

**ArX**

---

The ArX Revision Control System

---

Copyright © 2001 2002 Thomas Lord, Walter Landry, and the Regents of the University of California

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.1 or any later version published by the Free Software Foundation; with the Invariant Sections being:

(none)

with the Front-Cover Texts being

"visit [www.regexp.com](http://www.regexp.com)",

and with the Back-Cover Texts being

(none).

A copy of the license is included in the section entitled "GNU Free Documentation License".





# 1 Introducing ArX

ArX is a source code management, revision control, and configuration management tool.

## 1.1 Advantages of ArX

What makes ArX better than other revision control systems?

There are minor advantages, and a few major advantages.

The minor advantages are things like: regular, clean, interfaces; small code size; on-line help; features for browsing change sets with a web browser; software tools architecture. Some of these are arguably matters of taste – worth mentioning but not arguing for.

There are at least six major advantages to ArX: a file tree library of revisions, renaming, distributed repositories, robust and easy operation, configurations, and sophisticated branching and merging.

"a file tree library of revisions" means that all revisions can be made accessible in a (space efficient) library of revision trees. That means that you can use ordinary tools like `diff`, `find`, and `grep` to explore past revisions and revisions on other branches.

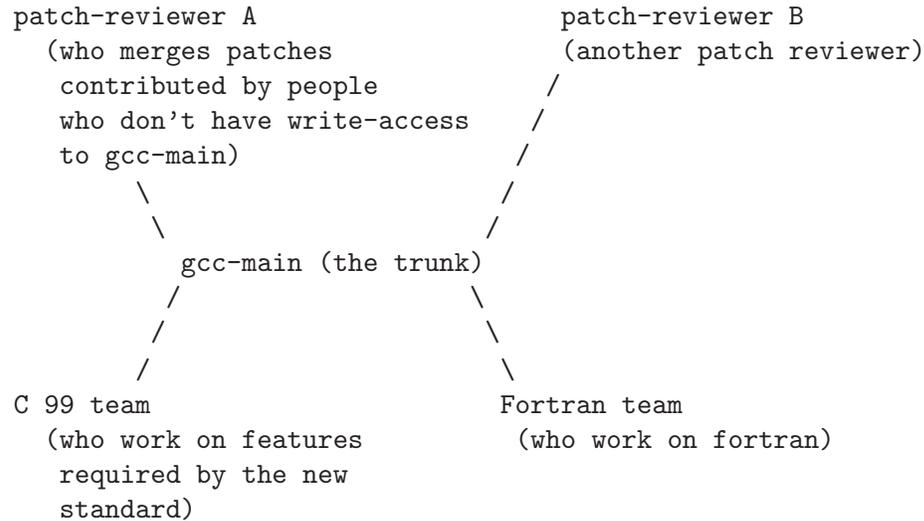
"renaming" means that you can rename files and directories and ArX keeps up just fine: it rearranges the patches between revisions in the corresponding way. And it's convenient: you can use ArX in a mode that doesn't require you to run `add`, `remove`, or `move` commands every time you add, delete, or rename a file or directory.

"robust and easy operation" means that ArX repository transactions have ACID properties and repositories are stored as ordinary unix files. Commits may be concurrent with reads and are atomic. There is no overhead from having to administer a relational or hash-table database. The transient locks held by ArX never become permanently wedged (requiring by-hand repair): they can be broken remotely to recover from interrupted commands.

"Distributed repositories" means that you never need write access to a repository in order to start branches from projects stored in that repository. You can store branches in any repository. That means every programmer can have private repositories for day-to-day work or as scratch areas for working out complex merges. Every sub-group working on a project can have their own repository. If two groups want to fork a project, but still loosely cooperate, they can each have their own trunk in their own repository – selectively merging changes between the two trunks. Without wishing to put too much hype on it, as far as ArX is concerned, there is exactly one repository database for the entire world – we all share it.

"configurations" are a mechanism for defining "meta-projects" that are a combination of multiple "sub-projects". For every meta-project, you need to be able to define *configuration rules*, which explain which revisions from what branches to check out to get a particular revision of the meta-project. ArX provides a flexible mechanism for that, allowing you to pick and choose pieces from various branches and repositories with varying degrees of specificity.

"Sophisticated branching and merging" has to do with ArX's support for asynchronous development on branches, coordinated by a shared trunk. Commonly, people arrange branches in a "star topology". There is a trunk in the middle, surrounded by branches, as in this (hypothetical) example:



And we might imagine many more branches than are actually drawn here.

Ideally, each of the surrounding branches will sometimes be merged into `gcc-main`. And as `gcc-main` changes, it will sometimes be merged back into the surrounding branches.

When two versions each merge with the other repeatedly, forming a merge without creating spurious conflicts is a tricky problem. (For a more detailed description of the problem, see [Chapter 14 \[Star Topology Branching and Merging\]](#), page 71.)

In CVS, the problem is solved by using tags and not generating conflicts for changes that appear to be redundant. Unfortunately, tags are expensive, work only within one repository, and worst of all, are complicated to use correctly. If you forget to create a tag at the right time (or worse, if your tagging operation is interleaved with other operations) you're hosed.

In ArX, history sensitivity is taken to the next level. ArX knows how to handle back-and-forth merging in a star topology automatically.

This is an advance in revision control similar to the invention of lockless operation in CVS. With lockless operation, you have a star topology with an archived development path in the middle, surrounded not by branches, but by working copies. Programmers can commit from the working copies to the trunk or merge from the trunk into working copies repeatedly, and everything works beautifully. The trunk is used to synchronize the multiple working copies.

ArX takes the next step. It allows those satellite working copies to be archived as long-lived branches. The trunk is used to synchronize the surrounding branches.

ArX automates the merge operation (with the `arx star-merge` command) and gives you control over the precedence-ordering of change sets (which branch takes priority, and which has changes that are rejected if they conflict). You don't have to use tags, and you don't have to figure out the right revision arguments to pass to an `update` command.

Each satellite has its own log history. Each can be used for progress tracking. A satellite can be used to maintain an alternative distribution that tracks the trunk, but sometimes takes the lead with additional features.

In addition to star topology merging, ArX provides some other fancy merging options too (e.g. see [Chapter 17 \[Multi-Branch Merging – The reconcile Command\]](#), page 81.)

## 1.2 Global Revision Control Done Right

The foundation of ArX is two commands: `mkpatch` and `dopatch`.

`mkpatch` computes a patch set describing the difference between two trees. `dopatch` applies a patch set to a tree, gracefully handling the cases when a patch doesn't apply cleanly.

Conceptually, `mkpatch` is similar to `diff -r` and `dopatch` is similar to `patch`. Unlike `diff` and `patch`, though, `mkpatch` and `dopatch` can handle the addition or removal of files and directories, the renaming of files and directories, files which are symbolic links, changes to file permissions, and binary files.

An ArX *repository* is a collection of full-source revisions, and patch sets. Brand new trees are represented as full-source revisions. Modified trees are stored as patch sets. Any revision can be reconstructed by retrieving a full-source revision, intermediate patch sets, and applying those patches.

ArX repositories have globally unique names and every revision in a repository has an easy to understand, easy to type name. Putting the two names together, ArX provides a global namespace for revisions.

On the basis of that global namespace, branches and merge operations can span repository boundaries. As far as ArX is concerned, all of the repositories you can access over the Internet are integrated into one gigantic repository – seamlessly integrated.

In ArX, there is no such thing as a "working directory". That is to say, there is no distinction between a tree that you download as a source distribution, and a tree that you check out from a repository. ArX is happy to work with both.

Every tree that ArX works with has a "patch log" – a record of all of the patches (in the global namespace) that have ever been applied to the tree, along with a record of the full-source revision that the tree started from. In ArX, a tree never belongs to a specific branch in some specific repository – instead: a source tree is considered to be a part of every branch in every repository for which it has a patch log. At any time, any tree can join any branch with which it has a common ancestor.

ArX is agile and flexible at handling patch sets. It provides `update` – a patch set manipulation operation that is logically equivalent to the traditional `update` operation of CVS; it provides `replay` – a history-sensitive merge operation equivalent to the `update` operation of Subversion (as documented in the Subversion manual); ArX also provides two rather more sophisticated merge operations – `reconcile` and `star-merge` – which handle very complex (yet quite realistic) branching and merging scenarios gracefully.

In general, the design philosophy of ArX is to be a very good librarian for whole-tree patch sets, and a very good mechanic for manipulating whole-tree patch sets. ArX is designed to stay out of your way while you hack, but come to your aid when you `commit`, `review`, `update`, or `merge`. The aim is simplicity, clarity of function, and flexibility.

## 1.3 Introducing ArX Project Trees

ArX manages "project trees". A project tree is a file system tree, usually containing the source code for a project.

What distinguishes a project tree from an ordinary tree is the presence of "ArX control files", primarily stored in a top-level subdirectory called `{arch}`.

The control files include information needed to keep an inventory of the tree, a "patch log" documenting the history of the tree, various default values for ArX commands applied to the tree, and a local cache of information to speed up some ArX commands.

When you distribute a tree, usually you will *include* all of the ArX control files – they are useful to others.

ArX has no special concept of a "working copy" (in other revision control systems, a working copy is a tree checked out from the revision control database, as contrasted with a tree from any other source). Any tree that contains ArX control files can be used as a "working copy". If you download a distribution for a project managed with ArX and unpack that distribution – you have a working copy.

For more information, see [Chapter 5 \[ArX Project Trees\]](#), page 25.

## 1.4 Introducing ArX Inventories

ArX keeps track of an inventory of the files in a project tree. For example, it can distinguish the files that are officially part of the tree from other files, such as scratch files or editor back-up files. The command:

```
% arx inventory --source
```

prints an inventory of files in a tree.

Every file has two names: its "location" (a path relative to the root of the tree) and its "tag" (a logical, location-independent name for the file). The command:

```
% arx inventory --source --tags
```

prints an inventory of files in a tree, showing the logical name of each file.

ArX uses tags to keep track, for example, of when files are renamed.

For more information, see [Chapter 6 \[ArX Project Inventories\]](#), page 27.

## 1.5 Introducing ArX Patch Sets

Experienced programmers should be familiar with the standard command `diff -c -r`, used to create a (standard) "patch set" describing the changes made between two copies of a tree. And they are familiar with the standard command `patch` – used to modify a tree according to the description of changes in a patch set.

Standard patch sets have limitations: they do not cleanly handle file or directory additions, removals, or renames, symbolic links, file permissions, or binary files.

ArX provides `mkpatch` and `dopatch`, similar to `diff -r` and `patch`, but without the limitations.

For more information, see [Appendix C \[ArX Patch Sets\], page 131](#).

## 1.6 Global Namespaces

ArX implements a global namespace of projects, taking into account the organization publishing a project, branching of projects, and versioning of branches.

ArX implements a global namespace of patch sets, building on the global namespace of projects.

Those namespaces give rise to the idea of a "development path". For example, the very first revision of the project ArX, might be called:

```
lord@regexps.com--ArX/ArX--0.1--base-0
```

The next three revisions might be called:

```
lord@regexps.com--ArX/ArX--0.1--patch-1
lord@regexps.com--ArX/ArX--0.1--patch-2
lord@regexps.com--ArX/ArX--0.1--patch-3
```

Each of those revision names is the name of a patch set that describes what changed in that revision, compared to the previous revision.

If ArX forked into a separate development paths, say "intl" (for internationalizing the code), there might be revisions such as:

```
lord@regexps.com--ArX/ArX--intl--0.1--patch-1
lord@regexps.com--ArX/ArX--intl--0.1--patch-2
```

ArX implements a global namespace of all user's of ArX (layered on top of email addresses). Every patch set has an associated log entry, with a `Creator:` line that contains a user's ArX id.

For more information, see [Chapter 7 \[The ArX Global Name-space of Users\], page 39](#), [Chapter 8 \[The ArX Global Name-space of Projects\], page 41](#), and [Chapter 11 \[Basic Revision Control\], page 51](#).

## 1.7 Introducing ArX Archives

ArX can manage repositories of revisions, storing those revisions as compressed tar files of patch sets or compressed tar files of complete trees.

ArX archives can be used remotely if they are made accessible by an ordinary FTP server, HTTP server with WebDAV, or sftp server. There is no need for a special ArX-specific server.

You can (safely, atomically) add a new revision to a repository (`arx commit`). You can reconstruct an arbitrary revision from the files in a repository (`arx get`). You can "branch" a development path to create a new, related development path (`arx create-branch`, `arx create-version`, and `arx tag`). Branches can cross repository boundaries, and to the user, ArX appears to integrate the two repositories into one.

Repositories can be migrated and, for read-only access, replicated.

ArX repositories have an easy-to-understand format, amendable to browsing by hand or with special-purpose interface programs.

For more information, see [Chapter 9 \[Archives\]](#), page 45.

## 1.8 Introducing ArX Patch Logs

In every project tree, ArX keeps a detailed "patch log": a record of what the original source tree was, along with a record for every patch applied to the tree (regardless of what branch the patch came from).

ArX log entries use RFC822-style formatting. Automatically generated headers record what files were changed by each patch, what repository the patch came from, what user created the patch, etc. The patch log is a very useful source of information for programmers.

ArX can automatically generate a GNU-style ChangeLog from its patch log.

When you "merge" two branches (combine the changes made in those branches), ArX uses the information in the patch log to avoid redundantly applying patches.

For more information, see [Chapter 13 \[Patch Logs and ChangeLogs\]](#), page 67.

## 1.9 Cheap Branching and Smart Merging

Creating a branch in **ArX** is inexpensive in both space and time.

The **ArX** commands for merging have many subtle features to help a merge go smoothly.

For more information, see [Chapter 12 \[Basic Branching and Merging\]](#), page 63, [Chapter 17 \[Multi-Branch Merging – The reconcile Command\]](#), page 81, and [Appendix F \[Idempotent Merging\]](#), page 143.

## 1.10 What Does It All Mean?

Putting all those features together, **ArX** is an elegant and more featureful replacement for older systems like **CVS**.

For example, using **ArX**, every programmer can conveniently have a private repository for day-to-day work rather than burdening a shared repository with per-user branches.

A project can be "multi-homed" – stored in multiple repositories – with different branches in each repository.

It is easy and convenient, when using **ArX**, to improve and clean-up projects by reorganizing the files and directories they contain: something that is quite awkward with **CVS**.

The fancy branching and merging commands of **ArX** make it convenient to do more development than ever in feature-specific branches, merging features as they are completed, generating a clean, complete, and isolated patch set for each new feature.

**ArX** is written in a software-tools style: it is made up of many small programs, each of which does one job well. The commands have very regular option and argument syntax and input and output formats. **ArX** is an excellent foundation for process automation and for layering under fancy graphical interfaces. **ArX** is self-documenting and extensible.



## 2 System Requirements

In order to use **ArX**, there are some software tools that you must already have available. These don't necessarily need to be on your `PATH` – **ArX** can use a separate `PATH` if you need it to.

**GNU Make** You will need **GNU Make** in order to build **ArX**.

**GNU Tar** You must have **GNU tar**. More specifically, you must have a version of `tar` that has options:

<code>-zcf F</code>	to create a gzip-compressed tar file called <code>F</code> , where <code>F</code> may be <code>'-'</code> , meaning to write the tar file to the standard-output stream
<code>-zxf F</code>	to extract files from a gzip-compressed tar file called <code>F</code> , where <code>F</code> may be <code>'-'</code> , meaning to read the tar file from the standard-input stream.
<code>-T -</code>	This option reads a list of files from standard input. Only those files are read or written to the archive -- others are ignored.
<code>-m</code>	When extracting files, don't restore modification times.

**GNU diff and GNU patch** After much deliberation, I've decided to go ahead and rely on the **GNU** versions of `diff` and `patch`. Specifically, you need a version of `diff` that can generate "unified format" output (option `-u`) and a version of `patch` that understands that format and that understand `--posix`. (It would be trivial to use "context diffs" and, thus, standard `diff` and `patch`, however, unified diffs are much easier to read, and I'm hoping that picking specific implementations of these critical sub-components will help contribute to the long-term stability of **ArX**.)

**Standard Posix Shell Tools** The package framework assumes that some standard Posix shell tools are available on your system. At the moment, `sh` must be installed as `/bin/sh`, but this will be corrected in a future release:

```
awk
cat
chmod
date
echo
find
fold
grep
head
ls
```

```

mkdir
printf
pwd
rm
sed
sh
tee
test
touch
tsort
wc
xargs

```

**ordinary -exec extension to find** Your version of `find` must be able to expand `{}` even in the context of a larger string. For example:

```
find . -exec echo ">>>{}<<<" ";"
```

should print a list of all files and directories, surrounded by `'>>>'` and `'<<<'`. GNU `find` has this property as do, I believe, most implementations. (Posix explicitly requires only that `{}` be expanded in isolation, leaving undefined the meaning of `{}` when embedded in a longer string.)

**The null Device** Your system must have `/dev/null`. Output directed to `/dev/null` should simply disappear from the universe, in the usual way.

**The BSD column program – maybe** If the program `column` is on your `PATH` (or the path being used by ArX), it should be a program which formats its input into columns and, with the option `-x`, fills columns before rows. If this program isn't found, it won't be used.

**A C and C++ compiler** To create ArX and related subprojects, you will need a C and a C++ compiler. At this point, the Gnu compilers have been the most tested.

## 3 Tutorial

In this tutorial, we are going to create an archive, put a simple program into it, and then create successive revisions and branches. This uses the bash shell. Different shells may require you to quote special characters differently.

### 3.1 Creating the first revision

We start with creating an archive. First, we have to tell arx who we are

```
% arx my-id 'J. Hacker <jhacker@foobar.org>'
```

You should also tell arx what kind of editor you like to use for editing logs etc. If you already have the \$EDITOR or \$ARCHEDITOR shell variables set, then you can skip this step. To set arx to use emacs, for example, you use the my-editor command

```
% arx my-editor emacs
```

and if you're a vi user, you can type

```
% arx my-editor vi
```

Now we create an archive

```
% arx make-archive jhacker@foobar.org--archive {archive}
```

This creates an archive in the directory {archive}. Normally, you should never need to look in that directory.

Next, we make that archive our default archive

```
% arx my-default-archive jhacker@foobar.org--archive
```

Now we want to make a revision library. You may not always need this, but it is often a nice feature if you have a lot of revisions that you want to look at.

```
% mkdir "{revisions}"
```

Then we tell arx about it

```
% arx my-revision-library "'pwd'/{revisions}"
```

Finally, we can start writing a program. We will use the simplest shell script, Hello world. First, we create a directory to house the project.

```
% mkdir hello
```

Then we create the program

```
% cd hello
% echo "echo Hello, World" > Hello
```

Now we are going to store this masterpiece in arx. We first have to create a log, letting ArX know that we are creating a new project. We will call the project "hello". Since it is the main line of development, we will call the branch "main". Since it is feature complete, we make the version 1.0.

```
% arx create-version hello--main--1.0
```

This opens the log file in the editor that you specified earlier. It also creates an {arch} sub-directory in the current directory. You should never need to look at things in this directory.

You can edit the log file, `++log.hello--main--1.0--jhacker@foobar.org--archive` in this case. You have to be sure to put something in the `Summary:` line. Also, you have to skip a line between headers (like `Summary:` and `Keywords:`) and the body. Otherwise the commit into arx won't work.

Alternately, we can just execute

```
% echo "Summary: initial" \  
> ./'arx create-version --non-interactive hello--main--1.0'
```

The `--non-interactive` option prevents arx from running the editor. Instead it just prints the name of the log file to standard output.

If you wanted to have a multi-line `Summary:`, you can continue by indenting the next line. As an example,

```
Summary: initial version of the foo-bar-baz program,  
        which does foo, bar and baz.
```

Now we have to register our file with arx

```
% arx add Hello
```

Finally, you can now archive the program in arx by running

```
% arx commit
```

## 3.2 Revisions

Now that you've archived the program, you've realized that you don't like the comma in "Hello, World", and you like it to output "Hello World". So you edit Hello, or just run

```
% echo "echo Hello World" > Hello
```

You need to make a log file, so you run either

```
% arx make-log
```

and edit the file `++log.hello--main--1.0--jhacker@foobar.org--archive`, or just run

```
% echo "Summary: Removed a comma" > 'arx make-log --non-interactive'
```

Then you commit the change

```
% arx commit
```

Now you get the idea for a bigger change. You want to make the program say goodbye. So you edit the program again

```
% echo "echo Hello World. Goodbye" > Hello
```

Since this is major change, you decide to bump the version number to 1.1. You do this just by using `create-version` again with a new version number

```
% arx create-version hello--main--1.1
```

and editing the file, or just running

```
% echo "Summary: Added goodbye" \  
> 'arx create-version --non-interactive hello--main--1.1'
```

Then you commit the change for the new version.

```
% arx commit
```

### 3.3 Branches

Now you've decided that you want to internationalize your program. So you want to create a german version. Since this version will have a lot of parts that won't be shared with the english version, you decide to make a branch. Your german is pretty rusty, and you've forgotten how to spell Auf Wiedersehen, so you decide to branch from the older, simpler version without goodbye. The simplest way to do this is to go into the parent directory and create a branch

```
% cd ..
% arx create-branch hello--main--1.0 hello--german--1.0 \
hello--german
```

This creates a directory hello-german and pops up an editor to edit the log file. You can also do it non-interactively

```
% echo "Summary: German version" \
> 'arx create-branch hello--main--1.0 hello--german--1.0 \
hello--german'
```

Now change into the directory and commit the new branch

```
% cd hello--german
% arx commit
```

Now you can edit the program, create a log, and commit a change as normal.

```
% echo "echo Guten Tag" > Hello
% echo "Summary: Translated" > 'arx make-log --non-interactive'
% arx commit
```

### 3.4 Wrapping up

Congratulations. You have now created a project, made revisions, changed the version number, and created a branch. You can view what versions are in the archive with

```
% arx versions
```

It should echo

```
hello--german--1.0
```

It only looked at the german line of development because your current directory (project tree) version is set to hello-german-1.0. We can look at the main line of development by specifying it explicitly.

```
% arx versions hello--main
```

and it should echo

```
hello--main--1.0
hello--main--1.1
```

We can see what revisions were made to hello-main-1.0

```
% arx revisions hello--main--1.0
base-0
patch-1
```

We can see what branches have been made

```
% arx branches
hello--german
hello--main
```

Finally, we can see what projects we have

```
% arx categories
hello
```

### 3.5 Revision Trees

Now suppose you have forgotten what things looked like in `hello-main-1.0`. A simple way to view them is with revision trees. It is a complete version of the source tree for a particular revision. For successive revisions, arx uses hard links to prevent unnecessary copying. To create a revision for `hello-main-1.0-base-0`, it is simply

```
% arx library-add hello--main--1.0--base-0
```

To view the original version of the file, then is just

```
% more ../{revisions\}/jhacker@foobar.org--archive/\
hello/hello--main/hello--main--1.0/\
hello--main--1.0--base-0/Hello
```

## 4 ArX Commands in General

Every command in ArX is accessed via the program ArX, using an ordinary sub-command syntax:

```
% arx <sub-command> <options> <parameters>
```

A list of all sub-commands can be obtained from:

```
% arx --help-commands
```

The most complete documentation for each command is available via its help message. The help messages in ArX are roughly comparable to traditional unix `man` pages. For example, try:

```
% arx make-archive --help
create a new archive directory
usage: make-archive [options] (-u | name) directory

-V --version          print version info
-h --help            display help

-r --readme file     save FILE as the README
                    for this archive

-u --update          update an existing archive
```

NAME is the global name for the archive. It must be an email address with a fully qualified domain name, optionally followed by "--" and a string of letters, digits, periods and dashes (but not two dashes in a row).

