

# Local Reasoning for the POSIX File System

Philippa Gardner, Gian Ntzik, and Adam Wright

Imperial College London

{p.gardner,gian.ntzik08,adam.wright07}@imperial.ac.uk

**Abstract.** We provide a program logic for specifying a core subset of the sequential POSIX file system, and for reasoning abstractly about client programs working with the file system.

**Keywords:** file systems, POSIX, local reasoning, separation logic.

## 1 Introduction

Local reasoning, in the style of separation logic, was introduced to reason about programs that manipulate the RAM memory model. Local reasoning has strong modular properties, which means that it scales. Many forms of *abstract* local reasoning have been introduced to specify structured data libraries: e.g. abstract predicates for linked lists [20], concurrent abstract predicates for abstract concurrent sets [7, 24, 23], and context logic for complex structured data such as the DOM [13]. Despite these advances, there are many other properties of real-world libraries that naturally resonate with this local-reasoning approach but have yet to be studied.

We study abstract local reasoning for the POSIX file system [2]. POSIX has an English specification which naturally describes commands which globally follow directory paths to locally update files or directories<sup>1</sup>. There has been much work on traditional reasoning techniques for specifying POSIX, such as the well-known Z specification [18]. However, the global path constraints associated with this work are substantial. Our aim is to use local reasoning to minimise the global path constraints. POSIX is an interesting test case for abstract local reasoning. It has enough emphasis on local update to suggest that the advantages of local reasoning might apply. However, the complexity of the data combined with concurrency, global paths and local update means that the application of local reasoning to this example is not straightforward.

Current work on abstract local reasoning cannot specify POSIX. For example, context logic works well for reasoning about sequential update of complex data, such as the W3C DOM library for XML update [13, 5, 22]. However, it has no mechanism for reasoning about global paths, it does not extend simply to concurrency, and it does not integrate well with ideas from separation logic. Concurrent abstract predicates [7, 24, 23] work well for reasoning abstractly about

---

<sup>1</sup> Both files and directories are called ‘files’ in POSIX. We use the term ‘entries’ to denote either directories or files.

simple concurrent data structures. However, they do not extend to complex data structures since the implementation details leak into the abstraction [12].

We introduce structural separation logic (SSL) for reasoning abstractly about complex structured data. SSL provides more fine-grained reasoning than context logic, leading to straightforward reasoning about disjoint concurrency and a natural integration with separation logic. Here, we demonstrate SSL by reasoning about POSIX. In [25, 12], we provide the general theory which relies substantially on ideas from the views framework [6]<sup>2</sup>. SSL combines fine-grained local reasoning about e.g. directory fragments with global path constraints about the overall structured data. The global path constraints limit the use of the frame rule in our sequential setting, and specify stability requirements of the environment in the concurrent setting. We illustrate our ideas using absolute linear paths, called paths in this paper. In future, it will be very interesting to study general paths (with the backwards `..` and symbolic links) as part of our abstract local reasoning agenda, since they walk right across the directory structure.

In this paper, we use SSL to reason about the sequential POSIX file system, demonstrating that our axioms correspond to the English description given in the POSIX standard. We identify a core subset of POSIX, which is both faithful to the standard and a natural subset with which to introduce our reasoning. We model various structures of the file-system state as standard heaps: file heaps mapping file identifiers (inodes) to bytes; and file-descriptor heaps mapping file descriptors to input/output related data. Separation logic can reason about these heap structures. We require SSL to reason about the directory structure, which we regard as a tree-shaped hierarchy<sup>3</sup>. SSL naturally integrates with separation logic, enabling us to reason about directories and the standard heap within the same logic. We demonstrate this integrated reasoning by verifying natural safety properties of a client software installer. Although we concentrate on sequential POSIX in this paper, our results immediately extend to POSIX with disjoint concurrency. In future, we will explore POSIX with shared-memory concurrency.

**Related Work.** There has been substantial work on formal specifications of file systems [15, 18, 9, 4], leading to a verification challenge by Joshi and Holzmann [17, 10]. It is not feasible to give a comprehensive account of this work in the space available; such an account will be in Ntzik’s thesis [19]. Here, we concentrate on demonstrating the advantages of local directory tree reasoning compared with first-order global tree reasoning and reasoning about heap structures with paths as addresses.

A natural question is whether we might use first-order reasoning as in [8], rather than local reasoning in the style of separation logic. For our program-logic application, first-order reasoning leads to scalability problems. Consider one case of a first-order specification of the `rename(p/a, p’/b)` command:

<sup>2</sup> Previous work on segment logic [14] was too complicated, because we needed views.

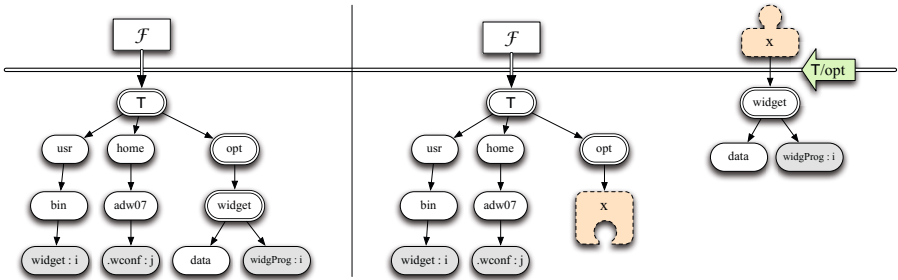
<sup>3</sup> In general, files and directories can be hard linked more than once. Most implementations only allow files to be linked more than once. This is a sensible choice as, for example, cycles generated by directory hard links are not detected by recursive traversal programs. We therefore regard POSIX as a tree-shaped hierarchy.

