). As I mentioned above, there are also relationships between revisions *across* these structures,




Figure 4.1: Relationships between files in working directory and filelogs in repository

and they are hierarchical in nature.

For every changeset in a repository, there is exactly one revision stored in the changelog. Each revision of the changelog contains a pointer to a single revision of the manifest. A revision of the manifest stores a pointer to a single revision of each filelog tracked when that changeset was created. These relationships are illustrated in figure 4.2.




Figure 4.2: Metadata relationships

As the illustration shows, there is *not* a “one to one” relationship between revisions in the changelog, manifest, or filelog. If the manifest hasn’t changed between two changesets, the changelog entries for those changesets will point to the same revision of the manifest. If a file that Mercurial tracks hasn’t changed between two changesets, the entry for that file in the two revisions of the manifest will point to the same revision of its filelog.

4.2 Safe, efficient storage

The underpinnings of changelogs, manifests, and filelogs are provided by a single structure called the *revlog*.

4.2.1 Efficient storage

The revlog provides efficient storage of revisions using a *delta* mechanism. Instead of storing a complete copy of a file for each revision, it stores the changes needed to transform an older revision into the new revision. For many kinds of file data, these deltas are typically a fraction of a percent of the size of a full copy of a file.

Some obsolete revision control systems can only work with deltas of text files. They must either store binary files as complete snapshots or encoded into a text representation, both of which are wasteful approaches. Mercurial can efficiently handle deltas of files with arbitrary binary contents; it doesn't need to treat text as special.

4.2.2 Safe operation

Mercurial only ever *appends* data to the end of a revlog file. It never modifies a section of a file after it has written it. This is both more robust and efficient than schemes that need to modify or rewrite data.

In addition, Mercurial treats every write as part of a *transaction* that can span a number of files. A transaction is *atomic*: either the entire transaction succeeds and its effects are all visible to readers in one go, or the whole thing is undone. This guarantee of atomicity means that if you're running two copies of Mercurial, where one is reading data and one is writing it, the reader will never see a partially written result that might confuse it.

The fact that Mercurial only appends to files makes it easier to provide this transactional guarantee. The easier it is to do stuff like this, the more confident you should be that it's done correctly.

4.2.3 Fast retrieval

Mercurial cleverly avoids a pitfall common to all earlier revision control systems: the problem of *inefficient retrieval*. Most revision control systems store the contents of a revision as an incremental series of modifications against a "snapshot". To reconstruct a specific revision, you must first read the snapshot, and then every one of the revisions between the snapshot and your target revision. The more history that a file accumulates, the more revisions you must read, hence the longer it takes to reconstruct a particular revision.




Figure 4.3: Snapshot of a revlog, with incremental deltas

The innovation that Mercurial applies to this problem is simple but effective. Once the cumulative amount of delta information stored since the last snapshot exceeds a fixed threshold, it stores a new snapshot (compressed, of course), instead of another delta. This makes it possible to reconstruct *any* revision of a file quickly. This approach works so well that it has since been copied by several other revision control systems.

Figure 4.3 illustrates the idea. In an entry in a revlog’s index file, Mercurial stores the range of entries from the data file that it must read to reconstruct a particular revision.

Aside: the influence of video compression

If you’re familiar with video compression or have ever watched a TV feed through a digital cable or satellite service, you may know that most video compression schemes store each frame of video as a delta against its predecessor frame. In addition, these schemes use “lossy” compression techniques to increase the compression ratio, so visual errors accumulate over the course of a number of inter-frame deltas.

Because it’s possible for a video stream to “drop out” occasionally due to signal glitches, and to limit the accumulation of artefacts introduced by the lossy compression process, video encoders periodically insert a complete frame (called a “key frame”) into the video stream; the next delta is generated against that frame. This means that if the video signal gets interrupted, it will resume once the next key frame is received. Also, the accumulation of encoding errors restarts anew with each key frame.

4.2.4 Identification and strong integrity

Along with delta or snapshot information, a revlog entry contains a cryptographic hash of the data that it represents. This makes it difficult to forge the contents of a revision, and easy to detect accidental corruption.

Hashes provide more than a mere check against corruption; they are used as the identifiers for revisions. The changeset identification hashes that you see as an end user are from revisions of the changelog. Although filelogs and the manifest also use hashes, Mercurial only uses these behind the scenes.

Mercurial verifies that hashes are correct when it retrieves file revisions and when it pulls changes from another repository. If it encounters an integrity problem, it will complain and stop whatever it’s doing.

In addition to the effect it has on retrieval efficiency, Mercurial’s use of periodic snapshots makes it more robust against partial data corruption. If a revlog becomes partly corrupted due to a hardware error or system bug, it’s often possible to reconstruct some or most revisions from the uncorrupted sections of the revlog, both before and after the corrupted section. This would not be possible with a delta-only storage model.

4.3 Revision history, branching, and merging

Every entry in a Mercurial revlog knows the identity of its immediate ancestor revision, usually referred to as its *parent*. In fact, a revision contains room for not one parent, but two. Mercurial uses a special hash, called the “null ID”, to represent the idea “there is no parent here”. This hash is simply a string of zeroes.

In figure 4.4, you can see an example of the conceptual structure of a revlog. Filelogs, manifests, and changelogs all have this same structure; they differ only in the kind of data stored in each delta or snapshot.

The first revision in a revlog (at the bottom of the image) has the null ID in both of its parent slots. For a “normal” revision, its first parent slot contains the ID of its parent revision, and its second contains the null ID, indicating that the revision has only one real parent. Any two revisions that have the same parent ID are branches. A revision that represents a merge between branches has two normal revision IDs in its parent slots.

4.4 The working directory

In the working directory, Mercurial stores a snapshot of the files from the repository as of a particular changeset.

The working directory “knows” which changeset it contains. When you update the working directory to contain a particular changeset, Mercurial looks up the appropriate revision of the manifest to find out which files it was tracking at the time that changeset was committed, and which revision of each file was then current. It then recreates a copy of each of those files, with the same contents it had when the changeset was committed.

The *dirstate* contains Mercurial’s knowledge of the working directory. This details which changeset the working directory is updated to, and all of the files that Mercurial is tracking in the working directory.

Just as a revision of a revlog has room for two parents, so that it can represent either a normal revision (with one parent) or a merge of two earlier revisions, the dirstate has slots for two parents. When you use the “hg update” command, the changeset that you update to is stored in the “first parent” slot, and the null ID in the second. When you “hg merge” with another changeset, the first parent remains unchanged, and the second parent is filled in with the changeset you’re merging with. The “hg parents” command tells you what the parents of the dirstate are.

4.4.1 What happens when you commit

The dirstate stores parent information for more than just book-keeping purposes. Mercurial uses the parents of the dirstate as *the parents of a new changeset* when you perform a commit.

Figure 4.5 shows the normal state of the working directory, where it has a single changeset as parent. That changeset is the *tip*, the newest changeset in the repository that has no children.

It’s useful to think of the working directory as “the changeset I’m about to commit”. Any files that you tell Mercurial that you’ve added, removed, renamed, or copied will be reflected in that changeset, as will modifications to any files that Mercurial is already tracking; the new changeset will have the parents of the working directory as its parents.

After a commit, Mercurial will update the parents of the working directory, so that the first parent is the ID of the new changeset, and the second is the null ID. This is shown in figure 4.6. Mercurial doesn’t touch any of the files in the working directory when you commit; it just modifies the dirstate to note its new parents.

4.4.2 Creating a new head

It’s perfectly normal to update the working directory to a changeset other than the current tip. For example, you might want to know what your project looked like last Tuesday, or you could be looking through changesets to see which one introduced a bug. In cases like this, the natural thing to do is update the working directory to the changeset you’re interested in, and then examine the files in the working directory directly to see their contents as they were when you committed that changeset. The effect of this is shown in figure 4.7.

Having updated the working directory to an older changeset, what happens if you make some changes, and then commit? Mercurial behaves in the same way as I outlined above. The parents of the working directory become the parents of the new changeset. This new changeset has no children, so it becomes the new tip. And the repository now contains two changesets that have no children; we call these *heads*. You can see the structure that this creates in figure 4.8.

Note: If you’re new to Mercurial, you should keep in mind a common “error”, which is to use the “hg pull” command without any options. By default, the “hg pull” command *does not* update the working directory, so you’ll bring new changesets into your repository, but the working directory will stay synced at the same changeset as before the pull. If you make some changes and commit afterwards, you’ll thus create a new head, because your working directory isn’t synced to whatever the current tip is.

I put the word “error” in quotes because all that you need to do to rectify this situation is “hg merge”, then “hg commit”. In other words, this almost never has negative consequences; it just surprises people. I’ll discuss other ways to avoid this behaviour, and why Mercurial behaves in this initially surprising way, later on.

4.4.3 Merging heads

When you run the “hg merge” command, Mercurial leaves the first parent of the working directory unchanged, and sets the second parent to the changeset you’re merging with, as shown in figure 4.9.

Mercurial also has to modify the working directory, to merge the files managed in the two changesets. Simplified a little, the merging process goes like this, for every file in the manifests of both changesets.

- If neither changeset has modified a file, do nothing with that file.

- If one changeset has modified a file, and the other hasn't, create the modified copy of the file in the working directory.
- If one changeset has removed a file, and the other hasn't (or has also deleted it), delete the file from the working directory.
- If one changeset has removed a file, but the other has modified the file, ask the user what to do: keep the modified file, or remove it?
- If both changesets have modified a file, invoke an external merge program to choose the new contents for the merged file. This may require input from the user.
- If one changeset has modified a file, and the other has renamed or copied the file, make sure that the changes follow the new name of the file.

There are more details—merging has plenty of corner cases—but these are the most common choices that are involved in a merge. As you can see, most cases are completely automatic, and indeed most merges finish automatically, without requiring your input to resolve any conflicts.

When you're thinking about what happens when you commit after a merge, once again the working directory is “the changeset I'm about to commit”. After the “hg merge” command completes, the working directory has two parents; these will become the parents of the new changeset.

Mercurial lets you perform multiple merges, but you must commit the results of each individual merge as you go. This is necessary because Mercurial only tracks two parents for both revisions and the working directory. While it would be technically possible to merge multiple changesets at once, the prospect of user confusion and making a terrible mess of a merge immediately becomes overwhelming.

4.5 Other interesting design features

In the sections above, I've tried to highlight some of the most important aspects of Mercurial's design, to illustrate that it pays careful attention to reliability and performance. However, the attention to detail doesn't stop there. There are a number of other aspects of Mercurial's construction that I personally find interesting. I'll detail a few of them here, separate from the “big ticket” items above, so that if you're interested, you can gain a better idea of the amount of thinking that goes into a well-designed system.

4.5.1 Clever compression

When appropriate, Mercurial will store both snapshots and deltas in compressed form. It does this by always *trying to* compress a snapshot or delta, but only storing the compressed version if it's smaller than the uncompressed version.

This means that Mercurial does “the right thing” when storing a file whose native form is compressed, such as a zip archive or a JPEG image. When these types of files are compressed a second time, the resulting file is usually bigger than the once-compressed form, and so Mercurial will store the plain zip or JPEG.

Deltas between revisions of a compressed file are usually larger than snapshots of the file, and Mercurial again does “the right thing” in these cases. It finds that such a delta exceeds the threshold at which it should store a complete snapshot of the file, so it stores the snapshot, again saving space compared to a naive delta-only approach.

Network recompression

When storing revisions on disk, Mercurial uses the “deflate” compression algorithm (the same one used by the popular zip archive format), which balances good speed with a respectable compression ratio. However, when transmitting revision data over a network connection, Mercurial uncompresses the compressed revision data.

If the connection is over HTTP, Mercurial recompresses the entire stream of data using a compression algorithm that gives a better compression ratio (the Burrows-Wheeler algorithm from the widely used bzip2 compression package). This combination of algorithm and compression of the entire stream (instead of a revision at a time) substantially reduces the number of bytes to be transferred, yielding better network performance over almost all kinds of network.

(If the connection is over `ssh`, Mercurial *doesn't* recompress the stream, because `ssh` can already do this itself.)

4.5.2 Read/write ordering and atomicity

Appending to files isn't the whole story when it comes to guaranteeing that a reader won't see a partial write. If you recall figure 4.2, revisions in the changelog point to revisions in the manifest, and revisions in the manifest point to revisions in filelogs. This hierarchy is deliberate.

A writer starts a transaction by writing filelog and manifest data, and doesn't write any changelog data until those are finished. A reader starts by reading changelog data, then manifest data, followed by filelog data.

Since the writer has always finished writing filelog and manifest data before it writes to the changelog, a reader will never read a pointer to a partially written manifest revision from the changelog, and it will never read a pointer to a partially written filelog revision from the manifest.

4.5.3 Concurrent access

The read/write ordering and atomicity guarantees mean that Mercurial never needs to *lock* a repository when it's reading data, even if the repository is being written to while the read is occurring. This has a big effect on scalability; you can have an arbitrary number of Mercurial processes safely reading data from a repository safely all at once, no matter whether it's being written to or not.

The lockless nature of reading means that if you're sharing a repository on a multi-user system, you don't need to grant other local users permission to *write* to your repository in order for them to be able to clone it or pull changes from it; they only need *read* permission. (This is *not* a common feature among revision control systems, so don't take it for granted! Most require readers to be able to lock a repository to access it safely, and this requires write permission on at least one directory, which of course makes for all kinds of nasty and annoying security and administrative problems.)

Mercurial uses locks to ensure that only one process can write to a repository at a time (the locking mechanism is safe even over filesystems that are notoriously hostile to locking, such as NFS). If a repository is locked, a writer will wait for a while to retry if the repository becomes unlocked, but if the repository remains locked for too long, the process attempting to write will time out after a while. This means that your daily automated scripts won't get stuck forever and pile up if a system crashes unnoticed, for example. (Yes, the timeout is configurable, from zero to infinity.)

Safe dirstate access

As with revision data, Mercurial doesn't take a lock to read the dirstate file; it does acquire a lock to write it. To avoid the possibility of reading a partially written copy of the dirstate file, Mercurial writes to a file with a unique name in the same directory as the dirstate file, then renames the temporary file atomically to dirstate. The file named dirstate is thus guaranteed to be complete, not partially written.

4.5.4 Avoiding seeks

Critical to Mercurial's performance is the avoidance of seeks of the disk head, since any seek is far more expensive than even a comparatively large read operation.

This is why, for example, the dirstate is stored in a single file. If there were a dirstate file per directory that Mercurial tracked, the disk would seek once per directory. Instead, Mercurial reads the entire single dirstate file in one step.

Mercurial also uses a "copy on write" scheme when cloning a repository on local storage. Instead of copying every revlog file from the old repository into the new repository, it makes a "hard link", which is a shorthand way to say "these two names point to the same file". When Mercurial is about to write to one of a revlog's files, it checks to see if the number of names pointing at the file is greater than one. If it is, more than one repository is using the file, so Mercurial makes a new copy of the file that is private to this repository.

A few revision control developers have pointed out that this idea of making a complete private copy of a file is not very efficient in its use of storage. While this is true, storage is cheap, and this method gives the highest performance while deferring most book-keeping to the operating system. An alternative scheme would most likely

reduce performance and increase the complexity of the software, each of which is much more important to the “feel” of day-to-day use.

4.5.5 Other contents of the dirstate

Because Mercurial doesn’t force you to tell it when you’re modifying a file, it uses the dirstate to store some extra information so it can determine efficiently whether you have modified a file. For each file in the working directory, it stores the time that it last modified the file itself, and the size of the file at that time.

When you explicitly “hg add”, “hg remove”, “hg rename” or “hg copy” files, Mercurial updates the dirstate so that it knows what to do with those files when you commit.

When Mercurial is checking the states of files in the working directory, it first checks a file’s modification time. If that has not changed, the file must not have been modified. If the file’s size has changed, the file must have been modified. If the modification time has changed, but the size has not, only then does Mercurial need to read the actual contents of the file to see if they’ve changed. Storing these few extra pieces of information dramatically reduces the amount of data that Mercurial needs to read, which yields large performance improvements compared to other revision control systems.




Figure 4.4:




Figure 4.5: The working directory can have two parents




Figure 4.6: The working directory gains new parents after a commit




Figure 4.7: The working directory, updated to an older changeset




Figure 4.8: After a commit made while synced to an older changeset




Figure 4.9: Merging two heads

Chapter 5

Mercurial in daily use

5.1 Telling Mercurial which files to track

Mercurial does not work with files in your repository unless you tell it to manage them. The “`hg status`” command will tell you which files Mercurial doesn’t know about; it uses a “?” to display such files.

To tell Mercurial to track a file, use the “`hg add`” command. Once you have added a file, the entry in the output of “`hg status`” for that file changes from “?” to “A”.

```
1 $ hg init add-example
2 $ cd add-example
3 $ echo a > a
4 $ hg status
5 ? a
6 $ hg add a
7 $ hg status
8 A a
9 $ hg commit -m 'Added one file'
10 $ hg status
```

After you run a “`hg commit`”, the files that you added before the commit will no longer be listed in the output of “`hg status`”. The reason for this is that “`hg status`” only tells you about “interesting” files—those that you have modified or told Mercurial to do something with—by default. If you have a repository that contains thousands of files, you will rarely want to know about files that Mercurial is tracking, but that have not changed. (You can still get this information; we’ll return to this later.)

Once you add a file, Mercurial doesn’t do anything with it immediately. Instead, it will take a snapshot of the file’s state the next time you perform a commit. It will then continue to track the changes you make to the file every time you commit, until you remove the file.

5.1.1 Explicit versus implicit file naming

A useful behaviour that Mercurial has is that if you pass the name of a directory to a command, every Mercurial command will treat this as “I want to operate on every file in this directory and its subdirectories”.

```
1 $ mkdir b
2 $ echo b > b/b
3 $ echo c > b/c
4 $ mkdir b/d
5 $ echo d > b/d/d
```

```
6 $ hg add b
7 adding b/b
8 adding b/c
9 adding b/d/d
10 $ hg commit -m 'Added all files in subdirectory'
```

Notice in this example that Mercurial printed the names of the files it added, whereas it didn't do so when we added the file named `a` in the earlier example.

What's going on is that in the former case, we explicitly named the file to add on the command line, so the assumption that Mercurial makes in such cases is that you know what you were doing, and it doesn't print any output.

However, when we *imply* the names of files by giving the name of a directory, Mercurial takes the extra step of printing the name of each file that it does something with. This makes it more clear what is happening, and reduces the likelihood of a silent and nasty surprise. This behaviour is common to most Mercurial commands.

5.1.2 Aside: Mercurial tracks files, not directories

Mercurial does not track directory information. Instead, it tracks the path to a file. Before creating a file, it first creates any missing directory components of the path. After it deletes a file, it then deletes any empty directories that were in the deleted file's path. This sounds like a trivial distinction, but it has one minor practical consequence: it is not possible to represent a completely empty directory in Mercurial.

Empty directories are rarely useful, and there are unintrusive workarounds that you can use to achieve an appropriate effect. The developers of Mercurial thus felt that the complexity that would be required to manage empty directories was not worth the limited benefit this feature would bring.

If you need an empty directory in your repository, there are a few ways to achieve this. One is to create a directory, then “`hg add`” a “hidden” file to that directory. On Unix-like systems, any file name that begins with a period (“.”) is treated as hidden by most commands and GUI tools. This approach is illustrated in figure 5.1.

```
1 $ hg init hidden-example
2 $ cd hidden-example
3 $ mkdir empty
4 $ touch empty/.hidden
5 $ hg add empty/.hidden
6 $ hg commit -m 'Manage an empty-looking directory'
7 $ ls empty
8 $ cd ..
9 $ hg clone hidden-example tmp
10 1 files updated, 0 files merged, 0 files removed, 0 files unresolved
11 $ ls tmp
12 empty
13 $ ls tmp/empty
```

Figure 5.1: Simulating an empty directory using a hidden file

Another way to tackle a need for an empty directory is to simply create one in your automated build scripts before they will need it.

5.2 How to stop tracking a file

Once you decide that a file no longer belongs in your repository, use the “`hg remove`” command; this deletes the file, and tells Mercurial to stop tracking it. A removed file is represented in the output of “`hg status`” with a “R”.

```

1 $ hg init remove-example
2 $ cd remove-example
3 $ echo a > a
4 $ mkdir b
5 $ echo b > b/b
6 $ hg add a b
7 adding b/b
8 $ hg commit -m 'Small example for file removal'
9 $ hg remove a
10 $ hg status
11 R a
12 $ hg remove b
13 removing b/b

```

After you “hg remove” a file, Mercurial will no longer track changes to that file, even if you recreate a file with the same name in your working directory. If you do recreate a file with the same name and want Mercurial to track the new file, simply “hg add” it. Mercurial will know that the newly added file is not related to the old file of the same name.

5.2.1 Removing a file does not affect its history

It is important to understand that removing a file has only two effects.

- It removes the current version of the file from the working directory.
- It stops Mercurial from tracking changes to the file, from the time of the next commit.

Removing a file *does not* in any way alter the *history* of the file.

If you update the working directory to a changeset in which a file that you have removed was still tracked, it will reappear in the working directory, with the contents it had when you committed that changeset. If you then update the working directory to a later changeset, in which the file had been removed, Mercurial will once again remove the file from the working directory.

5.2.2 Missing files

Mercurial considers a file that you have deleted, but not used “hg remove” to delete, to be *missing*. A missing file is represented with “!” in the output of “hg status”. Mercurial commands will not generally do anything with missing files.

```

1 $ hg init missing-example
2 $ cd missing-example
3 $ echo a > a
4 $ hg add a
5 $ hg commit -m 'File about to be missing'
6 $ rm a
7 $ hg status
8 ! a

```

If your repository contains a file that “hg status” reports as missing, and you want the file to stay gone, you can run “hg remove --after” at any time later on, to tell Mercurial that you really did mean to remove the file.

```
1 $ hg remove --after a
2 $ hg status
3 R a
```

On the other hand, if you deleted the missing file by accident, use “`hg revert filename`” to recover the file. It will reappear, in unmodified form.

```
1 $ hg revert a
2 $ cat a
3 a
4 $ hg status
```

5.2.3 Aside: why tell Mercurial explicitly to remove a file?

You might wonder why Mercurial requires you to explicitly tell it that you are deleting a file. Early during the development of Mercurial, it let you delete a file however you pleased; Mercurial would notice the absence of the file automatically when you next ran a “`hg commit`”, and stop tracking the file. In practice, this made it too easy to accidentally remove a file without noticing.

5.2.4 Useful shorthand—adding and removing files in one step

Mercurial offers a combination command, “`hg addremove`”, that adds untracked files and marks missing files as removed.

```
1 $ hg init addremove-example
2 $ cd addremove-example
3 $ echo a > a
4 $ echo b > b
5 $ hg addremove
6 adding a
7 adding b
```

The “`hg commit`” command also provides a `-A` option that performs this same add-and-remove, immediately followed by a commit.

```
1 $ echo c > c
2 $ hg commit -A -m 'Commit with addremove'
3 adding c
```

5.3 Copying files

Mercurial provides a “`hg copy`” command that lets you make a new copy of a file. When you copy a file using this command, Mercurial makes a record of the fact that the new file is a copy of the original file. It treats these copied files specially when you merge your work with someone else’s.

5.3.1 The results of copying during a merge

What happens during a merge is that changes “follow” a copy. To best illustrate what this means, let’s create an example. We’ll start with the usual tiny repository that contains a single file.

```

1 $ hg init my-copy
2 $ cd my-copy
3 $ echo line > file
4 $ hg add file
5 $ hg commit -m 'Added a file'

```

We need to do some work in parallel, so that we'll have something to merge. So let's clone our repository.

```

1 $ cd ..
2 $ hg clone my-copy your-copy
3 1 files updated, 0 files merged, 0 files removed, 0 files unresolved

```

Back in our initial repository, let's use the “hg copy” command to make a copy of the first file we created.

```

1 $ cd my-copy
2 $ hg copy file new-file

```

If we look at the output of the “hg status” command afterwards, the copied file looks just like a normal added file.

```

1 $ hg status
2 A new-file

```

But if we pass the -C option to “hg status”, it prints another line of output: this is the file that our newly-added file was copied *from*.

```

1 $ hg status -C
2 A new-file
3   file
4 $ hg commit -m 'Copied file'

```

Now, back in the repository we cloned, let's make a change in parallel. We'll add a line of content to the original file that we created.

```

1 $ cd ../your-copy
2 $ echo 'new contents' >> file
3 $ hg commit -m 'Changed file'

```

Now we have a modified `file` in this repository. When we pull the changes from the first repository, and merge the two heads, Mercurial will propagate the changes that we made locally to `file` into its copy, `new-file`.

```

1 $ hg pull ../my-copy
2 pulling from ../my-copy
3 searching for changes
4 adding changesets
5 adding manifests
6 adding file changes
7 added 1 changesets with 1 changes to 1 files (+1 heads)
8 (run 'hg heads' to see heads, 'hg merge' to merge)
9 $ hg merge
10 merging file and new-file

```

```
11 0 files updated, 1 files merged, 0 files removed, 0 files unresolved
12 (branch merge, don't forget to commit)
13 $ cat new-file
14 line
15 new contents
```

5.3.2 Why should changes follow copies?

This behaviour, of changes to a file propagating out to copies of the file, might seem esoteric, but in most cases it's highly desirable.

First of all, remember that this propagation *only* happens when you merge. So if you “hg copy” a file, and subsequently modify the original file during the normal course of your work, nothing will happen.

The second thing to know is that modifications will only propagate across a copy as long as the repository that you're pulling changes from *doesn't know* about the copy.

The reason that Mercurial does this is as follows. Let's say I make an important bug fix in a source file, and commit my changes. Meanwhile, you've decided to “hg copy” the file in your repository, without knowing about the bug or having seen the fix, and you have started hacking on your copy of the file.

If you pulled and merged my changes, and Mercurial *didn't* propagate changes across copies, your source file would now contain the bug, and unless you remembered to propagate the bug fix by hand, the bug would *remain* in your copy of the file.

By automatically propagating the change that fixed the bug from the original file to the copy, Mercurial prevents this class of problem. To my knowledge, Mercurial is the *only* revision control system that propagates changes across copies like this.

Once your change history has a record that the copy and subsequent merge occurred, there's usually no further need to propagate changes from the original file to the copied file, and that's why Mercurial only propagates changes across copies until this point, and no further.

5.3.3 How to make changes *not* follow a copy

If, for some reason, you decide that this business of automatically propagating changes across copies is not for you, simply use your system's normal file copy command (on Unix-like systems, that's `cp`) to make a copy of a file, then “hg add” the new copy by hand. Before you do so, though, please do reread section 5.3.2, and make an informed decision that this behaviour is not appropriate to your specific case.

5.3.4 Behaviour of the “hg copy” command

When you use the “hg copy” command, Mercurial makes a copy of each source file as it currently stands in the working directory. This means that if you make some modifications to a file, then “hg copy” it without first having committed those changes, the new copy will also contain the modifications you have made up until that point. (I find this behaviour a little counterintuitive, which is why I mention it here.)

The “hg copy” command acts similarly to the Unix `cp` command (you can use the “hg cp” alias if you prefer). The last argument is the *destination*, and all prior arguments are *sources*. If you pass it a single file as the source, and the destination does not exist, it creates a new file with that name.

```
1 $ mkdir k
2 $ hg copy a k
3 $ ls k
4 a
```

If the destination is a directory, Mercurial copies its sources into that directory.

```
1 $ mkdir d
2 $ hg copy a b d
3 $ ls d
4 a b
```

Copying a directory is recursive, and preserves the directory structure of the source.

```
1 $ hg copy c e
2 copying c/a/c to e/a/c
```

If the source and destination are both directories, the source tree is recreated in the destination directory.

```
1 $ hg copy c d
2 copying c/a/c to d/c/a/c
```

As with the “hg rename” command, if you copy a file manually and then want Mercurial to know that you’ve copied the file, simply use the `--after` option to “hg copy”.

```
1 $ cp a z
2 $ hg copy --after a z
```

5.4 Renaming files

It’s rather more common to need to rename a file than to make a copy of it. The reason I discussed the “hg copy” command before talking about renaming files is that Mercurial treats a rename in essentially the same way as a copy. Therefore, knowing what Mercurial does when you copy a file tells you what to expect when you rename a file.

When you use the “hg rename” command, Mercurial makes a copy of each source file, then deletes it and marks the file as removed.

```
1 $ hg rename a b
```

The “hg status” command shows the newly copied file as added, and the copied-from file as removed.

```
1 $ hg status
2 A b
3 R a
```

As with the results of a “hg copy”, we must use the `-C` option to “hg status” to see that the added file is really being tracked by Mercurial as a copy of the original, now removed, file.

```
1 $ hg status -C
2 A b
3   a
4 R a
```

As with “hg remove” and “hg copy”, you can tell Mercurial about a rename after the fact using the `--after` option. In most other respects, the behaviour of the “hg rename” command, and the options it accepts, are similar to the “hg copy” command.

5.4.1 Renaming files and merging changes

Since Mercurial’s rename is implemented as copy-and-remove, the same propagation of changes happens when you merge after a rename as after a copy.

If I modify a file, and you rename it to a new name, and then we merge our respective changes, my modifications to the file under its original name will be propagated into the file under its new name. (This is something you might expect to “simply work,” but not all revision control systems actually do this.)

Whereas having changes follow a copy is a feature where you can perhaps nod and say “yes, that might be useful,” it should be clear that having them follow a rename is definitely important. Without this facility, it would simply be too easy for changes to become orphaned when files are renamed.

5.4.2 Divergent renames and merging

The case of diverging names occurs when two developers start with a file—let’s call it `foo`—in their respective repositories.

```
1 $ hg clone orig anne
2 1 files updated, 0 files merged, 0 files removed, 0 files unresolved
3 $ hg clone orig bob
4 1 files updated, 0 files merged, 0 files removed, 0 files unresolved
```

Anne renames the file to `bar`.

```
1 $ cd anne
2 $ hg mv foo bar
3 $ hg ci -m 'Rename foo to bar'
```

Meanwhile, Bob renames it to `quux`.

```
1 $ cd ../bob
2 $ hg mv foo quux
3 $ hg ci -m 'Rename foo to quux'
```

I like to think of this as a conflict because each developer has expressed different intentions about what the file ought to be named.

What do you think should happen when they merge their work? Mercurial’s actual behaviour is that it always preserves *both* names when it merges changesets that contain divergent renames.

```
1 # See http://www.selenic.com/mercurial/bts/issue455
2 $ cd ../orig
3 $ hg pull -u ../anne
4 pulling from ../anne
5 searching for changes
6 adding changesets
7 adding manifests
8 adding file changes
9 added 1 changesets with 1 changes to 1 files
10 1 files updated, 0 files merged, 1 files removed, 0 files unresolved
11 $ hg pull ../bob
12 pulling from ../bob
13 searching for changes
14 adding changesets
15 adding manifests
```

```

16 adding file changes
17 added 1 changesets with 1 changes to 1 files (+1 heads)
18 (run 'hg heads' to see heads, 'hg merge' to merge)
19 $ hg merge
20 warning: detected divergent renames of foo to:
21   bar
22   quux
23 1 files updated, 0 files merged, 0 files removed, 0 files unresolved
24 (branch merge, don't forget to commit)
25 $ ls
26 bar quux

```

Notice that Mercurial does warn about the divergent renames, but it leaves it up to you to do something about the divergence after the merge.

5.4.3 Convergent renames and merging

Another kind of rename conflict occurs when two people choose to rename different *source* files to the same *destination*. In this case, Mercurial runs its normal merge machinery, and lets you guide it to a suitable resolution.

5.4.4 Other name-related corner cases

Mercurial has a longstanding bug in which it fails to handle a merge where one side has a file with a given name, while another has a directory with the same name. This is documented as [Mercurial bug no. 29](#).

```

1  $ hg init issue29
2  $ cd issue29
3  $ echo a > a
4  $ hg ci -Ama
5  adding a
6  $ echo b > b
7  $ hg ci -Amb
8  adding b
9  $ hg up 0
10 0 files updated, 0 files merged, 1 files removed, 0 files unresolved
11 $ mkdir b
12 $ echo b > b/b
13 $ hg ci -Amc
14 adding b/b
15 $ hg merge
16 abort: Is a directory: /tmp/issue29680qIY/issue29/b

```

5.5 Recovering from mistakes

Mercurial has some useful commands that will help you to recover from some common mistakes.

The “`hg revert`” command lets you undo changes that you have made to your working directory. For example, if you “`hg add`” a file by accident, just run “`hg revert`” with the name of the file you added, and while the file won’t be touched in any way, it won’t be tracked for adding by Mercurial any longer, either. You can also use “`hg revert`” to get rid of erroneous changes to a file.

It's useful to remember that the “hg revert” command is useful for changes that you have not yet committed. Once you've committed a change, if you decide it was a mistake, you can still do something about it, though your options may be more limited.

For more information about the “hg revert” command, and details about how to deal with changes you have already committed, see chapter 9.

Chapter 6

Collaborating with other people

As a completely decentralised tool, Mercurial doesn't impose any policy on how people ought to work with each other. However, if you're new to distributed revision control, it helps to have some tools and examples in mind when you're thinking about possible workflow models.

6.1 Mercurial's web interface

Mercurial has a powerful web interface that provides several useful capabilities.

For interactive use, the web interface lets you browse a single repository or a collection of repositories. You can view the history of a repository, examine each change (comments and diffs), and view the contents of each directory and file.

Also for human consumption, the web interface provides an RSS feed of the changes in a repository. This lets you “subscribe” to a repository using your favourite feed reader, and be automatically notified of activity in that repository as soon as it happens. I find this capability much more convenient than the model of subscribing to a mailing list to which notifications are sent, as it requires no additional configuration on the part of whoever is serving the repository.

The web interface also lets remote users clone a repository, pull changes from it, and (when the server is configured to permit it) push changes back to it. Mercurial's HTTP tunneling protocol aggressively compresses data, so that it works efficiently even over low-bandwidth network connections.

The easiest way to get started with the web interface is to use your web browser to visit an existing repository, such as the master Mercurial repository at <http://www.selenic.com/repo/hg?style=gitweb>.

If you're interested in providing a web interface to your own repositories, Mercurial provides two ways to do this. The first is using the “hg serve” command, which is best suited to short-term “lightweight” serving. See section 6.4 below for details of how to use this command. If you have a long-lived repository that you'd like to make permanently available, Mercurial has built-in support for the CGI (Common Gateway Interface) standard, which all common web servers support. See section 6.6 for details of CGI configuration.

6.2 Collaboration models

With a suitably flexible tool, making decisions about workflow is much more of a social engineering challenge than a technical one. Mercurial imposes few limitations on how you can structure the flow of work in a project, so it's up to you and your group to set up and live with a model that matches your own particular needs.

6.2.1 Factors to keep in mind

The most important aspect of any model that you must keep in mind is how well it matches the needs and capabilities of the people who will be using it. This might seem self-evident; even so, you still can't afford to forget it for a moment.

I once put together a workflow model that seemed to make perfect sense to me, but that caused a considerable amount of consternation and strife within my development team. In spite of my attempts to explain why we needed a complex set of branches, and how changes ought to flow between them, a few team members revolted. Even though they were smart people, they didn't want to pay attention to the constraints we were operating under, or face the consequences of those constraints in the details of the model that I was advocating.

Don't sweep foreseeable social or technical problems under the rug. Whatever scheme you put into effect, you should plan for mistakes and problem scenarios. Consider adding automated machinery to prevent, or quickly recover from, trouble that you can anticipate. As an example, if you intend to have a branch with not-for-release changes in it, you'd do well to think early about the possibility that someone might accidentally merge those changes into a release branch. You could avoid this particular problem by writing a hook that prevents changes from being merged from an inappropriate branch.

6.2.2 Informal anarchy

I wouldn't suggest an "anything goes" approach as something sustainable, but it's a model that's easy to grasp, and it works perfectly well in a few unusual situations.

As one example, many projects have a loose-knit group of collaborators who rarely physically meet each other. Some groups like to overcome the isolation of working at a distance by organising occasional "sprints". In a sprint, a number of people get together in a single location (a company's conference room, a hotel meeting room, that kind of place) and spend several days more or less locked in there, hacking intensely on a handful of projects.

A sprint is the perfect place to use the `hg serve` command, since `hg serve` does not require any fancy server infrastructure. You can get started with `hg serve` in moments, by reading section 6.4 below. Then simply tell the person next to you that you're running a server, send the URL to them in an instant message, and you immediately have a quick-turnaround way to work together. They can type your URL into their web browser and quickly review your changes; or they can pull a bugfix from you and verify it; or they can clone a branch containing a new feature and try it out.

The charm, and the problem, with doing things in an ad hoc fashion like this is that only people who know about your changes, and where they are, can see them. Such an informal approach simply doesn't scale beyond a handful people, because each individual needs to know about n different repositories to pull from.

6.2.3 A single central repository

For smaller projects migrating from a centralised revision control tool, perhaps the easiest way to get started is to have changes flow through a single shared central repository. This is also the most common "building block" for more ambitious workflow schemes.

Contributors start by cloning a copy of this repository. They can pull changes from it whenever they need to, and some (perhaps all) developers have permission to push a change back when they're ready for other people to see it.

Under this model, it can still often make sense for people to pull changes directly from each other, without going through the central repository. Consider a case in which I have a tentative bug fix, but I am worried that if I were to publish it to the central repository, it might subsequently break everyone else's trees as they pull it. To reduce the potential for damage, I can ask you to clone my repository into a temporary repository of your own and test it. This lets us put off publishing the potentially unsafe change until it has had a little testing.

In this kind of scenario, people usually use the `ssh` protocol to securely push changes to the central repository, as documented in section 6.5. It's also usual to publish a read-only copy of the repository over HTTP using CGI, as in section 6.6. Publishing over HTTP satisfies the needs of people who don't have push access, and those who want to use web browsers to browse the repository's history.

6.2.4 Working with multiple branches

Projects of any significant size naturally tend to make progress on several fronts simultaneously. In the case of software, it's common for a project to go through periodic official releases. A release might then go into "maintenance mode" for a while after its first publication; maintenance releases tend to contain only bug fixes, not new features. In

parallel with these maintenance releases, one or more future releases may be under development. People normally use the word “branch” to refer to one of these many slightly different directions in which development is proceeding.

Mercurial is particularly well suited to managing a number of simultaneous, but not identical, branches. Each “development direction” can live in its own central repository, and you can merge changes from one to another as the need arises. Because repositories are independent of each other, unstable changes in a development branch will never affect a stable branch unless someone explicitly merges those changes in.

Here’s an example of how this can work in practice. Let’s say you have one “main branch” on a central server.

```
1 $ hg init main
2 $ cd main
3 $ echo 'This is a boring feature.' > myfile
4 $ hg commit -A -m 'We have reached an important milestone!'
5 adding myfile
```

People clone it, make changes locally, test them, and push them back.

