Configuration Management

Object Oriented Programming

Cataldo Basile

<cataldo.basile@polito.it>

SoftEng

http://softeng.polito.it

Version 1.0.1 - March 2018

Versioning

Thesis.docx

ThesisFinal.docx

ThesisFinal Final.docx

ThesisFinalest Final.docx

ThesisFinalest FinalForsure.docx

ThesisFinalestF**k FinalForsure.docx

Issues

- What is the history of a document?
  - versioning
- What is the correct set of documents for a specific need?
  - configuration
- Who can access and change what?
  - change control
- How the system is obtained?
  - build

Configuration Management

- A discipline applying technical and administrative direction and surveillance to:
  - identify and document the functional and physical characteristics of a configuration item,
  - control changes to those characteristics,
  - record and report change processing and implementation status, and
  - verify compliance with specified requirements

[IEEE Std 828-2012]

Goals of CM

- Identify and manage parts of software
- Control access and changes to parts
- Allow to rebuild previous version of software

Terms

- Configuration item (CI)
- Configuration Management aggregate
- Configuration
- Version
- Baseline

Configuration Item (CI)

- *Aggregation of work products that is treated as a single entity in the configuration management process*
- CI (typically a file):
  - Has a name
  - All its version are numbered and kept
  - User decides to change version number with specific operation (commit)
  - It is possible to retrieve any previous version

Version

- The initial release or a re-release of a configuration item
- Instance of CI, e.g.
  - Req document 1.0
  - Req document 1.1

Version identification

- Procedures for version identification should define an unambiguous way of identifying component versions
- Basic techniques for component identification
  - Version numbering
  - Attribute-based identification

Version numbering

- Simple naming scheme uses a linear derivation
  e.g. V1, V1.1, V1.2, V2.1, V2.2 etc.
- Actual derivation structure is a tree or a network rather than a sequence
- Names are not meaningful
- Hierarchical naming scheme may be better

Configuration

- Set of CIs, each in a specific version

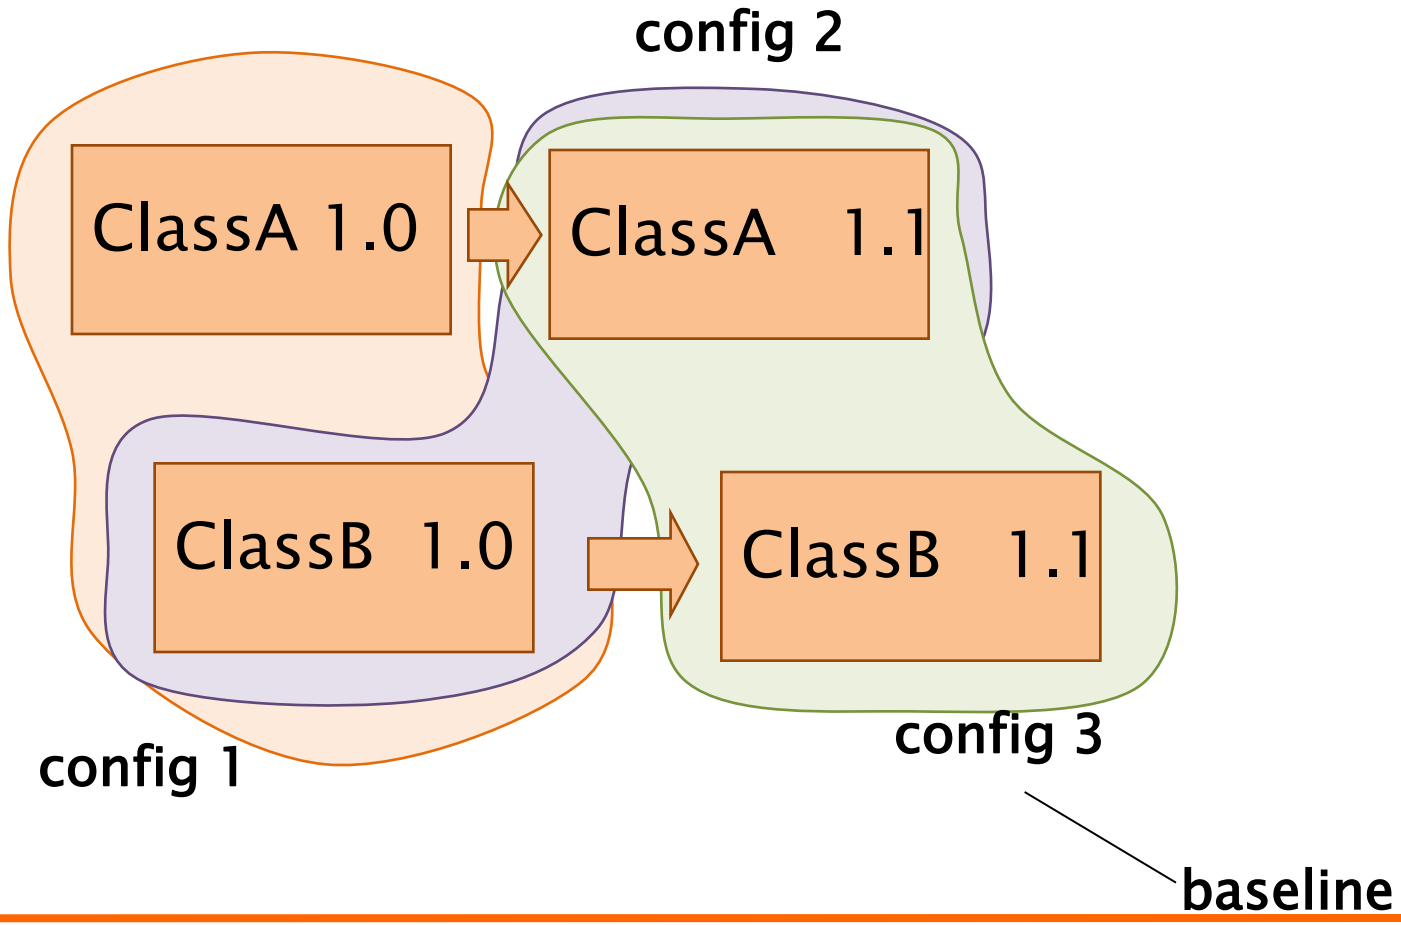

Configuration

- Snapshot of software at certain time
  - Various CIs, each in a certain version
  - Same CI may appear in different configurations
  - Also configuration has version