$$\begin{aligned} & \{ \text{resolve}(p, d[t + a[t']]) \wedge \text{resolve}(p', d'[t'' \wedge \neg \text{exists}(b)]) \wedge \neg \exists p''. p' = p/a/p'' \} \\ & \quad \text{rename}(p/a, p'/b) \\ & \{ \text{resolve}(p, d[t]) \wedge \text{resolve}(p', d'[t'' + b[t']]) \wedge \neg \exists p''. p' = p/a/p'' \} \end{aligned}$$

In the precondition, the assertion  $\text{resolve}(p, d[t + a[t']])$  states that path  $p$  resolves to the directory  $d$  containing the subdirectory  $a$  and list  $t$  of unknown entries. The assertion  $\text{resolve}(p', d'[t'' \wedge \neg \text{exists}(b)])$  states that path  $p'$  resolves to directory  $d'$  with no  $b$  entry. Finally, the assertion  $\neg \exists p''. p' = p/a/p''$  is a path constraint, stating that path  $p'$  cannot be a descendant of  $p/a$  which would be an error case in POSIX. In the postcondition, the assertions state that the directory  $a$  has gone from  $d$ , and a new directory  $b$  has been created under directory  $d'$  with the contents of the old  $a$ .

Now consider program  $\text{rename}(p/a, p'/b) ; \text{rename}(p''/c, p'''/d)$ . In this case, we need path constraints in the precondition stating the following properties: path  $p'$  is not a descendant of path  $p/a$ ;  $p'''$  is not a descendant of  $p''/c$ ; and, in addition,  $p''/c$  is not a descendant of  $p/a$  since directory  $a$  has been removed. These syntactic path checks mushroom as more rename commands are added. Hence, this style of reasoning does not scale. Those familiar with separation logic might recognise that this example is analogous to Reynolds' original list example for justifying separation logic [21].

A completely different approach, used in much of the work on the formal specification of file systems in Z [18] and other methods [15], is to treat paths as heap addresses. Define the set of heaps as  $\text{PATHS} \stackrel{\text{fin}}{\mapsto} \text{BYTES} \cup \mathcal{P}(\text{FNAMES})$ , mapping paths to byte sequences in the file case or sets of names in the directory case. This approach requires significant global constraints: for example, in the specification of  $\text{rename}(p/a, p/a')$  not only we would replace  $p/a$  with  $p/a'$  in the heap, but also every descendant  $p/a/p'$  with  $p/a'/p'$  in order to preserve path consistency.



**Fig. 1.** The left-hand diagram represents a complete directory; the right, the same directory instrumented with abstract addresses

## 2 Example Specifications

We focus on the sequential POSIX file system in this paper, in particular studying a core fragment of 16 commands. Although this fragment is small, it includes most of the primitive commands that manipulate the structure and perform

input-output (IO) and a large proportion of the file system commands can be implemented using them.

First consider the English description of the `rmdir` command<sup>4</sup>:

`[r := rmdir(path)]` Remove the directory identified by `path` and set `r` to 0.

The directory must be empty.

Intuitively, this command traverses the *global* structure, using the global path `path` to identify the location of the update. It then performs *local* update, removing the empty directory whilst leaving the rest of the file system unchanged.

We capture this combination of global traversal and local update using structural separation logic. Consider figure 1. The left-hand side illustrates part of a standard structured heap, consisting of a heap cell with address  $\mathcal{F}$  whose structured value is a complete directory tree. The right-hand side illustrates part of an abstract heap, consisting of the heap cell  $\mathcal{F}$  whose value is now an incomplete directory tree with *body address*  $x$ , and an abstract heap cell with *abstract address*  $x$  whose value is a pair consisting of a *path promise* ‘ $\top/opt$ ’ and, in this case, a complete subdirectory. The path promise provides the stability condition that body address  $x$  must be at the end of ‘ $\top/opt$ ’. This promise to  $x$  allows us to reason locally about the abstract heap cell  $x$  whilst retaining the knowledge of its location in the global structure. Notice that the two heaps illustrated in figure 1 describe the same state, in that they differ only in the instrumentation added by abstract addresses. We move from the left-hand to the right-hand view of the state using *abstract allocation* which creates a new abstract heap cell containing the subdirectory; the converse is *abstract deallocation*.

With structured separation logic (SSL), we can reason about such abstract heap cells. For example, the assertion  $subdir(\alpha @ P, A[\emptyset])$  describes the ownership of the abstract heap cell with address given by logical variable  $\alpha$  and value given by the path promise  $P$  and empty subdirectory  $A[\emptyset]$ . Using this assertion, we are able to provide an axiomatic specification of the `rmdir` command:

$$\begin{aligned} & \{expr(\mathit{path}, P/A) \wedge var(\mathit{r}, -) * \mathcal{E} * subdir(\alpha @ P, A[\emptyset])\} \\ & \quad \mathit{r} := \mathit{rmdir}(\mathit{path}) \\ & \{var(\mathit{r}, 0) * \mathcal{E} * subdir(\alpha @ P, \emptyset)\} \end{aligned}$$

In the precondition, the assertion for the abstract heap cell at  $\alpha$  describes the subdirectory resource necessary for the command to succeed, plus the stable information that the subdirectory is to be found under global path  $P$ . In addition, the precondition contains assertions about path expressions and variables. The *expression assertion*  $expr(\mathit{path}, P/A)$  is a pure assertion which states that expression `path` has logical expression  $P/A$  as its value. This logical expression describes an arbitrary path  $P$  followed by the directory or file name  $A$ . The *variable assertion*  $var(\mathit{r}, -)$  states that the program variable `r` has some

