The Software Heritage Filesystem (SwhFS): Integrating Source Code Archival with Development

Thibault Allançon
EPITA and Inria
Paris, France
thibault.allancon@epita.fr

Antoine Pietri
Inria
Paris, France
antoine.pietri@inria.fr

Stefano Zacchiroli
Université de Paris and Inria
Paris, France
zack@irif.fr

Abstract—**We introduce the Software Heritage filesystem (SwhFS), a user-space filesystem that integrates large-scale open source software archival with development workflows. SwhFS provides a POSIX filesystem view of Software Heritage, the largest public archive of software source code and version control system (VCS) development history.**

Using SwhFS, developers can quickly "checkout" any of the 2 billion commits archived by Software Heritage, even after they disappear from their previous known location and without incurring the performance cost of repository cloning. SwhFS works across unrelated repositories and different VCS technologies. Other source code artifacts archived by Software Heritage—individual source code files and trees, releases, and branches—can also be accessed using common programming tools and custom scripts, as if they were locally available.

A screencast of SwhFS is available online at dx.doi.org/10.5281/zenodo.4531411.

Index Terms—**source code, FUSE, filesystem, open source, version control system, digital libraries, digital preservation**

I. INTRODUCTION

Distributed version control (DVCS) is the state-of-the-art for source code management in software development. With DVCSs the full development history is replicated by each developer and code changes are exchanged among peers.

DVCS *per se* does not guarantee *code availability* though. Public repositories can disappear: they can be made private or moved; they can be targeted by (potentially bogus) copyright takedown notices; their hosting platform might shutdown. Specific commits can also disappear from repositories, e.g., due to their history being rewritten. The commits will still be available from other copies of the repository, but there is no easy way to find them.

Large-scale source code archival is a crucial part of a solution to this problem, especially for the vast corpus of free/open source software. Software Heritage [2], [7] is the largest archive of public code, having archived at the time of writing almost 10 billion (B) unique source code files and 2 B unique commits from more than 150 M projects spanning multiple hosting platforms and VCS technologies.

Source code artifacts in the Software Heritage archive (source code file or tree, commit, release, etc.) are stored in a global Merkle DAG and identified by content-based persistent identifiers called SWHIDs [6], e.g., swh:1:rev:9d76c0b163675505d1a901e5fe5249a2c55609bc. Given a SWHID, one can browse what's behind it via a Web UI.[1] SWHIDs for source code that is not locally available can be obtained in various ways: searching archived software by origin URL or metadata; finding SWHIDs mentioned in scientific papers [5], Wikidata, or software bills of materials [15]; deriving SWHIDs from other VCS references (e.g., "9d76c0b163675505d1a901e5fe5249a2c55609bc" in the previous example is a Git-compatible commit identifier).

Contributions and use cases: We introduce the *Software Heritage Filesystem (SwhFS)*,[2] a user-space filesystem that gives access to the Software Heritage archive as a POSIX filesystem. SwhFS integrates with development workflows by ensuring code availability when desired artifacts disappear from their last known location. Provided a reference to the desired code can be obtained, the corresponding file, directory, or commit can be used as if it were locally available.

Even for still-available code, SwhFS offers a convenient way to explore huge code bases [9], [10] that require significant time to be fully retrieved over the network, large amounts of disk space, and high I/O costs when switching branches.

II. RELATED WORK

Other VCS filesystems have been proposed in the past. RepoFS [12] exposes a Git repository as a FUSE [1], [16] filesystem, making different trade-offs than SwhFS. RepoFS needs a *local* Git repository, while SwhFS loads it lazily over the network. This makes SwhFS more suitable for quick exploration, and RepoFS more suitable for repository mining.

GitOD [13] and VFS for Git [10] expose remote Git repositories as local filesystems, loading them lazily in order to reduce disk and bandwidth usage for huge code bases, retaining compatibility with standard Git tools.

GitFS [11] and FigFS [8] also offer virtual filesystem interfaces on top of Git. FigFS is now abandoned. GitFS is not, but focuses on the use case of authoring new commits.

All related works reviewed thus far are Git-specific and have a single-repository scope. SwhFS is VCS-agnostic and spans the entire Software Heritage archive, giving access to hundreds of million code repositories in a unified view.