Once the main branch reaches a release milestone, you can use the “hg tag” command to give a permanent name to the milestone revision.

```
1 $ hg tag v1.0
2 $ hg tip
3 changeset: 1:c630d6d2a0b6
4 tag: tip
5 user: Bryan O'Sullivan <bos@serpentine.com>
6 date: Mon Dec 10 19:54:19 2007 +0000
7 summary: Added tag v1.0 for changeset b74ad3349882
8
9 $ hg tags
10 tip 1:c630d6d2a0b6
11 v1.0 0:b74ad3349882
```

Let’s say some ongoing development occurs on the main branch.

```
1 $ cd ../main
2 $ echo 'This is exciting and new!' >> myfile
3 $ hg commit -m 'Add a new feature'
4 $ cat myfile
5 This is a boring feature.
6 This is exciting and new!
```

Using the tag that was recorded at the milestone, people who clone that repository at any time in the future can use “hg update” to get a copy of the working directory exactly as it was when that tagged revision was committed.

```
1 $ cd ..
2 $ hg clone -U main main-old
3 $ cd main-old
4 $ hg update v1.0
5 1 files updated, 0 files merged, 0 files removed, 0 files unresolved
6 $ cat myfile
7 This is a boring feature.
```

In addition, immediately after the main branch is tagged, someone can then clone the main branch on the server to a new “stable” branch, also on the server.

```

1 $ cd ..
2 $ hg clone -rv1.0 main stable
3 requesting all changes
4 adding changesets
5 adding manifests
6 adding file changes
7 added 1 changesets with 1 changes to 1 files
8 1 files updated, 0 files merged, 0 files removed, 0 files unresolved

```

Someone who needs to make a change to the stable branch can then clone *that* repository, make their changes, commit, and push their changes back there.

```

1 $ hg clone stable stable-fix
2 1 files updated, 0 files merged, 0 files removed, 0 files unresolved
3 $ cd stable-fix
4 $ echo 'This is a fix to a boring feature.' > myfile
5 $ hg commit -m 'Fix a bug'
6 $ hg push
7 pushing to /tmp/branchingOTQG3Y/stable
8 searching for changes
9 adding changesets
10 adding manifests
11 adding file changes
12 added 1 changesets with 1 changes to 1 files

```

Because Mercurial repositories are independent, and Mercurial doesn't move changes around automatically, the stable and main branches are *isolated* from each other. The changes that you made on the main branch don't "leak" to the stable branch, and vice versa.

You'll often want all of your bugfixes on the stable branch to show up on the main branch, too. Rather than rewrite a bugfix on the main branch, you can simply pull and merge changes from the stable to the main branch, and Mercurial will bring those bugfixes in for you.

```

1 $ cd ../main
2 $ hg pull ../stable
3 pulling from ../stable
4 searching for changes
5 adding changesets
6 adding manifests
7 adding file changes
8 added 1 changesets with 1 changes to 1 files (+1 heads)
9 (run 'hg heads' to see heads, 'hg merge' to merge)
10 $ hg merge
11 merging myfile
12 0 files updated, 1 files merged, 0 files removed, 0 files unresolved
13 (branch merge, don't forget to commit)
14 $ hg commit -m 'Bring in bugfix from stable branch'
15 $ cat myfile
16 This is a fix to a boring feature.
17 This is exciting and new!

```

The main branch will still contain changes that are not on the stable branch, but it will also contain all of the bugfixes from the stable branch. The stable branch remains unaffected by these changes.

6.2.5 Feature branches

For larger projects, an effective way to manage change is to break up a team into smaller groups. Each group has a shared branch of its own, cloned from a single “master” branch used by the entire project. People working on an individual branch are typically quite isolated from developments on other branches.




Figure 6.1: Feature branches

When a particular feature is deemed to be in suitable shape, someone on that feature team pulls and merges from the master branch into the feature branch, then pushes back up to the master branch.

6.2.6 The release train

Some projects are organised on a “train” basis: a release is scheduled to happen every few months, and whatever features are ready when the “train” is ready to leave are allowed in.

This model resembles working with feature branches. The difference is that when a feature branch misses a train, someone on the feature team pulls and merges the changes that went out on that train release into the feature branch, and the team continues its work on top of that release so that their feature can make the next release.

6.2.7 The Linux kernel model

The development of the Linux kernel has a shallow hierarchical structure, surrounded by a cloud of apparent chaos. Because most Linux developers use `git`, a distributed revision control tool with capabilities similar to Mercurial, it’s useful to describe the way work flows in that environment; if you like the ideas, the approach translates well across tools.

At the center of the community sits Linus Torvalds, the creator of Linux. He publishes a single source repository that is considered the “authoritative” current tree by the entire developer community. Anyone can clone Linus’s tree, but he is very choosy about whose trees he pulls from.

Linus has a number of “trusted lieutenants”. As a general rule, he pulls whatever changes they publish, in most cases without even reviewing those changes. Some of those lieutenants are generally agreed to be “maintainers”, responsible for specific subsystems within the kernel. If a random kernel hacker wants to make a change to a subsystem that they want to end up in Linus’s tree, they must find out who the subsystem’s maintainer is, and ask that maintainer to take their change. If the maintainer reviews their changes and agrees to take them, they’ll pass them along to Linus in due course.

Individual lieutenants have their own approaches to reviewing, accepting, and publishing changes; and for deciding when to feed them to Linus. In addition, there are several well known branches that people use for different purposes. For example, a few people maintain “stable” repositories of older versions of the kernel, to which they apply critical fixes as needed. Some maintainers publish multiple trees: one for experimental changes; one for changes that they are about to feed upstream; and so on. Others just publish a single tree.

This model has two notable features. The first is that it’s “pull only”. You have to ask, convince, or beg another developer to take a change from you, because there are almost no trees to which more than one person can push, and there’s no way to push changes into a tree that someone else controls.

The second is that it's based on reputation and acclaim. If you're an unknown, Linus will probably ignore changes from you without even responding. But a subsystem maintainer will probably review them, and will likely take them if they pass their criteria for suitability. The more "good" changes you contribute to a maintainer, the more likely they are to trust your judgment and accept your changes. If you're well-known and maintain a long-lived branch for something Linus hasn't yet accepted, people with similar interests may pull your changes regularly to keep up with your work.

Reputation and acclaim don't necessarily cross subsystem or "people" boundaries. If you're a respected but specialised storage hacker, and you try to fix a networking bug, that change will receive a level of scrutiny from a network maintainer comparable to a change from a complete stranger.

To people who come from more orderly project backgrounds, the comparatively chaotic Linux kernel development process often seems completely insane. It's subject to the whims of individuals; people make sweeping changes whenever they deem it appropriate; and the pace of development is astounding. And yet Linux is a highly successful, well-regarded piece of software.

6.2.8 Pull-only versus shared-push collaboration

A perpetual source of heat in the open source community is whether a development model in which people only ever pull changes from others is "better than" one in which multiple people can push changes to a shared repository.

Typically, the backers of the shared-push model use tools that actively enforce this approach. If you're using a centralised revision control tool such as Subversion, there's no way to make a choice over which model you'll use: the tool gives you shared-push, and if you want to do anything else, you'll have to roll your own approach on top (such as applying a patch by hand).

A good distributed revision control tool, such as Mercurial, will support both models. You and your collaborators can then structure how you work together based on your own needs and preferences, not on what contortions your tools force you into.

6.2.9 Where collaboration meets branch management

Once you and your team set up some shared repositories and start propagating changes back and forth between local and shared repos, you begin to face a related, but slightly different challenge: that of managing the multiple directions in which your team may be moving at once. Even though this subject is intimately related to how your team collaborates, it's dense enough to merit treatment of its own, in chapter 8.

6.3 The technical side of sharing

The remainder of this chapter is devoted to the question of serving data to your collaborators.

6.4 Informal sharing with "hg serve"

Mercurial's "hg serve" command is wonderfully suited to small, tight-knit, and fast-paced group environments. It also provides a great way to get a feel for using Mercurial commands over a network.

Run "hg serve" inside a repository, and in under a second it will bring up a specialised HTTP server; this will accept connections from any client, and serve up data for that repository until you terminate it. Anyone who knows the URL of the server you just started, and can talk to your computer over the network, can then use a web browser or Mercurial to read data from that repository. A URL for a "hg serve" instance running on a laptop is likely to look something like `http://my-laptop.local:8000/`.

The "hg serve" command is *not* a general-purpose web server. It can do only two things:

- Allow people to browse the history of the repository it's serving, from their normal web browsers.
- Speak Mercurial's wire protocol, so that people can "hg clone" or "hg pull" changes from that repository.

In particular, “hg serve” won’t allow remote users to *modify* your repository. It’s intended for read-only use.

If you’re getting started with Mercurial, there’s nothing to prevent you from using “hg serve” to serve up a repository on your own computer, then use commands like “hg clone”, “hg incoming”, and so on to talk to that server as if the repository was hosted remotely. This can help you to quickly get acquainted with using commands on network-hosted repositories.

6.4.1 A few things to keep in mind

Because it provides unauthenticated read access to all clients, you should only use “hg serve” in an environment where you either don’t care, or have complete control over, who can access your network and pull data from your repository.

The “hg serve” command knows nothing about any firewall software you might have installed on your system or network. It cannot detect or control your firewall software. If other people are unable to talk to a running “hg serve” instance, the second thing you should do (*after* you make sure that they’re using the correct URL) is check your firewall configuration.

By default, “hg serve” listens for incoming connections on port 8000. If another process is already listening on the port you want to use, you can specify a different port to listen on using the `-p` option.

Normally, when “hg serve” starts, it prints no output, which can be a bit unnerving. If you’d like to confirm that it is indeed running correctly, and find out what URL you should send to your collaborators, start it with the `-v` option.

6.5 Using the Secure Shell (ssh) protocol

You can pull and push changes securely over a network connection using the Secure Shell (ssh) protocol. To use this successfully, you may have to do a little bit of configuration on the client or server sides.

If you’re not familiar with ssh, it’s a network protocol that lets you securely communicate with another computer. To use it with Mercurial, you’ll be setting up one or more user accounts on a server so that remote users can log in and execute commands.

(If you *are* familiar with ssh, you’ll probably find some of the material that follows to be elementary in nature.)

6.5.1 How to read and write ssh URLs

An ssh URL tends to look like this:

```
ssh://bos@hg.serpentine.com:22/hg/hgbook
```

1. The “ssh://” part tells Mercurial to use the ssh protocol.
2. The “bos@” component indicates what username to log into the server as. You can leave this out if the remote username is the same as your local username.
3. The “hg.serpentine.com” gives the hostname of the server to log into.
4. The “:22” identifies the port number to connect to the server on. The default port is 22, so you only need to specify this part if you’re *not* using port 22.
5. The remainder of the URL is the local path to the repository on the server.

There’s plenty of scope for confusion with the path component of ssh URLs, as there is no standard way for tools to interpret it. Some programs behave differently than others when dealing with these paths. This isn’t an ideal situation, but it’s unlikely to change. Please read the following paragraphs carefully.

Mercurial treats the path to a repository on the server as relative to the remote user’s home directory. For example, if user `foo` on the server has a home directory of `/home/foo`, then an ssh URL that contains a path component of `bar` *really* refers to the directory `/home/foo/bar`.

If you want to specify a path relative to another user's home directory, you can use a path that starts with a tilde character followed by the user's name (let's call them `otheruser`), like this.

```
1 ssh://server/~otheruser/hg/repo
```

And if you really want to specify an *absolute* path on the server, begin the path component with two slashes, as in this example.

```
1 ssh://server//absolute/path
```

6.5.2 Finding an ssh client for your system

Almost every Unix-like system comes with OpenSSH preinstalled. If you're using such a system, run `which ssh` to find out if the `ssh` command is installed (it's usually in `/usr/bin`). In the unlikely event that it isn't present, take a look at your system documentation to figure out how to install it.

On Windows, you'll first need to choose download a suitable ssh client. There are two alternatives.

- Simon Tatham's excellent PuTTY package [\[Tat\]](#) provides a complete suite of ssh client commands.
- If you have a high tolerance for pain, you can use the Cygwin port of OpenSSH.

In either case, you'll need to edit your `Mercurial.ini` file to tell Mercurial where to find the actual client command. For example, if you're using PuTTY, you'll need to use the `plink` command as a command-line ssh client.

```
1 [ui]  
2 ssh = C:/path/to/plink.exe -ssh -i "C:/path/to/my/private/key"
```

Note: The path to `plink` shouldn't contain any whitespace characters, or Mercurial may not be able to run it correctly (so putting it in `C:\Program Files` is probably not a good idea).

6.5.3 Generating a key pair

To avoid the need to repetitively type a password every time you need to use your ssh client, I recommend generating a key pair. On a Unix-like system, the `ssh-keygen` command will do the trick. On Windows, if you're using PuTTY, the `puttygen` command is what you'll need.

When you generate a key pair, it's usually *highly* advisable to protect it with a passphrase. (The only time that you might not want to do this is when you're using the ssh protocol for automated tasks on a secure network.)

Simply generating a key pair isn't enough, however. You'll need to add the public key to the set of authorised keys for whatever user you're logging in remotely as. For servers using OpenSSH (the vast majority), this will mean adding the public key to a list in a file called `authorized_keys` in their `.ssh` directory.

On a Unix-like system, your public key will have a `.pub` extension. If you're using `puttygen` on Windows, you can save the public key to a file of your choosing, or paste it from the window it's displayed in straight into the `authorized_keys` file.

6.5.4 Using an authentication agent

An authentication agent is a daemon that stores passphrases in memory (so it will forget passphrases if you log out and log back in again). An ssh client will notice if it's running, and query it for a passphrase. If there's no authentication agent running, or the agent doesn't store the necessary passphrase, you'll have to type your passphrase every time Mercurial tries to communicate with a server on your behalf (e.g. whenever you pull or push changes).

The downside of storing passphrases in an agent is that it's possible for a well-prepared attacker to recover the plain text of your passphrases, in some cases even if your system has been power-cycled. You should make your own judgment as to whether this is an acceptable risk. It certainly saves a lot of repeated typing.

On Unix-like systems, the agent is called `ssh-agent`, and it's often run automatically for you when you log in. You'll need to use the `ssh-add` command to add passphrases to the agent's store. On Windows, if you're using PuTTY, the `pageant` command acts as the agent. It adds an icon to your system tray that will let you manage stored passphrases.

6.5.5 Configuring the server side properly

Because `ssh` can be fiddly to set up if you're new to it, there's a variety of things that can go wrong. Add Mercurial on top, and there's plenty more scope for head-scratching. Most of these potential problems occur on the server side, not the client side. The good news is that once you've gotten a configuration working, it will usually continue to work indefinitely.

Before you try using Mercurial to talk to an `ssh` server, it's best to make sure that you can use the normal `ssh` or `putty` command to talk to the server first. If you run into problems with using these commands directly, Mercurial surely won't work. Worse, it will obscure the underlying problem. Any time you want to debug `ssh`-related Mercurial problems, you should drop back to making sure that plain `ssh` client commands work first, *before* you worry about whether there's a problem with Mercurial.

The first thing to be sure of on the server side is that you can actually log in from another machine at all. If you can't use `ssh` or `putty` to log in, the error message you get may give you a few hints as to what's wrong. The most common problems are as follows.

- If you get a “connection refused” error, either there isn't an SSH daemon running on the server at all, or it's inaccessible due to firewall configuration.
- If you get a “no route to host” error, you either have an incorrect address for the server or a seriously locked down firewall that won't admit its existence at all.
- If you get a “permission denied” error, you may have mistyped the username on the server, or you could have mistyped your key's passphrase or the remote user's password.

In summary, if you're having trouble talking to the server's `ssh` daemon, first make sure that one is running at all. On many systems it will be installed, but disabled, by default. Once you're done with this step, you should then check that the server's firewall is configured to allow incoming connections on the port the `ssh` daemon is listening on (usually 22). Don't worry about more exotic possibilities for misconfiguration until you've checked these two first.

If you're using an authentication agent on the client side to store passphrases for your keys, you ought to be able to log into the server without being prompted for a passphrase or a password. If you're prompted for a passphrase, there are a few possible culprits.

- You might have forgotten to use `ssh-add` or `pageant` to store the passphrase.
- You might have stored the passphrase for the wrong key.

If you're being prompted for the remote user's password, there are another few possible problems to check.

- Either the user's home directory or their `.ssh` directory might have excessively liberal permissions. As a result, the `ssh` daemon will not trust or read their `authorized_keys` file. For example, a group-writable home or `.ssh` directory will often cause this symptom.
- The user's `authorized_keys` file may have a problem. If anyone other than the user owns or can write to that file, the `ssh` daemon will not trust or read it.

In the ideal world, you should be able to run the following command successfully, and it should print exactly one line of output, the current date and time.

```
ssh myserver date
```

If, on your server, you have login scripts that print banners or other junk even when running non-interactive commands like this, you should fix them before you continue, so that they only print output if they're run interactively. Otherwise these banners will at least clutter up Mercurial's output. Worse, they could potentially cause problems with running Mercurial commands remotely. Mercurial makes tries to detect and ignore banners in non-interactive `ssh` sessions, but it is not foolproof. (If you're editing your login scripts on your server, the usual way to see if a login script is running in an interactive shell is to check the return code from the command `tty -s`.)

Once you've verified that plain old `ssh` is working with your server, the next step is to ensure that Mercurial runs on the server. The following command should run successfully:

```
1 ssh myserver hg version
```

If you see an error message instead of normal “`hg version`” output, this is usually because you haven't installed Mercurial to `/usr/bin`. Don't worry if this is the case; you don't need to do that. But you should check for a few possible problems.

- Is Mercurial really installed on the server at all? I know this sounds trivial, but it's worth checking!
- Maybe your shell's search path (usually set via the `PATH` environment variable) is simply misconfigured.
- Perhaps your `PATH` environment variable is only being set to point to the location of the `hg` executable if the login session is interactive. This can happen if you're setting the path in the wrong shell login script. See your shell's documentation for details.
- The `PYTHONPATH` environment variable may need to contain the path to the Mercurial Python modules. It might not be set at all; it could be incorrect; or it may be set only if the login is interactive.

If you can run “`hg version`” over an `ssh` connection, well done! You've got the server and client sorted out. You should now be able to use Mercurial to access repositories hosted by that username on that server. If you run into problems with Mercurial and `ssh` at this point, try using the `--debug` option to get a clearer picture of what's going on.

6.5.6 Using compression with `ssh`

Mercurial does not compress data when it uses the `ssh` protocol, because the `ssh` protocol can transparently compress data. However, the default behaviour of `ssh` clients is *not* to request compression.

Over any network other than a fast LAN (even a wireless network), using compression is likely to significantly speed up Mercurial's network operations. For example, over a WAN, someone measured compression as reducing the amount of time required to clone a particularly large repository from 51 minutes to 17 minutes.

Both `ssh` and `plink` accept a `-C` option which turns on compression. You can easily edit your `hgrc` to enable compression for all of Mercurial's uses of the `ssh` protocol.

```
1 [ui]  
2 ssh = ssh -C
```

If you use `ssh`, you can configure it to always use compression when talking to your server. To do this, edit your `.ssh/config` file (which may not yet exist), as follows.

```
1 Host hg  
2     Compression yes  
3     HostName hg.example.com
```

This defines an alias, `hg`. When you use it on the `ssh` command line or in a Mercurial `ssh`-protocol URL, it will cause `ssh` to connect to `hg.example.com` and use compression. This gives you both a shorter name to type and compression, each of which is a good thing in its own right.

6.6 Serving over HTTP using CGI

Depending on how ambitious you are, configuring Mercurial’s CGI interface can take anything from a few moments to several hours.

We’ll begin with the simplest of examples, and work our way towards a more complex configuration. Even for the most basic case, you’re almost certainly going to need to read and modify your web server’s configuration.

Note: Configuring a web server is a complex, fiddly, and highly system-dependent activity. I can’t possibly give you instructions that will cover anything like all of the cases you will encounter. Please use your discretion and judgment in following the sections below. Be prepared to make plenty of mistakes, and to spend a lot of time reading your server’s error logs.

6.6.1 Web server configuration checklist

Before you continue, do take a few moments to check a few aspects of your system’s setup.

1. Do you have a web server installed at all? Mac OS X ships with Apache, but many other systems may not have a web server installed.
2. If you have a web server installed, is it actually running? On most systems, even if one is present, it will be disabled by default.
3. Is your server configured to allow you to run CGI programs in the directory where you plan to do so? Most servers default to explicitly disabling the ability to run CGI programs.

If you don’t have a web server installed, and don’t have substantial experience configuring Apache, you should consider using the `lighttpd` web server instead of Apache. Apache has a well-deserved reputation for baroque and confusing configuration. While `lighttpd` is less capable in some ways than Apache, most of these capabilities are not relevant to serving Mercurial repositories. And `lighttpd` is undeniably *much* easier to get started with than Apache.

6.6.2 Basic CGI configuration

On Unix-like systems, it’s common for users to have a subdirectory named something like `public_html` in their home directory, from which they can serve up web pages. A file named `foo` in this directory will be accessible at a URL of the form `http://www.example.com/username/foo`.

To get started, find the `hgweb.cgi` script that should be present in your Mercurial installation. If you can’t quickly find a local copy on your system, simply download one from the master Mercurial repository at <http://www.selenic.com/repo/hg/raw-file/tip/hgweb.cgi>.

You’ll need to copy this script into your `public_html` directory, and ensure that it’s executable.

```
1 cp ../hgweb.cgi ~/public_html
2 chmod 755 ~/public_html/hgweb.cgi
```

The `755` argument to `chmod` is a little more general than just making the script executable: it ensures that the script is executable by anyone, and that “group” and “other” write permissions are *not* set. If you were to leave those write permissions enabled, Apache’s `suexec` subsystem would likely refuse to execute the script. In fact, `suexec` also insists that the *directory* in which the script resides must not be writable by others.

```
1 chmod 755 ~/public_html
```

What could *possibly* go wrong?

Once you've copied the CGI script into place, go into a web browser, and try to open the URL <http://myhostname/~myuser/hgweb.cgi>, but brace yourself for instant failure. There's a high probability that trying to visit this URL will fail, and there are many possible reasons for this. In fact, you're likely to stumble over almost every one of the possible errors below, so please read carefully. The following are all of the problems I ran into on a system running Fedora 7, with a fresh installation of Apache, and a user account that I created specially to perform this exercise.

Your web server may have per-user directories disabled. If you're using Apache, search your config file for a `UserDir` directive. If there's none present, per-user directories will be disabled. If one exists, but its value is disabled, then per-user directories will be disabled. Otherwise, the string after `UserDir` gives the name of the subdirectory that Apache will look in under your home directory, for example `public_html`.

Your file access permissions may be too restrictive. The web server must be able to traverse your home directory and directories under your `public_html` directory, and read files under the latter too. Here's a quick recipe to help you to make your permissions more appropriate.

```
1 chmod 755 ~
2 find ~/public_html -type d -print0 | xargs -0r chmod 755
3 find ~/public_html -type f -print0 | xargs -0r chmod 644
```

The other possibility with permissions is that you might get a completely empty window when you try to load the script. In this case, it's likely that your access permissions are *too permissive*. Apache's `suexec` subsystem won't execute a script that's group- or world-writable, for example.

Your web server may be configured to disallow execution of CGI programs in your per-user web directory. Here's Apache's default per-user configuration from my Fedora system.

```
1 <Directory /home/*/public_html>
2     AllowOverride FileInfo AuthConfig Limit
3     Options MultiViews Indexes SymLinksIfOwnerMatch IncludesNoExec
4     <Limit GET POST OPTIONS>
5         Order allow,deny
6         Allow from all
7     </Limit>
8     <LimitExcept GET POST OPTIONS>
9         Order deny,allow
10        Deny from all
11    </LimitExcept>
12 </Directory>
```

If you find a similar-looking `Directory` group in your Apache configuration, the directive to look at inside it is `Options`. Add `ExecCGI` to the end of this list if it's missing, and restart the web server.

If you find that Apache serves you the text of the CGI script instead of executing it, you may need to either uncomment (if already present) or add a directive like this.

```
1 AddHandler cgi-script .cgi
```

The next possibility is that you might be served with a colourful Python traceback claiming that it can't import a mercurial-related module. This is actually progress! The server is now capable of executing your CGI script. This error is only likely to occur if you're running a private installation of Mercurial, instead of a system-wide version. Remember that the web server runs the CGI program without any of the environment variables that you take for granted in an interactive session. If this error happens to you, edit your copy of `hgweb.cgi` and follow the directions inside it to correctly set your `PYTHONPATH` environment variable.

Finally, you are *certain* to be served with another colourful Python backtrace: this one will complain that it can't find `/path/to/repository`. Edit your `hgweb.cgi` script and replace the `/path/to/repository` string with the complete path to the repository you want to serve up.

At this point, when you try to reload the page, you should be presented with a nice HTML view of your repository's history. Whew!

Configuring `lighttpd`

To be exhaustive in my experiments, I tried configuring the increasingly popular `lighttpd` web server to serve the same repository as I described with Apache above. I had already overcome all of the problems I outlined with Apache, many of which are not server-specific. As a result, I was fairly sure that my file and directory permissions were good, and that my `hgweb.cgi` script was properly edited.

Once I had Apache running, getting `lighttpd` to serve the repository was a snap (in other words, even if you're trying to use `lighttpd`, you should read the Apache section). I first had to edit the `mod_access` section of its config file to enable `mod_cgi` and `mod_userdir`, both of which were disabled by default on my system. I then added a few lines to the end of the config file, to configure these modules.

```
1 userdir.path = "public_html"
2 cgi.assign = ( ".cgi" => "" )
```

With this done, `lighttpd` ran immediately for me. If I had configured `lighttpd` before Apache, I'd almost certainly have run into many of the same system-level configuration problems as I did with Apache. However, I found `lighttpd` to be noticeably easier to configure than Apache, even though I've used Apache for over a decade, and this was my first exposure to `lighttpd`.

6.6.3 Sharing multiple repositories with one CGI script

The `hgweb.cgi` script only lets you publish a single repository, which is an annoying restriction. If you want to publish more than one without wracking yourself with multiple copies of the same script, each with different names, a better choice is to use the `hgwebdir.cgi` script.

The procedure to configure `hgwebdir.cgi` is only a little more involved than for `hgweb.cgi`. First, you must obtain a copy of the script. If you don't have one handy, you can download a copy from the master Mercurial repository at <http://www.selenic.com/repo/hg/raw-file/tip/hgwebdir.cgi>.

You'll need to copy this script into your `public_html` directory, and ensure that it's executable.

```
1 cp ../hgwebdir.cgi ~/public_html
2 chmod 755 ~/public_html ~/public_html/hgwebdir.cgi
```

With basic configuration out of the way, try to visit <http://myhostname/~myuser/hgwebdir.cgi> in your browser. It should display an empty list of repositories. If you get a blank window or error message, try walking through the list of potential problems in section 6.6.2.

The `hgwebdir.cgi` script relies on an external configuration file. By default, it searches for a file named `hgweb.config` in the same directory as itself. You'll need to create this file, and make it world-readable. The format of the file is similar to a Windows "ini" file, as understood by Python's `ConfigParser` [Pyt] module.

The easiest way to configure `hgwebdir.cgi` is with a section named `collections`. This will automatically publish *every* repository under the directories you name. The section should look like this:

```
1 [collections]
2 /my/root = /my/root
```

Mercurial interprets this by looking at the directory name on the *right* hand side of the "=" sign; finding repositories in that directory hierarchy; and using the text on the *left* to strip off matching text from the names it will actually list in the web interface. The remaining component of a path after this stripping has occurred is called a "virtual path".

Given the example above, if we have a repository whose local path is `/my/root/this/repo`, the CGI script will strip the leading `/my/root` from the name, and publish the repository with a virtual path of `this/repo`. If the base URL for our CGI script is `http://myhostname/~myuser/hgwebdir.cgi`, the complete URL for that repository will be `http://myhostname/~myuser/hgwebdir.cgi/this/repo`.

If we replace `/my/root` on the left hand side of this example with `/my`, then `hgwebdir.cgi` will only strip off `/my` from the repository name, and will give us a virtual path of `root/this/repo` instead of `this/repo`.

The `hgwebdir.cgi` script will recursively search each directory listed in the `collections` section of its configuration file, but it will not recurse into the repositories it finds.

The `collections` mechanism makes it easy to publish many repositories in a “fire and forget” manner. You only need to set up the CGI script and configuration file one time. Afterwards, you can publish or unpublish a repository at any time by simply moving it into, or out of, the directory hierarchy in which you’ve configured `hgwebdir.cgi` to look.

Explicitly specifying which repositories to publish

In addition to the `collections` mechanism, the `hgwebdir.cgi` script allows you to publish a specific list of repositories. To do so, create a `paths` section, with contents of the following form.

```
1 [paths]
2 repo1 = /my/path/to/some/repo
3 repo2 = /some/path/to/another
```

In this case, the virtual path (the component that will appear in a URL) is on the left hand side of each definition, while the path to the repository is on the right. Notice that there does not need to be any relationship between the virtual path you choose and the location of a repository in your filesystem.

If you wish, you can use both the `collections` and `paths` mechanisms simultaneously in a single configuration file.

Note: If multiple repositories have the same virtual path, `hgwebdir.cgi` will not report an error. Instead, it will behave unpredictably.

6.6.4 Downloading source archives

Mercurial’s web interface lets users download an archive of any revision. This archive will contain a snapshot of the working directory as of that revision, but it will not contain a copy of the repository data.

By default, this feature is not enabled. To enable it, you’ll need to add an `allow_archive` item to the `[web]` section of your `hgrc`.

6.6.5 Web configuration options

Mercurial’s web interfaces (the “`hg serve`” command, and the `hgweb.cgi` and `hgwebdir.cgi` scripts) have a number of configuration options that you can set. These belong in a section named `[web]`.

`allow_archive` Determines which (if any) archive download mechanisms Mercurial supports. If you enable this feature, users of the web interface will be able to download an archive of whatever revision of a repository they are viewing. To enable the archive feature, this item must take the form of a sequence of words drawn from the list below.

`bz2` A tar archive, compressed using `bzip2` compression. This has the best compression ratio, but uses the most CPU time on the server.

`gz` A tar archive, compressed using `gzip` compression.

`zip` A zip archive, compressed using `LZW` compression. This format has the worst compression ratio, but is widely used in the Windows world.

If you provide an empty list, or don't have an `allow_archive` entry at all, this feature will be disabled. Here is an example of how to enable all three supported formats.

```
1 [web]
2 allow_archive = bz2 gz zip
```

`allowpull` Boolean. Determines whether the web interface allows remote users to “hg pull” and “hg clone” this repository over HTTP. If set to `no` or `false`, only the “human-oriented” portion of the web interface is available.

`contact` String. A free-form (but preferably brief) string identifying the person or group in charge of the repository. This often contains the name and email address of a person or mailing list. It often makes sense to place this entry in a repository's own `.hg/hgrc` file, but it can make sense to use in a global `hgrc` if every repository has a single maintainer.

`maxchanges` Integer. The default maximum number of changesets to display in a single page of output.

`maxfiles` Integer. The default maximum number of modified files to display in a single page of output.

`stripes` Integer. If the web interface displays alternating “stripes” to make it easier to visually align rows when you are looking at a table, this number controls the number of rows in each stripe.

`style` Controls the template Mercurial uses to display the web interface. Mercurial ships with two web templates, named `default` and `gitweb` (the latter is much more visually attractive). You can also specify a custom template of your own; see chapter 11 for details. Here, you can see how to enable the `gitweb` style.

```
1 [web]
2 style = gitweb
```

`templates` Path. The directory in which to search for template files. By default, Mercurial searches in the directory in which it was installed.

If you are using `hgwebdir.cgi`, you can place a few configuration items in a `[web]` section of the `hgweb.config` file instead of a `hgrc` file, for convenience. These items are `motd` and `style`.

Options specific to an individual repository

A few `[web]` configuration items ought to be placed in a repository's local `.hg/hgrc`, rather than a user's or global `hgrc`.

`description` String. A free-form (but preferably brief) string that describes the contents or purpose of the repository.

`name` String. The name to use for the repository in the web interface. This overrides the default name, which is the last component of the repository's path.

Options specific to the “hg serve” command

Some of the items in the `[web]` section of a `hgrc` file are only for use with the “hg serve” command.

`accesslog` Path. The name of a file into which to write an access log. By default, the “hg serve” command writes this information to standard output, not to a file. Log entries are written in the standard “combined” file format used by almost all web servers.

`address` String. The local address on which the server should listen for incoming connections. By default, the server listens on all addresses.

`errorlog` Path. The name of a file into which to write an error log. By default, the “hg serve” command writes this information to standard error, not to a file.

`ipv6` Boolean. Whether to use the IPv6 protocol. By default, IPv6 is not used.

`port` Integer. The TCP port number on which the server should listen. The default port number used is 8000.

Choosing the right `hgrc` file to add `[web]` items to

It is important to remember that a web server like Apache or `lighttpd` will run under a user ID that is different to yours. CGI scripts run by your server, such as `hgweb.cgi`, will usually also run under that user ID.

If you add `[web]` items to your own personal `hgrc` file, CGI scripts won't read that `hgrc` file. Those settings will thus only affect the behaviour of the "`hg serve`" command when you run it. To cause CGI scripts to see your settings, either create a `hgrc` file in the home directory of the user ID that runs your web server, or add those settings to a system-wide `hgrc` file.

Chapter 7

File names and pattern matching

Mercurial provides mechanisms that let you work with file names in a consistent and expressive way.

7.1 Simple file naming

Mercurial uses a unified piece of machinery “under the hood” to handle file names. Every command behaves uniformly with respect to file names. The way in which commands work with file names is as follows.

If you explicitly name real files on the command line, Mercurial works with exactly those files, as you would expect.

```
1 $ hg add COPYING README examples/simple.py
```

When you provide a directory name, Mercurial will interpret this as “operate on every file in this directory and its subdirectories”. Mercurial traverses the files and subdirectories in a directory in alphabetical order. When it encounters a subdirectory, it will traverse that subdirectory before continuing with the current directory.

```
1 $ hg status src
2 ? src/main.py
3 ? src/watcher/_watcher.c
4 ? src/watcher/watcher.py
5 ? src/xyzzy.txt
```

7.2 Running commands without any file names

Mercurial’s commands that work with file names have useful default behaviours when you invoke them without providing any file names or patterns. What kind of behaviour you should expect depends on what the command does. Here are a few rules of thumb you can use to predict what a command is likely to do if you don’t give it any names to work with.

- Most commands will operate on the entire working directory. This is what the “hg add” command does, for example.
- If the command has effects that are difficult or impossible to reverse, it will force you to explicitly provide at least one name or pattern (see below). This protects you from accidentally deleting files by running “hg remove” with no arguments, for example.

It’s easy to work around these default behaviours if they don’t suit you. If a command normally operates on the whole working directory, you can invoke it on just the current directory and its subdirectories by giving it the name “.”.

```

1 $ cd src
2 $ hg add -n
3 adding ../MANIFEST.in
4 adding ../examples/performant.py
5 adding ../setup.py
6 adding main.py
7 adding watcher/_watcher.c
8 adding watcher/watcher.py
9 adding xyzzy.txt
10 $ hg add -n .
11 adding main.py
12 adding watcher/_watcher.c
13 adding watcher/watcher.py
14 adding xyzzy.txt

```

Along the same lines, some commands normally print file names relative to the root of the repository, even if you’re invoking them from a subdirectory. Such a command will print file names relative to your subdirectory if you give it explicit names. Here, we’re going to run “hg status” from a subdirectory, and get it to operate on the entire working directory while printing file names relative to our subdirectory, by passing it the output of the “hg root” command.

```

1 $ hg status
2 A COPYING
3 A README
4 A examples/simple.py
5 ? MANIFEST.in
6 ? examples/performant.py
7 ? setup.py
8 ? src/main.py
9 ? src/watcher/_watcher.c
10 ? src/watcher/watcher.py
11 ? src/xyzzy.txt
12 $ hg status `hg root`
13 A ../COPYING
14 A ../README
15 A ../examples/simple.py
16 ? ../MANIFEST.in
17 ? ../examples/performant.py
18 ? ../setup.py
19 ? main.py
20 ? watcher/_watcher.c
21 ? watcher/watcher.py
22 ? xyzzy.txt

```

7.3 Telling you what’s going on

The “hg add” example in the preceding section illustrates something else that’s helpful about Mercurial commands. If a command operates on a file that you didn’t name explicitly on the command line, it will usually print the name of the file, so that you will not be surprised what’s going on.

The principle here is of *least surprise*. If you’ve exactly named a file on the command line, there’s no point in repeating it back at you. If Mercurial is acting on a file *implicitly*, because you provided no names, or a directory, or a pattern (see below), it’s safest to tell you what it’s doing.

For commands that behave this way, you can silence them using the `-q` option. You can also get them to print the name of every file, even those you’ve named explicitly, using the `-v` option.

7.4 Using patterns to identify files

In addition to working with file and directory names, Mercurial lets you use *patterns* to identify files. Mercurial’s pattern handling is expressive.

On Unix-like systems (Linux, MacOS, etc.), the job of matching file names to patterns normally falls to the shell. On these systems, you must explicitly tell Mercurial that a name is a pattern. On Windows, the shell does not expand patterns, so Mercurial will automatically identify names that are patterns, and expand them for you.

To provide a pattern in place of a regular name on the command line, the mechanism is simple:

```
1 syntax:patternbody
```

That is, a pattern is identified by a short text string that says what kind of pattern this is, followed by a colon, followed by the actual pattern.

Mercurial supports two kinds of pattern syntax. The most frequently used is called `glob`; this is the same kind of pattern matching used by the Unix shell, and should be familiar to Windows command prompt users, too.

When Mercurial does automatic pattern matching on Windows, it uses `glob` syntax. You can thus omit the “`glob:`” prefix on Windows, but it’s safe to use it, too.

The `re` syntax is more powerful; it lets you specify patterns using regular expressions, also known as `regexps`.

By the way, in the examples that follow, notice that I’m careful to wrap all of my patterns in quote characters, so that they won’t get expanded by the shell before Mercurial sees them.

7.4.1 Shell-style `glob` patterns

This is an overview of the kinds of patterns you can use when you’re matching on `glob` patterns.

The “`*`” character matches any string, within a single directory.

```
1 $ hg add 'glob:*.py'
2 adding main.py
```

The “`***`” pattern matches any string, and crosses directory boundaries. It’s not a standard Unix `glob` token, but it’s accepted by several popular Unix shells, and is very useful.

```
1 $ cd ..
2 $ hg status 'glob:***.py'
3 A examples/simple.py
4 A src/main.py
5 ? examples/performant.py
6 ? setup.py
7 ? src/watcher/watcher.py
```

The “`?`” pattern matches any single character.

```
1 $ hg status 'glob:***.?'
2 ? src/watcher/_watcher.c
```

The “[” character begins a *character class*. This matches any single character within the class. The class ends with a “]” character. A class may contain multiple *ranges* of the form “a-f”, which is shorthand for “abcdef”.

```
1 $ hg status 'glob:**[nr-t]'
```

```
2 ? MANIFEST.in
```

```
3 ? src/xyzyzy.txt
```

If the first character after the “[” in a character class is a “!”, it *negates* the class, making it match any single character not in the class.

A “{” begins a group of subpatterns, where the whole group matches if any subpattern in the group matches. The “,” character separates subpatterns, and “}” ends the group.

```
1 $ hg status 'glob:*. {in,py}'
```

```
2 ? MANIFEST.in
```

```
3 ? setup.py
```

Watch out!

Don’t forget that if you want to match a pattern in any directory, you should not be using the “*” match-any token, as this will only match within one directory. Instead, use the “***” token. This small example illustrates the difference between the two.

```
1 $ hg status 'glob:*.py'
```

```
2 ? setup.py
```

```
3 $ hg status 'glob:***.py'
```

```
4 A examples/simple.py
```

```
5 A src/main.py
```

```
6 ? examples/performant.py
```

```
7 ? setup.py
```

```
8 ? src/watcher/watcher.py
```

7.4.2 Regular expression matching with re patterns

Mercurial accepts the same regular expression syntax as the Python programming language (it uses Python’s regexp engine internally). This is based on the Perl language’s regexp syntax, which is the most popular dialect in use (it’s also used in Java, for example).

I won’t discuss Mercurial’s regexp dialect in any detail here, as regexps are not often used. Perl-style regexps are in any case already exhaustively documented on a multitude of web sites, and in many books. Instead, I will focus here on a few things you should know if you find yourself needing to use regexps with Mercurial.

A regexp is matched against an entire file name, relative to the root of the repository. In other words, even if you’re already in subdirectory `foo`, if you want to match files under this directory, your pattern must start with “`foo/`”.