---

<sup>4</sup> This description only presents the case when the operation succeeds. When the command fails, for example if `path` does not identify an existing file or directory, the result is to assign -1 to `r` and set the global variable `errno` to `ENOENT`. We give the full error specifications in our technical report [11]. In this paper we discuss such cases only when required by an example.

arbitrary value, and  $\mathcal{E}$  describes the extra variable resource necessary for `path` to be evaluated. This follows the standard variables-as-resource approach [3]

The postcondition states that variable `r` now has value 0, whilst the abstract heap cell  $\alpha$  is empty with the path promise  $P$ . Notice that we do *not* remove the abstract heap cell  $\alpha$ . If the axiom destroyed this cell, the associated  $\alpha$  body address (which must exist in some data in the frame) would have no matching cell address. This would break the stability of the system, where a cell address always matches with a body address at the appropriate path promise. The additional variable resource predicate  $\mathcal{E}$  is unchanged between the pre and postconditions.

The specification of `rmdir` is *small* in that the precondition intuitively describes local ownership of the minimum resource needed to safely run the command: the variables `r` and those needed to evaluate `path` given by  $\mathcal{E}$ ; and the abstract cell address  $\alpha$  with the subdirectory being updated. It also describes the global information that only (incomplete) directories satisfying path promise  $P$  associated with  $\alpha$  can be framed on. To illustrate this, consider the complete directory  $dir(\mathcal{F}, \top[C+D[A[\emptyset]]])$ , the path  $P = \top/D/A$  and the proof derivation:

```

{  $expr(\text{path}, \top/D/A) \wedge var(\mathbf{r}, -) * \mathcal{E} * dir(\mathcal{F}, \top[C+D[A[\emptyset]]])$  }
// abstract allocation
{  $expr(\text{path}, \top/D/A) \wedge var(\mathbf{r}, -) * \mathcal{E} * \exists \alpha. (dir(\mathcal{F}, \top[C+D[\alpha]]) * subdir(\alpha_{\otimes \top/D}, A[\emptyset]))$  }
// existential elimination and frame rule and apply the axiom
{  $expr(\text{path}, \top/D/A) \wedge var(\mathbf{r}, -) * \mathcal{E} * subdir(\alpha_{\otimes \top/D}, A[\emptyset])$  }
 $\mathbf{r} := rmdir(\text{path})$ 
{  $var(\mathbf{r}, 0) * \mathcal{E} * subdir(\alpha_{\otimes \top/D}, \emptyset)$  }
// existential, frame rule reapplication
{  $var(\mathbf{r}, 0) * \mathcal{E} * \exists \alpha. (dir(\mathcal{F}, \top[C+D[\alpha]]) * subdir(\alpha_{\otimes \top/D}, \emptyset))$  }
// abstract deallocation
{  $var(\mathbf{r}, 0) * \mathcal{E} * dir(\mathcal{F}, \top[C+D[\emptyset]])$  }

```

The initial precondition contains the assertion  $dir(\mathcal{F}, \top[C+D[A[\emptyset]]])$  describing a complete directory tree at the file-system root  $\top$ , with arbitrary contents captured by the logical variable  $C$  and a directory named  $D$  that contains the empty directory  $A$ . This precondition does not match the precondition of `rmdir`, and so we take the following steps. First, we abstractly allocate a new abstract heap cell containing the  $A$  directory, existentially quantifying the abstract address  $\alpha$  to ensure that the address is fresh. Then, we apply the standard Hoare logic existential elimination to set aside the existential binding of  $\alpha$ , and use the frame rule to set aside the resource that `rmdir` does not need. We are now in a position to match `rmdir`'s precondition, where  $\top/D$  is  $P$ . After applying the axiom we can reintroduce the resource and binding set aside with frame and existential elimination, and abstractly deallocate the cell with address  $\alpha$ .

Now consider the `unlink` command and its English specification:

$[\mathbf{r} := \text{unlink}(\text{path})]$  Remove the link to the file identified by `path`. Using SSL, we can formalise the English specification in a similar fashion, with the following small axiom:

$$\begin{aligned} & \{  $expr(\text{path}, P/A) \wedge var(\mathbf{r}, -) * \mathcal{E} * subdir(\alpha_{\otimes P}, A : I)$  \} \\ & \quad \mathbf{r} := \text{unlink}(\text{path}) \\ & \{  $var(\mathbf{r}, 0) * \mathcal{E} * subdir(\alpha_{\otimes P}, \emptyset)$  \} \end{aligned}$$

In the precondition,  $\text{subdir}(\alpha @ P, A : I)$  states that a file named  $A$  is found at abstract cell address  $\alpha$  at the end of path  $P$ . The file data is not included in the precondition, but can be found at file inode  $I$ . When the last link to a file is removed, the file will no longer be accessible by any path, and we assume garbage collection will remove any associated file data.

Finally, consider the `stat` command, which returns meta-data about the file or directory identified by the path argument. In this paper, we take that meta-data to be just the file type,  $D$  for directory and  $F$  for file. There is one axiom for each file type; the directory case is:

$$\left\{ \begin{array}{l} \text{expr}(\text{path}, P/A) \wedge \text{var}(t, -) * \mathcal{E} * \text{subdir}(\alpha @ P, A[\beta]) \\ \mathbf{t} := \text{stat}(\text{path}) \\ \text{var}(t, D) * \mathcal{E} * \text{subdir}(\alpha @ P, A[\beta]) \end{array} \right\}$$

Notice that the specification uses *body address*  $\beta$  in  $A[\beta]$  to specify that the content of  $A$  is not changed by the command. It does not need more detailed knowledge of the contents of  $A$  since the command does not require this knowledge to determine that the entry is a directory.