To the best of our knowledge SwhFS is the first attempt to integrate large-scale source code archival into development workflows via the filesystem interface.

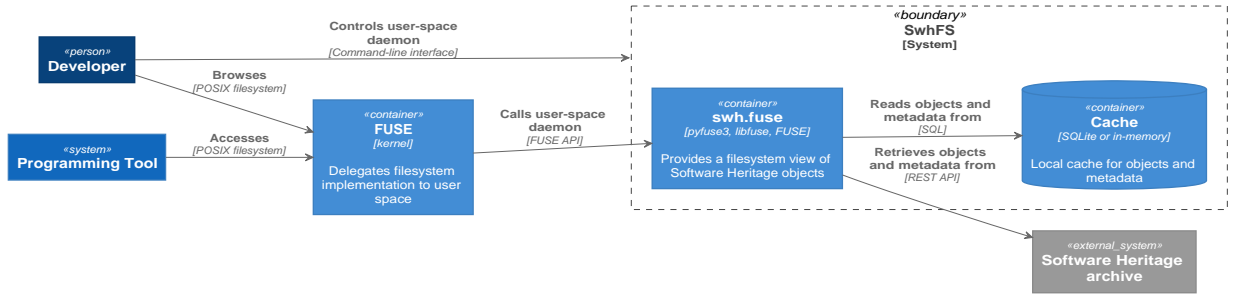
[1] https://archive.softwareheritage.org/

[2] https://docs.softwareheritage.org/devel/swh-fuse/

Fig. 1. Architecture of the Software Heritage virtual user-space filesystem (SwhFS)

III. DESIGN

Figure 1 shows the SwhFS architecture as a C4 diagram [4].

a) Front-end: POSIX filesystem and FUSE: Users control SwhFS by starting/stopping its user-space daemon via a convenient CLI interface. While running, SwhFS provides a filesystem view of all the source code artifacts archived by Software Heritage. This view is exposed as a POSIX filesystem that can be browsed using file-navigation tools and accessed by programming tools like text editors, IDEs, compilers, debuggers, custom shell scripts, etc.

According to the FUSE design [16], the POSIX filesystem interface seen by SwhFS clients is implemented by the OS kernel, which delegates decisions about the content of virtual files and directories to a user-space daemon, communicating with it via the FUSE API [1]. The SwhFS daemon is implemented in the swh.fuse Python module, using the pyfuse3 bindings to the C FUSE API.

b) Filesystem layout: The filesystem layout implemented by swh.fuse consists of (1) a set of entry points and (2) filesystem representations of Software Heritage objects. *Entry points* are located just below the SwhFS mount point; most notably archive/ allows browsing the Software Heritage archive by known SWHIDs while origin/ allows doing so by the URLs of archived projects (see Section IV for examples).

Filesystem representations of Software Heritage objects depend on the object type:

- *Source code files* (SWHID type: cnt) are represented as regular POSIX files;
- *Source code directories* (dir type) as directories whose content matches what was archived;
- *Commits and releases* (rev and rel) as directories containing a root sub-directory pointing to the source code tree plus auxiliary files containing metadata such as commit message and ancestry, timestamps, author information, etc.;
- *Repository snapshots* (snp) as directories containing one entry for each repository branch (master, v1.0, bug-6531, etc.), each pointing to the filesystem representation of the corresponding commit or release.

Relationships between accessed archived objects are represented using symbolic links. For instance, the source tree of a commit object archive/swh:1:rev:.../root is a symbolic link to a directory object located under archive/, e.g., "-> ../../swh:1:dir:...". This approach avoids duplications in the virtual filesystem, reifying on disk the sharing of Merkle structures.

The walkthrough of Section IV presents additional details about the SwhFS filesystem layout. A full layout specification is included with SwhFS documentation.

c) Backend: a Software Heritage ↔ FUSE adapter: swh.fuse is an adapter between the Software Heritage REST API[3] (SWH API) and the FUSE API. When filesystem entities that represent archived objects are accessed, swh.fuse uses the SWH API to retrieve data and metadata about them. Obtained results are then used to assemble the expected filesystem layout and return it to the kernel via the FUSE API. For example, when a source tree object is listed using readdir(), SwhFS invokes the /directory endpoint of the Software Heritage API and return to the kernel its directory entries using a lazy iterator; when a source code file is open()-ed, /content/raw is called and byte chunks of the retrieved blob are returned to the kernel piecemeal.