One thing to note, if you’re familiar with Perl-style regexps, is that Mercurial’s are *rooted*. That is, a regexp starts matching against the beginning of a string; it doesn’t look for a match anywhere within the string. To match anywhere in a string, start your pattern with “`.*`”.

7.5 Filtering files

Not only does Mercurial give you a variety of ways to specify files; it lets you further winnow those files using *filters*. Commands that work with file names accept two filtering options.

- `-I`, or `--include`, lets you specify a pattern that file names must match in order to be processed.

- `-X`, or `--exclude`, gives you a way to *avoid* processing files, if they match this pattern.

You can provide multiple `-I` and `-X` options on the command line, and intermix them as you please. Mercurial interprets the patterns you provide using glob syntax by default (but you can use regexps if you need to).

You can read a `-I` filter as “process only the files that match this filter”.

```
1 $ hg status -I '*.in'
2 ? MANIFEST.in
```

The `-X` filter is best read as “process only the files that don’t match this pattern”.

```
1 $ hg status -X '**.py' src
2 ? src/watcher/_watcher.c
3 ? src/xyzzzy.txt
```

7.6 Ignoring unwanted files and directories

XXX.

7.7 Case sensitivity

If you’re working in a mixed development environment that contains both Linux (or other Unix) systems and Macs or Windows systems, you should keep in the back of your mind the knowledge that they treat the case (“N” versus “n”) of file names in incompatible ways. This is not very likely to affect you, and it’s easy to deal with if it does, but it could surprise you if you don’t know about it.

Operating systems and filesystems differ in the way they handle the *case* of characters in file and directory names. There are three common ways to handle case in names.

- Completely case insensitive. Uppercase and lowercase versions of a letter are treated as identical, both when creating a file and during subsequent accesses. This is common on older DOS-based systems.
- Case preserving, but insensitive. When a file or directory is created, the case of its name is stored, and can be retrieved and displayed by the operating system. When an existing file is being looked up, its case is ignored. This is the standard arrangement on Windows and MacOS. The names `f00` and `F00` identify the same file. This treatment of uppercase and lowercase letters as interchangeable is also referred to as *case folding*.
- Case sensitive. The case of a name is significant at all times. The names `f00` and `F00` identify different files. This is the way Linux and Unix systems normally work.

On Unix-like systems, it is possible to have any or all of the above ways of handling case in action at once. For example, if you use a USB thumb drive formatted with a FAT32 filesystem on a Linux system, Linux will handle names on that filesystem in a case preserving, but insensitive, way.

7.7.1 Safe, portable repository storage

Mercurial’s repository storage mechanism is *case safe*. It translates file names so that they can be safely stored on both case sensitive and case insensitive filesystems. This means that you can use normal file copying tools to transfer a Mercurial repository onto, for example, a USB thumb drive, and safely move that drive and repository back and forth between a Mac, a PC running Windows, and a Linux box.

7.7.2 Detecting case conflicts

When operating in the working directory, Mercurial honours the naming policy of the filesystem where the working directory is located. If the filesystem is case preserving, but insensitive, Mercurial will treat names that differ only in case as the same.

An important aspect of this approach is that it is possible to commit a changeset on a case sensitive (typically Linux or Unix) filesystem that will cause trouble for users on case insensitive (usually Windows and MacOS) users. If a Linux user commits changes to two files, one named `myfile.c` and the other named `MyFile.C`, they will be stored correctly in the repository. And in the working directories of other Linux users, they will be correctly represented as separate files.

If a Windows or Mac user pulls this change, they will not initially have a problem, because Mercurial's repository storage mechanism is case safe. However, once they try to “`hg update`” the working directory to that changeset, or “`hg merge`” with that changeset, Mercurial will spot the conflict between the two file names that the filesystem would treat as the same, and forbid the update or merge from occurring.

7.7.3 Fixing a case conflict

If you are using Windows or a Mac in a mixed environment where some of your collaborators are using Linux or Unix, and Mercurial reports a case folding conflict when you try to “`hg update`” or “`hg merge`”, the procedure to fix the problem is simple.

Just find a nearby Linux or Unix box, clone the problem repository onto it, and use Mercurial's “`hg rename`” command to change the names of any offending files or directories so that they will no longer cause case folding conflicts. Commit this change, “`hg pull`” or “`hg push`” it across to your Windows or MacOS system, and “`hg update`” to the revision with the non-conflicting names.

The changeset with case-conflicting names will remain in your project's history, and you still won't be able to “`hg update`” your working directory to that changeset on a Windows or MacOS system, but you can continue development unimpeded.

<p>Note: Prior to version 0.9.3, Mercurial did not use a case safe repository storage mechanism, and did not detect case folding conflicts. If you are using an older version of Mercurial on Windows or MacOS, I strongly recommend that you upgrade.</p>

Chapter 8

Managing releases and branchy development

Mercurial provides several mechanisms for you to manage a project that is making progress on multiple fronts at once. To understand these mechanisms, let's first take a brief look at a fairly normal software project structure.

Many software projects issue periodic “major” releases that contain substantial new features. In parallel, they may issue “minor” releases. These are usually identical to the major releases off which they're based, but with a few bugs fixed.

In this chapter, we'll start by talking about how to keep records of project milestones such as releases. We'll then continue on to talk about the flow of work between different phases of a project, and how Mercurial can help you to isolate and manage this work.

8.1 Giving a persistent name to a revision

Once you decide that you'd like to call a particular revision a “release”, it's a good idea to record the identity of that revision. This will let you reproduce that release at a later date, for whatever purpose you might need at the time (reproducing a bug, porting to a new platform, etc).

```
1 $ hg init mytag
2 $ cd mytag
3 $ echo hello > myfile
4 $ hg commit -A -m 'Initial commit'
5 adding myfile
```

Mercurial lets you give a permanent name to any revision using the “hg tag” command. Not surprisingly, these names are called “tags”.

```
1 $ hg tag v1.0
```

A tag is nothing more than a “symbolic name” for a revision. Tags exist purely for your convenience, so that you have a handy permanent way to refer to a revision; Mercurial doesn't interpret the tag names you use in any way. Neither does Mercurial place any restrictions on the name of a tag, beyond a few that are necessary to ensure that a tag can be parsed unambiguously. A tag name cannot contain any of the following characters:

- Colon (ASCII 58, “:”)
- Carriage return (ASCII 13, “\r”)

- Newline (ASCII 10, “\n”)

You can use the “hg tags” command to display the tags present in your repository. In the output, each tagged revision is identified first by its name, then by revision number, and finally by the unique hash of the revision.

```
1 $ hg tags
2 tip                1:476755c10e3c
3 v1.0              0:b6b1b0bfe2e9
```

Notice that `tip` is listed in the output of “hg tags”. The `tip` tag is a special “floating” tag, which always identifies the newest revision in the repository.

In the output of the “hg tags” command, tags are listed in reverse order, by revision number. This usually means that recent tags are listed before older tags. It also means that `tip` is always going to be the first tag listed in the output of “hg tags”.

When you run “hg log”, if it displays a revision that has tags associated with it, it will print those tags.

```
1 $ hg log
2 changeset: 1:476755c10e3c
3 tag:       tip
4 user:      Bryan O'Sullivan <bos@serpentine.com>
5 date:      Mon Dec 10 19:54:39 2007 +0000
6 summary:   Added tag v1.0 for changeset b6b1b0bfe2e9
7
8 changeset: 0:b6b1b0bfe2e9
9 tag:       v1.0
10 user:     Bryan O'Sullivan <bos@serpentine.com>
11 date:     Mon Dec 10 19:54:39 2007 +0000
12 summary:  Initial commit
13
```

Any time you need to provide a revision ID to a Mercurial command, the command will accept a tag name in its place. Internally, Mercurial will translate your tag name into the corresponding revision ID, then use that.

```
1 $ echo goodbye > myfile2
2 $ hg commit -A -m 'Second commit'
3 adding myfile2
4 $ hg log -r v1.0
5 changeset: 0:b6b1b0bfe2e9
6 tag:       v1.0
7 user:      Bryan O'Sullivan <bos@serpentine.com>
8 date:      Mon Dec 10 19:54:39 2007 +0000
9 summary:   Initial commit
10
```

There’s no limit on the number of tags you can have in a repository, or on the number of tags that a single revision can have. As a practical matter, it’s not a great idea to have “too many” (a number which will vary from project to project), simply because tags are supposed to help you to find revisions. If you have lots of tags, the ease of using them to identify revisions diminishes rapidly.

For example, if your project has milestones as frequent as every few days, it’s perfectly reasonable to tag each one of those. But if you have a continuous build system that makes sure every revision can be built cleanly, you’d be introducing a lot of noise if you were to tag every clean build. Instead, you could tag failed builds (on the assumption that they’re rare!), or simply not use tags to track buildability.

If you want to remove a tag that you no longer want, use “hg tag --remove”.

```

1 $ hg tag --remove v1.0
2 $ hg tags
3 tip                                3:aea218a7dc9f

```

You can also modify a tag at any time, so that it identifies a different revision, by simply issuing a new “hg tag” command. You’ll have to use the `-f` option to tell Mercurial that you *really* want to update the tag.

```

1 $ hg tag -r 1 v1.1
2 $ hg tags
3 tip                                4:21447ec9e7dd
4 v1.1                              1:476755c10e3c
5 $ hg tag -r 2 v1.1
6 abort: a tag named v1.1 already exists (use -f to force)
7 $ hg tag -f -r 2 v1.1
8 $ hg tags
9 tip                                5:043fd2a3319a
10 v1.1                             2:436d47f30256

```

There will still be a permanent record of the previous identity of the tag, but Mercurial will no longer use it. There’s thus no penalty to tagging the wrong revision; all you have to do is turn around and tag the correct revision once you discover your error.

Mercurial stores tags in a normal revision-controlled file in your repository. If you’ve created any tags, you’ll find them in a file named `.hgtags`. When you run the “hg tag” command, Mercurial modifies this file, then automatically commits the change to it. This means that every time you run “hg tag”, you’ll see a corresponding changeset in the output of “hg log”.

```

1 $ hg tip
2 changeset: 5:043fd2a3319a
3 tag:      tip
4 user:     Bryan O’Sullivan <bos@serpentine.com>
5 date:     Mon Dec 10 19:54:39 2007 +0000
6 summary:  Added tag v1.1 for changeset 436d47f30256
7

```

8.1.1 Handling tag conflicts during a merge

You won’t often need to care about the `.hgtags` file, but it sometimes makes its presence known during a merge. The format of the file is simple: it consists of a series of lines. Each line starts with a changeset hash, followed by a space, followed by the name of a tag.

If you’re resolving a conflict in the `.hgtags` file during a merge, there’s one twist to modifying the `.hgtags` file: when Mercurial is parsing the tags in a repository, it *never* reads the working copy of the `.hgtags` file. Instead, it reads the *most recently committed* revision of the file.

An unfortunate consequence of this design is that you can’t actually verify that your merged `.hgtags` file is correct until *after* you’ve committed a change. So if you find yourself resolving a conflict on `.hgtags` during a merge, be sure to run “hg tags” after you commit. If it finds an error in the `.hgtags` file, it will report the location of the error, which you can then fix and commit. You should then run “hg tags” again, just to be sure that your fix is correct.

8.1.2 Tags and cloning

You may have noticed that the “hg clone” command has a `-r` option that lets you clone an exact copy of the repository as of a particular changeset. The new clone will not contain any project history that comes after the revision you specified. This has an interaction with tags that can surprise the unwary.

Recall that a tag is stored as a revision to the `.hgtags` file, so that when you create a tag, the changeset in which it's recorded necessarily refers to an older changeset. When you run `"hg clone -r foo"` to clone a repository as of tag `foo`, the new clone *will not contain the history that created the tag* that you used to clone the repository. The result is that you'll get exactly the right subset of the project's history in the new repository, but *not* the tag you might have expected.

8.1.3 When permanent tags are too much

Since Mercurial's tags are revision controlled and carried around with a project's history, everyone you work with will see the tags you create. But giving names to revisions has uses beyond simply noting that revision `4237e45506ee` is really `v2.0.2`. If you're trying to track down a subtle bug, you might want a tag to remind you of something like "Anne saw the symptoms with this revision".

For cases like this, what you might want to use are *local* tags. You can create a local tag with the `-l` option to the `"hg tag"` command. This will store the tag in a file called `.hg/localtags`. Unlike `.hgtags`, `.hg/localtags` is not revision controlled. Any tags you create using `-l` remain strictly local to the repository you're currently working in.

8.2 The flow of changes—big picture vs. little

To return to the outline I sketched at the beginning of a chapter, let's think about a project that has multiple concurrent pieces of work under development at once.

There might be a push for a new "main" release; a new minor bugfix release to the last main release; and an unexpected "hot fix" to an old release that is now in maintenance mode.

The usual way people refer to these different concurrent directions of development is as "branches". However, we've already seen numerous times that Mercurial treats *all of history* as a series of branches and merges. Really, what we have here is two ideas that are peripherally related, but which happen to share a name.

- "Big picture" branches represent the sweep of a project's evolution; people give them names, and talk about them in conversation.
- "Little picture" branches are artefacts of the day-to-day activity of developing and merging changes. They expose the narrative of how the code was developed.

8.3 Managing big-picture branches in repositories

The easiest way to isolate a "big picture" branch in Mercurial is in a dedicated repository. If you have an existing shared repository—let's call it `myproject`—that reaches a "1.0" milestone, you can start to prepare for future maintenance releases on top of version 1.0 by tagging the revision from which you prepared the 1.0 release.

```
1 $ cd myproject
2 $ hg tag v1.0
```

You can then clone a new shared `myproject-1.0.1` repository as of that tag.

```
1 $ cd ..
2 $ hg clone myproject myproject-1.0.1
3 2 files updated, 0 files merged, 0 files removed, 0 files unresolved
```

Afterwards, if someone needs to work on a bug fix that ought to go into an upcoming 1.0.1 minor release, they clone the `myproject-1.0.1` repository, make their changes, and push them back.

```

1 $ hg clone myproject-1.0.1 my-1.0.1-bugfix
2 2 files updated, 0 files merged, 0 files removed, 0 files unresolved
3 $ cd my-1.0.1-bugfix
4 $ echo 'I fixed a bug using only echo!' >> myfile
5 $ hg commit -m 'Important fix for 1.0.1'
6 $ hg push
7 pushing to /tmp/branch-repo9eWu3d/myproject-1.0.1
8 searching for changes
9 adding changesets
10 adding manifests
11 adding file changes
12 added 1 changesets with 1 changes to 1 files

```

Meanwhile, development for the next major release can continue, isolated and unabated, in the `myproject` repository.

```

1 $ cd ..
2 $ hg clone myproject my-feature
3 2 files updated, 0 files merged, 0 files removed, 0 files unresolved
4 $ cd my-feature
5 $ echo 'This sure is an exciting new feature!' > mynewfile
6 $ hg commit -A -m 'New feature'
7 adding mynewfile
8 $ hg push
9 pushing to /tmp/branch-repo9eWu3d/myproject
10 searching for changes
11 adding changesets
12 adding manifests
13 adding file changes
14 added 1 changesets with 1 changes to 1 files

```

8.4 Don't repeat yourself: merging across branches

In many cases, if you have a bug to fix on a maintenance branch, the chances are good that the bug exists on your project's main branch (and possibly other maintenance branches, too). It's a rare developer who wants to fix the same bug multiple times, so let's look at a few ways that Mercurial can help you to manage these bugfixes without duplicating your work.

In the simplest instance, all you need to do is pull changes from your maintenance branch into your local clone of the target branch.

```

1 $ cd ..
2 $ hg clone myproject myproject-merge
3 3 files updated, 0 files merged, 0 files removed, 0 files unresolved
4 $ cd myproject-merge
5 $ hg pull ../myproject-1.0.1
6 pulling from ../myproject-1.0.1
7 searching for changes
8 adding changesets
9 adding manifests
10 adding file changes

```

```
11 added 1 changesets with 1 changes to 1 files (+1 heads)
12 (run 'hg heads' to see heads, 'hg merge' to merge)
```

You'll then need to merge the heads of the two branches, and push back to the main branch.

```
1 $ hg merge
2 1 files updated, 0 files merged, 0 files removed, 0 files unresolved
3 (branch merge, don't forget to commit)
4 $ hg commit -m 'Merge bugfix from 1.0.1 branch'
5 $ hg push
6 pushing to /tmp/branch-repo9eWu3d/myproject
7 searching for changes
8 adding changesets
9 adding manifests
10 adding file changes
11 added 2 changesets with 1 changes to 1 files
```

8.5 Naming branches within one repository

In most instances, isolating branches in repositories is the right approach. Its simplicity makes it easy to understand; and so it's hard to make mistakes. There's a one-to-one relationship between branches you're working in and directories on your system. This lets you use normal (non-Mercurial-aware) tools to work on files within a branch/repository.

If you're more in the "power user" category (*and* your collaborators are too), there is an alternative way of handling branches that you can consider. I've already mentioned the human-level distinction between "small picture" and "big picture" branches. While Mercurial works with multiple "small picture" branches in a repository all the time (for example after you pull changes in, but before you merge them), it can *also* work with multiple "big picture" branches.

The key to working this way is that Mercurial lets you assign a persistent *name* to a branch. There always exists a branch named `default`. Even before you start naming branches yourself, you can find traces of the `default` branch if you look for them.

As an example, when you run the "hg commit" command, and it pops up your editor so that you can enter a commit message, look for a line that contains the text "HG: branch default" at the bottom. This is telling you that your commit will occur on the branch named `default`.

To start working with named branches, use the "hg branches" command. This command lists the named branches already present in your repository, telling you which changeset is the tip of each.

```
1 $ hg tip
2 changeset: 0:746e67aa5ce7
3 tag:       tip
4 user:      Bryan O'Sullivan <bos@serpentine.com>
5 date:      Mon Dec 10 19:54:20 2007 +0000
6 summary:   Initial commit
7
8 $ hg branches
9 default                                0:746e67aa5ce7
```

Since you haven't created any named branches yet, the only one that exists is `default`.

To find out what the "current" branch is, run the "hg branch" command, giving it no arguments. This tells you what branch the parent of the current changeset is on.

```
1 $ hg branch
2 default
```

To create a new branch, run the “hg branch” command again. This time, give it one argument: the name of the branch you want to create.

```
1 $ hg branch foo
2 marked working directory as branch foo
3 $ hg branch
4 foo
```

After you’ve created a branch, you might wonder what effect the “hg branch” command has had. What do the “hg status” and “hg tip” commands report?

```
1 $ hg status
2 $ hg tip
3 changeset: 0:746e67aa5ce7
4 tag:      tip
5 user:     Bryan O’Sullivan <bos@serpentine.com>
6 date:     Mon Dec 10 19:54:20 2007 +0000
7 summary:  Initial commit
8
```

Nothing has changed in the working directory, and there’s been no new history created. As this suggests, running the “hg branch” command has no permanent effect; it only tells Mercurial what branch name to use the *next* time you commit a changeset.

When you commit a change, Mercurial records the name of the branch on which you committed. Once you’ve switched from the default branch to another and committed, you’ll see the name of the new branch show up in the output of “hg log”, “hg tip”, and other commands that display the same kind of output.

```
1 $ echo 'hello again' >> myfile
2 $ hg commit -m 'Second commit'
3 $ hg tip
4 changeset: 1:01a75295776d
5 branch:   foo
6 tag:      tip
7 user:     Bryan O’Sullivan <bos@serpentine.com>
8 date:     Mon Dec 10 19:54:20 2007 +0000
9 summary:  Second commit
10
```

The “hg log”-like commands will print the branch name of every changeset that’s not on the default branch. As a result, if you never use named branches, you’ll never see this information.

Once you’ve named a branch and committed a change with that name, every subsequent commit that descends from that change will inherit the same branch name. You can change the name of a branch at any time, using the “hg branch” command.

```
1 $ hg branch
2 foo
3 $ hg branch bar
4 marked working directory as branch bar
5 $ echo new file > newfile
6 $ hg commit -A -m 'Third commit'
7 adding newfile
8 $ hg tip
```

```

9 changeset: 2:ace59737a6c7
10 branch:   bar
11 tag:      tip
12 user:     Bryan O'Sullivan <bos@serpentine.com>
13 date:     Mon Dec 10 19:54:21 2007 +0000
14 summary:  Third commit
15

```

In practice, this is something you won't do very often, as branch names tend to have fairly long lifetimes. (This isn't a rule, just an observation.)

8.6 Dealing with multiple named branches in a repository

If you have more than one named branch in a repository, Mercurial will remember the branch that your working directory is on when you start a command like “hg update” or “hg pull -u”. It will update the working directory to the tip of this branch, no matter what the “repo-wide” tip is. To update to a revision that's on a different named branch, you may need to use the `-C` option to “hg update”.

This behaviour is a little subtle, so let's see it in action. First, let's remind ourselves what branch we're currently on, and what branches are in our repository.

```

1 $ hg parents
2 changeset: 2:ace59737a6c7
3 branch:   bar
4 tag:      tip
5 user:     Bryan O'Sullivan <bos@serpentine.com>
6 date:     Mon Dec 10 19:54:21 2007 +0000
7 summary:  Third commit
8
9 $ hg branches
10 bar                2:ace59737a6c7
11 foo                1:01a75295776d (inactive)
12 default            0:746e67aa5ce7 (inactive)

```

We're on the `bar` branch, but there also exists an older “hg foo” branch.

We can “hg update” back and forth between the tips of the `foo` and `bar` branches without needing to use the `-C` option, because this only involves going backwards and forwards linearly through our change history.

```

1 $ hg update foo
2 0 files updated, 0 files merged, 1 files removed, 0 files unresolved
3 $ hg parents
4 changeset: 1:01a75295776d
5 branch:   foo
6 user:     Bryan O'Sullivan <bos@serpentine.com>
7 date:     Mon Dec 10 19:54:20 2007 +0000
8 summary:  Second commit
9
10 $ hg update bar
11 1 files updated, 0 files merged, 0 files removed, 0 files unresolved
12 $ hg parents
13 changeset: 2:ace59737a6c7
14 branch:   bar

```

```
15 tag:          tip
16 user:         Bryan O'Sullivan <bos@serpentine.com>
17 date:         Mon Dec 10 19:54:21 2007 +0000
18 summary:      Third commit
19
```

If we go back to the `foo` branch and then run “`hg update`”, it will keep us on `foo`, not move us to the tip of `bar`.

```
1 $ hg update foo
2 0 files updated, 0 files merged, 1 files removed, 0 files unresolved
3 $ hg update
4 0 files updated, 0 files merged, 0 files removed, 0 files unresolved
```

Committing a new change on the `foo` branch introduces a new head.

```
1 $ echo something > somefile
2 $ hg commit -A -m 'New file'
3 adding somefile
4 $ hg heads
5 changeset:   3:56e2f990c337
6 branch:      foo
7 tag:         tip
8 parent:      1:01a75295776d
9 user:        Bryan O'Sullivan <bos@serpentine.com>
10 date:        Mon Dec 10 19:54:21 2007 +0000
11 summary:     New file
12
13 changeset:   2:ace59737a6c7
14 branch:      bar
15 user:        Bryan O'Sullivan <bos@serpentine.com>
16 date:        Mon Dec 10 19:54:21 2007 +0000
17 summary:     Third commit
18
```

We can no longer update from `foo` to `bar` without going “sideways” in history, so Mercurial forces us to provide the `-C` option to “`hg update`”.

```
1 $ hg update bar
2 abort: update spans branches, use 'hg merge' or 'hg update -C' to lose changes
3 $ hg update -C bar
4 1 files updated, 0 files merged, 1 files removed, 0 files unresolved
```

8.7 Branch names and merging

As you’ve probably noticed, merges in Mercurial are not symmetrical. Let’s say our repository has two heads, 17 and 23. If I “`hg update`” to 17 and then “`hg merge`” with 23, Mercurial records 17 as the first parent of the merge, and 23 as the second. Whereas if I “`hg update`” to 23 and then “`hg merge`” with 17, it records 23 as the first parent, and 17 as the second.

This affects Mercurial’s choice of branch name when you merge. After a merge, Mercurial will retain the branch name of the first parent when you commit the result of the merge. If your first parent’s branch name is `foo`, and you merge with `bar`, the branch name will still be `foo` after you merge.

It's not unusual for a repository to contain multiple heads, each with the same branch name. Let's say I'm working on the `foo` branch, and so are you. We commit different changes; I pull your changes; I now have two heads, each claiming to be on the `foo` branch. The result of a merge will be a single head on the `foo` branch, as you might hope.

But if I'm working on the `bar` branch, and I merge work from the `foo` branch, the result will remain on the `bar` branch.

```
1 $ hg branch
2 bar
3 $ hg merge
4 1 files updated, 0 files merged, 0 files removed, 0 files unresolved
5 (branch merge, don't forget to commit)
6 $ hg commit -m 'Merge'
7 $ hg tip
8 changeset: 4:f170fbc1577e
9 branch:    bar
10 tag:       tip
11 parent:    2:ace59737a6c7
12 parent:    3:56e2f990c337
13 user:      Bryan O'Sullivan <bos@serpentine.com>
14 date:      Mon Dec 10 19:54:22 2007 +0000
15 summary:   Merge
16
```

To give a more concrete example, if I'm working on the `bleeding-edge` branch, and I want to bring in the latest fixes from the `stable` branch, Mercurial will choose the “right” (bleeding-edge) branch name when I pull and merge from `stable`.

8.8 Branch naming is generally useful

You shouldn't think of named branches as applicable only to situations where you have multiple long-lived branches cohabiting in a single repository. They're very useful even in the one-branch-per-repository case.

In the simplest case, giving a name to each branch gives you a permanent record of which branch a changeset originated on. This gives you more context when you're trying to follow the history of a long-lived branchy project.

If you're working with shared repositories, you can set up a `pretxnchangegroup` hook on each that will block incoming changes that have the “wrong” branch name. This provides a simple, but effective, defence against people accidentally pushing changes from a “bleeding edge” branch to a “stable” branch. Such a hook might look like this inside the shared repo's `hgrc`.

```
1 [hooks]
2 pretxnchangegroup.branch = hg heads --template 'branches ' | grep mybranch
```

Chapter 9

Finding and fixing your mistakes

To err might be human, but to really handle the consequences well takes a top-notch revision control system. In this chapter, we'll discuss some of the techniques you can use when you find that a problem has crept into your project. Mercurial has some highly capable features that will help you to isolate the sources of problems, and to handle them appropriately.

9.1 Erasing local history

9.1.1 The accidental commit

I have the occasional but persistent problem of typing rather more quickly than I can think, which sometimes results in me committing a changeset that is either incomplete or plain wrong. In my case, the usual kind of incomplete changeset is one in which I've created a new source file, but forgotten to “hg add” it. A “plain wrong” changeset is not as common, but no less annoying.

9.1.2 Rolling back a transaction

In section 4.2.2, I mentioned that Mercurial treats each modification of a repository as a *transaction*. Every time you commit a changeset or pull changes from another repository, Mercurial remembers what you did. You can undo, or *roll back*, exactly one of these actions using the “hg rollback” command. (See section 9.1.4 for an important caveat about the use of this command.)

Here's a mistake that I often find myself making: committing a change in which I've created a new file, but forgotten to “hg add” it.

```
1 $ hg status
2 M a
3 $ echo b > b
4 $ hg commit -m 'Add file b'
```

Looking at the output of “hg status” after the commit immediately confirms the error.

```
1 $ hg status
2 ? b
3 $ hg tip
4 changeset: 1:8a557de9a817
5 tag:      tip
6 user:     Bryan O'Sullivan <bos@serpentine.com>
7 date:     Mon Dec 10 19:54:38 2007 +0000
```

```
8 summary:      Add file b
9
```

The commit captured the changes to the file a, but not the new file b. If I were to push this changeset to a repository that I shared with a colleague, the chances are high that something in a would refer to b, which would not be present in their repository when they pulled my changes. I would thus become the object of some indignation.

However, luck is with me—I’ve caught my error before I pushed the changeset. I use the “hg rollback” command, and Mercurial makes that last changeset vanish.

```
1 $ hg rollback
2 rolling back last transaction
3 $ hg tip
4 changeset:  0:15f839321f52
5 tag:        tip
6 user:       Bryan O’Sullivan <bos@serpentine.com>
7 date:       Mon Dec 10 19:54:38 2007 +0000
8 summary:    First commit
9
10 $ hg status
11 M a
12 ? b
```

Notice that the changeset is no longer present in the repository’s history, and the working directory once again thinks that the file a is modified. The commit and rollback have left the working directory exactly as it was prior to the commit; the changeset has been completely erased. I can now safely “hg add” the file b, and rerun my commit.

```
1 $ hg add b
2 $ hg commit -m 'Add file b, this time for real'
```

9.1.3 The erroneous pull

It’s common practice with Mercurial to maintain separate development branches of a project in different repositories. Your development team might have one shared repository for your project’s “0.9” release, and another, containing different changes, for the “1.0” release.

Given this, you can imagine that the consequences could be messy if you had a local “0.9” repository, and accidentally pulled changes from the shared “1.0” repository into it. At worst, you could be paying insufficient attention, and push those changes into the shared “0.9” tree, confusing your entire team (but don’t worry, we’ll return to this horror scenario later). However, it’s more likely that you’ll notice immediately, because Mercurial will display the URL it’s pulling from, or you will see it pull a suspiciously large number of changes into the repository.

The “hg rollback” command will work nicely to expunge all of the changesets that you just pulled. Mercurial groups all changes from one “hg pull” into a single transaction, so one “hg rollback” is all you need to undo this mistake.

9.1.4 Rolling back is useless once you’ve pushed

The value of the “hg rollback” command drops to zero once you’ve pushed your changes to another repository. Rolling back a change makes it disappear entirely, but *only* in the repository in which you perform the “hg rollback”. Because a rollback eliminates history, there’s no way for the disappearance of a change to propagate between repositories.

If you’ve pushed a change to another repository—particularly if it’s a shared repository—it has essentially “escaped into the wild,” and you’ll have to recover from your mistake in a different way. What will happen if you push a

changeset somewhere, then roll it back, then pull from the repository you pushed to, is that the changeset will reappear in your repository.

(If you absolutely know for sure that the change you want to roll back is the most recent change in the repository that you pushed to, *and* you know that nobody else could have pulled it from that repository, you can roll back the changeset there, too, but you really should really not rely on this working reliably. If you do this, sooner or later a change really will make it into a repository that you don't directly control (or have forgotten about), and come back to bite you.)

9.1.5 You can only roll back once

Mercurial stores exactly one transaction in its transaction log; that transaction is the most recent one that occurred in the repository. This means that you can only roll back one transaction. If you expect to be able to roll back one transaction, then its predecessor, this is not the behaviour you will get.

```
1 $ hg rollback
2 rolling back last transaction
3 $ hg rollback
4 no rollback information available
```

Once you've rolled back one transaction in a repository, you can't roll back again in that repository until you perform another commit or pull.

9.2 Reverting the mistaken change

If you make a modification to a file, and decide that you really didn't want to change the file at all, and you haven't yet committed your changes, the "hg revert" command is the one you'll need. It looks at the changeset that's the parent of the working directory, and restores the contents of the file to their state as of that changeset. (That's a long-winded way of saying that, in the normal case, it undoes your modifications.)

Let's illustrate how the "hg revert" command works with yet another small example. We'll begin by modifying a file that Mercurial is already tracking.

```
1 $ cat file
2 original content
3 $ echo unwanted change >> file
4 $ hg diff file
5 diff -r 7365fa775ae5 file
6 --- a/file      Mon Dec 10 19:54:27 2007 +0000
7 +++ b/file      Mon Dec 10 19:54:27 2007 +0000
8 @@ -1,1 +1,2 @@ original content
9  original content
10 +unwanted change
```

If we don't want that change, we can simply "hg revert" the file.

```
1 $ hg status
2 M file
3 $ hg revert file
4 $ cat file
5 original content
```

The "hg revert" command provides us with an extra degree of safety by saving our modified file with a .orig extension.

```
1 $ hg status
2 ? file.orig
3 $ cat file.orig
4 original content
5 unwanted change
```

Here is a summary of the cases that the “hg revert” command can deal with. We will describe each of these in more detail in the section that follows.

- If you modify a file, it will restore the file to its unmodified state.
- If you “hg add” a file, it will undo the “added” state of the file, but leave the file itself untouched.
- If you delete a file without telling Mercurial, it will restore the file to its unmodified contents.
- If you use the “hg remove” command to remove a file, it will undo the “removed” state of the file, and restore the file to its unmodified contents.

9.2.1 File management errors

The “hg revert” command is useful for more than just modified files. It lets you reverse the results of all of Mercurial’s file management commands—“hg add”, “hg remove”, and so on.

If you “hg add” a file, then decide that in fact you don’t want Mercurial to track it, use “hg revert” to undo the add. Don’t worry; Mercurial will not modify the file in any way. It will just “unmark” the file.

```
1 $ echo oops > oops
2 $ hg add oops
3 $ hg status oops
4 A oops
5 $ hg revert oops
6 $ hg status
7 ? oops
```

Similarly, if you ask Mercurial to “hg remove” a file, you can use “hg revert” to restore it to the contents it had as of the parent of the working directory.

```
1 $ hg remove file
2 $ hg status
3 R file
4 $ hg revert file
5 $ hg status
6 $ ls file
7 file
```

This works just as well for a file that you deleted by hand, without telling Mercurial (recall that in Mercurial terminology, this kind of file is called “missing”).

```
1 $ rm file
2 $ hg status
3 ! file
4 $ hg revert file
5 $ ls file
6 file
```

If you revert a “hg copy”, the copied-to file remains in your working directory afterwards, untracked. Since a copy doesn’t affect the copied-from file in any way, Mercurial doesn’t do anything with the copied-from file.

```
1 $ hg copy file new-file
2 $ hg revert new-file
3 $ hg status
4 ? new-file
```

A slightly special case: reverting a rename

If you “hg rename” a file, there is one small detail that you should remember. When you “hg revert” a rename, it’s not enough to provide the name of the renamed-to file, as you can see here.

```
1 $ hg rename file new-file
2 $ hg revert new-file
3 $ hg status
4 ? new-file
```

As you can see from the output of “hg status”, the renamed-to file is no longer identified as added, but the renamed-from file is still removed! This is counter-intuitive (at least to me), but at least it’s easy to deal with.

```
1 $ hg revert file
2 no changes needed to file
3 $ hg status
4 ? new-file
```

So remember, to revert a “hg rename”, you must provide *both* the source and destination names.

(By the way, if you rename a file, then modify the renamed-to file, then revert both components of the rename, when Mercurial restores the file that was removed as part of the rename, it will be unmodified. If you need the modifications in the renamed-to file to show up in the renamed-from file, don’t forget to copy them over.)

These fiddly aspects of reverting a rename arguably constitute a small bug in Mercurial.

9.3 Dealing with committed changes

Consider a case where you have committed a change *a*, and another change *b* on top of it; you then realise that change *a* was incorrect. Mercurial lets you “back out” an entire changeset automatically, and building blocks that let you reverse part of a changeset by hand.

Before you read this section, here’s something to keep in mind: the “hg backout” command undoes changes by *adding* history, not by modifying or erasing it. It’s the right tool to use if you’re fixing bugs, but not if you’re trying to undo some change that has catastrophic consequences. To deal with those, see section [9.4](#).

9.3.1 Backing out a changeset

The “hg backout” command lets you “undo” the effects of an entire changeset in an automated fashion. Because Mercurial’s history is immutable, this command *does not* get rid of the changeset you want to undo. Instead, it creates a new changeset that *reverses* the effect of the to-be-undone changeset.

The operation of the “hg backout” command is a little intricate, so let’s illustrate it with some examples. First, we’ll create a repository with some simple changes.

```
1 $ hg init myrepo
2 $ cd myrepo
```

```

3 $ echo first change >> myfile
4 $ hg add myfile
5 $ hg commit -m 'first change'
6 $ echo second change >> myfile
7 $ hg commit -m 'second change'

```

The “hg backout” command takes a single changeset ID as its argument; this is the changeset to back out. Normally, “hg backout” will drop you into a text editor to write a commit message, so you can record why you’re backing the change out. In this example, we provide a commit message on the command line using the `-m` option.

9.3.2 Backing out the tip changeset

We’re going to start by backing out the last changeset we committed.

```

1 $ hg backout -m 'back out second change' tip
2 reverting myfile
3 changeset 2:e02eba531f95 backs out changeset 1:c3e45317eb42
4 $ cat myfile
5 first change

```

You can see that the second line from `myfile` is no longer present. Taking a look at the output of “hg log” gives us an idea of what the “hg backout” command has done.

```

1 $ hg log --style compact
2 2[tip] e02eba531f95 2007-12-10 19:53 +0000 bos
3   back out second change
4
5 1 c3e45317eb42 2007-12-10 19:53 +0000 bos
6   second change
7
8 0 f3db226c9812 2007-12-10 19:53 +0000 bos
9   first change
10

```

Notice that the new changeset that “hg backout” has created is a child of the changeset we backed out. It’s easier to see this in figure 9.1, which presents a graphical view of the change history. As you can see, the history is nice and linear.

9.3.3 Backing out a non-tip change

If you want to back out a change other than the last one you committed, pass the `--merge` option to the “hg backout” command.

```

1 $ cd ..
2 $ hg clone -r1 myrepo non-tip-repo
3 requesting all changes
4 adding changesets
5 adding manifests
6 adding file changes
7 added 2 changesets with 2 changes to 1 files
8 1 files updated, 0 files merged, 0 files removed, 0 files unresolved
9 $ cd non-tip-repo

```




Figure 9.1: Backing out a change using the “hg backout” command

This makes backing out any changeset a “one-shot” operation that’s usually simple and fast.

```
1 $ echo third change >> myfile
2 $ hg commit -m 'third change'
3 $ hg backout --merge -m 'back out second change' 1
4 reverting myfile
5 changeset 3:dc2f0a481aab backs out changeset 1:c3e45317eb42
6 merging with changeset 2:6beeebc9c7f6
7 merging myfile
8 0 files updated, 1 files merged, 0 files removed, 0 files unresolved
9 (branch merge, don't forget to commit)
```

If you take a look at the contents of `myfile` after the backout finishes, you’ll see that the first and third changes are present, but not the second.

```
1 $ cat myfile
2 first change
3 third change
```

As the graphical history in figure 9.2 illustrates, Mercurial actually commits *two* changes in this kind of situation (the box-shaped nodes are the ones that Mercurial commits automatically). Before Mercurial begins the backout process, it first remembers what the current parent of the working directory is. It then backs out the target changeset, and commits that as a changeset. Finally, it merges back to the previous parent of the working directory, and commits the result of the merge.

The result is that you end up “back where you were”, only with some extra history that undoes the effect of the changeset you wanted to back out.

Always use the `--merge` option

In fact, since the `--merge` option will do the “right thing” whether or not the changeset you’re backing out is the tip (i.e. it won’t try to merge if it’s backing out the tip, since there’s no need), you should *always* use this option when you run the “hg backout” command.




