

Scuola di Scienze
Corso di Laurea in Informatica

TgFuseFs: How High School Students Can Write a Filesystem Prototype

Tesi di laurea

Relatore:

**Chiar.mo Prof.
Renzo Davoli**

Correlatori:

**Dott. Marco Sbaraglia
Dott. Michael Lodi**

Presentata da:

Riccardo Maffei



Information about licenses and credits can be found in Appendix A.

*To those I love and
those who love me...*

Abstract

Italian high school students who are majoring in Computer Science usually study subjects like programming, databases, networks, system engineering, electronics and operating systems. While most of these subjects let the students practice with projects, operating systems usually is taught in a more theoretical way because practical projects either are too simple to be didactically useful or require too many prerequisites. Hence, components like filesystems are only studied in theory from an high level point of view.

We think that building a filesystem prototype could be considered active learning and could improve the operating systems learning experience. For this reason in this work we will show how fifth year students with very few prerequisites can build their first working prototype of a remote filesystem in userspace using Python, FUSE and Telegram.

Since the activity is designed for high school students, the prototype won't be perfect but we will present some of the issues that students should be aware of and more advanced students should address.

Contents

Introduction	1
1 General Concepts	3
1.1 Requirements	3
1.2 Python	4
1.2.1 asyncio	4
1.3 FUSE and libfuse	8
1.3.1 pyfuse3	10
1.4 Telegram	11
1.4.1 Telethon	11
2 Implementation	13
2.1 Data Structures	13
2.1.1 Inode	13
2.1.2 Superblock	14
2.1.3 DirectoryData	15
2.1.4 FileData	16
2.2 Wrapper	16
2.2.1 Connection	16
2.2.2 Helpers	17
2.2.3 Serialization	18
2.2.4 Public API	19
2.3 mktgfs.py	22
2.4 Filesystem	24
2.4.1 Operations	24
2.4.2 Initialization and Mount	38
2.5 Testing	40
3 Issues and Improvements	43
3.1 Data Storage	43
3.2 Synchronization	43
3.3 Caching	44
3.4 Serialization	44
3.5 Additional Improvements	44
4 Conclusions and Future Works	45

Appendix A Licenses and Credits	47
A.1 Third-Party Content	47
Acknowledgments	49
References	51

List of Figures

1.1	A flow-chart diagram showing how the event loop manages tasks. . . .	5
1.2	A flow-chart diagram showing how FUSE works.	8
1.3	A Telegram Desktop native app screenshot shown on a laptop.	11
2.1	Example of testing session.	41
2.2	Telegram chat after tests.	41

List of Listings

1	Coroutines definition syntax.	4
2	Example of task creation to schedule a coroutine on the loop.	5
3	Example of two different coroutines.	6
4	Example of coroutine affected by TOC/TOU bug	7
5	Example of deadlock due to non-reentrancy of <code>asyncio</code> locks	7
6	Example of a <code>pyfuse3</code> -based filesystem implementation structure	10
7	Example of how to start the <code>Telethon</code> client.	12
8	Example of inode implementation.	13
9	Example of superblock implementation.	14
10	Example of directory implementation.	15
11	Example of file implementation.	16
12	Example of <code>TgFuseWrapper</code> initializer implementation.	16
13	Example of <code>_get_pinned_file</code> implementation.	17
14	Example of <code>_upload_data</code> implementation.	18
15	Example of <code>_download_data</code> implementation.	18
16	Example of <code>_pickle_and_save</code> implementation.	19
17	Example of <code>_unpickle</code> implementation.	19
18	Example of <code>write_superblock</code> implementation.	20
19	Example of <code>read_superblock</code> implementation.	20
20	Example of <code>write_data</code> implementation.	21
21	Example of <code>read_data</code> implementation.	21
22	Example of <code>delete_data</code> implementation.	21
23	Example of <code>mktgfs.py</code> implementation.	22
24	Example of <code>TgFuseFs</code> initializer implementation.	24
25	Example of <code>getattr</code> and <code>_getattr</code> implementation.	25
26	Example of <code>setattr</code> implementation.	26
27	Example of <code>lookup</code> and <code>_lookup</code> implementation.	27
28	Example of <code>opendir</code> implementation.	28
29	Example of <code>readdir</code> implementation.	29
30	Example of <code>releasedir</code> implementation.	29
31	Example of <code>mkdir</code> and <code>mknod</code> implementation.	30
32	Example of <code>_create</code> implementation.	31
33	Example of <code>_update_directory</code> implementation.	32
34	Example of <code>open</code> implementation.	33
35	Example of <code>read</code> implementation.	34
36	Example of <code>write</code> implementation.	34
37	Example of <code>release</code> implementation.	35
38	Example of <code>rmdir</code> and <code>unlink</code> implementation.	36

39	Example of <code>_remove</code> implementation.	36
40	Example of <code>forget</code> implementation.	37
41	Example of <code>close</code> implementation.	38
42	Example of mount script implementation.	39
43	Unlink handler test from the documentation.	40

Introduction

Italian high school students who are majoring in Computer Science usually study subjects like programming, databases, networks, system engineering, electronics and operating systems. While most of these subjects let students practice with projects, operating systems usually is taught in a more theoretical way because practical projects either are too simple to be didactically useful or require too many prerequisites. Hence, components like filesystems are only studied in theory from an high level point of view.

We think that building a filesystem prototype could be a practical activity that falls under the definition of active learning presented by [Bowell and Eison\[1\]](#) because it “involves students in doing things and thinking about the things they are doing”. According to a study by [Freeman et al.\[2\]](#), “active learning increases student performance across the STEM disciplines” and for this reason we think that building a filesystem prototype could improve the operating systems learning experience.

Even though there are other kind of schools teaching a subset of the subjects listed above, we think that students of Computer Science at “istituti tecnici”¹ usually have a more comprehensive computing education. In particular we think that the subject called “Tecnologie e progettazione di sistemi informatici e di telecomunicazioni”² could teach filesystems through practical projects during its 4 hours³ weekly schedule following the ministerial guidelines⁴.

In this work we will show how to implement a remote filesystem in userspace using Python, FUSE and Telegram.

Telegram lets user upload an unlimited number of files (up to 1.5GB each) as message attachments and supports a “self” chat where users can keep their files and messages. We will use that chat as the actual storage for filesystem data and metadata. The protocol used to interact with Telegram APIs is called `MTPProto`; however, we are not going to use it directly but we will use a library called `Teleton` which will make the interaction easier for students.

Writing filesystems usually requires kernel programming experience that high school students usually lack of. On the other hand writing a filesystem in userspace lets students use almost every language and library they want and also reduce the effort needed for debugging and avoids system crashes which may occur during the development of kernel modules. For this reason we are going to use `libfuse v3.x` through its Python bindings `pyfuse3`.

Since the activity is designed for high school students, the prototype won't be

¹Vocational technical high schools. Lit.: “technical institutes”

²Lit.: “Technologies and design of computer systems and telecommunications”

³Attachment C4 of the decree of the President of the Republic “D.P.R. 15 marzo 2010, n. 88”^[3]

⁴Attachment A.2 of the ministerial directive “Direttiva 16 gennaio 2012, n. 4”^[4]

perfect; in chapter 3 we will present some of the issues and improvements that students should be aware of.

Chapter 1

General Concepts

1.1 Requirements

This work assumes the students are enrolled in the fifth year of a Computer Science high school program. For this assumption, students should satisfy at least the following requisites:

- 3+ year of programming experience (in this work we use Python 3.7)
- very basic knowledge of concurrent programming concepts such as:
 - interleaving
 - locking and mutual exclusion
- strong experience and knowledge of UNIX/Linux environment
- OS knowledge including some theoretical knowledge of basic filesystem implementation concepts

This whole work assumes the development is done on Ubuntu (or any Debian derivative) with the following packages:

- Python 3.7
- fuse3
- libfuse3 v3.4.x or later
- libfuse3-dev development packages.

The following Python libraries are also required:

- pyfuse3
- Telethon

For obvious reasons, a Telegram account is needed too.

1.2 Python

In high school programming classes student are usually taught programming languages such as Python, C++ or Java. For this filesystem implementation we decided to use Python because there are many useful libraries that will ease the development. Even though we will not use multithreading, both `Telethon` and `pyfuse3` use asynchronous programming through `asyncio`. In the following section we will briefly introduce some of its basic concept, more can be found in the official documentation[5].

1.2.1 asyncio

`asyncio`[5] is an asynchronous programming module introduced in Python 3.4. It consists of several components but students will need to get familiar with:

- event loop
- coroutines
- `asyncio` locks

1.2.1.1 Event Loop

The event loop manages and distributes the execution of different tasks which are used to run coroutines in event loops. As shown in figure 1.1, when a task needs to wait something to finish, such as an I/O call, the task is suspended and another one is executed. This works in a similar way to cooperative schedulers. When the blocking call ends, the task will be pushed back into the loop queue and will be eventually executed resuming the coroutine from where it was interrupted.

More about coroutines and the loop in the following sections.

1.2.1.2 Coroutines

Coroutines are a more generalized form of subroutines which can be entered, exited, and resumed at many different points. The basic syntax to define a coroutine is to define a function with `async` before `def` as shown in listing 1

```
1 async def coroutine():  
2     # code here  
3     ...
```

Listing 1: Coroutines definition syntax.