Normally the archive directory must not already exist, but there is an exception:

If `-u` or `--update` is specified, then attempt to bring an existing archive up to date with respect to this version of arx and update meta-info, such as the README file.

If `-u` or `--update` is specified, the archive name must not be specified.

There is a great deal of regularity among commands regarding option names and parameter syntax. Hopefully, you'll pick this up as you learn the various commands.

## 4.1 The ArX Commands

Here is a synopsis of all ArX commands: the output of `arx --help-commands`:

```

arx sub-commands
-----

* User Commands

    my-id : print or change your id
my-default-archive : print or change your default archive
    gui-diff : print or change path to your GUI-based diff tool

    my-browser : print or change your default arch browser
    my-editor : print or change your default arch editor
    my-guidiff : print or change your default arch guidiff

register-archive : record the location of an archive
whereis-archive : print the registered location of an archive
archives : report registered archives and their locations

* Project Tree Commands

    tree-version : print the default version for a project tree

* Project Tree Inventory Commands

    inventory : inventory a source tree

tagging-method : print or change a project tree tagging method
tree-lint : audit a source tree

    add : add an explicit inventory tag
    delete : remove an explicit inventory tag
    move : move an explicit inventory tag

explicit-default : set the default explicit tag for a directory

```

\* Patch Set Commands

mkpatch : compare two source trees and create a patch tree  
mkpatch-files : compare select files in two source trees and create a patch  
dopatch : apply a patch tree to a source tree

patch-report : generate a report from a patch set

\* Archive Commands

make-archive : create a new archive directory

categories : list the categories in an archive  
branches : list the branches in an archive category  
versions : list the versions in an archive branch  
revisions : list the revisions in an archive version

cat-archive-log : print the contents of an archive log entry

archive-cache-revision : cache a full source tree in an archive  
archive-cached-revisions : list full source trees cached in an archive  
archive-uncache-revision : remove a cached revision from an archive

category-readme : print the =README of a category  
branch-readme : print the =README of a branch  
version-readme : print the =README of a version

\* Patch Log Commands

```
make-log : initialize a new log file entry

logs : list patch logs in a project tree
add-log : add a version patch log to a project tree
remove-log : remove a version patch log from a project tree

log-ls : print version patches in a project tree
cat-log : print the contents of a project tree log entry
log-header-field : filter a header field from a log entry

changelog : generate a change log from a patch log
log-for-merge : generate a log entry body for a merge

merge-points : report where two branches have been merged
new-on-branch : list tree patches new to a branch
```

\* Archive Transaction Commands

```
commit : archive a revision
get : construct a project tree for a revision
push-mirror : update a read-only repository mirror
```

\* Multi-project Configuration Commands

```
build-config : instantiate a multi-project tree
update-config : update a multi-project tree
replay-config : replay a multi-project tree

record-config : record a revision-specific configuration
show-config : show the revision frontier of a configuration

config-history : report the history of a configuration
```

\* Commands for Branching and Merging

```
update : merge local changes with the latest revision
replay : merge the latest revision with local changes
delta-patch : compute and apply an arbitrary patch
star-merge : merge of mutually merged branches
make-sync-tree : prepare a tree that will synchronize branches

tag : create a continuation revision (tag or branch)

create-branch : create a project tree for a new branch
create-version : make the current tree into a new version

join-branch : join-branch a sibling branch

whats-missing : print patches missing from a project tree
reconcile : filter which plans a replay-based multi-branch merge
```

\* Local Cache Commands

```
what-changed : compare project tree to cached pristine
file-diffs : compare file with cached pristine revision
file-undo : undo file changes from cached pristine revision

undo : temporarily undo changes in working directory
redo : redo undone changes in working directory

add-pristine : locally cache a pristine revision
delete-pristine : remove a pristine trees from a project tree
pristines : list the pristine trees in a project tree
```

\* Revision Library Commands

```
my-revision-library : print or set your revision library path

    library-find : find a revision in a revision library
    library-add : add a revision to a revision library
    library-remove : remove a revision from a revision library

    library-archives : list the archives in the revision library
library-categories : list categories in the revision library
    library-branches : list branches in the revision library
    library-versions : list versions in the revision library
    library-revisions : list revisions in the revision library

    library-log : print a log message from the library
    library-file : find a file in a revision library

touched-files-prereqs : print prereqs of a revision
```

\* Web Related Commands

```
    patch-set-web : create or update a patch-set web
update-distributions : build or update an FTP area

distribution-name : revision name -> distribution name
```

\* Notification Commands

```
        notify : trigger actions for changes to an archive
        my-notifier : print or set your default notify directory

mail-new-categories : send email notices about new categories
  mail-new-branches : send email notices about new branches
  mail-new-versions : send email notices about new versions
  mail-new-revisions : send email notices about new revisions

        notify-library : add new revisions to the library
        notify-browser : add new revisions to the browser

push-new-revisions : send email notices about new revisions

        sendmail-mailx : send email with sendmail ala POSIX mailx
```

For help with a command, try: arx command --help



## 5 ArX Project Trees

One of the central organizing concepts of **ArX** is the project tree. For the most part, a project tree is an ordinary tree of files, directories and symbolic links – all files that you create and maintain as part of your project.

In addition to the tree content you create, **ArX** itself maintains some control files in a project tree. You should regard these extra control files as part of your project's content. For example, if you distribute source for a program managed by **ArX**, you will ordinarily want to *include* the control files: they are useful to other people.

### 5.1 Initializing a Project Tree

To initialize a project tree for the first time, from the root of the tree, use the command:

```
% arx create-version (project name)
```

That will create a subdirectory, at the root of the tree, called `{arch}`. It will also open an editor for entering your initial log file. Most **ArX** control files are stored in the tree rooted at `{arch}`. You should never create, remove, or modify files there by hand.

When you have multiple project trees for related projects it is good practice to make them sibling directories. This is because **ArX** sometimes caches information in project trees and those caches can speed up some operations. When looking for information in a cache, **ArX** looks not only in the current project tree, but in any sibling project trees. For example, while working on **ArX**, I might have several copies of **ArX**, each for working on a different set of features:

```
% cd ~/wd
% ls
ArX      ArX-branches  ArX-inventory  ArX-reporting
```

In this manual, `~/wd` always refers to a "directory of project trees". (The convention means the same thing as, but is less cumbersome than `#{PROJECT_TREES}` – there is no requirement that your directory of project trees be called `~/wd`, or that you only have one such directory.)



## 6 ArX Project Inventories

In a project tree, some of the files and directories are "part of the source" – they are of interest to ArX. Other files and directories may be scratch files, editor back-up files, and temporary or intermediate files generated by programs. Those other files should be ignored by most ArX commands.

This chapter discusses how ArX recognizes which files to pay attention to, and which to ignore.

ArX has flexible facilities for keeping track of all of the files and directories in your project: for taking "inventories" of your project tree. It has these facilities for three reasons:

**Distinguishing Source** ArX uses a project inventory to distinguish files and directories which are part of your project from other files and directories which are temporary files, scratch files, editor backup-files, and so forth.

Additionally, ArX permits you to overlay projects: store more than one project at a single root. When you do that, ArX uses inventories to sort out which files and directories belong to each project. (The topic of overlays, however, is deferred until a later chapter.)

**Recognizing Renames** Every file or directory in an ArX inventory has two names. One name is simply the location (path) of the file relative to the root of the project tree. The other name is a "logical name" for the file: a name that remains the same regardless of where in the project tree the file is located. When ArX compares two versions of a project tree, it uses logical names to discover when files or directories have been moved, renamed, deleted, or added.

### 6.1 Choices Regarding Inventories

For each project tree, you have a choice to make regarding how project inventories work. The options are described briefly here, then in more detail in the sections that follow.

**Explicit Inventories** The most familiar (and default) option is to use an explicit inventory. Whenever you add, delete or rename a file, you must inform ArX of that fact explicitly. For example, after adding the file `foo.c`, you have to tell ArX

```
% arx add foo.c
```

If you rename `foo.c` to `bar.c`, then you must use ArX to do so

```
% arx move foo.c bar.c
```

Finally, if you delete `foo.c`, you must do so with ArX

```
% arx delete foo.c
```

**Naming Conventions** Another option is to simply use naming conventions. ArX will search your tree for files matching certain naming patterns, and consider all of those files to be source files.

When you use only naming conventions to take an inventory, the logical name of a file and its location name are exactly the same. For that reason, if you rename a file, ArX will think you deleted a file with the old name, and added a file with the new name. If you delete a file, then add a file with the same name, ArX will think that the new file is a modified form of the old file. None of these limitations are fatal. ArX will still work, but they do

limit the effectiveness of ArX at branching and merging. ("Branching" and "merging" are topics of a later chapter.)

**Implicit Inventories** A third option combines some of the advantages of using naming conventions with some of the advantages of explicit inventories: implicit inventories. When you use an implicit inventory, every file that passes the naming conventions is considered source. You **may** explicitly add, delete, and rename files – allowing ArX to precisely track renames for those files and directories. You also **may** store a file tag (the "logical name" of a file) *in* any file. If you don't explicitly tag a file, and use an implicit inventory, ArX will search for those embedded tags and use them to precisely detect new files, deleted files, and renamed files.

Each of the three options is called a tagging method.

There is some advice at the end of this chapter about how to choose among the three tagging methods.

## 6.2 Specifying a Tagging Method

If you never explicitly specify a tagging method, ArX will use explicit inventories. You choose it specifically with this command issued in a project tree:

```
% arx tagging-method explicit
```

Similarly, to use either naming conventions or implicit inventories, use one of the commands:

```
% arx tagging-method names
% arx tagging-method implicit
```

To find out what method a given project tree uses, use the same command with no argument:

```
% arx tagging-method
names
```

## 6.3 The inventory Command

The command `arx inventory` is used to print a list of source files. It has many options, including options to print other kinds of file lists (such as a list of all editor backup files, or a list of all files which are not source):

```
% cd source-tree

% arx inventory --source
hello.c
hello.h
library
library/buffer.c
library/buffer.h
...
```

contrasted with:

```
% cd source-tree

% ls
hello.c  hello.c.~1~  hello.h  library
```

(Notice that `hello.c.~1~` is not included in the inventory of source files.)

The naming conventions used by ArX are as follows:

**Control Files** A control file *is* part of the source, but control files are not included in the output of `arx inventory` unless the `--all` flag is used. Control file and directory names match any of these patterns:

```
.arch-project-tree
.arch-ids
.owned.*
.common
{arch}
```

**Junk Files** A junk file is *not* part of the source. A junk file or directory name matches the pattern:

```
,*
```

or if it contains any of the characters:

```
<space>
<tab>
<newline>
[
]

?
\
```

Note that if a directory name matches that pattern, then none of the contents of the directory are part of the source, regardless of their names.

Junk files are listed by the command:

```
% arx inventory --junk
```

ArX sometimes creates junk files and directories of its own. When it does, those files and directories have names that match the pattern:

```
,,*
```

You should avoid creating files and directories with names that match that pattern. ArX will freely delete files and directories with names that match `,,*` whenever it needs to re-use such a name.

Usually, ArX will delete any junk file it creates before the command that created the junk file terminates. Sometimes, though, when a command fails, ArX will leave behind junk files or directories matching `,,*`. This is a debugging feature, likely to be removed in a future

release. For now, whenever you find such a file (and are confident it isn't being used by a currently running command), you are free to delete it.

**Backup Files** If a file is not a junk file, it may be a backup file. Backup files are not part of the source. They match any of the patterns:

```
*~
*.bak
*.modified
*.orig
*.original
*.rej
*.rejects
```

Backup files are listed by the command:

```
% arx inventory --backups
```

**Precious Files** If a file is not a control file, junk file, or backup file, it might be a precious file. Precious files are not part of the source. For all intents and purposes, ArX treats precious, backup, and unrecognized files the same.

Precious files and directories match one of these patterns:

```
++
.gdbinit
=build*
=install*
CVS
RCS
TAGS
```

Of course, precious files can be listed by the command:

```
% arx inventory --precious
```

Sometimes ArX will create its own precious files – usually to save some information that you might not want to lose. When it does, it creates a file or directory matching the pattern:

```
+++
```

You should avoid creating such filenames yourself. ArX won't ever delete such a file – but if one happens to get in the way of an ArX command, that command will fail with an error.

**Source Files** If a file is not a control or junk file, it might be an ordinary source file. Source files are, of course, the files that ArX stores in an archive (along with control files).

Source files must match the pattern:

```
[=a-zA-Z0-9]*
```

but must *not* match any of the patterns:

```
*.o
*.core
core
```

Ordinary source files are listed by:

```
% arx inventory --source
```

Some files which are ArX control files are counted as source even though they don't match the patterns above. However, these files are not listed by default. All source files (ordinary source plus control files) are listed by:

```
% arx inventory --source --all
```

**Unrecognized Files** Any file that doesn't fall into the above categories is an unrecognized file. Unrecognized files can be listed by the command:

```
% arx inventory --unrecognized
```

**WARNING** The basic pattern for source files is:

```
[=a-zA-Z]*
```

however, you should restrict yourself to file names that do not contain spaces. Filenames containing spaces are likely to trigger bugs in the current release of ArX.

## 6.4 Using an Explicit Inventory

Explicit inventories are the default, but if you want to set it anyway, then use this command:

```
% arx tagging-method explicit
```

Note that you must use that command from within a working directory tree that has already been initialized.

When using explicit designation, it is (ordinarily) necessary to add every file and directory in the source to the explicit list using the command:

```
% arx add FILE
```

If FILE is a directory, that will create FILE/.arch\_ids/=id. If it is a regular file or symbolic link, it will create (in the same directory) .arch\_ids/FILE.id. In either case, the file created will contain an obscure string known as an "inventory tag" (inventory tags are explained in more detail below).

If you remove a regular file or symbolic link, you must use the command:

```
% arx delete FILE
```

It will remove FILE and its inventory tag.

In order to remove a directory, you must yourself remove the .arch\_ids subdirectory. That will also implicitly remove the inventory tags of any files that ArX thinks are stored in that directory.

If you rename a regular file or symbolic link, you can use the command:

```
% arx move OLD-NAME NEW-NAME
```

to move the file and its inventory tag.

If you rename a directory, its inventory tag (and the tags for all files and subdirectories it contains) move with it automatically (because the `.arch_ids` subdirectory has moved).

When you run `arx inventory` in a working directory using explicit designation, only explicitly designated source files are listed. If you would rather see a list of all files passing the naming conventions for source files, use:

```
% arx inventory --source --names
```

If you are importing a project into ArX, it may be convenient to add everything that matches the naming conventions. An idiom for doing this is

```
$ arx inventory --names --source --both | xargs arx add
```

Then you can clean up by explicitly adding or deleting files.

You should also read about `tree-lint` later in this chapter.

## 6.5 Using an Implicit Inventory

To use implicit tagging, use the following command in your working directory:

```
% arx tagging-method implicit
```

When implicit tagging is used, every file that passes the naming conventions is treated as source. If a file or directory has an explicit tag (created with `add`), ArX will use that explicit tag to recognize when a file has moved. If a file (but not a directory or symbolic link) lacks an explicit tag, ArX will look for a tag in the file itself.

A tag within a file has one of two forms. It may be either:

```
<punct><basename><spaces>-<spaces><tag>
```

where `<punct>` is an arbitrary string of punctuation and spaces, `<basename>` is the basename of the file, and `<tag>` an inventory tag for the file. Or:

```
<punct>tag:<spaces><tag>
```

In either case, `<tag>` should be unique among the files within a directory. A tag within a file must occur within the first 1024 bytes of the file.

One convention for source files is to add a comment to the top of every file, briefly stating the purpose of the file:

```
/* hello.c - 'main' for the hello world program
...

```

or:

```
/* tag: 'main' for the hello world program
...

```

This may cause problems if you rename the file. You'll probably want to change the name and/or the description, but ArX will think that the old file was deleted and a new one created.

Another possible convention is to use a string identifying the author and the time the file was first created (or first tagged):

```
/* tag: joe.hacker@gnu.org Thu Nov 29 17:25:15 PST 2001
...
```

If you use the `basename` form of an implicit tag, and actually rename a file (rather than simply move it between directories), you do need to remember to update the tag line to reflect the new `basename`.

When you use implicit tagging, it is ok if a file lacks any tag at all, either explicit or implicit. In that case, if you rename the file, `ArX` will think you've deleted the old file and added a new one – but aside from that, everything will work normally.

**CAUTION:** Leading and trailing spaces around an inventory tag are not considered part of the tag. Within a tag, every non-graphical character is replaced by `_`. For example, you write the that tag:

```
'main' for the hello    world program
```

the actual inventory tag is:

```
'main'_for_the_hello____world_program
```

It is possible that a future release of `ArX` will slightly change the rule – so that multiple spaces and tabs are replaced by a single `_`.

## 6.6 Recognizing Renames – Inventory Tags

If you are using naming conventions only to recognize source files, then if you rename a directory or file, `ArX` will conclude that you have deleted the old file, and created a new file.

If you are using an explicit source inventory, `ArX` will always recognize when a directory is renamed (presuming that the `.arch_ids` subdirectory is preserved), and it will recognize when a file is renamed if you use `move` (rather than `delete` and `add`). Of course, `ArX` can be fooled if you swap two files without swapping their inventory tags.

If you are using an implicit inventory, `ArX` will never recognize when an untagged file is renamed (it will think "delete" and "add"). If a file is tagged explicitly, `ArX` will recognize when the file is added, deleted, or renamed – just as when using an explicit inventory. If a file is not tagged explicitly, but has an embedded tag, `ArX` will recognize when the file is added, deleted or moved.

## 6.7 Keeping Things Neat and Tidy

The command:

```
% arx tree-lint
```

is useful for keeping things neat and tidy.

If you use explicit tagging, it will tell you of any tags for which the corresponding file does not exist. It will tell you of any files that pass the naming conventions, but for which no explicit tag exists.

If you use implicit tagging, it will tell you of any files for which no tag can be found – either explicit or implicit. It will tell you of any explicit tags for which the corresponding file does not exist.

In either case, or if you are using naming conventions only, `tree-lint` will tell you of any files that don't fit the naming conventions at all.

Finally, if you use explicit or implicit tagging, `tree-lint` will check for cases where multiple files use the same tag. If any two files do have the same tag, you **must** correct that, either by editing the tag (if it is in the file itself) or by using `delete` and `add` to replace a duplicated explicit tag.

## 6.8 The Inventory Tag Abstraction in Detail

When `ArX` considers the files and directories in a working directory it builds a one-to-one index mapping path names (relative to the root of the working directory tree) to inventory tags.

The inventory tag of a file is its "logical identity". The path is the position of that identity within the particular working dir.

You can see the inventory tag for each source file with the command:

```
% arx inventory --source --tags
```

When `ArX` compares two project trees, it bases the comparison on logical identities. If both trees have a file with a particular inventory tag, but the files are in different positions, then `ArX` considers the file to have been moved or renamed. Similarly, if an inventory tag is present in one tree, but missing in the other, then `ArX` considers the file to have been added or deleted.

If you use naming conventions only, the inventory tag of each file is the same as its path. Thus, when using the `names` tagging method, `ArX` never recognizes that a file has been moved or renamed.

When you use the `explicit` tagging method, inventory tags are stored in the `.arch-ids` directories. There is a file in `.arch-ids` for each tagged file (and one file for the directory containing `.arch-ids`), and those files contain the tags.

When you use the `implicit` tagging method, tags in `.arch-ids` directories take precedence (if they exist). If a file is not explicitly tagged, `ArX` searches for the inventory tag in

the file itself (as described earlier in the chapter). Finally, if a file is not tagged at all, then its path is used as the inventory tag.

## 6.9 A Warning About Changing Tagging Methods

Be cautious when changing tagging methods for directories already checked-in to an ArX revision control archive.

For example, if you change from the tagging method `names` to `explicit`, then the inventory tag for every file will change. ArX will think that you've deleted all of the files in the old tree, and added all of the files in the new tree.

## 6.10 Other Ways to Tag Files

In some situations, it isn't convenient to explicitly tag every file or to add an implicit tag to every file.

You can supply a default tag for every file that doesn't have an explicit tag with the command:

```
% arx explicit-default TAG-PREFIX
```

After that, every file in that directory which lacks an explicit tag will have the tag:

```
TAG-PREFIX__BASENAME
```

where `BASENAME` is the basename of the file. Default tags created in this way take precedence over implicit tags embedded in files. You can find out the default tag for a directory with:

```
% arx explicit-default  
TAG-PREFIX
```

and remove the default with:

```
% arx explicit-default --delete
```

You can also specify a default tag which has *lower* precedence than implicit tags:

```
% arx explicit-default --weak TAG-PREFIX
```

and view that default:

```
% arx explicit-default --weak
```

or delete it:

```
% arx explicit-default --weak --delete
```

## 6.11 Telling tree-lint to Shut Up

When using implicit tags, you may sometimes have a directory with many files that have no tag (either explicit or implicit), but not want those files to appear in a report of untagged files generated by `tree-lint`. There are two ways to tell `tree-lint` to shut-up about such files:

One is to provide a default explicit tag or weak default explicit tag using `arx explicit-default`, as described above.

The second method is to label the directory as "don't care" directory – which means that `tree-lint` shouldn't complain about untagged files. You can do that with:

```
% arx explicit-default --dont-care set
```

or remove the "don't care" flag with:

```
% arx explicit-default --delete --dont-care
```

You can find out whether the "don't care" flag is set in a given directory with:

```
% arx explicit-default --dont-care
```

## 6.12 Which Tagging Method Should You Use?

Given the choice of the `names`, `explicit`, and `implicit` tagging conventions, which one should you choose?