Figure 9.2: Automated backout of a non-tip change using the “hg backout” command

9.3.4 Gaining more control of the backout process

While I’ve recommended that you always use the `--merge` option when backing out a change, the “hg backout” command lets you decide how to merge a backout changeset. Taking control of the backout process by hand is something you will rarely need to do, but it can be useful to understand what the “hg backout” command is doing for you automatically. To illustrate this, let’s clone our first repository, but omit the backout change that it contains.

```

1 $ cd ..
2 $ hg clone -r1 myrepo newrepo
3 requesting all changes
4 adding changesets
5 adding manifests
6 adding file changes
7 added 2 changesets with 2 changes to 1 files
8 1 files updated, 0 files merged, 0 files removed, 0 files unresolved
9 $ cd newrepo

```

As with our earlier example, We’ll commit a third changeset, then back out its parent, and see what happens.

```

1 $ echo third change >> myfile
2 $ hg commit -m 'third change'
3 $ hg backout -m 'back out second change' 1
4 reverting myfile
5 changeset 3:dc2f0a481aab backs out changeset 1:c3e45317eb42
6 the backout changeset is a new head - do not forget to merge
7 (use "backout --merge" if you want to auto-merge)

```

Our new changeset is again a descendant of the changeset we backout out; it's thus a new head, *not* a descendant of the changeset that was the tip. The “hg backout” command was quite explicit in telling us this.

```
1 $ hg log --style compact
2 3[tip]:1 dc2f0a481aab 2007-12-10 19:53 +0000 bos
3   back out second change
4
5 2 6beeebc9c7f6 2007-12-10 19:53 +0000 bos
6   third change
7
8 1 c3e45317eb42 2007-12-10 19:53 +0000 bos
9   second change
10
11 0 f3db226c9812 2007-12-10 19:53 +0000 bos
12   first change
13
```

Again, it's easier to see what has happened by looking at a graph of the revision history, in figure 9.3. This makes it clear that when we use “hg backout” to back out a change other than the tip, Mercurial adds a new head to the repository (the change it committed is box-shaped).




Figure 9.3: Backing out a change using the “hg backout” command

After the “hg backout” command has completed, it leaves the new “backout” changeset as the parent of the working directory.

```
1 $ hg parents
2 changeset: 3:dc2f0a481aab
3 tag: tip
4 parent: 1:c3e45317eb42
5 user: Bryan O'Sullivan <bos@serpentine.com>
6 date: Mon Dec 10 19:53:16 2007 +0000
7 summary: back out second change
8
```

Now we have two isolated sets of changes.

```
1 $ hg heads
2 changeset: 3:dc2f0a481aab
3 tag: tip
4 parent: 1:c3e45317eb42
5 user: Bryan O'Sullivan <bos@serpentine.com>
6 date: Mon Dec 10 19:53:16 2007 +0000
7 summary: back out second change
8
9 changeset: 2:6beeebc9c7f6
10 user: Bryan O'Sullivan <bos@serpentine.com>
11 date: Mon Dec 10 19:53:16 2007 +0000
12 summary: third change
13
```

Let's think about what we expect to see as the contents of `myfile` now. The first change should be present, because we've never backed it out. The second change should be missing, as that's the change we backed out. Since the history graph shows the third change as a separate head, we *don't* expect to see the third change present in `myfile`.

```
1 $ cat myfile
2 first change
```

To get the third change back into the file, we just do a normal merge of our two heads.

```
1 $ hg merge
2 merging myfile
3 0 files updated, 1 files merged, 0 files removed, 0 files unresolved
4 (branch merge, don't forget to commit)
5 $ hg commit -m 'merged backout with previous tip'
6 $ cat myfile
7 first change
8 third change
```

Afterwards, the graphical history of our repository looks like figure 9.4.

9.3.5 Why “hg backout” works as it does

Here's a brief description of how the “hg backout” command works.

1. It ensures that the working directory is “clean”, i.e. that the output of “hg status” would be empty.
2. It remembers the current parent of the working directory. Let's call this changeset `orig`
3. It does the equivalent of a “hg update” to sync the working directory to the changeset you want to back out. Let's call this changeset `backout`
4. It finds the parent of that changeset. Let's call that changeset `parent`.
5. For each file that the `backout` changeset affected, it does the equivalent of a “hg revert -r parent” on that file, to restore it to the contents it had before that changeset was committed.
6. It commits the result as a new changeset. This changeset has `backout` as its parent.
7. If you specify `--merge` on the command line, it merges with `orig`, and commits the result of the merge.




Figure 9.4: Manually merging a backout change

An alternative way to implement the “hg backout” command would be to “hg export” the to-be-backed-out changeset as a diff, then use the `--reverse` option to the `patch` command to reverse the effect of the change without fiddling with the working directory. This sounds much simpler, but it would not work nearly as well.

The reason that “hg backout” does an update, a commit, a merge, and another commit is to give the merge machinery the best chance to do a good job when dealing with all the changes *between* the change you’re backing out and the current tip.

If you’re backing out a changeset that’s 100 revisions back in your project’s history, the chances that the `patch` command will be able to apply a reverse diff cleanly are not good, because intervening changes are likely to have “broken the context” that `patch` uses to determine whether it can apply a patch (if this sounds like gibberish, see [12.4](#) for a discussion of the `patch` command). Also, Mercurial’s merge machinery will handle files and directories being renamed, permission changes, and modifications to binary files, none of which `patch` can deal with.

9.4 Changes that should never have been

Most of the time, the “hg backout” command is exactly what you need if you want to undo the effects of a change. It leaves a permanent record of exactly what you did, both when committing the original changeset and when you cleaned up after it.

On rare occasions, though, you may find that you’ve committed a change that really should not be present in the repository at all. For example, it would be very unusual, and usually considered a mistake, to commit a software project’s object files as well as its source files. Object files have almost no intrinsic value, and they’re *big*, so they increase the size of the repository and the amount of time it takes to clone or pull changes.

Before I discuss the options that you have if you commit a “brown paper bag” change (the kind that’s so bad that you want to pull a brown paper bag over your head), let me first discuss some approaches that probably won’t work.

Since Mercurial treats history as accumulative—every change builds on top of all changes that preceded it—you generally can’t just make disastrous changes disappear. The one exception is when you’ve just committed a change, and it hasn’t been pushed or pulled into another repository. That’s when you can safely use the “hg rollback” command, as I detailed in section 9.1.2.

After you’ve pushed a bad change to another repository, you *could* still use “hg rollback” to make your local copy of the change disappear, but it won’t have the consequences you want. The change will still be present in the remote repository, so it will reappear in your local repository the next time you pull.

If a situation like this arises, and you know which repositories your bad change has propagated into, you can *try* to get rid of the change from *every* one of those repositories. This is, of course, not a satisfactory solution: if you miss even a single repository while you’re expunging, the change is still “in the wild”, and could propagate further.

If you’ve committed one or more changes *after* the change that you’d like to see disappear, your options are further reduced. Mercurial doesn’t provide a way to “punch a hole” in history, leaving changesets intact.

XXX This needs filling out. The hg-replay script in the examples directory works, but doesn’t handle merge changesets. Kind of an important omission.

9.4.1 Protect yourself from “escaped” changes

If you’ve committed some changes to your local repository and they’ve been pushed or pulled somewhere else, this isn’t necessarily a disaster. You can protect yourself ahead of time against some classes of bad changeset. This is particularly easy if your team usually pulls changes from a central repository.

By configuring some hooks on that repository to validate incoming changesets (see chapter 10), you can automatically prevent some kinds of bad changeset from being pushed to the central repository at all. With such a configuration in place, some kinds of bad changeset will naturally tend to “die out” because they can’t propagate into the central repository. Better yet, this happens without any need for explicit intervention.

For instance, an incoming change hook that verifies that a changeset will actually compile can prevent people from inadvertently “breaking the build”.

9.5 Finding the source of a bug

While it’s all very well to be able to back out a changeset that introduced a bug, this requires that you know which changeset to back out. Mercurial provides an invaluable extension, called `bisect`, that helps you to automate this process and accomplish it very efficiently.

The idea behind the `bisect` extension is that a changeset has introduced some change of behaviour that you can identify with a simple binary test. You don’t know which piece of code introduced the change, but you know how to test for the presence of the bug. The `bisect` extension uses your test to direct its search for the changeset that introduced the code that caused the bug.

Here are a few scenarios to help you understand how you might apply this extension.

- The most recent version of your software has a bug that you remember wasn’t present a few weeks ago, but you don’t know when it was introduced. Here, your binary test checks for the presence of that bug.
- You fixed a bug in a rush, and now it’s time to close the entry in your team’s bug database. The bug database requires a changeset ID when you close an entry, but you don’t remember which changeset you fixed the bug in. Once again, your binary test checks for the presence of the bug.
- Your software works correctly, but runs 15% slower than the last time you measured it. You want to know which changeset introduced the performance regression. In this case, your binary test measures the performance of your software, to see whether it’s “fast” or “slow”.
- The sizes of the components of your project that you ship exploded recently, and you suspect that something changed in the way you build your project.

From these examples, it should be clear that the `bisect` extension is not useful only for finding the sources of bugs. You can use it to find any “emergent property” of a repository (anything that you can’t find from a simple text search of the files in the tree) for which you can write a binary test.

We’ll introduce a little bit of terminology here, just to make it clear which parts of the search process are your responsibility, and which are Mercurial’s. A *test* is something that *you* run when `bisect` chooses a changeset. A *probe* is what `bisect` runs to tell whether a revision is good. Finally, we’ll use the word “bisect”, as both a noun and a verb, to stand in for the phrase “search using the `bisect` extension”.

One simple way to automate the searching process would be simply to probe every changeset. However, this scales poorly. If it took ten minutes to test a single changeset, and you had 10,000 changesets in your repository, the exhaustive approach would take on average 35 *days* to find the changeset that introduced a bug. Even if you knew that the bug was introduced by one of the last 500 changesets, and limited your search to those, you’d still be looking at over 40 hours to find the changeset that introduced your bug.

What the *bisect* extension does is use its knowledge of the “shape” of your project’s revision history to perform a search in time proportional to the *logarithm* of the number of changesets to check (the kind of search it performs is called a dichotomic search). With this approach, searching through 10,000 changesets will take less than two hours, even at ten minutes per test. Limit your search to the last 500 changesets, and it will take less than an hour.

The `bisect` extension is aware of the “branchy” nature of a Mercurial project’s revision history, so it has no problems dealing with branches, merges, or multiple heads in a repository. It can prune entire branches of history with a single probe, which is how it operates so efficiently.

9.5.1 Using the `bisect` extension

Here’s an example of `bisect` in action. To keep the core of Mercurial simple, `bisect` is packaged as an extension; this means that it won’t be present unless you explicitly enable it. To do this, edit your `hgrc` and add the following section header (if it’s not already present):

```
1 [extensions]
```

Then add a line to this section to enable the extension:

```
1 hgbisect =
```

Note: That’s right, there’s a “h” at the front of the name of the `bisect` extension. The reason is that Mercurial is written in Python, and uses a standard Python package called `bisect`. If you omit the “h” from the name “`hgbisect`”, Mercurial will erroneously find the standard Python `bisect` package, and try to use it as a Mercurial extension. This won’t work, and Mercurial will crash repeatedly until you fix the spelling in your `hgrc`. Ugh.

Now let’s create a repository, so that we can try out the `bisect` extension in isolation.

```
1 $ hg init mybug
2 $ cd mybug
```

We’ll simulate a project that has a bug in it in a simple-minded way: create trivial changes in a loop, and nominate one specific change that will have the “bug”. This loop creates 35 changesets, each adding a single file to the repository. We’ll represent our “bug” with a file that contains the text “i have a gub”.

```
1 $ buggy_change=22
2 $ for (( i = 0; i < 35; i++ )); do
3 >   if [[ $i = $buggy_change ]]; then
4 >     echo 'i have a gub' > myfile$i
5 >     hg commit -q -m 'buggy changeset'
```

```

6 > else
7 >     echo 'nothing to see here, move along' > myfile$i
8 >     hg commit -q -A -m 'normal changeset'
9 > fi
10 > done

```

The next thing that we'd like to do is figure out how to use the `bisect` extension. We can use Mercurial's normal built-in help mechanism for this.

```

1 $ hg help bisect
2 hg bisect [help|init|reset|next|good|bad]
3
4 Dichotomic search in the DAG of changesets
5
6 This extension helps to find changesets which cause problems.
7 To use, mark the earliest changeset you know introduces the problem
8 as bad, then mark the latest changeset which is free from the problem
9 as good. Bisect will update your working directory to a revision for
10 testing. Once you have performed tests, mark the working directory
11 as bad or good and bisect will either update to another candidate
12 changeset or announce that it has found the bad revision.
13
14 Note: bisect expects bad revisions to be descendants of good revisions.
15 If you are looking for the point at which a problem was fixed, then make
16 the problem-free state "bad" and the problematic state "good."
17
18 For subcommands see "hg bisect help"
19
20 use "hg -v help bisect" to show global options
21 $ hg bisect help
22 list of subcommands for the bisect extension
23
24 bad      mark revision as bad and update to the next revision to test
25 good     mark revision as good and update to the next revision to test
26 help     show help for a given bisect subcommand or all subcommands
27 init     start a new bisection
28 next     find and update to the next revision to test
29 reset    finish a bisection

```

The `bisect` extension works in steps. Each step proceeds as follows.

1. You run your binary test.
 - If the test succeeded, you tell `bisect` by running the “`hg bisect good`” command.
 - If it failed, use the “`hg bisect bad`” command to let the `bisect` extension know.
2. The extension uses your information to decide which changeset to test next.
3. It updates the working directory to that changeset, and the process begins again.

The process ends when `bisect` identifies a unique changeset that marks the point where your test transitioned from “succeeding” to “failing”.

To start the search, we must run the “`hg bisect init`” command.

```

1 $ hg bisect init
2 abort: bisect directory already exists
3

```

In our case, the binary test we use is simple: we check to see if any file in the repository contains the string “i have a gub”. If it does, this changeset contains the change that “caused the bug”. By convention, a changeset that has the property we’re searching for is “bad”, while one that doesn’t is “good”.

Most of the time, the revision to which the working directory is synced (usually the tip) already exhibits the problem introduced by the buggy change, so we’ll mark it as “bad”.

```

1 $ hg bisect bad

```

Our next task is to nominate a changeset that we know *doesn’t* have the bug; the `bisect` extension will “bracket” its search between the first pair of good and bad changesets. In our case, we know that revision 10 didn’t have the bug. (I’ll have more words about choosing the first “good” changeset later.)

```

1 $ hg bisect good 10
2 Testing changeset 22:baeda37353e3 (24 changesets remaining, ~4 tests)
3 0 files updated, 0 files merged, 12 files removed, 0 files unresolved

```

Notice that this command printed some output.

- It told us how many changesets it must consider before it can identify the one that introduced the bug, and how many tests that will require.
- It updated the working directory to the next changeset to test, and told us which changeset it’s testing.

We now run our test in the working directory. We use the `grep` command to see if our “bad” file is present in the working directory. If it is, this revision is bad; if not, this revision is good.

```

1 $ if grep -q 'i have a gub' *
2 > then
3 >   result=bad
4 > else
5 >   result=good
6 > fi
7 $ echo this revision is $result
8 this revision is bad
9 $ hg bisect $result
10 Testing changeset 16:f410b84ecalc (12 changesets remaining, ~3 tests)
11 0 files updated, 0 files merged, 6 files removed, 0 files unresolved

```

This test looks like a perfect candidate for automation, so let’s turn it into a shell function.

```

1 $ mytest() {
2 >   if grep -q 'i have a gub' *
3 >   then
4 >     result=bad
5 >   else
6 >     result=good
7 >   fi
8 >   echo this revision is $result
9 >   hg bisect $result
10 > }

```

We can now run an entire test step with a single command, `mytest`.

```
1 $ mytest
2 this revision is good
3 Testing changeset 19:f4dcc543ad84 (6 changesets remaining, ~2 tests)
4 3 files updated, 0 files merged, 0 files removed, 0 files unresolved
```

A few more invocations of our canned test step command, and we're done.

```
1 $ mytest
2 this revision is good
3 Testing changeset 20:67d386ac9c6a (3 changesets remaining, ~1 tests)
4 1 files updated, 0 files merged, 0 files removed, 0 files unresolved
5 $ mytest
6 this revision is good
7 Testing changeset 21:18b8376e673d (2 changesets remaining, ~1 tests)
8 1 files updated, 0 files merged, 0 files removed, 0 files unresolved
9 $ mytest
10 this revision is good
11 The first bad revision is:
12 changeset: 22:baeda37353e3
13 user:      Bryan O'Sullivan <bos@serpentine.com>
14 date:      Mon Dec 10 19:54:17 2007 +0000
15 summary:   buggy changeset
16
```

Even though we had 40 changesets to search through, the `bisect` extension let us find the changeset that introduced our “bug” with only five tests. Because the number of tests that the `bisect` extension grows logarithmically with the number of changesets to search, the advantage that it has over the “brute force” search approach increases with every changeset you add.

9.5.2 Cleaning up after your search

When you're finished using the `bisect` extension in a repository, you can use the “`hg bisect reset`” command to drop the information it was using to drive your search. The extension doesn't use much space, so it doesn't matter if you forget to run this command. However, `bisect` won't let you start a new search in that repository until you do a “`hg bisect reset`”.

```
1 $ hg bisect reset
```

9.6 Tips for finding bugs effectively

9.6.1 Give consistent input

The `bisect` extension requires that you correctly report the result of every test you perform. If you tell it that a test failed when it really succeeded, it *might* be able to detect the inconsistency. If it can identify an inconsistency in your reports, it will tell you that a particular changeset is both good and bad. However, it can't do this perfectly; it's about as likely to report the wrong changeset as the source of the bug.

9.6.2 Automate as much as possible

When I started using the `bisect` extension, I tried a few times to run my tests by hand, on the command line. This is an approach that I, at least, am not suited to. After a few tries, I found that I was making enough mistakes that I was having to restart my searches several times before finally getting correct results.

My initial problems with driving the `bisect` extension by hand occurred even with simple searches on small repositories; if the problem you're looking for is more subtle, or the number of tests that `bisect` must perform increases, the likelihood of operator error ruining the search is much higher. Once I started automating my tests, I had much better results.

The key to automated testing is twofold:

- always test for the same symptom, and
- always feed consistent input to the “`hg bisect`” command.

In my tutorial example above, the `grep` command tests for the symptom, and the `if` statement takes the result of this check and ensures that we always feed the same input to the “`hg bisect`” command. The `mytest` function marries these together in a reproducible way, so that every test is uniform and consistent.

9.6.3 Check your results

Because the output of a `bisect` search is only as good as the input you give it, don't take the changeset it reports as the absolute truth. A simple way to cross-check its report is to manually run your test at each of the following changesets:

- The changeset that it reports as the first bad revision. Your test should still report this as bad.
- The parent of that changeset (either parent, if it's a merge). Your test should report this changeset as good.
- A child of that changeset. Your test should report this changeset as bad.

9.6.4 Beware interference between bugs

It's possible that your search for one bug could be disrupted by the presence of another. For example, let's say your software crashes at revision 100, and worked correctly at revision 50. Unknown to you, someone else introduced a different crashing bug at revision 60, and fixed it at revision 80. This could distort your results in one of several ways.

It is possible that this other bug completely “masks” yours, which is to say that it occurs before your bug has a chance to manifest itself. If you can't avoid that other bug (for example, it prevents your project from building), and so can't tell whether your bug is present in a particular changeset, the `bisect` extension cannot help you directly. Instead, you'll need to manually avoid the changesets where that bug is present, and do separate searches “around” it.

A different problem could arise if your test for a bug's presence is not specific enough. If you check for “my program crashes”, then both your crashing bug and an unrelated crashing bug that masks it will look like the same thing, and mislead `bisect`.

9.6.5 Bracket your search lazily

Choosing the first “good” and “bad” changesets that will mark the end points of your search is often easy, but it bears a little discussion nevertheless. From the perspective of `bisect`, the “newest” changeset is conventionally “bad”, and the older changeset is “good”.

If you're having trouble remembering when a suitable “good” change was, so that you can tell `bisect`, you could do worse than testing changesets at random. Just remember to eliminate contenders that can't possibly exhibit the bug (perhaps because the feature with the bug isn't present yet) and those where another problem masks the bug (as I discussed above).

Even if you end up “early” by thousands of changesets or months of history, you will only add a handful of tests to the total number that `bisect` must perform, thanks to its logarithmic behaviour.

Chapter 10

Handling repository events with hooks

Mercurial offers a powerful mechanism to let you perform automated actions in response to events that occur in a repository. In some cases, you can even control Mercurial’s response to those events.

The name Mercurial uses for one of these actions is a *hook*. Hooks are called “triggers” in some revision control systems, but the two names refer to the same idea.

10.1 An overview of hooks in Mercurial

Here is a brief list of the hooks that Mercurial supports. We will revisit each of these hooks in more detail later, in section 10.8.

`changegroup` This is run after a group of changesets has been brought into the repository from elsewhere.

`commit` This is run after a new changeset has been created in the local repository.

`incoming` This is run once for each new changeset that is brought into the repository from elsewhere. Notice the difference from `changegroup`, which is run once per *group* of changesets brought in.

`outgoing` This is run after a group of changesets has been transmitted from this repository.

`prechangegroup` This is run before starting to bring a group of changesets into the repository.

`precommit` Controlling. This is run before starting a commit.

`preoutgoing` Controlling. This is run before starting to transmit a group of changesets from this repository.

`pretag` Controlling. This is run before creating a tag.

`pretxnchangegroup` Controlling. This is run after a group of changesets has been brought into the local repository from another, but before the transaction completes that will make the changes permanent in the repository.

`pretxncommit` Controlling. This is run after a new changeset has been created in the local repository, but before the transaction completes that will make it permanent.

`preupdate` Controlling. This is run before starting an update or merge of the working directory.

`tag` This is run after a tag is created.

`update` This is run after an update or merge of the working directory has finished.

Each of the hooks whose description begins with the word “Controlling” has the ability to determine whether an activity can proceed. If the hook succeeds, the activity may proceed; if it fails, the activity is either not permitted or undone, depending on the hook.

10.2 Hooks and security

10.2.1 Hooks are run with your privileges

When you run a Mercurial command in a repository, and the command causes a hook to run, that hook runs on *your* system, under *your* user account, with *your* privilege level. Since hooks are arbitrary pieces of executable code, you should treat them with an appropriate level of suspicion. Do not install a hook unless you are confident that you know who created it and what it does.

In some cases, you may be exposed to hooks that you did not install yourself. If you work with Mercurial on an unfamiliar system, Mercurial will run hooks defined in that system's global `hgrc` file.

If you are working with a repository owned by another user, Mercurial can run hooks defined in that user's repository, but it will still run them as "you". For example, if you "hg pull" from that repository, and its `.hg/hgrc` defines a local `outgoing` hook, that hook will run under your user account, even though you don't own that repository.

Note: This only applies if you are pulling from a repository on a local or network filesystem. If you're pulling over http or ssh, any `outgoing` hook will run under whatever account is executing the server process, on the server.

XXX To see what hooks are defined in a repository, use the "hg config hooks" command. If you are working in one repository, but talking to another that you do not own (e.g. using "hg pull" or "hg incoming"), remember that it is the other repository's hooks you should be checking, not your own.

10.2.2 Hooks do not propagate

In Mercurial, hooks are not revision controlled, and do not propagate when you clone, or pull from, a repository. The reason for this is simple: a hook is a completely arbitrary piece of executable code. It runs under your user identity, with your privilege level, on your machine.

It would be extremely reckless for any distributed revision control system to implement revision-controlled hooks, as this would offer an easily exploitable way to subvert the accounts of users of the revision control system.

Since Mercurial does not propagate hooks, if you are collaborating with other people on a common project, you should not assume that they are using the same Mercurial hooks as you are, or that theirs are correctly configured. You should document the hooks you expect people to use.

In a corporate intranet, this is somewhat easier to control, as you can for example provide a "standard" installation of Mercurial on an NFS filesystem, and use a site-wide `hgrc` file to define hooks that all users will see. However, this too has its limits; see below.

10.2.3 Hooks can be overridden

Mercurial allows you to override a hook definition by redefining the hook. You can disable it by setting its value to the empty string, or change its behaviour as you wish.

If you deploy a system- or site-wide `hgrc` file that defines some hooks, you should thus understand that your users can disable or override those hooks.

10.2.4 Ensuring that critical hooks are run

Sometimes you may want to enforce a policy that you do not want others to be able to work around. For example, you may have a requirement that every changeset must pass a rigorous set of tests. Defining this requirement via a hook in a site-wide `hgrc` won't work for remote users on laptops, and of course local users can subvert it at will by overriding the hook.

Instead, you can set up your policies for use of Mercurial so that people are expected to propagate changes through a well-known "canonical" server that you have locked down and configured appropriately.

One way to do this is via a combination of social engineering and technology. Set up a restricted-access account; users can push changes over the network to repositories managed by this account, but they cannot log into the account

and run normal shell commands. In this scenario, a user can commit a changeset that contains any old garbage they want.

When someone pushes a changeset to the server that everyone pulls from, the server will test the changeset before it accepts it as permanent, and reject it if it fails to pass the test suite. If people only pull changes from this filtering server, it will serve to ensure that all changes that people pull have been automatically vetted.

10.3 Care with `pretxn` hooks in a shared-access repository

If you want to use hooks to do some automated work in a repository that a number of people have shared access to, you need to be careful in how you do this.

Mercurial only locks a repository when it is writing to the repository, and only the parts of Mercurial that write to the repository pay attention to locks. Write locks are necessary to prevent multiple simultaneous writers from scribbling on each other's work, corrupting the repository.

Because Mercurial is careful with the order in which it reads and writes data, it does not need to acquire a lock when it wants to read data from the repository. The parts of Mercurial that read from the repository never pay attention to locks. This lockless reading scheme greatly increases performance and concurrency.

With great performance comes a trade-off, though, one which has the potential to cause you trouble unless you're aware of it. To describe this requires a little detail about how Mercurial adds changesets to a repository and reads those changes.

When Mercurial *writes* metadata, it writes it straight into the destination file. It writes file data first, then manifest data (which contains pointers to the new file data), then changelog data (which contains pointers to the new manifest data). Before the first write to each file, it stores a record of where the end of the file was in its transaction log. If the transaction must be rolled back, Mercurial simply truncates each file back to the size it was before the transaction began.

When Mercurial *reads* metadata, it reads the changelog first, then everything else. Since a reader will only access parts of the manifest or file metadata that it can see in the changelog, it can never see partially written data.

Some controlling hooks (`pretxncommit` and `pretxnchangegroup`) run when a transaction is almost complete. All of the metadata has been written, but Mercurial can still roll the transaction back and cause the newly-written data to disappear.

If one of these hooks runs for long, it opens a window of time during which a reader can see the metadata for changesets that are not yet permanent, and should not be thought of as “really there”. The longer the hook runs, the longer that window is open.

10.3.1 The problem illustrated

In principle, a good use for the `pretxnchangegroup` hook would be to automatically build and test incoming changes before they are accepted into a central repository. This could let you guarantee that nobody can push changes to this repository that “break the build”. But if a client can pull changes while they're being tested, the usefulness of the test is zero; an unsuspecting someone can pull untested changes, potentially breaking their build.

The safest technological answer to this challenge is to set up such a “gatekeeper” repository as *unidirectional*. Let it take changes pushed in from the outside, but do not allow anyone to pull changes from it (use the `preoutgoing` hook to lock it down). Configure a `changegroup` hook so that if a build or test succeeds, the hook will push the new changes out to another repository that people *can* pull from.

In practice, putting a centralised bottleneck like this in place is not often a good idea, and transaction visibility has nothing to do with the problem. As the size of a project—and the time it takes to build and test—grows, you rapidly run into a wall with this “try before you buy” approach, where you have more changesets to test than time in which to deal with them. The inevitable result is frustration on the part of all involved.

An approach that scales better is to get people to build and test before they push, then run automated builds and tests centrally *after* a push, to be sure all is well. The advantage of this approach is that it does not impose a limit on the rate at which the repository can accept changes.

10.4 A short tutorial on using hooks

It is easy to write a Mercurial hook. Let's start with a hook that runs when you finish a “hg commit”, and simply prints the hash of the changeset you just created. The hook is called `commit`.

```
1 $ hg init hook-test
2 $ cd hook-test
3 $ echo '[hooks]' >> .hg/hgrc
4 $ echo 'commit = echo committed $HG_NODE' >> .hg/hgrc
5 $ cat .hg/hgrc
6 [hooks]
7 commit = echo committed $HG_NODE
8 $ echo a > a
9 $ hg add a
10 $ hg commit -m 'testing commit hook'
11 committed 0d001c09cc4c4327b19b3d80bd47dfe53e8d7bc3
```

Figure 10.1: A simple hook that runs when a changeset is committed

All hooks follow the pattern in example 10.1. You add an entry to the `[hooks]` section of your `hgrc`. On the left is the name of the event to trigger on; on the right is the action to take. As you can see, you can run an arbitrary shell command in a hook. Mercurial passes extra information to the hook using environment variables (look for `HG_NODE` in the example).

10.4.1 Performing multiple actions per event

Quite often, you will want to define more than one hook for a particular kind of event, as shown in example 10.2. Mercurial lets you do this by adding an *extension* to the end of a hook's name. You extend a hook's name by giving the name of the hook, followed by a full stop (the “.” character), followed by some more text of your choosing. For example, Mercurial will run both `commit.foo` and `commit.bar` when the `commit` event occurs.

```
1 $ echo 'commit.when = echo -n "date of commit: "; date' >> .hg/hgrc
2 $ echo a >> a
3 $ hg commit -m 'i have two hooks'
4 committed 3f91730660135818354e79a2356bb164759b1b1a
5 date of commit: Mon Dec 10 19:54:30 GMT 2007
```

Figure 10.2: Defining a second `commit` hook

To give a well-defined order of execution when there are multiple hooks defined for an event, Mercurial sorts hooks by extension, and executes the hook commands in this sorted order. In the above example, it will execute `commit.bar` before `commit.foo`, and `commit` before both.

It is a good idea to use a somewhat descriptive extension when you define a new hook. This will help you to remember what the hook was for. If the hook fails, you'll get an error message that contains the hook name and extension, so using a descriptive extension could give you an immediate hint as to why the hook failed (see section 10.4.2 for an example).

10.4.2 Controlling whether an activity can proceed

In our earlier examples, we used the `commit` hook, which is run after a commit has completed. This is one of several Mercurial hooks that run after an activity finishes. Such hooks have no way of influencing the activity itself.

Mercurial defines a number of events that occur before an activity starts; or after it starts, but before it finishes. Hooks that trigger on these events have the added ability to choose whether the activity can continue, or will abort.

The `pretxncommit` hook runs after a commit has all but completed. In other words, the metadata representing the changeset has been written out to disk, but the transaction has not yet been allowed to complete. The `pretxncommit` hook has the ability to decide whether the transaction can complete, or must be rolled back.

If the `pretxncommit` hook exits with a status code of zero, the transaction is allowed to complete; the commit finishes; and the `commit` hook is run. If the `pretxncommit` hook exits with a non-zero status code, the transaction is rolled back; the metadata representing the changeset is erased; and the `commit` hook is not run.

```
1 $ cat check_bug_id
2 #!/bin/sh
3 # check that a commit comment mentions a numeric bug id
4 hg log -r $1 --template {desc} | grep -q "\<bug *[0-9]"
5 $ echo 'pretxncommit.bug_id_required = ./check_bug_id $HG_NODE' >> .hg/hgrc
6 $ echo a >> a
7 $ hg commit -m 'i am not mentioning a bug id'
8 transaction abort!
9 rollback completed
10 abort: pretxncommit.bug_id_required hook exited with status 1
11 $ hg commit -m 'i refer you to bug 666'
12 committed fa8e5405d879d86b46515af6b92054150dc43bc8
13 date of commit: Mon Dec 10 19:54:30 GMT 2007
```

Figure 10.3: Using the `pretxncommit` hook to control commits

The hook in example 10.3 checks that a commit comment contains a bug ID. If it does, the commit can complete. If not, the commit is rolled back.

10.5 Writing your own hooks

When you are writing a hook, you might find it useful to run Mercurial either with the `-v` option, or the `verbose` config item set to “true”. When you do so, Mercurial will print a message before it calls each hook.

10.5.1 Choosing how your hook should run

You can write a hook either as a normal program—typically a shell script—or as a Python function that is executed within the Mercurial process.

Writing a hook as an external program has the advantage that it requires no knowledge of Mercurial’s internals. You can call normal Mercurial commands to get any added information you need. The trade-off is that external hooks are slower than in-process hooks.

An in-process Python hook has complete access to the Mercurial API, and does not “shell out” to another process, so it is inherently faster than an external hook. It is also easier to obtain much of the information that a hook requires by using the Mercurial API than by running Mercurial commands.

If you are comfortable with Python, or require high performance, writing your hooks in Python may be a good choice. However, when you have a straightforward hook to write and you don’t need to care about performance (probably the majority of hooks), a shell script is perfectly fine.

10.5.2 Hook parameters

Mercurial calls each hook with a set of well-defined parameters. In Python, a parameter is passed as a keyword argument to your hook function. For an external program, a parameter is passed as an environment variable.

Whether your hook is written in Python or as a shell script, the hook-specific parameter names and values will be the same. A boolean parameter will be represented as a boolean value in Python, but as the number 1 (for “true”) or 0 (for “false”) as an environment variable for an external hook. If a hook parameter is named `foo`, the keyword argument for a Python hook will also be named `foo`, while the environment variable for an external hook will be named `HG_FOO`.

10.5.3 Hook return values and activity control

A hook that executes successfully must exit with a status of zero if external, or return boolean “false” if in-process. Failure is indicated with a non-zero exit status from an external hook, or an in-process hook returning boolean “true”. If an in-process hook raises an exception, the hook is considered to have failed.

For a hook that controls whether an activity can proceed, zero/false means “allow”, while non-zero/true/exception means “deny”.

10.5.4 Writing an external hook

When you define an external hook in your `hgrc` and the hook is run, its value is passed to your shell, which interprets it. This means that you can use normal shell constructs in the body of the hook.

An executable hook is always run with its current directory set to a repository’s root directory.

Each hook parameter is passed in as an environment variable; the name is upper-cased, and prefixed with the string “HG_”.

With the exception of hook parameters, Mercurial does not set or modify any environment variables when running a hook. This is useful to remember if you are writing a site-wide hook that may be run by a number of different users with differing environment variables set. In multi-user situations, you should not rely on environment variables being set to the values you have in your environment when testing the hook.

10.5.5 Telling Mercurial to use an in-process hook

The `hgrc` syntax for defining an in-process hook is slightly different than for an executable hook. The value of the hook must start with the text “python:”, and continue with the fully-qualified name of a callable object to use as the hook’s value.

The module in which a hook lives is automatically imported when a hook is run. So long as you have the module name and `PYTHONPATH` right, it should “just work”.

The following `hgrc` example snippet illustrates the syntax and meaning of the notions we just described.

```
1 [hooks]
2 commit.example = python:mymodule.submodule.myhook
```

When Mercurial runs the `commit.example` hook, it imports `mymodule.submodule`, looks for the callable object named `myhook`, and calls it.

10.5.6 Writing an in-process hook

The simplest in-process hook does nothing, but illustrates the basic shape of the hook API:

```
1 def myhook(ui, repo, **kwargs):
2     pass
```

The first argument to a Python hook is always a `mercurial.ui.ui` object. The second is a repository object; at the moment, it is always an instance of `mercurial.localrepo.localrepository`. Following these two arguments are other keyword arguments. Which ones are passed in depends on the hook being called, but a hook can ignore arguments it doesn't care about by dropping them into a keyword argument dict, as with `**kwargs` above.

10.6 Some hook examples

10.6.1 Writing meaningful commit messages

It's hard to imagine a useful commit message being very short. The simple `pretxncommit` hook of figure 10.4 will prevent you from committing a changeset with a message that is less than ten bytes long.

```
1 $ cat .hg/hgrc
2 [hooks]
3 pretxncommit.msglen = test 'hg tip --template {desc} | wc -c' -ge 10
4 $ echo a > a
5 $ hg add a
6 $ hg commit -A -m 'too short'
7 transaction abort!
8 rollback completed
9 abort: pretxncommit.msglen hook exited with status 1
10 $ hg commit -A -m 'long enough'
```

Figure 10.4: A hook that forbids overly short commit messages

10.6.2 Checking for trailing whitespace

An interesting use of a commit-related hook is to help you to write cleaner code. A simple example of “cleaner code” is the dictum that a change should not add any new lines of text that contain “trailing whitespace”. Trailing whitespace is a series of space and tab characters at the end of a line of text. In most cases, trailing whitespace is unnecessary, invisible noise, but it is occasionally problematic, and people often prefer to get rid of it.

You can use either the `precommit` or `pretxncommit` hook to tell whether you have a trailing whitespace problem. If you use the `precommit` hook, the hook will not know which files you are committing, so it will have to check every modified file in the repository for trailing white space. If you want to commit a change to just the file `foo`, but the file `bar` contains trailing whitespace, doing a check in the `precommit` hook will prevent you from committing `foo` due to the problem with `bar`. This doesn't seem right.

Should you choose the `pretxncommit` hook, the check won't occur until just before the transaction for the commit completes. This will allow you to check for problems only the exact files that are being committed. However, if you entered the commit message interactively and the hook fails, the transaction will roll back; you'll have to re-enter the commit message after you fix the trailing whitespace and run “`hg commit`” again.

Figure 10.5 introduces a simple `pretxncommit` hook that checks for trailing whitespace. This hook is short, but not very helpful. It exits with an error status if a change adds a line with trailing whitespace to any file, but does not print any information that might help us to identify the offending file or line. It also has the nice property of not paying attention to unmodified lines; only lines that introduce new trailing whitespace cause problems.

The example of figure 10.6 is much more complex, but also more useful. It parses a unified diff to see if any lines add trailing whitespace, and prints the name of the file and the line number of each such occurrence. Even better, if the change adds trailing whitespace, this hook saves the commit comment and prints the name of the save file before exiting and telling Mercurial to roll the transaction back, so you can use “`hg commit -l filename`” to reuse the saved commit message once you've corrected the problem.

```

1 $ cat .hg/hgrc
2 [hooks]
3 pretxncommit.whitespace = hg export tip | (! egrep -q '^\.+.*[ \t]$')
4 $ echo 'a ' > a
5 $ hg commit -A -m 'test with trailing whitespace'
6 adding a
7 transaction abort!
8 rollback completed
9 abort: pretxncommit.whitespace hook exited with status 1
10 $ echo 'a' > a
11 $ hg commit -A -m 'drop trailing whitespace and try again'

```

Figure 10.5: A simple hook that checks for trailing whitespace

As a final aside, note in figure 10.6 the use of perl’s in-place editing feature to get rid of trailing whitespace from a file. This is concise and useful enough that I will reproduce it here.

```

1 perl -pi -e 's,                s+$,, ' filename

```

10.7 Bundled hooks

Mercurial ships with several bundled hooks. You can find them in the `hgext` directory of a Mercurial source tree. If you are using a Mercurial binary package, the hooks will be located in the `hgext` directory of wherever your package installer put Mercurial.

10.7.1 `acl`—access control for parts of a repository

The `acl` extension lets you control which remote users are allowed to push changesets to a networked server. You can protect any portion of a repository (including the entire repo), so that a specific remote user can push changes that do not affect the protected portion.

This extension implements access control based on the identity of the user performing a push, *not* on who committed the changesets they’re pushing. It makes sense to use this hook only if you have a locked-down server environment that authenticates remote users, and you want to be sure that only specific users are allowed to push changes to that server.

Configuring the `acl` hook

In order to manage incoming changesets, the `acl` hook must be used as a `pretxnchangegroup` hook. This lets it see which files are modified by each incoming changeset, and roll back a group of changesets if they modify “forbidden” files. Example:

```

1 [hooks]
2 pretxnchangegroup.acl = python:hgext.acl.hook

```

The `acl` extension is configured using three sections.

The `[acl]` section has only one entry, `sources`, which lists the sources of incoming changesets that the hook should pay attention to. You don’t normally need to configure this section.