Inside coroutines it's possible to `await` on other coroutines or any `awaitable`. Coroutines should be scheduled and cannot be called directly; calling `coroutine()` will just return a coroutine object. To be actually executed they must be `awaited`

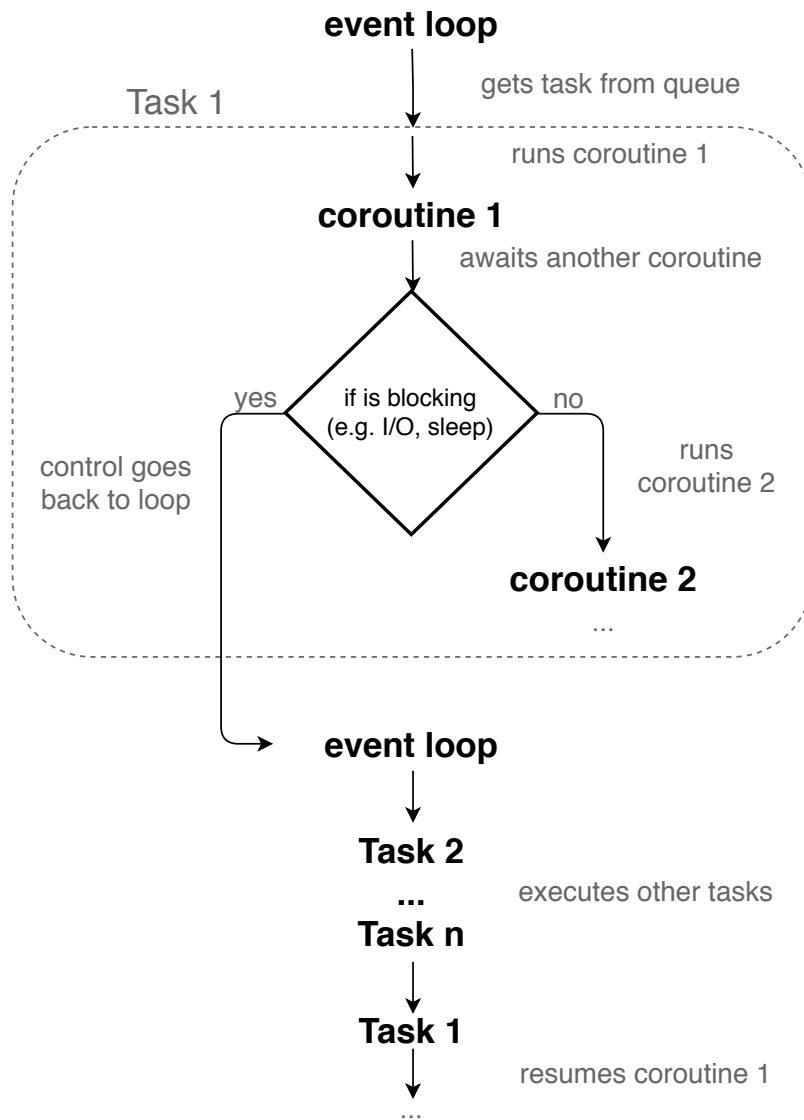


Figure 1.1: A flow-chart diagram showing how the event loop manages tasks.

from another running coroutine or must be scheduled on the loop for example creating a task (see listing 2) that will be executed concurrently to other tasks already scheduled on the loop.

```

1  async def coroutine():
2      # code here
3      ...
4  ...
5  task = asyncio.create_task(coroutine())
  
```

Listing 2: Example of task creation to schedule a coroutine on the loop.

1.2.1.3 asyncio Locks

Students should be already familiar with concurrent programming, in particular with race conditions and basic synchronization primitives. In this work we will disable multithreading for the sake of simplicity; however, due to the nature of coroutines there could still be interleaving and race conditions.

Let's take a look at the example shown in listing 3.

```
1 async def first():
2     print('First start')
3     await asyncio.sleep(10)
4     print('First end')
5
6 async def second():
7     print('Second start')
8     await asyncio.sleep(10)
9     print('Second end')
```

Listing 3: Example of two different coroutines.

Let's say that both `first` and `second` are scheduled on the loop at the same time. A beginner student may think that without concurrency there will not be interleaving but that's false. When **something blocking** like `sleep` is awaited a context switch happens. The currently running task (wrapping the coroutine) is suspended and the control goes back to the loop which may resume a different task as explained in section 1.2.1.1. One possible output may be something like this:

```
First start
Second start
First end
Second end
```

Note that awaiting something doesn't necessarily imply that a (task) context switch will happen. The context switch will happen if and when a blocking call is executed (such as I/O, sleep etc...).

Because of this behavior, access to shared data may lead to race conditions and bugs such as the TOC/TOU[6] shown in listing 4.

```

1  async def coroutine(self):
2      # check if some file is not open
3      if not self.file.is_open:
4          # the file is not open here
5          # await something
6          await asyncio.sleep(10)
7          # the file could be already open here!!
8      ...

```

Listing 4: Example of coroutine affected by TOC/TOU bug

To avoid this kind of bugs we could use some advanced synchronization techniques but they are too advanced for high school students and, therefore, they are out of the scope of this work. In this implementation we will use simpler yet inefficient techniques such as using mutual exclusion around critical sections using locks which students should be already familiar with. This will lead to performance issues as described in chapter 3.

Beware that `asyncio.Locks` are not reentrant, this means that the execution of `coroutine` shown in listing 5 will lead to deadlock.

```

1  async def coroutine(self):
2      # acquire lock
3      with self.lock:
4          # call something which requires the same lock
5          await self.same()
6
7  async def same(self):
8      # acquire lock
9      with self.lock: # <-- deadlock
10     ...

```

Listing 5: Example of deadlock due to non-reentrancy of `asyncio` locks

1.3 FUSE and libfuse

FUSE[7] (Filesystem in Userspace) is a software interface that lets non-privileged users to write their own filesystem writing only userspace code.

FUSE main components are:

- the FUSE kernel module (`fuse.ko`)
- the userspace library (`libfuse.*`)[8]
- a mount utility (`fusermount`)

To write a filesystem, either real or virtual, we should write a userspace program that will handle all the filesystem operations we want to support. To do so, the program will be linked to `libfuse` and implement the operations callbacks such as `getattr`, `open` etc... (more information in chapter 2 and section 2.4.1)

Let's say that we have implemented a FUSE filesystem in a program called `hello`, we mount it on `/tmp/fuse` and try to list that directory. As shown in figure 1.2 the following will happen:

1. `ls` will send the request to the kernel through system calls.
2. in kernelspace `VFS` will check which filesystem is mounted at that location and redirects the request to the right module. In this scenario: `FUSE`.
3. the FUSE kernel module will redirect the request to the `hello` program (which has been registered in the kernel module on launch).
4. the program (in userspace) will handle the request through the implemented callbacks and sends the answer all the way back.

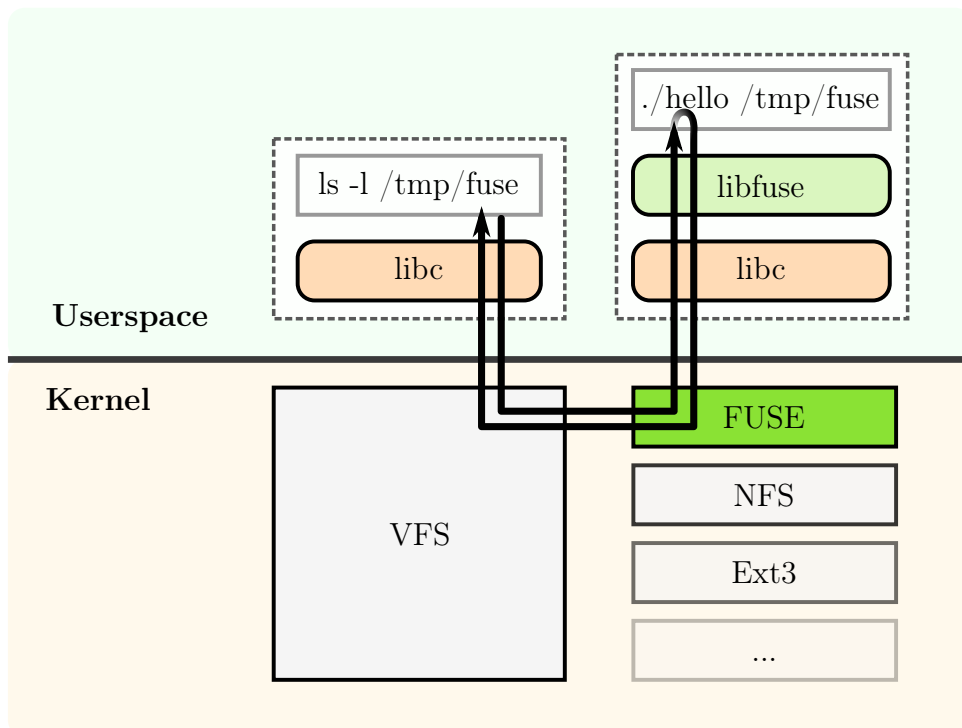


Figure 1.2: A flow-chart diagram showing how FUSE works.

Writing a FUSE-based filesystem has both advantages and disadvantages. The advantages of FUSE over standard kernel implementation are:

- Kernel development experience is not required.
- Both development and usage can be done by non-privileged users.
- Developing in userspace means no system crash will occur in case of something goes wrong.
- The filesystem can be developed in almost any preferred programming language thanks to the several bindings available.
- Any IDE, debugger and library can be used.
- Clean API to implement operations.
- Out-of-the-box user isolation.
- Avoid potential license issues like those happened to ZFS[9][10].

On the other hand there are disadvantages such as:

- Lower performance
- Higher requirements for the target system (such as `libfuse` installed)
- Not the best option when needing a multi user filesystem.

1.3.1 pyfuse3

pyfuse3[11][12] is a library that implements libfuse bindings for Python and it's compatible with both Trio and asyncio. Writing a filesystem consist in subclassing the pyfuse3.Operations class and implementing the request handler for the operation we want to support.

The FUSE kernel module will call those handlers as described in section 1.3 and in the official documentation[12]. In listing 6 is shown a simple empty structure (pay attention to the sixth line which enables the support for asyncio) while in chapter 2 we will present the prototype implementation details.

```
1 import pyfuse3
2 import pyfuse3_asyncio
3 import asyncio
4
5 # enable asyncio support
6 pyfuse3_asyncio.enable()
7
8 # implement subclass
9 class ExampleFS(pyfuse3.Operations):
10     # implement request handlers here such as open(), read() etc...
11
12     async def open(self, inode_n, flags, ctx):
13         ...
14
15     async def read(self, fh, off, size):
16         ...
17
18     async def write(self, fd, offset, buf):
19         ...
20
21     ...
22
23 # init the operations handler and FUSE options
24 operations = ExampleFS()
25 fuse_options = set(pyfuse3.default_options)
26 fuse_options.add('fname=examplefs')
27 # run pyfuse through asyncio
28 asyncio.get_event_loop().run_until_complete(pyfuse3.main())
29 # error handling is omitted
```

Listing 6: Example of a pyfuse3-based filesystem implementation structure

1.4 Telegram

Telegram[13] is a popular cloud-based instant messaging and VoIP platform with a focus on security and speed. It supports several platforms and provides mobile and desktop native clients.



Figure 1.3: A Telegram Desktop native app screenshot shown on a laptop.