The commands discussed so far are enough to implement the POSIX command `r := remove(path)`. According to its POSIX description, this command removes the file or empty directory identified by the `path` argument. In figure 2 we implement `remove` and derive its specification. Notice that the derived specification exactly matches the English description obtained from POSIX. Following the same process, we can use the core fragment of this paper to “discover” formal specifications of many more complex commands of POSIX.

```

{ expr(path, P/A) ∧ var(x, -) * E * subdir(α @ P, (A : I ∨ A[∅])) }
r := remove(path) ≜ local t {
  t := stat(path);
  { ∃T. expr(path, P/A) ∧ var(x, -) * var(t, T) * E * subdir(α @ P, (A : I ∧ T = F) ∨ (A[∅] ∧ T = D)) }
  if t = F
  { expr(path, P/A) ∧ var(x, -) * E * subdir(α @ P, A : I) }
  r := unlink(path);
  { var(x, 0) * E * subdir(α @ P, ∅) }
  else if t = D
  { expr(path, P/A) ∧ var(x, -) * E * subdir(α @ P, A[∅]) }
  r := rmdir(path);
  { var(x, 0) * E * subdir(α @ P, ∅) }
  else r := -1;
  { var(x, 0) * var(t, -) * E * subdir(α @ P, ∅) }
}
{ var(x, 0) * subdir(α @ P, ∅) }

```

Fig. 2. An implementation of `remove` and the derived specification

### 3 File System Specification

We provide an axiomatic specification of our sequential POSIX commands using SSL.

### 3.1 Abstract Program State

An abstract program state comprises: an abstract *file-system heap*, which represents the directory tree and associated files, as might intuitively reside on a hard disk; a *process heap*, which represents the computer memory during execution; and a *variable store*, which represents the values of program variables.

**File-System Heaps.** Abstract file-system heaps are abstract heaps whose cells contain partial directories. Directories are defined using a set of *inodes* INODES, ranged over by  $\iota, \kappa, \dots$ , and a set of file names FNAMES, ranged over by  $A, B, \dots$ , for naming directories and files. Both sets are defined as in POSIX. Our partial directories are instrumented by body addresses (context holes), drawn from the countably infinite set of abstract addresses ABSADDRS, ranged over by  $x, y, z, \dots$ , with  $(\text{FNAMES} \cup \{\mathcal{F}\} \cup \text{INODES}) \cap \text{ABSADDRS} = \emptyset$  where  $\mathcal{F}$  is the distinguished address of the root directory.

**Definition 1 (Directories).** *The set of unrooted directories, UDIRS, is:*

$$ud ::= \emptyset \mid a : \iota \mid a[ud] \mid ud + ud \mid x$$

where  $\emptyset$  is the empty list of entries,  $a : \iota$  is a file link associating file name  $a \in \text{FNAMES}$  with inode  $\iota \in \text{INODES}$ ,  $a[ud]$  is a directory named  $a$  containing unrooted abstract directory  $ud$ ,  $+$  is directory composition and  $x \in \text{ABSADDRS}$  is a body address. The directories have sibling-unique names, body addresses are unique, and  $+$  is commutative and associative with identity  $\emptyset$ .

There is a distinguished  $\top \notin \text{FNAMES}$  representing the root directory of the file-system tree. The set of rooted directories, RDIRS, is defined as  $\text{RDIRS} \triangleq \{\top[ud] \mid ud \in \text{UDIRS}\}$ . The set of directories,  $d \in \text{DIRS}$ , is defined by  $\text{DIRS} \triangleq \text{UDIRS} \cup \text{RDIRS}$ . Each directory entry has a type DETYPES  $\triangleq \{F, D\}$ , where  $F$  denotes a hard link to a file and  $D$  a directory.

Each body address can be replaced by entries via *context application*.

**Definition 2 (Context application).** *The addresses function,  $\text{addr} : \text{DIRS} \rightarrow \mathcal{P}(\text{ABSADDRS})$  describes the set of body addresses in a directory. Context application is the function  $\circ : \text{ABSADDRS} \rightarrow (\text{DIRS} \rightarrow \text{UDIRS}) \rightarrow \text{DIRS}$  defined by:*

$$d_1 \circ_x ud_2 = \begin{cases} d_1[ud_2/x] & x \in \text{addr}(d_1) \wedge \text{addr}(d_1) \cap \text{addr}(ud_2) \subseteq \{x\} \\ \text{undefined} & \text{otherwise} \end{cases}$$

where  $d_1[ud_2/x]$  is the substitution of  $ud_2$  for  $x$  in  $d_1$ . The function is defined only if the result is in DIRS.

Many POSIX commands refer to entries in the file system tree by *absolute linear paths* through the directory tree. General paths (with  $\dots$  and symbolic links) are complex but we should be able to handle general paths using a combination of promises and *obligations* discussed in the conclusions: the abstract address  $x$  will have the promise that the part of the path in the context is stable, and the obligation to keep the part of the path in the context stable.

**Definition 3 (Paths and Resolution).** *The set of relative paths, RELPATHS, is defined by:*

$$rp ::= \epsilon \mid a \mid rp/rp$$

where  $a \in \text{FNAMES}$  and the path composition  $/$  is associative with identity  $\epsilon$ . The set of absolute paths is  $\text{ABPATHS} \triangleq \{\top\} \cup \{\top/rp \mid rp \in \text{RELPATHS}\}$ . The set of abstract paths is  $\text{ABSPATHS} = \{p/x \mid p \in \text{ABPATHS}, x \in \text{ABSADDRS}\}$ . The set of paths,  $p \in \text{PATHS}$ , is  $\text{PATHS} \triangleq \text{RELPATHS} \cup \text{ABSPATHS}$ .