`serve` Control incoming changesets that are arriving from a remote repository over http or ssh. This is the default value of `sources`, and usually the only setting you’ll need for this configuration item.

```

1 $ cat .hg/hgrc
2 [hooks]
3 pretxncommit.whitespace = .hg/check_whitespace.py
4 $ echo 'a ' >> a
5 $ hg commit -A -m 'add new line with trailing whitespace'
6 a, line 2: trailing whitespace added
7 commit message saved to .hg/commit.save
8 transaction abort!
9 rollback completed
10 abort: pretxncommit.whitespace hook exited with status 1
11 $ sed -i 's, *$,,' a
12 $ hg commit -A -m 'trimmed trailing whitespace'
13 a, line 2: trailing whitespace added
14 commit message saved to .hg/commit.save
15 transaction abort!
16 rollback completed
17 abort: pretxncommit.whitespace hook exited with status 1

```

Figure 10.6: A better trailing whitespace hook

pull Control incoming changesets that are arriving via a pull from a local repository.

push Control incoming changesets that are arriving via a push from a local repository.

bundle Control incoming changesets that are arriving from another repository via a bundle.

The `[acl.allow]` section controls the users that are allowed to add changesets to the repository. If this section is not present, all users that are not explicitly denied are allowed. If this section is present, all users that are not explicitly allowed are denied (so an empty section means that all users are denied).

The `[acl.deny]` section determines which users are denied from adding changesets to the repository. If this section is not present or is empty, no users are denied.

The syntaxes for the `[acl.allow]` and `[acl.deny]` sections are identical. On the left of each entry is a glob pattern that matches files or directories, relative to the root of the repository; on the right, a user name.

In the following example, the user `docwriter` can only push changes to the `docs` subtree of the repository, while `intern` can push changes to any file or directory except `source/sensitive`.

```

1 [acl.allow]
2 docs/** = docwriter
3
4 [acl.deny]
5 source/sensitive/** = intern

```

Testing and troubleshooting

If you want to test the `acl` hook, run it with Mercurial's debugging output enabled. Since you'll probably be running it on a server where it's not convenient (or sometimes possible) to pass in the `--debug` option, don't forget that you can enable debugging output in your `hgrc`:

```

1 [ui]
2 debug = true

```

With this enabled, the `acl` hook will print enough information to let you figure out why it is allowing or forbidding pushes from specific users.

10.7.2 `bugzilla`—integration with Bugzilla

The `bugzilla` extension adds a comment to a Bugzilla bug whenever it finds a reference to that bug ID in a commit comment. You can install this hook on a shared server, so that any time a remote user pushes changes to this server, the hook gets run.

It adds a comment to the bug that looks like this (you can configure the contents of the comment—see below):

```
1 Changeset aad8b264143a, made by Joe User <joe.user@domain.com> in
2 the frobnitz repository, refers to this bug.
3
4 For complete details, see
5 http://hg.domain.com/frobnitz?cmd=changeset;node=aad8b264143a
6
7 Changeset description:
8     Fix bug 10483 by guarding against some NULL pointers
```

The value of this hook is that it automates the process of updating a bug any time a changeset refers to it. If you configure the hook properly, it makes it easy for people to browse straight from a Bugzilla bug to a changeset that refers to that bug.

You can use the code in this hook as a starting point for some more exotic Bugzilla integration recipes. Here are a few possibilities:

- Require that every changeset pushed to the server have a valid bug ID in its commit comment. In this case, you'd want to configure the hook as a `pretxncommit` hook. This would allow the hook to reject changes that didn't contain bug IDs.
- Allow incoming changesets to automatically modify the *state* of a bug, as well as simply adding a comment. For example, the hook could recognise the string "fixed bug 31337" as indicating that it should update the state of bug 31337 to "requires testing".

Configuring the `bugzilla` hook

You should configure this hook in your server's `hgrc` as an `incoming` hook, for example as follows:

```
1 [hooks]
2 incoming.bugzilla = python:hgext.bugzilla.hook
```

Because of the specialised nature of this hook, and because Bugzilla was not written with this kind of integration in mind, configuring this hook is a somewhat involved process.

Before you begin, you must install the MySQL bindings for Python on the host(s) where you'll be running the hook. If this is not available as a binary package for your system, you can download it from [\[Dus\]](#).

Configuration information for this hook lives in the `[bugzilla]` section of your `hgrc`.

version The version of Bugzilla installed on the server. The database schema that Bugzilla uses changes occasionally, so this hook has to know exactly which schema to use. At the moment, the only version supported is 2.16.

host The hostname of the MySQL server that stores your Bugzilla data. The database must be configured to allow connections from whatever host you are running the `bugzilla` hook on.

user The username with which to connect to the MySQL server. The database must be configured to allow this user to connect from whatever host you are running the `bugzilla` hook on. This user must be able to access and modify Bugzilla tables. The default value of this item is `bugs`, which is the standard name of the Bugzilla user in a MySQL database.

password The MySQL password for the user you configured above. This is stored as plain text, so you should make sure that unauthorised users cannot read the `hgrc` file where you store this information.

db The name of the Bugzilla database on the MySQL server. The default value of this item is `bugs`, which is the standard name of the MySQL database where Bugzilla stores its data.

notify If you want Bugzilla to send out a notification email to subscribers after this hook has added a comment to a bug, you will need this hook to run a command whenever it updates the database. The command to run depends on where you have installed Bugzilla, but it will typically look something like this, if you have Bugzilla installed in `/var/www/html/bugzilla`:

```
1 cd /var/www/html/bugzilla && ./processmail %s nobody@nowhere.com
```

The Bugzilla `processmail` program expects to be given a bug ID (the hook replaces “%s” with the bug ID) and an email address. It also expects to be able to write to some files in the directory that it runs in. If Bugzilla and this hook are not installed on the same machine, you will need to find a way to run `processmail` on the server where Bugzilla is installed.

Mapping committer names to Bugzilla user names

By default, the `bugzilla` hook tries to use the email address of a changeset’s committer as the Bugzilla user name with which to update a bug. If this does not suit your needs, you can map committer email addresses to Bugzilla user names using a `[usermap]` section.

Each item in the `[usermap]` section contains an email address on the left, and a Bugzilla user name on the right.

```
1 [usermap]
2 jane.user@example.com = jane
```

You can either keep the `[usermap]` data in a normal `hgrc`, or tell the `bugzilla` hook to read the information from an external `usermap` file. In the latter case, you can store `usermap` data by itself in (for example) a user-modifiable repository. This makes it possible to let your users maintain their own `usermap` entries. The main `hgrc` file might look like this:

```
1 # regular hgrc file refers to external usermap file
2 [bugzilla]
3 usermap = /home/hg/repos/userdata/bugzilla-usermap.conf
```

While the `usermap` file that it refers to might look like this:

```
1 # bugzilla-usermap.conf - inside a hg repository
2 [usermap]
3 stephanie@example.com = steph
```

Configuring the text that gets added to a bug

You can configure the text that this hook adds as a comment; you specify it in the form of a Mercurial template. Several `hgrc` entries (still in the `[bugzilla]` section) control this behaviour.

strip The number of leading path elements to strip from a repository's path name to construct a partial path for a URL. For example, if the repositories on your server live under `/home/hg/repos`, and you have a repository whose path is `/home/hg/repos/app/tests`, then setting `strip` to 4 will give a partial path of `app/tests`. The hook will make this partial path available when expanding a template, as `webroot`.

template The text of the template to use. In addition to the usual changeset-related variables, this template can use `hgweb` (the value of the `hgweb` configuration item above) and `webroot` (the path constructed using `strip` above).

In addition, you can add a `baseurl` item to the `[web]` section of your `hgrc`. The `bugzilla` hook will make this available when expanding a template, as the base string to use when constructing a URL that will let users browse from a Bugzilla comment to view a changeset. Example:

```
1 [web]
2 baseurl = http://hg.domain.com/
```

Here is an example set of `bugzilla` hook config information.

```
1 [bugzilla]
2 host = bugzilla.example.com
3 password = mypassword
4 version = 2.16
5 # server-side repos live in /home/hg/repos, so strip 4 leading
6 # separators
7 strip = 4
8 hgweb = http://hg.example.com/
9 usermap = /home/hg/repos/notify/bugzilla.conf
10 template = Changeset {node|short}, made by {author} in the {webroot}
11     repo, refers to this bug.                                nFor complete details, see
12     {hgweb}{webroot}?cmd=changeset;node={node|short}         nChangeset
13     description:                                             n          t{desc|tabindent}
```

Testing and troubleshooting

The most common problems with configuring the `bugzilla` hook relate to running Bugzilla's `processmail` script and mapping committer names to user names.

Recall from section 10.7.2 above that the user that runs the Mercurial process on the server is also the one that will run the `processmail` script. The `processmail` script sometimes causes Bugzilla to write to files in its configuration directory, and Bugzilla's configuration files are usually owned by the user that your web server runs under.

You can cause `processmail` to be run with the suitable user's identity using the `sudo` command. Here is an example entry for a `sudoers` file.

```
1 hg_user = (httpd_user) NOPASSWD: /var/www/html/bugzilla/processmail-wrapper %s
```

This allows the `hg_user` user to run a `processmail-wrapper` program under the identity of `httpd_user`.

This indirection through a wrapper script is necessary, because `processmail` expects to be run with its current directory set to wherever you installed Bugzilla; you can't specify that kind of constraint in a `sudoers` file. The contents of the wrapper script are simple:

```
1 #!/bin/sh
2 cd `dirname $0` && ./processmail "$1" nobody@example.com
```

It doesn't seem to matter what email address you pass to `processmail`.

If your `[usermap]` is not set up correctly, users will see an error message from the `bugzilla` hook when they push changes to the server. The error message will look like this:

```
cannot find bugzilla user id for john.q.public@example.com
```

What this means is that the committer's address, `john.q.public@example.com`, is not a valid Bugzilla user name, nor does it have an entry in your `[usermap]` that maps it to a valid Bugzilla user name.

10.7.3 `notify`—send email notifications

Although Mercurial's built-in web server provides RSS feeds of changes in every repository, many people prefer to receive change notifications via email. The `notify` hook lets you send out notifications to a set of email addresses whenever changesets arrive that those subscribers are interested in.

As with the `bugzilla` hook, the `notify` hook is template-driven, so you can customise the contents of the notification messages that it sends.

By default, the `notify` hook includes a diff of every changeset that it sends out; you can limit the size of the diff, or turn this feature off entirely. It is useful for letting subscribers review changes immediately, rather than clicking to follow a URL.

Configuring the `notify` hook

You can set up the `notify` hook to send one email message per incoming changeset, or one per incoming group of changesets (all those that arrived in a single pull or push).

```
[hooks]
# send one email per group of changes
changegroup.notify = python:hgext.notify.hook
# send one email per change
incoming.notify = python:hgext.notify.hook
```

Configuration information for this hook lives in the `[notify]` section of a `hgrc` file.

test By default, this hook does not send out email at all; instead, it prints the message that it *would* send. Set this item to `false` to allow email to be sent. The reason that sending of email is turned off by default is that it takes several tries to configure this extension exactly as you would like, and it would be bad form to spam subscribers with a number of “broken” notifications while you debug your configuration.

config The path to a configuration file that contains subscription information. This is kept separate from the main `hgrc` so that you can maintain it in a repository of its own. People can then clone that repository, update their subscriptions, and push the changes back to your server.

strip The number of leading path separator characters to strip from a repository's path, when deciding whether a repository has subscribers. For example, if the repositories on your server live in `/home/hg/repos`, and `notify` is considering a repository named `/home/hg/repos/shared/test`, setting `strip` to 4 will cause `notify` to trim the path it considers down to `shared/test`, and it will match subscribers against that.

template The template text to use when sending messages. This specifies both the contents of the message header and its body.

maxdiff The maximum number of lines of diff data to append to the end of a message. If a diff is longer than this, it is truncated. By default, this is set to 300. Set this to 0 to omit diffs from notification emails.

sources A list of sources of changesets to consider. This lets you limit `notify` to only sending out email about changes that remote users pushed into this repository via a server, for example. See section 10.8.3 for the sources you can specify here.

If you set the `baseurl` item in the `[web]` section, you can use it in a template; it will be available as `webroot`. Here is an example set of `notify` configuration information.

```
1 [notify]
2 # really send email
3 test = false
4 # subscriber data lives in the notify repo
5 config = /home/hg/repos/notify/notify.conf
6 # repos live in /home/hg/repos on server, so strip 4 "/" chars
7 strip = 4
8 template = X-Hg-Repo: {webroot} n
9 Subject: {webroot}: {desc|firstline|strip} n
10 From: {author} n
11 n
12 changeset {node|short} in {root} n
13 details: {baseurl}{webroot}?cmd=changeset;node={node|short} n
14 description: n
15 t{desc|tabindent|strip}
16
17 [web]
18 baseurl = http://hg.example.com/
```

This will produce a message that looks like the following:

```
1 X-Hg-Repo: tests/slave
2 Subject: tests/slave: Handle error case when slave has no buffers
3 Date: Wed, 2 Aug 2006 15:25:46 -0700 (PDT)
4
5 changeset 3cba9bfe74b5 in /home/hg/repos/tests/slave
6 details: http://hg.example.com/tests/slave?cmd=changeset;node=3cba9bfe74b5
7 description:
8     Handle error case when slave has no buffers
9 diffs (54 lines):
10
11 diff -r 9d95df7cf2ad -r 3cba9bfe74b5 include/tests.h
12 --- a/include/tests.h      Wed Aug 02 15:19:52 2006 -0700
13 +++ b/include/tests.h      Wed Aug 02 15:25:26 2006 -0700
14 @@ -212,6 +212,15 @@ static __inline__ void test_headers(void *h)
15 [...snip...]
```

Testing and troubleshooting

Do not forget that by default, the `notify` extension *will not send any mail* until you explicitly configure it to do so, by setting `test` to `false`. Until you do that, it simply prints the message it *would* send.

10.8 Information for writers of hooks

10.8.1 In-process hook execution

An in-process hook is called with arguments of the following form:

```

1 def myhook(ui, repo, **kwargs):
2     pass

```

The `ui` parameter is a `mercurial.ui.ui` object. The `repo` parameter is a `mercurial.localrepo.localrepository` object. The names and values of the `**kwargs` parameters depend on the hook being invoked, with the following common features:

- If a parameter is named `node` or `parentN`, it will contain a hexadecimal changeset ID. The empty string is used to represent “null changeset ID” instead of a string of zeroes.
- If a parameter is named `url`, it will contain the URL of a remote repository, if that can be determined.
- Boolean-valued parameters are represented as Python `bool` objects.

An in-process hook is called without a change to the process’s working directory (unlike external hooks, which are run in the root of the repository). It must not change the process’s working directory, or it will cause any calls it makes into the Mercurial API to fail.

If a hook returns a boolean “false” value, it is considered to have succeeded. If it returns a boolean “true” value or raises an exception, it is considered to have failed. A useful way to think of the calling convention is “tell me if you fail”.

Note that changeset IDs are passed into Python hooks as hexadecimal strings, not the binary hashes that Mercurial’s APIs normally use. To convert a hash from hex to binary, use the `mercurial.node.bin` function.

10.8.2 External hook execution

An external hook is passed to the shell of the user running Mercurial. Features of that shell, such as variable substitution and command redirection, are available. The hook is run in the root directory of the repository (unlike in-process hooks, which are run in the same directory that Mercurial was run in).

Hook parameters are passed to the hook as environment variables. Each environment variable’s name is converted in upper case and prefixed with the string “HG_”. For example, if the name of a parameter is “node”, the name of the environment variable representing that parameter will be “HG_NODE”.

A boolean parameter is represented as the string “1” for “true”, “0” for “false”. If an environment variable is named `HG_NODE`, `HG_PARENT1` or `HG_PARENT2`, it contains a changeset ID represented as a hexadecimal string. The empty string is used to represent “null changeset ID” instead of a string of zeroes. If an environment variable is named `HG_URL`, it will contain the URL of a remote repository, if that can be determined.

If a hook exits with a status of zero, it is considered to have succeeded. If it exits with a non-zero status, it is considered to have failed.

10.8.3 Finding out where changesets come from

A hook that involves the transfer of changesets between a local repository and another may be able to find out information about the “far side”. Mercurial knows *how* changes are being transferred, and in many cases *where* they are being transferred to or from.

Sources of changesets

Mercurial will tell a hook what means are, or were, used to transfer changesets between repositories. This is provided by Mercurial in a Python parameter named `source`, or an environment variable named `HG_SOURCE`.

`serve` Changesets are transferred to or from a remote repository over http or ssh.

`pull` Changesets are being transferred via a pull from one repository into another.

`push` Changesets are being transferred via a push from one repository into another.

`bundle` Changesets are being transferred to or from a bundle.

Where changes are going—remote repository URLs

When possible, Mercurial will tell a hook the location of the “far side” of an activity that transfers changeset data between repositories. This is provided by Mercurial in a Python parameter named `url`, or an environment variable named `HG_URL`.

This information is not always known. If a hook is invoked in a repository that is being served via `http` or `ssh`, Mercurial cannot tell where the remote repository is, but it may know where the client is connecting from. In such cases, the URL will take one of the following forms:

- `remote:ssh:ip-address`—remote `ssh` client, at the given IP address.
- `remote:http:ip-address`—remote `http` client, at the given IP address. If the client is using `SSL`, this will be of the form `remote:https:ip-address`.
- Empty—no information could be discovered about the remote client.

10.9 Hook reference

10.9.1 `changegroup`—after remote changesets added

This hook is run after a group of pre-existing changesets has been added to the repository, for example via a “`hg pull`” or “`hg unbundle`”. This hook is run once per operation that added one or more changesets. This is in contrast to the `incoming` hook, which is run once per changeset, regardless of whether the changesets arrive in a group.

Some possible uses for this hook include kicking off an automated build or test of the added changesets, updating a bug database, or notifying subscribers that a repository contains new changes.

Parameters to this hook:

`node` A changeset ID. The changeset ID of the first changeset in the group that was added. All changesets between this and `tip`, inclusive, were added by a single “`hg pull`”, “`hg push`” or “`hg unbundle`”.

`source` A string. The source of these changes. See section 10.8.3 for details.

`url` A URL. The location of the remote repository, if known. See section 10.8.3 for more information.

See also: `incoming` (section 10.9.3), `prechangegroup` (section 10.9.5), `pretxnchangegroup` (section 10.9.9)

10.9.2 `commit`—after a new changeset is created

This hook is run after a new changeset has been created.

Parameters to this hook:

`node` A changeset ID. The changeset ID of the newly committed changeset.

`parent1` A changeset ID. The changeset ID of the first parent of the newly committed changeset.

`parent2` A changeset ID. The changeset ID of the second parent of the newly committed changeset.

See also: `precommit` (section 10.9.6), `pretxncommit` (section 10.9.10)

10.9.3 `incoming`—after one remote changeset is added

This hook is run after a pre-existing changeset has been added to the repository, for example via a “`hg push`”. If a group of changesets was added in a single operation, this hook is called once for each added changeset.

You can use this hook for the same purposes as the `changegroup` hook (section 10.9.1); it’s simply more convenient sometimes to run a hook once per group of changesets, while other times it’s handier once per changeset.

Parameters to this hook:

`node` A changeset ID. The ID of the newly added changeset.

`source` A string. The source of these changes. See section 10.8.3 for details.

`url` A URL. The location of the remote repository, if known. See section 10.8.3 for more information.

See also: `changegroup` (section 10.9.1) `prechangegroup` (section 10.9.5), `pretxnchangegroup` (section 10.9.9)

10.9.4 outgoing—after changesets are propagated

This hook is run after a group of changesets has been propagated out of this repository, for example by a “hg push” or “hg bundle” command.

One possible use for this hook is to notify administrators that changes have been pulled.

Parameters to this hook:

`node` A changeset ID. The changeset ID of the first changeset of the group that was sent.

`source` A string. The source of the of the operation (see section 10.8.3). If a remote client pulled changes from this repository, `source` will be `serve`. If the client that obtained changes from this repository was local, `source` will be `bundle`, `pull`, or `push`, depending on the operation the client performed.

`url` A URL. The location of the remote repository, if known. See section 10.8.3 for more information.

See also: `preoutgoing` (section 10.9.7)

10.9.5 prechangegroup—before starting to add remote changesets

This controlling hook is run before Mercurial begins to add a group of changesets from another repository.

This hook does not have any information about the changesets to be added, because it is run before transmission of those changesets is allowed to begin. If this hook fails, the changesets will not be transmitted.

One use for this hook is to prevent external changes from being added to a repository. For example, you could use this to “freeze” a server-hosted branch temporarily or permanently so that users cannot push to it, while still allowing a local administrator to modify the repository.

Parameters to this hook:

`source` A string. The source of these changes. See section 10.8.3 for details.

`url` A URL. The location of the remote repository, if known. See section 10.8.3 for more information.

See also: `changegroup` (section 10.9.1), `incoming` (section 10.9.3), `pretxnchangegroup` (section 10.9.9)

10.9.6 precommit—before starting to commit a changeset

This hook is run before Mercurial begins to commit a new changeset. It is run before Mercurial has any of the metadata for the commit, such as the files to be committed, the commit message, or the commit date.

One use for this hook is to disable the ability to commit new changesets, while still allowing incoming changesets. Another is to run a build or test, and only allow the commit to begin if the build or test succeeds.

Parameters to this hook:

`parent1` A changeset ID. The changeset ID of the first parent of the working directory.

`parent2` A changeset ID. The changeset ID of the second parent of the working directory.

If the commit proceeds, the parents of the working directory will become the parents of the new changeset.

See also: `commit` (section 10.9.2), `pretxncommit` (section 10.9.10)

10.9.7 **preoutgoing**—before starting to propagate changesets

This hook is invoked before Mercurial knows the identities of the changesets to be transmitted.

One use for this hook is to prevent changes from being transmitted to another repository.

Parameters to this hook:

source A string. The source of the operation that is attempting to obtain changes from this repository (see section 10.8.3). See the documentation for the `source` parameter to the `outgoing` hook, in section 10.9.4, for possible values of this parameter.

url A URL. The location of the remote repository, if known. See section 10.8.3 for more information.

See also: `outgoing` (section 10.9.4)

10.9.8 **pretag**—before tagging a changeset

This controlling hook is run before a tag is created. If the hook succeeds, creation of the tag proceeds. If the hook fails, the tag is not created.

Parameters to this hook:

local A boolean. Whether the tag is local to this repository instance (i.e. stored in `.hg/localtags`) or managed by Mercurial (stored in `.hgtags`).

node A changeset ID. The ID of the changeset to be tagged.

tag A string. The name of the tag to be created.

If the tag to be created is revision-controlled, the `precommit` and `pretxncommit` hooks (sections 10.9.2 and 10.9.10) will also be run.

See also: `tag` (section 10.9.12)

10.9.9 **pretxnchangegroup**—before completing addition of remote changesets

This controlling hook is run before a transaction—that manages the addition of a group of new changesets from outside the repository—completes. If the hook succeeds, the transaction completes, and all of the changesets become permanent within this repository. If the hook fails, the transaction is rolled back, and the data for the changesets is erased.

This hook can access the metadata associated with the almost-added changesets, but it should not do anything permanent with this data. It must also not modify the working directory.

While this hook is running, if other Mercurial processes access this repository, they will be able to see the almost-added changesets as if they are permanent. This may lead to race conditions if you do not take steps to avoid them.

This hook can be used to automatically vet a group of changesets. If the hook fails, all of the changesets are “rejected” when the transaction rolls back.

Parameters to this hook:

node A changeset ID. The changeset ID of the first changeset in the group that was added. All changesets between this and `tip`, inclusive, were added by a single “`hg pull`”, “`hg push`” or “`hg unbundle`”.

source A string. The source of these changes. See section 10.8.3 for details.

url A URL. The location of the remote repository, if known. See section 10.8.3 for more information.

See also: `changegroup` (section 10.9.1), `incoming` (section 10.9.3), `prechangegroup` (section 10.9.5)

10.9.10 `pretxncommit`—before completing commit of new changeset

This controlling hook is run before a transaction—that manages a new commit—completes. If the hook succeeds, the transaction completes and the changeset becomes permanent within this repository. If the hook fails, the transaction is rolled back, and the commit data is erased.

This hook can access the metadata associated with the almost-new changeset, but it should not do anything permanent with this data. It must also not modify the working directory.

While this hook is running, if other Mercurial processes access this repository, they will be able to see the almost-new changeset as if it is permanent. This may lead to race conditions if you do not take steps to avoid them.

Parameters to this hook:

`node` A changeset ID. The changeset ID of the newly committed changeset.

`parent1` A changeset ID. The changeset ID of the first parent of the newly committed changeset.

`parent2` A changeset ID. The changeset ID of the second parent of the newly committed changeset.

See also: `precommit` (section 10.9.6)

10.9.11 `preupdate`—before updating or merging working directory

This controlling hook is run before an update or merge of the working directory begins. It is run only if Mercurial's normal pre-update checks determine that the update or merge can proceed. If the hook succeeds, the update or merge may proceed; if it fails, the update or merge does not start.

Parameters to this hook:

`parent1` A changeset ID. The ID of the parent that the working directory is to be updated to. If the working directory is being merged, it will not change this parent.

`parent2` A changeset ID. Only set if the working directory is being merged. The ID of the revision that the working directory is being merged with.

See also: `update` (section 10.9.13)

10.9.12 `tag`—after tagging a changeset

This hook is run after a tag has been created.

Parameters to this hook:

`local` A boolean. Whether the new tag is local to this repository instance (i.e. stored in `.hg/localtags`) or managed by Mercurial (stored in `.hgtags`).

`node` A changeset ID. The ID of the changeset that was tagged.

`tag` A string. The name of the tag that was created.

If the created tag is revision-controlled, the `commit` hook (section 10.9.2) is run before this hook.

See also: `pretag` (section 10.9.8)

10.9.13 `update`—after updating or merging working directory

This hook is run after an update or merge of the working directory completes. Since a merge can fail (if the external `hgmerge` command fails to resolve conflicts in a file), this hook communicates whether the update or merge completed cleanly.

`error` A boolean. Indicates whether the update or merge completed successfully.

parent1 A changeset ID. The ID of the parent that the working directory was updated to. If the working directory was merged, it will not have changed this parent.

parent2 A changeset ID. Only set if the working directory was merged. The ID of the revision that the working directory was merged with.

See also: `preupdate` (section [10.9.11](#))

Chapter 11

Customising the output of Mercurial

Mercurial provides a powerful mechanism to let you control how it displays information. The mechanism is based on templates. You can use templates to generate specific output for a single command, or to customise the entire appearance of the built-in web interface.

11.1 Using precanned output styles

Packaged with Mercurial are some output styles that you can use immediately. A style is simply a precanned template that someone wrote and installed somewhere that Mercurial can find.

Before we take a look at Mercurial’s bundled styles, let’s review its normal output.

```
1 $ hg log -r1
2 changeset: 1:66687f9c6bca
3 tag:      mytag
4 user:     Bryan O'Sullivan <bos@serpentine.com>
5 date:     Mon Dec 10 19:54:40 2007 +0000
6 summary:  added line to end of <<hello>> file.
7
```

This is somewhat informative, but it takes up a lot of space—five lines of output per changeset. The `compact` style reduces this to three lines, presented in a sparse manner.

```
1 $ hg log --style compact
2 3[tip] e6bc04c0642a 2007-12-10 19:54 +0000 bos
3   Added tag v0.1 for changeset 558b9b187012
4
5 2[v0.1] 558b9b187012 2007-12-10 19:54 +0000 bos
6   Added tag mytag for changeset 66687f9c6bca
7
8 1[mytag] 66687f9c6bca 2007-12-10 19:54 +0000 bos
9   added line to end of <<hello>> file.
10
11 0 461d417cd5d5 2007-12-10 19:54 +0000 bos
12   added hello
13
```

The `changelog` style hints at the expressive power of Mercurial’s templating engine. This style attempts to follow the GNU Project’s changelog guidelines[RS].

```

1 $ hg log --style changelog
2 2007-12-10 Bryan O'Sullivan <bos@serpentine.com>
3
4     * .hgtags:
5     Added tag v0.1 for changeset 558b9b187012
6     [e6bc04c0642a] [tip]
7
8     * .hgtags:
9     Added tag mytag for changeset 66687f9c6bca
10    [558b9b187012] [v0.1]
11
12    * goodbye, hello:
13    added line to end of <<hello>> file.
14
15    in addition, added a file with the helpful name (at least i hope
16    that some might consider it so) of goodbye.
17    [66687f9c6bca] [mytag]
18
19    * hello:
20    added hello
21    [461d417cd5d5]
22

```

You will not be shocked to learn that Mercurial's default output style is named `default`.

11.1.1 Setting a default style

You can modify the output style that Mercurial will use for every command by editing your `hgrc` file, naming the style you would prefer to use.

```

1 [ui]
2 style = compact

```

If you write a style of your own, you can use it by either providing the path to your style file, or copying your style file into a location where Mercurial can find it (typically the `templates` subdirectory of your Mercurial install directory).

11.2 Commands that support styles and templates

All of Mercurial's "log-like" commands let you use styles and templates: "hg incoming", "hg log", "hg outgoing", and "hg tip".

As I write this manual, these are so far the only commands that support styles and templates. Since these are the most important commands that need customisable output, there has been little pressure from the Mercurial user community to add style and template support to other commands.

11.3 The basics of templating

At its simplest, a Mercurial template is a piece of text. Some of the text never changes, while other parts are *expanded*, or replaced with new text, when necessary.

Before we continue, let's look again at a simple example of Mercurial's normal output.

```

1 $ hg log -r1
2 changeset: 1:66687f9c6bca
3 tag:      mytag
4 user:      Bryan O'Sullivan <bos@serpentine.com>
5 date:      Mon Dec 10 19:54:40 2007 +0000
6 summary:   added line to end of <<hello>> file.
7

```

Now, let's run the same command, but using a template to change its output.

```

1 $ hg log -r1 --template 'i saw a changeset\n'
2 i saw a changeset

```

The example above illustrates the simplest possible template; it's just a piece of static text, printed once for each changeset. The `--template` option to the “hg log” command tells Mercurial to use the given text as the template when printing each changeset.

Notice that the template string above ends with the text “\n”. This is an *escape sequence*, telling Mercurial to print a newline at the end of each template item. If you omit this newline, Mercurial will run each piece of output together. See section 11.5 for more details of escape sequences.

A template that prints a fixed string of text all the time isn't very useful; let's try something a bit more complex.

```

1 $ hg log --template 'i saw a changeset: {desc}\n'
2 i saw a changeset: Added tag v0.1 for changeset 558b9b187012
3 i saw a changeset: Added tag mytag for changeset 66687f9c6bca
4 i saw a changeset: added line to end of <<hello>> file.
5
6 in addition, added a file with the helpful name (at least i hope that some might consider it so) of goodbye
7 i saw a changeset: added hello

```

As you can see, the string “{desc}” in the template has been replaced in the output with the description of each changeset. Every time Mercurial finds text enclosed in curly braces (“{” and “}”), it will try to replace the braces and text with the expansion of whatever is inside. To print a literal curly brace, you must escape it, as described in section 11.5.

11.4 Common template keywords

You can start writing simple templates immediately using the keywords below.

author String. The unmodified author of the changeset.

branches String. The name of the branch on which the changeset was committed. Will be empty if the branch name was default.

date Date information. The date when the changeset was committed. This is *not* human-readable; you must pass it through a filter that will render it appropriately. See section 11.6 for more information on filters. The date is expressed as a pair of numbers. The first number is a Unix UTC timestamp (seconds since January 1, 1970); the second is the offset of the committer's timezone from UTC, in seconds.

desc String. The text of the changeset description.

files List of strings. All files modified, added, or removed by this changeset.

```

1 $ hg log -r1 --template 'author: {author}\n'
2 author: Bryan O'Sullivan <bos@serpentine.com>
3 $ hg log -r1 --template 'desc:\n{desc}\n'
4 desc:
5 added line to end of <<hello>> file.
6
7 in addition, added a file with the helpful name (at least i hope that some might consider it so) of goodbye
8 $ hg log -r1 --template 'files: {files}\n'
9 files: goodbye hello
10 $ hg log -r1 --template 'file_adds: {file_adds}\n'
11 file_adds: goodbye
12 $ hg log -r1 --template 'file_dels: {file_dels}\n'
13 file_dels:
14 $ hg log -r1 --template 'node: {node}\n'
15 node: 66687f9c6bcab073158d9b32b907ebc118d583a3
16 $ hg log -r1 --template 'parents: {parents}\n'
17 parents:
18 $ hg log -r1 --template 'rev: {rev}\n'
19 rev: 1
20 $ hg log -r1 --template 'tags: {tags}\n'
21 tags: mytag

```

Figure 11.1: Template keywords in use

`file_adds` List of strings. Files added by this changeset.

`file_dels` List of strings. Files removed by this changeset.

`node` String. The changeset identification hash, as a 40-character hexadecimal string.

`parents` List of strings. The parents of the changeset.

`rev` Integer. The repository-local changeset revision number.

`tags` List of strings. Any tags associated with the changeset.

A few simple experiments will show us what to expect when we use these keywords; you can see the results in figure 11.1.

As we noted above, the `date` keyword does not produce human-readable output, so we must treat it specially. This involves using a *filter*, about which more in section 11.6.

```

1 $ hg log -r1 --template 'date: {date}\n'
2 date: 1197316480.00
3 $ hg log -r1 --template 'date: {date|isodate}\n'
4 date: 2007-12-10 19:54 +0000

```

11.5 Escape sequences

Mercurial's templating engine recognises the most commonly used escape sequences in strings. When it sees a backslash ("`\`") character, it looks at the following character and substitutes the two characters with a single replacement, as described below.

`\\` Backslash, “\”, ASCII 134.
`\n` Newline, ASCII 12.
`\r` Carriage return, ASCII 15.
`\t` Tab, ASCII 11.
`\v` Vertical tab, ASCII 13.
`\{` Open curly brace, “{”, ASCII 173.
`\}` Close curly brace, “}”, ASCII 175.

As indicated above, if you want the expansion of a template to contain a literal “\”, “{”, or “}” character, you must escape it.

11.6 Filtering keywords to change their results

Some of the results of template expansion are not immediately easy to use. Mercurial lets you specify an optional chain of *filters* to modify the result of expanding a keyword. You have already seen a common filter, `isodate`, in action above, to make a date readable.

Below is a list of the most commonly used filters that Mercurial supports. While some filters can be applied to any text, others can only be used in specific circumstances. The name of each filter is followed first by an indication of where it can be used, then a description of its effect.

`addbreaks` Any text. Add an XHTML “`<br/>`” tag before the end of every line except the last. For example, “`foo\nbar`” becomes “`foo<br/>\nbar`”.

`age` date keyword. Render the age of the date, relative to the current time. Yields a string like “10 minutes”.

`basename` Any text, but most useful for the `files` keyword and its relatives. Treat the text as a path, and return the basename. For example, “`foo/bar/baz`” becomes “`baz`”.

`date` date keyword. Render a date in a similar format to the Unix `date` command, but with timezone included. Yields a string like “`Mon Sep 04 15:13:13 2006 -0700`”.

`domain` Any text, but most useful for the `author` keyword. Finds the first string that looks like an email address, and extract just the domain component. For example, “`Bryan O’Sullivan <bos@serpentine.com>`” becomes “`serpentine.com`”.

`email` Any text, but most useful for the `author` keyword. Extract the first string that looks like an email address. For example, “`Bryan O’Sullivan <bos@serpentine.com>`” becomes “`bos@serpentine.com`”.

`escape` Any text. Replace the special XML/XHTML characters “`&`”, “`<`” and “`>`” with XML entities.

`fill68` Any text. Wrap the text to fit in 68 columns. This is useful before you pass text through the `tabindent` filter, and still want it to fit in an 80-column fixed-font window.

`fill76` Any text. Wrap the text to fit in 76 columns.

`firstline` Any text. Yield the first line of text, without any trailing newlines.

`hgdate` date keyword. Render the date as a pair of readable numbers. Yields a string like “`1157407993 25200`”.

`isodate` date keyword. Render the date as a text string in ISO 8601 format. Yields a string like “`2006-09-04 15:13:13 -0700`”.

`obfuscate` Any text, but most useful for the `author` keyword. Yield the input text rendered as a sequence of XML entities. This helps to defeat some particularly stupid screen-scraping email harvesting spambots.

`person` Any text, but most useful for the `author` keyword. Yield the text before an email address. For example, “Bryan O’Sullivan <bos@serpentine.com>” becomes “Bryan O’Sullivan”.

`rfc822date` `date` keyword. Render a date using the same format used in email headers. Yields a string like “Mon, 04 Sep 2006 15:13:13 -0

`short` Changeset hash. Yield the short form of a changeset hash, i.e. a 12-byte hexadecimal string.

`shortdate` `date` keyword. Render the year, month, and day of the date. Yields a string like “2006-09-04”.

`strip` Any text. Strip all leading and trailing whitespace from the string.

`tabindent` Any text. Yield the text, with every line except the first starting with a tab character.

`urlescape` Any text. Escape all characters that are considered “special” by URL parsers. For example, `foo bar` becomes `foo%20bar`.

`user` Any text, but most useful for the `author` keyword. Return the “user” portion of an email address. For example, “Bryan O’Sullivan <bos@serpentine.com>” becomes “bos”.

Note: If you try to apply a filter to a piece of data that it cannot process, Mercurial will fail and print a Python exception. For example, trying to run the output of the `desc` keyword into the `isodate` filter is not a good idea.

11.6.1 Combining filters

It is easy to combine filters to yield output in the form you would like. The following chain of filters tidies up a description, then makes sure that it fits cleanly into 68 columns, then indents it by a further 8 characters (at least on Unix-like systems, where a tab is conventionally 8 characters wide).

```
1 $ hg log -r1 --template 'description:\n\t{desc|strip|fill68|tabindent}\n'
2 description:
3     added line to end of <<hello>> file.
4
5     in addition, added a file with the helpful name (at least i hope
6     that some might consider it so) of goodbye.
```

Note the use of “\t” (a tab character) in the template to force the first line to be indented; this is necessary since `tabindent` indents all lines *except* the first.

Keep in mind that the order of filters in a chain is significant. The first filter is applied to the result of the keyword; the second to the result of the first filter; and so on. For example, using `fill68|tabindent` gives very different results from `tabindent|fill68`.

11.7 From templates to styles

A command line template provides a quick and simple way to format some output. Templates can become verbose, though, and it’s useful to be able to give a template a name. A style file is a template with a name, stored in a file.