Telegram, which is based on the `MTPROTO`[14] protocol, has a lot of functionalities such as bots, groups, channels, scheduled messages, reminders, stickers and secret chats but we will only use a little subset including the “self” (or “Saved Messages”) chat and the file upload functionality.

Telegram let’s its users upload an unlimited number of files with the maximum size of 1.5GB each.

The special chat called “Saved Messages” is a chat where people can save all kinds of messages supported by Telegram. For example users can send a message (text, file, document, media etc) directly to that chat or forward there any message from other chats.

1.4.1 Telethon

`Telethon`[15][16] is a Pure Python 3 `MTPROTO` library to interact with Telegram APIs based on `asyncio`.

To interact with Telegram APIs the first requirement is to get an API key (an ID-hash pair) from <https://my.telegram.org/> then the client must be started in one of the many documented way[16]. An example is shown in listing 7.

```

1 from telethon import TelegramClient
2 import asyncio
3
4 # ID-hash pair from my.telegram.org
5 api_id = 12345
6 api_hash = '0123456789abcdef0123456789abcdef'
7
8 # init the client
9 client = TelegramClient('SESSIONNAME', api_id, api_hash)
10 # start the client with the given phone number.
11 client.start('+390000000000')
12
13 # example coroutine
14 async def coroutine1():
15     # send a message to myself
16     await client.send_message('me', 'Hello world!')
17
18 # run the coroutine
19 asyncio.get_event_loop().run_until_complete(coroutine1())

```

Listing 7: Example of how to start the Telethon client.

The library provides a lot of functionalities, in chapter 2 and section 2.2 we will explain those we need to use.

Chapter 2

Implementation

In this chapter we will show an example of implementation for the filesystem prototype. In particular we will show the code along with comments and an explanation of the major choice made.

2.1 Data Structures

In this section we will present the implementation of the main data structure used in this filesystem prototype.

All the data structures are implemented in a file called `model.py`.

2.1.1 Inode

The inode contains all the file metadata along with some other utility information.

```
7 class Inode:
8     def __init__(self, number):
9         # inode attributes like 'stat'
10        self.attributes = EntryAttributes()
11        # pointer to data (message id) or data when inlining
12        self.data_pointer = 0
13        # init the number in attributes
14        self.attributes.st_ino = number
15
16    def is_directory(self):
17        # check mode and return
18        return stat.S_ISDIR(self.attributes.st_mode)
19
20    def is_regular_file(self):
21        # check mode and return
22        return stat.S_ISREG(self.attributes.st_mode)
```

Listing 8: Example of inode implementation.

The attribute `attributes` is of type `pyfuse3.EntityAttributes`[12] and contains all the metadata usually returned by a `stat(1)` call such as the inode number, the size, the mode, timestamps etc...

The attribute `data_pointer` is a pointer to the actual file data. In this prototype we used a only a single pointer to a Telegram message which contains the file data as attachment. More considerations about this choice in chapter 3.

2.1.1.1 `is_directory(self)` and `is_regular_file(self)`

This prototype supports only regular files and directories. These two utility methods are useful to check whether the inode belong to a file or a directory. They are implemented using the `stat` library and the mode attribute.

2.1.2 Superblock

The superblock contains metadata of the filesystem and all the inodes.

```
25 class Superblock:
26     def __init__(self, inode_n):
27         # init inode dict with the root inode and free set
28         self.inodes = {1: Inode(1)}
29         self.free_set = set(range(2, inode_n + 1))
30
31     def get_new_inode(self):
32         # if there are free inodes
33         if self.free_set:
34             # get a free inode number
35             free_n = self.free_set.pop()
36             # add a new inode with the given number to the dict
37             self.inodes[free_n] = Inode(free_n)
38             # return the new inode
39             return self.inodes[free_n]
40         # return None otherwise
41         return None
42
43     def free_inode(self, number):
44         # delete the inode from the dictionary
45         del self.inodes[number]
46         # add the inode number to the free set
47         self.free_set.add(number)
48
49     def get_inode_by_number(self, number):
50         # return the inode if exist, None otherwise
51         return self.inodes.get(number, None)
```

Listing 9: Example of superblock implementation.

For easier access the inodes are stored in a dictionary but in order to reduce the size of the superblock (that needs to be uploaded) unused inodes are not saved at all and only their numbers are stored in a set. More considerations about this choice in chapter 3. The inode number 1 is automatically added for the filesystem root directory.

2.1.2.1 `get_new_inode(self)`

This method is used to get a new empty inode. It removes one inode number from the free set and creates a new empty `Inode` with the same number and than adds it to the dictionary. Returns `None` if no inode is available.

2.1.2.2 `free_inode(self, number)`

This method frees the inode with the given number. The actual inode is removed from the dictionary and its number is added in the free set.

2.1.2.3 `get_inode_by_number(self, number)`

This method returns the inode with the given number if exists. `None` otherwise.

2.1.3 `DirectoryData`

UNIX well known philosophy says that “*everything is a file or a process*”[17]. Directories are just files whose content is a list of other files. The `DirectoryData` class wraps exactly that list.

```
54 class DirectoryData:
55     def __init__(self, self_inode_n, parent_inode_n):
56         # init directory entries with '.' and '..'
57         self.entries = {b'.': self_inode_n, b'..': parent_inode_n}
58
59     def __len__(self):
60         return len(self.entries)
```

Listing 10: Example of directory implementation.

The “list” is implemented as a dictionary where the key is the name of the entry in bytes (as requested by the documentation[12]) and the value is the number of the inode. The dictionary is always prefilled with `'.'` and `'..'` respectively the current and parent directory entry.

2.1.3.1 `__len__(self)`

This method is invoked when `len` is called with a `DirectoryData` object as a parameter to get the size of the directory. Since this is implemented as the size of the dictionary, an empty directory will always have a size of 2.

2.1.4 FileData

The FileData class wraps the actual data of a file.

```
63 class FileData:
64     def __init__(self, initial_data=b''):
65         self.raw_data = initial_data
66
67     def __len__(self):
68         return len(self.raw_data)
```

Listing 11: Example of file implementation.

In this prototype it could be replaced by using `bytes` directly; however, this class will become useful when other feature will be added to the filesystem. The data is saved as `bytes` and it's prefilled as empty.

2.1.4.1 __len__(self)

This method is invoked when `len` is called with a `FileData` object as a parameter to get the size of the file. This is implemented simply calling `len` on the raw bytes.

2.2 Wrapper

The class `TgFuseWrapper`, implemented in the file `wrapper.py`, contains all the utility methods needed to interact with the remote storage, in this case Telegram. This class should be the only piece of code interacting directly with Telegram and exposes a little API (its methods) used by the other modules of the filesystem. In this way changing the storage should be as easy as reimplementing this class without any other modification to the rest of the code.

The class has only 2 static constants, `api_id` and `api_hash`, which contains the Telegram API keys.

2.2.1 Connection

The class initializer shown in listing 12 takes care of the connection to Telegram using the Telethon library.

```
15 def __init__(self, number):
16     # create and save the client with default api keys
17     self.client = TelegramClient('anon', TgFuseWrapper.api_id,
18     ↪ TgFuseWrapper.api_hash)
19     # start the client
20     self.client.start(number)
```

Listing 12: Example of `TgFuseWrapper` initializer implementation.

The attribute `client` is a `TelegramClient` created and started connecting to the given number. As documented[16], `start` works even outside of coroutines.

2.2.2 Helpers

The following methods are helpers to accomplish common task such as download and upload of raw data.

2.2.2.1 `_get_pinned_file(self)`

This method, shown in listing 13, retrieves the file attached to the pinned message as bytes if exists. None otherwise.

```
21 async def _get_pinned_file(self):
22     # WORKAROUND: to get pinned messages in PRIVATE chats
23     # get full chat object which contains the pinned message id
24     full = await self.client(GetFullUserRequest('me'))
25     # get the message from its id
26     message = await self.client.get_messages('me',
27     ↪     ids=full.pinned_msg_id)
27     # get and return the attached file as bytes (None if missing)
28     return await self.client.download_media(message, file=bytes)
```

Listing 13: Example of `_get_pinned_file` implementation.

In order to get pinned messages for private chats we cannot use the high level API so the chat full information is retrieved using the low-level raw API. Once the message is retrieved, the attached media is downloaded in-memory and returned as bytes if any.

2.2.2.2 `_upload_data(self, data, caption, old_to_delete)`

This method, shown in listing 14, uploads the given data and returns the message.

```

30 async def _upload_data(self, data, caption=None, old_to_delete=None):
31     # upload given data
32     message = await self.client.send_file('me',
33                                         file=data,
34                                         caption=caption,
35                                         force_document=True)
36     # if old to delete is provided then try to delete that message
37     if old_to_delete is not None:
38         await self.delete_data(old_to_delete)
39     # return the new message
40     return message

```

Listing 14: Example of `_upload_data` implementation.

When the `old_to_delete` parameter is provided this upload is intended as a replacement and the message with the given id is deleted. This method accepts a `caption` which is useful for debug purposes to visualize in Telegram a caption below the uploaded message.

2.2.2.3 `_download_data(self, message_id)`

This method, shown in listing 15, downloads the data from the message with the given id and returns it as bytes if exists. Returns None otherwise.

```

45 async def _download_data(self, message_id):
46     # get message by id from 'me'
47     message = await self.client.get_messages('me', ids=message_id)
48     # get and return the attached file as bytes (None if missing)
49     return await self.client.download_media(message, file=bytes)

```

Listing 15: Example of `_download_data` implementation.

2.2.3 Serialization

Both data and metadata (such as the superblock shown in section 2.1.2) must be eventually saved somehow. While there are more secure and efficient way to do that (more considerations in chapter 3), for sake of simplicity in this prototype we serialized the whole object using the `pickle` library.

2.2.3.1 `_pickle_and_save(self, obj, caption, old_to_delete)`

This method, shown in listing 16, serializes the given object to bytes using `pickle` and then uploads the data with the other given parameters.

```

90 async def _pickle_and_save(self, obj, caption=None,
    ↪ old_to_delete=None):
91     # pickle the given object
92     pickled = pickle.dumps(obj)
93     # upload the pickled data and return the message
94     return await self._upload_data(pickled, caption, old_to_delete)

```

Listing 16: Example of `_pickle_and_save` implementation.