Baseline

- Configuration in stable, frozen form
  - Not all configurations are baselines
  - Any further change / development will produce new version(s) of CI(s), will not modify baseline
- Types of baselines
  - Development – for internal use
  - Product – for delivery

CHANGE CONTROL

Repository

- A collection of all software-related artifacts belonging to a system
- The location/format in which such a collection is stored

Typical case

- Team develops software
- Many people need to access different parts of software
  - Common repository (shared folder),
  - Everybody can read/write documents/files

File system limitations

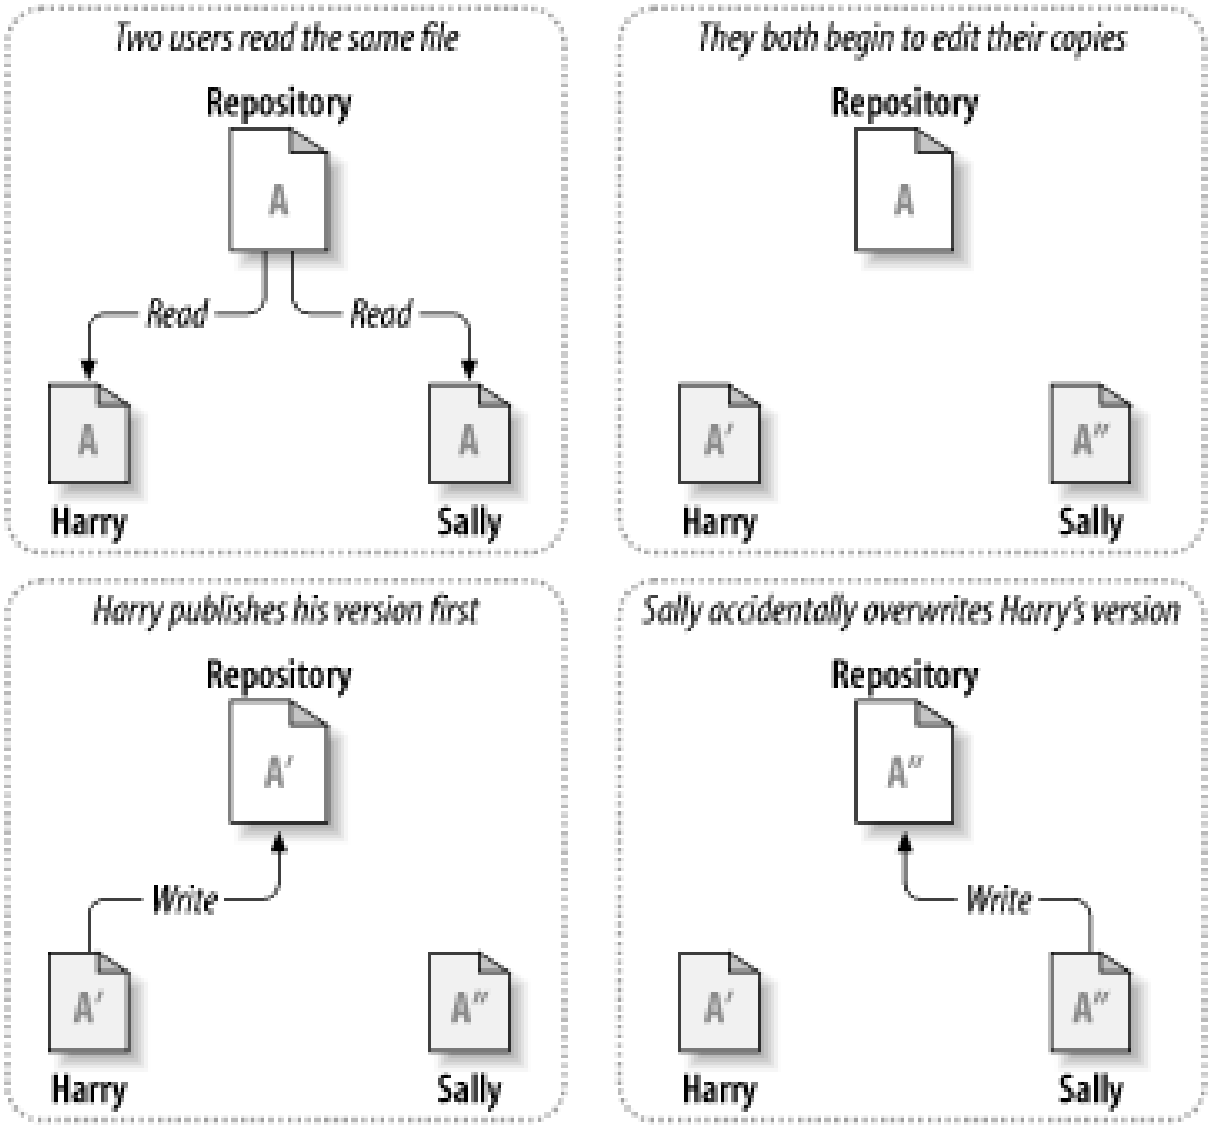

Check-in / check-out

- Check-out
  - Extraction of CI from repository
    - with goal of changing it or not
    - After checkout next users are notified
- Check-in (or commit)
  - Insertion of CI under control

Check-in / check-out - scenarios

- Lock-modify-unlock (or serialization)
  - Only one developer can change at a time
- Copy-modify-merge
  - Many change in parallel, then merge