More than that, using a style file unlocks the power of Mercurial’s templating engine in ways that are not possible using the command line `--template` option.

```

1 $ hg log -r1 --template '{author}\n'
2 Bryan O'Sullivan <bos@serpentine.com>
3 $ hg log -r1 --template '{author|domain}\n'
4 serpentine.com
5 $ hg log -r1 --template '{author|email}\n'
6 bos@serpentine.com
7 $ hg log -r1 --template '{author|obfuscate}\n' | cut -c-76
8 &#66;&#114;&#121;&#97;&#110;&#32;&#79;&#39;&#83;&#117;&#108;&#108;&#105;&#11
9 $ hg log -r1 --template '{author|person}\n'
10 Bryan O'Sullivan
11 $ hg log -r1 --template '{author|user}\n'
12 bos
13 $ hg log -r1 --template 'looks almost right, but actually garbage: {date}\n'
14 looks almost right, but actually garbage: 1197316480.00
15 $ hg log -r1 --template '{date|age}\n'
16 1 second
17 $ hg log -r1 --template '{date|date}\n'
18 Mon Dec 10 19:54:40 2007 +0000
19 $ hg log -r1 --template '{date|hgdate}\n'
20 1197316480 0
21 $ hg log -r1 --template '{date|isodate}\n'
22 2007-12-10 19:54 +0000
23 $ hg log -r1 --template '{date|rfc822date}\n'
24 Mon, 10 Dec 2007 19:54:40 +0000
25 $ hg log -r1 --template '{date|shortdate}\n'
26 2007-12-10
27 $ hg log -r1 --template '{desc}\n' | cut -c-76
28 added line to end of <<hello>> file.
29
30 in addition, added a file with the helpful name (at least i hope that some m
31 $ hg log -r1 --template '{desc|addbreaks}\n' | cut -c-76
32 added line to end of <<hello>> file.<br/>
33 <br/>
34 in addition, added a file with the helpful name (at least i hope that some m
35 $ hg log -r1 --template '{desc|escape}\n' | cut -c-76
36 added line to end of &lt;&lt;hello>> file.
37
38 in addition, added a file with the helpful name (at least i hope that some m
39 $ hg log -r1 --template '{desc|fill68}\n'
40 added line to end of <<hello>> file.
41
42 in addition, added a file with the helpful name (at least i hope
43 that some might consider it so) of goodbye.
44 $ hg log -r1 --template '{desc|fill76}\n'
45 added line to end of <<hello>> file.
46
47 in addition, added a file with the helpful name (at least i hope that some
48 might consider it so) of goodbye.
49 $ hg log -r1 --template '{desc|firstline}\n'
50 added line to end of <<hello>> file.
51 $ hg log -r1 --template '{desc|strip}\n' | cut -c-76
52 added line to end of <<hello>> file.
53
54 in addition, added a file with the helpful name (at least i hope that some m
55 $ hg log -r1 --template '{desc|tabindent}\n' | expand | cut -c-76
56 added line to end of <<hello>> file.
57
58 in addition, added a file with the helpful name (at least i hope tha
59 $ hg log -r1 --template '{node}\n'
60 66687f9c6bcab073158d9b32b907ebc118d583a3

```

11.7.1 The simplest of style files

Our simple style file contains just one line:

```
1 $ echo 'changeset = "rev: {rev}\n" ' > rev
2 $ hg log -l1 --style ./rev
3 rev: 3
```

This tells Mercurial, “if you’re printing a changeset, use the text on the right as the template”.

11.7.2 Style file syntax

The syntax rules for a style file are simple.

- The file is processed one line at a time.
- Leading and trailing white space are ignored.
- Empty lines are skipped.
- If a line starts with either of the characters “#” or “;”, the entire line is treated as a comment, and skipped as if empty.
- A line starts with a keyword. This must start with an alphabetic character or underscore, and can subsequently contain any alphanumeric character or underscore. (In regexp notation, a keyword must match `[A-Za-z_][A-Za-z0-9_]*`.)
- The next element must be an “=” character, which can be preceded or followed by an arbitrary amount of white space.
- If the rest of the line starts and ends with matching quote characters (either single or double quote), it is treated as a template body.
- If the rest of the line *does not* start with a quote character, it is treated as the name of a file; the contents of this file will be read and used as a template body.

11.8 Style files by example

To illustrate how to write a style file, we will construct a few by example. Rather than provide a complete style file and walk through it, we’ll mirror the usual process of developing a style file by starting with something very simple, and walking through a series of successively more complete examples.

11.8.1 Identifying mistakes in style files

If Mercurial encounters a problem in a style file you are working on, it prints a terse error message that, once you figure out what it means, is actually quite useful.

```
1 $ cat broken.style
2 changeset =
```

Notice that `broken.style` attempts to define a `changeset` keyword, but forgets to give any content for it. When instructed to use this style file, Mercurial promptly complains.

```
1 $ hg log -r1 --style broken.style
2 abort: broken.style:1: parse error
```

This error message looks intimidating, but it is not too hard to follow.

- The first component is simply Mercurial’s way of saying “I am giving up”.

```
1 abort: broken.style:1: parse error
```

- Next comes the name of the style file that contains the error.

```
1 abort: broken.style:1: parse error
```

- Following the file name is the line number where the error was encountered.

```
1 abort: broken.style:1: parse error
```

- Finally, a description of what went wrong.

```
1 abort: broken.style:1: parse error
```

The description of the problem is not always clear (as in this case), but even when it is cryptic, it is almost always trivial to visually inspect the offending line in the style file and see what is wrong.

11.8.2 Uniquely identifying a repository

If you would like to be able to identify a Mercurial repository “fairly uniquely” using a short string as an identifier, you can use the first revision in the repository.

```
1 $ hg log -r0 --template '{node}'
2 1eb106e681ac8e686b89705b5f7548febae34ded
```

This is not guaranteed to be unique, but it is nevertheless useful in many cases.

- It will not work in a completely empty repository, because such a repository does not have a revision zero.
- Neither will it work in the (extremely rare) case where a repository is a merge of two or more formerly independent repositories, and you still have those repositories around.

Here are some uses to which you could put this identifier:

- As a key into a table for a database that manages repositories on a server.
- As half of a $\{repository\ ID, revision\ ID\}$ tuple. Save this information away when you run an automated build or other activity, so that you can “replay” the build later if necessary.

11.8.3 Mimicking Subversion’s output

Let’s try to emulate the default output format used by another revision control tool, Subversion.

```
1 $ svn log -r9653
2 -----
3 r9653 | sean.hefty | 2006-09-27 14:39:55 -0700 (Wed, 27 Sep 2006) | 5 lines
4
5 On reporting a route error, also include the status for the error,
6 rather than indicating a status of 0 when an error has occurred.
7
```

```

8 Signed-off-by: Sean Hefty <sean.hefty@intel.com>
9
10 -----

```

Since Subversion’s output style is fairly simple, it is easy to copy-and-paste a hunk of its output into a file, and replace the text produced above by Subversion with the template values we’d like to see expanded.

```

1 $ cat svn.template
2 r{rev} | {author|user} | {date|isodate} ({date|rfc822date})
3
4 {desc|strip|fill176}
5
6 -----

```

There are a few small ways in which this template deviates from the output produced by Subversion.

- Subversion prints a “readable” date (the “Wed, 27 Sep 2006” in the example output above) in parentheses. Mercurial’s templating engine does not provide a way to display a date in this format without also printing the time and time zone.
- We emulate Subversion’s printing of “separator” lines full of “-” characters by ending the template with such a line. We use the templating engine’s `header` keyword to print a separator line as the first line of output (see below), thus achieving similar output to Subversion.
- Subversion’s output includes a count in the header of the number of lines in the commit message. We cannot replicate this in Mercurial; the templating engine does not currently provide a filter that counts the number of items it is passed.

It took me no more than a minute or two of work to replace literal text from an example of Subversion’s output with some keywords and filters to give the template above. The style file simply refers to the template.

```

1 $ cat svn.style
2 header = '-----\n\n'
3 changeset = svn.template

```

We could have included the text of the template file directly in the style file by enclosing it in quotes and replacing the newlines with “

n” sequences, but it would have made the style file too difficult to read. Readability is a good guide when you’re trying to decide whether some text belongs in a style file, or in a template file that the style file points to. If the style file will look too big or cluttered if you insert a literal piece of text, drop it into a template instead.

Chapter 12

Managing change with Mercurial Queues

12.1 The patch management problem

Here is a common scenario: you need to install a software package from source, but you find a bug that you must fix in the source before you can start using the package. You make your changes, forget about the package for a while, and a few months later you need to upgrade to a newer version of the package. If the newer version of the package still has the bug, you must extract your fix from the older source tree and apply it against the newer version. This is a tedious task, and it's easy to make mistakes.

This is a simple case of the “patch management” problem. You have an “upstream” source tree that you can't change; you need to make some local changes on top of the upstream tree; and you'd like to be able to keep those changes separate, so that you can apply them to newer versions of the upstream source.

The patch management problem arises in many situations. Probably the most visible is that a user of an open source software project will contribute a bug fix or new feature to the project's maintainers in the form of a patch.

Distributors of operating systems that include open source software often need to make changes to the packages they distribute so that they will build properly in their environments.

When you have few changes to maintain, it is easy to manage a single patch using the standard `diff` and `patch` programs (see section 12.4 for a discussion of these tools). Once the number of changes grows, it starts to make sense to maintain patches as discrete “chunks of work,” so that for example a single patch will contain only one bug fix (the patch might modify several files, but it's doing “only one thing”), and you may have a number of such patches for different bugs you need fixed and local changes you require. In this situation, if you submit a bug fix patch to the upstream maintainers of a package and they include your fix in a subsequent release, you can simply drop that single patch when you're updating to the newer release.

Maintaining a single patch against an upstream tree is a little tedious and error-prone, but not difficult. However, the complexity of the problem grows rapidly as the number of patches you have to maintain increases. With more than a tiny number of patches in hand, understanding which ones you have applied and maintaining them moves from messy to overwhelming.

Fortunately, Mercurial includes a powerful extension, Mercurial Queues (or simply “MQ”), that massively simplifies the patch management problem.

12.2 The prehistory of Mercurial Queues

During the late 1990s, several Linux kernel developers started to maintain “patch series” that modified the behaviour of the Linux kernel. Some of these series were focused on stability, some on feature coverage, and others were more speculative.

The sizes of these patch series grew rapidly. In 2002, Andrew Morton published some shell scripts he had been using to automate the task of managing his patch queues. Andrew was successfully using these scripts to manage hundreds (sometimes thousands) of patches on top of the Linux kernel.

12.2.1 A patchwork quilt

In early 2003, Andreas Gruenbacher and Martin Quinson borrowed the approach of Andrew’s scripts and published a tool called “patchwork quilt” [AG], or simply “quilt” (see [Gru05] for a paper describing it). Because quilt substantially automated patch management, it rapidly gained a large following among open source software developers.

Quilt manages a *stack of patches* on top of a directory tree. To begin, you tell quilt to manage a directory tree, and tell it which files you want to manage; it stores away the names and contents of those files. To fix a bug, you create a new patch (using a single command), edit the files you need to fix, then “refresh” the patch.

The refresh step causes quilt to scan the directory tree; it updates the patch with all of the changes you have made. You can create another patch on top of the first, which will track the changes required to modify the tree from “tree with one patch applied” to “tree with two patches applied”.

You can *change* which patches are applied to the tree. If you “pop” a patch, the changes made by that patch will vanish from the directory tree. Quilt remembers which patches you have popped, though, so you can “push” a popped patch again, and the directory tree will be restored to contain the modifications in the patch. Most importantly, you can run the “refresh” command at any time, and the topmost applied patch will be updated. This means that you can, at any time, change both which patches are applied and what modifications those patches make.

Quilt knows nothing about revision control tools, so it works equally well on top of an unpacked tarball or a Subversion repository.

12.2.2 From patchwork quilt to Mercurial Queues

In mid-2005, Chris Mason took the features of quilt and wrote an extension that he called Mercurial Queues, which added quilt-like behaviour to Mercurial.

The key difference between quilt and MQ is that quilt knows nothing about revision control systems, while MQ is *integrated* into Mercurial. Each patch that you push is represented as a Mercurial changeset. Pop a patch, and the changeset goes away.

Because quilt does not care about revision control tools, it is still a tremendously useful piece of software to know about for situations where you cannot use Mercurial and MQ.

12.3 The huge advantage of MQ

I cannot overstate the value that MQ offers through the unification of patches and revision control.

A major reason that patches have persisted in the free software and open source world—in spite of the availability of increasingly capable revision control tools over the years—is the *agility* they offer.

Traditional revision control tools make a permanent, irreversible record of everything that you do. While this has great value, it’s also somewhat stifling. If you want to perform a wild-eyed experiment, you have to be careful in how you go about it, or you risk leaving unneeded—or worse, misleading or destabilising—traces of your missteps and errors in the permanent revision record.

By contrast, MQ’s marriage of distributed revision control with patches makes it much easier to isolate your work. Your patches live on top of normal revision history, and you can make them disappear or reappear at will. If you don’t like a patch, you can drop it. If a patch isn’t quite as you want it to be, simply fix it—as many times as you need to, until you have refined it into the form you desire.

As an example, the integration of patches with revision control makes understanding patches and debugging their effects—and their interplay with the code they’re based on—*enormously* easier. Since every applied patch has an associated changeset, you can use “hg log filename” to see which changesets and patches affected a file. You can use the bisect extension to binary-search through all changesets and applied patches to see where a bug got introduced or fixed. You can use the “hg annotate” command to see which changeset or patch modified a particular line of a source file. And so on.

12.4 Understanding patches

Because MQ doesn't hide its patch-oriented nature, it is helpful to understand what patches are, and a little about the tools that work with them.

The traditional Unix `diff` command compares two files, and prints a list of differences between them. The `patch` command understands these differences as *modifications* to make to a file. Take a look at figure 12.1 for a simple example of these commands in action.

```
1 $ echo 'this is my first line' > oldfile
2 $ echo 'my first line is here' > newfile
3 $ diff -u oldfile newfile > tiny.patch
4 $ cat tiny.patch
5 --- oldfile      2007-12-10 19:54:33.000000000 +0000
6 +++ newfile      2007-12-10 19:54:33.000000000 +0000
7 @@ -1 +1 @@
8 -this is my first line
9 +my first line is here
10 $ patch < tiny.patch
11 patching file oldfile
12 $ cat oldfile
13 my first line is here
```

Figure 12.1: Simple uses of the `diff` and `patch` commands

The type of file that `diff` generates (and `patch` takes as input) is called a “patch” or a “diff”; there is no difference between a patch and a diff. (We'll use the term “patch”, since it's more commonly used.)

A patch file can start with arbitrary text; the `patch` command ignores this text, but MQ uses it as the commit message when creating changesets. To find the beginning of the patch content, `patch` searches for the first line that starts with the string “`diff -`”.

MQ works with *unified* diffs (`patch` can accept several other diff formats, but MQ doesn't). A unified diff contains two kinds of header. The *file header* describes the file being modified; it contains the name of the file to modify. When `patch` sees a new file header, it looks for a file with that name to start modifying.

After the file header comes a series of *hunks*. Each hunk starts with a header; this identifies the range of line numbers within the file that the hunk should modify. Following the header, a hunk starts and ends with a few (usually three) lines of text from the unmodified file; these are called the *context* for the hunk. If there's only a small amount of context between successive hunks, `diff` doesn't print a new hunk header; it just runs the hunks together, with a few lines of context between modifications.

Each line of context begins with a space character. Within the hunk, a line that begins with “-” means “remove this line,” while a line that begins with “+” means “insert this line.” For example, a line that is modified is represented by one deletion and one insertion.

We will return to some of the more subtle aspects of patches later (in section 12.6), but you should have enough information now to use MQ.

12.5 Getting started with Mercurial Queues

Because MQ is implemented as an extension, you must explicitly enable before you can use it. (You don't need to download anything; MQ ships with the standard Mercurial distribution.) To enable MQ, edit your `~/.hgrc` file, and add the lines in figure 12.5.

Once the extension is enabled, it will make a number of new commands available. To verify that the extension is working, you can use “`hg help`” to see if the “`hg qinit`” command is now available; see the example in figure 12.3.

```
1 [extensions]
2 hgext.mq =
```

Figure 12.2: Contents to add to `~/.hgrc` to enable the MQ extension

```
1 $ hg help qinit
2 hg qinit [-c]
3
4 init a new queue repository
5
6     The queue repository is unversioned by default. If -c is
7     specified, qinit will create a separate nested repository
8     for patches (qinit -c may also be run later to convert
9     an unversioned patch repository into a versioned one).
10    You can use qcommit to commit changes to this queue repository.
11
12 options:
13
14 -c --create-repo  create queue repository
15
16 use "hg -v help qinit" to show global options
```

Figure 12.3: How to verify that MQ is enabled

You can use MQ with *any* Mercurial repository, and its commands only operate within that repository. To get started, simply prepare the repository using the “`hg qinit`” command (see figure 12.4). This command creates an empty directory called `.hg/patches`, where MQ will keep its metadata. As with many Mercurial commands, the “`hg qinit`” command prints nothing if it succeeds.

```
1 $ hg init mq-sandbox
2 $ cd mq-sandbox
3 $ echo 'line 1' > file1
4 $ echo 'another line 1' > file2
5 $ hg add file1 file2
6 $ hg commit -m'first change'
7 $ hg qinit
```

Figure 12.4: Preparing a repository for use with MQ

12.5.1 Creating a new patch

To begin work on a new patch, use the “`hg qnew`” command. This command takes one argument, the name of the patch to create. MQ will use this as the name of an actual file in the `.hg/patches` directory, as you can see in figure 12.5.

Also newly present in the `.hg/patches` directory are two other files, `series` and `status`. The `series` file lists all of the patches that MQ knows about for this repository, with one patch per line. Mercurial uses the `status` file for internal book-keeping; it tracks all of the patches that MQ has *applied* in this repository.

```

1 $ hg tip
2 changeset: 0:d02ac12a4c2b
3 tag: tip
4 user: Bryan O'Sullivan <bos@serpentine.com>
5 date: Mon Dec 10 19:54:35 2007 +0000
6 summary: first change
7
8 $ hg qnew first.patch
9 $ hg tip
10 changeset: 1:3e0e88db723b
11 tag: qtip
12 tag: first.patch
13 tag: tip
14 tag: qbase
15 user: Bryan O'Sullivan <bos@serpentine.com>
16 date: Mon Dec 10 19:54:35 2007 +0000
17 summary: [mq]: first.patch
18
19 $ ls .hg/patches
20 first.patch series status

```

Figure 12.5: Creating a new patch

Note: You may sometimes want to edit the `series` file by hand; for example, to change the sequence in which some patches are applied. However, manually editing the `status` file is almost always a bad idea, as it's easy to corrupt MQ's idea of what is happening.

Once you have created your new patch, you can edit files in the working directory as you usually would. All of the normal Mercurial commands, such as “`hg diff`” and “`hg annotate`”, work exactly as they did before.

12.5.2 Refreshing a patch

When you reach a point where you want to save your work, use the “`hg qrefresh`” command (figure 12.5) to update the patch you are working on. This command folds the changes you have made in the working directory into your patch, and updates its corresponding changeset to contain those changes.

You can run “`hg qrefresh`” as often as you like, so it's a good way to “checkpoint” your work. Refresh your patch at an opportune time; try an experiment; and if the experiment doesn't work out, “`hg revert`” your modifications back to the last time you refreshed.

12.5.3 Stacking and tracking patches

Once you have finished working on a patch, or need to work on another, you can use the “`hg qnew`” command again to create a new patch. Mercurial will apply this patch on top of your existing patch. See figure 12.8 for an example. Notice that the patch contains the changes in our prior patch as part of its context (you can see this more clearly in the output of “`hg annotate`”).

So far, with the exception of “`hg qnew`” and “`hg qrefresh`”, we've been careful to only use regular Mercurial commands. However, MQ provides many commands that are easier to use when you are thinking about patches, as illustrated in figure 12.9:

```

1 $ echo 'line 2' >> file1
2 $ hg diff
3 diff -r 3e0e88db723b file1
4 --- a/file1      Mon Dec 10 19:54:35 2007 +0000
5 +++ b/file1      Mon Dec 10 19:54:36 2007 +0000
6 @@ -1,1 +1,2 @@ line 1
7   line 1
8 +line 2
9 $ hg qrefresh
10 $ hg diff
11 $ hg tip --style=compact --patch
12 1[qtip,first.patch,tip,qbase]   bbf804e0b836   2007-12-10 19:54 +0000   bos
13   [mq]: first.patch
14
15 diff -r d02ac12a4c2b -r bbf804e0b836 file1
16 --- a/file1      Mon Dec 10 19:54:35 2007 +0000
17 +++ b/file1      Mon Dec 10 19:54:36 2007 +0000
18 @@ -1,1 +1,2 @@ line 1
19   line 1
20 +line 2
21

```

Figure 12.6: Refreshing a patch

- The “hg qseries” command lists every patch that MQ knows about in this repository, from oldest to newest (most recently *created*).
- The “hg qapplied” command lists every patch that MQ has *applied* in this repository, again from oldest to newest (most recently applied).

12.5.4 Manipulating the patch stack

The previous discussion implied that there must be a difference between “known” and “applied” patches, and there is. MQ can manage a patch without it being applied in the repository.

An *applied* patch has a corresponding changeset in the repository, and the effects of the patch and changeset are visible in the working directory. You can undo the application of a patch using the “hg qpop” command. MQ still *knows about*, or manages, a popped patch, but the patch no longer has a corresponding changeset in the repository, and the working directory does not contain the changes made by the patch. Figure 12.10 illustrates the difference between applied and tracked patches.

You can reapply an unapplied, or popped, patch using the “hg qpush” command. This creates a new changeset to correspond to the patch, and the patch’s changes once again become present in the working directory. See figure 12.11 for examples of “hg qpop” and “hg qpush” in action. Notice that once we have popped a patch or two patches, the output of “hg qseries” remains the same, while that of “hg qapplied” has changed.

12.5.5 Pushing and popping many patches

While “hg qpush” and “hg qpop” each operate on a single patch at a time by default, you can push and pop many patches in one go. The `-a` option to “hg qpush” causes it to push all unapplied patches, while the `-a` option to “hg qpop” causes it to pop all applied patches. (For some more ways to push and pop many patches, see section 12.7 below.)

```

1 $ echo 'line 3' >> file1
2 $ hg status
3 M file1
4 $ hg qrefresh
5 $ hg tip --style=compact --patch
6 1[qtip,first.patch,tip,qbase] e6769d3ff988 2007-12-10 19:54 +0000 bos
7 [mq]: first.patch
8
9 diff -r d02ac12a4c2b -r e6769d3ff988 file1
10 --- a/file1      Mon Dec 10 19:54:35 2007 +0000
11 +++ b/file1      Mon Dec 10 19:54:36 2007 +0000
12 @@ -1,1 +1,3 @@ line 1
13     line 1
14 +line 2
15 +line 3
16

```

Figure 12.7: Refresh a patch many times to accumulate changes

12.5.6 Safety checks, and overriding them

Several MQ commands check the working directory before they do anything, and fail if they find any modifications. They do this to ensure that you won't lose any changes that you have made, but not yet incorporated into a patch. Figure 12.13 illustrates this; the “hg qnew” command will not create a new patch if there are outstanding changes, caused in this case by the “hg add” of file3.

Commands that check the working directory all take an “I know what I'm doing” option, which is always named `-f`. The exact meaning of `-f` depends on the command. For example, “hg qnew `-f`” will incorporate any outstanding changes into the new patch it creates, but “hg qpop `-f`” will revert modifications to any files affected by the patch that it is popping. Be sure to read the documentation for a command's `-f` option before you use it!

12.5.7 Working on several patches at once

The “hg qrefresh” command always refreshes the *topmost* applied patch. This means that you can suspend work on one patch (by refreshing it), pop or push to make a different patch the top, and work on *that* patch for a while.

Here's an example that illustrates how you can use this ability. Let's say you're developing a new feature as two patches. The first is a change to the core of your software, and the second—layered on top of the first—changes the user interface to use the code you just added to the core. If you notice a bug in the core while you're working on the UI patch, it's easy to fix the core. Simply “hg qrefresh” the UI patch to save your in-progress changes, and “hg qpop” down to the core patch. Fix the core bug, “hg qrefresh” the core patch, and “hg qpush” back to the UI patch to continue where you left off.

12.6 More about patches

MQ uses the GNU `patch` command to apply patches, so it's helpful to know a few more detailed aspects of how `patch` works, and about patches themselves.

12.6.1 The strip count

If you look at the file headers in a patch, you will notice that the pathnames usually have an extra component on the front that isn't present in the actual path name. This is a holdover from the way that people used to generate patches

```

1 $ hg qnew second.patch
2 $ hg log --style=compact --limit=2
3 2[qtip,second.patch,tip] 5e99adb53675 2007-12-10 19:54 +0000 bos
4 [mq]: second.patch
5
6 1[first.patch,qbase] e6769d3ff988 2007-12-10 19:54 +0000 bos
7 [mq]: first.patch
8
9 $ echo 'line 4' >> file1
10 $ hg qrefresh
11 $ hg tip --style=compact --patch
12 2[qtip,second.patch,tip] 0927fbdce91c 2007-12-10 19:54 +0000 bos
13 [mq]: second.patch
14
15 diff -r e6769d3ff988 -r 0927fbdce91c file1
16 --- a/file1      Mon Dec 10 19:54:36 2007 +0000
17 +++ b/file1      Mon Dec 10 19:54:36 2007 +0000
18 @@ -1,3 +1,4 @@ line 1
19     line 1
20     line 2
21     line 3
22 +line 4
23
24 $ hg annotate file1
25 0: line 1
26 1: line 2
27 1: line 3
28 2: line 4

```

Figure 12.8: Stacking a second patch on top of the first

(people still do this, but it’s somewhat rare with modern revision control tools).

Alice would unpack a tarball, edit her files, then decide that she wanted to create a patch. So she’d rename her working directory, unpack the tarball again (hence the need for the rename), and use the `-r` and `-N` options to `diff` to recursively generate a patch between the unmodified directory and the modified one. The result would be that the name of the unmodified directory would be at the front of the left-hand path in every file header, and the name of the modified directory would be at the front of the right-hand path.

Since someone receiving a patch from the Alices of the net would be unlikely to have unmodified and modified directories with exactly the same names, the `patch` command has a `-p` option that indicates the number of leading path name components to strip when trying to apply a patch. This number is called the *strip count*.

An option of “`-p1`” means “use a strip count of one”. If `patch` sees a file name `foo/bar/baz` in a file header, it will strip `foo` and try to patch a file named `bar/baz`. (Strictly speaking, the strip count refers to the number of *path separators* (and the components that go with them) to strip. A strip count of one will turn `foo/bar` into `bar`, but `/foo/bar` (notice the extra leading slash) into `foo/bar`.)

The “standard” strip count for patches is one; almost all patches contain one leading path name component that needs to be stripped. Mercurial’s “`hg diff`” command generates path names in this form, and the “`hg import`” command and MQ expect patches to have a strip count of one.

If you receive a patch from someone that you want to add to your patch queue, and the patch needs a strip count other than one, you cannot just “`hg qimport`” the patch, because “`hg qimport`” does not yet have a `-p` option (see [Mercurial bug no. 311](#)). Your best bet is to “`hg qnew`” a patch of your own, then use “`patch -pN`” to

```

1 $ hg qseries
2 first.patch
3 second.patch
4 $ hg qapplied
5 first.patch
6 second.patch

```

Figure 12.9: Understanding the patch stack with “hg qseries” and “hg qapplied”




Figure 12.10: Applied and unapplied patches in the MQ patch stack

apply their patch, followed by “hg addremove” to pick up any files added or removed by the patch, followed by “hg qrefresh”. This complexity may become unnecessary; see [Mercurial bug no. 311](#) for details.

12.6.2 Strategies for applying a patch

When `patch` applies a hunk, it tries a handful of successively less accurate strategies to try to make the hunk apply. This falling-back technique often makes it possible to take a patch that was generated against an old version of a file, and apply it against a newer version of that file.

First, `patch` tries an exact match, where the line numbers, the context, and the text to be modified must apply exactly. If it cannot make an exact match, it tries to find an exact match for the context, without honouring the line numbering information. If this succeeds, it prints a line of output saying that the hunk was applied, but at some *offset* from the original line number.

If a context-only match fails, `patch` removes the first and last lines of the context, and tries a *reduced* context-only match. If the hunk with reduced context succeeds, it prints a message saying that it applied the hunk with a *fuzz factor* (the number after the fuzz factor indicates how many lines of context `patch` had to trim before the patch applied).

When neither of these techniques works, `patch` prints a message saying that the hunk in question was rejected. It saves rejected hunks (also simply called “rejects”) to a file with the same name, and an added `.rej` extension. It also saves an unmodified copy of the file with a `.orig` extension; the copy of the file without any extensions will contain any changes made by hunks that *did* apply cleanly. If you have a patch that modifies `foo` with six hunks, and one of them fails to apply, you will have: an unmodified `foo.orig`, a `foo.rej` containing one hunk, and `foo`, containing the changes made by the five successful five hunks.

12.6.3 Some quirks of patch representation

There are a few useful things to know about how `patch` works with files.

```

1 $ hg qapplied
2 first.patch
3 second.patch
4 $ hg qpop
5 Now at: first.patch
6 $ hg qseries
7 first.patch
8 second.patch
9 $ hg qapplied
10 first.patch
11 $ cat file1
12 line 1
13 line 2
14 line 3

```

Figure 12.11: Modifying the stack of applied patches

```

1 $ hg qpush -a
2 applying second.patch
3 Now at: second.patch
4 $ cat file1
5 line 1
6 line 2
7 line 3
8 line 4

```

Figure 12.12: Pushing all unapplied patches

- This should already be obvious, but `patch` cannot handle binary files.
- Neither does it care about the executable bit; it creates new files as readable, but not executable.
- `patch` treats the removal of a file as a diff between the file to be removed and the empty file. So your idea of “I deleted this file” looks like “every line of this file was deleted” in a patch.
- It treats the addition of a file as a diff between the empty file and the file to be added. So in a patch, your idea of “I added this file” looks like “every line of this file was added”.
- It treats a renamed file as the removal of the old name, and the addition of the new name. This means that renamed files have a big footprint in patches. (Note also that Mercurial does not currently try to infer when files have been renamed or copied in a patch.)
- `patch` cannot represent empty files, so you cannot use a patch to represent the notion “I added this empty file to the tree”.

12.6.4 Beware the fuzz

While applying a hunk at an offset, or with a fuzz factor, will often be completely successful, these inexact techniques naturally leave open the possibility of corrupting the patched file. The most common cases typically involve applying a patch twice, or at an incorrect location in the file. If `patch` or “`hg qpush`” ever mentions an offset or fuzz factor, you should make sure that the modified files are correct afterwards.

```
1 $ echo 'file 3, line 1' >> file3
2 $ hg qnew add-file3.patch
3 $ hg qnew -f add-file3.patch
4 abort: patch "add-file3.patch" already exists
```

Figure 12.13: Forcibly creating a patch

It's often a good idea to refresh a patch that has applied with an offset or fuzz factor; refreshing the patch generates new context information that will make it apply cleanly. I say “often,” not “always,” because sometimes refreshing a patch will make it fail to apply against a different revision of the underlying files. In some cases, such as when you're maintaining a patch that must sit on top of multiple versions of a source tree, it's acceptable to have a patch apply with some fuzz, provided you've verified the results of the patching process in such cases.

12.6.5 Handling rejection

If “hg qpush” fails to apply a patch, it will print an error message and exit. If it has left `.rej` files behind, it is usually best to fix up the rejected hunks before you push more patches or do any further work.

If your patch *used to* apply cleanly, and no longer does because you've changed the underlying code that your patches are based on, Mercurial Queues can help; see section 12.8 for details.

Unfortunately, there aren't any great techniques for dealing with rejected hunks. Most often, you'll need to view the `.rej` file and edit the target file, applying the rejected hunks by hand.

If you're feeling adventurous, Neil Brown, a Linux kernel hacker, wrote a tool called `wiggle` [Bro], which is more vigorous than `patch` in its attempts to make a patch apply.

Another Linux kernel hacker, Chris Mason (the author of Mercurial Queues), wrote a similar tool called `mpatch` [Mas], which takes a simple approach to automating the application of hunks rejected by `patch`. The `mpatch` command can help with four common reasons that a hunk may be rejected:

- The context in the middle of a hunk has changed.
- A hunk is missing some context at the beginning or end.
- A large hunk might apply better—either entirely or in part—if it was broken up into smaller hunks.
- A hunk removes lines with slightly different content than those currently present in the file.

If you use `wiggle` or `mpatch`, you should be doubly careful to check your results when you're done. In fact, `mpatch` enforces this method of double-checking the tool's output, by automatically dropping you into a merge program when it has done its job, so that you can verify its work and finish off any remaining merges.

12.7 Getting the best performance out of MQ

MQ is very efficient at handling a large number of patches. I ran some performance experiments in mid-2006 for a talk that I gave at the 2006 EuroPython conference [O'S06]. I used as my data set the Linux 2.6.17-mm1 patch series, which consists of 1,738 patches. I applied these on top of a Linux kernel repository containing all 27,472 revisions between Linux 2.6.12-rc2 and Linux 2.6.17.

On my old, slow laptop, I was able to “hg qpush -a” all 1,738 patches in 3.5 minutes, and “hg qpop -a” them all in 30 seconds. (On a newer laptop, the time to push all patches dropped to two minutes.) I could “hg qrefresh” one of the biggest patches (which made 22,779 lines of changes to 287 files) in 6.6 seconds.

Clearly, MQ is well suited to working in large trees, but there are a few tricks you can use to get the best performance of it.

First of all, try to “batch” operations together. Every time you run “hg qpush” or “hg qpop”, these commands scan the working directory once to make sure you haven’t made some changes and then forgotten to run “hg qrefresh”. On a small tree, the time that this scan takes is unnoticeable. However, on a medium-sized tree (containing tens of thousands of files), it can take a second or more.

The “hg qpush” and “hg qpop” commands allow you to push and pop multiple patches at a time. You can identify the “destination patch” that you want to end up at. When you “hg qpush” with a destination specified, it will push patches until that patch is at the top of the applied stack. When you “hg qpop” to a destination, MQ will pop patches until the destination patch is at the top.

You can identify a destination patch using either the name of the patch, or by number. If you use numeric addressing, patches are counted from zero; this means that the first patch is zero, the second is one, and so on.

12.8 Updating your patches when the underlying code changes

It’s common to have a stack of patches on top of an underlying repository that you don’t modify directly. If you’re working on changes to third-party code, or on a feature that is taking longer to develop than the rate of change of the code beneath, you will often need to sync up with the underlying code, and fix up any hunks in your patches that no longer apply. This is called *rebasing* your patch series.

The simplest way to do this is to “hg qpop -a” your patches, then “hg pull” changes into the underlying repository, and finally “hg qpush -a” your patches again. MQ will stop pushing any time it runs across a patch that fails to apply during conflicts, allowing you to fix your conflicts, “hg qrefresh” the affected patch, and continue pushing until you have fixed your entire stack.

This approach is easy to use and works well if you don’t expect changes to the underlying code to affect how well your patches apply. If your patch stack touches code that is modified frequently or invasively in the underlying repository, however, fixing up rejected hunks by hand quickly becomes tiresome.

It’s possible to partially automate the rebasing process. If your patches apply cleanly against some revision of the underlying repo, MQ can use this information to help you to resolve conflicts between your patches and a different revision.

The process is a little involved.

1. To begin, “hg qpush -a” all of your patches on top of the revision where you know that they apply cleanly.
2. Save a backup copy of your patch directory using “hg qsave -e -c”. This prints the name of the directory that it has saved the patches in. It will save the patches to a directory called `.hg/patches.N`, where *N* is a small integer. It also commits a “save changeset” on top of your applied patches; this is for internal book-keeping, and records the states of the `series` and `status` files.
3. Use “hg pull” to bring new changes into the underlying repository. (Don’t run “hg pull -u”; see below for why.)
4. Update to the new tip revision, using “hg update -C” to override the patches you have pushed.
5. Merge all patches using “hg qpush -m -a”. The `-m` option to “hg qpush” tells MQ to perform a three-way merge if the patch fails to apply.

During the “hg qpush -m”, each patch in the `series` file is applied normally. If a patch applies with fuzz or rejects, MQ looks at the queue you “hg qsave”d, and performs a three-way merge with the corresponding changeset. This merge uses Mercurial’s normal merge machinery, so it may pop up a GUI merge tool to help you to resolve problems.

When you finish resolving the effects of a patch, MQ refreshes your patch based on the result of the merge.

At the end of this process, your repository will have one extra head from the old patch queue, and a copy of the old patch queue will be in `.hg/patches.N`. You can remove the extra head using “hg qpop -a -n `patches.N`” or “hg strip”. You can delete `.hg/patches.N` once you are sure that you no longer need it as a backup.

12.9 Identifying patches

MQ commands that work with patches let you refer to a patch either by using its name or by a number. By name is obvious enough; pass the name `foo.patch` to “`hg qpush`”, for example, and it will push patches until `foo.patch` is applied.

As a shortcut, you can refer to a patch using both a name and a numeric offset; `foo.patch-2` means “two patches before `foo.patch`”, while `bar.patch+4` means “four patches after `bar.patch`”.

Referring to a patch by index isn’t much different. The first patch printed in the output of “`hg qseries`” is patch zero (yes, it’s one of those start-at-zero counting systems); the second is patch one; and so on

MQ also makes it easy to work with patches when you are using normal Mercurial commands. Every command that accepts a changeset ID will also accept the name of an applied patch. MQ augments the tags normally in the repository with an eponymous one for each applied patch. In addition, the special tags `qbase` and `qtip` identify the “bottom-most” and topmost applied patches, respectively.

These additions to Mercurial’s normal tagging capabilities make dealing with patches even more of a breeze.

- Want to patchbomb a mailing list with your latest series of changes?

```
1 hg email qbase:qtip
```

(Don’t know what “patchbombing” is? See section [14.4](#).)

- Need to see all of the patches since `foo.patch` that have touched files in a subdirectory of your tree?

```
1 hg log -r foo.patch:qtip subdir
```

Because MQ makes the names of patches available to the rest of Mercurial through its normal internal tag machinery, you don’t need to type in the entire name of a patch when you want to identify it by name.

Another nice consequence of representing patch names as tags is that when you run the “`hg log`” command, it will display a patch’s name as a tag, simply as part of its normal output. This makes it easy to visually distinguish applied patches from underlying “normal” revisions. Figure [12.14](#) shows a few normal Mercurial commands in use with applied patches.

12.10 Useful things to know about

There are a number of aspects of MQ usage that don’t fit tidily into sections of their own, but that are good to know. Here they are, in one place.

- Normally, when you “`hg qpop`” a patch and “`hg qpush`” it again, the changeset that represents the patch after the pop/push will have a *different identity* than the changeset that represented the hash beforehand. See section [B.1.13](#) for information as to why this is.
- It’s not a good idea to “`hg merge`” changes from another branch with a patch changeset, at least if you want to maintain the “patchiness” of that changeset and changesets below it on the patch stack. If you try to do this, it will appear to succeed, but MQ will become confused.

12.11 Managing patches in a repository

Because MQ’s `.hg/patches` directory resides outside a Mercurial repository’s working directory, the “underlying” Mercurial repository knows nothing about the management or presence of patches.

This presents the interesting possibility of managing the contents of the patch directory as a Mercurial repository in its own right. This can be a useful way to work. For example, you can work on a patch for a while, “`hg qrefresh`” it, then “`hg commit`” the current state of the patch. This lets you “roll back” to that version of the patch later on.

```

1 $ hg qapplied
2 first.patch
3 second.patch
4 $ hg log -r qbase:qtip
5 changeset: 1:6942ab25b40a
6 tag: first.patch
7 tag: qbase
8 user: Bryan O'Sullivan <bos@serpentine.com>
9 date: Mon Dec 10 19:54:33 2007 +0000
10 summary: [mq]: first.patch
11
12 changeset: 2:e20ef0f44138
13 tag: qtip
14 tag: second.patch
15 tag: tip
16 user: Bryan O'Sullivan <bos@serpentine.com>
17 date: Mon Dec 10 19:54:34 2007 +0000
18 summary: [mq]: second.patch
19
20 $ hg export second.patch
21 # HG changeset patch
22 # User Bryan O'Sullivan <bos@serpentine.com>
23 # Date 1197316474 0
24 # Node ID e20ef0f441384d3125463a3d3321a046f08b747c
25 # Parent 6942ab25b40a286d857368d0beedb287ed112e6f
26 [mq]: second.patch
27
28 diff -r 6942ab25b40a -r e20ef0f44138 other.c
29 --- /dev/null Thu Jan 01 00:00:00 1970 +0000
30 +++ b/other.c Mon Dec 10 19:54:34 2007 +0000
31 @@ -0,0 +1,1 @@
32 +double u;

```

Figure 12.14: Using MQ's tag features to work with patches

You can then share different versions of the same patch stack among multiple underlying repositories. I use this when I am developing a Linux kernel feature. I have a pristine copy of my kernel sources for each of several CPU architectures, and a cloned repository under each that contains the patches I am working on. When I want to test a change on a different architecture, I push my current patches to the patch repository associated with that kernel tree, pop and push all of my patches, and build and test that kernel.