As part of adaptation, swh.fuse takes care and abstracts over details such as pagination, encoding issues, inode and file description allocation, etc.

d) Performance optimizations: Remote filesystems can be notoriously painful to use, due to network overhead. To make things worse for SwhFS, the backend technology stack (REST APIs and long-term archival storage) was not designed with filesystem-level performances in mind. In order to make it fast enough for practical use, SwhFS implements several optimizations.

Caching: Several *on-disk caches* are stored in SQLite databases to avoid repeating remote API calls. The *blob cache* stores raw source code file contents. The *metadata cache* stores the metadata of any kind of looked up object. Using these two caches it is possible to navigate any part of the Software Heritage archive that has been accessed in the past, even while disconnected from the network.

Merkle properties and the fact that SwhFS is read-only make cache invalidation unneeded, because no cached object can ever change. It is also always possible to purge objects from these caches without introducing inconsistencies.

---

[3]https://archive.softwareheritage.org/api/

In-memory caching is used for directories, as most filesystems do. The *direntry cache* maps directories to their entries, so that frequently accessed directories can be listed without even incurring SQLite query costs.

Compressed in-memory graph representation: When accessing commits and release objects, developers often need to quickly explore their commit history, *à la* `git log`. Exploring commit histories via the core SWH API would be too slow for that, as one HTTP call per commit is needed, times tens or hundreds of thousand commits for large repositories, with no parallelization opportunity, given that commit identifiers are discovered incrementally.

To address this issue we built upon `swh-graph`,[4] a compressed in-memory representation of the Software Heritage archive, obtained applying webgraph compression techniques [3]. The compressed structure of the global VCS graph comprising 17 B nodes and 200 B edges fits in $\approx$100 GiB of RAM and can be visited with amortized traversal cost close to a single memory access per traversed edge.

We have extended the SWH API to expose the `swh-graph` graph traversal API. For commit objects, SwhFS invokes it to perform a graph visit of all reachable commits and then populate several *history summary views* which are available under the `history/by-date/`, `history/by-hash/`, and `history/by-page/` directories (see Section IV for details). Even for huge repositories like the Linux kernel, retrieving the full commit history (as a list of SWHIDs) requires just a few tens of seconds, including network transfer.

Asynchronicity: In order to maximize I/O throughput SwhFS is implemented in asynchronous style: all blocking operations yield control to other coroutines instead of blocking. Additionally, SWH API calls are delegated to a shared thread pool that can keep persistent HTTP connections to the remote API backend, rather than establishing new ones at each call.

IV. WALKTHROUGH

This section shows how to use SwhFS via concrete examples. A screencast video is available online at dx.doi.org/10.5281/zenodo.4531411.

a) Installation: SwhFS is implemented in Python, released under the GPL3 license, and distributed via PyPI. It can be installed from there running `pip install swh.fuse`.

SwhFS development happens on the Software Heritage forge,[5] where issues and patches can be submitted.

b) Setup and teardown: Like with any filesystems, SwhFS must be "mounted" before use and "unmounted" afterwards. Users should first mount the Software Heritage archive as a whole and then browse archived objects looking up their SWHIDs below the `archive/` entry-point. To mount the Software Heritage archive, use the `swh fs mount` command:

```
$ mkdir swhfs
$ swh fs mount swhfs/  # mount the archive
$ ls -F swhfs/  # list entry points
archive/  cache/  origin/  README
```

[4]https://docs.softwareheritage.org/devel/swh-graph/
[5]https://forge.softwareheritage.org/source/swh-fuse/

By default SwhFS daemonizes in background and logs to syslog; it can be kept in foreground, logging to the console, by passing `-f/--foreground` to `mount`.

To unmount SwhFS use `swh fs umount PATH`. Note that, since SwhFS is a *user-space* filesystem, (un)mounting are not privileged operations, any user can do it.