Lock-Modify-Unlock

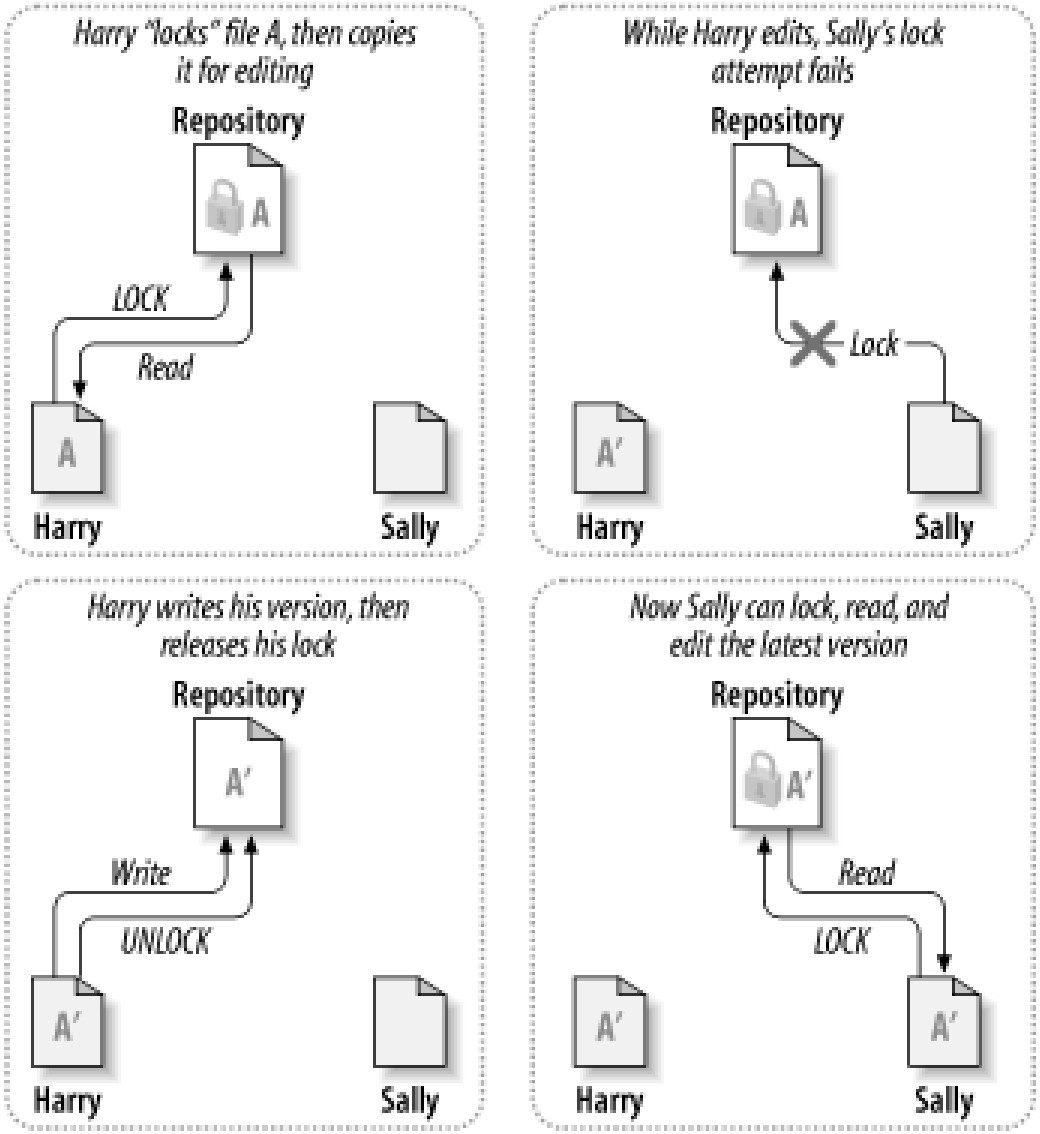

Lock-Modify-Unlock

- Pro
  - Conflicts are impossible
- Cons
  - No parallel work is possible, large delays can be induced
  - Developers can possibly forget to unlock so blocking the whole team