The path resolution function  $\text{resolve} : \text{PATHS} \times \text{DIRS} \rightarrow \text{DIRS}$  is defined by:

$$\begin{aligned} \text{resolve}(a, d + a : \iota) &= a : \iota & \text{resolve}(a/rp, d_1 + a[d_2]) &= \text{resolve}(rp, d_2) & \text{if } rp \neq \epsilon \\ \text{resolve}(a, d_1 + a[d_2]) &= a[d_2] & \text{resolve}(\top, \top[d]) &= \top[d] \\ \text{resolve}(x, x + d) &= x & \text{resolve}(\top/rp, \top[d]) &= \text{resolve}(rp, d) & \text{if } rp \neq \epsilon \end{aligned}$$

In all other cases, the result is undefined.

A file-system heap is the union of three finite partial functions: from distinguished address  $\mathcal{F}$  to the root directory which might be partial; from abstract addresses to absolute paths (expressing where the corresponding body address lies) and directories; and from inodes to byte sequences representing file contents. We construct file-system heaps in two phases: first, we define *pre-file-system heaps*; then, we define well-formedness conditions to give the full definition.

**Definition 4 (Pre-file-system Heap).** Let  $\text{BYTES}$  be the set of finite byte sequences. A pre-file-system heap,  $pfs \in \text{PREFS}$ , is a function in the set

$$(\{\mathcal{F}\} \rightarrow \{\epsilon\} \times \text{RDIRS}) \sqcup (\text{ABSADDRS} \xrightarrow{\text{fn}} \text{ABPATHS} \times \text{DIRS}) \sqcup (\text{INODES} \xrightarrow{\text{fn}} \text{BYTES})$$

Let  $\text{inodes}(d)$  denote the set of all inodes occurring in directory  $d$ . A pre-file-system-heap,  $pfs$ , is complete if:  $\text{dom}(pfs) \cap \text{ABSADDRS} = \emptyset$ ;  $pfs(\mathcal{F}) = (\epsilon, rd)$ ;  $\text{addrs}(rd) = \emptyset$ ; and  $\text{inodes}(rd) \subseteq \text{dom}(pfs)$ <sup>5</sup>.

Pre-file-system heaps may use abstract addresses incorrectly. For example, two separate partial directories at different addresses may contain the same body address, or the path promises may not correctly identify the location of the directory. We define a *collapse relation*, with which we give a well-formedness condition that ensures addresses are used correctly. The collapse relation intuitively states that we can connect a cell address to the matching body address with context application, if the paths match, as illustrated in figure 3.




Fig. 3. Collapse relation

<sup>5</sup> Complete pre-file-system-heaps are thus simple DAGs, with sharing occurring only at the leaves in the sense that two separate file names can point to the same inode.



**Definition 5 (Collapse Relation).** *The one-step collapse relation,  $\downarrow \subseteq \text{PREFS} \times \text{PREFS}$ , relates  $pfs_1 \downarrow pfs_2$  if and only if there is some address  $addr \in \text{ABSADDRS} \cup \{\mathcal{F}\}$  and unique  $y \in \text{ABSADDRS}$  such that:*

1.  $pfs_1(addr) = (p, d)$  and  $pfs_1(y) = (p_y, d_y)$ ;
2.  $y \in \text{addrs}(d)$ ;
3. there is some  $q \in \text{PATHS}$  such that  $p_y = p/q$ ;
4.  $\text{resolve}(q, d) = y$ ;
5.  $pfs_2 = pfs_1[addr \mapsto (p, d \circ_y d_y)]/y^6$ .

Let  $\downarrow^*$  be the reflexive, transitive closure of  $\downarrow$ .

Using collapse, we can detect all pre-file-system heaps that use invalid addressing. Given  $pfs$ , the correct use of abstract addressing falls into three cases:

1.  $pfs$  is complete, and is thus trivially uses abstract addresses correctly.
2.  $pfs$  uses abstract addresses, but is related via collapse to a complete abstract file system. In this case, the complete system it is related to must be unique (see [25] for details).
3.  $pfs$  uses abstract addresses, but is not immediately related to a complete file system. However, at least one other pre-abstract file system  $pfs'$  can be found such that the union of the two *does* collapse to a complete file system (as in case 2). In this case,  $pfs$  is a *partial* file-system heap, missing some data, but still using abstract addressing in a consistent way.

With the collapse relation, we can now define file-system heaps.

**Definition 6 (File-system Heaps).** *The set of file-system heaps, FS ranged by  $fs$ , is defined as:*

$$\text{FS} = \{ pfs \in \text{PREFS} \mid \exists pfs', pfs'' \in \text{PREFS}. pfs \sqcup pfs' \downarrow^* pfs'' \wedge pfs'' \text{ is complete} \}$$

**Process Heaps.** The *process heap* represents the contents of the heap during program execution. It contains structures used for controlling access to files and directories: *open file descriptions* and *directory streams*. An open file description is a record holding information that controls file accesses: the inode and current offset of an open file. It is used to support the commands **read**, **write**, **lseek** and **close**. The heap addresses of open file descriptions, in POSIX terminology called *file descriptors*, are given by the set **OFADDRS** and ranged by  $f, g, \dots$

A directory stream is an abstract data structure that captures the set of the entries in a given directory and supports the **opendir**, **readdir** and **closedir** commands. For example, when **opendir(p)** is used, a fresh *directory stream address* from the set **DSADDRS** is allocated and mapped to a directory stream, which provides a snapshot of the entry names in the directory given by path  $p$ . Here, we deviate from POSIX. The **readdir** command returns the names of entries contained within a directory. POSIX allows a high degree of non-determinism when using **readdir** on a directory whilst modifying its contents: one may see some changes; all changes; or none. Specifying the full behaviour is possible, but complex. To aid comprehension, we chose a *snapshot semantics*.