2.2.3.2 `_unpickle(self, file)`

This method, shown in listing 17, deserializes the object from the given file using `pickle` and then returns it. At the moment is just a wrapper around `pickle.loads` but it's useful in case the serialization technique changes.

```

96 def _unpickle(self, file):
97     # load (may be changed in case of new serialization method)
98     return pickle.loads(file)

```

Listing 17: Example of `_unpickle` implementation.

2.2.4 Public API

The following methods are those that are called from the other modules. In case the storage changes (from Telegram to something different) the class should be rewritten keeping the signatures of these methods intact. In this way only minor changes to the other modules may be required.

2.2.4.1 `write_superblock(self, superblock, should_replace)`

This method, shown in listing 18, writes the given superblock to Telegram.

```

51 async def write_superblock(self, superblock, should_replace=False):
52     # WORKAROUND: to get pinned messages in PRIVATE chats
53     # if the old superblock should be replaced get its message id,
54     #   ↪ None otherwise
55     old_id = (await
56     #   ↪ self.client(GetFullUserRequest('me')).pinned_msg_id if
57     #   ↪ should_replace else None
58     # pickle and save the superblock
59     message = await self._pickle_and_save(superblock, 'superblock',
60     #   ↪ old_id)
61     # pin the message containing the superblock data
62     await message.pin()
63     # return the message id
64     return message.id

```

Listing 18: Example of write_superblock implementation.

This is done by serializing and uploading the superblock and then pinning the message so that it can be retrieved easier later. In case the old one should be replaced, the old message id is retrieved and passed as old_to_remove.

2.2.4.2 read_superblock(self)

This method, shown in listing 19, reads the superblock from Telegram.

```

62 async def read_superblock(self):
63     # get the pinned file if exist, None otherwise
64     file = await self._get_pinned_file()
65     # the superblock, None as default
66     superblock = None
67     # if existed
68     if file is not None:
69         # unpickle the superblock
70         superblock = self._unpickle(file)
71     # return the superblock or None
72     return superblock

```

Listing 19: Example of read_superblock implementation.

This is done by getting the pinned file and then deserializing the superblock from that file.

2.2.4.3 write_data(self, data, caption, old_to_delete)

This method, shown in listing 20, writes the given data to Telegram and returns the message id.


```

74 async def write_data(self, data, caption=None, old_to_delete=None):
75     # pickle and save data then return the message id
76     return (await self._pickle_and_save(data, caption,
    ↪     old_to_delete)).id

```

Listing 20: Example of write_data implementation.

This is done by serializing the object (FileData or DirectoryData) and uploading it to Telegram. If necessary the old_to_delete is replaced and a caption is added.

2.2.4.4 read_data(self, message_id)

This method, shown in listing 21, reads and returns the data from Telegram from the message with the given id if exist. None otherwise.

```

78 async def read_data(self, message_id):
79     # get the file by id
80     file = await self._download_data(message_id)
81     # the directory data, None as default
82     data = None
83     # if existed
84     if file is not None:
85         # unpickle the data
86         data = self._unpickle(file)
87     # return the retrieved data or None
88     return data

```

Listing 21: Example of read_data implementation.

This is done by downloading the file attached to the message and then deserializing the object (FileData or DirectoryData).

2.2.4.5 delete_data(self, message_id)

This method, shown in listing 22, deletes the data from Telegram from the message with the given id.

```

42 async def delete_data(self, message_id):
43     await self.client.delete_messages('me', message_id)

```

Listing 22: Example of delete_data implementation.

The implementation is straightforward, just call Telethon to remove the message with the given id.

2.3 mktgfs.py

This prototype, implemented in the file `tgfuse.py`, mounts an **existing** filesystem from Telegram to a local mountpoint. An empty filesystem should be somehow created before mounting it.

```
1 import asyncio
2 import os
3 import sys
4 from time import time_ns
5
6 from wrapper import *
7
8
9 async def make():
10     # Create an empty superblock
11     s = Superblock(1000)
12     # get the inode reserved for the root directory
13     ino = s.get_inode_by_number(1)
14     # init its attributes
15     entry = ino.attributes
16     entry.st_mode = (stat.S_IFDIR | 0o755)
17     stamp = time_ns()
18     entry.st_atime_ns = stamp
19     entry.st_ctime_ns = stamp
20     entry.st_mtime_ns = stamp
21     entry.st_gid = os.getgid()
22     entry.st_uid = os.getuid()
23     # create and init an empty directory
24     dd = DirectoryData(1, 1)
25     entry.st_size = len(dd)
26     # write the directory
27     message_id = await w.write_data(dd, 'ROOT directory data')
28     # save the pointer in the inode
29     ino.data_pointer = message_id
30     # write the superblock
31     await w.write_superblock(s)
32
33
34 # create a wrapper with the argument phone number
35 w = TgFuseWrapper(sys.argv[1])
36 # run the coroutine on the loop
37 asyncio.get_event_loop().run_until_complete(make())
```

Listing 23: Example of `mktgfs.py` implementation.

The file `mktgfs.py` shown in listing 23 is used to create an empty filesystem with

1000 maximum inodes and only an empty root directory. This script is extremely useful even during development because lets students easily clean the remote filesystem which could have been corrupted by bugs or other reasons.

A `TgFuseWrapper` is created with the phone number provided as parameter and then the `make` coroutine is run. The coroutine creates a new superblock and initializes the attributes for the root directory inode. The empty directory is then uploaded and its id is saved in the inode `data_pointer`. Finally the superblock is written to Telegram.

Note that the parent inode is set as itself, this will be overwritten with the actual parent of the mountpoint by the kernel.

2.4 Filesystem

The actual filesystem implementation is part of the file `tgfuse.py`. In this section we will describe our implementation example of some of the operation handlers.

2.4.1 Operations

The class `TgFuseFs`, extension of `pyfuse3.Operations`, implements the core of the filesystem operations.

Most of the requirements used in the following sections are described in the [pyfuse3 documentation\[12\]](#).

Note: we assume `asyncio` support is enabled through `pyfuse3_asyncio.enable()`.

2.4.1.1 Attributes and Initializer

```
22 def __init__(self, number: str, *args, **kwargs):
23     super().__init__(*args, **kwargs)
24     # create the wrapper
25     self.wrapper = TgFuseWrapper(number)
26     # synchronously get the superblock
27     self.superblock = asyncio.get_event_loop().run_until_complete(
28         ↪ self.wrapper.read_superblock())
29     assert self.superblock is not None
30     # create a mutex lock for the superblock
31     self.sb_lock = asyncio.Lock()
32     # init a counter
33     self.counter = itertools.count()
34     # create open files and dirs dict
35     self.open_dirs = {}
36     self.open_files = {}
37     # create a lookup counter dict defaulting to 0
38     self.lookup_counters = defaultdict(int)
39     # create a list deferred removal
40     self.deferred = []
```

Listing 24: Example of `TgFuseFs` initializer implementation.

The initializer shown in listing 24 creates a `TgFuseWrapper` and synchronously retrieves the superblock then initializes some utility attributes.

The attribute `sb_lock` is an `asyncio.Lock` used as mutex for operations over the superblock.

The attribute `counter` is just an integer counter used to avoid collisions of directories file handles.

The attribute `open_dirs` is a dictionary that holds data for open directories. The keys are the integer file handles returned by an `opendir` call while the values are tuples `(inode, dd)` where `inode` is a reference to the directory inode and

`dd` is a reference to the directory data (cloned at `opendir` time as explained in section 2.4.1.5).

The attribute `open_files` is a dictionary that holds data for open files. The keys are the file handles returned by an `open` call (inode numbers) while the values are tuples (`inode`, `fd`, `dirty`, `count`) where `inode` is a reference to the file inode, `fd` is a reference to the file data (used as cache until the file is closed by all clients), `dirty` is a flag that is true if the cache has been written to and `count` is an integer which counts how many time the file has been opened (and not closed yet).

The attribute `lookup_counters` is a `defaultdict` which holds lookup counters for inodes as requested by the documentation. The keys are the inode numbers while the values are integers defaulting to zero.

The attribute `deferred` is a list of inodes whose deletion has been deferred by the `unlink` handler. See section 2.4.1.8.

2.4.1.2 `getattr(self, inode_n, ctx)`

The `getattr` request handler should return the attributes of the inode with the given number if exist.

More detailed requirements can be found on the `pyfuse3` request handlers documentation.

```
41 async def getattr(self, inode_n, ctx):
42     # acquire superblock mutex and call internal method
43     async with self.sb_lock:
44         return await self._getattr(inode_n, ctx)
45
46 async def _getattr(self, inode_n, ctx=None):
47     # try to get the requested inode and if succeeded
48     inode = self.superblock.get_inode_by_number(inode_n)
49     if inode is not None:
50         # return the requested attributes
51         return inode.attributes
52     # otherwise raise ENOENT
53     raise FUSEError(errno.ENOENT)
```

Listing 25: Example of `getattr` and `_getattr` implementation.

Listing 25 shows our example of `getattr` and `_getattr` implementation.

The internal `_getattr` method assumes the superblock lock has already been acquired. Tries to get the inode with the requested number and if succeeded returns its attributes. Otherwise, the inode doesn't exist and `ENOENT` is raised. This method is useful and can be called internally by other handlers.

The method `getattr` is the actual handler. It simply acquires the superblock lock and then calls the internal `_getattr` method returning its result.

2.4.1.3 `setattr(self, inode_n, new_attr, fields, fh, ctx)`

The `setattr` request handler should change the attributes of the inode with the given number to the new provided values and then return all the attributes.

More detailed requirements can be found on the `pyfuse3` request handlers documentation.

```
55 async def setattr(self, inode_n, new_attr, fields, fh, ctx):
56     # acquire superblock mutex
57     async with self.sb_lock:
58         # try to get the requested inode
59         inode = self.superblock.get_inode_by_number(inode_n)
60         # get its attributes
61         attr = inode.attributes
62         # update the required fields
63         attr.st_atime_ns = new_attr.st_atime_ns if
64         ↪ fields.update_atime else attr.st_atime_ns
65         attr.st_mtime_ns = new_attr.st_mtime_ns if
66         ↪ fields.update_mtime else attr.st_mtime_ns
67         attr.st_ctime_ns = new_attr.st_ctime_ns if
68         ↪ fields.update_ctime else time_ns() # now if not
69         ↪ specified
70         attr.st_mode = new_attr.st_mode if fields.update_mode else
71         ↪ attr.st_mode
72         attr.st_uid = new_attr.st_uid if fields.update_uid else
73         ↪ attr.st_uid
74         attr.st_gid = new_attr.st_gid if fields.update_gid else
75         ↪ attr.st_gid
76         attr.st_size = new_attr.st_size if fields.update_size else
77         ↪ attr.st_size
78         # return the attributes
79         return attr
```

Listing 26: Example of `setattr` implementation.