Managing patches in a repository makes it possible for multiple developers to work on the same patch series without colliding with each other, all on top of an underlying source base that they may or may not control.

12.11.1 MQ support for patch repositories

MQ helps you to work with the `.hg/patches` directory as a repository; when you prepare a repository for working with patches using “`hg qinit`”, you can pass the `-c` option to create the `.hg/patches` directory as a Mercurial repository.

Note: If you forget to use the `-c` option, you can simply go into the `.hg/patches` directory at any time and run `“hg init”`. Don’t forget to add an entry for the status file to the `.hgignore` file, though (`“hg qinit -c”` does this for you automatically); you *really* don’t want to manage the status file.

As a convenience, if MQ notices that the `.hg/patches` directory is a repository, it will automatically `“hg add”` every patch that you create and import.

Finally, MQ provides a shortcut command, `“hg qcommit”`, that runs `“hg commit”` in the `.hg/patches` directory. This saves some cumbersome typing.

12.11.2 A few things to watch out for

MQ’s support for working with a repository full of patches is limited in a few small respects.

MQ cannot automatically detect changes that you make to the patch directory. If you `“hg pull”`, manually edit, or `“hg update”` changes to patches or the `series` file, you will have to `“hg qpop -a”` and then `“hg qpush -a”` in the underlying repository to see those changes show up there. If you forget to do this, you can confuse MQ’s idea of which patches are applied.

12.12 Third party tools for working with patches

Once you’ve been working with patches for a while, you’ll find yourself hungry for tools that will help you to understand and manipulate the patches you’re dealing with.

The `diffstat` command [Dic] generates a histogram of the modifications made to each file in a patch. It provides a good way to “get a sense of” a patch—which files it affects, and how much change it introduces to each file and as a whole. (I find that it’s a good idea to use `diffstat`’s `-p` option as a matter of course, as otherwise it will try to do clever things with prefixes of file names that inevitably confuse at least me.)

The `patchutils` package [Wau] is invaluable. It provides a set of small utilities that follow the “Unix philosophy;” each does one useful thing with a patch. The `patchutils` command I use most is `filterdiff`, which extracts subsets from a patch file. For example, given a patch that modifies hundreds of files across dozens of directories, a single invocation of `filterdiff` can generate a smaller patch that only touches files whose names match a particular glob pattern. See section 13.9.2 for another example.

12.13 Good ways to work with patches

Whether you are working on a patch series to submit to a free software or open source project, or a series that you intend to treat as a sequence of regular changesets when you’re done, you can use some simple techniques to keep your work well organised.

Give your patches descriptive names. A good name for a patch might be `rework-device-alloc.patch`, because it will immediately give you a hint what the purpose of the patch is. Long names shouldn’t be a problem; you won’t be typing the names often, but you *will* be running commands like `“hg qapplied”` and `“hg qtop”` over and over. Good naming becomes especially important when you have a number of patches to work with, or if you are juggling a number of different tasks and your patches only get a fraction of your attention.

Be aware of what patch you’re working on. Use the `“hg qtop”` command and skim over the text of your patches frequently—for example, using `“hg tip -p”`—to be sure of where you stand. I have several times worked on and `“hg qrefresh”`ed a patch other than the one I intended, and it’s often tricky to migrate changes into the right patch after making them in the wrong one.

For this reason, it is very much worth investing a little time to learn how to use some of the third-party tools I described in section 12.12, particularly `diffstat` and `filterdiff`. The former will give you a quick idea of what changes your patch is making, while the latter makes it easy to splice hunks selectively out of one patch and into another.

```

1 $ diffstat -p1 remove-redundant-null-checks.patch
2 drivers/char/agp/sgi-agp.c      |    5 +----
3 drivers/char/hvcs.c             |   11 +++++-----
4 drivers/message/fusion/mptfc.c  |    6 +-----
5 drivers/message/fusion/mptsas.c |    3 +--
6 drivers/net/fs_enet/fs_enet-mii.c |    3 +--
7 drivers/net/wireless/ipw2200.c  |   22 ++++++-----
8 drivers/scsi/libata-scsi.c      |    4 +---
9 drivers/video/au1100fb.c        |    3 +--
10 8 files changed, 19 insertions(+), 38 deletions(-)
11 $ filterdiff -i '*/video/*' remove-redundant-null-checks.patch
12 --- a/drivers/video/au1100fb.c~remove-redundant-null-checks-before-free-in-drivers
13 +++ a/drivers/video/au1100fb.c
14 @@ -743,8 +743,7 @@ void __exit au1100fb_cleanup(void)
15 {
16     driver_unregister(&au1100fb_driver);
17
18     if (drv_info.opt_mode)
19         kfree(drv_info.opt_mode);
20 +     kfree(drv_info.opt_mode);
21 }
22
23 module_init(au1100fb_init);

```

Figure 12.15: The diffstat, filterdiff, and lsdiff commands

12.14 MQ cookbook

12.14.1 Manage “trivial” patches

Because the overhead of dropping files into a new Mercurial repository is so low, it makes a lot of sense to manage patches this way even if you simply want to make a few changes to a source tarball that you downloaded.

Begin by downloading and unpacking the source tarball, and turning it into a Mercurial repository.

```

1 $ download netplug-1.2.5.tar.bz2
2 $ tar jxf netplug-1.2.5.tar.bz2
3 $ cd netplug-1.2.5
4 $ hg init
5 $ hg commit -q --addremove --message netplug-1.2.5
6 $ cd ..
7 $ hg clone netplug-1.2.5 netplug
8 18 files updated, 0 files merged, 0 files removed, 0 files unresolved

```

Continue by creating a patch stack and making your changes.

```

1 $ cd netplug
2 $ hg qinit
3 $ hg qnew -m 'fix build problem with gcc 4' build-fix.patch
4 $ perl -pi -e 's/int addr_len/socklen_t addr_len/' netlink.c
5 $ hg qrefresh
6 $ hg tip -p

```

```

7 changeset: 1:fdef4c3a0798
8 tag:      qtip
9 tag:      build-fix.patch
10 tag:      tip
11 tag:      qbase
12 user:     Bryan O'Sullivan <bos@serpentine.com>
13 date:     Mon Dec 10 19:54:34 2007 +0000
14 summary:  fix build problem with gcc 4
15
16 diff -r 625168521d89 -r fdef4c3a0798 netlink.c
17 --- a/netlink.c      Mon Dec 10 19:54:34 2007 +0000
18 +++ b/netlink.c      Mon Dec 10 19:54:34 2007 +0000
19 @@ -275,7 +275,7 @@ netlink_open(void)
20         exit(1);
21     }
22
23 -    int addr_len = sizeof(addr);
24 +    socklen_t addr_len = sizeof(addr);
25
26     if (getsockname(fd, (struct sockaddr *) &addr, &addr_len) == -1) {
27         do_log(LOG_ERR, "Could not get socket details: %m");
28     }

```

Let's say a few weeks or months pass, and your package author releases a new version. First, bring their changes into the repository.

```

1 $ hg qpop -a
2 Patch queue now empty
3 $ cd ..
4 $ download netplug-1.2.8.tar.bz2
5 $ hg clone netplug-1.2.5 netplug-1.2.8
6 18 files updated, 0 files merged, 0 files removed, 0 files unresolved
7 $ cd netplug-1.2.8
8 $ hg locate -0 | xargs -0 rm
9 $ cd ..
10 $ tar jxf netplug-1.2.8.tar.bz2
11 $ cd netplug-1.2.8
12 $ hg commit --addremove --message netplug-1.2.8

```

The pipeline starting with “hg locate” above deletes all files in the working directory, so that “hg commit”'s --addremove option can actually tell which files have really been removed in the newer version of the source.

Finally, you can apply your patches on top of the new tree.

```

1 $ cd ../netplug
2 $ hg pull ../netplug-1.2.8
3 pulling from ../netplug-1.2.8
4 searching for changes
5 adding changesets
6 adding manifests
7 adding file changes
8 added 1 changesets with 12 changes to 12 files
9 (run 'hg update' to get a working copy)

```

```

10 $ hg qpush -a
11 applying build-fix.patch
12 Now at: build-fix.patch

```

12.14.2 Combining entire patches

MQ provides a command, “hg qfold” that lets you combine entire patches. This “folds” the patches you name, in the order you name them, into the topmost applied patch, and concatenates their descriptions onto the end of its description. The patches that you fold must be unapplied before you fold them.

The order in which you fold patches matters. If your topmost applied patch is *foo*, and you “hg qfold” *bar* and *quux* into it, you will end up with a patch that has the same effect as if you applied first *foo*, then *bar*, followed by *quux*.

12.14.3 Merging part of one patch into another

Merging *part* of one patch into another is more difficult than combining entire patches.

If you want to move changes to entire files, you can use *filterdiff*’s *-i* and *-x* options to choose the modifications to snip out of one patch, concatenating its output onto the end of the patch you want to merge into. You usually won’t need to modify the patch you’ve merged the changes from. Instead, MQ will report some rejected hunks when you “hg qpush” it (from the hunks you moved into the other patch), and you can simply “hg qrefresh” the patch to drop the duplicate hunks.

If you have a patch that has multiple hunks modifying a file, and you only want to move a few of those hunks, the job becomes more messy, but you can still partly automate it. Use “*lsdiff -nvv*” to print some metadata about the patch.

```

1 $ lsdiff -nvv remove-redundant-null-checks.patch
2 22      File #1      a/drivers/char/agp/sgi-agp.c
3      24      Hunk #1      static int __devinit agp_sgi_init(void)
4 37      File #2      a/drivers/char/hvcs.c
5      39      Hunk #1      static struct tty_operations hvcs_ops =
6      53      Hunk #2      static int hvcs_alloc_index_list(int n)
7 69      File #3      a/drivers/message/fusion/mptfc.c
8      71      Hunk #1      mptfc_GetFcDevPage0(MPT_ADAPTER *ioc, in
9 85      File #4      a/drivers/message/fusion/mptsas.c
10     87      Hunk #1      mptsas_probe_hba_phys(MPT_ADAPTER *ioc)
11 98      File #5      a/drivers/net/fs_enet/fs_enet-mii.c
12    100      Hunk #1      static struct fs_enet_mii_bus *create_bu
13 111     File #6      a/drivers/net/wireless/ipw2200.c
14    113      Hunk #1      static struct ipw_fw_error *ipw_alloc_er
15    126      Hunk #2      static ssize_t clear_error(struct device
16    140      Hunk #3      static void ipw_irq_tasklet(struct ipw_p
17    150      Hunk #4      static void ipw_pci_remove(struct pci_de
18 164     File #7      a/drivers/scsi/libata-scsi.c
19    166      Hunk #1      int ata_cmd_ioctl(struct scsi_device *sc
20 178     File #8      a/drivers/video/au1100fb.c
21    180      Hunk #1      void __exit au1100fb_cleanup(void)

```

This command prints three different kinds of number:

- (in the first column) a *file number* to identify each file modified in the patch;
- (on the next line, indented) the line number within a modified file where a hunk starts; and

- (on the same line) a *hunk number* to identify that hunk.

You'll have to use some visual inspection, and reading of the patch, to identify the file and hunk numbers you'll want, but you can then pass them to `filterdiff`'s `--files` and `--hunks` options, to select exactly the file and hunk you want to extract.

Once you have this hunk, you can concatenate it onto the end of your destination patch and continue with the remainder of section [12.14.2](#).

12.15 Differences between quilt and MQ

If you are already familiar with quilt, MQ provides a similar command set. There are a few differences in the way that it works.

You will already have noticed that most quilt commands have MQ counterparts that simply begin with a “q”. The exceptions are quilt's `add` and `remove` commands, the counterparts for which are the normal Mercurial “`hg add`” and “`hg remove`” commands. There is no MQ equivalent of the quilt `edit` command.

Chapter 13

Advanced uses of Mercurial Queues

While it's easy to pick up straightforward uses of Mercurial Queues, use of a little discipline and some of MQ's less frequently used capabilities makes it possible to work in complicated development environments.

In this chapter, I will use as an example a technique I have used to manage the development of an Infiniband device driver for the Linux kernel. The driver in question is large (at least as drivers go), with 25,000 lines of code spread across 35 source files. It is maintained by a small team of developers.

While much of the material in this chapter is specific to Linux, the same principles apply to any code base for which you're not the primary owner, and upon which you need to do a lot of development.

13.1 The problem of many targets

The Linux kernel changes rapidly, and has never been internally stable; developers frequently make drastic changes between releases. This means that a version of the driver that works well with a particular released version of the kernel will not even *compile* correctly against, typically, any other version.

To maintain a driver, we have to keep a number of distinct versions of Linux in mind.

- One target is the main Linux kernel development tree. Maintenance of the code is in this case partly shared by other developers in the kernel community, who make “drive-by” modifications to the driver as they develop and refine kernel subsystems.
- We also maintain a number of “backports” to older versions of the Linux kernel, to support the needs of customers who are running older Linux distributions that do not incorporate our drivers. (To *backport* a piece of code is to modify it to work in an older version of its target environment than the version it was developed for.)
- Finally, we make software releases on a schedule that is necessarily not aligned with those used by Linux distributors and kernel developers, so that we can deliver new features to customers without forcing them to upgrade their entire kernels or distributions.

13.1.1 Tempting approaches that don't work well

There are two “standard” ways to maintain a piece of software that has to target many different environments.

The first is to maintain a number of branches, each intended for a single target. The trouble with this approach is that you must maintain iron discipline in the flow of changes between repositories. A new feature or bug fix must start life in a “pristine” repository, then percolate out to every backport repository. Backport changes are more limited in the branches they should propagate to; a backport change that is applied to a branch where it doesn't belong will probably stop the driver from compiling.

The second is to maintain a single source tree filled with conditional statements that turn chunks of code on or off depending on the intended target. Because these “*ifdefs*” are not allowed in the Linux kernel tree, a manual or

automatic process must be followed to strip them out and yield a clean tree. A code base maintained in this fashion rapidly becomes a rat's nest of conditional blocks that are difficult to understand and maintain.

Neither of these approaches is well suited to a situation where you don't "own" the canonical copy of a source tree. In the case of a Linux driver that is distributed with the standard kernel, Linus's tree contains the copy of the code that will be treated by the world as canonical. The upstream version of "my" driver can be modified by people I don't know, without me even finding out about it until after the changes show up in Linus's tree.

These approaches have the added weakness of making it difficult to generate well-formed patches to submit upstream.

In principle, Mercurial Queues seems like a good candidate to manage a development scenario such as the above. While this is indeed the case, MQ contains a few added features that make the job more pleasant.

13.2 Conditionally applying patches with guards

Perhaps the best way to maintain sanity with so many targets is to be able to choose specific patches to apply for a given situation. MQ provides a feature called "guards" (which originates with quilt's `guards` command) that does just this. To start off, let's create a simple repository for experimenting in.

```
1 $ hg qinit
2 $ hg qnew hello.patch
3 $ echo hello > hello
4 $ hg add hello
5 $ hg qrefresh
6 $ hg qnew goodbye.patch
7 $ echo goodbye > goodbye
8 $ hg add goodbye
9 $ hg qrefresh
```

This gives us a tiny repository that contains two patches that don't have any dependencies on each other, because they touch different files.

The idea behind conditional application is that you can "tag" a patch with a *guard*, which is simply a text string of your choosing, then tell MQ to select specific guards to use when applying patches. MQ will then either apply, or skip over, a guarded patch, depending on the guards that you have selected.

A patch can have an arbitrary number of guards; each one is *positive* ("apply this patch if this guard is selected") or *negative* ("skip this patch if this guard is selected"). A patch with no guards is always applied.

13.3 Controlling the guards on a patch

The "`hg qguard`" command lets you determine which guards should apply to a patch, or display the guards that are already in effect. Without any arguments, it displays the guards on the current topmost patch.

```
1 $ hg qguard
2 goodbye.patch: unguarded
```

To set a positive guard on a patch, prefix the name of the guard with a "+".

```
1 $ hg qguard +foo
2 $ hg qguard
3 goodbye.patch: +foo
```

To set a negative guard on a patch, prefix the name of the guard with a "-".

```

1 $ hg qguard hello.patch -quux
2 $ hg qguard hello.patch
3 hello.patch: -quux

```

Note: The “hg qguard” command *sets* the guards on a patch; it doesn’t *modify* them. What this means is that if you run “hg qguard +a +b” on a patch, then “hg qguard +c” on the same patch, the *only* guard that will be set on it afterwards is +c.

Mercurial stores guards in the `series` file; the form in which they are stored is easy both to understand and to edit by hand. (In other words, you don’t have to use the “hg qguard” command if you don’t want to; it’s okay to simply edit the `series` file.)

```

1 $ cat .hg/patches/series
2 hello.patch #-quux
3 goodbye.patch #+foo

```

13.4 Selecting the guards to use

The “hg qselect” command determines which guards are active at a given time. The effect of this is to determine which patches MQ will apply the next time you run “hg qpush”. It has no other effect; in particular, it doesn’t do anything to patches that are already applied.

With no arguments, the “hg qselect” command lists the guards currently in effect, one per line of output. Each argument is treated as the name of a guard to apply.

```

1 $ hg qpop -a
2 Patch queue now empty
3 $ hg qselect
4 no active guards
5 $ hg qselect foo
6 number of unguarded, unapplied patches has changed from 1 to 2
7 $ hg qselect
8 foo

```

In case you’re interested, the currently selected guards are stored in the `guards` file.

```

1 $ cat .hg/patches/guards
2 foo

```

We can see the effect the selected guards have when we run “hg qpush”.

```

1 $ hg qpush -a
2 applying hello.patch
3 applying goodbye.patch
4 Now at: goodbye.patch

```

A guard cannot start with a “+” or “-” character. The name of a guard must not contain white space, but most other characters are acceptable. If you try to use a guard with an invalid name, MQ will complain:

```

1 $ hg qselect +foo
2 abort: guard '+foo' starts with invalid character: '+'

```

Changing the selected guards changes the patches that are applied.

```
1 $ hg qselect quux
2 number of guarded, applied patches has changed from 0 to 2
3 $ hg qpop -a
4 Patch queue now empty
5 $ hg qpush -a
6 patch series already fully applied
```

You can see in the example below that negative guards take precedence over positive guards.

```
1 $ hg qselect foo bar
2 number of unguarded, unapplied patches has changed from 0 to 2
3 $ hg qpop -a
4 no patches applied
5 $ hg qpush -a
6 applying hello.patch
7 applying goodbye.patch
8 Now at: goodbye.patch
```

13.5 MQ's rules for applying patches

The rules that MQ uses when deciding whether to apply a patch are as follows.

- A patch that has no guards is always applied.
- If the patch has any negative guard that matches any currently selected guard, the patch is skipped.
- If the patch has any positive guard that matches any currently selected guard, the patch is applied.
- If the patch has positive or negative guards, but none matches any currently selected guard, the patch is skipped.

13.6 Trimming the work environment

In working on the device driver I mentioned earlier, I don't apply the patches to a normal Linux kernel tree. Instead, I use a repository that contains only a snapshot of the source files and headers that are relevant to Infiniband development. This repository is 1% the size of a kernel repository, so it's easier to work with.

I then choose a "base" version on top of which the patches are applied. This is a snapshot of the Linux kernel tree as of a revision of my choosing. When I take the snapshot, I record the changeset ID from the kernel repository in the commit message. Since the snapshot preserves the "shape" and content of the relevant parts of the kernel tree, I can apply my patches on top of either my tiny repository or a normal kernel tree.

Normally, the base tree atop which the patches apply should be a snapshot of a very recent upstream tree. This best facilitates the development of patches that can easily be submitted upstream with few or no modifications.

13.7 Dividing up the series file

I categorise the patches in the `series` file into a number of logical groups. Each section of like patches begins with a block of comments that describes the purpose of the patches that follow.

The sequence of patch groups that I maintain follows. The ordering of these groups is important; I'll describe why after I introduce the groups.

- The “accepted” group. Patches that the development team has submitted to the maintainer of the Infiniband subsystem, and which he has accepted, but which are not present in the snapshot that the tiny repository is based on. These are “read only” patches, present only to transform the tree into a similar state as it is in the upstream maintainer’s repository.
- The “rework” group. Patches that I have submitted, but that the upstream maintainer has requested modifications to before he will accept them.
- The “pending” group. Patches that I have not yet submitted to the upstream maintainer, but which we have finished working on. These will be “read only” for a while. If the upstream maintainer accepts them upon submission, I’ll move them to the end of the “accepted” group. If he requests that I modify any, I’ll move them to the beginning of the “rework” group.
- The “in progress” group. Patches that are actively being developed, and should not be submitted anywhere yet.
- The “backport” group. Patches that adapt the source tree to older versions of the kernel tree.
- The “do not ship” group. Patches that for some reason should never be submitted upstream. For example, one such patch might change embedded driver identification strings to make it easier to distinguish, in the field, between an out-of-tree version of the driver and a version shipped by a distribution vendor.

Now to return to the reasons for ordering groups of patches in this way. We would like the lowest patches in the stack to be as stable as possible, so that we will not need to rework higher patches due to changes in context. Putting patches that will never be changed first in the `series` file serves this purpose.

We would also like the patches that we know we’ll need to modify to be applied on top of a source tree that resembles the upstream tree as closely as possible. This is why we keep accepted patches around for a while.

The “backport” and “do not ship” patches float at the end of the `series` file. The backport patches must be applied on top of all other patches, and the “do not ship” patches might as well stay out of harm’s way.

13.8 Maintaining the patch series

In my work, I use a number of guards to control which patches are to be applied.

- “Accepted” patches are guarded with `accepted`. I enable this guard most of the time. When I’m applying the patches on top of a tree where the patches are already present, I can turn this patch off, and the patches that follow it will apply cleanly.
- Patches that are “finished”, but not yet submitted, have no guards. If I’m applying the patch stack to a copy of the upstream tree, I don’t need to enable any guards in order to get a reasonably safe source tree.
- Those patches that need reworking before being resubmitted are guarded with `rework`.
- For those patches that are still under development, I use `devel`.
- A backport patch may have several guards, one for each version of the kernel to which it applies. For example, a patch that backports a piece of code to 2.6.9 will have a `2.6.9` guard.

This variety of guards gives me considerable flexibility in determining what kind of source tree I want to end up with. For most situations, the selection of appropriate guards is automated during the build process, but I can manually tune the guards to use for less common circumstances.

13.8.1 The art of writing backport patches

Using MQ, writing a backport patch is a simple process. All such a patch has to do is modify a piece of code that uses a kernel feature not present in the older version of the kernel, so that the driver continues to work correctly under that older version.

A useful goal when writing a good backport patch is to make your code look as if it was written for the older version of the kernel you’re targeting. The less obtrusive the patch, the easier it will be to understand and maintain. If you’re writing a collection of backport patches to avoid the “rat’s nest” effect of lots of `#ifdefs` (hunks of source code that are only used conditionally) in your code, don’t introduce version-dependent `#ifdefs` into the patches. Instead, write several patches, each of which makes unconditional changes, and control their application using guards.

There are two reasons to divide backport patches into a distinct group, away from the “regular” patches whose effects they modify. The first is that intermingling the two makes it more difficult to use a tool like the `patchbomb` extension to automate the process of submitting the patches to an upstream maintainer. The second is that a backport patch could perturb the context in which a subsequent regular patch is applied, making it impossible to apply the regular patch cleanly *without* the earlier backport patch already being applied.

13.9 Useful tips for developing with MQ

13.9.1 Organising patches in directories

If you’re working on a substantial project with MQ, it’s not difficult to accumulate a large number of patches. For example, I have one patch repository that contains over 250 patches.

If you can group these patches into separate logical categories, you can if you like store them in different directories; MQ has no problems with patch names that contain path separators.

13.9.2 Viewing the history of a patch

If you’re developing a set of patches over a long time, it’s a good idea to maintain them in a repository, as discussed in section 12.11. If you do so, you’ll quickly discover that using the “`hg diff`” command to look at the history of changes to a patch is unworkable. This is in part because you’re looking at the second derivative of the real code (a diff of a diff), but also because MQ adds noise to the process by modifying time stamps and directory names when it updates a patch.

However, you can use the `extdiff` extension, which is bundled with Mercurial, to turn a diff of two versions of a patch into something readable. To do this, you will need a third-party package called `patchutils` [Wau]. This provides a command named `interdiff`, which shows the differences between two diffs as a diff. Used on two versions of the same diff, it generates a diff that represents the diff from the first to the second version.

You can enable the `extdiff` extension in the usual way, by adding a line to the `[extensions]` section of your `hgrc`.

```
1 [extensions]
2 extdiff =
```

The `interdiff` command expects to be passed the names of two files, but the `extdiff` extension passes the program it runs a pair of directories, each of which can contain an arbitrary number of files. We thus need a small program that will run `interdiff` on each pair of files in these two directories. This program is available as `hg-interdiff` in the `examples` directory of the source code repository that accompanies this book.

```
1 #!/usr/bin/env python
2 #
3 # Adapter for using interdiff with mercurial’s extdiff extension.
4 #
5 # Copyright 2006 Bryan O’Sullivan <bos@serpentine.com>
```

```

6  #
7  # This software may be used and distributed according to the terms of
8  # the GNU General Public License, incorporated herein by reference.
9
10 import os, sys
11
12 def walk(base):
13     # yield all non-directories below the base path.
14     for root, dirs, files in os.walk(base):
15         for f in files:
16             path = os.path.join(root, f)
17             yield path[len(base)+1:], path
18
19 # create list of unique file names under both directories.
20 files = dict(walk(sys.argv[1]))
21 files.update(walk(sys.argv[2]))
22 files = files.keys()
23 files.sort()
24
25 def name(base, f):
26     # interdiff requires two files; use /dev/null if one is missing.
27     path = os.path.join(base, f)
28     if os.path.exists(path):
29         return path
30     return '/dev/null'
31
32 ret = 0
33
34 for f in files:
35     if os.system('interdiff "%s" "%s"' % (name(sys.argv[1], f),
36                                           name(sys.argv[2], f))):
37         ret = 1
38
39 sys.exit(ret)

```

With the `hg-interdiff` program in your shell's search path, you can run it as follows, from inside an MQ patch directory:

```

1 hg extdiff -p hg-interdiff -r A:B my-change.patch

```

Since you'll probably want to use this long-winded command a lot, you can get `hgext` to make it available as a normal Mercurial command, again by editing your `hgrc`.

```

1 [extdiff]
2 cmd.interdiff = hg-interdiff

```

This directs `hgext` to make an `interdiff` command available, so you can now shorten the previous invocation of “`hg extdiff`” to something a little more wieldy.

```

1 hg interdiff -r A:B my-change.patch

```

Note: The `interdiff` command works well only if the underlying files against which versions of a patch are generated remain the same. If you create a patch, modify the underlying files, and then regenerate the patch, `interdiff` may not produce useful output.

The `extdiff` extension is useful for more than merely improving the presentation of MQ patches. To read more about it, go to section [14.2](#).

Chapter 14

Adding functionality with extensions

While the core of Mercurial is quite complete from a functionality standpoint, it's deliberately shorn of fancy features. This approach of preserving simplicity keeps the software easy to deal with for both maintainers and users.

However, Mercurial doesn't box you in with an inflexible command set: you can add features to it as *extensions* (sometimes known as *plugins*). We've already discussed a few of these extensions in earlier chapters.

- Section 3.3 covers the `fetch` extension; this combines pulling new changes and merging them with local changes into a single command, “`hg fetch`”.
- The `bisect` extension adds an efficient pruning search for changes that introduced bugs, and we documented it in chapter 9.5.
- In chapter 10, we covered several extensions that are useful for hook-related functionality: `acl` adds access control lists; `bugzilla` adds integration with the Bugzilla bug tracking system; and `notify` sends notification emails on new changes.
- The Mercurial Queues patch management extension is so invaluable that it merits two chapters and an appendix all to itself. Chapter 12 covers the basics; chapter 13 discusses advanced topics; and appendix B goes into detail on each command.

In this chapter, we'll cover some of the other extensions that are available for Mercurial, and briefly touch on some of the machinery you'll need to know about if you want to write an extension of your own.

- In section 14.1, we'll discuss the possibility of *huge* performance improvements using the `inotify` extension.

14.1 Improve performance with the `inotify` extension

Are you interested in having some of the most common Mercurial operations run as much as a hundred times faster? Read on!

Mercurial has great performance under normal circumstances. For example, when you run the “`hg status`” command, Mercurial has to scan almost every directory and file in your repository so that it can display file status. Many other Mercurial commands need to do the same work behind the scenes; for example, the “`hg diff`” command uses the status machinery to avoid doing an expensive comparison operation on files that obviously haven't changed.

Because obtaining file status is crucial to good performance, the authors of Mercurial have optimised this code to within an inch of its life. However, there's no avoiding the fact that when you run “`hg status`”, Mercurial is going to have to perform at least one expensive system call for each managed file to determine whether it's changed since the last time Mercurial checked. For a sufficiently large repository, this can take a long time.

To put a number on the magnitude of this effect, I created a repository containing 150,000 managed files. I timed “`hg status`” as taking ten seconds to run, even when *none* of those files had been modified.

Many modern operating systems contain a file notification facility. If a program signs up to an appropriate service, the operating system will notify it every time a file of interest is created, modified, or deleted. On Linux systems, the kernel component that does this is called `inotify`.

Mercurial's `inotify` extension talks to the kernel's `inotify` component to optimise “`hg status`” commands. The extension has two components. A daemon sits in the background and receives notifications from the `inotify` subsystem. It also listens for connections from a regular Mercurial command. The extension modifies Mercurial's behaviour so that instead of scanning the filesystem, it queries the daemon. Since the daemon has perfect information about the state of the repository, it can respond with a result instantaneously, avoiding the need to scan every directory and file in the repository.

Recall the ten seconds that I measured plain Mercurial as taking to run “`hg status`” on a 150,000 file repository. With the `inotify` extension enabled, the time dropped to 0.1 seconds, a factor of *one hundred* faster.

Before we continue, please pay attention to some caveats.

- The `inotify` extension is Linux-specific. Because it interfaces directly to the Linux kernel's `inotify` subsystem, it does not work on other operating systems.
- It should work on any Linux distribution that was released after early 2005. Older distributions are likely to have a kernel that lacks `inotify`, or a version of `glibc` that does not have the necessary interfacing support.
- Not all filesystems are suitable for use with the `inotify` extension. Network filesystems such as NFS are a non-starter, for example, particularly if you're running Mercurial on several systems, all mounting the same network filesystem. The kernel's `inotify` system has no way of knowing about changes made on another system. Most local filesystems (e.g. `ext3`, `XFS`, `ReiserFS`) should work fine.

The `inotify` extension is not yet shipped with Mercurial as of May 2007, so it's a little more involved to set up than other extensions. But the performance improvement is worth it!

The extension currently comes in two parts: a set of patches to the Mercurial source code, and a library of Python bindings to the `inotify` subsystem.

Note: There are *two* Python `inotify` binding libraries. One of them is called `pyinotify`, and is packaged by some Linux distributions as `python-inotify`. This is *not* the one you'll need, as it is too buggy and inefficient to be practical.

To get going, it's best to already have a functioning copy of Mercurial installed.

Note: If you follow the instructions below, you'll be replacing any existing installation of Mercurial that you might already have. This is a “bleeding edge” Mercurial code. Don't say you weren't warned!

1. Clone the Python `inotify` binding repository. Build and install it.

```
1 hg clone http://hg.kublai.com/python/inotify
2 cd inotify
3 python setup.py build --force
4 sudo python setup.py install --skip-build
```

2. Clone the `crew` Mercurial repository. Clone the `inotify` patch repository so that Mercurial Queues will be able to apply patches to your copy of the `crew` repository.

```
1 hg clone http://hg.intevation.org/mercurial/crew
2 hg clone crew inotify
3 hg clone http://hg.kublai.com/mercurial/patches/inotify inotify/.hg/patches
```

3. Make sure that you have the Mercurial Queues extension, `mq`, enabled. If you've never used MQ, read section 12.5 to get started quickly.
4. Go into the `inotify` repo, and apply all of the `inotify` patches using the `-a` option to the “`hg qpush`” command.

```
1 cd inotify
2 hg qpush -a
```

If you get an error message from “hg qpush”, you should not continue. Instead, ask for help.

5. Build and install the patched version of Mercurial.

```
1 python setup.py build --force
2 sudo python setup.py install --skip-build
```

Once you’ve build a suitably patched version of Mercurial, all you need to do to enable the `inotify` extension is add an entry to your `hgrc`.

```
1 [extensions]
2 inotify =
```

When the `inotify` extension is enabled, Mercurial will automatically and transparently start the status daemon the first time you run a command that needs status in a repository. It runs one status daemon per repository.

The status daemon is started silently, and runs in the background. If you look at a list of running processes after you’ve enabled the `inotify` extension and run a few commands in different repositories, you’ll thus see a few `hg` processes sitting around, waiting for updates from the kernel and queries from Mercurial.

The first time you run a Mercurial command in a repository when you have the `inotify` extension enabled, it will run with about the same performance as a normal Mercurial command. This is because the status daemon needs to perform a normal status scan so that it has a baseline against which to apply later updates from the kernel. However, *every* subsequent command that does any kind of status check should be noticeably faster on repositories of even fairly modest size. Better yet, the bigger your repository is, the greater a performance advantage you’ll see. The `inotify` daemon makes status operations almost instantaneous on repositories of all sizes!

If you like, you can manually start a status daemon using the “hg inserve” command. This gives you slightly finer control over how the daemon ought to run. This command will of course only be available when the `inotify` extension is enabled.

When you’re using the `inotify` extension, you should notice *no difference at all* in Mercurial’s behaviour, with the sole exception of status-related commands running a whole lot faster than they used to. You should specifically expect that commands will not print different output; neither should they give different results. If either of these situations occurs, please report a bug.

14.2 Flexible diff support with the `extdiff` extension

Mercurial’s built-in “hg diff” command outputs plaintext unified diffs.

```
1 $ hg diff
2 diff -r 1824bf2dd15e myfile
3 --- a/myfile      Mon Dec 10 19:54:28 2007 +0000
4 +++ b/myfile      Mon Dec 10 19:54:28 2007 +0000
5 @@ -1,1 +1,2 @@ The first line.
6  The first line.
7 +The second line.
```

If you would like to use an external tool to display modifications, you’ll want to use the `extdiff` extension. This will let you use, for example, a graphical diff tool.

The `extdiff` extension is bundled with Mercurial, so it’s easy to set up. In the `[extensions]` section of your `hgrc`, simply add a one-line entry to enable the extension.

```

1 [extensions]
2 extdiff =

```

This introduces a command named “hg extdiff”, which by default uses your system’s diff command to generate a unified diff in the same form as the built-in “hg diff” command.

```

1 $ hg extdiff
2 --- a.1824bf2dd15e/myfile      2007-12-10 19:54:28.000000000 +0000
3 +++ /tmp/extdiff9pHeZ6/a/myfile      2007-12-10 19:54:28.000000000 +0000
4 @@ -1 +1,2 @@
5   The first line.
6 +The second line.

```

The result won’t be exactly the same as with the built-in “hg diff” variations, because the output of diff varies from one system to another, even when passed the same options.

As the “making snapshot” lines of output above imply, the “hg extdiff” command works by creating two snapshots of your source tree. The first snapshot is of the source revision; the second, of the target revision or working directory. The “hg extdiff” command generates these snapshots in a temporary directory, passes the name of each directory to an external diff viewer, then deletes the temporary directory. For efficiency, it only snapshots the directories and files that have changed between the two revisions.

Snapshot directory names have the same base name as your repository. If your repository path is /quux/bar/foo, then foo will be the name of each snapshot directory. Each snapshot directory name has its changeset ID appended, if appropriate. If a snapshot is of revision a631aca1083f, the directory will be named foo.a631aca1083f. A snapshot of the working directory won’t have a changeset ID appended, so it would just be foo in this example. To see what this looks like in practice, look again at the “hg extdiff” example above. Notice that the diff has the snapshot directory names embedded in its header.

The “hg extdiff” command accepts two important options. The -p option lets you choose a program to view differences with, instead of diff. With the -o option, you can change the options that “hg extdiff” passes to the program (by default, these options are “-Npru”, which only make sense if you’re running diff). In other respects, the “hg extdiff” command acts similarly to the built-in “hg diff” command: you use the same option names, syntax, and arguments to specify the revisions you want, the files you want, and so on.

As an example, here’s how to run the normal system diff command, getting it to generate context diffs (using the -c option) instead of unified diffs, and five lines of context instead of the default three (passing 5 as the argument to the -C option).

```

1 $ hg extdiff -o -Nprc5
2 *** a.1824bf2dd15e/myfile      Mon Dec 10 19:54:28 2007
3 --- /tmp/extdiff9pHeZ6/a/myfile      Mon Dec 10 19:54:28 2007
4 *****
5 *** 1 ****
6 --- 1,2 ----
7   The first line.
8 + The second line.

```

Launching a visual diff tool is just as easy. Here’s how to launch the kdiff3 viewer.

```

1 hg extdiff -p kdiff3 -o ''

```

If your diff viewing command can’t deal with directories, you can easily work around this with a little scripting. For an example of such scripting in action with the mq extension and the interdiff command, see section 13.9.2.

14.2.1 Defining command aliases

It can be cumbersome to remember the options to both the “hg extdiff” command and the diff viewer you want to use, so the `extdiff` extension lets you define *new* commands that will invoke your diff viewer with exactly the right options.

All you need to do is edit your `hgrc`, and add a section named `[extdiff]`. Inside this section, you can define multiple commands. Here’s how to add a `kdiff3` command. Once you’ve defined this, you can type “hg kdiff3” and the `extdiff` extension will run `kdiff3` for you.

```
1 [extdiff]
2 cmd.kdiff3 =
```

If you leave the right hand side of the definition empty, as above, the `extdiff` extension uses the name of the command you defined as the name of the external program to run. But these names don’t have to be the same. Here, we define a command named “hg wibble”, which runs `kdiff3`.

```
1 [extdiff]
2 cmd.wibble = kdiff3
```

You can also specify the default options that you want to invoke your diff viewing program with. The prefix to use is “`opts.`”, followed by the name of the command to which the options apply. This example defines a “hg vimdiff” command that runs the `vim` editor’s `DirDiff` extension.

```
1 [extdiff]
2 cmd.vimdiff = vim
3 opts.vimdiff = -f '+next' '+execute "DirDiff" argv(0) argv(1)'
```

14.3 Cherry-picking changes with the `transplant` extension

Need to have a long chat with Brendan about this.

14.4 Send changes via email with the `patchbomb` extension

Many projects have a culture of “change review”, in which people send their modifications to a mailing list for others to read and comment on before they commit the final version to a shared repository. Some projects have people who act as gatekeepers; they apply changes from other people to a repository to which those others don’t have access.

Mercurial makes it easy to send changes over email for review or application, via its `patchbomb` extension. The extension is so named because changes are formatted as patches, and it’s usual to send one changeset per email message. Sending a long series of changes by email is thus much like “bombing” the recipient’s inbox, hence “patchbomb”.

As usual, the basic configuration of the `patchbomb` extension takes just one or two lines in your `hgrc`.

```
1 [extensions]
2 patchbomb =
```

Once you’ve enabled the extension, you will have a new command available, named “hg email”.

The safest and best way to invoke the “hg email” command is to *always* run it first with the `-n` option. This will show you what the command *would* send, without actually sending anything. Once you’ve had a quick glance over the changes and verified that you are sending the right ones, you can rerun the same command, with the `-n` option removed.

The “hg email” command accepts the same kind of revision syntax as every other Mercurial command. For example, this command will send every revision between 7 and `tip`, inclusive.

```
hg email -n 7:tip
```

You can also specify a *repository* to compare with. If you provide a repository but no revisions, the “hg email” command will send all revisions in the local repository that are not present in the remote repository. If you additionally specify revisions or a branch name (the latter using the `-b` option), this will constrain the revisions sent.