<sup>6</sup> That is,  $pfs_2$  is equal to the function obtained from  $pfs_1$  by removing  $y$  from the domain, mapping  $addr$  to  $(p, d \circ_y d_y)$ , and leaving the other mappings the same.

**Definition 7 (Process Heaps).** A process heap, denoted  $\text{ph} \in \text{PH}$ , is a partial function in the set  $(\text{DSADDRS} \xrightarrow{\text{fin}} \mathcal{P}(\text{FNAMES})) \sqcup (\text{OFADDRS} \xrightarrow{\text{fin}} (\text{INODES} \times \mathbb{N}))$

**Variable Stores.** Variables are assigned values through a *variable store*,  $\sigma : \text{VARS} \rightarrow \text{VALUES}$ , with the set of variable stores denoted  $\Sigma$ . Variables are dynamically typed, with values drawn from the set:

$$\text{VALUES} \triangleq \mathbb{Z} \uplus \{\mathbf{true}, \mathbf{false}\} \uplus \text{RELPATHS} \uplus \text{ABPATHS} \uplus \text{BYTES} \uplus \text{INODES} \\ \uplus \text{OFADDRS} \uplus \text{DSADDRS} \uplus \text{DETTYPES}$$

**Definition 8 (Abstract Program States).** Given the sets of file-system heaps  $\text{FS}$ , process heaps  $\text{PH}$  and variable stores  $\Sigma$ , the set of abstract program states,  $as \in \text{ASTATES}$ , is defined as:  $\text{ASTATES} \triangleq \text{FS} \times \text{PH} \times \Sigma$ .

### 3.2 Programming Language

We define a standard imperative sequential WHILE language with calls to POSIX commands. *Program expressions* are used as the rvalue of assignments and as parameters to control-flow commands. They consist of the standard literals, variable lookup, arithmetic and boolean operations, and path concatenation  $\text{Expr}/\text{Expr}$ . Expression evaluation  $[[\cdot]]_\sigma : \text{EXPR} \rightarrow \Sigma \rightarrow \text{VALUES}$  is mostly standard<sup>7</sup>. Our core POSIX commands can be classified into *structural* commands that manipulate the file system structure, *primitive IO* commands that read and write the contents of files, and *state* commands for querying the type of files.

**Definition 9 (Core Fragment & Programming Language).** The core POSIX fragment consists of structural commands  $\mathbb{C}_{\text{Str}} \in \text{COMM}_{\text{Str}}$ , IO commands  $\mathbb{C}_{\text{IO}} \in \text{COMM}_{\text{IO}}$ , and state commands  $\mathbb{C}_{\text{Stat}} \in \text{COMM}_{\text{Stat}}$ :

$$\begin{aligned} \mathbb{C}_{\text{Str}} ::= & \text{r} := \text{mkdir}(\text{path}) \mid \text{r} := \text{rmdir}(\text{path}) \mid \text{r} := \text{link}(\text{existing}, \text{new}) \\ & \mid \text{r} := \text{unlink}(\text{path}) \mid \text{r} := \text{rename}(\text{old}, \text{new}) \\ \mathbb{C}_{\text{IO}} ::= & \text{dir} := \text{opendir}(\text{path}) \mid \text{fn} := \text{readdir}(\text{dir}) \mid \text{closedir}(\text{dir}) \\ & \mid \text{fd} := \text{open}(\text{path}, \text{flags}) \mid \text{buffer} := \text{read}(\text{fd}, \text{size}) \\ & \mid \text{size} := \text{write}(\text{fd}, \text{buffer}) \\ & \mid \text{offset}' := \text{lseek}(\text{fd}, \text{offset}, \text{whence}) \mid \text{close}(\text{fd}) \\ \mathbb{C}_{\text{Stat}} ::= & \text{t} := \text{stat}(\text{path}) \end{aligned}$$

The commands,  $\mathbb{C} \in \text{COMM}$ , of the programming language are:

$$\begin{aligned} \mathbb{C} ::= & \text{var} := \text{Expr} \mid \text{local var in } \mathbb{C} \mid \text{if Expr then } \mathbb{C} \text{ else } \mathbb{C} \\ & \mid \text{while Expr do } \mathbb{C} \mid \text{skip} \mid \mathbb{C}; \mathbb{C} \mid \mathbb{C}_{\text{Str}} \mid \mathbb{C}_{\text{IO}} \mid \mathbb{C}_{\text{Stat}} \end{aligned}$$

In POSIX, the commands are specified as C function interfaces. Here, we adapt them to a simple imperative programming style for simplicity. Details relating to the semantics of C are thus abstracted. We have formally specified all the commands of this fragment using SSL in [11]. Here, we present specifications for those commands that are used in our examples.

<sup>7</sup> Concatenation:  $[[\text{Expr}/\text{Expr}']]_\sigma \triangleq [[\text{Expr}]]_\sigma / [[\text{Expr}']]_\sigma$  iff  $[[\text{Expr}']]_\sigma \notin \text{ABPATHS}$ .

### 3.3 Assertions

We describe assertions for reasoning about POSIX programs.<sup>8</sup> Analogous to programs using variables and expressions, assertions use logical variables and expressions. Logical variables are mapped to values by a logical environment,  $e \in \text{LEnv}$ , extending program values with directories, paths, abstract addresses, and sets of these values. Logical expressions, denoted by  $E, E'$ , are defined and evaluated similarly to program expressions, disallowing program variables. We denote logical variables with block capitals  $A, B, X, Y, \dots$ , except for abstract address variables denoted  $\alpha, \beta, \dots$ .