The `explicit` method is the default. It requires manually informing ArX that a particular file should be under version control. Both `names` and `implicit` try to guess what kind of files should be archived, and which shouldn't. Unless you are very careful, and, for example, don't include any generated files in your source directory, `names` and `implicit` will accidentally add unwanted files to your archive.

The `names` method is best for project trees that you don't control, and for which the maintainer does not include file tags (either explicit or implicit). For such trees, the `names` method will always work, but if you want to use the `explicit` or `implicit` method, you'll have to add file tags yourself. It also works reasonably well for scripts (such as perl, python, or shell), because there are no object files that can be accidentally included in the archive.

The `implicit` method is, for some, the most convenient. You just get in the habit of adding a `tag:` line to the bottom of each new file and doing a single `arx add` for each directory. After those steps, you can rename files and directories freely – without having to remember to tell ArX in a separate command.

On the other hand, the `implicit` method has two limitations. One limitation is that you must accept the possibility of accidentally adding new files to the inventory. Any file you create that passes the naming conventions counts as source. The other, closely related, limitation is that if you use `implicit` inventories, you will **never** want to compile a program in its own source directory. When you compile a program, that creates intermediate files

and executables. Many of those files will almost certainly pass the naming conventions for source – so ArX will wrongly include them in a source inventory. You might want to include a safeguard in your `configure` scripts that causes them to refuse to compile my programs in the source tree.

## 6.13 Altering the Naming Conventions

The file `{arch}/=tagging-method` defines the naming conventions used for a particular project tree. By editing that file, you can establish naming conventions that are different from the defaults, which are described above.

That file can contain blank lines and comments (lines beginning with `#`) and directives, one per line. The permissible directives are:

```
implicit
explicit
names
    specify the tagging method to use for this tree

exclude RE
junk RE
backup RE
precious RE
unrecognized RE
source RE
    specify a regular expression to use for the indicated
    category of files.
```

Regular expressions are specified in Posix ERE syntax (the same syntax used by `egrep`, `grep -E`, and `awk`) and have default values which implement the naming conventions described above.

The `exclude` pattern should match a subset of files matched by the `source` pattern. Files which match `exclude` are printed by:

```
% arx inventory --source --control
```

but not printed by:

```
% arx inventory --source
```

Although you can define your own naming conventions, there are some minor limitations:

The file names `.` and `..` are always ignored by `inventory`.

File names which contain non-printing characters, spaces, or any of the globbing characters (`*`, `[`, `]`, `\`, `?`) are always placed in the category `unrecognized`. This is so that tools which operate on project trees can safely presume that no source file has a name that includes these characters.

File names which begin with `,,` are always placed in the category `junk`. This is so that tools which operate on a project tree can safely destroy or create files beginning with `,,`

The default naming conventions are given by:

```
exclude ^(.arch-ids|\{arch\})$
junk ^(.*)$
backup ^.*(~|\.[0-9]+~|\.bak|\.orig|\.rej|\.original|\.modified|\.reject)$
precious ^(\+.*|\.gdbinit|=build\.*|=install\.*|CVS|CVS\.adm|RCS|RCSLOG|SCCS|TA
unrecognized ^(.*\.(o|a|so|core)|core)$
source ^([_a-zA-Z0-9]*|\.arch-ids|\{arch\}|\.arch-project-tree)$
```

## 7 The ArX Global Name-space of Users

For various purposes (such as labelling the author of log messages), ArX maintains a global name-space of users. Every user of ArX has an associated user ID, which is (ideally) globally unique.

An id string has two parts: a free-form part, and a unique-id part. The unique-id part is an email address with a fully-qualified domain name. That part of your id string should be unique to you in the world. Here is an example:

```

      Joe Hacker <joe.hacker@gnu.org>
        ^           ^
        |           |
    free form part  unique id
  
```

The free-form part must match the regexp:

```
[[[:alnum:][:space:][:punct:]]]*
```

and the unique-id part must match the regexp:

```
<[-.[:alnum:]]+@[-.[:alnum:]]+\.[-.[:alnum:]]+>
```

You should only need to set your ID once, which you can do with the command:

```
% arx my-id ID-STRING
```

You can check your id with:

```
% arx my-id
Joe Hacker <joe.hacker@gnu.org>
```

or:

```
% arx my-id --uid
joe.hacker@gnu.org
```

Clearly it is a good idea to use a real email address (belonging to you) for the id string, but there is nothing that requires this.



## 8 The ArX Global Name-space of Projects

Another central organizing concept of ArX is a global namespace of projects.

This chapter describes how projects are named, and how the names are applied to project trees.

### 8.1 The Structure of Project Names

The ArX project name-space is designed to reflect the intricacies of how projects evolve in the real world. For example, two or more different organizations may separately develop and distribute a given project. A project can split into multiple development paths. Each path typically evolves through a series of versions. ArX provides a way to precisely name all of these different instances of a given project.

An ArX project name has four basic parts, introduced here, and explained in detail below.

**An Archive Name** The archive name identifies the organization that distributes particular versions of a project. Some example archive names might be:

```
hackerlab@regexps.com--hackerlab
hurd@gnu.org--hurd-team
```

**The Category Name** The category name is a generic name for the project. It is what people usually think of as a "project name". Some example category names are:

```
ArX
gcc
rx
```

**The Branch Label** When a project splits into multiple development paths (even if only temporarily – as a convenience to the maintainers), each path is given a branch label. Some example branch labels are:

```
development
experimental
```

**The Version Number** ArX uses fairly simple version numbers, consisting of a major and minor version number:

```
1.2
2.0
```

The four parts of a project name fit together this way:

```
<archive-name>/<category>[--<branch-label>]--<version-number>
```

As in the example:

```
hackerlab@regexps.com--hackerlab/ArX--development--2.0
      ^             ^             ^             ^
      |             |             |             |
archive name    category    branch label    version
```

Notice that the archive name is separated from the rest of the project name by /. The other parts of the project name are separated by --.

When you use ArX, you often abbreviate project names. For example, you can leave off the archive name and a default archive will be presumed:

```
ArX--development--2.0
```

or you can leave off the version number and, depending on context, that means either "the latest version" or "all versions":

```
ArX--development
```

If you have one branch which is "primary branch", you can leave out the branch label:

```
ArX--2.0
hackerlab@regexps.com--hackerlab/ArX--2.0
ArX
hackerlab@regexps.com--hackerlab/ArX
```

Those last two names are also sometimes used to mean "every branch of ArX" or "every branch of ArX at the hackerlab archive", respectively.

## 8.2 Archive Names

An archive name designates an organization that develops and/or distributes software. Archive names should be globally unique.

Later in the manual, we'll see that archive names are specifically used to identify revision control archives.

An archive name consists of an email address (with complete hostname), followed by -, followed by an additional string of numbers, letters and dashes. Choose an email address which is exclusively yours (or your project's). That way, your archive name(s) will be globally unique. Here is an example:

```
joe.hacker@gnu.org--test-archive
```

If your organization is going to have more than one revision control archive, you'll use more than one archive name:

```
joe.hacker@gnu.org--gcc-archive
joe.hacker@gnu.org--guile-archive
joe.hacker@gnu.org--2001
```

You can choose an archive name to use as the default for all ArX commands. When you run a command without explicitly specifying a archive, the default is used:

```
% arx my-default-archive joe.hacker@gnu.org--2001
```

To find out the current default:

```
% arx my-default-archive
joe.hacker@gnu.org--2001
```

In general, ArX sub-commands accept the option `-A` to specify a non-default archive:

```
% arx my-default-archive -A joe.hacker@gnu.org--test-archive
joe.hacker@gnu.org--test-archive
```

### 8.3 Category Names and Branch Labels

The category of a project name identifies, generally, what the project is. The category is the same no matter who is distributing the project, or which development path is being considered.

The branch label is optional. It can be used to distinguish alternative development paths for a given project.

The category and branch label must match the regexp:

```
[[[:alpha:]]([[:alnum:]]*(-[[:alnum:]]+)?)*
```

or in other words, they must begin with a letter and consist entirely of digits, letters, and dashes – but must not contain two dashes in a row, and must not end with a dash.

### 8.4 Version Numbers

In a somewhat arbitrary but extremely traditional way, branches are divided into a series of versions.

The version number must match the regexp:

```
[[[:digit:]]+\.[[:digit:]]+
```

or in other words, it must consist of two strings of digits, separated by a single period. The first string of digits is called the major version number and the second string of digits is called the minor version number.

**Note:** *Version numbers* are not *revision numbers*. In other words, when the ArX revision control system stores multiple snapshots of the development of your project, it does *not* assign a new version number to each snapshot. Instead, each project version is further subdivided into something called "patch levels", which are explained in detail later in the manual.

## 8.5 Labelling Project Trees

Every project tree may be labeled with a project name using the `set-tree-version` command, as in this example:

```
% cd ~/wd/ArX
% arx set-tree-version ArX--1.0
```

That project name becomes the default for ArX operations within that tree.

You can find out the project name of a tree with `tree-version`:

```
% arx tree-version
joe.hacker@gnu.org--2001/ArX--1.0
```

Notice that ArX used the default archive (returned by `my-default-archive`) when `set-tree-version` was invoked. You can also specify an archive explicitly, as in these two examples:

```
% arx set-tree-version joe.hacker@gnu.org--alt/ArX--1.0
```

or

```
% arx set-tree-version -A joe.hacker@gnu.org--alt ArX--1.0
```

## 8.6 Combining Project Trees

You can, in fact, combine project trees: storing the files and directories from multiple projects under a common root. This can be useful if you have separately maintained projects that, nevertheless, are tightly integrated.

## 9 Archives

This chapter discusses archives – places where revisions can be stored and shared. Archives are used to keep a detailed history of how a project evolves. They are used to help a team of developers stay "in sync" – in agreement about what the official, latest revision of a project really is. Archives are used to help coordinate divergent development paths (branches) and to merge changes between branches.

### 9.1 Archive Names Revisited

Every archive has a name which *should* be globally unique. These names were previously introduced in the context of project names generally (see [Section 8.2 \[Archive Names\]](#), page 42.)

To briefly review, an archive name consists of an email address (with complete hostname), followed by "-", followed by an additional string of numbers, letters and dashes (but not two dashes in a row). Choose an email address which is exclusively yours (or your project's). That way, your archive name will be globally unique. Here is an example:

```
joe.hacker@gnu.org--test-archive
```

### 9.2 Creating a New Archive

To create a new archive on a local file system, use the `make-archive` command:

```
% arx make-archive NAME LOCATION
```

The `NAME` is the name for the archive. The `LOCATION` is a name for a directory that will be created to hold the archive. For example:

```
% arx make-archive \  
    joe.hacker@gnu.org--test-archive \  
    ~/{test-archive}
```

## 9.3 Mapping Archive Names to Locations

ArX maintains a mapping of archive names to archive locations separately for each user, in the directory `~/.arch-params`.

You can see the list of archives for which you have a recorded location by using the `archives` command, as in this example:

```
% arx archives
joe.hacker@gnu.org--test-archive
/home/joe/{test-archive}
```

The command `register-archive` is used to add, remove, or update the location of an archive:

```
% arx register-archive ARCHIVE-NAME LOCATION
```

records a new archive location or changes an old one. The command:

```
% arx register-archive -d ARCHIVE-NAME
```

removes the record of an archive location.

If you want to know the registered location of a particular archive, you can use:

```
% arx whereis-archive ARCHIVE-NAME
LOCATION
```

## 9.4 Remote Archives

There are four ways to set up a remote archive: `ftp`, `wu-ftp`, `sftp`, and `http`. The `wu-ftp` method works around some bugs in `wu-ftpd`'s implementation of the `ftp` protocol.

The first thing that you have to do is set up the (s)ftp or `http` server on the remote machine. ArX does NOT have to be installed.

### 9.4.1 HTTP

There are two ways that `http` access can work.

#### 9.4.1.1 Webdav

This is the recommended and most reliable way. For `apache`, this usually involves installing the `mod_dav` modules. This will work with `apache` 1.3 or later. Then you have to add something like the following to the `conf` file for `apache`:

```

<Directory /home/*/public_html>
  DAV On
  AllowOverride FileInfo AuthConfig Limit
  Options MultiViews Indexes SymLinksIfOwnerMatch IncludesNoExec
  <Limit GET POST OPTIONS PROPFIND>
    Order allow,deny
    Allow from all
  </Limit>
  <Limit PUT DELETE PATCH PROPPATCH MKCOL COPY MOVE LOCK UNLOCK>
    Order deny,allow
    Deny from all
  </Limit>
</Directory>

```

You might have to change the first line of that to make it point to where your archives are.

### 9.4.1.2 Explicit lists

This should only be used if you are unable to install webdav support on your server. You do this by creating a listing file. One idiom to do this is with the command

```

find . -type d -exec sh -c "(cd ; ls | sed -e 's/~/^/' > .listing)" \;

```

this is CTRL-M ---!

You will have to run this script each time you update the archive.

## 9.4.2 SFTP

Setting up the sftp server generally requires just setting up an ssh server. On Debian systems, there is a bug in the default config file (see <http://bugs.debian.org/cgi-bin/bugreport.cgi?bug=141979>).

## 9.4.3 Accessing the Archives

You can access the archive by specifying the uri:

```

ftp://[user@]host/dir           # passive ftp
wu-ftp://[user@]host/dir       # passive ftp with wu-ftp
sftp://[user@]host/dir         # the ssh ftp protocol
http[s]://[user@]host[:port]/dir # webdav

```

ArX saves the location in your `.arch-params` directory, and it is made unreadable to anyone but you. Pretty worthless security, so don't rely on it. Currently, https is not implemented, so the only way to create secure remote writeable archives is with sftp.

As a concrete example, to register an anonymous ftp server, the command would look something like

```

% arx register-archive joe.hacker@gnu.org--test-archive \
  ftp://anonymous@ftp.gnu.org/users/joe/{test-archive}

```

and then you can treat it like a local, read-only archive. For an http server, it would look like

```
% arx register-archive joe.hacker@gnu.org--test-archive \  
http://www.gnu.org/~joe/{test-archive}
```

## 9.5 Your Default Archive

Using `ArX` is generally simplified by setting a default archive – the archive to use by default when no other is specified.

The command:

```
% arx my-default-archive [options] [archive]
```

can be used to set or check your default archive. (This command was previously introduced. See [Section 8.2 \[Archive Names\]](#), page 42.)

## 10 Development Paths

An ArX archive is organized around development paths. Each development path is a succession of revisions, each (usually) differing from the previous revision by a simple patch.

Every development path has a project name with version number, such as:

```
hello--devo--1.0
```

Project names were previously discussed (see [Chapter 8 \[The ArX Global Name-space of Projects\]](#), page 41).

As we'll see in the next chapter, each development path is further subdivided into specific revisions, each having a patch level name, as in these examples:

```
hello--devo--1.0--patch-3
hello--devo--1.0--patch-4
hello--devo--1.0--patch-5
```

### 10.1 Creating a Development Path

Supposing that you have set your default archive, you can create a new development path with the `create-branch` or `create-version` commands.

### 10.2 Examining an Archive

You can review what's in an archive using the commands:

```
% arx categories
    - print a list of categories

% arx branches CATEGORY-NAME
    - print a list of branches within a category

% arx versions BRANCH-NAME
    -print a list of versions within a branch
```

### 10.3 Fully Qualified Version Names

Every version is associated with a particular archive. Ordinarily, if you write a simple version name, your default archive is presumed. However, wherever a version name is called for, you can use a fully qualified version name of the form:

ARCHIVE-NAME/VERSION-NAME

as in the example:

`joe.hacker@gnu.org--test-archive/hello--devo--1.0`

That can be useful when operating on a remote archive that is not your default archive. For example, you could use the name in the example to retrieve the latest revision from Joe Hacker's archive, in preparation for creating a branch of that version in your local archive.

## 11 Basic Revision Control

This chapter introduces the fundamental operations for storing revisions in archives, retrieving them, doing clever things with patches, and managing project trees for archived projects.

### 11.1 The First Revision

When beginning a new project, the first step is to check in the very first revision of the tree. From the top level directory of the project tree, you can do this with `create-version`. For example

```
% arx create-version CATEGORY--BRANCH--VERSION
```

This will set up all of the things behind the stage needed to commit the tree. It will not actually modify the archive. It will also pop up an editor (set with `my-editor`) to edit a log file template. It will look like

```
Summary:
```

```
Keywords:
```

The details of what should go in a log message are project specific. To make a multi-line `Summary:` or `Keywords:` line, just indent it. Also, you have to skip a space after the `Keywords:` line before. So a complete log entry might look like

```
Summary: Added foo, removed bar.  Rearranged and split up
       the baz directories.
```

```
Keywords: foo, bar
```

```
Added foo, which required reworking lart.c to handle
a more general case.
```

```
Removed bar, because it is now subsumed by foo.
```

It could also be as simple as

```
Summary: Added foo, removed bar.
```

In general, the format of a log message uses RFC822 style headers followed by a free-form body. The only required header is the `Summary:.` A body is not required. When ArX stores a log message, it will add some headers of its own.

**Add files** By default, you have to explicitly add files to the list of files that are considered to be source (and thus should be archived). So if you have files `foo.c`, `bar.py`, and `baz.pl` in the project tree, but you only want to archive `foo.c` and `bar.py`, then issue the command

```
% arx add foo.c bar.py
```

You can check which files have been marked as source with

```
% arx inventory
```

See [Chapter 6 \[ArX Project Inventories\]](#), page 27 for more information about how ArX decides what to archive.

**Archive the Tree** Again, in the root of the project tree, use the `commit` command to archive the tree:

```
% arx commit
```

`commit` will make all of the necessary categories, branches, and versions, store the tree in the archive (as a compressed tar file), and update the patch log of the working directory to reflect the commit.

After committing a revision to, say, `hello--devo--1.0`, you can use the command `revisions` to see a list of archived revisions:

```
% arx revisions hello--devo--1.0
base-0
```

`base-0` is the "patch level name" for the first revision. Patch levels are described in greater detail below.

## 11.2 Successive Revisions

Suppose that you have continued to edit your working directory after checking in the initial revision. Successive revisions can be checked in with the command `commit`, issued at the root of the working directory. As before, you must first have a log file. You can use the `make-log` function to prepare a log message for each commit:

```
% arx make-log
```

This will pop up an editor to modify the log file template. To commit the changes, you once again use `commit`.

```
% arx commit
```

After several commits, you will have a number of patch levels:

```
% arx revisions hello--devo--1.0
base-0
patch-1
patch-2
patch-3
...
```

As a point of interest, the base revision is stored as a compressed tar file of the entire tree. Each `patch-N` is stored as a compressed tar file of just a patch set that describes how to derive that revision from the previous revision.

## 11.3 Patch Levels

Within a version, there is a sequence of revisions, each of which is a different patch level. Every patch level has a patch level name, derived from the version name by adding a suffix:

```
hello--devo--1.0--base-0
hello--devo--1.0--patch-1
hello--devo--1.0--patch-2
hello--devo--1.0--patch-3
...
```

Just as with version names, there is also such a thing as a fully qualified patch level name:

```
joe.hacker@gnu.org--test-archive/hello--devo--1.0--patch-3
```

It is perhaps worth mentioning that patch level names can be sorted easily using:

```
sort -t - -k 1,1 -k 2,2n
```

or in reverse:

```
sort -t - -k 1,1r -k 2,2rn
```

## 11.4 Tagging

You can also give alternative names for revisions. For example, you can use the tag `release-candidate` to name whatever revision people should download for testing purposes, regardless of what branch, version, or patch level the release candidate happens to be on.

In ArX, tags are revisions on ordinary branches. For example, suppose we are developing `ArX--devo--0.5` and want to create a tag `release-candidate` to mark revisions which "early adopters" should test. This is as simple as

```
% arx tag ArX--devo--0.5--patch-37 \
      ArX--release-candidate--0.5
```

Note that `source-revision` and `tag-revision` may be abbreviated. For example, to tag the most recent revision of `ArX--devo`:

```
% arx tag ArX--devo--0.5 ArX--release-candidate--0.5
```

or

```
% arx tag ArX--devo ArX--release-candidate
```

After such a command, you can retrieve the tagged revision in the ordinary way:

```
% arx get ArX--release-candidate--0.5
```

Note that when you `get` a tag that way, the default version of the resulting project tree is the tag's version, not the tagged version (see [Section 8.5 \[Labelling Project Trees\]](#), page 44).

You can always update a tag, making it point to a later revision, again using the `tag` command:

```
% arx tag ArX--devo--0.5--patch-53 \
    ArX--release-candidate--0.5
```

You can see the history of a tag in the usual way, too:

```
% arx revisions --summary ArX--release-candidate--0.5
base-0
    tag of joe.hacker@gnu.org--ArX/ArX--devo--0.5--patch-37
patch-1
    tag of joe.hacker@gnu.org--ArX/ArX--devo--0.5--patch-53
```

## 11.5 Development Phases

**Note:** Phased development may be removed in the future. Its' status is uncertain.

Development within a version may be optionally divided into four phases: base revision, pre-patches, version revision, and post-patches. Conceptually, "pre-patches" are revisions made before the version is "done". "Post-patches" are revisions made after the version is done (e.g. "bug fix patches").

At the beginning of a development path is the base revision (called `base-0`). Between the pre-patch and post-patch phases is the version revision (called `version-0`). Post-patches have names like `versionfix-1`, `versionfix-2`, etc.

So the total sequence of revisions within a development path has the form:

The Initial Revision:

```
base-0
```

Pre-Patches:

```
patch-1
patch-2
patch-3
...
patch-N
```

The Version Revision:

```
version-0
```

Post-Patches:

```
versionfix-1
versionfix-2
versionfix-3
...
```

Typically, the `version-0` revision is what would be released under the version number and the post-patch revisions are fixes made after the release.

The ordinary `commit` command creates pre-patches (`patch-N` revisions).

To create the version revision, use `commit --seal`:

```
% arx commit --seal
```

After the version revision exists, ordinary `commit` will no longer work. To create a post-patch revision, use `commit --fix`:

```
% arx commit --fix
```

If you are familiar with other revision control systems, a four-phased development process may at first seem somewhat arbitrary and needlessly complicated. Two points are worth mentioning:

First, phased development is entirely optional. Nothing requires you to ever `--seal` a version, and if you never seal a version, you never need to use `--fix`. Instead, every revision (after `base-0`) will be a `patch-N` revision.

Second, phased development is a handy way to prevent accidents when organizing more complex projects. Sealing a version is a way to set a flag that says, in effect, "ordinary development in this version has stopped – you might not really want to make a new revision here." You can make a new revision if you insist (by specifying `--fix`), but the requirement that you insist helps alert you to the fact that your new revision might need to be merged with later versions of the same project; or that that version you are revising has already been released. Phased development is a bookkeeping convenience: for the most part, `ArX` treats all revisions equally, regardless of their phase.