Copy-Modify-Merge

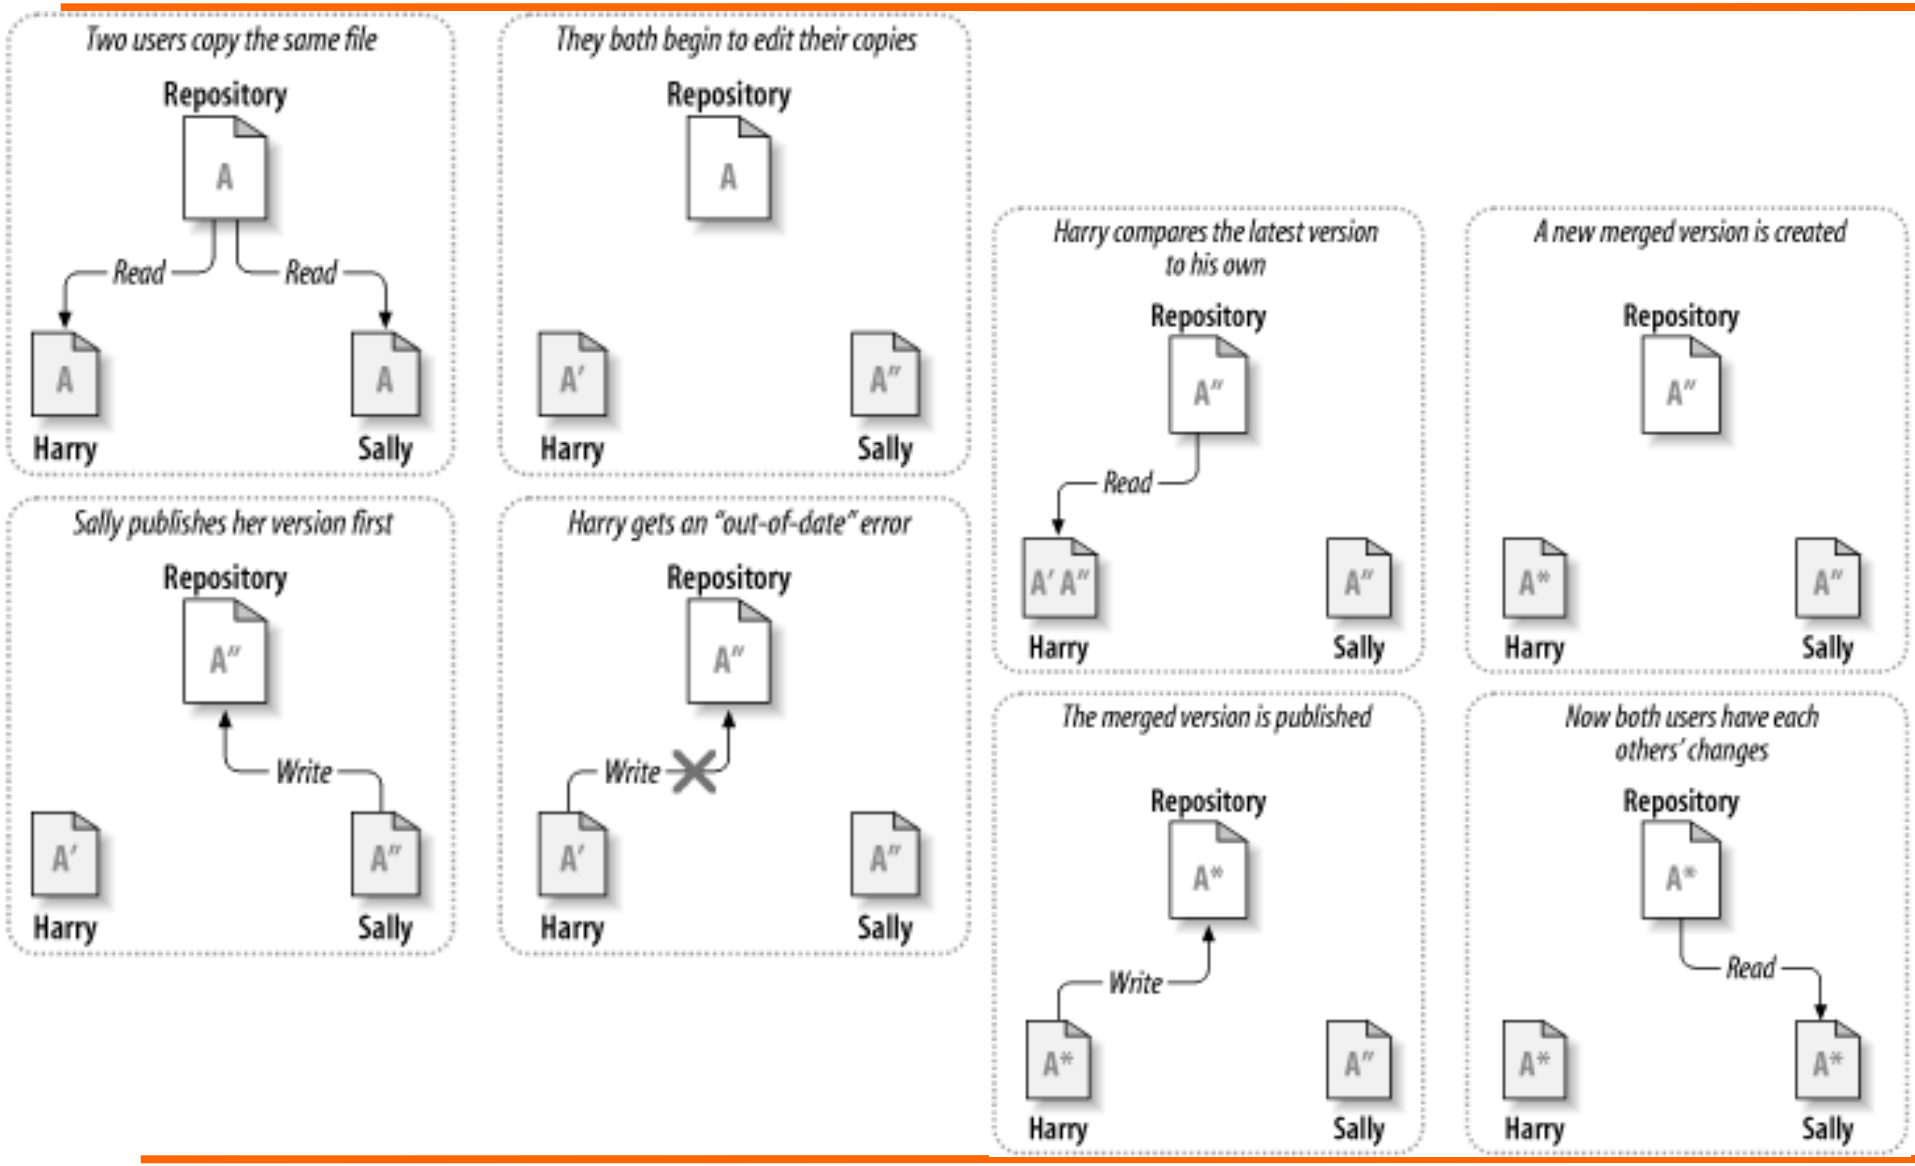

Copy-Modify-Merge

- Pros
  - More flexible
  - Several developers can work in parallel
  - No developer can block others
- Con
  - Requires care to resolve the conflicts

Tools

- CM + VM
  - RCS
  - CVS
  - SCCS
  - PCVS
  - Subversion
  - BitKeeper
  - Git

VERSION CONTROL WITH SUBVERSION

What is Subversion

- Free/open-source version control system:
  - it manages any collection of files and directories over time in a central repository;
  - it remembers every change ever made to your files and directories;
  - it can access its repository across networks

Features

- Directory versioning and true version history
- Atomic commits
- Metadata versioning
- Several topologies of network access
- Consistent data handling
- Branching and tagging
- Usable by other applications and languages