It’s perfectly safe to run the “hg email” command without the names of the people you want to send to: if you do this, it will just prompt you for those values interactively. (If you’re using a Linux or Unix-like system, you should have enhanced *readline*-style editing capabilities when entering those headers, too, which is useful.)

When you are sending just one revision, the “hg email” command will by default use the first line of the changeset description as the subject of the single email message it sends.

If you send multiple revisions, the “hg email” command will usually send one message per changeset. It will preface the series with an introductory message, in which you should describe the purpose of the series of changes you’re sending.

14.4.1 Changing the behaviour of patchbombs

Not every project has exactly the same conventions for sending changes in email; the `patchbomb` extension tries to accommodate a number of variations through command line options.

- You can write a subject for the introductory message on the command line using the `-s` option. This takes one argument, the text of the subject to use.
- To change the email address from which the messages originate, use the `-f` option. This takes one argument, the email address to use.
- The default behaviour is to send unified diffs (see section 12.4 for a description of the format), one per message. You can send a binary bundle instead with the `-b` option.
- Unified diffs are normally prefaced with a metadata header. You can omit this, and send unadorned diffs, with the `--plain` option.
- Diffs are normally sent “inline”, in the same body part as the description of a patch. This makes it easiest for the largest number of readers to quote and respond to parts of a diff, as some mail clients will only quote the first MIME body part in a message. If you’d prefer to send the description and the diff in separate body parts, use the `-a` option.
- Instead of sending mail messages, you can write them to an `mbx`-format mail folder using the `-m` option. That option takes one argument, the name of the file to write to.
- If you would like to add a `diffstat`-format summary to each patch, and one to the introductory message, use the `-d` option. The `diffstat` command displays a table containing the name of each file patched, the number of lines affected, and a histogram showing how much each file is modified. This gives readers a qualitative glance at how complex a patch is.

Appendix A

Command reference

A.1 “hg add”—add files at the next commit

--include, also -I

--exclude, also -X

--dry-run, also -n

A.2 “hg diff”—print changes in history or working directory

Show differences between revisions for the specified files or directories, using the unified diff format. For a description of the unified diff format, see section [12.4](#).

By default, this command does not print diffs for files that Mercurial considers to contain binary data. To control this behaviour, see the `-a` and `--git` options.

A.2.1 Options

--nodates option

Omit date and time information when printing diff headers.

--ignore-blank-lines, also -B

Do not print changes that only insert or delete blank lines. A line that contains only whitespace is not considered blank.

--include, also -I

Include files and directories whose names match the given patterns.

--exclude, also -X

Exclude files and directories whose names match the given patterns.

--text, also -a

If this option is not specified, “hg diff” will refuse to print diffs for files that it detects as binary. Specifying `-a` forces “hg diff” to treat all files as text, and generate diffs for all of them.

This option is useful for files that are “mostly text” but have a few embedded NUL characters. If you use it on files that contain a lot of binary data, its output will be incomprehensible.

--ignore-space-change, also -b

Do not print a line if the only change to that line is in the amount of white space it contains.

--git, also -g

Print git-compatible diffs. XXX reference a format description.

--show-function, also -p

Display the name of the enclosing function in a hunk header, using a simple heuristic. This functionality is enabled by default, so the -p option has no effect unless you change the value of the `showfunc` config item, as in the following example.

```
1  $ echo '[diff]' >> $HGRC
2  $ echo 'showfunc = False' >> $HGRC
3  $ hg diff
4  diff -r 212644e6881d myfile.c
5  --- a/myfile.c      Mon Dec 10 19:54:23 2007 +0000
6  +++ b/myfile.c      Mon Dec 10 19:54:23 2007 +0000
7  @@ -1,4 +1,4 @@
8   int myfunc()
9   {
10  -     return 1;
11  +     return 10;
12  }
13  $ hg diff -p
14  diff -r 212644e6881d myfile.c
15  --- a/myfile.c      Mon Dec 10 19:54:23 2007 +0000
16  +++ b/myfile.c      Mon Dec 10 19:54:23 2007 +0000
17  @@ -1,4 +1,4 @@ int myfunc()
18   int myfunc()
19   {
20  -     return 1;
21  +     return 10;
22  }
```

--rev, also -r

Specify one or more revisions to compare. The “hg diff” command accepts up to two -r options to specify the revisions to compare.

1. Display the differences between the parent revision of the working directory and the working directory.
2. Display the differences between the specified changeset and the working directory.
3. Display the differences between the two specified changesets.

You can specify two revisions using either two -r options or revision range notation. For example, the two revision specifications below are equivalent.

```
1 hg diff -r 10 -r 20
2 hg diff -r10:20
```

When you provide two revisions, Mercurial treats the order of those revisions as significant. Thus, “`hg diff -r10:20`” will produce a diff that will transform files from their contents as of revision 10 to their contents as of revision 20, while “`hg diff -r20:10`” means the opposite: the diff that will transform files from their revision 20 contents to their revision 10 contents. You cannot reverse the ordering in this way if you are diffing against the working directory.

--ignore-all-space, also -w

A.3 “hg version”—print version and copyright information

This command displays the version of Mercurial you are running, and its copyright license. There are four kinds of version string that you may see.

- The string “unknown”. This version of Mercurial was not built in a Mercurial repository, and cannot determine its own version.
- A short numeric string, such as “1.1”. This is a build of a revision of Mercurial that was identified by a specific tag in the repository where it was built. (This doesn’t necessarily mean that you’re running an official release; someone else could have added that tag to any revision in the repository where they built Mercurial.)
- A hexadecimal string, such as “875489e31abe”. This is a build of the given revision of Mercurial.
- A hexadecimal string followed by a date, such as “875489e31abe+20070205”. This is a build of the given revision of Mercurial, where the build repository contained some local changes that had not been committed.

A.3.1 Tips and tricks

Why do the results of “hg diff” and “hg status” differ?

When you run the “`hg status`” command, you’ll see a list of files that Mercurial will record changes for the next time you perform a commit. If you run the “`hg diff`” command, you may notice that it prints diffs for only a *subset* of the files that “`hg status`” listed. There are two possible reasons for this.

The first is that “`hg status`” prints some kinds of modifications that “`hg diff`” doesn’t normally display. The “`hg diff`” command normally outputs unified diffs, which don’t have the ability to represent some changes that Mercurial can track. Most notably, traditional diffs can’t represent a change in whether or not a file is executable, but Mercurial records this information.

If you use the `--git` option to “`hg diff`”, it will display `git`-compatible diffs that *can* display this extra information.

The second possible reason that “`hg diff`” might be printing diffs for a subset of the files displayed by “`hg status`” is that if you invoke it without any arguments, “`hg diff`” prints diffs against the first parent of the working directory. If you have run “`hg merge`” to merge two changesets, but you haven’t yet committed the results of the merge, your working directory has two parents (use “`hg parents`” to see them). While “`hg status`” prints modifications relative to *both* parents after an uncommitted merge, “`hg diff`” still operates relative only to the first parent. You can get it to print diffs relative to the second parent by specifying that parent with the `-r` option. There is no way to print diffs relative to both parents.

Generating safe binary diffs

If you use the `-a` option to force Mercurial to print diffs of files that are either “mostly text” or contain lots of binary data, those diffs cannot subsequently be applied by either Mercurial’s `hg import` command or the system’s `patch` command.

If you want to generate a diff of a binary file that is safe to use as input for `hg import`, use the `“hg diff”-git` option when you generate the patch. The system `patch` command cannot handle binary patches at all.

Appendix B

Mercurial Queues reference

B.1 MQ command reference

For an overview of the commands provided by MQ, use the command `hg help mq`.

B.1.1 “`hg qapplied`”—print applied patches

The `hg qapplied` command prints the current stack of applied patches. Patches are printed in oldest-to-newest order, so the last patch in the list is the “top” patch.

B.1.2 “`hg qcommit`”—commit changes in the queue repository

The `hg qcommit` command commits any outstanding changes in the `.hg/patches` repository. This command only works if the `.hg/patches` directory is a repository, i.e. you created the directory using `hg qinit -c` or ran `hg init` in the directory after running `hg qinit`.

This command is shorthand for `hg commit --cwd .hg/patches`.

B.1.3 “`hg qdelete`”—delete a patch from the `series` file

The `hg qdelete` command removes the entry for a patch from the `series` file in the `.hg/patches` directory. It does not pop the patch if the patch is already applied. By default, it does not delete the patch file; use the `-f` option to do that.

Options:

`-f` Delete the patch file.

B.1.4 “`hg qdiff`”—print a diff of the topmost applied patch

The `hg qdiff` command prints a diff of the topmost applied patch. It is equivalent to `hg diff -r-2:-1`.

B.1.5 “`hg qfold`”—merge (“fold”) several patches into one

The `hg qfold` command merges multiple patches into the topmost applied patch, so that the topmost applied patch makes the union of all of the changes in the patches in question.

The patches to fold must not be applied; `hg qfold` will exit with an error if any is. The order in which patches are folded is significant; `hg qfold a b` means “apply the current topmost patch, followed by a, followed by b”.

The comments from the folded patches are appended to the comments of the destination patch, with each block of comments separated by three asterisk (“`***`”) characters. Use the `-e` option to edit the commit message for the combined patch/changeset after the folding has completed.

Options:

- e Edit the commit message and patch description for the newly folded patch.
- l Use the contents of the given file as the new commit message and patch description for the folded patch.
- m Use the given text as the new commit message and patch description for the folded patch.

B.1.6 “hg qheader”—display the header/description of a patch

The “hg qheader” command prints the header, or description, of a patch. By default, it prints the header of the topmost applied patch. Given an argument, it prints the header of the named patch.

B.1.7 “hg qimport”—import a third-party patch into the queue

The “hg qimport” command adds an entry for an external patch to the `series` file, and copies the patch into the `.hg/patches` directory. It adds the entry immediately after the topmost applied patch, but does not push the patch.

If the `.hg/patches` directory is a repository, “hg qimport” automatically does an “hg add” of the imported patch.

B.1.8 “hg qinit”—prepare a repository to work with MQ

The “hg qinit” command prepares a repository to work with MQ. It creates a directory called `.hg/patches`.

Options:

- c Create `.hg/patches` as a repository in its own right. Also creates a `.hgignore` file that will ignore the status file.

When the `.hg/patches` directory is a repository, the “hg qimport” and “hg qnew” commands automatically “hg add” new patches.

B.1.9 “hg qnew”—create a new patch

The “hg qnew” command creates a new patch. It takes one mandatory argument, the name to use for the patch file. The newly created patch is created empty by default. It is added to the `series` file after the current topmost applied patch, and is immediately pushed on top of that patch.

If “hg qnew” finds modified files in the working directory, it will refuse to create a new patch unless the `-f` option is used (see below). This behaviour allows you to “hg qrefresh” your topmost applied patch before you apply a new patch on top of it.

Options:

- f Create a new patch if the contents of the working directory are modified. Any outstanding modifications are added to the newly created patch, so after this command completes, the working directory will no longer be modified.
- m Use the given text as the commit message. This text will be stored at the beginning of the patch file, before the patch data.

B.1.10 “hg qnext”—print the name of the next patch

The “hg qnext” command prints the name of the next patch in the `series` file after the topmost applied patch. This patch will become the topmost applied patch if you run “hg qpush”.

B.1.11 “hg qpop”—pop patches off the stack

The “hg qpop” command removes applied patches from the top of the stack of applied patches. By default, it removes only one patch.

This command removes the changesets that represent the popped patches from the repository, and updates the working directory to undo the effects of the patches.

This command takes an optional argument, which it uses as the name or index of the patch to pop to. If given a name, it will pop patches until the named patch is the topmost applied patch. If given a number, “hg qpop” treats the number as an index into the entries in the series file, counting from zero (empty lines and lines containing only comments do not count). It pops patches until the patch identified by the given index is the topmost applied patch.

The “hg qpop” command does not read or write patches or the series file. It is thus safe to “hg qpop” a patch that you have removed from the series file, or a patch that you have renamed or deleted entirely. In the latter two cases, use the name of the patch as it was when you applied it.

By default, the “hg qpop” command will not pop any patches if the working directory has been modified. You can override this behaviour using the -f option, which reverts all modifications in the working directory.

Options:

- a Pop all applied patches. This returns the repository to its state before you applied any patches.
- f Forcibly revert any modifications to the working directory when popping.
- n Pop a patch from the named queue.

The “hg qpop” command removes one line from the end of the status file for each patch that it pops.

B.1.12 “hg qprev”—print the name of the previous patch

The “hg qprev” command prints the name of the patch in the series file that comes before the topmost applied patch. This will become the topmost applied patch if you run “hg qpop”.

B.1.13 “hg qpush”—push patches onto the stack

The “hg qpush” command adds patches onto the applied stack. By default, it adds only one patch.

This command creates a new changeset to represent each applied patch, and updates the working directory to apply the effects of the patches.

The default data used when creating a changeset are as follows:

- The commit date and time zone are the current date and time zone. Because these data are used to compute the identity of a changeset, this means that if you “hg qpop” a patch and “hg qpush” it again, the changeset that you push will have a different identity than the changeset you popped.
- The author is the same as the default used by the “hg commit” command.
- The commit message is any text from the patch file that comes before the first diff header. If there is no such text, a default commit message is used that identifies the name of the patch.

If a patch contains a Mercurial patch header (XXX add link), the information in the patch header overrides these defaults.

Options:

- a Push all unapplied patches from the series file until there are none left to push.
- l Add the name of the patch to the end of the commit message.
- m If a patch fails to apply cleanly, use the entry for the patch in another saved queue to compute the parameters for a three-way merge, and perform a three-way merge using the normal Mercurial merge machinery. Use the resolution of the merge as the new patch content.

-n Use the named queue if merging while pushing.

The “hg qpush” command reads, but does not modify, the `series` file. It appends one line to the “hg status” file for each patch that it pushes.

B.1.14 “hg qrefresh”—update the topmost applied patch

The “hg qrefresh” command updates the topmost applied patch. It modifies the patch, removes the old changeset that represented the patch, and creates a new changeset to represent the modified patch.

The “hg qrefresh” command looks for the following modifications:

- Changes to the commit message, i.e. the text before the first diff header in the patch file, are reflected in the new changeset that represents the patch.
- Modifications to tracked files in the working directory are added to the patch.
- Changes to the files tracked using “hg add”, “hg copy”, “hg remove”, or “hg rename”. Added files and copy and rename destinations are added to the patch, while removed files and rename sources are removed.

Even if “hg qrefresh” detects no changes, it still recreates the changeset that represents the patch. This causes the identity of the changeset to differ from the previous changeset that identified the patch.

Options:

- e Modify the commit and patch description, using the preferred text editor.
- m Modify the commit message and patch description, using the given text.
- l Modify the commit message and patch description, using text from the given file.

B.1.15 “hg qrename”—rename a patch

The “hg qrename” command renames a patch, and changes the entry for the patch in the `series` file.

With a single argument, “hg qrename” renames the topmost applied patch. With two arguments, it renames its first argument to its second.

B.1.16 “hg qrestore”—restore saved queue state

XXX No idea what this does.

B.1.17 “hg qsave”—save current queue state

XXX Likewise.

B.1.18 “hg qseries”—print the entire patch series

The “hg qseries” command prints the entire patch series from the `series` file. It prints only patch names, not empty lines or comments. It prints in order from first to be applied to last.

B.1.19 “hg qtop”—print the name of the current patch

The “hg qtop” prints the name of the topmost currently applied patch.

B.1.20 “hg qunapplied”—print patches not yet applied

The “hg qunapplied” command prints the names of patches from the `series` file that are not yet applied. It prints them in order from the next patch that will be pushed to the last.

B.1.21 “hg qversion”

The “hg qversion” command prints the version of MQ that is in use.

B.1.22 “hg strip”—remove a revision and descendants

The “hg strip” command removes a revision, and all of its descendants, from the repository. It undoes the effects of the removed revisions from the repository, and updates the working directory to the first parent of the removed revision.

The “hg strip” command saves a backup of the removed changesets in a bundle, so that they can be reapplied if removed in error.

Options:

- b Save unrelated changesets that are intermixed with the stripped changesets in the backup bundle.
- f If a branch has multiple heads, remove all heads. XXX This should be renamed, and use -f to strip revs when there are pending changes.
- n Do not save a backup bundle.

B.2 MQ file reference

B.2.1 The series file

The `series` file contains a list of the names of all patches that MQ can apply. It is represented as a list of names, with one name saved per line. Leading and trailing white space in each line are ignored.

Lines may contain comments. A comment begins with the “#” character, and extends to the end of the line. Empty lines, and lines that contain only comments, are ignored.

You will often need to edit the `series` file by hand, hence the support for comments and empty lines noted above. For example, you can comment out a patch temporarily, and “hg qpush” will skip over that patch when applying patches. You can also change the order in which patches are applied by reordering their entries in the `series` file.

Placing the `series` file under revision control is also supported; it is a good idea to place all of the patches that it refers to under revision control, as well. If you create a patch directory using the -c option to “hg qinit”, this will be done for you automatically.

B.2.2 The status file

The `status` file contains the names and changeset hashes of all patches that MQ currently has applied. Unlike the `series` file, this file is not intended for editing. You should not place this file under revision control, or modify it in any way. It is used by MQ strictly for internal book-keeping.

Appendix C

Installing Mercurial from source

C.1 On a Unix-like system

If you are using a Unix-like system that has a sufficiently recent version of Python (2.3 or newer) available, it is easy to install Mercurial from source.

1. Download a recent source tarball from <http://www.selenic.com/mercurial/download>.
2. Unpack the tarball:

```
1  gzip -dc mercurial-version.tar.gz | tar xf -
```

3. Go into the source directory and run the installer script. This will build Mercurial and install it in your home directory.

```
1  cd mercurial-version
2  python setup.py install --force --home=$HOME
```

Once the install finishes, Mercurial will be in the `bin` subdirectory of your home directory. Don't forget to make sure that this directory is present in your shell's search path.

You will probably need to set the `PYTHONPATH` environment variable so that the Mercurial executable can find the rest of the Mercurial packages. For example, on my laptop, I have set it to `/home/bos/lib/python`. The exact path that you will need to use depends on how Python was built for your system, but should be easy to figure out. If you're uncertain, look through the output of the installer script above, and see where the contents of the `mercurial` directory were installed to.

C.2 On Windows

Building and installing Mercurial on Windows requires a variety of tools, a fair amount of technical knowledge, and considerable patience. I very much *do not recommend* this route if you are a “casual user”. Unless you intend to hack on Mercurial, I strongly suggest that you use a binary package instead.

If you are intent on building Mercurial from source on Windows, follow the “hard way” directions on the Mercurial wiki at <http://www.selenic.com/mercurial/wiki/index.cgi/WindowsInstall>, and expect the process to involve a lot of fiddly work.

Appendix D

Open Publication License

Version 1.0, 8 June 1999

D.1 Requirements on both unmodified and modified versions

The Open Publication works may be reproduced and distributed in whole or in part, in any medium physical or electronic, provided that the terms of this license are adhered to, and that this license or an incorporation of it by reference (with any options elected by the author(s) and/or publisher) is displayed in the reproduction.

Proper form for an incorporation by reference is as follows:

Copyright (c) *year* by *author's name or designee*. This material may be distributed only subject to the terms and conditions set forth in the Open Publication License, vx.y or later (the latest version is presently available at <http://www.opencontent.org/openpub/>).

The reference must be immediately followed with any options elected by the author(s) and/or publisher of the document (see section D.6).

Commercial redistribution of Open Publication-licensed material is permitted.

Any publication in standard (paper) book form shall require the citation of the original publisher and author. The publisher and author's names shall appear on all outer surfaces of the book. On all outer surfaces of the book the original publisher's name shall be as large as the title of the work and cited as possessive with respect to the title.

D.2 Copyright

The copyright to each Open Publication is owned by its author(s) or designee.

D.3 Scope of license

The following license terms apply to all Open Publication works, unless otherwise explicitly stated in the document.

Mere aggregation of Open Publication works or a portion of an Open Publication work with other works or programs on the same media shall not cause this license to apply to those other works. The aggregate work shall contain a notice specifying the inclusion of the Open Publication material and appropriate copyright notice.

Severability. If any part of this license is found to be unenforceable in any jurisdiction, the remaining portions of the license remain in force.

No warranty. Open Publication works are licensed and provided "as is" without warranty of any kind, express or implied, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose or a warranty of non-infringement.

D.4 Requirements on modified works

All modified versions of documents covered by this license, including translations, anthologies, compilations and partial documents, must meet the following requirements:

1. The modified version must be labeled as such.
2. The person making the modifications must be identified and the modifications dated.
3. Acknowledgement of the original author and publisher if applicable must be retained according to normal academic citation practices.
4. The location of the original unmodified document must be identified.
5. The original author's (or authors') name(s) may not be used to assert or imply endorsement of the resulting document without the original author's (or authors') permission.

D.5 Good-practice recommendations

In addition to the requirements of this license, it is requested from and strongly recommended of redistributors that:

1. If you are distributing Open Publication works on hardcopy or CD-ROM, you provide email notification to the authors of your intent to redistribute at least thirty days before your manuscript or media freeze, to give the authors time to provide updated documents. This notification should describe modifications, if any, made to the document.
2. All substantive modifications (including deletions) be either clearly marked up in the document or else described in an attachment to the document.
3. Finally, while it is not mandatory under this license, it is considered good form to offer a free copy of any hardcopy and CD-ROM expression of an Open Publication-licensed work to its author(s).

D.6 License options

The author(s) and/or publisher of an Open Publication-licensed document may elect certain options by appending language to the reference to or copy of the license. These options are considered part of the license instance and must be included with the license (or its incorporation by reference) in derived works.

- A To prohibit distribution of substantively modified versions without the explicit permission of the author(s). “Substantive modification” is defined as a change to the semantic content of the document, and excludes mere changes in format or typographical corrections.

To accomplish this, add the phrase “Distribution of substantively modified versions of this document is prohibited without the explicit permission of the copyright holder.” to the license reference or copy.

- B To prohibit any publication of this work or derivative works in whole or in part in standard (paper) book form for commercial purposes is prohibited unless prior permission is obtained from the copyright holder.

To accomplish this, add the phrase “Distribution of the work or derivative of the work in any standard (paper) book form is prohibited unless prior permission is obtained from the copyright holder.” to the license reference or copy.

Bibliography

- [AG] Jean Delvare Andreas Gruenbacher, Martin Quinson. Patchwork quilt. <http://savannah.nongnu.org/projects/quilt>.
- [BI] Ronald Oussoren Bob Ippolito. Universal macpython. <http://bob.pythonmac.org/archives/2006/04/10/python-and-universal-binaries-on-mac-os-x/>.
- [Bro] Neil Brown. wiggle—apply conflicting patches. <http://cgi.cse.unsw.edu.au/~neilb/source/wiggle/>.
- [Dic] Thomas Dickey. diffstat—make a histogram of diff output. <http://dickey.his.com/diffstat/diffstat.html>.
- [Dus] Andy Dustman. Mysql for python. <http://sourceforge.net/projects/mysql-python>.
- [Gru05] Andreas Gruenbacher. How to survive with many patches (introduction to quilt). <http://www.suse.de/~agruen/quilt.pdf>, June 2005.
- [Mas] Chris Mason. mpatch—help solve patch rejects. <http://oss.oracle.com/~mason/mpatch/>.
- [O’S06] Bryan O’Sullivan. Achieving high performance in mercurial. In *EuroPython Conference*, July 2006. XXX.
- [Pyt] Python.org. ConfigParser—configuration file parser. <http://docs.python.org/lib/module-ConfigParser.html>.
- [RS] GNU Project volunteers Richard Stallman. Gnu coding standards—change logs. http://www.gnu.org/prep/standards/html_node/Change-Logs.html.
- [Tat] Simon Tatham. Putty—open source ssh client for windows. <http://www.chiark.greenend.org.uk/~sgtatham/putty/>.
- [Wau] Tim Waugh. patchutils—programs that operate on patch files. <http://cyberelk.net/tim/patchutils/>.

Index

- `.hg/hgrc` file, 71, 107
- `.hg/localtags` file, 82, 123, 124
- `.hg/patches.N` directory, 147
- `.hg/patches` directory, 139, 148–150, 173, 174
- `.hg/store/data` directory, 35
- `.hgitignore` file, 150, 174
- `.hgrc` file, 18, 33
- `.hgtags` file, 81, 82, 123, 124
- `.orig` file, 144
- `.rej` file, 144, 146
- `.ssh/config` file, 66
- `.ssh` directory, 64, 65
- EMAIL environment variable, 18
- HGMERGE environment variable, 29, 31
- HGUSER environment variable, 18
- HG_NODE environment variable, 109, 120
- HG_PARENT1 environment variable, 120
- HG_PARENT2 environment variable, 120
- HG_SOURCE environment variable, 120
- HG_URL environment variable, 120, 121
- Mercurial.ini configuration file, 64
- PATH environment variable, 66
- PYTHONPATH environment variable, 66, 68, 111, 178
- acl extension, 113–115, 163
- acl hook, 113
- addbreaks template filter, 130
- addremove command, 50, 144
- add command, 42, 47–49, 52, 55, 73, 74, 89, 90, 92, 142, 150, 154, 169, 174, 176
 - dry-run option, 169
 - exclude option, 169
 - include option, 169
 - I option, 169
 - X option, 169
 - n option, 169
- age template filter, 130
- annotate command, 137, 140
- authorized_keys file, 64, 65
- author template keyword, 128, 130, 131
 - domain filter, 130
 - email filter, 130
 - person filter, 131
 - user filter, 131
- backout command, 93–99
 - merge option, 94–96, 98
 - m option, 94
- basename template filter, 130
- bisect command, 102, 104, 105
- bisect extension, 2, 100–105, 137, 163
- branches command, 84
- branches template keyword, 128
- branch command, 84, 85
- bugzilla extension, 115–118, 163
- bugzilla hook, 115, 116
- bundle command, 122
- changegroup hook, 106, 108, 121–123
- chmod system command, 67
- clone command, 12, 17, 62, 63, 71, 81, 82
 - r option, 81
- commit command, 18–20, 29, 39, 47, 50, 84, 109, 112, 148, 150, 152, 173, 175
 - addremove option, 152
 - A option, 50
 - l option, 112
 - u option, 18
- commit hook, 106, 109, 110, 121, 122, 124
- config command, 107
- convert command (conver extension), 9
- convert extension, 9
- conver extension
 - convert command, 9
- copy command, 42, 50–53, 93, 176
 - after option, 53
- cp command, 52
- cp system command, 52
- date template filter, 130
- date template keyword, 128, 130, 131
 - age filter, 130
 - date filter, 130
 - hgdate filter, 130
 - isodate filter, 130, 131
 - rfc822date filter, 131
 - shortdate filter, 131
- desc template keyword, 128, 131
- diffstat command
 - p option, 150

- diffstat system command, 150, 151, 168
- diff command, 17, 18, 20, 140, 143, 160, 163, 165, 166, 169–173
 - exclude option, 169
 - git option, 169–171
 - ignore-all-space option, 171
 - ignore-blank-lines option, 169
 - ignore-space-change option, 170
 - include option, 169
 - nodates option, 169
 - rev option, 170
 - show-function option, 170
 - text option, 169
 - B option, 169
 - C option, 166
 - I option, 169
 - N option, 143
 - X option, 169
 - a option, 169, 172
 - b option, 170
 - c option, 166
 - g option, 170
 - p option, 170
 - r option, 143, 170, 171
 - w option, 171
- diff system command, 136, 138, 143, 166
- domain template filter, 130
- email command (patchbomb extension), 167, 168
 - plain option, 168
 - a option, 168
 - b option, 168
 - d option, 168
 - f option, 168
 - m option, 168
 - n option, 167
 - s option, 168
- email template filter, 130
- escape template filter, 130
- export command, 99
- extdiff command (extdiff extension), 161, 166, 167
 - o option, 166
 - p option, 166
- extdiff extension, 160, 162, 165, 167
 - extdiff command, 161, 166, 167
 - o option, 166
 - p option, 166
- fetch command, 33
- fetch command (fetch extension), 163
- fetch extension, 33, 163
 - fetch command, 163
- file_adds template keyword, 129
- file_dels template keyword, 129
- files template keyword, 128, 130
- fill68 template filter, 130
- fill76 template filter, 130
- filterdiff command
 - files option, 154
 - hunks option, 154
 - i option, 153
 - x option, 153
- filterdiff system command, 150, 151, 153, 154
- firstline template filter, 130
- foo command, 86
- git system command, 61, 170, 171
- grep system command, 103, 105
- guards file, 157
- header template keyword, 135
- heads command, 27
- help command, 11, 12, 138, 173
- hg-interdiff file, 160, 161
- hgdate template filter, 130
- hgext extension, 161
- hgmerge system command, 29, 31, 124
- hgrc configuration file, 66, 70–72, 88, 101, 107, 109, 111, 114–118, 127, 160, 161, 165, 167
- hgrc file
 - acl.allow section, 114
 - acl.deny section, 114
 - acl section, 113
 - bundle entry, 114
 - pull entry, 114
 - push entry, 114
 - serve entry, 113
 - sources entry, 113
 - bugzilla section, 115, 116
 - db entry, 116
 - host entry, 115
 - notify entry, 116
 - password entry, 116
 - usermap entry, 116
 - user entry, 116
 - version entry, 115
- diff section
 - showfunc entry, 170
- extdiff section, 167
- extensions section, 33, 160, 165
- hooks section, 109
- notify section, 118
 - config entry, 118
 - maxdiff entry, 118
 - sources entry, 118
 - strip entry, 118
 - template entry, 118
 - test entry, 118, 119

- ui section
 - username entry, 18
 - verbose entry, 110
- usermap section, 116–118
- web section, 70–72, 117, 119
 - accesslog entry, 71
 - address entry, 71
 - allow_archive entry, 70, 71
 - allowpull entry, 71
 - baseurl entry, 117, 119
 - contact entry, 71
 - description entry, 71
 - errorlog entry, 71
 - ipv6 entry, 72
 - maxchanges entry, 71
 - maxfiles entry, 71
 - motd entry, 71
 - name entry, 71
 - port entry, 72
 - stripes entry, 71
 - style entry, 71
 - templates entry, 71
- hgweb.cgi file, 67–70, 72
- hgweb.config file, 69, 71
- hgwebdir.cgi file, 69–71
- hg system command, 66
- import command, 143, 172
- incoming command, 21, 63, 107, 127
- incoming hook, 106, 115, 121–123
- init command, 150, 173
- inotify extension, 163–165
 - inserve command, 165
- inserve command (inotify extension), 165
- interdiff system command, 160, 162, 166
- isodate template filter, 130, 131
- kdiff3 system command, 30, 34, 166, 167
- locate command, 152
- log command, 13–16, 19, 20, 80, 81, 85, 94, 127, 128, 137, 148
 - patch option, 16
 - rev option, 14, 16, 17
 - template option, 128, 131
 - p option, 16
 - r option, 14, 17
- lsdiff command, 153
- lsdiff system command, 151
- mercurial.localrepo module
 - localrepository class, 112, 120
- mercurial.node module
 - bin function, 120
- mercurial.ui module
 - ui class, 112, 120
- merge command, 28, 33, 39, 40, 78, 87, 148, 171
- merge system command, 31, 32
- mpatch system command, 146
- mq extension, 164, 166
 - qapplied command, 141, 144, 150, 173
 - qcommit command, 150, 173
 - qdelete command, 173
 - qdel command
 - f option, 173
 - qdiff command, 173
 - qfold command, 153, 173
 - e option, 173, 174
 - l option, 174
 - m option, 174
 - qguard command, 156, 157
 - qheader command, 174
 - qimport command, 143, 174
 - qinit command, 138, 139, 149, 173, 174, 177
 - c option, 149, 150, 173, 174, 177
 - qnew command, 139, 140, 142, 143, 174
 - f option, 142, 174
 - m option, 174
 - qnext command, 174
 - qpop command, 141, 142, 147, 148, 175
 - a option, 141, 146, 147, 150, 175
 - f option, 142, 175
 - n option, 147, 175
 - qprev command, 175
 - qpush command, 141, 142, 145–148, 153, 157, 164, 165, 174–177
 - a option, 141, 146, 147, 150, 164, 175
 - l option, 175
 - m option, 147, 175
 - n option, 176
 - qrefresh command, 140, 142, 144, 146–148, 150, 153, 174, 176
 - e option, 176
 - l option, 176
 - m option, 176
 - qrename command, 176
 - qrestore command, 176
 - qsave command, 147, 176
 - c option, 147
 - e option, 147
 - qselect command, 157
 - qseries command, 141, 144, 148, 176
 - qtop command, 150, 176
 - qunapplied command, 176
 - qversion command, 177
- node template keyword, 129
 - short filter, 131
- notify extension, 118, 119, 163

- obfuscate template filter, 131
- outgoing command, 23, 127
- outgoing hook, 106, 107, 122, 123
- pageant system command, 65
- parents command, 22, 28, 29, 39, 171
- parents template keyword, 129
- patchbomb extension, 160, 167, 168
 - email command, 167, 168
 - plain option, 168
 - a option, 168
 - b option, 168
 - d option, 168
 - f option, 168
 - m option, 168
 - n option, 167
 - s option, 168
- patchutils package, 150, 160
- patch command, 143
 - reverse option, 99
 - p option, 143
- patch system command, 99, 136, 138, 142–146, 172
- perl system command, 113
- person template filter, 131
- plink system command, 64, 66
- prechangegroup hook, 106, 121–123
- precommit hook, 106, 112, 121–124
- preoutgoing hook, 106, 108, 122, 123
- pretag hook, 106, 123, 124
- pretxnchangegroup hook, 88, 106, 108, 113, 121–123
- pretxncommit hook, 106, 108, 110, 112, 115, 121–124
- preupdate hook, 106, 124, 125
- pull command, 21–23, 26, 33, 39, 62, 71, 78, 86, 90, 107, 121, 123, 147, 150
 - u option, 22
- push command, 23, 78, 121–123
- puttygen system command, 64
- putty system command, 65
- qapplied command (mq extension), 141, 144, 150, 173
- qcommit command (mq extension), 150, 173
- qdelete command (mq extension), 173
- qdel command (mq extension)
 - f option, 173
- qdiff command (mq extension), 173
- qfold command, 173
- qfold command (mq extension), 153, 173
 - e option, 173, 174
 - l option, 174
 - m option, 174
- qguard command, 157
- qguard command (mq extension), 156, 157
- qheader command (mq extension), 174
- qimport command (mq extension), 143, 174
- qinit command, 150, 173
- qinit command (mq extension), 138, 139, 149, 173, 174, 177
 - c option, 149, 150, 173, 174, 177
- qnew command, 142
- qnew command (mq extension), 139, 140, 142, 143, 174
 - f option, 142, 174
 - m option, 174
- qnext command (mq extension), 174
- qpop command, 142, 146, 147, 150
- qpop command (mq extension), 141, 142, 147, 148, 175
 - a option, 141, 146, 147, 150, 175
 - f option, 142, 175
 - n option, 147, 175
- qprev command (mq extension), 175
- qpush command, 146, 147, 150
- qpush command (mq extension), 141, 142, 145–148, 153, 157, 164, 165, 174–177
 - a option, 141, 146, 147, 150, 164, 175
 - l option, 175
 - m option, 147, 175
 - n option, 176
- qrefresh command (mq extension), 140, 142, 144, 146–148, 150, 153, 174, 176
 - e option, 176
 - l option, 176
 - m option, 176
- qrename command (mq extension), 176
- qrestore command (mq extension), 176
- qsave command, 147
- qsave command (mq extension), 147, 176
 - c option, 147
 - e option, 147
- qselect command (mq extension), 157
- qseries command (mq extension), 141, 144, 148, 176
- qtop command (mq extension), 150, 176
- qunapplied command (mq extension), 176
- qversion command (mq extension), 177
- remove command, 42, 48, 49, 53, 73, 92, 154, 176
 - after option, 49
- rename command, 42, 53, 78, 93, 176
 - after option, 53
- revert command, 50, 55, 56, 91–93, 98, 140
- rev template keyword, 129
- rfc822date template filter, 131
- rollback command, 89, 90, 100
- root command, 74
- sed system command, 17
- series file, 139, 140, 147, 150, 157–159, 173–177
- serve command, 57, 58, 62, 63, 70–72
 - p option, 63
- shortdate template filter, 131

- short template filter, 131
- ssh-add system command, 65
- ssh-agent system command, 65
- ssh-keygen system command, 64
- ssh command
 - C option, 66
- ssh system command, 41, 58, 64–66
- status command, 17, 18, 20, 47–49, 51, 53, 74, 85, 89, 93, 98, 163, 164, 171, 176
 - C option, 51, 53
- status file, 139, 140, 147, 150, 174, 175, 177
- strip command, 147, 177
 - b option, 177
 - f option, 177
 - n option, 177
- strip template filter, 131
- sudo system command, 117
- tabindent template filter, 130, 131
- tabindent template keyword, 131
- tags command, 80, 81
- tags template keyword, 129
- tag command, 59, 79–82
 - f option, 81
 - l option, 82
- tag hook, 106, 123, 124
- tar system command, 70
- tip command, 20, 21, 85, 127, 150
 - p option, 150
- transplant extension, 167
- unbundle command, 121, 123
- update command, 21, 22, 27, 28, 33, 39, 59, 78, 86, 87, 98, 147, 150
 - C option, 86, 87, 147
- update hook, 106, 124
- urlescape template filter, 131
- user template filter, 131
- version command, 11, 66, 171
- vim system command, 167
- wiggle system command, 146
- zip system command, 70

configuration file

- Mercurial.ini (Windows), 64
- hgrc (Linux/Unix), 66, 70–72, 88, 101, 107, 109, 111, 114–118, 127, 160, 161, 165, 167

environment variables

- EMAIL, 18
- HGMERGE, 29, 31
- HGUSER, 18
- HG_NODE, 109, 120
- HG_PARENT1, 120
- HG_PARENT2, 120
- HG_SOURCE, 120
- HG_URL, 120, 121
- PATH, 66
- PYTHONPATH, 66, 68, 111, 178

global options

- debug option, 66, 114
- exclude option, 77
- include option, 76
- quiet option, 17
- verbose option, 12, 16, 17
- I option, 76, 77
- X option, 77
- q option, 17, 75
- v option, 12, 16, 17, 63, 75, 110

hooks

- acl, 113
- bugzilla, 115, 116
- changegroup, 106, 108, 121–123
- commit, 106, 109, 110, 121, 122, 124
- incoming, 106, 115, 121–123
- outgoing, 106, 107, 122, 123
- prechangegroup, 106, 121–123
- precommit, 106, 112, 121–124
- preoutgoing, 106, 108, 122, 123
- pretag, 106, 123, 124
- pretxnchangegroup, 88, 106, 108, 113, 121–123
- pretxncommit, 106, 108, 110, 112, 115, 121–124
- preupdate, 106, 124, 125
- tag, 106, 123, 124
- update, 106, 124

Mercurial bug database

- bug 29, 55
- bug 311, 143, 144

tags

- tip, 121, 123
- special tag names
 - qbase, 148
 - qtip, 148

template filters

- addbreaks, 130
- age, 130
- basename, 130
- date, 130
- domain, 130
- email, 130
- escape, 130
- fill68, 130
- fill76, 130

- firstline, 130
- hgdate, 130
- isodate, 130, 131
- obfuscate, 131
- person, 131
- rfc822date, 131
- shortdate, 131
- short, 131
- strip, 131
- tabindent, 130, 131
- urlescape, 131
- user, 131

template keywords

- author, 128, 130, 131
- branches, 128
- date, 128, 130, 131
- desc, 128, 131
- file_adds, 129
- file_dels, 129
- files, 128, 130
- header, 135
- node, 129
- parents, 129
- rev, 129
- tabindent, 131
- tags, 129