## 11.6 Getting a Revision

To retrieve a revision from an archive, use `arx get`:

```
% arx get REVISION DIR
```

as in:

```
% arx get hello--devo--1.0--patch-4 hello
```

to retrieve the revision `patch-4` and store it in the new directory `hello`.

An abbreviation can be used to obtain the most recent revision of a version:

```
% arx get hello--devo--1.0 hello
```

or to obtain the most recent revision of the highest-numbered version:

```
% arx get hello--devo hello
```

A fully-qualified name can be used to obtain a revision from someplace other than the default archive:

```
% arx get joe.hacker@gnu.org--test-archive/hello--devo
```

## 11.7 Optimizing Archives for get

The way that `get` ordinarily works is that it searches backwards from the desired revision to find the nearest full-source base revision. It gets the compressed tar file for that base revision and creates a source tree. Then, for each intermediate patch level, it gets a compressed tar file of the patch set, uncompresses and un-tars the patch-set, and applies the patch-set to the source tree.

If there are many intermediate patch-sets, that process can be slow. In such cases, you can ask ArX to cache a full-source copy of an arbitrary revision, with the command:

```
% arx archive-cache-revision [ARCHIVE/]REVISION
```

That command first builds the requested revision, then it builds a compressed tar file of the revision, then it stores the tar file back in the archive. Subsequent attempts to `get` the same revision (or any later revision) will use the cached tree.

To remove a previously cached tree, use:

```
% arx archive-uncache-revision [ARCHIVE/]REVISION
```

For each user, ArX also maintains a "client side" cache of revisions that can speed up `get` (and other operations). See [Chapter 22 \[Revision Tree Libraries\]](#), page 99 for more details.

## 11.8 Finding Out What Changed

Before performing a `commit`, you might want to check to see what has actually changed – that is, find out exactly what patch set your `commit` will create.

You can do that with the command `what-changed`:

```
% arx what-changed
[...patch set report...]
```

`what-changed` computes a patch set between your modified project tree and the latest patch level for which your project tree is up-to-date. In other words, it tells you what changes have been made to your tree compared to the tree in the archive.

`what-changed` leaves behind a directory containing the patch set and patch report, which you can usefully browse by hand.

```
% ls
, ,what-changed.ArX--devo--0.5--patch-14--lord@regexprs.com--ArX-1
[...]
```

The default output of `what-changed` is formatted for use with the `outline` mode of **GNU Emacs**. The emacs mode bundled in the distribution makes it easy to browse the differences.

You can ask `what-changed` to also generate an HTML-formatted report with:

```
% arx what-changed --url
URL of patch report
```

The HTML report has links to each context diff, added, and removed file. The output of the command is a `file:` method URL, which can be used by browsers

```
% mozilla -remote "openURL('arx what-changed --url')"
```

A simpler solution is to set your browser using the `my-browser` command. Then you can invoke `what-changed` with the `--new-browser` or `--new-tab` options. For example,

```
% arx what-changed --new-tab
```

will open the report in a new tab of your browser.

## 11.9 The `whats-missing` Command

Suppose that more than one programmer is checking revisions into a version, Alice and Bob for example.

Alice and Bob both start with working directories and both make some changes. Alice commits several changes. Now Bob's working directory has fallen behind archived development path.

The command `whats-missing` can be used to tell Bob which patches he is missing:

```
% arx whats-missing --summary
patch-N
    summary of patch N
patch-N+1
    summary of patch N+1
...
```

For each patch that Bob is missing, `whats-missing --summary` prints the name of the patch and the contents of the `Summary:` header from its log message.

The `whats-missing` command is explained in greater detail in a later chapter (see [Chapter 13 \[Patch Logs and ChangeLogs\]](#), page 67).

## 11.10 Update

So Bob is behind by a few patches, but also has his own modifications.

Diagrammatically, we have something like:

```

Patch Levels      Bob's Working
in the           Directory
Archive:
-----

base-0

patch-1
patch-2
patch-3
patch-4 -----> bob-0 (Bob's initial working directory)
patch-5           |
patch-6           |
patch-7           V
patch-8           bob-1 (Bob's working dir with changes)

```

Bob is missing patches five through eight.

The command `update` can be used to fix the situation:

```
% arx update OLD-DIR NEW-DIR
```

as in:

```
% arx update bob-1 bob-2
```

`update` works in several steps. First, it gets a copy of the latest revision (`patch-8` in this case). It also gets a copy of the revision from which `OLD-DIR` is derived (`patch-4` in this case). Then it uses `mkpatch` to compute the differences between `OLD-DIR` and its source revision, and applies those differences (using `dopatch`) to the latest revision.

In the example, `update` will create the directory `bob-2`, with the source:

```
delta(patch-4, bob-1) [patch-8]
```

(For information about this notation, see [Appendix B \[The Theory of Patches and Revisions\]](#), page 123.)

Applying that patch might cause conflicts. In that case, `update` will print a message telling Bob to look for `.rej` files.

As a convenience, `update` also copies all non-source, non-junk files from `OLD-DIR` to `NEW-DIR` (see [Chapter 6 \[ArX Project Inventories\]](#), page 27).

Of course, if Bob only wanted to "partly update", he could do that with an extra parameter to `update`, as in the example:

```
# update, but only up to patch level 6:
#
% arx update bob-1 bob-2 hello--devo--1.1--patch-6
```

Finally, `update` can replace `OLD-DIR` with the updated directory if given the `--in-place` flag:

```
% arx update --in-place OLD-DIR
```

## 11.11 Replay

`update` isn't the only way to catch-up with a development path. Another option is `replay`:

```
% arx replay old-dir new-dir
```

Using the same example:

```
Patch Levels      Bob's Working
in the            Directory
Archive:
-----
base-0

patch-1
patch-2
patch-3
patch-4 -----> bob-0 (Bob's initial working directory)
patch-5           |
patch-6           |
patch-7           V
patch-8           bob-1 (Bob's working dir with changes)
```

and the command:

```
% arx replay bob-1 bob-2
```

`replay` will first copy `bob-1` to create `bob-2`. Then it will apply each missing patch in succession until `bob-2` is up-to-date, or until a merge conflict occurs.

Thus, if no conflict occurs, `replay` computes:

```
patch-8 [ patch-7 [ patch-6 [ patch-5 [ bob-1 ]]]]
```

If a conflict occurred in, say, `patch-6`, then `replay` would compute:

```
patch-6 [ patch-5 [ bob-1 ]]
```

and after fixing the conflict, Bob could use a second `replay` command to apply patches seven and eight.

As with `update`, `replay` also copies all non-source, non-junk files from `OLD-DIR` to `NEW-DIR` (see [Chapter 6 \[ArX Project Inventories\]](#), page 27).

Of course, if Bob only wanted to "partly replay", he could do that with an extra parameter to `replay`, as in the example:

```
# replay, but only up to patch level 6:
#

% arx replay bob-1 bob-2 hello--devo--1.1--patch-6
```

You can use `replay` to modify an existing directory rather than creating a new directory:

```
% arx replay --in-place DIR [REVISION]
```

but be careful: if `DIR` contains precious local changes, and conflicts occur, or if you simply decide the `replay` wasn't a good idea, you'll have to do some work to revert the `replay`.

## 11.12 The Next Version

In simple situations, a version like `hello--devo--1.0` will be followed by the next version: `hello--devo--1.1` or `hello--devo--2.0`, for example.

Such a version is called a continuation of the previous version and it is created with `create-version`.

```
% arx create-version NEXT-VERSION
[...edit log message...]
% arx commit
```

as in:

```
% arx create-version hello--devo--1.1
[...edit log message...]
% arx commit
```

Those commands create the `base-0` revision of the new version but instead of storing complete source for the base revision, they store a pointer to the older revision which the base revision is equal to.





## 12 Basic Branching and Merging

When a single development path splits into two paths, that's called a branch. Typically, a branch is followed by a merge – adding the changes made in a branch back to the branch from which it diverged.

This chapter explains how to create a branch and the simplest way to merge changes from two branches. Later chapters will explain fancier techniques, useful in more complex situations.

### 12.1 Creating a Branch

You can create branches with the `create-branch` command. It leaves you with a working directory for the new branch. The sequence of commands is simply

```
% arx create-branch OLD-REVISION NEW-BRANCH-VERSION WORKING-DIR
[...edit log message...]
% cd WORKING-DIR
% arx commit
```

`create-branch` will invoke an editor (specified by `my-editor`) to modify the log template.

For example, to create a branch `hello--experimental--1.0` from the latest revision of `hello--devo--1.0`, use:

```
% arx create-branch hello--devo--1.0 \
                        hello--experimental--1.0 \
                        wd
[...edit log message...]
% cd wd
% arx commit
```

If you make modifications to a project tree, but then decide that those modifications should go into a new branch, you can use the `--in-place` option to `create-branch`. For example,

```
% arx get hello--devo--1.0 wd
% cd wd
[...edit files...]
% arx create-branch --in-place hello--experimental--1.0
[...edit log file...]
% arx commit
```

## 12.2 Distributed Branches

There is no requirement that a branch be stored in the same archive as the revision from which it branched. For example, you can create a private archive, and store some branches there – only merging those changes back into the shared archive when they are ready.

Here's a tip: make your private archive your default archive. Use fully-qualified version and revision names when getting or committing revisions in the shared archive. This makes it less likely that you'll accidentally make unintended changes to the shared archive.

## 12.3 `whats-missing` Revisited

If you have a project tree for a branch, you might want to know what has happened in the version from which you branched.

The `whats-missing` command is used for this. In a working directory for a branch, use:

```
% arx whats-missing --summary ORIGINAL-VERSION
```

where `ORIGINAL-VERSION` is the version name of the version from which you branched. Actually, `ORIGINAL-VERSION` can be any version for which your project tree has a patch log.

The `whats-missing` command is explained in greater detail in the next chapter (see [Chapter 13 \[Patch Logs and ChangeLogs\]](#), page 67).

## 12.4 update and replay Revisited

Similarly, `update` and `replay` work for any version for which a project tree has a patch log, such as a version from which a branch occurred:

```
% arx update OLD-DIR NEW-DIR [archive/]VERSION
```

```
% arx replay OLD-DIR NEW-DIR [archive/]VERSION
```

## 12.5 Merging After a Branch

The simplest use of branching and merging is this: you have one development path, call it the "trunk". You form a branch from that development path, which we'll just call "the branch".

To make some changes, you do your work on the branch: check out the latest revision from the branch, make changes, commit, make more changes, commit again, etc.

As you work, you might sometimes need to "catch up" to changes made to the trunk. You can do that by using `update` or `replay`.

When you're done, and the branch is fully up-to-date with the trunk, you can check out the latest branch revision, then commit that version to the trunk. All of the changes that you made on the branch will be summarized into a single patch.

There are more complicated and more realistic uses of branches. These are the subjects of the next several chapters.



## 13 Patch Logs and ChangeLogs

Every project tree has an associated patch log: a collection of log entries for each `commit` or `import` in the history of that patch tree. When you commit a new revision, the log entry you write is saved in two places: it is saved in the archive as a plain text file (for browsing and as a record of complex ancestry relationships), and in the project tree itself (for browsing and to control history sensitive merging).

Logs are organized by version. The command:

```
% arx logs
```

lists all of the version names for which a project tree has a patch log.

The command:

```
% arx log-ls [[ARCHIVE/]VERSION-NAME]
```

lists all of the patch levels for which a tree has log entries (for revisions in the indicated version). With the `summary` flag:

```
% arx log-ls --summary [[ARCHIVE/]VERSION-NAME]
```

the **Summary:** header of each log entry is printed.

To see the complete text of an entry, use:

```
% arx cat-log [ARCHIVE/]VERSION--PATCH-LEVEL
```

## 13.1 Branches and Patch Logs

When you form a branch, project trees on the branch have (at least) two patch logs: one for the original development path, and one for the branch itself. When you merge changes from one branch to another, so long as both branches have the same project category, the merged tree has patch logs for both branches. ("Project categories" were introduced in [Chapter 8 \[The ArX Global Name-space of Projects\]](#), page 41).)

## 13.2 Comparing Patch Logs to Archives

You can find out if an archive contains patches that haven't yet been applied to your project tree with this command:

```
% arx whats-missing [[ARCHIVE/]VERSION ...]
<list of missing patches>
```

That command compares the patch log stored in the archived with the patch log found in the project tree and prints the list of missing patches. There may be missing patches if your tree is not up-to-date with respect to the archive, or if when your tree was created, some patches were skipped.

You can see the `Summary:` line of each missing patch with:

```
% arx whats-missing --summary [[ARCHIVE/]VERSION ...]
```

If you want the list to contain fully-qualified patch level names, use the `--full` option.

If you want to know where branch A stands in relation to branch B, one way to find out is with:

```
% arx get A dir

% cd dir

% arx whats-missing B
```

(It is possible to obtain the same information without having to `get` a revision from branch A, using commands already introduced, plus some other shell commands. The details are left as an exercise for the interested reader.)

### 13.3 ChangeLogs

The command `arx changelog` generates a GNU-style `ChangeLog` file from a patch log:

```
% arx changelog
```

or

```
% arx changelog [ARCHIVE/]VERSION
```

The `ChangeLog` file generated for ArX, for example, might begin:

```
# do not edit -- automatically generated by ArX changelog
#
# tag: automatic-ChangeLog--lord@regexps.com--ArX-1/ArX--devo--0.5
#

2001-12-17  Tom Lord <lord@regexps.com>

Summary:
  'update' and 'replay' output format and bug fixes

  'update' and 'replay' -- structured output and updated argument
  processing for reasonable defaults.

  'replay': copy precious files before (not after) applying patches
  so they are carried along with directory renames in patch sets.

  'dopatch': don't pipe 'arx heading' into 'arx body-indent'.

modified files:
ChangeLog src/ArX/=TODO
src/ArX/branching-and-merging/replay.sh
src/ArX/branching-and-merging/update.sh
src/ArX/patch-sets/dopatch.sh

2001-12-17  Tom Lord <lord@regexps.com>

Summary:

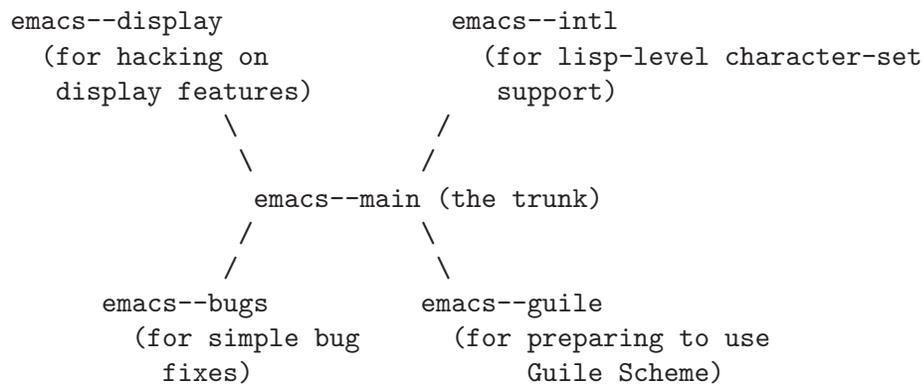
[...]
```

If someone has a full distribution of your work, they can generate a `ChangeLog` as needed. You can also include `ChangeLogs` in tarball distributions using the `--changelog` option for `update-distributions`. That will make a directory `ChangeLog.d`, and the `ChangeLogs` will go in `ChangeLog.d/ARCHIVE/REVISION`. If `ChangeLog.d` already exists, then ArX will not put `Changelogs` into that tree.



## 14 Star Topology Branching and Merging

A common way to use branches is to form a star topology. At the center is a trunk or "primary development path". A number of branches form a "star" around the trunk. There might be one branch for each developer or sub-team; one branch for each large task; one branch for each change reviewer; one branch for each category of tasks, etc. Here's a (fictional) example of using a star topology to work on various aspects of GNU `emacs`:



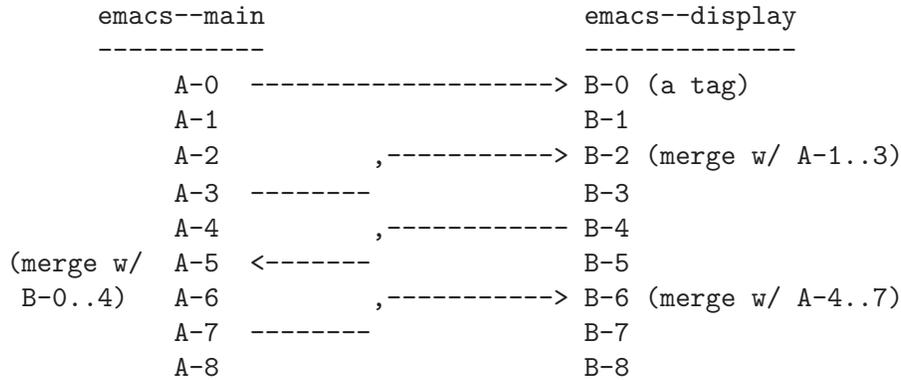
When using a star topology, developers make changes on the branches, merging those changes into `emacs-main` at significant milestones. Periodically, they merge the collected changes on `emacs-main` back to the development branches. The trunk is the official development sources. Work takes place on the branches, and the developers use the trunk to stay in-sync.

This results in a situation where each branch has merged with `emacs-main` multiple times, and `emacs-main` with each branch several times. In such situations, simple `update` and `replay` are inadequate for performing merges.

This chapter illustrates the problem with using `update` and `replay` for merging in a star topology, explains the solution in general terms, and finally presents `star-merge`, a command that implements the solution.

## 14.1 The Star Topology Merge Problem

Consider the trunk and one of its branches after several revisions and merges on each:



Now, suppose we have a working directory that was checked-out from A-8 and perhaps contains local modifications. Our goal is to merge B-8 with that working directory, giving a revision that is up-to-date with both branches and that includes any local changes.

```
% arx update --in-place --dir WD emacs--display
```

modifies WD such that:

```
WD := delta(B-4, WD) [B-8]
```

Well, what has changed in WD since B-4? B-4 was up-to-date up to A-3, but WD contains additional changes from A-4..A-8 – all of those changes are included in the delta computed by `update`.

`update` applies that delta to B-8. Unfortunately, B-8 already includes the changes in A-4..7 – so conflicts are guaranteed.

We can't use `replay` either:

```
% arx replay --in-place --dir WD emacs--display
```

tries to apply all the patches missing from B. For WD, those patches are B-5..B-8. Unfortunately, patch B-6 includes all the changes from A-4..7 (perhaps with additional changes or conflict resolutions). WD already has all of those same changes, so again, conflicts are guaranteed.

## 14.2 Solving the Star Topology Merge Problem

In the case of the example above, there are three reasonable ways to solve the merge problem, depending on the relative precedence we want to give the two development paths and the working directory.

If we want to give precedence to the changes already in the working directory, we can perform the merge by computing:

```
WD := delta (A-7, B-8) [WD]
```

If we want to give precedence to the changes in `emacs--display`, we can perform the merge by computing:

```
WD := delta (A-7, WD) [B-8]
```

Finally, we might want to merge in two steps: first merging `emacs--display` with A-8 (giving precedence to `emacs--display`), then adding the local changes of WD to that:

```
tmp := delta (A-7, B-8) [A-8]
WD := delta (A, WD) [tmp]
```

There are in fact six different ways of precedence-ordering the two branches and WD however one of the above solutions works for each of the six possibilities.

Three different solutions are needed if WD was checked out from `emacs--display` and our goal is to merge in changes from `emacs--main`. These are:

```
1:      WD := delta (A-7, A-8) [WD]

2:      tmp := delta (A-7, A-8) [B-8]
        WD := delta (B-8, WD) [tmp]

3:      WD := delta (A-7, WD) [A-8]
        (this is an ordinary 'update')
```

In the two two-step solutions, legitimate conflicts can occur while building the `tmp` tree: its necessary to resolve such conflicts by hand before performing the second step.

Suitably abstracted, those six merge techniques are sufficient to solve all star topology merging problems with a choice of branch and local change precedence without generating any spurious merge conflicts.

### 14.3 The star-merge Command

The `star-merge` command figures out how to solve a star topology merge problem and performs the merge. Its syntax is:

```
% arx star-merge [--in-place] A B C [output-dir]
```

where `A`, `B`, and `C` are the names of two branches and a working directory (in any order). (If your working directory name can be mistaken for a branch name, and most of them can, you should prefix the directory name with `./`).

With the `--in-place` option, the working directory is directly modified or else completely replaced with the result. Otherwise, the merge is stored in a new directory (`output-dir`).

The order of arguments determines the precedence of changes. `A`, `B` and `C` have a common ancestor revision – call it `X`. The order of arguments means (conceptually): start with revision `X`, add the changes from `A`, then the changes from `B`, then the changes from `C`. The specific ordering of changes determines, when there are conflicts, which branch's code is automatically incorporated, and which is left as `.rej` files.

If the merge requires two steps, and conflicts occur in the first step, `star-merge` will stop after the first step leaving the output directory in an intermediate state. After fixing the conflicts, you can complete such a merge with:

```
% arx star-merge --finish [directory]
```

**NOTE:** In the current release, there is a limitation on star merge. Suppose that `B` is a branch of `A` and that `B` and `A` have never been previously merged. In order to merge them, you must use a working directory checked out from `B` (the branch version) not from `A` (the branched-from version). Alternately, you can use the `join-branch` command on the directory checked out from `A`.

This restriction will be removed in the next release.





## 15 Writing Log Entries for Merges

Suppose you have used `star-merge` or some other commands to merge two or more branches. You are ready to commit those changes, but first you have to write a log message.

All of the changes you've just merged are already explained in existing log messages. Your new log message can simply summarize the changes, and point to the more detailed log entries. `ArX` has commands which help with this.

The command:

```
% arx new-on-branch [--reverse] [--dir .] VERSION
```

reports a list of patches found in a project tree, but not previously merged with `VERSION`. Those patches are the ones that were added by your merging activity.

You can print a summary of the changes made by those patches by using `xargs` and the `log-ls` command. For example, the main development path for `ArX 1.0` is called `ArX--devo--1.0`. I develop new features on various branches. When it comes time to write a log message before committing a merge to `ArX--devo`, I use:

```
% arx new-on-branch --reverse ArX--devo--1.0 \
  | xargs arx log-ls --full --summary

lord@regexps.com--2002/ArX--lord--1.0--patch-14
  'merge-points' fixes and 'new-on-branch' speed-up
lord@regexps.com--2002/ArX--lord--1.0--patch-13
  Use --dir consistently.
lord@regexps.com--2002/ArX--lord--1.0--patch-12
  output format touch-ups
```

That idiom is captured, along with a little bit of formatting, by the command `log-for-merge`:

```
% arx log-for-merge ArX--devo--1.0
```

Patches applied:

```
* lord@regexps.com--2002/ArX--lord--1.0--patch-14
  'merge-points' fixes and 'new-on-branch' speed-up
* lord@regexps.com--2002/ArX--lord--1.0--patch-13
  Use --dir consistently.
* lord@regexps.com--2002/ArX--lord--1.0--patch-12
  output format touch-ups
```

(An advantage of using the formate generated by `log-for-merge` is that it is understood by other `ArX` commands which automatically format web pages from `ChangeLogs`.)