Architecture

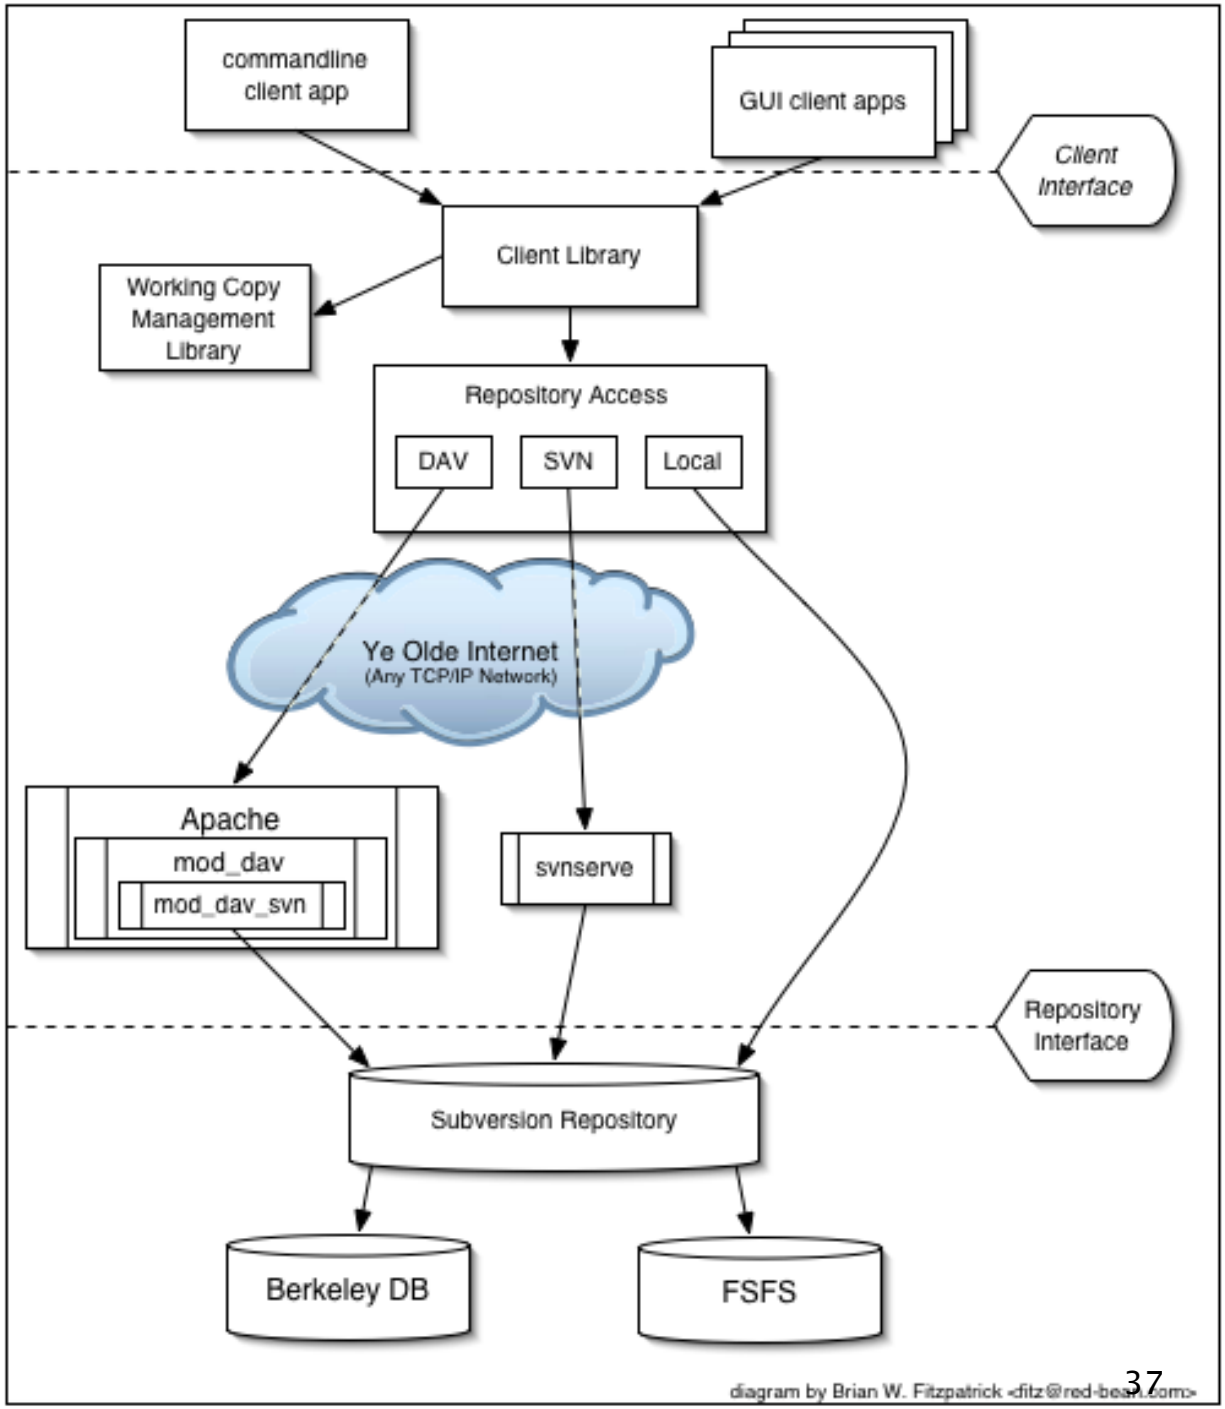

The repository

- Central store of data
- It stores information in the form of a file system
- Any number of clients connect to the repository, and then
  - read (**update**) or
  - write (**commit**) to these files.

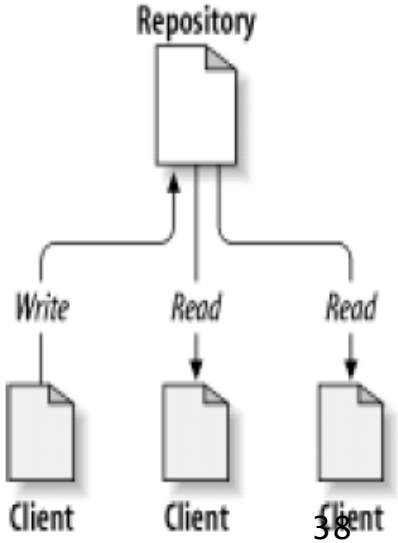

The working copy (WC)

- Ordinary directory tree on your local system, containing a copy of the repository files (checkout)
- Subversion will never incorporate other people's changes (update), nor make your own changes available to others (commit), until you explicitly tell it to do so.