The configuration file `~/.swh/config/global.yml` is read if present. Its main use case is inserting a per-user authentication token for the SWH API, which might be needed in case of heavy use to bypass the default API rate limit.

c) Source code browsing: Here is a SwhFS Hello World:

```
$ cd swhfs/
$ cat archive/swh:1:cnt:c839dea9e8e6f0528b4\
68214348fee8669b305b2
#include <stdio.h>

int main(void) {
    printf("Hello, World!\n");
}
```

Given the SWHID of a source code file, we can directly access it via the filesystem. We can do the same with entire source code directories. Here is the historical Apollo 11 source code, where we can find interesting comments about the antenna during landing:

```
$ cd archive/swh:1:dir:1fee702c7e6d14395bbf\
5ac3598e73bcbf97b030
$ ls | wc -l
127
$ grep -i antenna THE_LUNAR_LANDING.s | cut -f 5
# IS THE LR ANTENNA IN POSITION 1 YET
# BRANCH IF ANTENNA ALREADY IN POSITION 1
```

We can checkout the commit of a more modern codebase, like jQuery, and count its JavaScript lines of code (SLOC):

```
$ cd archive/swh:1:rev:9d76c0b163675505d1a9\
01e5fe5249a2c55609bc
$ ls -F
history/  meta.json@  parent@  parents/  root@
$ find root/src/ \
    -type f -name '*.js' | xargs cat | wc -l
10136
```

d) Commit history browsing: `meta.json` contains complete commit metadata, e.g.:

```
$ jq .author.name,.date,.message meta.json
"Michal Golebiowski-Owczarek"
"2020-03-02T23:02:42+01:00"
"Prevent collision with Object.prototype ..."
```

Commit history can be browsed commit-by-commit digging into `parent(s)/` directories or, more efficiently, using the history summaries located under `history/`:

```
$ ls -f history/by-page/000/ | wc -l
6469
$ ls -f history/by-page/000/ | head -n 2
swh:1:rev:358b769a00c3a09a...
swh:1:rev:4a7fc8544e2020c7...
```

The jQuery commit at hand is preceded by 6469 commits, which can be listed in "`git log`" order via the `by-page` view.

The `by-hash` and `by-date` views, inspired by RepoFS [12], list commits sharded by commit identifier and timestamp:

```
$ ls history/by-hash/00/ | head -n 1
swh:1:rev:0018f7700bf8004d...
$ ls -F history/by-date/
2006/  2007/  2008/  ...  2018/  2019/  2020/
$ ls -f history/by-date/2020/01/08/
swh:1:rev:437f389a24a6bef...
$ jq .date history/by-date/2020/01/08/*/meta.json
"2020-01-08T00:35:55+01:00"
```

Note that to populate the `by-date` view, metadata about all commits in the history are needed. To avoid blocking, metadata are retrieved asynchronously in the background, populating the view incrementally. The hidden `by-date/.status` file provides a progress report and is removed upon completion.

e) Repository snapshots and branches: Snapshot objects keep track of where each branch and release (or "tag") pointed to at archival time. Here is an example with the Unix history repository [14], which uses historical Unix releases as nested branch names:

```
$ cd archive/swh:1:snp:2ca5d6eff8f04a671c0d\
5b13646cede522c64b7d/refs/heads
$ ls -f | wc -l ; ls -f | grep Bell
40
Bell-32V-Snapshot-Development
Bell-Release
$ cd Bell-Release
$ jq .message,.date meta.json
"Bell 32V release ..."
"1979-05-02T23:26:55-05:00"
$ grep core root/usr/src/games/fortune.c
    printf("Memory fault -- core dumped\n");
```

Two of the 40 top-level branches correspond to Bell Labs releases. We can dig into the UNIX/32V `fortune` implementation instantly, without having to clone a 1.6 GiB repository.

f) Software origins: Software can also be explored by the URL it was archived from, using the `origin/` entry point:

```
$ cd origin/
$ cd https$ ls
2015-07-09/  2016-02-23/  2016-03-28/  ...
$ ls -F 2015-07-09/
meta.json  snapshot@
```

we can see a list of all archival crawls of the Linux kernel repository made by Software Heritage, and then navigate to the state of the repository as it was in 2015 (as a snapshot object). Note that one needs to use the exact origin URL and percent-encode it. To help with that, the companion `swh web search` CLI tool is available:

```
$ swh web search "torvalds linux" \
--limit 1 --url-encode | cut -f1
https
```