## 16 Arbitrary Patching with delta-patch

A general purpose, though low-level tool for merging is the `delta-patch` command:

```
% arx delta-patch FROM TO UPON OUTPUT-DIR
```

computes a patch set from `FROM` to `TO` and applies that patch set to `UPON` storing the result in `OUTPUT-DIR`.

The arguments `FROM`, `TO`, and `UPON` can be the names of (any) revision, or the directory names of project trees.

If `UPON` is a project tree, you can modify that tree directly with the `--in-place` option:

```
% arx delta-patch --in-place FROM TO UPON
```

**Note:** You pretty much have to know what you're doing to use this command. Normally, you should use `star-merge`, `update` or `replay`.



## 17 Multi-Branch Merging – The reconcile Command

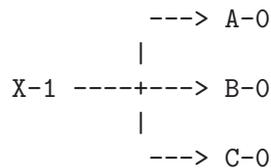
History sensitive merging with `replay` can avoid some avoidable merge conflicts, but not all. One example is a class of merge problems that we'll name "Repeated Multi-Branch Merging": the problem of merging several branches when each of the branches have previously merged with some of the others. Although this kind of merging seems arcane, it can, in fact, easily arise in quite realistic situations (for example, when simultaneously supporting multiple releases of a single project).

Below is an example to illustrate the problem. The set-up in this example is a bit long, but each step along the way is perfectly reasonable, and the end result is quite a tangled knot. The pay-off will be seeing how to cope with the resulting mess.

### 17.1 The Repeated Multi-Branch Merge Problem

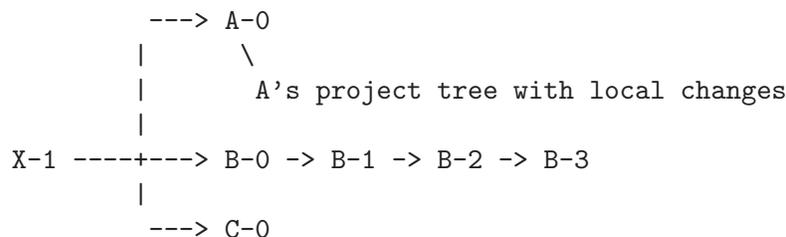
Imagine that we start with a particular version of a particular branch, call it X. We'll begin at a particular revision in that branch version: X-1 (for the purposes of this explanation, calling the revision X-1 is much less cumbersome than using a real revision name like `foo--mumble--3.5--patch-24`).

Three programmers each form their own branch from X-1: call them A, B, and C:

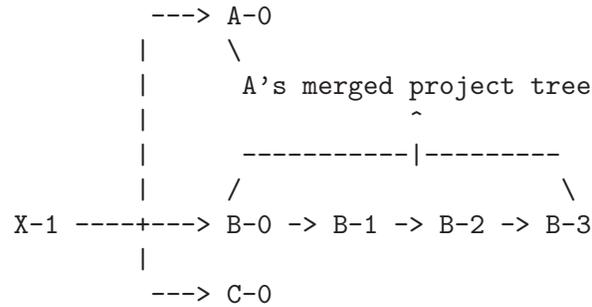


The plan here is develop on each branch, then merge the changes together to create a new revision of X.

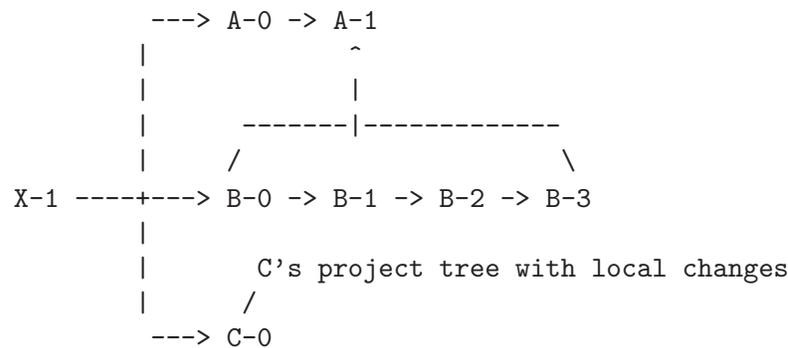
Programmer B starts off, and creates a series of revisions. Simultaneously, A creates a project tree and starts making local changes for his first revision:



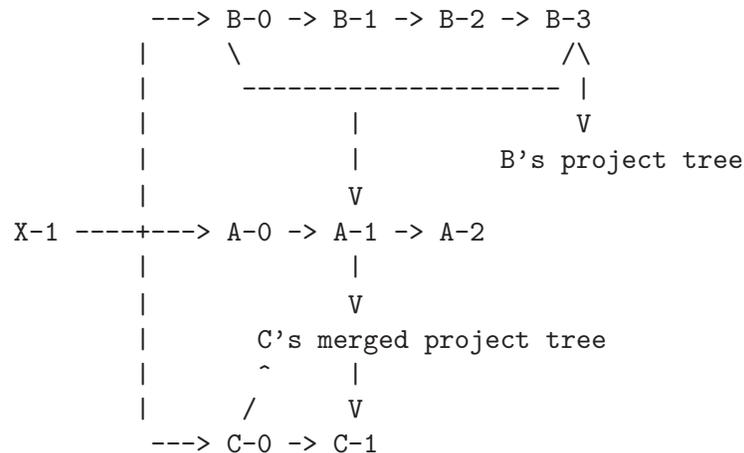
Programmer A wants to develop on top of those three patches from B, and so forms a merge. At this stage, A can do a simple `update` or `replay` to create a merged project tree:



Let's assume that A's merge involved some conflicts: B's code has been slightly rearranged in the merged tree. Now A can check in that revision. Meanwhile, C starts work:



C decides it would be a good idea to merge with the feature's found in A-1. In doing so, he'll also be picking up B-0..3. Once again, a simple `update` (or `replay`) is sufficient at this point, though to keep things interesting, we'll again assume that there are conflicts to resolve during the update. And meanwhile, by the way, B works on his next patch, and A commits a new revision:



We're nearly done with the set-up: B decides to merge in the changes in A-2. This is a slightly interesting merge (although not the primary topic of this chapter). The common ancestor of A-2 and B's project tree is B-3. We previously assumed that when A updated against B-3 there were conflicts that had to be resolved by hand. B has a choice. He can use `update` against A-2 to create a new tree:



## 17.2 The Challenge

Whew. What an (unfortunately plausible) mess. Now for the challenge:

```
Create X-3, which is up-to-date with A-4, B-5, and C-1
```

There is no one right answer to the challenge: no elegant solution that is guaranteed to avoid merge conflicts. Indeed, there are many ways to perform the merge which differ in terms of what conflicts they'll produce. The goal of ArX is to arm programmers with plenty of tools to understand the situation, explore, generate and apply patches effectively, and find a reasonable solution with the greatest degree possible of automated assistance.

## 17.3 The Simple update Solutions

Simple `update` gives us a whole collection of simplistic solutions. For example, X could update against A, then B, then C or:

```
intermediate-1 := delta (A-4, X-1) [X-2]
intermediate-2 := delta (B-5, X-1) [intermediate-1]
X-3-candidate  := delta (C-1, X-1) [intermediate-2]
```

That `update` path has some problems, though. `delta (A-4, X-1)` includes the changes in `delta (A-1, X-1)`, and so does `delta (B-5, X-1)`. So creating `intermediate-2` will involve redundant patching and plenty of opportunities for conflicts. Similar problems occur when creating `X-3-candidate`.

X could try doing the updates in a different order, but similar problems will still occur.

## 17.4 The Simple replay Solutions

X could `replay` the branches in some order. Suppose he `replays` A, then B, then C:

```
intermediate-1 := A-4 [ A-3 [ A-2 [ A-1 [ A-0 [ X-2 ] ] ] ] ]
intermediate-2 := B-5 [ B-4 [ intermediate-1 ] ]
X-3-candidate  := C-1 [ C-0 [ intermediate-2 ] ]
```

History sensitivity helped a bit there: `replay` knows better than to apply `B-0..3` – eliminating one source of needless conflicts. Still, when we `replay` `B-5` and `C-1`, there will be plenty of conflicts to make up for that.

It's also worth mentioning that that this solution involves applying nine different patches: we can do better. By differently ordering the `replay` solution, we get by with fewer patches (`replay` C first, then A, then B, for example). Figuring out the best order in which to apply patches is, ultimately, the subject of this chapter:

## 17.5 The reconcile Solution

Suppose that X asks, of the tree X-2:

```
% arx whats-missing A B C
```

the answer is:

```
A-0
A-1
A-2
A-3
A-4
B-0
B-1
B-2
B-3
B-4
B-5
C-0
C-1
```

X can also ask the more interesting question:

```
% arx whats-missing --merges A B C
```

which will answer not only what patches are missing, but what patches include other patches:

```
A-0 A-0
A-1 A-1
A-1 B-0
A-1 B-1
A-1 B-2
A-1 B-3
A-2 A-2
A-3 A-3
A-4 A-4
B-0 B-0
B-1 B-1
B-2 B-2
B-3 B-3
B-4 B-4
B-5 B-5
```

```

B-5 A-0
B-5 A-1
B-5 A-2
C-0 C-0
C-1 C-1
C-1 A-0
C-1 A-1
C-1 B-0
C-1 B-1
C-1 B-2
C-1 B-3

```

X can pipe that list into a filtering command, `arx reconcile`, which does some magic (the trick is revealed below):

```

% arx whats-missing --merges A B C \
| arx reconcile
C-0
C-1
B-4
B-5
A-3
A-4

```

which means that X can perform the merge with just:

```
A-4 [ A-3 [ B-5 [ B-4 [ C-1 [ C-0 [ X-2 ]]]]]]
```

There is still a potential source of conflicts – when applying B-5 in this case – but the patch set is as small as possible (six patches instead of our earlier nine), and the sources of conflicts are as few as possible.

How did `reconcile` find that solution? What’s the magic? Conceptually, `reconcile` works in two steps.

First, `reconcile` computes a subset of all the patches: the necessary patches. The set of necessary patches is the smallest set of patches which, applied in some order, is sufficient to bring the tree up to date. (Proof that there is a unique smallest set of patches with that property is left as an exercise for the interested reader.)

Second, `reconcile` repeatedly selects the next "necessary" patch to apply, until none are left. At each step of this loop, candidates for the next patch to apply are the patches all of whose prerequisites are in place. Of those, the next patch is the one that comes first in the first column of the input to `reconcile`.

So, you don’t believe this obscure command is useful in real life? See **Even/Odd Versions** in [Appendix A \[Implementing Development Policies\]](#), page 113.





## 18 Reverting Development

Suppose that you are working along, making revisions, when you realize that the last few revisions are just bad, and you want to revert them. You don't want to completely erase the development, since there may be pieces that you would want to use in a different context. But you do want to make it look as all of those changes were reversed.

There is a command for this situation:

```
% arx make-sync-tree good-revision version [dir]
```

This creates a project tree for `version` whose source code is exactly the same as `good-revision` but which has all of the current patch logs. You can then `commit` that tree, reverting the patches after `good-revision`.

However, `make-sync-tree` is more general in application. Suppose that you have two branches: a primary development line and a developer branch from that trunk. For one reason or another, none of those changes which are on the developer branch but not the trunk are needed – perhaps similar changes have recently been made on the trunk.

Now the developer wants to merge the trunk back to his branch, but the desired outcome of this merge is quite trivial: he wants the source in the branch to wind up looking exactly like the source in the trunk, but the patch log in the branch to have all log entries from both the trunk and the branch. In short, it should appear that after diverging, the branch was later edited to agree with the trunk on how to make the branch's changes.

In this case, the command looks like

```
% arx make-sync-tree from-revision to-version [dir]
```

This creates a project tree for `to-version` whose source code is exactly the same as `from-revision` but whose patch log is a combination of the latest patch log from `to-version` and the patch log of `from-revision`.

The project tree thus created can be committed to `to-version`, resulting in a perfect synchronization between two branches.



## 19 Multi-Tree Projects

Large projects are often most naturally divided into a number of independently maintained sub-projects. For example, ArX itself consists of six separate projects:

```
./src
    A generic top-level source file for combining
    projects.

./src/arx
    ArX itself. That is, the 'arx' command and all of
    the shell scripts that implement it.

./src/file-utils
    Generic utilities for operating on files and
    directories.

./src/ftp-utils
    A scriptable FTP client.

./src/hackerlab
    The hackerlab C library.

./src/text-utils
    Generic utilities for operating on text files.
```

In an ArX source tree, each of those sub-directories is a separate project tree (each has its own {arch} subdirectory; each is in a separate category in the repository).

To put all the pieces together, use the ArX configuration commands:

(TO BE DOCUMENTED:)

Config file format, e.g.:

```
./configs/regexps.com/devo.ArX:

#
# Check out an ArX distribution from the devo branches.
# Latest revisions.
#

./src                lord@regexps.com--2002/package-framework--devo
./src/ArX            lord@regexps.com--2002/ArX--devo
./src/file-utils     lord@regexps.com--2002/file-utils--devo
./src/ftp-utils      lord@regexps.com--2002/ftp-utils--devo
./src/hackerlab      lord@regexps.com--2002/hackerlab--devo
./src/shell-utils    lord@regexps.com--2002/shell-utils--devo
./src/text-utils     lord@regexps.com--2002/text-utils--devo
```

and

```
% arx build-config --help
instantiate a multi-project tree
usage: build-config [options] config-name

-V --version          print version info
-h --help            display help

-d --dir DIR         cd to DIR first
--config-dir CFG    get config from directory CFG

--silent            no output
--quiet            only output errors
--default-output    default output
--report           slightly verbose output
--verbose          maximal output
--debug            debugging output
```

Build the named configuration based on the specification in the "configs" subdirectory at the root of the project tree.

and

```
% arx update-config --help
update a multi-project tree
usage: update-config [options] config-name

-V --version          print version info
-h --help            display help

--silent            no output
--quiet            only output errors
--default-output    default output
--report           slightly verbose output
--verbose          maximal output
--debug            debugging output

-d --dir DIR         cd to DIR first
--config-dir CFG    get config from directory CFG

--force            pass the --force flag to update
```

Update the named configuration based on the specification in the CFG directory (or "configs" subdirectory at the root of the project tree containing DIR).

and

```
% arx replay-config --help
replay a multi-project tree
usage: replay-config [options] config-name

-V --version          print version info
-h --help            display help

-d --dir DIR         cd to DIR first
--config-dir CFG     get config from directory CFG

--silent            no output
--quiet            only output errors
--default-output    default output
--report           slightly verbose output
--verbose          maximal output
--debug           debugging output
```

Update the named configuration based on the specification in the CFG directory (or the "configs" subdirectory at the root of the project tree in DIR).

and

```
% arx record-config --help
record a revision-specific configuration
usage: record-config [options] config-name new-config-name

-V --version          print version info
-h --help            display help

--silent            no output
--quiet            only output errors
--default-output    default output
--report           slightly verbose output
--verbose          maximal output
--debug           debugging output

-d --dir DIR         cd to DIR first
--config-dir CFG     get config from directory CFG

-f --force           overwrite an existing config
```

Record a new configuration based on the old one.

and

```
% arx show-config --help
show the revision frontier of a configuration
usage: show-config [options] config-name
```

```
-V --version          print version info
-h --help            display help

-d --dir DIR         cd to DIR first
--config-dir CFG     get config from directory CFG
```

Print the revision frontier specified by the named configuration in the CFG directory (or "configs" subdirectory at the root of the project tree containing DIR).

## 20 ArX Distribution Tools

To be documented:

```
% arx update-distributions --help
build or update a tar ball
usage: update-distributions [options] [archive/]revision dist-name dir
```

-V --version	print version info
-h --help	display help
--control	include control files in the tar file.
--changelog	include ChangeLog's in the tar file.
--cache DIR	cache directory for locally cached revisions
--config NAME	after checking out the revision, pass NAME to "build-config" to complete the source tree.
--silent	no output
--quiet	only output errors
--default-output	default output
--report	slightly verbose output
--verbose	maximal output
--debug	debugging output

Ensure that DIR exists and is populated with the indicated revision from the archive.

The tar file will be named:

```
DIST-NAME-#{VERSION}#{SUFFIX}
```

where VERSION is the version number of the revision being packaged, and SUFFIX is derived from the patch-level name of the revision.

The value of #{SUFFIX} depends on the revision's patch level, as in these examples:

revision:	tar file:
ArX--release--1.0--base-0	ArX-1.0pre0.tar.gz
ArX--release--1.0--patch-3	ArX-1.0pre3.tar.gz
ArX--release--1.0--version-0	ArX-1.0.tar.gz
ArX--release--1.0--versionfix-5	ArX-1.0.5.tar.gz

The top-level directory in the tar file has the same name as the tar file (but without the ".tar.gz" suffix).



## 21 The Pristine Revision Cache

Each project tree contains a cache of pristine revisions: revisions extracted from an archive and stored without modifications. This cache is used to speed up operations such as `commit`.

It is occasionally useful to know what revisions are cached in a given tree, and to add or remove cached revisions.

*commands to document*

```
% arx pristines --help
```

```
% arx add-pristine --help
```

```
% arx delete-pristine --help
```



## 22 Revision Tree Libraries

For many purposes, it is useful to have a library containing pristine trees of a large number of revisions – for example, all of the revisions in a particular version. To be practical, though, such a library must be represented in a space-efficient way.

Unix hard-links provide a natural way to store such a library. Each successive revision in a series is a copy of the previous, but with unmodified files shared via hard-links.

ArX provides commands to help you build, maintain, and browse such a library.

### 22.1 Your Revision Library Location

To begin a new revision library, first create a new directory (DIR) and then register its location:

```
% arx my-revision-library DIR
```

You can check the location of your library with:

```
% arx my-revision-library
```

or unregister it with:

```
% arx my-revision-library -d
```

### 22.2 Revision Library Format

A revision library has subdirectories of the form:

```
ARCHIVE-NAME/CATEGORY/BRANCH/VERSION/REVISION/
```

Each REVISION directory contains the complete source of a particular revision, along with some supplementary subdirectories and files:

```
REVISION/,,patch-set/
```

```
The patch set that creates this revision from  
its ancestor (unless the revision is a full-source  
base revision).
```

Although the permissions on files in the revision library are determined by patch sets, **you must never modify files in the revision library**. Doing so will cause odd errors and failures in various ArX commands.

## 22.3 Adding a Revision to the Library

You can add a selected revision to your revision library with:

```
% arx library-add REVISION
```

That will add not only `REVISION` to the library, but all directly preceding revisions (recursively) which are either from the same archive, or from other archives which are also in the same library. ("Directly preceding" in this case means the previous patch level, for ordinary revisions, or the tagged revision (for revisions which are tags).)

## 22.4 Finding a Revision in the Library

You can find a particular revision in the library with `library-find`:

```
% arx library-find REVISION  
PATH-TO-REVISION
```

The output is an absolute path name to the library directory containing the revision. (Once again, you must not modify files in that directory.)

## 22.5 Removing a Revision from the Library

To remove a particular revision from the library, use:

```
% arx library-remove REVISION
```

Be aware of the following limitation in the current release: suppose that you add three successive revisions, `A`, `B`, and `C`. Then you remove `B`, then re-add `B`. Now there is a chance that the file sharing between `B` and `C` will be less than optimal, causing your library to be larger than it needs to be. (You can fix this by then removing and re-adding `C`.)

## 22.6 Listing Library Contents

The command `library-archives` lists all archives with records in the library:

```
% arx library-archives
ARCHIVE-NAME
ARCHIVE-NAME
...
```

Similarly, you can list categories, branches, versions, or revisions:

```
% arx library-categories [ARCHIVE]
% arx library-branches [ARCHIVE/CATEGORY]
% arx library-versions [ARCHIVE/BRANCH]
% arx library-revisions [ARCHIVE/VERSION]
```

## 22.7 Individual Files in the Revision Library

You can locate an individual file in a revision library with:

```
% archive library-file FILE [REVISION]
PATH
```

or obtain its contents with:

```
% archive cat-library-file FILE [REVISION]
...file contents...
```

Both commands accept the options `--tag` and `--this`. With `--tag`, the argument `FILE` is interpreted as an inventory tag, and the file with that tag is found.

With `--this`, `FILE` is interpreted as a file relative to the current directory, which should be part of a project tree. The file's inventory tag is computed and the corresponding file found in `REVISION`.

## 22.8 Determining Patch Set Prerequisites

```
% arx touched-files-prereqs REVISION
```

That command looks at the patch set for `REVISION` and at all preceding patch sets in the same version (it searches your library rather than your repository for this purpose). It reports the list of patches that touch overlapping sets of files and directories – in other words, it tells you what patches can be applied independently of others. The command has an option to exclude from consideration file names matching a certain pattern (e.g. `=README` or `ChangeLog`). It has an option to exclude from the output list patches which have already been applied to a given project tree. It has an option to report the specific files which are overlapped.

## 23 ArX Triggers

In some circumstances, it is very useful to trigger actions upon the detection of changes to an archive. For example, you might want to send an email notification whenever new revisions are checked in.

**Note:** Notifiers are being revamped. This section is expected to be obsolete soon.

The command `notify` provides this capability:

```
% arx notify NOTIFICATION-DIR
```

That command reads configuration files found in `NOTIFICATION-DIR`. Those configuration files explain what parts of which repositories to monitor, and what actions to take upon changes. When changes are detected, `notify` invokes actions.

This chapter explains `notify` in general. It documents the format of a `NOTIFICATION-DIR` and the configuration files. It explains what `notify` does. Later in the manual is documentation explaining how to use `notify` to achieve certain specific effects.

### 23.1 The Four notify Configuration Files

```
|=rules.archives| |=rules.categories| |=rules.branches| |=rules.versions|
```

`notify` is controlled by four configuration files, found at the root of a `NOTIFICATION-DIR`. These are:

```
=rules.archives
    Specifies actions to take whenever new categories
    are added to specific archives.

=rules.categories
    Specifies actions to take whenever new branches
    are added to specific categories.

=rules.branches
    Specifies actions to take whenever new versions
    are added to specific branches.

=rules.versions
    Specifies actions to take whenever new revisions
    are added to specific versions.
```

Within these files, blank lines are permitted and comments begin with `'#'` and continue to the end of line.