Listing 26 shows our example of implementation.

After acquiring the superblock lock, the inode and its attributes are retrieved and then only the required values are updated (`field`, instance of `SetattrFields`, specifies which fields should be updated). The only exception is `st_ctime_ns`: when no new value is provided it's set to the current time (to indicate that the inode metadata was changed). Finally all the inode attributes are returned (both the changed and unchanged values).

2.4.1.4 `lookup(self, parent_inode_n, name, ctx)`

The `lookup` request handler should look up a directory entry by name and get its attributes.

More detailed requirements can be found on the `pyfuse3` request handlers documentation.

```
73 async def lookup(self, parent_inode_n, name, ctx):
74     # acquire superblock mutex
75     async with self.sb_lock:
76         # try to lookup (may raise exceptions)
77         result = await self._lookup(parent_inode_n, name, ctx)
78         # increase the lookup count (if no exceptions occurred)
79         self.lookup_counters[result.st_ino] += 1
80         return result
81
82 async def _lookup(self, parent_inode_n, name, ctx=None):
83     # try to get the requested inode
84     parent_inode =
85     ↪ self.superblock.get_inode_by_number(parent_inode_n)
86     # get directory data
87     dd = await self.wrapper.read_data(parent_inode.data_pointer)
88     # try to get the entry inode number and if succeeded
89     entry_inode_n = dd.entries.get(name, None)
90     if entry_inode_n is not None:
91         # FOUND: get attr and return
92         return await self._getattr(entry_inode_n, ctx)
93     # entry not found: raise ENOENT
94     raise FUSEError(errno.ENOENT)
```

Listing 27: Example of `lookup` and `_lookup` implementation.

Listing 27 shows our example of `lookup` and `_lookup` implementation.

The internal `_lookup` method assumes the superblock lock has already been acquired. The parent directory's `DirectoryData` is downloaded from Telegram and the entry with the given name is retrieved if exists. If the entry existed its attributes are returned calling `_getattr`; otherwise, `ENOENT` is raised. This method is useful and can be called internally by other handlers.

The method `lookup` is the actual handler. It acquires the superblock lock and calls the internal `_lookup` method. If the call succeeded the lookup counter for the entry is increased as required by the documentation and the result is returned. Otherwise, the exception (in this case `FUSEError(errno.ENOENT)`) is propagated.

2.4.1.5 `opendir(self, inode_n, ctx)`, `readdir(self, fh, start_id, token)` and `releasedir(self, fh)`

The handlers `opendir`, `readdir` and `releasedir` are used to access directories. In particular `opendir` should open the directory with the given inode number and return a file handle, `readdir` should list some entries from the open directory with

the given file handle and `releasedir` should release the open directory with the given file handle.

Instead of returning the directory entries directly, the `readdir` handler should call `readdir_reply` for each directory entry and stop when there are no more entries or when the call returns `False`. Multiple `readdir` calls may be used to read a directory, to keep track of the position an identifier `next_id` is passed to the reply and is passed back to the next `readdir` call (`start_id`). If entries are added or removed between calls, they may or may not be returned. However, they must not cause other entries to be skipped or returned more than once.

More detailed requirements can be found on the `pyfuse3` request handlers documentation.

In order to comply with the requirement to be resilient against concurrent edits of the directory, a simple caching approach is used in this prototype example: when the directory is opened with `opendir` the current status of the directory is cached and reused for each subsequent `readdir` call until the directory is closed with `releasedir`. The `next_id` passed to `readdir_reply` and passed back to `readdir` is just the index of the last read entry incremented by one (treating the cached `DirectoryData` entries as a list).

```
95 async def opendir(self, inode_n, ctx):
96     # acquire superblock mutex
97     async with self.sb_lock:
98         # try to get the requested inode
99         inode = self.superblock.get_inode_by_number(inode_n)
100        # if completely failed
101        if inode is None:
102            # raise ENOENT
103            raise FUSEError(errno.ENOENT)
104        # else if succeeded but is not a directory
105        elif not stat.S_ISDIR(inode.attributes.st_mode):
106            # raise ENOTDIR
107            raise FUSEError(errno.ENOTDIR)
108        # else, it's a directory, good!
109        # get directory data
110        dd = await self.wrapper.read_data(inode.data_pointer)
111        # get next counter value as fh
112        fh = next(self.counter)
113        # save current DD status in open dirs (cache) and return the
114        ↪ key
115        self.open_dirs[fh] = (inode, dd)
116        return fh
```

Listing 28: Example of `opendir` implementation.

Listing 28 shows our example of `opendir` implementation.

After acquiring the superblock lock, if the inode with the given number doesn't

exist `ENOENT` is raised. If it exists but isn't a directory `ENOTDIR` is raised instead. If everything is right the `DirectoryData` is downloaded from Telegram and it's saved along with a reference to the inode as a tuple `(inode, dd)` in the `open_dirs` dictionary using the next value of the internal counter `counter` as key. That key is used as file handle and is returned.

```
117 async def readdir(self, fh, start_id, token):
118     # acquire superblock mutex
119     async with self.sb_lock:
120         # get directory data from cache and its inode
121         (dir_inode, dd) = self.open_dirs[fh]
122         # init an index as start id
123         index = start_id
124         # iterate over the entries from the given offset
125         for (name, inode_n) in list(dd.entries.items())[start_id:]:
126             # get attributes for the current inode
127             attr = await self._getattr(inode_n)
128             # reply and if necessary stop the iteration
129             if not readdir_reply(token, name, attr, index + 1):
130                 break
131             # increase lookup counter only if not '.' nor '..'
132             if name != b'.' and name != b'..':
133                 self.lookup_counters[inode_n] += 1
134             index += 1
135             # update atime
136             dir_inode.attributes.st_atime_ns = time_ns()
```

Listing 29: Example of `readdir` implementation.

Listing 29 shows our example of `readdir` implementation.

After acquiring the superblock lock, the cache is accessed and an `index` is initialized as the given `start_id` (0 if the first call, the next entry position otherwise). The entries are treated as a list and are iterated over starting from `start_id` position (note that the cached data never changes). For each entry the attributes are retrieved and `readdir_reply` is called with the entry's name, attributes and the index incremented by one. If the call returned `False` the iteration is stopped; otherwise, the lookup counter for the entry is incremented unless it's `'.'` or `'..'` as requested by the documentation. The index is increased in each iteration. Before returning, the handler updates the directory access timestamp.

```
138 async def releasedir(self, fh):
139     # remove the cached DD status
140     del self.open_dirs[fh]
```

Listing 30: Example of `releasedir` implementation.

Listing 30 shows our example of `releasedir` implementation.

The cached data is simply removed from `open_dirs`.

2.4.1.6 `mkdir(self, parent_inode_n, name, mode, ctx)` and `mknod(self, parent_inode_n, name, mode, rdev, ctx)`

The handlers `mkdir` and `mknod` are used to create new directories and files (the prototype presented in this work supports only regular files).

The new file or directory must be created with the given `name` and `mode` inside the directory with `parent_inode_n` inode number. These methods should return the attributes of the newly created file or directory.

More detailed requirements can be found on the `pyfuse3` request handlers documentation.

```
166 async def mkdir(self, parent_inode_n, name, mode, ctx):
167     # acquire superblock mutex and call internal method
168     async with self.sb_lock:
169         return await self._create(parent_inode_n, name, mode, ctx)
170
171 async def mknod(self, parent_inode_n, name, mode, rdev, ctx):
172     # acquire superblock mutex and call internal method
173     async with self.sb_lock:
174         # special file are not supported: rdev is ignored
175         return await self._create(parent_inode_n, name, mode, ctx)
```

Listing 31: Example of `mkdir` and `mknod` implementation.

Listing 31 shows our example of `mkdir` and `mknod` implementation.

Since the creation of directories and files is very similar, these two methods just acquire the superblock lock and then call an internal method `_create`. That method takes care of the actual creation of the new file or directory.

```

177 async def _create(self, parent_inode_n, name, mode, ctx):
178     # get a new inode for the new file or directory
179     new_ino = self.superblock.get_new_inode()
180     # if none (no more inodes) raise ENOSPC
181     if new_ino is None:
182         raise FUSEError(errno.ENOSPC)
183     # init metadata
184     new_ino.attributes.st_mode = mode
185     new_ino.attributes.st_uid = ctx.uid
186     new_ino.attributes.st_gid = ctx.gid
187     now = time_ns()
188     new_ino.attributes.st_atime_ns = now
189     new_ino.attributes.st_mtime_ns = now
190     new_ino.attributes.st_ctime_ns = now
191     # if we are creating a directory
192     if new_ino.is_directory():
193         # create DirectoryData with only add '.' and '..'
194         data = DirectoryData(new_ino.attributes.st_ino,
195                               ↪ parent_inode_n)
196     # if we are creating a regular file
197     elif new_ino.is_regular_file():
198         # create empty FileData
199         data = FileData()
200     else:
201         # we don't implement this case:
202         # free the inode and raise ENOSYS
203         self.superblock.free_inode(new_ino.attributes.st_ino)
204         raise FUSEError(errno.ENOSYS)
205     # write the file data and save the id as pointer
206     new_ino.data_pointer = await self.wrapper.write_data(data)
207     # save the size
208     new_ino.attributes.st_size = len(data)
209     # update parent directory
210     await self._update_directory(parent_inode_n, [('+', name,
211                                               ↪ new_ino.attributes.st_ino)])
212     # increase the lookup counter and return the attributes
213     self.lookup_counters[new_ino.attributes.st_ino] += 1
214     return new_ino.attributes

```

Listing 32: Example of `_create` implementation.

Listing 32 shows our example of `_create` implementation.