Cell Assertions		Directory Assertions	
Directory Tree	$dir(\mathcal{F}, \phi)$	Empty Entry	$\emptyset$
Subdirectory	$subdir(\alpha @ E, \phi)$	File Type Entry	$E : I$
File	$file(I, E)$	Directory Type Entry	$E[\phi]$
File Descriptor	$fd(X, I, E)$	File System Root	$\top[\phi]$
Directory Stream	$ds(X, E)$	Logical Expression	$E$
Heap	$ptr(E, E)$	Entry List	$\phi + \phi$
Variable	$var(\mathbf{var}, E)$	Context Application	$\phi \circ_{\alpha} \phi$
Expression	$expr(\text{Expr}, E)$	Path Resolution	$@E$

**Fig. 4.** Assertion language

Assertions,  $P, Q \in \text{ASRTS}$ , are constructed from: the standard first-order logic connectives and quantifiers; the *separating conjunction* of separation logic,  $P \star Q$ , and its unit,  $\text{emp}$ ; and the cell assertions of figure 4 which describe file-system heaps, process heaps and variable stores. Key is the subdirectory assertion,  $subdir(\alpha @ E, \phi)$ , which combines *local* information  $\phi$  about the partial directory at  $\alpha$ , and *global* information about the environment using path promise  $E$ . It states that, at abstract cell address given by  $\alpha$ , there is a partial subdirectory satisfying directory assertion  $\phi$  (to be explained) which can be rejoined with the main directory using body address  $\alpha$  which must be at the end of path expression  $E$ . The splitting and joining of partial directories gives rise to novel allocation and deallocation axioms, discussed in Section 3.4.

The file assertion,  $file(I, E)$ , describes the file with inode address given by the logical variable  $I$  and contents given by the byte sequence described by logical expression  $E$ . The next three cell assertions describe elements of the process heap and are directly lifted from definition 7. The final two describe the contents of the variable store. The assertion  $var(\mathbf{var}, E)$ , describes program variable  $\mathbf{var}$  with its value given by the logical expression  $E$ . Our core program commands accept parameters given by program expressions. The pure assertion  $expr(\text{Expr}, E)$  states that the program expression  $\text{Expr}$  evaluates to the value of the logical expression  $E$ . The evaluation requires that we own all the variables used in the expression. Since in an arbitrary expression the variables are unknown, we will

<sup>8</sup> We have been asked whether ramified separation logic for reasoning about dags might be worth exploring [16]. It uses the *sepish* connective to say that there is possibly some shared dag structure, but where it is not determined. Here, the dag structure is fully determined at the leaves, so ramified separation logic is not appropriate.

typically use this assertion in conjunction with an *exact* assertion  $\mathcal{E}$ , leading to assertions of the form  $\text{expr}(\text{Expr}, E) \wedge \mathcal{E}$ , where  $\mathcal{E}$  captures all the variable resource required to evaluate  $\text{Expr}$ .

Directory assertions,  $\phi, \psi \in \text{DIRASRTS}$ , are constructed from the standard first-order connectives and quantifiers, and the directory assertions of figure 4 describing the structure of directories, context application and path resolution. Most have been directly lifted from the structure of directories (definition 1). *Context application*,  $\phi \circ_\alpha \psi$ , taken from context logic, describes a directory tree that can be separated into an partial directory satisfying  $\phi$ , with abstract body address  $\alpha$  bound in the assertion, and a partial directory satisfying  $\psi$ . The assertion  $@E$  describes directories in which the path given by  $E$  resolves.

**Definition 10 (Derived Assertions).** *The standard first-order logic assertions are derived from  $\Rightarrow$  and **false**. Additionally, we define the following:*

$$\begin{aligned} \diamond\phi &\triangleq \mathbf{true} \circ_\alpha \phi & \spadesuit\phi &\triangleq \mathbf{true} + \phi \\ \text{complete} &\triangleq \neg\exists\alpha. \diamond\alpha & \text{top\_complete} &\triangleq \neg\exists\alpha. \spadesuit\alpha \\ \text{entry}(A) &\triangleq A[\mathbf{true}] \vee \exists I. (A : I) & \text{top}(\phi) &\triangleq \phi \wedge \text{top\_complete} \\ \text{can\_create}(A) &\triangleq (\neg \spadesuit\text{entry}(A)) \wedge \text{top\_complete} \\ \text{names}(S) &\triangleq \forall A. (A \in S \iff \spadesuit\text{entry}(A)) \wedge \text{top\_complete} \end{aligned}$$

The assertion  $\diamond\phi$  is read “somewhere  $\phi$ ”, and describes directories containing some directory satisfying  $\phi$ . The assertion  $\spadesuit\phi$  is similar, restricted to siblings. The assertion *complete* describes directories that do not contain any abstract body addresses and thus no subdirectory is missing; *top\_complete* is similar, but restricted to siblings. The assertion *top*( $\phi$ ) states that the directory entries satisfy  $\phi$ , and that no sibling entries have been split away. The assertion *can\_create*( $A$ ) states that an entry named  $A$  can be safely created at the current sibling level (used for commands that create new entries such as `mkdir`). Finally, *names*( $S$ ) states that every name in the set  $S$  is present as an entry.