The file `=rules.versions` contains lines with the format:

```
VERSION      ACTION
```

where `VERSION` is a fully-qualified version name, and `action` a string (with no whitespace) of the form:

COMMAND:THUNK

Whenever new revisions are detected in `VERSION`, `notify` invokes:

```
arx COMMAND THUNK RVN ...
```

where `COMMAND` and `THUNK` are from the configuration file, and the trailing `RVN` arguments are a list of fully qualified revision names of newly detected revisions.

The file `=rules.branches` is similar. It contains lines of the form:

```
BRANCH COMMAND:THUNK VERSION-ACTION
```

Whenever new versions are created in `BRANCH`, `notify` takes two actions. First, it invokes:

```
arx COMMAND THUNK VSN ...
```

Second, it adds new rules to `=rules.versions` of the form:

```
VSN VERSION-ACTION
```

The files `=rules.categories` and `=rules.archives` are similar. The former contains lines of the form:

```
CATEGORY COMMAND:THUNK BRANCH-ACTION VERSION-ACTION
```

and the latter lines of the form:

```
ARCHIVE COMMAND:THUNK CATEGORY-ACT BRANCH-ACT VERSION-ACT
```

In addition to the configuration files, `notify` maintains its private state in `NOTIFY-DIR`. In particular, it keeps track of what notifications have already been sent, in order to try to avoid sending redundant notifications.

## 23.2 Triggers for Everything in a Repository

Configuring triggers for every category, branch, and version in an archive is quite simple: you only need to write the `=rules.archives` file – other configuration files will be automatically generated from that.

The initial `=rules.archives` file should contain lines of the form:

```
ARCHIVE COMMAND:THUNK CATEGORY-ACT BRANCH-ACT VERSION-ACT
```

for each set of actions you want to trigger. After creating the config file, run `notify` in that directory. The first run will detect all categories in the repository, invoke `COMMAND` for those, and create a `rules.categories` file with lines:

```
CATEGORY CATEGORY-ACT BRANCH-ACT VERSION-ACT
```

Then `notify` will read that config file, detect all the branches in the repository, invoke `CATEGORY-ACT`, and create `=rules.branches`.

Finally, `=rules.versions` will be automatically created in a similar way.

Whenever a new category, branch, or version is added to an archive, the rules files will be automatically updated to reflect the addition (when `notify` is run, of course).

### 23.3 Robustness Issues with Triggers

Unfortunately, some fundamental physical properties of the universe make it impossible for `notify` to guarantee that actions will be invoked only once for each new category, branch, version, or revision. A (presumably rare) well timed interrupt or system failure can cause `notify` to invoke actions more than once for a given change to the archive.

Consequently, actions should be designed to be robust against that eventuality.

`notify` is designed to be safe in the face of concurrent invocations: if two processes invoke `notify` at the same time, everything should work wonderfully with no resultant redundant actions or NOTIFY-DIR corruptions.

### 23.4 Scheduling Triggers with cron

One strategy for using `notify` is to run it from a `cron` job.

### 23.5 Scheduling Triggers Directly

These commands (and only these commands) modify ArX repositories in ways that are significant to `notify`:

```
make-category
make-branch
make-version
import
commit
tag
```

Corresponding to those commands, are four others:

```
% arx my-category-action COMMAND ARGUMENTS...

% arx my-branch-action COMMAND ARGUMENTS...

% arx my-version-action COMMAND ARGUMENTS...

% arx my-revision-action COMMAND ARGUMENTS...
```

which in those invocations, set an action to be taken upon successful completion of an archive transaction, and in these invocations:

```
% arx my-category-action
```

```
% arx my-branch-action

% arx my-version-action

% arx my-revision-action
```

report the actions to be taken.

The actions have a fairly obvious meaning. After creating a new branch, for example, the command `make-branch` will invoke:

```
% COMMAND ARGUMENTS... BRANCH
```

where `BRANCH` is a fully qualified branch name.

A useful action to take is to invoke `notify` directly. Another useful action (though imposing a greater system administration burden) is to invoke a script which sends mail to some address, informing the recipient of the repository change. The recipient can be, for example, an automated process which invokes `notify`.

When the repository-modifying commands invoke one of these actions, they run the action as a `daemon`: an asynchronous process with no controlling terminal, with standard input redirected from `/dev/null`, and with `stdout` and `stderr` saved in a temporary file. If the command exits with non-0 status, the output from the command is sent to the user in an email notice.

Once again, the damn constraints of physics, particularly the prohibition against "action at a distance", impose a constraint: it is impossible to guarantee the `COMMAND` will be invoked and in (hopefully rare) circumstances, it might not be. `COMMANDs` should be chosen with that constraint in mind.

## 24 Using Triggers

This chapter describes two ways to use triggers: for sending email about changes to an archive, and for keeping a revision library up to date (see [Chapter 22 \[Revision Tree Libraries\]](#), page 99).

**Note:** Notifiers are being revamped. This section is expected to be obsolete soon.

### 24.1 Sending Email Notices About Repository Changes

ArX provides four commands designed for use as triggers that send email about changes to a repository. (See [Chapter 23 \[ArX Triggers\]](#), page 103.) These are:

```
arx mail-new-categories
arx mail-new-branches
arx mail-new-versions
arx mail-new-revisions
```

You can create a triggers for sending email with a `notify =rules.archives` configuration rule such as this (here split over multiple lines to fit the printed page – in practice, this rule would be on one line):

```
joe.hacker--2002@gnu.org
  mail-new-categories:joe@gnu.org
  mail-new-branches:joe@gnu.org
  mail-new-versions:joe@gnu.org
  mail-new-revisions:joe@gnu.org
```

## 24.2 Updating a Revision Library

You can use triggers to keep a revision library automatically up-to-date with a rule like this:

```
joe.hacker--2002@gnu.org
  noop:
  noop:
  noop:
  library-add:--ignore-extra-args
```

## 25 Automatic Triggers

**Note:** Notifiers are being revamped. This section is expected to be obsolete soon.

There are exactly six commands which modify repositories:

```
arx commit
arx import
arx make-category
arx make-branch
arx make-version
arx tag
```

Those commands can automatically invoke `arx notify` if you have a default `notify` directory set. You can set such a default with:

```
% arx my-notifier DIR
```

check your current default with:

```
% arx my-notifier
```

or delete your default setting with:

```
% arx my-notifier -d
```



## 26 Graphical User Interface

At present there is no real GUI for ArX. For emacs lovers, there is an incomplete emacs mode. It should be in the `src/user-interface/emacs` directory, along with instructions on how to install it.

You can also use your favorite gui diff when using `file-diff`. First, you must set your gui diff program using `my-guidiff`

```
arx my-guidiff tkdiff
```

Then you invoke `file-diffs` with the `--gui` option. That is, when using the default revision

```
arx file-diffs --gui file
```



## Appendix A Implementing Development Policies

Different projects have different policies for managing the "main development path" as distinguished from various kinds of release (such as candidate releases, experimental releases, and stable releases).

**ArX** is flexible enough to allow many such policies to be implemented in a direct way. Here are some examples and hints about using **ArX**.

### A.1 Milestone/Numbered Versions

One development policy is based on milestones. The team works on one milestone goal at a time. When they have just about reached that milestone, they make candidate releases of the milestone for people to test. After fixing bugs, a milestone release is made.

After several milestones, the milestone release gets a version number, and becomes a version release.

Beginning from scratch on the first milestone, the developers create the first development path – in which to work on reaching the first milestone:

```
mozilla--devo--1.0
-----
base-0
patch-1
patch-2
patch-3
patch-4
```

After some time, they are ready to make some candidate releases and enter a bug-fixing cycle. They'll use tags for that – making the bug fixes in the `devo` branch:

```
mozilla--devo--1.0
-----
base-0
patch-1
patch-2
patch-3
patch-4 ----->base-0 (tag -- 1st release candidate)
patch-5 .----->patch-1 (tag -- 2nd candidate)
patch-6 ----->patch-3 (tag -- 3rd candidate)
patch-7 |
patch-8 -----
```

In the diagram, patches 5..8 in the `devo` branch are fixes made based on reports from candidate releases. We'll suppose that after `patch-8`, the first milestone appears to be stable. The developers want to make the milestone release, and begin work on milestone 2:



```

|
| mozilla-1.0
| -----
->base-0 (tag)
    version-0 (versioned release)

```

That pattern of **ArX** usage is very simple because the developers stay in sync, following a strict cycle of working on a milestone, making candidate releases and bug-fixing, making a milestone releases and starting on the next milestone. Less synchronized development can be much more complicated, as illustrated in the next example:

## A.2 Even/Odd Versions

Another, more intricate policy is based on version numbers. Odd numbered versions are the "leading-edge" of development – often unstable, but having the very latest sources. Even numbered versions are "stable" – lagging behind the leading-edge, but containing only code known to work reasonably well.

The **ArX** concept of a "continuation version" is ideal for this, because there is no requirement that a continuation version be a continuation of the immediately preceding version.

A series of diagrams can help to illustrate this usage of **ArX** and some of the subtleties that can arise. We'll start with just the initial leading-edge development path

```

linux--0.1
-----
base-0
patch-1
patch-2
patch-3
patch-4
version-0

```

At that point, the developers decide to make the first stable release:

```

linux--0.1                linux--0.2
-----                -----
base-0                    ----> base-0 (continuation)
patch-1                   |   patch-1
patch-2                   |   version-0
patch-3                   |
patch-4                   |
version-0 -----

```

We can suppose that `patch-1` of `linux--0.2` is just a quick change to the top level `README` file, and that after that – the new stable version is sealed (creating `version-0`) and released. We never particularly want to merge `patch-1` of `linux-0.2` back on to the odd-numbered versions.

Naturally, some bugs are detected in the stable release – three in rapid succession. The developers fix these on the *leading edge* branch, planning to merge them with the stable tree later:

```

linux--0.1          linux--0.2
-----
base-0              ----> base-0 (continuation)
patch-1            |   patch-1
patch-2            |   version-0
patch-3            |
patch-4            |
version-0          -----
versionfix-1
versionfix-2
versionfix-3

```

They wait a week for some testing to occur. The bug fixes are looking ok, so the developers decide to start work on the next leading-edge version:

```

linux--0.1          linux--0.2
-----
base-0              ----> base-0 (continuation)
patch-1            |   patch-1
patch-2            |   version-0
patch-3            |
patch-4            |
version-0          -----
versionfix-1
versionfix-2
--versionfix-3
|
|
| linux--0.3
| -----
->base-0 (continuation)
  patch-1

```

Just as the developers are about to merge the bug fixes with the stable version, one more bug report trickles in. Fortunately, it's a trivial bug – so the developers are confident about making the fix in the leading edge, but immediately releasing it in the stable version.

Here's a catch, though – the latest leading edge version (`linux--0.3`) has already diverged from the stable version because of `patch-1`. The developers definitely don't want `patch-1` of 0.3 in the the stable 0.2 yet, so they fix the newly reported bug in 0.1, then merge all four bug fixes with the stable tree in the usual way. Meanwhile, separate work also continues on 0.3:

```

linux--0.1          linux--0.2
-----
base-0              ----> base-0 (continuation)
patch-1            |   patch-1
patch-2            |   version-0

```

```

    patch-3          |      ->versionfix-1 (merge)
    patch-4          |      |
    version-0 ----- |      |
    versionfix-1\    |      |
    versionfix-2 |-----|
--versionfix-3 |
| versionfix-4/
|
|
|
| linux--0.3
| -----
->base-0 (continuation)
  patch-1
  patch-2
  patch-3

```

It's worth noting at that point that `linux--0.3` is missing a bug-fix (`versionfix-4`) from `linux--0.1`:

```

% cd ~/wd/linux--0.3

% arx whats-missing linux--0.1
versionfix-4

```

We'll eventually do an update to pick up that bug fix, but first, let's make the situation more complicated.

Suppose, some mailing lists get wind of `patch-3` in `linux--0.3`. Soon `Slashdot` and `Newsforge` pick up the story. It turns out that `patch-3` is a very desirable feature and the implementation appears to be clean and stable. People are clammering for its appearance in a stable release, and the developers happen to think it is a good idea.

So, they start the next stable revision:

```

    linux--0.1          linux--0.2
    -----          -----
    base-0              ----> base-0 (continuation)
    patch-1            |      patch-1
    patch-2            |      version-0
    patch-3            |      ->versionfix-1 (merge)
    patch-4            |      |
    version-0 ----- |      |
    versionfix-1\     |      |
    versionfix-2 |-----|      linux--0.4
--versionfix-3 |      |      -----
| versionfix-4/          ---->base-0 (continuation)
|
|
|

```

```

| linux--0.3
| -----
->base-0 (continuation)
  patch-1
  patch-2
  patch-3

```

They only want patch-3 of 0.3 for 0.4 – not anything else. That’s a job for replay --exact:

```

      linux--0.1                linux--0.2
      -----                -----
base-0                ----> base-0 (continuation)
patch-1                |    patch-1
patch-2                |    version-0
patch-3                |    ->versionfix-1 (merge)
patch-4                |    |    |
version-0 -----    |    |
versionfix-1\         |    |
versionfix-2 |-----|    linux--0.4
--versionfix-3 |    |    -----
| versionfix-4/    |    ---->base-0 (continuation)
|                |    ==>patch-1 (replay --exact merge)
|                |    version-0
|                |    .
| linux--0.3        |    |
| -----          |    |
->base-0 (continuation) |    |
  patch-1            |    |
  patch-2            |    |
  patch-3-----    |    |
  ...                |

```

After merging the much-desired patch-3 of 0.3 into 0.4, the developers seal 0.4 and make the stable release.

Let’s suppose that development on 0.3 continues for a while. After some time, the developers decide that 0.3 has acquired enough new features. They want to do two things: start 0.5, and start getting the stable 0.6 release ready:

```

      linux--0.1                linux--0.2
      -----                -----
base-0                ----> base-0 (continuation)
patch-1                |    patch-1
patch-2                |    version-0
patch-3                |    ->versionfix-1 (merge)
patch-4                |    |    |
version-0 -----    |    |
versionfix-1\         |    |
versionfix-2 |-----|    linux--0.4

```

```

--versionfix-3 |           | -----
| versionfix-4/           | ---->base-0 (continuation)
|                         | ==>patch-1 (replay --exact merge)
|                         |   version-0
|                         | .
|                         | |
| linux--0.3             | |
| -----               | |
->base-0 (continuation)  | |
  patch-1                 | |
  patch-2                 | |
  patch-3-----         | |
  ...                     | |
-version-0                | |
|                         | |
|                         | |
|                         | |
| linux--0.5             | |
| -----               | |
->base-0                   | |

```

Notice that they haven't created the base revision for 0.6 yet. There's a choice here. They could make 0.6 a continuation of 0.4, then merge in all the patches they're missing from 0.3. On the other hand, they could make 0.6 a continuation of some 0.3 and pick up all those missing patches "the easy way".

But, oops, when someone checks out `version-0` from 0.3 and runs `whats-missing`, they find out that the 0.3 branch never picked up `versionfix-4` from 0.1. That's easily fixed by updating a 0.3 tree against 0.1, and checking in the resulting merge to 0.3. The resulting merge is also the revision that will become the base revision for 0.6:

```

linux--0.1             linux--0.2
-----             -----
base-0                 ----> base-0 (continuation)
patch-1                |   patch-1
patch-2                |   version-0
patch-3                |   ->versionfix-1 (merge)
patch-4                |   |
version-0 -----     |   |
versionfix-1\          |   |
versionfix-2 |-----  |   linux--0.4
--versionfix-3 |       |   -----
| versionfix-4/       |   ---->base-0 (continuation)
|                   |   ==>patch-1 (replay --exact merge)
|                   |   version-0
|                   |   .
|                   |   |

```

```

| linux--0.3      |          |
| -----      |          |
->base-0 (continuation) |
  patch-1      |          |
  patch-2      |          |
  patch-3-----
  ...          |          |
-version-0      V          | linux--0.6
| versionfix-1 (0.1 update)----->base-0 (continuation)
|                                     version-0
|
| linux--0.5
| -----
->base-0
  patch-1
  patch-2

```

Meanwhile, new work continues on 0.5.

But now, the 0.5 tree is in an interesting state. If a developer checks out the latest 0.5 and asks:

```

% arx whats-missing linux--0.1
linux--0.1--versionfix-4

```

If the developer asks `whats-missing` from 0.3, the answer is:

```

% arx whats-missing linux--0.3
linux--0.3--versionfix-1

```

If those two patches were unrelated – there would be no problem: simply update from both branches and check the result into 0.5.

In fact, though, `versionfix-1` from 0.3 is really the same change as `versionfix-4` from 0.1 (look back at how `versionfix-1` was created). Let’s also suppose that when the fix was merged into 0.3, a slight change had to be made – to resolve a merge conflict.

So if the developer just blindly updates from 0.1, then from 0.3, the second update will result in new conflicts. That might not be so bad if we’re only talking about two patches – but if we were talking about 20 or 200, a lot of needless work would be called for.

Fortunately, ArX can help. First, the developer gets the latest 0.5 revision:

```

% arx get linux--0.5 ~/wd/linux--0.5

```

Then gets a list of all patches missing from 0.1 and 0.3:

```

% cd ~/wd/linux--0.5

% arx whats-missing --full linux--0.1 linux--0.3
archive@kernel.org--primary/linux--0.1--versionfix-4
archive@kernel.org--primary/linux--0.3--versionfix-1

```

That list can be piped into the `reconcile` tool:

```
% ... | arx reconcile
archive@kernel.org--primary/linux--0.3--versionfix-1
```

What happened? `reconcile` figured out that `versionfix-1` from 0.3 already includes `versionfix-4` from 0.1 – there’s no need to apply both patches. So `patch-plan` reported the list of patches that *do* need to be applied, and in this case, there’s only one.

In a more complicated situation, `patch-plan` would print a list of patches in the order they should be applied. In general, it will be a subset of the patches in its input, but applying that subset will have the same effect as applying all of the patches (but hopefully with fewer conflicts).

The developer uses `arx replay --list` to process that list, finally winding up with:

```
linux--0.1                linux--0.2
-----
base-0                    ----> base-0 (continuation)
patch-1                   |   patch-1
patch-2                   |   version-0
patch-3                   |   ->versionfix-1 (merge)
patch-4                   |   |   |
version-0 -----       |   |
versionfix-1\             |   |
versionfix-2 |-----   |   linux--0.4
--versionfix-3 |         |   -----
| versionfix-4/         |   ---->base-0 (continuation)
|                   |   ==>patch-1 (replay --exact merge)
|                   |   version-0
|                   |   .
| linux--0.3           |   |
| -----             |   .
->base-0 (continuation) |   |
patch-1               |   .
patch-2               |   |
patch-3-----
...                   |   linux--0.6
-version-0            V   -----
| versionfix-1 (0.1 update)----->base-0 (continuation)
|                   |   version-0
|                   |
| linux--0.5           |   |
| -----             |   |
->base-0               |   |
patch-1               |   |
patch-2               V
patch-3 (0.1/0.3 reconciliation)
```

Now if someone gets the latest revision of 0.5 and asks:

```
% arx whats-missing --full linux--0.1 linux--0.3  
[no output]
```

Isn't reconcile handy?

## Appendix B The Theory of Patches and Revisions

This appendix briefly explains "patch sets" and "revision control" in the abstract.

### B.1 The Theory of Patches

A patch set is an expression of the differences between two revisions of a tree of files (usually, primarily, text files). A patch set tells you what files and directories have been added or removed between the two revisions, what files have been renamed, what files have changed. For files added or removed, a patch set tells you the complete contents of those files. For files modified, a patch set contains a description of the changes in the form of a context diff (see the man page for `diff(1)`). If a file is a symbolic link, and the link target has changed, the patch set records that fact. If a regular file or directory is replaced by a symbolic link (or vice versa) the patch records that fact. Finally, if any files have had their permissions or modification times changed – a patch set records that too.

Some notation will be helpful. A shorter name for "patch set" is delta. Let's suppose that A and B are two revisions of a source tree. Then:

```
delta (A, B)
```

is the name for a patch set describing the differences between A and B.

Anytime you make a series of changes to a tree, perhaps using shell utilities and text editors, the entire series of changes can be summarized as a single patch set. In a sense, a patch set (or "delta") is the fundamental editing operation, in terms of which all others, and all combinations of others, can be described.

You can apply a patch – which means to make the changes it describes. In our notation:

```
delta (A, B) [A] == B
```

says "the patch set describing the differences between A and B, when applied to A, gives B". A patch set can also be applied "in reverse":

```
delta (A, B) {B} == A
```

"the delta from A to B, applied in reverse to B, gives A".

A patch set can also be applied (or reverse-applied) to a tree which is not the same as either A or B. For example, suppose that we have a tree `A_prime` which is similar to A, but has some slight differences. Then:

```
delta (A, B) [A_prime] ~== B_prime
```

where `~==` means "approximately equals". When a patch set is applied to a tree which is not one of the trees used to compute the patch set, the edits might or might not go well. For example, if the patch set wants to modify a file F, but `A_prime` doesn't contain F, the patch can't be applied perfectly. Similarly, if the patch set wants to modify F, but the version of F in `A_prime` is already very different from the version in A, then those edits can't be done automatically.

Nevertheless, for the kinds of changes people typically make to trees of source code, approximately applied patch sets are very useful. For example, suppose we start with tree *A*, and create two revisions *B\_one* and *B\_two*. Then:

$$\text{delta (A, B\_one) [B\_two]} \sim == \text{delta (A, B\_two) [B\_one]}$$

To make that more concrete: if programmer Alice makes a set of changes to give us *B\_one*, and programmer Bob makes a set of changes to give us *B\_two*, then the patch sets between *A* and those two revisions of *B* give us a way to get *B\_three* – a tree that contains *both* Alice's and Bob's changes. Even when a patch set can't be perfectly applied in that way, it can often be applied to do "80%" of the work, making it much easier to finish merging the two sets of changes by hand.

Incidentally, patch sets have a useful algebraic property if we think of them as functions that can be composed. Using the notation  $F \circ G$  to mean "the function *F* composed with the function *G*":

For all trees, *A*, *B*, and *C*:

$$\text{delta (B, C) } \circ \text{ delta (A, B) } == \text{ delta (A, C)}$$

so to build *C* from *A*, we can use:

$$\text{delta (A, C) [ A ] } == C$$

but that is the same as:

$$(\text{delta (B, C) } \circ \text{ delta (A, B)}) [ A ] == C$$

or, in other words:

$$\text{delta (B, C) [ delta (A, B) [ A ] ] } == C$$

The algebraic property suggests that if I want to apply:

$$\text{delta (B, C) } \circ \text{ delta (A, B)}$$

I can save time by instead applying:

$$\text{delta (A, C)}$$

and if I want to have a record of all three:

$$\begin{aligned} &\text{delta (A, B)} \\ &\text{delta (B, C)} \\ &\text{delta (A, C)} \end{aligned}$$

I can save space by storing only:

$$\begin{aligned} &\text{delta (A, B)} \\ &\text{delta (B, C)} \end{aligned}$$

And if we're applying patch sets to trees that might need to be touched up by hand, and I want to apply  $\text{delta (A, C)}$ , then I have a choice between applying:

`delta (A, C)`

and have just one, possibly large set of errors to clean-up by hand, or applying:

`delta (A, B) then delta (B, C)`

and having two, but possibly smaller sets of errors to clean-up by hand.