The internal `_create` method assumes the superblock lock has already been acquired. If there is no available inode then `ENOSPC` is raised. Otherwise, a new inode is retrieved from the superblock and its metadata is initialized. If the method is handling the creation of a directory, then the data is set as an empty `DirectoryData` (with only `'.'` and `'..'`), if the method is handling the creation of a regular file,

then the data is set as an empty `FileData`, otherwise, the method is handling the creation of an unsupported type then the inode is freed and `ENOSYS` is raised. During development it's useful to add a custom caption here to help visualize in the chat what's happening. The data is uploaded to Telegram, the id is saved in the inode `data_pointer` and the size is updated. The new entry is added to the parent directory calling the internal `_update_directory` method then the lookup counter for the new file or directory is increased as required by the documentation just before returning the attributes.

```

142 async def _update_directory(self, dir_ino_n, entries):
143     # get directory data
144     dir_ino = self.superblock.get_inode_by_number(dir_ino_n)
145     dir_dd = await self.wrapper.read_data(dir_ino.data_pointer)
146     # for each entry
147     for (action, name, inode_n) in entries:
148         # if action is an add ('+')
149         if action == '+':
150             # add the new entry
151             dir_dd.entries[name] = inode_n
152         # otherwise is a delete
153         else:
154             assert dir_dd.entries[name] == inode_n
155             # remove the entry
156             del dir_dd.entries[name]
157     # update parent directory size
158     dir_ino.attributes.st_size = len(dir_dd)
159     # update ctime and mtime
160     now = time_ns()
161     dir_ino.attributes.st_ctime_ns = now
162     dir_ino.attributes.st_mtime_ns = now
163     # write and save parent directory data
164     dir_ino.data_pointer = await self.wrapper.write_data(dir_dd,
        ↪ old_to_delete=dir_ino.data_pointer)

```

Listing 33: Example of `_update_directory` implementation.

Listing 33 shows our example of `_update_directory` implementation.

The internal `_update_directory` method assumes the superblock lock has already been acquired. This method receives the inode number of the parent directory and a list of tuples (`action`, `name`, `inode_n`) where `action` is the action to perform ('+' if the entry should be added or something else if it should be removed), `name` is the entry name and `inode_n` is the entry inode number. The old `DirectoryData` is downloaded from Telegram and the given entries are added or removed. The inode metadata is updated (size, timestamps...) and the new data is uploaded back to Telegram replacing the old one.

```

2.4.1.7 open(self, file_inode_n, flags, ctx),
        read(self, fh, off, size),
        write(self, fh, off, buf) and
        release(self, fh)

```

The handlers `open`, `read`, `write` and `release` are used to access files.

In particular `open` should open the file with the given inode number and flags and return a `FileInfo` instance with the file handle, `read` should read the given amount of bytes from the open file with the given file handle starting from the given offset, `write` should write the given buffer to the open file with the given file handle at the given offset and return the amount of bytes written and `release` should release the open file with the given file handle.

More detailed requirements can be found on the `pyfuse3` request handlers documentation.

In this prototype example the `open` flags are ignored and the built-in cache is not enabled. In order to speed up the file access, a simple caching approach is used: when a file is opened with `open` the file is downloaded from Telegram and reused for any subsequent `read` or `write` call. If the file is already opened, `open` will access the in-memory copy of the file instead of downloading it from Telegram. When the file is released for the last time (one for each `open` call) the file is actually uploaded to Telegram if necessary.

```

214 async def open(self, file_inode_n, flags, ctx):
215     # acquire superblock mutex
216     async with self.sb_lock:
217         # try to get the file and open count from local open cache
218         (inode, fd, dirty, count) = self.open_files.get(file_inode_n,
219     ↪ (None, None, False, 0))
219     # if failed then populate it
220     if fd is None:
221         # get the inode
222         inode = self.superblock.get_inode_by_number(file_inode_n)
223         # get file data from Telegram
224         fd = await self.wrapper.read_data(inode.data_pointer)
225         # dirty already false and count already 0
226     # update (or create) entry in cache
227     self.open_files[file_inode_n] = (inode, fd, dirty, count + 1)
228     # return the inode number as fh
229     return FileInfo(fh=file_inode_n)

```

Listing 34: Example of `open` implementation.

Listing 34 shows our example of `open` implementation.

After acquiring the superblock lock, the method tries to access the cache to get the tuple `(inode, fd, dirty, count)` where `inode` is a reference to the file inode, `fd` is the cached `FileData`, `dirty` is a flag that is true if the cache has been written

to and `count` is a counter of how many times the file has been opened (and not closed yet). If the cache was empty the `FileData` is downloaded from Telegram and the inode is retrieved (Note that the default value for `dirty` is `False` and for `count` is 0). The cache is then populated (or updated) with the values and the counter incremented by one. A new `FileInfo` with the inode number as file handle is returned. Note that access to the cache must be done in mutual exclusion and this is automatically achieved because is done after acquiring the superblock lock.

```

231 async def read(self, fh, off, size):
232     # acquire superblock mutex
233     async with self.sb_lock:
234         # get inode and fd from the cache
235         (inode, fd, _, _) = self.open_files[fh]
236         # update atime and return the requested data
237         inode.attributes.st_atime_ns = time_ns()
238         return fd.raw_data[off:off + size]

```

Listing 35: Example of read implementation.

Listing 35 shows our example of read implementation.

After acquiring the superblock lock, the inode and `FileData` are retrieved from the cache. The file access timestamp is updated and the requested data is returned.

```

240 async def write(self, fh, off, buf):
241     # acquire superblock mutex
242     async with self.sb_lock:
243         # get inode and fd from the cache
244         (inode, fd, _, count) = self.open_files[fh]
245         # if the resulting file would be too big raise EFBIG
246         if max(off + len(buf), len(fd)) > 1.5E9:
247             raise FUSEError(errno.EFBIG)
248         # write to cache, set as dirty and update size
249         fd.raw_data = fd.raw_data[:off] + buf + fd.raw_data[off +
250             ↪ len(buf):]
251         self.open_files[fh] = (inode, fd, True, count)
252         inode.attributes.st_size = len(fd)
253         # update timestamps and return the amount of data written
254         now = time_ns()
255         inode.attributes.st_ctime_ns = now
256         inode.attributes.st_mtime_ns = now
257         return len(buf)

```

Listing 36: Example of write implementation.

Listing 36 shows our example of write implementation.

After acquiring the superblock lock, the inode, `FileData` and counter are retrieved from the cache. If the file would become too big (current implementation

uses 1 telegram attachment) `EFBIG` is raised. Otherwise, the given buffer is written to the `FileData` and the `dirty` flag is updated as `True`. The `inode` size and timestamps are updated and the length of the buffer is returned.

```
258 async def release(self, fh):
259     # acquire superblock mutex
260     async with self.sb_lock:
261         # get inode and fd from the cache
262         (inode, fd, dirty, count) = self.open_files[fh]
263         # if is the last "open"
264         if count == 1:
265             # if the cache is dirty (invalid remote data)
266             if dirty:
267                 # write (replace) data and save pointer
268                 inode.data_pointer = await self.wrapper.write_data(
                ↪ fd, old_to_delete=inode.data_pointer)
269                 # delete from cache
270                 del self.open_files[fh]
271             # otherwise just reduce the counter
272         else:
273             self.open_files[fh] = (inode, fd, dirty, count - 1)
```

Listing 37: Example of `release` implementation.

Listing 37 shows our example of `release` implementation.

After acquiring the superblock lock, the cache is accessed and if the counter is not 1 it's simply decreased. Otherwise, this means that this is the last `release` call and the cached data should be removed. If the `dirty` flag is `True` the `FileData` is uploaded to Telegram replacing the old one and the new id is saved in the `inode` `data_pointer`.

2.4.1.8 `rmdir(self, parent_inode_n, name, ctx)` and `unlink(self, parent_inode_n, name, ctx)`

The handlers `rmdir` and `unlink` are used to delete directories and files (the prototype presented in this work doesn't support links).

The file or directory with the given `name` should be removed from the directory with `parent_inode_n` inode number. As required by the documentation, the entry should be removed from the parent directory but the deletion of the actual data and inode may be deferred until the inode is not known to the kernel anymore.

More detailed requirements can be found on the `pyfuse3` request handlers documentation.

```

275 async def rmdir(self, parent_inode_n, name, ctx):
276     # acquire superblock mutex
277     async with self.sb_lock:
278         await self._remove(parent_inode_n, name, ctx, is_dir=True)
279
280 async def unlink(self, parent_inode_n, name, ctx):
281     # acquire superblock mutex
282     async with self.sb_lock:
283         await self._remove(parent_inode_n, name, ctx)

```

Listing 38: Example of `rmdir` and `unlink` implementation.

Listing 38 shows our example of `rmdir` and `unlink` implementation.

Since the deletion of directories and files is very similar, these two methods just acquire the superblock lock and then call an internal method `_remove`. That method takes care of the actual deletion of the file or directory.

```

285 async def _remove(self, parent_inode_n, name, ctx, is_dir=False):
286     # try to lookup (may raise ENOENT)
287     attr = await self._lookup(parent_inode_n, name, ctx)
288     # if we are removing a directory
289     if is_dir:
290         # ensure it's actually a directory otherwise raise ENOENT
291         if not stat.S_ISDIR(attr.st_mode):
292             raise FUSEError(errno.ENOTDIR)
293         # else if is not empty (only . and ..) raise ENOTEMPTY
294         elif attr.st_size > 2:
295             raise FUSEError(errno.ENOTEMPTY)
296     # if the lookup count is zero
297     if self.lookup_counters[attr.st_ino] == 0:
298         # get the inode
299         inode = self.superblock.get_inode_by_number(attr.st_ino)
300         # remove the file data and free its inode
301         await self.wrapper.delete_data(inode.data_pointer)
302         self.superblock.free_inode(attr.st_ino)
303     else:
304         # defer deletion
305         self.deferred.append(attr.st_ino)
306     # update the parent directory
307     await self._update_directory(parent_inode_n, [('-', name,
↵ attr.st_ino)])

```

Listing 39: Example of `_remove` implementation.

Listing 39 shows our example of `_remove` implementation.

The internal `_remove` method assumes the superblock lock has already been acquired. The attributes are retrieved calling `_lookup` (which may raise `ENOENT`)

then if the method is called to remove a directory, is checked if it's actually a directory (raising `ENOTDIR` otherwise) and if it's empty (raising `ENOTDIR` otherwise). As required by the documentation, if the lookup counter for the inode is 0, then the data and the inode is immediately deleted; otherwise, the deletion is deferred adding the inode to the `deferred` list. In both cases the entry is removed from the parent directory. The actual deletion will be done when the counter reaches 0, see section 2.4.1.9.