Revisions

- Each time the repository accepts a commit, it creates a new state of the file system tree, called a revision.
- Global revision numbers: each revision is assigned a progressive unique natural number (previous revision + 1)
  - A freshly created repository has revision 0 (zero)
- The whole repo gets a new revision number
  - Revision *N* represents the state of the repository after the *N*th commit.

Mixed revisions

- Suppose you have a working copy entirely at revision 10. You edit the file foo.html and then perform an svn commit, which creates revision 15 in the repository.
- Therefore the only safe thing the Subversion client can do is mark the one file—foo.html—as being at revision 15. The rest of the working copy remains at revision 10. This is a mixed revision.
- Only by running svn update can the latest changes be downloaded, and the whole working copy be marked as revision 15.
- Memento:
  - Every time you run svn commit, your working copy ends up with some mixture of revisions: the things you just committed are marked as having larger working revisions than everything else.

Basic Procedure

- Create working copy from a repository
  - **svn checkout** *<repository>*

  When ready…
- Synchronize contents of WC with repo
  - **svn update**

  Work on WC
- Possibly add new files
  - **svn add** *<file list>*
- Push work to repository
  - **svn commit -m** *"<Log message>"*

Commit Log Message

- Structure of the message

```
<type>(<scope>): <subject>
<body>
<footer>
```

- Example

```
fix(middleware): ensure Range headers
adhere more closely to RFC 2616

Added one new dependency, use `range-
parser` (Express dependency) to compute
range. It is more well-tested in the
wild.

Fixes #2310
```

http://karma-runner.github.io/1.0/dev/git-commit-msg.html

Conflicts

- A conflict arise, upon commit, if the file has been updated in the meanwhile
  - N: the revision (BASE) that was modified
    - the repo revision at the time of last update
  - M: the current revision (HEAD) in the repository ($\geq$N)
- A conflict occurs if:
  - M $>$ N and
  - Contents of revisions M and N differ

Conflicts

- Subversion places three extra unversioned files in the working copy:
  - **`filename.mine`** : the local file as it existed in the working copy before the update
    - This file has only the latest local changes in it.
  - **`filename.r`***`OLDREV`* : the file that was the BASE revision before the update.
    - The file checked out before any local edit.
  - **`filename.r`***`NEWREV`* : the file that Subversion client just received from the server upon update.
    - The HEAD revision of the repository.
- The original file contains a mix version of HEAD (**`.r`***`NEW`*) and BASE (**`.mine`**) with change markers

Conflict example

- You and Sally both edit file **sandwich.txt** at the same time. Sally commits her changes, and when you go to update your working copy, you get a conflict

```
$ svn update
Conflict discovered in 'sandwich.txt'.
Select: (p)postpone,(df)diff-full,(e)edit,
        (h)elp for more options : p
C  sandwich.txt
Updated to revision 2.
```

Conflict example

- In your working copy you get

```
$ ls
sandwich.txt
sandwich.txt.mine
sandwich.txt.r1
sandwich.txt.r2
```

- You're going to have to edit **`sandwich.txt`** to resolve the conflicts

Conflict example

- The contents of the file **`sandwich.txt`** is

```
Top piece of bread
Mayonnaise
Lettuce
<<<<<<< .mine
Salami
Mortadella
Prosciutto
=======
Sauerkraut
Grilled Chicken
>>>>>>> .r2
Creole Mustard
Bottom piece of bread
```

Changes your made in the conflicting area

Changes Sally previously committed in the area

Conflict example

- The updated file **`sandwich.txt`** you create and saved is

```
Top piece of bread
Mayonnaise
Lettuce
Mortadella
Prosciutto
Grilled Chicken
Creole Mustard
Bottom piece of bread
```

Conflict example

- Once the conflict has been composed you ought to signal it has been resolved

```
$ svn resolve --accept working sandwich.txt
Resolved conflicted state of 'sandwich.txt'
$ svn commit -m "Pick and choosen."
```