V. CONCLUSION

We introduced SwhFS, a user-space filesystem that allows browsing source code artifacts archived by Software Heritage as a POSIX filesystem. SwhFS integrates archival of public code with development workflows, allowing to quickly "checkout" any archived source code file, tree, commit, or repository without incurring the full repository cloning costs. SwhFS works across unrelated repositories and different version control technologies. SwhFS gives access to more than 9 B source code files and 2 B commits, archived by Software Heritage from more than 140 M projects, growing daily.

Future work: We plan to increase SwhFS throughput by making the SWH API queryable for multiple objects at once. Doing so might be enough to address mining software repository (MSR) use cases, which are currently out of scope for SwhFS. We are also studying the feasibility of integration with stock Git tools, making commands like `git log` and `blame` work within SwhFS mounts.

Exposing in the filesystem layout additional metadata available from Software Heritage—like licensing information and project metadata—is a low-hanging fruit that can provide added value on top of the already exposed VCS information.

REFERENCES

[1] Filesystem in userspace (FUSE). https://github.com/libfuse/libfuse. Accessed 2020-09-29.
[2] Jean-François Abramatic, Roberto Di Cosmo, and Stefano Zacchiroli. Building the universal archive of source code. *Communications of the ACM*, 61(10):29–31, September 2018.
[3] Paolo Boldi, Antoine Pietri, Sebastiano Vigna, and Stefano Zacchiroli. Ultra-large-scale repository analysis via graph compression. In *SANER 2020: The 27th IEEE International Conference on Software Analysis, Evolution and Reengineering*. IEEE, 2020.
[4] Simon Brown. The C4 model for software architecture. https://c4model.com/, 2018. Accessed 2020-11-16.
[5] Roberto Di Cosmo. Announcing biblatex-software: software citation made easy. *ACM SIGSOFT Softw. Eng. Notes*, 45(4):22–23, 2020.
[6] Roberto Di Cosmo, Morane Gruenpeter, and Stefano Zacchiroli. Identifiers for digital objects: the case of software source code preservation. In *iPRES 2018: the 15th Intl. Conference on Digital Preservation*, 2018.
[7] Roberto Di Cosmo and Stefano Zacchiroli. Software Heritage: Why and how to preserve software source code. In *iPRES 2017: the 14th International Conference on Digital Preservation*, 2017.
[8] Reilly Grant. Filesystem interface for the git version control system-final report. Technical report, University of Pennsylvania, 2009. https://www.stwing.upenn.edu/~rm/figfs/final.pdf, accessed 2020-11-20.
[9] Microsoft. Scalar. https://github.com/microsoft/scalar. Accessed 2020-09-29.
[10] Microsoft. VFS for Git. https://vfsforgit.org/. Accessed 2020-09-29.
[11] Presslabs. gitfs: Version controlled file system. https://github.com/presslabs/gitfs, 2014. Accessed 2020-11-17.
[12] Vitalis Salis and Diomidis Spinellis. RepoFS: File system view of Git repositories. *SoftwareX*, 9:288–292, 2019.
[13] Jonatan Schroeder. GitOD: An on demand distributed file system approach to version control. In *CTS 2012: International Conference on Collaboration Technologies and Systems*, pages 613–615. IEEE, 2012.
[14] Diomidis Spinellis. A repository of Unix history and evolution. *Empirical Software Engineering*, 22(3):1372–1404, 2017.
[15] Kate Stewart, Phil Odence, and Esteban Rockett. Software package data exchange (SPDX) specification. *IFOSS L. Rev.*, 2:191, 2010.
[16] Bharath Kumar Reddy Vangoor, Vasily Tarasov, and Erez Zadok. To FUSE or not to FUSE: performance of user-space file systems. In *15th USENIX Conference on File and Storage Technologies, FAST 2017*, pages 59–72. USENIX Association, 2017.