2.4.1.9 `forget(self, inode_list)`

The `forget` handler is used to decrease the lookup counters for the given inodes by the given amount. `inode_list` is a list of tuples (`inode`, `amount`) where `inode` is the inode number and `amount` is the amount to be subtracted. If the lookup counter reaches 0 deferred deletion should be performed.

```

309 async def forget(self, inode_list):
310     # acquire superblock mutex
311     async with self.sb_lock:
312         # iterate over the list
313         for (inode_n, amount) in inode_list:
314             # decrease the counter by the given amount
315             self.lookup_counters[inode_n] -= amount
316             # if the lookup count is 0 and the removal is deferred
317             if self.lookup_counters[inode_n] == 0 and inode_n in
                 ↪ self.deferred:
318                 # remove the file data and free its inode
319                 await self.wrapper.delete_data(self.superblock,
                 ↪ .get_inode_by_number(inode_n).data_pointer)
320                 self.superblock.free_inode(inode_n)
321                 # remove it from the deferred list
322                 self.deferred.remove(inode_n)

```

Listing 40: Example of `forget` implementation.

Listing 40 shows our example of `forget` implementation.

After acquiring the superblock lock, lookup counters for all the given inodes are decreased by the given amount. If the counter reaches 0 and the deletion was deferred then it's performed at this point and the inode is removed from the `deferred` list.

2.4.1.10 Cleanup

After the `pyfuse3.main` ends (e.g. when the file system is unmounted) there are some additional steps to perform. The superblock must be updated on Telegram also, as specified in the documentation, the filesystem needs to take care to clean up inodes that at that point still have non-zero lookup count because it may not have received all the `forget` calls.

```

324 async def close(self):
325     # we assume this is called only when not running the fs
326     # for this reason we don't need to acquire the mutex
327     # force-forget any deferred inode with non-zero lookup count
328     await self.forget(
329         [(inode_n, self.lookup_counters[inode_n]) for inode_n in
330          ↪ self.deferred]
331     )
332     # write the superblock
333     await self.wrapper.write_superblock(self.superblock, True)

```

Listing 41: Example of close implementation.

Listing 41 shows our example of close implementation. The method `forget` is explicitly called with the current lookup count for every deferred inode then the superblock is uploaded replacing the old one.

2.4.2 Initialization and Mount

One of the responsibilities of the file `tgfuse.py` is to initialize and mount the filesystem. As shown in listing 42 some arguments such as the mountpoint, phone number and a debug flag are parsed. `pyfuse` is initialized with a `TgFuseFs` instance (linked to the phone number), the mountpoint and other options. Adding the debug option is extremely recommended because it will make `pyfuse` print which handlers are called and their debug information. The `pyfuse3.main` function is executed on the loop until completed and then the `tgfusefs.close` method is called to clean and sync the filesystem. Only a basic error handling is done here and should be improved in a production-grade implementation.

```

335 def main():
336     # parse arguments from command line
337     options = parse_args()
338     # instance a TgFuseFs with the given number
339     tgfusefs = TgFuseFs(options.phone_number)
340     # add fuse options including debug if necessary
341     fuse_options = set(pyfuse3.default_options)
342     fuse_options.add('fsname=tgfuse')
343     if options.debug_fuse:
344         fuse_options.add('debug')
345     # init pyfuse3 with our filesystem implementation and options
346     pyfuse3.init(tgfusefs, options.mountpoint, fuse_options)
347     loop = asyncio.get_event_loop()
348     try:
349         # run pyfuse3.main and then tgfusefs.close
350         loop.run_until_complete(pyfuse3.main())
351         loop.run_until_complete(tgfusefs.close())
352     except:
353         pyfuse3.close(unmount=True)
354         raise
355     finally:
356         loop.close()
357     # close pyfuse3 normally
358     pyfuse3.close()
359
360 def parse_args():
361     parser = ArgumentParser()
362     parser.add_argument('mountpoint', type=str,
363                         help='Where to mount the file system')
364     parser.add_argument('phone_number', type=str,
365                         help='Phone number like +XXXXXXXXXXXX')
366     parser.add_argument('--debug-fuse', action='store_true',
367                         default=False,
368                         help='Enable FUSE debugging output')
369     return parser.parse_args()
370
371
372
373 if __name__ == '__main__':
374     main()

```

Listing 42: Example of mount script implementation.

2.5 Testing

Students have not enough experience to do extensive testing, test performance or prove correctness. Students should mainly do empirical tests manually verifying the expected results while looking at the Telegram chat evolving.

Some of the tests may be:

- mount and unmount the filesystem
- create a file
- create a directory
- change metadata and permissions
- list a directory
- execute stat
- read and write to files
- try to remove files and directory testing edge cases
- ...

Students are invited to perform tests during the development as soon as the functionalities are ready. If the storage on Telegram gets corrupted it should be cleaned using the `mktgfs.py` script, see section 2.3.

Figure 2.1 shows a simple example of testing session and figure 2.2 shows the chat at the end of the session.

Other useful tests are provided by the `pyfuse3` documentation[12] in the “common gotchas” section. For example if the script shown on listing 43 raises an error this means that there probably is a bug in the implementation of the `unlink` handler and the file content is removed immediately instead of deferring the actual deletion to the `forget` handler.

```
1 # assuming the filesystem is mounted at mnt
2 with open('mnt/file_one', 'w+') as fh1:
3     fh1.write('foo')
4     fh1.flush()
5     with open('mnt/file_one', 'a') as fh2:
6         os.unlink('mnt/file_one')
7         assert 'file_one' not in os.listdir('mnt')
8         fh2.write('bar')
9     os.close(os.dup(fh1.fileno()))
10    fh1.seek(0)
11    assert fh1.read() == 'foobar'
```

Listing 43: Unlink handler test from the documentation.

```
Terminal: tests x +
riccardo@tesi-VM:~$ ls mountpoint
riccardo@tesi-VM:~$ echo 42 > mountpoint/test
riccardo@tesi-VM:~$ cat mountpoint/test
42
riccardo@tesi-VM:~$ stat mountpoint/test
  File: mountpoint/test
  Dim.: 3                Blocchi: 0                Blocco di IO: 4096  file regolare
Device: 31h/49d Inode: 3                Coll.: 1
Accesso : (0664/-rw-rw-r--) Uid: ( 1000/riccardo)  Gid: ( 1000/riccardo)
Accesso  : 2019-12-17 23:59:51.743907568 +0100
Modifica : 2019-12-17 23:59:44.792823140 +0100
Cambio   : 2019-12-17 23:59:44.792823140 +0100
Creazione: -
riccardo@tesi-VM:~$ mkdir mountpoint/test_dir
riccardo@tesi-VM:~$ ls -la mountpoint
totale 4
drwxr-xr-x  1 riccardo riccardo   4 dic 18 00:00 .
drwxr-xr-x 19 riccardo riccardo 4096 dic 17 23:58 ..
-rw-rw-r--  1 riccardo riccardo   3 dic 17 23:59 test
drwxrwxr-x  1 riccardo riccardo   2 dic 18 00:00 test_dir
riccardo@tesi-VM:~$ touch mountpoint/test_dir/to_delete
riccardo@tesi-VM:~$ rmdir mountpoint/test_dir/
rmdir: rimozione di 'mountpoint/test_dir/' non riuscita: Directory non vuota
riccardo@tesi-VM:~$ rm mountpoint/test mountpoint/test_dir/to_delete
riccardo@tesi-VM:~$ rmdir mountpoint/test_dir
riccardo@tesi-VM:~$
```

Figure 2.1: Example of testing session.

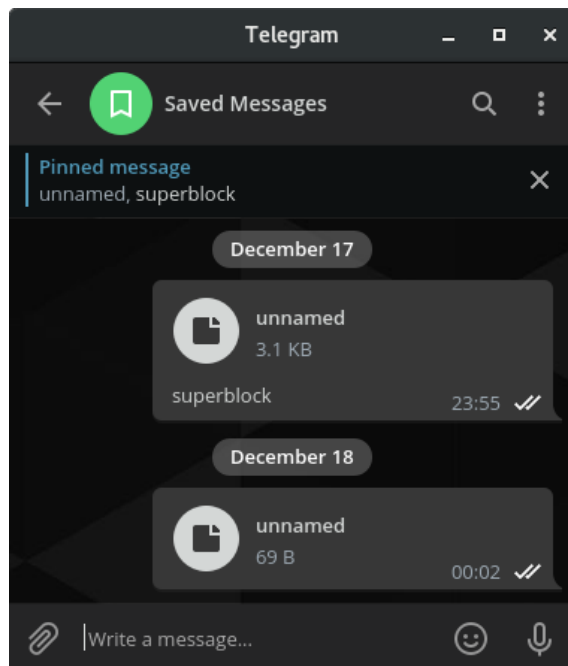


Figure 2.2: Telegram chat after tests.

Chapter 3

Issues and Improvements

The prototype presented in this work has been designed to be understood by fifth year high school students who are majoring in Computer Science. Even if it works, the prototype is a toy and has issues and limitations that students should be aware of.

3.1 Data Storage

In the current prototype implementation each inode has a single data pointer and each file (or directory) is stored completely in a single Telegram attachment. This rises two issues:

1. The maximum file size (or size of the directory content) is 1.5GB.
2. Every time a small part of the file needs to be downloaded, the whole file must be retrieved from Telegram.

There are many possible solutions, one of them would be to divide the file in data blocks and upload them in different Telegram attachments. More advanced techniques could use multi-level data block indexing like those implemented in `ext` filesystems[18].

3.2 Synchronization

In order to avoid race conditions the filesystem needs to use some kind of synchronization technique when accessing shared data structures. The current implementation uses simple locks which are easily understandable by students; however, as a side effect this solution serializes most of the handler calls degrading the overall performance. In a production-grade filesystem this issue should be addressed.

It should be noted that FUSE and the VFS layer already provide some basic locking for example preventing a file creation and a file rename at the same time in the same directory. More about this can be found in the `pyfuse3` documentation[12].

3.3 Caching

Network latency affects the filesystem performance. In the presented implementation files are kept in memory as long as they are open in order to speed up reads and writes. An useful improvement would be to implement a more solid caching mechanism and enabling the kernel write-back cache if available.