Typical work cycle

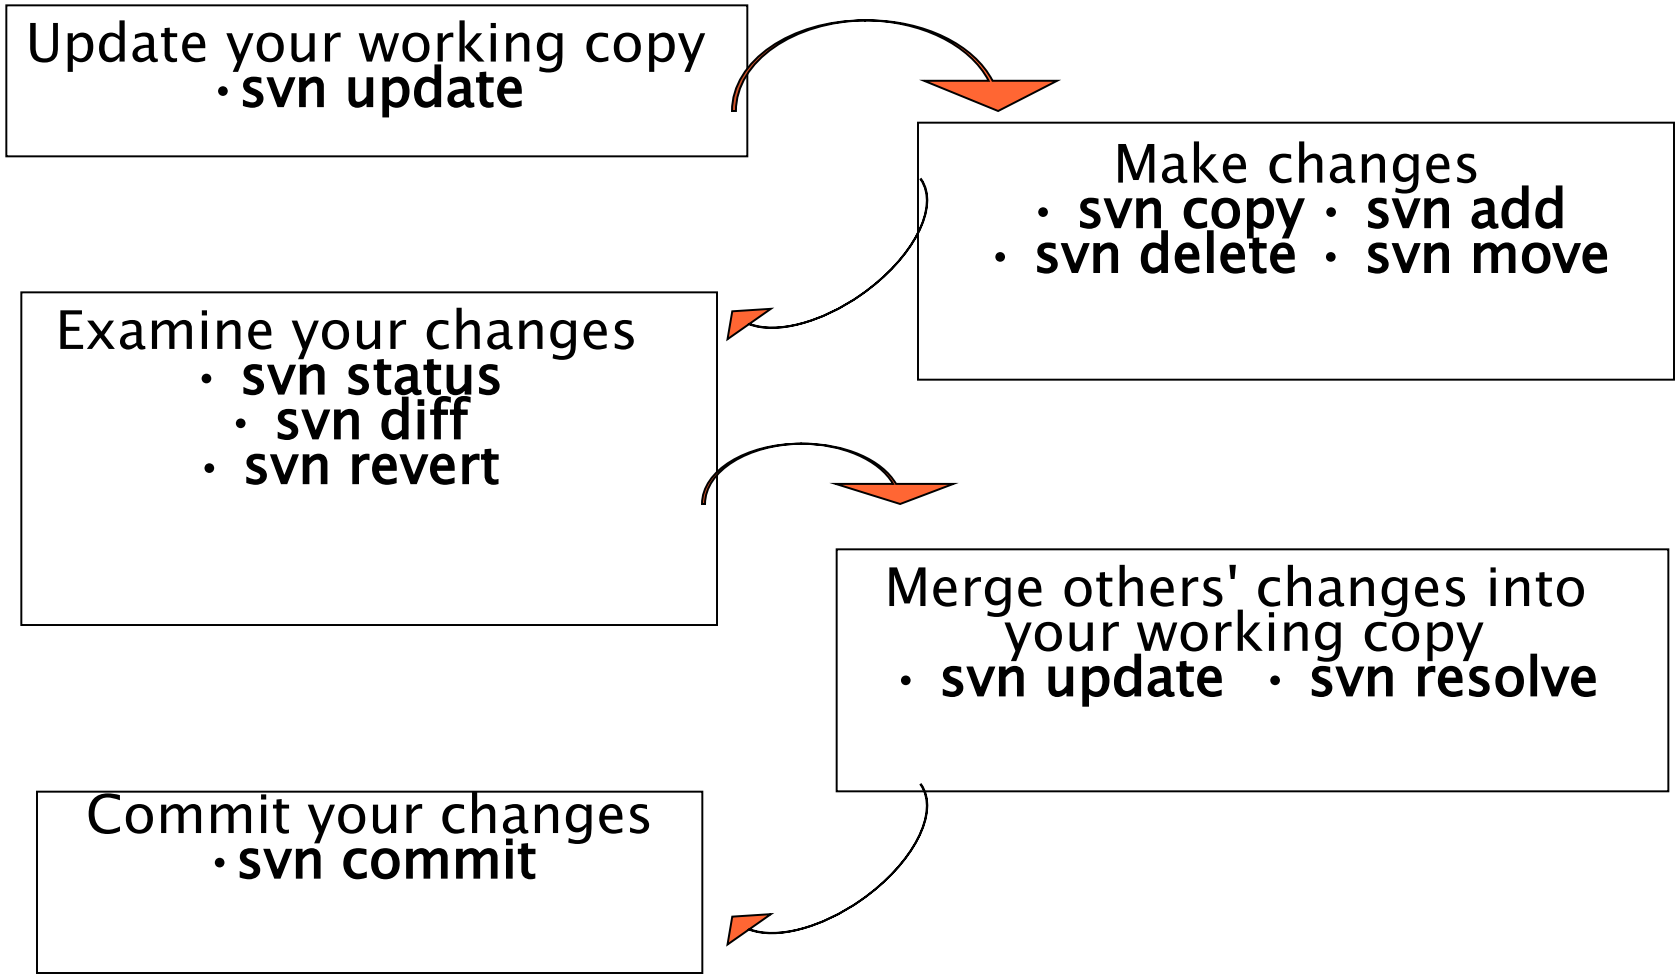

Branches: general concept

- Line of development that exists independently of another line, yet still shares a common history if you look far enough back in time.
- A branch always takes life as a copy of something, and moves on from there, independently generating its own history

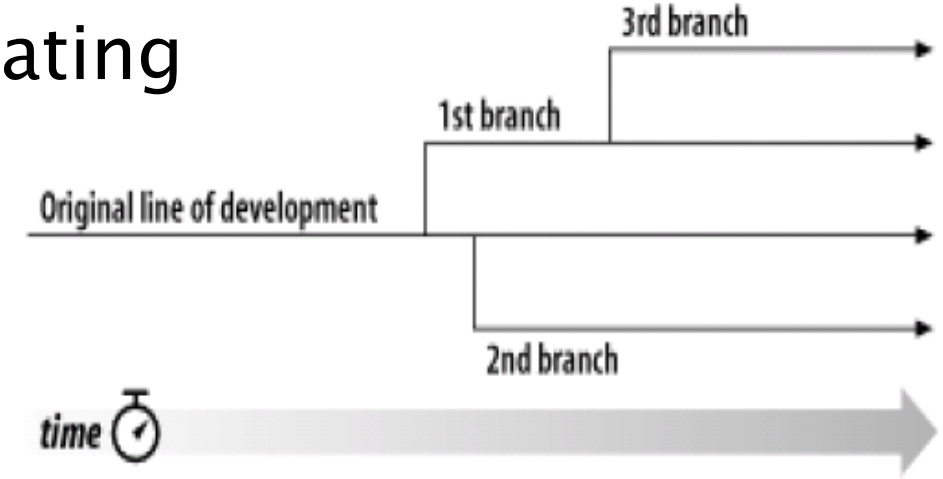

Branches: motivation

- Branches allow working in isolation form the main branch
  - Several new features or fixes can be developed independently and concurrently
  - When work is complete, they can be merged into the main branch
- Branches may represent different configurations, e.g. by platform

Branches in Subversion

- Unlike many other version control systems, Subversion's branches exist as normal filesystem directories in the repository, not in an extra dimension. These directories just happen to carry some extra historical information.
- Subversion has no internal concept of a branch—only copies. When you copy a directory, the resulting directory is only a “branch” because you attach that meaning to it. You may think of the directory differently, or treat it differently, but to Subversion it's just an ordinary directory that happens to have been created by copying.