Patch sets have many uses, but three important ones are:

**Compression** One use is *compression*. If **A** and **B** are large trees, and the differences between them small, then the patch set between them will be much smaller than either tree. You can save disk space by not storing both **A** and **B**, but instead, storing only **A** (or only **B**) along with `delta (A, B)`. That makes patch sets an ideal form for space-efficient archival of multiple revisions of a tree.

Similarly, if someone has downloaded a copy of **A** (or **B**) and they want a copy of **B** (or **A**), they can save download time and bandwidth by downloading only `delta (A, B)`. That makes patch sets a good way to distribute multiple revisions of a tree.

**Inspection** Another use for patch sets is *inspection*. If a patch set is stored in a human-readable format, it provides a useful way to quickly see precisely what has changed between two revisions of a tree. For example, a patch set is handy for reviewing the changes made by programmers to a large source tree.

**Combining Separate Efforts** For some kinds of trees, patch sets are good at merging (combining) changes made by people working separately (as in the example of Alice and Bob, above). This is especially true of program source code. That makes patch sets a very handy tool for making a team of programmers more effective – allowing the work separately up to a point, then combine their efforts by creating and applying patch sets.

## B.2 The Theory of Revisions

Suppose that we start with a tree, **A0**, and make a set of changes resulting in the tree **A1**:

**A0**  
**A1**

We can repeat that process several times:

**A0**  
**A1**  
**A2**  
**A3**  
...

Each instance of the tree is called a revision. Between each revision and its successor, we can compute a patch set:

`delta (A0,A1) [A0] == A1`  
`delta (A1,A2) [A1] == A2`  
`delta (A2,A3) [A2] == A3`  
...

Something we can usefully do is create an archive of revisions. We might store the first tree verbatim, and every successive tree as a delta:

```
A0: "complete copy of tree"
A1: delta (A0,A1)
A2: delta (A1,A2)
A3: delta (A2,A3)
...
```

If we want to retrieve an  $A_n$ , we start with  $A_0$  and apply the first  $n$  deltas:

```
delta (An-1, An) [delta (An-2, An-1) [...[A0]...]] = An
```

or, making our notation more concise:

```
An [ An-1 [ An-2 [ ... [A0] ...]]] = An
```

Each revision in a series like this is called a patch level. The entire series is called a development path.

At any point along the way, we might make a copy of some  $A_n$ , and start a new development path. For example, we might copy  $A_2$  to form  $B_0$ :

```
A0
A1
A2 -----> B0
A3                B1
...                B2
                   ...
```

When we have multiple, related development paths, each is called a branch. The tree we copied to start a new branch (e.g.  $A_2$ ) is called a branch point.

If we're building an archive, we can store  $B_0$  as a pointer to the  $A_0$  development path, and every successive revision of the  $B_0$  path as an ordinary delta:

```
A0                B0: "equal to A2"
A1: delta (A0,A1)  B1: delta (B0, B1)
A2: delta (A1,A2)  B2: delta (B1, B2)
A3: delta (A2,A3)  B3: delta (B2, B3)
...                ...
```

To make all this more concrete, imagine that the  $A_0$  development path is successive revisions of a program we're working on. Alice wants to add a very complicated feature. Rather than make many small changes to the  $A_0$  development path, she makes a branch, the  $B_0$  development path – and works on the complicated feature there. Each new revision of the  $B_0$  branch is a small step on the way to the complicated feature.

Eventually, Alice is done with the feature – but meanwhile, the  $A_0$  development path has added several changes of its own. What we'd like to do next is to make a revision of the tree that has both sets of changes – from both development paths.

```

A0                                B0: "equal to A2"
A1: delta (A0,A1)                 B1: delta (B0, B1)
A2: delta (A1,A2)                 B2: delta (B1, B2)
A3: delta (A2,A3)                 B3: delta (B2, B3)
A4: delta (A3,A4)                 B4: delta (B3, B4)

```

How can we make revision A5, which includes all the changes made on both branches?

The "theory of patches" gives us several possible solutions.

One solution is to make this tree:

```
B4 [ B3 [ B2 [ B1 [ A4 ] ] ] ]
```

That solution is called replaying the patches from the B0 branch against the A0 branch.

That *might* work reasonably, but patch set B1 wasn't formed from A4 – it was formed from B0 which is the same as A2. So when we apply B1 to A4, there might be problems that have to be resolved "by hand". The same will happen again when we apply B2, B3, and B4. In some situations, the risk and complexity of doing all that work by hand is worth it – but not in other situations. What other options do we have?

Another solution is to make this tree:

```
A4 [ A3 [ B4 ] ]
```

That solution is also replaying: replaying the patches from the A0 branch against the B0 branch. The same kind of problem might occur (having to fix things up by hand), but in this case, we're only applying two patches instead of four – so this might be a simpler solution.

Here's a third solution:

```
delta (A2, B4) [A4]
```

That solution is based on the fact that:

```
B0 == A2
```

and the algebraic property that:

```
B4 o B3 o B2 o B1 == delta (B0, B4)
```

That solution is called updating the B0 development path with respect to the A0 path. The difference between "replaying" and "updating" is a little bit subtle. When we "replay" from another development path, that means that we take all patches we're missing from that other path, and apply them in order. When we "update" from another development path, that means we take the latest revision on that other path, and apply to it a delta between the branch point and our own most-recent revision.

Applying the patch during an update certainly can fail to work perfectly – it might require fixing up by-hand. On the other hand, an update only ever applies one patch; often,

therefore, the amount of by-hand repairs is minimized. "Nine times out of ten," updating is the preferred technique for joining two previously branched revisions.

There are other, more obscure solutions too. To choose one arbitrarily, we might try building:

```
B4 [ A3 [ B2 [ A4 [ B1 [ B3 [ A2 ]]]]]]
```

A bizarre solution like that is so rare it doesn't really have name – but "one time in ten thousand" – it's the solution that works best.

No matter what solution we choose, if we store the resulting revision back on the B0 path, we'll wind up with:

```
A0                                B0: "equal to A2"
A1: delta (A0,A1)                 B1: delta (B0, B1)
A2: delta (A1,A2)                 B2: delta (B1, B2)
A3: delta (A2,A3)                 B3: delta (B2, B3)
A4: delta (A3,A4)                 B4: delta (B3, B4)
                                   B5: delta (B4, B5) "has changes A3, A4"
```

We can store that same revision back on the A0 development path:

```
A0                                B0: "equal to A2"
A1: delta (A0,A1)                 B1: delta (B0, B1)
A2: delta (A1,A2)                 B2: delta (B1, B2)
A3: delta (A2,A3)                 B3: delta (B2, B3)
A4: delta (A3,A4)                 B4: delta (B3, B4)
A5: delta (A5, A4) ==             B5: delta (B4, B5) "has changes A3, A4"
```

A5 and B5 are called a merge point. For all practical purposes, a merge point is also a branch point – since, using the example, B5 and A5 are equal, just the the two revisions of the original branch point (A2 and B0) were equal. If additional development happens on the two branches, we no longer have to worry about merging all changes since A2 and B0; we can instead just merge only the changes since A5 and B5.

### **B.3 What is a Revision Control System?**

So what is a revision control system?

A revision control system is a set of tools for computing and applying patch sets, for archiving patch sets, for distributing patch sets, and for helping to merge changes on the basis of patch sets.

A revision control system has to come up with a reasonable way of naming and cataloging revisions. It has to be able to represent branch points and help with merges. When merges occur, a good revision control system should help figure out what patches to apply to which revisions in order to minimize hand-editing.



## Appendix C ArX Patch Sets

It is often extremely useful to compare two project trees (usually for the same project) and figure out exactly what has changed between them. A record of such changes is called a patch set or a delta.

If you have a patch set between an "old tree" and a "new tree", you can "apply the patch" to the old tree to get the new tree – in other words, you can automatically make the editing changes described by a patch set. If you have some third tree, you can apply the patch to get an approximation of making the same changes to that third tree. (see [Appendix B \[The Theory of Patches and Revisions\]](#), page 123).

ArX includes sophisticated tools for creating and applying patch sets. In general, you will not use the commands in this section. You will use higher level commands like `update`, `replay`, or `star-merge`.

### C.1 `mkpatch`

`mkpatch` computes a patch set describing the differences between two trees. The basic command syntax is:

```
% mkpatch ORIGINAL MODIFIED DESTINATION
```

which compares the trees `ORIGINAL` and `MODIFIED`.

`mkpatch` creates a new directory, `DESTINATION`, and stores the patch set there.

When `mkpatch` compares trees, it uses inventory tags. For example, it considers two directories or two files to be "the same directory (or file)" if they have the same tag – regardless of where each is located in its respective tree.

A patch set produced by `mkpatch` describes what files and directories have been added or removed, which have been renamed, which files have been changed (and how they have been changed), and what file permissions have changed (and how). When regular text files are compared, `mkpatch` produces a context diff describing the differences. `mkpatch` can compare binary files (saving complete copies of the old and new versions if they differ) and symbolic links (saving the old and new link targets, if they differ).

A detailed description of the format of a patch set is provided in an appendix (see [Appendix D \[The ArX Patch Set Format\]](#), page 137).

## C.2 dopatch

`dopatch` is used to apply a patch set to `tree`:

```
% dopatch patch-set tree
```

If `tree` is exactly the same as the the "original" tree seen by `mkpatch`, then the effect is to modify `tree` so that it is exactly the same as the the "modified" tree seen by `mkpatch`, with one exception (explained below).

"Exactly the same" means that the directory structure is the same, symbolic link targets are the same, the contents of regular files are the same, and file permissions are the same. Modification times, files with multiple (hard) links, and file ownership are not reliably preserved.

The exception to the "exactly the same" rule is that if the patch requires that files or directories be removed from `tree`, those files and directories will be saved in a subdirectory of `tree` with an eye-splitting name matching the pattern:

```
++removed-by-dopatch-PATCH--DATE
```

where `PATCH` is the name of the patch-set directory and `DATE` a timestamp.

`dopatch` also supports reverse patching.

## C.3 Inexact Patching

What if a tree patched by `dopatch` is not exactly the same as the original tree seen by `mkpatch`?

Below is a brief description of what to expect. Complete documentation of the `dopatch` process is included with the source code.

`dopatch` takes an inventory of the tree being patched. It uses inventory tags to decide which files and directories expected by the patch set are present or missing from the tree, and to figure out where each file and directory is located in the tree.

**Simple Patches** If the patch set contains an ordinary patch or metadata patch for a link, directory or file, and that file is present in the tree, `dopatch` applies the patch in the ordinary way. If the patch applies cleanly, the modified file, link, or directory is left in place.

If a simple patch fails to apply cleanly, `dopatch` will always leave behind a `.orig` file (the file originally in the tree being patched, without any changes) and a `.rej` file (the part of the patch that could not be applied).

If the patch was a context diff, `dopatch` will also leave behind the file itself – partially patched.

If an (unsuccessful) patch was for a binary file, no partially-patched file will be left. Instead, there will be:

```
.orig  -- the file originally in the tree being patched,
        without modifications.
```

```
.rej    -- a complete copy of the file from the modified tree,
         with permissions copied from '.orig'.

.patch-orig -- a complete copy of the file from the original
              tree seen by 'mkpatch', with permissions
              retained from that original

         -or-

         the symbolic link from the original tree seen
         by 'mkpatch' with permissions as in the original
         tree.
```

If an (unsuccessful) patch was for a symbolic link, no partially patched file will be left. Instead there will be:

```
.orig    -- the unmodified file from the original tree

.rej     -- a symbolic link with the target intended by the
         patch and permissions copied from .orig

.patch-orig -- a complete copy of the file from the original
              tree seen by 'mkpatch', with permissions
              retained from that original

         -or-

         the symbolic link from the original tree seen
         by 'mkpatch' with permissions as in the original
         tree.
```

### Patches for Missing Files

All patches for missing files and directories are stored in a subdirectory of the root of the tree being patched called

```
==missing-file-patches-PATCH-DATE
```

where PATCH is the basename of the patch set directory and DATE a time-stamp.

### Directory Rearrangements and New Directories

Directories are added, deleted, and rearranged much as you would expect, even if you don't know it's what you'd expect.

Suppose that when mkpatch was called the ORIGINAL tree had:

Directory or file:	Tag:
a/x.c	tag_1
a/bar.c	tag_2

but the MODIFIED tree had:

```

a/x.c          tag_1
a/y.c          tag_2

```

with changes to both files. The patch will want to rename the file with tag `tag_2` to `y.c`, and change the contents of the files with tags `tag_1` and `tag_2`.

Suppose, for example, that you have a tree with:

```

a/foo.c        tag_1
a/zip.c        tag_2

```

and then you apply the patch to that tree. After the patch, you'll be left with:

```

a/foo.c        tag_1
a/y.c (was zip.c) tag_2

```

with patches made to the contents of both files.

Here's a sample of some subtleties and ways of handling conflicts:

Suppose that the original tree seen by `mkpatch` has:

```

Directory or file:      Tag:
./a                     tag_a
./a/b                   tag_b
./a/b/c                 tag_c

```

and that the modified directory has:

```

./a                     tag_a
./a/c                   tag_c
./a/c/b                 tag_b

```

Finally, suppose that the tree has:

```

./x                     tag_a
./x/b                   tag_b
./x/c                   tag_new_directory
./x/c/b                 tag_diffent_file_named_b
./x/c/q                 tag_c

```

When patch gets done with the tree, it will have:

```

./x                     tag_a
    Since the patch doesn't do anything
    to change the directory with tag_a.

./x/c.orig              tag_new_directory
./x/c.rej               tag_c
    Since the patch wants to make the
    directory with tag_c a subdirectory named "c"
    of the directory with tag_a, but the tree
    already had a different directory there,
    with the tag tag_new_directory.

```

```
./x/c.rej/b                                tag_b
Since the patch wants to rename the directory
with tag_b to be a subdirectory named "b"
of the directory with tag_c.
```

```
./x/c.orig/b                               tag_diffent_file_named_b
Since the patch made new changes to this file,
it stayed with its parent directory.
```



## Appendix D The ArX Patch Set Format

An ArX patch set is a directory containing a number of files and subdirectories. Each is described below.

**Files:**

```
orig-dirs-index
mod-dirs-index
orig-files-index
mod-files-index
```

**Format:**

```
<file><space><tag>
```

**Sorting:**

```
sort -k 2
```

These contain indexes for all files and directories added, removed, or modified between the two trees.

**Files:**

```
original-only-dir-metadata
modified-only-dir-metadata
```

**Format:**

```
<metadata><space><name>
```

**Sorting:**

```
sort -k 2
```

The field `<metadata>` contains literal output from the program `file-metadata` given the options `--permissions`. Some example output is:

```
--permissions 777
```

That output is also suitable for use as options and option arguments to the program `set-file-metadata`. Future releases of ArX might add additional flags (beside just `permissions`).

The list records the file permissions for all directories present in only one of the two trees.

**Directories:**

```
removed-files-archive
new-files-archive
```

Each of these directories contains complete copies of all files that occur in only the original tree (`removed-files-archive`) or modified tree (`new-files-archive`). Each saved file is

archived at the same relative location it had in its source tree, with permissions (at least) preserved.

**Directory:**

patches

This directory contains a tree whose directory structure is a subset of the directory structure of the modified tree. It contains modification data for directories and files common to both trees.

For a file stored in the modified tree at the path `new_name`, the `patches` directory may contain:

`new_name.link-orig`

The original file is a symbolic link.  
 'new\_name.link-orig' is a text file containing the target of that link plus a final newline.

This file is only present if link target has changed, or if the link was replaced by a regular file.

`new_name.link-mod`

The modified file is a symbolic link and this file is a text file containing the target for the link plus a final newline.

This file is only present if the link target has changed, or if the link replaces a regular file.

`new_name.original`

This is a complete copy of the file from the original tree, preserving (at least) permissions.

This file is only present if the file was replaced by a symbolic link, or if the file contents can not be handled by 'diff(1)'.  
 handled by 'diff(1)'.  
 handled by 'diff(1)'.

`new_name.modified`

This is a complete copy of the file from the modified tree, preserving (at least) permissions.

This file is only present if the file replaces a symbolic link, or if the file contents can not be handled by 'diff(1)'.

`new_name.patch`

This is a standard context diff between the original file and modified file. One popular version of diff ('GNU diff') generates non-standard context diffs by omitting one copy of lines of context that are identical between the original and modified file, so for now, '.patch' files may have the same bug. Fortunately, the only popular version of 'patch' ('GNU patch') is tolerant of receiving such input.

```
new_name.meta-orig  
new_name.meta-mod
```

File metadata (currently only permissions) changed between the two versions of the file. These files contain output from the 'file-metadata' program with the flags '--symlink --permissions', suitable for comparison to similar output, and for use as options and option arguments to 'set-file-metadata'.

These files are also included if a regular file has replaced a symbolic link or vice versa.

```
new_name/=dir-meta-orig  
new_name/=dir-meta-mod
```

Directory metadata (currently only permissions) changed between the two versions of the directory containing these files. These files contain output from the 'file-metadata' program with the flags '--symlink --permissions', suitable for comparison to similar output, and for use as options and option arguments to 'set-file-metadata'.

**Note:** If a regular file (or symbolic link) replaces a directory, or vice versa, this is recorded as a file (or link) removed (or added) in one tree and added (or removed) in the other.



## Appendix E The ArX Archive Format

An ArX archive is a directory containing a number of files and subdirectories. Its structure is described in this appendix.

### E.1 Directory Structure

Each category, branch, version, and revision are given a separate directory. These are nested. From the root of the archive:

```
CATEGORY/
  CATEGORY/BRANCH/
    CATEGORY/BRANCH/VERSION/
      CATEGORY/BRANCH/VERSION/REVISION/
```

For example, version 1.0 of the devo branch of the ArX category might have these directories:

```
ArX/
  ArX/ArX--devo/
    ArX/ArX--devo--1.0/
      ArX/ArX--devo--1.0/base-0/ # the base revision
      ArX/ArX--devo--1.0/patch-1/ # pre-patch revisions
      ArX/ArX--devo--1.0/patch-2/ # ...
      ...
      ArX/ArX--devo--1.0/version-0/ # the version revision
      ArX/ArX--devo--1.0/versionfix-1/ # post-patch revisions
      ArX/ArX--devo--1.0/versionfix-2/ # ...
      ...
```

### E.2 Within a Revision Directory

Every revision directory contains the log message for that revision as a plain-text file:

**File:** log

**Format:** RFC822-style headers plus body

If a revision directory is a full-source revision (typically a `base-0` revision) it will contain a compressed tar file of the entire revision tree. The name of the tar file is the full name (sans archive name) of the revision:

**File:** REVISION.tar.gz

**Format:** gzip-compressed tar file containing a complete source tree, rooted in a single top-level directory named REVISION.

**Example:** ArX-devo-1.0-base-0.tar.gz

```
contains a full-source tree rooted at
'ArX--devo--1.0--base-0'
```

If a revision directory is a patch (not a full-source revision), then the revision directory contains a compressed tar file of the patch set (see [Appendix D \[The ArX Patch Set Format\]](#), page 137):

**File:** REVISION.patches.tar.gz

**Format:** gzip-compressed tar file containing a patch set, rooted in a single top-level directory named REVISION.

**Example:** ArX-devel-1.0-patch-1.tar.gz

contains a patch-set tree rooted at 'ArX--devo--1.0--patch-1'

**Note:** Every revision is either a full source revision or a patch revision. Thus, every revision directory contains exactly one of the two files:

REVISION.tar.gz

REVISION.patches.tar.gz

If a revision is a continuation revision (a tag of some other revision)

## Appendix F Idempotent Merging

[Idempotent merging is not implemented in the current release.]

Let's suppose that we have a main development path, with several branches:

```

main
----
base-0
patch-1
patch-2-----> branch-a
patch-3          |
patch-4          |
patch-5          |-----> branch-b
                  |
                  |-----> branch-c

```

What happens if each of the three branches checks in a revision that is an **update** against the **main** branch. In other words, each branch will have a delta that summarizes patches 3..5 of the main branch.

If we try to **replay** from two or more of those branches, we'll wind up replaying several of those deltas that summarize 3..5. Those patches will be redundant and will likely generate merge conflicts.

**Update** won't do much better in this case. The common ancestor of all three branches is their **base-0** revision, which is the same as **patch-2** on **main**.

Now suppose I start with the latest revision on **branch-a**, which includes patches 3..5 of the main branch, plus some changes specific to **branch-a**. And **branch-b** is similar – it has 3..5 from **main** some changes specific to **branch-b**.

If I **update** my **branch-A** revision against **branch-B**, the A revision is compared to the common ancestor. In essence:

```
update_patch = delta (branch-a--latest, main--patch-2)
```

Note that the update patch contains all the changes needed for 3..5 from **main**. **update** will apply that patch to the latest revision of **branch-B**:

```
update_a_from_b = update_patch [ branch-b--latest ]
```

but **branch-b--latest** already includes patches 3..5 from **main**. There's a good chance the merge will have conflicts.

When such messes occur, the **reconcile** command, introduced in the previous chapter, can help you out of them. But wouldn't it be better to avoid such problems in the first place?

## F.1 The i-merge Command

A tool for tackling the problem directly is an idempotent merge:

```
% arx i-merge [ --update [ARCHIVE/]REVISION
                | --replay [ARCHIVE/]REVISION ]
                [ARCHIVE/]SOURCE-REVISION
                directory
```

That creates a project tree in `DIRECTORY` by using `arx get` to obtain the `SOURCE-REVISION`, then applying each of the specified `update` and `replay` commands, in the order specified.

If any merge conflicts occur, the command issues an error, and leaves the partially merged directory, along with an explanation of where it left off.

If the command succeeds, though, the project tree will be left in a special state which permits the use of the `--idempotent` flag to commit.

```
% arx commit --idempotent
```

which, in fact, creates two new revisions. The first revision created is the intermediate directory, containing only `SOURCE-REVISION` plus the series of `updates` and `replays` you specified – no other changes. The log message for this revision is automatically generated, and has the special header `idempotent-merge:` with the list of patches applied.