3.4 Serialization

Another aspect that could be improved is the serialization. The naive implementation shown in this work uses the `pickle` library to serialize the whole object which is quick and easy but comes with some downsides:

- It's language-specific. Therefore writing software in different languages to access the filesystem can be difficult.
- It's not memory efficient.
- It has serious security implications, for example, malicious data can lead to arbitrary code execution[19].

A solution could be to create a custom binary format for each data structure.

The `Superblock` serialization could be further improved. Unused inodes are not stored at all (and for this reason are not serialized) but there should be a way to keep track of them. The presented implementation uses a set of integers to store the numbers of unused inodes which is more efficient than keeping the inodes but it still has a large memory footprint[20]. An improvement would be to not serialize that set at all and just reconstruct it parsing the inode list during deserialization. This would require to serialize only the maximum inode number.

3.5 Additional Improvements

A lot of additional improvements can be done. Some of them may be:

- support links and other types.
- support other handlers (e.g.: `rename`).
- implement a more robust security and error handling¹ (e.g.: too big directory, invalid flags, wrong permissions etc...).
- implement a more robust failure handling (e.g.: prevent corruption after crash).
- implement better file handles (e.g.: prevent them from growing indefinitely).

¹some scenarios are already checked by FUSE, other should be implemented in the handlers

Chapter 4

Conclusions and Future Works

In this work, we showed how to implement a remote filesystem in userspace using Python, FUSE, and Telegram, as a project example for high school students.

In chapter 1 we presented the requirements along with some advanced Python functionalities, the main libraries used in the prototype and their basic usage. We advise teachers to have students practice beforehand with those functionalities and libraries, for example through some small exercises.

In chapter 2 we showed our example of filesystem prototype presenting and explaining functionalities and major implementation choices. In the same chapter we showed an example of testing that we suggest students do as soon as possible while developing in order to spot and fix bugs and improve their engagement.

In chapter 3 we explained some of the issues and limitations of this prototype. It's extremely important that students understand those issues and the reasons why a production-grade filesystem should be implemented in a different way. Those students who will continue studying Computer Science at university may find it useful to try writing a new implementation and fix the issues using their new operating systems knowledge and experience.

The prototype shown in this work is just an example of how students can write their own first filesystem prototype at school. Teachers can tailor their own versions to their students' needs and learning goals. For example teachers could create a version where most of the filesystem is already implemented and ask students to implement only the missing parts, providing and discussing specifications together. In more advanced classes teachers and students together could discuss and define new and more advanced specifications, in order to co-construct developing strategies and implementation choices.


In the future, with more time and resources, this work could be further improved not only from the technical point of view but also in terms of computing education. From a practical point of view, a teaching unit can be created. In particular, the activities should be detailed providing ready-to-use material and tools to teachers, deciding times and methods of teaching and providing an effective way to assess students learning and help them recover if necessary. From a research point of view, we assumed that this practical approach could improve the operating systems learning experience but that should be proved. A study should be conducted on real students, designing and executing an experiment through which the learning results

of students that experience this practical approach are compared to those of another group of students (the control group) that follow a more traditional approach.

Appendix A

Licenses and Credits

This work is the thesis for the Bachelor’s degree in Computer Science presented by Riccardo Maffei¹.



Except otherwise noted, this work is licensed under the Creative Commons Attribution-ShareAlike 4.0 International License².



Except otherwise noted, the code is licensed under the GNU Affero General Public License³ either version 3 or any later version.

A.1 Third-Party Content

The image shown in figure 1.2 is named “*Structural diagram of Filesystem in Userspace*”⁴ and is licensed by Sven⁵ under the Creative Commons Attribution-ShareAlike 3.0 International License⁶.

The image shown in figure 1.3 is property of Telegram⁷ and is used here with permission obtained through the support.

The code shown in listing 43 is a slightly edited version of the test provided by the `pyfuse3` documentation[12] in the “common gotchas” section and is used here with the author’s permission.

¹<https://orcid.org/0000-0002-6392-9701>

²To view a copy of this license, visit <http://creativecommons.org/licenses/by-sa/4.0/>

³To view a copy of this license, visit <https://www.gnu.org/licenses/agpl-3.0.html>

⁴Source: https://commons.wikimedia.org/wiki/File:FUSE_structure.svg

⁵User page: <https://commons.wikimedia.org/wiki/User:Sven>

⁶To view a copy of this license, visit <https://creativecommons.org/licenses/by-sa/3.0/>

⁷Source: <https://desktop.telegram.org/>

Acknowledgments

I want to thank my thesis supervisor, Professor Davoli, and my co-supervisors, Dr. Sbaraglia and Dr. Lodi, for their patience and precious support during the development of this thesis. I also want to thank all the Professors who have fed my hunger for knowledge and fueled my desire to become a Professor.

Thanks to my friends, classmates and all my teammates with whom I created great projects and wonderful memories. Thanks to Federico, Fox, Francesco, Giosuè, Nicholas, Oleksandr and Samuele.

A special thanks goes to my girl best friend, teammate and partner in crime Teresa. We spent days and nights studying, coding, debugging or looking for a solution to issues together. She helped me study and made me pass exams I wouldn't otherwise. She's been a great friend and supported me while my life was falling apart. Thank you Teresa, "sei la piccola stella che porto nei momenti in cui non ho luce".[21]

Thanks to that little part of my family that decided not to stab me in the back and to everybody who did their best to help me while I was bleeding.

Finally, a thanks goes to those who deserve one but may not like to be directly mentioned here.

References

- [1] Charles C. Bonwell and James A. Eison. *Active Learning: Creating Excitement in the Classroom*. ASHE-ERIC Higher Education Report No. 1, 1991. ISBN: 1-878380-08-7.
- [2] Scott Freeman et al. “Active learning increases student performance in science, engineering, and mathematics”. In: *Proceedings of the National Academy of Sciences* 111.23 (2014), pp. 8410–8415. ISSN: 0027-8424. DOI: [10.1073/pnas.1319030111](https://doi.org/10.1073/pnas.1319030111). eprint: <https://www.pnas.org/content/111/23/8410.full.pdf>. URL: <https://www.pnas.org/content/111/23/8410>.
- [3] Presidente della Repubblica. “Decreto del Presidente della Repubblica 14 marzo 2010, n. 88, Regolamento recante norme per il riordino degli istituti tecnici a norma dell’articolo 64, comma 4, del decreto-legge 25 giugno 2008, n. 112, convertito, con modificazioni, dalla legge 6 agosto 2008, n. 133.” In: *Gazzetta Ufficiale Serie Generale n. 137 - Suppl. Ordinario n. 128* (June 15, 2010), pp. 43–89. URL: <https://www.gazzettaufficiale.it/eli/gu/2010/06/15/137/so/128/sg/pdf>.
- [4] Ministero dell’istruzione, dell’università e della ricerca. “Direttiva 16 gennaio 2012, n. 4, Adozione delle Linee guida per il passaggio al nuovo ordinamento degli Istituti tecnici a norma dell’articolo 8, comma 3, del decreto del Presidente della Repubblica 15 marzo 2010, n. 88 - Secondo biennio e quinto anno.” In: *Gazzetta Ufficiale Serie Generale n. 76 - Suppl. Ordinario n. 60* (Mar. 30, 2012), pp. 1–295. URL: <https://www.gazzettaufficiale.it/eli/gu/2012/03/30/76/so/60/sg/pdf>.
- [5] *asyncio documentation*. <https://docs.python.org/3/library/asyncio.html>. [Online; accessed 15-November-2019].
- [6] Wikipedia contributors. *Time-of-check to time-of-use — Wikipedia, The Free Encyclopedia*. https://en.wikipedia.org/w/index.php?title=Time-of-check_to_time-of-use&oldid=925302898. [Online; accessed 15-November-2019]. 2019.
- [7] *FUSE kernel.org documentation*. <https://www.kernel.org/doc/Documentation/filesystems/fuse.txt>. [Online; accessed 20-November-2019].
- [8] *libfuse GitHub repository*. <https://github.com/libfuse/libfuse>. [Online; accessed 15-November-2019].
- [9] Bradley M. Kuhn and Karen M. Sandler. *GPL Violations Related to Combining ZFS and Linux*. <https://sfconservancy.org/blog/2016/feb/25/zfs-and-linux/>. [Online; accessed 21-November-2019]. Feb. 25, 2016.

- [10] Eben Moglen and Mishi Choudhary. *The Linux Kernel, CDDL and Related Issues*. <https://www.softwarefreedom.org/resources/2016/linux-kernel-cddl.html>. [Online; accessed 21-November-2019]. Feb. 26, 2016.
- [11] *pyfuse3 GitHub repository*. <https://github.com/libfuse/pyfuse3>. [Online; accessed 15-November-2019].
- [12] *pyfuse3 documentation*. <http://www.rath.org/pyfuse3-docs/>. [Online; accessed 15-November-2019].
- [13] *Telegram official website*. <https://telegram.org>. [Online; accessed 11-November-2019].
- [14] *MTPProto documentation*. <https://core.telegram.org/mproto>. [Online; accessed 11-November-2019].
- [15] *Telethon GitHub repository*. <https://github.com/LonamiWebs/Telethon>. [Online; accessed 9-November-2019].
- [16] *Telethon documentation*. <https://docs.telethon.dev/en/latest/>. [Online; accessed 9-November-2019].
- [17] Wikipedia contributors. *Everything is a file* — *Wikipedia, The Free Encyclopedia*. https://en.wikipedia.org/w/index.php?title=Everything_is_a_file&oldid=921228251. [Online; accessed 21-November-2019]. 2019.
- [18] Rémy Card, Theodore Ts'o, and Stephen Tweedie. “Design and Implementation of the Second Extended Filesystem”. In: *Proceedings of the First Dutch International Symposium on Linux*. 1994. ISBN: 9036703859.
- [19] Wikipedia contributors. *Arbitrary code execution* — *Wikipedia, The Free Encyclopedia*. https://en.wikipedia.org/w/index.php?title=Arbitrary_code_execution&oldid=931153642. [Online; accessed 4-January-2020]. 2019.
- [20] M. Gorelick and I. Ozsvald. *High Performance Python: Practical Performant Programming for Humans*. O'Reilly Media, 2014. ISBN: 9781449361778.
- [21] Ultimo. *Piccola stella*. 2019.