Branches in Subversion

You create a branche with **svn copy**:

```
$ svn copy /trunk \
/calc/branches/my-calc-branch \
```

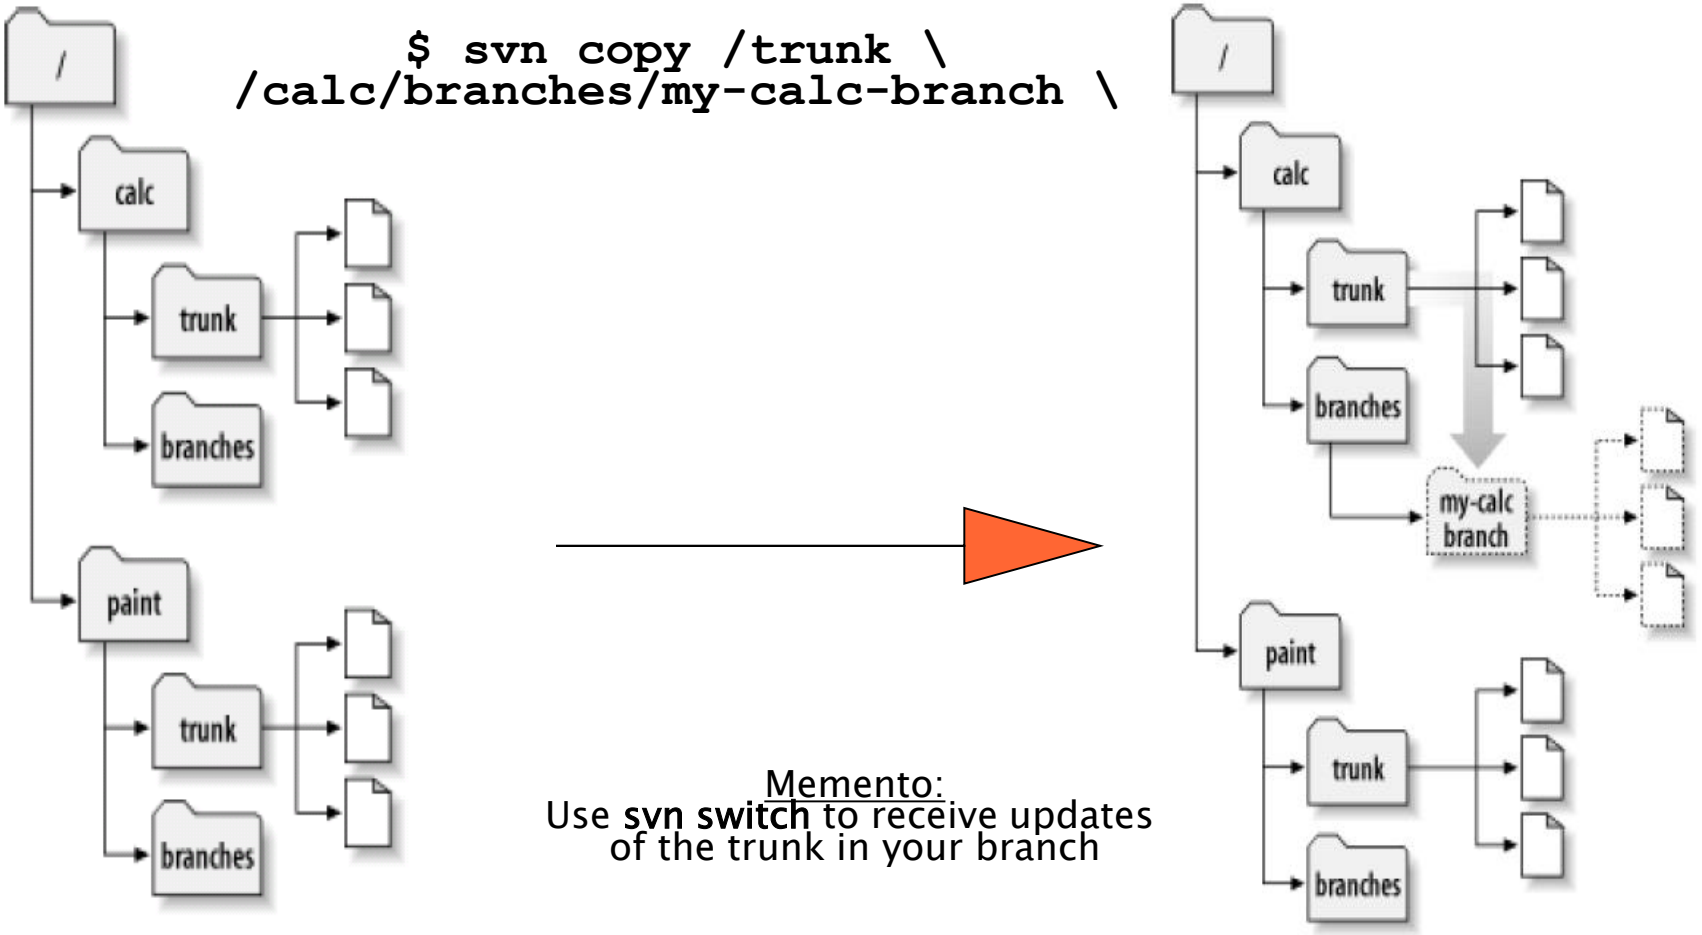

Memento:
Use **svn switch** to receive updates
of the trunk in your branch

Subversion repo structure

- To use branches a repository contains two top-level folders:
  - **`trunk`**: contains the main branch
  - **`branches`**: contain the branches
    - one sub-folder for each branch

Merge

- When work is done in a branch, it must be merged into the *trunk*. This is done by **`svn merge`** command.
  - Similar to **`svn diff`** command, instead of printing the differences to your terminal, however, it applies them directly to your working copy as local modifications. Svn diff command ignores ancestry, svn merge does not.
  - A better name for the command might have been svn diffand-apply, because that's all that happens: two repository trees are compared, and the differences are applied to a working copy.
- Conflicts may be produced by **`svn merge`**:
  - Conflicts are solved in the usual way

Semantic Versioning

- Product numbering based on

  MAJOR.MINOR.PATCH

- Increment
  - MAJOR: when you make (possibly incompatible) API changes,
  - MINOR: when you add functionality in a backwards-compatible manner, and
  - PATCH: when you make backwards-compatible bug fixes.

http://semver.org

References and Further Readings

- IEEE STD 1042 – 1987 IEEE guide to software configuration management
- IEEE STD 828-2012: IEEE Standard for Configuration Management in Systems and Software Engineering
- B.Collins-Sussman, B.W.Fitzpatrick C.M.Pilato. Version Control with Subversion: For Subversion 1.7, 2011
- Semantic Versioning. **http://semver.org**