The second revision contains the log message you wrote, plus any subsequent changes you made.

When `reply` wants to apply a patch set, it checks to see if it is an idempotent patch set. If it is, and all of the patches included in the patch set are missing from the tree being patched, `replay` proceeds in the usual way: by applying the set of deltas in the patch set.

If some of the patches included in an idempotent merge have already been applied to the tree being patched, then `replay` applies only those patches not already included.

One possible policy is that every branch should merge only from the main branch, and should always merge from the main branch using an idempotent update:

```
% cd ~/wd

% arx i-merge --update main branch-a branch-a-merged
[...]
```

Each branch will then contain a number of idempotent patch sets, as in this example:

branch-a	branch-b
-----	-----
base-0	base-0
patch-1	patch-1
patch-2..."idempotent merge w/main patches 2,3"	patch-2
patch-3 "idempotent merge w/main patch 2"...	patch-3
patch-4	patch-4
patch-5..."idempotent merge w/main patch 4"	patch-5

```

patch-6          "idempotent merge w/main patch 3"...patch-6
patch-7..."idempotent merge w/main patch 5"          patch-7
patch-8          patch-8
                "idempotent merge w/main patches 4,5"...patch-9
                patch-10

```

What if we want to form a merge of these two branches?

## F.2 idempotent Merges and the replay Command

We can start with a project tree for the latest revision of `branch-a`

```

% arx get ~/wd/branch-a
% cd ~/wd/branch-a

```

`Branch-a` does not already have a patch log for `branch-b`, though the two branches have a common ancestor, so `add-sibling-log` will solve that problem:

```

% arx add-sibling-log branch-b
[....]

```

Now we can find out what the merge needs to do:

```

% arx whats-missing branch-b
patch-1
patch-2
patch-3
patch-4
patch-5
patch-6
patch-7
patch-8
patch-9
patch-10

```

If we use `replay`:

```

% cd ~/wd
% arx replay ~/wd/branch-a ~/wd/branch-a-merged branch-b

```

These patches will be applied:

```

branch-b/patch-1
branch-b/patch-2
branch-b/patch-4
branch-b/patch-5
branch-b/patch-7
branch-b/patch-8
branch-b/patch-10

```

Patches 3, 6, 9 are skipped (though their log entries are added to the project tree) because `branch-a` already has all of the patches those idempotent patch sets include.

### F.3 idempotent Patch Sets and the update Command

What will `update` do? It computes a patch between the project tree being updated and the common ancestor:

```
update_patch = delta (branch-a--patch-8, branch-b--base-0)
```

then applies that to the update revision:

```
branch-a-mege = update_patch [ branch-b--patch-10 ]
```

The idempotent revisions don't help there. However....

### F.4 idempotent Patch Sets and Partial Updates

The `--partial` flag to the `update` command takes advantage of idempotent patch sets. As before, we assume these revisions:

branch-a	branch-b
-----	-----
base-0	base-0
patch-1	patch-1
patch-2..."idempotent merge w/main patches 2,3"	patch-2
patch-3 "idempotent merge w/main patch 2"...	patch-3
patch-4	patch-4
patch-5..."idempotent merge w/main patch 4"	patch-5
patch-6 "idempotent merge w/main patch 3"...	patch-6
patch-7..."idempotent merge w/main patch 5"	patch-7
patch-8	patch-8
	"idempotent merge w/main patches 4,5"...
	patch-9
	patch-10

We have a tree for `branch-a--patch-8` and it's ancestor on `branch-b` is `base-0`. So:

```
% cd ~/wd/branch-a--patch-8
% arx whats-missing branch-b
patch-1
patch-2
patch-3
patch-4
patch-5
[...]
patch-10
```

If we use:

```
% arx update --partial ~/wd/branch-a--patch-8 \
~/wd/branch-a-b-partial-update \
branch-b
```

then `update` uses not the latest revision of `branch-b`, but the revision just prior to the oldest idempotent patch set that `branch-a` does not yet have. In this case, `patch-3` of `branch-b` is the oldest idempotent patch that `branch-a` is missing. So:

```
% arx update --partial ~/wd/branch-a--patch-8 \  
    ~/wd/branch-a-b-partial-update \  
    branch-b
```

is equivalent to:

```
% arx update --partial ~/wd/branch-a--patch-8 \  
    ~/wd/branch-a-b-partial-update \  
    branch-b--patch-2
```

After which:

```
% cd ~/wd/branch-a-b-partial-update  
% arx whats-missing branch-b  
patch-3  
patch-4  
patch-5  
[...]  
patch-10
```

When the earliest missing patch prior to a partial update is an idempotent patch set, and the tree already has all of the patches included in that patch set, `update` simply adds the log message for the idempotent patch to the tree being patched, and stops.

If the tree being patched has *none* of the patches included in that patch set, `update` updates against the idempotent revision in the normal way.

Finally, if the tree being patch has *some* but not all of the patches included in the idempotent patch set, `update` gives up with an error and suggests that you use `replay` to apply the idempotent patch set. (In a future release, `update` will do something more intelligent in this case).

The upshot of this is that you can merge branches A and B with a series of partial updates and replays:

```
# update against patch-level 2  
#  
% arx update --partial ~/wd/branch-a--patch-8  
    ~/wd/branch-a-b-partial-update  
    branch-b  
  
# replay patch-level 3  
# -- we already have all the patches included in the  
# idempotent branch-b patch, "patch-3" - so all this  
# really does is install a log message.  
#  
% arx replay ~/wd/branch-a-b-partial-update  
    ~/wd/branch-a-b-partial-update-2  
    branch-b--patch-3
```

```
# update against patch-level 5
#
% arx update --partial ~/wd/branch-a-b-partial-update-2
                        ~/wd/branch-a-b-partial-update-3
                        branch-b

% arx replay ~/wd/branch-a-b-partial-update-3
             ~/wd/branch-a-b-partial-update-4
             branch-b--patch-5

% arx update --partial ~/wd/branch-a-b-partial-update-5
                        ~/wd/branch-a-b-partial-update-6
                        branch-b

% arx replay ~/wd/branch-a-b-partial-update-6
             ~/wd/branch-a-b-partial-update-7
             branch-b--patch-9

% arx update --partial ~/wd/branch-a-b-partial-update-9
                        ~/wd/branch-a-b-merged
                        branch-b
```

# Appendix G The GNU Free Documentation License

GNU Free Documentation License  
Version 1.1, March 2000

Copyright (C) 2000 Free Software Foundation, Inc.  
59 Temple Place, Suite 330, Boston, MA 02111-1307 USA  
Everyone is permitted to copy and distribute verbatim copies  
of this license document, but changing it is not allowed.

## 0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other written document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

## 1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you".

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (For example, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, whose contents can be viewed and edited directly and straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup has been designed to thwart or discourage subsequent modification by readers is not Transparent. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML designed for human modification. Opaque formats include PostScript, PDF, proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

## 2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

## 3. COPYING IN QUANTITY

If you publish printed copies of the Document numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a publicly-accessible computer-network location containing a complete Transparent copy of the Document, free of added material, which the general network-using public has access to download anonymously at no charge using public-standard network protocols. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

#### 4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has less than five).
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.
- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice

- giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
  - H. Include an unaltered copy of this License.
  - I. Preserve the section entitled "History", and its title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
  - J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
  - K. In any section entitled "Acknowledgements" or "Dedications", preserve the section's title, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
  - L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
  - M. Delete any section entitled "Endorsements". Such a section may not be included in the Modified Version.
  - N. Do not retitle any existing section as "Endorsements" or to conflict in title with any Invariant Section.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties--for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

## 5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections entitled "History" in the various original documents, forming one section entitled "History"; likewise combine any sections entitled "Acknowledgements", and any sections entitled "Dedications". You must delete all sections entitled "Endorsements."

## 6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

## 7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, does not as a whole count as a Modified Version of the Document, provided no compilation copyright is claimed for the compilation. Such a compilation is called an "aggregate", and this License does not apply to the other self-contained works thus compiled with the Document, on account of their being thus compiled, if they are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one quarter of the entire aggregate, the Document's Cover Texts may be placed on covers that surround only the Document within the aggregate. Otherwise they must appear on covers around the whole aggregate.

## 8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License provided that you also include the original English version of this License. In case of a disagreement between the translation and the original English version of this License, the original English version will prevail.

## 9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

#### 10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

#### ADDENDUM: How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

Copyright (c) YEAR YOUR NAME.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.1 or any later version published by the Free Software Foundation; with the Invariant Sections being LIST THEIR TITLES, with the Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST. A copy of the license is included in the section entitled "GNU Free Documentation License".

If you have no Invariant Sections, write "with no Invariant Sections" instead of saying which ones are invariant. If you have no Front-Cover Texts, write "no Front-Cover Texts" instead of "Front-Cover Texts being LIST"; likewise for Back-Cover Texts.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.



## Appendix H The GNU General Public License

The Hackerlab C Library is free software: you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation (and reproduced below).

This software is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License (reproduced below) for more details.

### GNU GENERAL PUBLIC LICENSE Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.

59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

#### Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

GNU GENERAL PUBLIC LICENSE  
TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.

b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an

announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

- a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

#### NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

## END OF TERMS AND CONDITIONS

## How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

```
<one line to give the program's name and a brief idea of what it
does.>
Copyright (C) <year> <name of author>
```

```
This program is free software; you can redistribute it and/or
modify it under the terms of the GNU General Public License as
published by the Free Software Foundation; either version 2 of
the License, or (at your option) any later version.
```

```
This program is distributed in the hope that it will be useful,
but WITHOUT ANY WARRANTY; without even the implied warranty of
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
GNU General Public License for more details.
```

```
You should have received a copy of the GNU General Public
License along with this program; if not, write to the Free
Software Foundation, Inc., 59 Temple Place, Suite 330, Boston,
MA 02111-1307 USA
```

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

Gnomovision version 69, Copyright (C) year name of author  
Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type  
'show w'. This is free software, and you are welcome to  
redistribute it under certain conditions; type 'show c' for  
details.

The hypothetical commands 'show w' and 'show c' should show the  
appropriate parts of the General Public License. Of course, the  
commands you use may be called something other than 'show w' and  
'show c'; they could even be mouse-clicks or menu items--whatever  
suits your program.

You should also get your employer (if you work as a programmer) or  
your school, if any, to sign a "copyright disclaimer" for the  
program, if necessary. Here is a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the  
program 'Gnomovision' (which makes passes at compilers) written by  
James Hacker.

<signature of Ty Coon>, 1 April 1989  
Ty Coon, President of Vice

This General Public License does not permit incorporating your  
program into proprietary programs. If your program is a subroutine  
library, you may consider it more useful to permit linking  
proprietary applications with the library. If this is what you want  
to do, use the GNU Library General Public License instead of this  
License.



## Appendix I Preliminary Data Sheet for ArX

### Package:

ArX (a revision control system)

### Function:

ArX performs revision control offering fancy features for branching and merging. It provides distributed repositories (repositories spread over multiple hosts) and a global (world-wide) namespace for branches and revisions.

### Key Features:

ArX is simple, small, and featureful.

Distributed repositories are especially appropriate for projects developed in the "open source" style -- with geographically distributed developers and sub-teams, loosely and flexibly cooperating.

ArX's fancy merging features are ideal for projects supporting multiple concurrent releases and/or performing a lot of development in separate branches.

PDF and HTML documentation are included.

### Licensing:

ArX is distributed under the terms of the GNU General Public License, Version 2, as published by the Free Software Foundation.

Prerequisites:

- standard C and C++ compiler
- Posix libc.a (standard C library)
- GNU Make
- GNU tar
- GNU patch
- most of the standard Posix shell, text, and file utilities

Recommended and Disrecommended Applications:

ArX is brand new. The recommended application at this stage is evaluation, help with porting, and help with testing.

Given resources for robust testing, ArX could be made ready for heavy-duty, mission critical application within months -- but absent those resources, ArX suffers from the risks associated with brand-new software.

If your projects are characterized by distributed development, consider evaluating ArX and finding ways to contribute to a polished, commercial-quality release.

Limitations:

The primary limitation is that ArX needs porting and testing.

A secondary limitation is that ArX likes to operate on whole-trees, not individual files. This limitation can be overcome with further development, though it remains to be seen if user's truly miss such features.

Size:

The core of ArX is around 40K lines of code, mostly shell and awk scripts.

ArX relies on the Hackerlab C library, which adds considerably to the overall code size. It is possible to eliminate this dependency, but not necessarily desirable.

Performance:

ArX has some nice performance characteristics. At the moment, ArX can still be quite slow, but there are many possible ways to speed it up.

Repositories are stored in compressed format and tree-deltas are efficiently represented. This both saves disk space and reduces network traffic.

On the client side, ArX makes heavy use of caching to speed up some operations and make detached operation possible.



## 27 Indexes



## Short Contents

1	Introducing ArX . . . . .	3
2	System Requirements . . . . .	11
3	Tutorial . . . . .	13
4	ArX Commands in General . . . . .	17
5	ArX Project Trees . . . . .	25
6	ArX Project Inventories . . . . .	27
7	The ArX Global Name-space of Users . . . . .	39
8	The ArX Global Name-space of Projects . . . . .	41
9	Archives . . . . .	45
10	Development Paths . . . . .	49
11	Basic Revision Control . . . . .	51
12	Basic Branching and Merging . . . . .	63
13	Patch Logs and ChangeLogs . . . . .	67
14	Star Topology Branching and Merging . . . . .	71
15	Writing Log Entries for Merges . . . . .	77
16	Arbitrary Patching with delta-patch . . . . .	79
17	Multi-Branch Merging – The reconcile Command . . . . .	81
18	Reverting Development . . . . .	89
19	Multi-Tree Projects . . . . .	91
20	ArX Distribution Tools . . . . .	95
21	The Pristine Revision Cache . . . . .	97
22	Revision Tree Libraries . . . . .	99
23	ArX Triggers . . . . .	103
24	Using Triggers . . . . .	107
25	Automatic Triggers . . . . .	109
26	Graphical User Interface . . . . .	111
A	Implementing Development Policies . . . . .	113
B	The Theory of Patches and Revisions . . . . .	123
C	ArX Patch Sets . . . . .	131
D	The ArX Patch Set Format . . . . .	137
E	The ArX Archive Format . . . . .	141
F	Idempotent Merging . . . . .	143
G	The GNU Free Documentation License . . . . .	149
H	The GNU General Public License . . . . .	159
I	Preliminary Data Sheet for ArX . . . . .	169

27 Indexes ..... 173

# Table of Contents

<b>1</b>	<b>Introducing ArX</b> .....	<b>3</b>
1.1	Advantages of ArX .....	3
1.2	Global Revision Control Done Right .....	5
1.3	Introducing ArX Project Trees .....	6
1.4	Introducing ArX Inventories .....	6
1.5	Introducing ArX Patch Sets .....	7
1.6	Global Namespaces .....	7
1.7	Introducing ArX Archives .....	8
1.8	Introducing ArX Patch Logs .....	8
1.9	Cheap Branching and Smart Merging .....	9
1.10	What Does It All Mean? .....	9
<b>2</b>	<b>System Requirements</b> .....	<b>11</b>
<b>3</b>	<b>Tutorial</b> .....	<b>13</b>
3.1	Creating the first revision .....	13
3.2	Revisions .....	14
3.3	Branches .....	15
3.4	Wrapping up .....	15
3.5	Revision Trees .....	16
<b>4</b>	<b>ArX Commands in General</b> .....	<b>17</b>
4.1	The ArX Commands .....	18
<b>5</b>	<b>ArX Project Trees</b> .....	<b>25</b>
5.1	Initializing a Project Tree .....	25
<b>6</b>	<b>ArX Project Inventories</b> .....	<b>27</b>
6.1	Choices Regarding Inventories .....	27
6.2	Specifying a Tagging Method .....	28
6.3	The inventory Command .....	28
6.4	Using an Explicit Inventory .....	31
6.5	Using an Implicit Inventory .....	32
6.6	Recognizing Renames – Inventory Tags .....	33
6.7	Keeping Things Neat and Tidy .....	34
6.8	The Inventory Tag Abstraction in Detail .....	34
6.9	A Warning About Changing Tagging Methods .....	35
6.10	Other Ways to Tag Files .....	35
6.11	Telling tree-lint to Shut Up .....	36
6.12	Which Tagging Method Should You Use? .....	36
6.13	Altering the Naming Conventions .....	37

<b>7</b>	<b>The ArX Global Name-space of Users . . . . .</b>	<b>39</b>
<b>8</b>	<b>The ArX Global Name-space of Projects . . .</b>	<b>41</b>
	8.1 The Structure of Project Names . . . . .	41
	8.2 Archive Names . . . . .	42
	8.3 Category Names and Branch Labels . . . . .	43
	8.4 Version Numbers . . . . .	43
	8.5 Labelling Project Trees . . . . .	44
	8.6 Combining Project Trees . . . . .	44
<b>9</b>	<b>Archives . . . . .</b>	<b>45</b>
	9.1 Archive Names Revisited . . . . .	45
	9.2 Creating a New Archive . . . . .	45
	9.3 Mapping Archive Names to Locations . . . . .	46
	9.4 Remote Archives . . . . .	46
	9.4.1 HTTP . . . . .	46
	9.4.1.1 Webdav . . . . .	46
	9.4.1.2 Explicit lists . . . . .	47
	9.4.2 SFTP . . . . .	47
	9.4.3 Accessing the Archives . . . . .	47
	9.5 Your Default Archive . . . . .	48
<b>10</b>	<b>Development Paths . . . . .</b>	<b>49</b>
	10.1 Creating a Development Path . . . . .	49
	10.2 Examining an Archive . . . . .	49
	10.3 Fully Qualified Version Names . . . . .	50
<b>11</b>	<b>Basic Revision Control . . . . .</b>	<b>51</b>
	11.1 The First Revision . . . . .	51
	11.2 Successive Revisions . . . . .	52
	11.3 Patch Levels . . . . .	53
	11.4 Tagging . . . . .	53
	11.5 Development Phases . . . . .	54
	11.6 Getting a Revision . . . . .	55
	11.7 Optimizing Archives for get . . . . .	56
	11.8 Finding Out What Changed . . . . .	56
	11.9 The whats-missing Command . . . . .	57
	11.10 Update . . . . .	58
	11.11 Replay . . . . .	59
	11.12 The Next Version . . . . .	60
<b>12</b>	<b>Basic Branching and Merging . . . . .</b>	<b>63</b>
	12.1 Creating a Branch . . . . .	63
	12.2 Distributed Branches . . . . .	64
	12.3 whats-missing Revisited . . . . .	64
	12.4 update and replay Revisited . . . . .	65
	12.5 Merging After a Branch . . . . .	65

<b>13</b>	<b>Patch Logs and ChangeLogs . . . . .</b>	<b>67</b>
13.1	Branches and Patch Logs . . . . .	68
13.2	Comparing Patch Logs to Archives . . . . .	68
13.3	ChangeLogs . . . . .	69
<b>14</b>	<b>Star Topology Branching and Merging . . . . .</b>	<b>71</b>
14.1	The Star Topology Merge Problem . . . . .	72
14.2	Solving the Star Topology Merge Problem . . . . .	73
14.3	The star-merge Command . . . . .	74
<b>15</b>	<b>Writing Log Entries for Merges . . . . .</b>	<b>77</b>
<b>16</b>	<b>Arbitrary Patching with delta-patch . . . . .</b>	<b>79</b>
<b>17</b>	<b>Multi-Branch Merging – The reconcile Command . . . . .</b>	<b>81</b>
17.1	The Repeated Multi-Branch Merge Problem . . . . .	81
17.2	The Challenge . . . . .	84
17.3	The Simple update Solutions . . . . .	84
17.4	The Simple replay Solutions . . . . .	84
17.5	The reconcile Solution . . . . .	85
<b>18</b>	<b>Reverting Development . . . . .</b>	<b>89</b>
<b>19</b>	<b>Multi-Tree Projects . . . . .</b>	<b>91</b>
<b>20</b>	<b>ArX Distribution Tools . . . . .</b>	<b>95</b>
<b>21</b>	<b>The Pristine Revision Cache . . . . .</b>	<b>97</b>
<b>22</b>	<b>Revision Tree Libraries . . . . .</b>	<b>99</b>
22.1	Your Revision Library Location . . . . .	99
22.2	Revision Library Format . . . . .	99
22.3	Adding a Revision to the Library . . . . .	100
22.4	Finding a Revision in the Library . . . . .	100
22.5	Removing a Revision from the Library . . . . .	100
22.6	Listing Library Contents . . . . .	101
22.7	Individual Files in the Revision Library . . . . .	101
22.8	Determining Patch Set Prerequisites . . . . .	102

<b>23</b>	<b>ArX Triggers .....</b>	<b>103</b>
23.1	The Four notify Configuration Files .....	103
23.2	Triggers for Everything in a Repository .....	104
23.3	Robustness Issues with Triggers .....	105
23.4	Scheduling Triggers with cron .....	105
23.5	Scheduling Triggers Directly .....	105
<b>24</b>	<b>Using Triggers .....</b>	<b>107</b>
24.1	Sending Email Notices About Repository Changes .....	107
24.2	Updating a Revision Library .....	108
<b>25</b>	<b>Automatic Triggers .....</b>	<b>109</b>
<b>26</b>	<b>Graphical User Interface .....</b>	<b>111</b>
<b>Appendix A</b>	<b>Implementing Development Policies</b>	<b>113</b>
	.....	
A.1	Milestone/Numbered Versions .....	113
A.2	Even/Odd Versions .....	115
<b>Appendix B</b>	<b>The Theory of Patches and</b>	
	<b>Revisions .....</b>	<b>123</b>
B.1	The Theory of Patches .....	123
B.2	The Theory of Revisions .....	125
B.3	What is a Revision Control System? .....	129
<b>Appendix C</b>	<b>ArX Patch Sets .....</b>	<b>131</b>
C.1	mkpatch .....	131
C.2	dopatch .....	132
C.3	Inexact Patching .....	132
<b>Appendix D</b>	<b>The ArX Patch Set Format ...</b>	<b>137</b>
<b>Appendix E</b>	<b>The ArX Archive Format .....</b>	<b>141</b>
E.1	Directory Structure .....	141
E.2	Within a Revision Directory .....	141
<b>Appendix F</b>	<b>Idempotent Merging .....</b>	<b>143</b>
F.1	The i-merge Command .....	144
F.2	idempotent Merges and the replay Command .....	145
F.3	idempotent Patch Sets and the update Command .....	146
F.4	idempotent Patch Sets and Partial Updates .....	146

Appendix G	The GNU Free Documentation License .....	149
Appendix H	The GNU General Public License .....	159
Appendix I	Preliminary Data Sheet for ArX .....	169
27	Indexes .....	173