### 3.4 Program Logic

We describe our program logic for reasoning about our core fragment of sequential POSIX, comprising standard rules from separation logic, axioms for specifying the POSIX commands (Figure 5), and *abstract allocation and deallocation* axioms 11. The abstract allocation and deallocation axioms are similar to normal heap allocation and deallocation axioms, but instead of introducing and deleting fresh heap cells, they introduce and delete abstract heap cells in order to split and recombine partial directories. They are essential for our local reasoning about directories, and are only possible due to the recent technological advances of the views framework [6]. For uniformity, we give these as axioms over the *id* command, which has no operational effect. It is a technical device to enable the small axioms in Figure 5 to be used whenever required.

**Definition 11 (Abstract allocation and deallocation axioms).** *The axioms for abstract allocation and abstract deallocation are, respectively:*

$$\begin{aligned} &\{ \text{subdir}(\alpha_{@P}, (\phi_1 \wedge @Q/\beta) \circ_\beta \phi_2) \} \\ &\quad \text{id} \\ &\{ \exists\gamma. (\text{subdir}(\alpha_{@P}, (\phi_1 \wedge @Q/\beta) \circ_\beta \gamma) * \text{subdir}(\gamma_{@P/Q}, \phi_2)) \} \end{aligned}$$

$$\{\exists\gamma. (subdir(\alpha@P, (\phi_1 \wedge @Q/\beta) \circ_\beta \gamma) \star subdir(\gamma@P/Q, \phi_2))\}$$

$$\quad \quad \quad id$$

$$\{subdir(\alpha@P, (\phi_1 \wedge @Q/\beta) \circ_\beta \phi_2)\}$$

The first axiom is *abstract allocation*. The precondition states that there is a partial directory at cell  $\alpha$  with path promise  $P$ . This partial directory can be viewed as an application of two separate parts: the context directory described by  $\phi_1$  which contains a relative path  $Q$  ending in body address  $\beta$ , applied via  $\beta$  to the subdirectory described by  $\phi_2$ . The postcondition states that directory really can be separated into its two subparts: the subdirectory satisfying  $\phi_2$  is “allocated” into its own abstract heap cell  $\gamma$  whose corresponding body address is at absolute path  $P/Q$ ; and the context directory at  $\alpha$  satisfying  $\phi_1$  with  $\gamma$  replacing  $\beta$ . *Abstract deallocation* is the converse: if we know that  $\gamma$  is at the end of path  $Q$  in a directory that is itself at the end of path  $P$ , it is safe to combine the two using context application.

We justify the abstract allocation and deallocation axioms by referring to the collapse relation in Definition 5. Abstract allocation is the assertion equivalent of “expanding” by one step, in that the result introduces one additional abstract address, but still collapses to the same complete heap. Deallocation is the equivalent of a single collapse step, and will still result in the same complete file system. Therefore, whilst abstract (de)allocation changes the abstract addressing in use by a file system, it does not change the underlying file system.

Figure 5 provides the axioms for specifying the commands used in our software installer example, plus the axioms for `rename` as it is the most challenging command. The complete set of axioms is given in [11]. Each axiom must be *stable* with respect to both abstract addresses and path promises. Axioms cannot introduce or remove abstract addresses, and must not invalidate any path promises that have been issued. Commands that alter paths (for example, `rename`) ensure this later point by requiring that the subdirectories described by the preconditions contain no abstract body addresses. Commands such as `rename` and `stat` have multiple axioms, each covering a different behaviour specified in POSIX depending on the precondition state<sup>9</sup>.

Consider the `mkdir(path)` command. According to its POSIX description, it creates a new empty directory identified by `path`. An existing entry with the same name must not already exist. In our precondition, `path` evaluates to a path of the form  $P/B/A$ . The subdirectory assertion  $subdir(\alpha@P, B[C \wedge can\_create(A)])$  states that the subdirectory  $B$  must be at abstract address  $\alpha$  found at the end of path  $P$  with contents  $C$  where the predicate  $can\_create(A)$  (definition 10) states that it is safe to create a new entry  $A$ . In the postcondition, the assertion  $subdir(\alpha@P, B[C + A[\emptyset]])$  states that the empty directory  $A$  has indeed been created. Note that in the case we create a new directory directly under the root, in the path expression  $P$  will be an empty path and  $B$  will be  $\tau$ .

Now consider `link(existing, new)`, which creates a new hard link with path `new` to the file identified by the path `existing`. Its first axiom is similar to that of `mkdir`. In the precondition, it has two subdirectory assertions, one

<sup>9</sup> The preconditions in such cases are mutually exclusive.

$$\begin{array}{l}
\left\{ \begin{array}{l} \text{expr}(\text{path}, P/B/A) \wedge \text{var}(\mathbf{r}, -) * \mathcal{E} \\ * \text{subdir}(\alpha @ P, B[C \wedge \text{can\_create}(A)]) \end{array} \right\} \\
\left\{ \text{var}(\mathbf{r}, 0) * \mathcal{E} * \text{subdir}(\alpha @ P, B[C + A[\emptyset]]) \right\} \\
\left\{ \begin{array}{l} \text{expr}(\text{path}, P/A) \wedge \text{var}(\mathbf{r}, -) \\ * \mathcal{E} * \text{subdir}(\alpha @ P, A[\emptyset]) \end{array} \right\} \\
\left\{ \begin{array}{l} \mathbf{r} := \text{mkdir}(\text{path}) \\ \text{var}(\mathbf{r}, 0) * \mathcal{E} * \text{subdir}(\alpha @ P, \emptyset) \end{array} \right\} \\
\left\{ \begin{array}{l} \text{expr}(\text{old}, P/D/A) \wedge \text{expr}(\text{new}, P/D/B) \\ \wedge \text{var}(\mathbf{r}, -) * \mathcal{E} * \text{subdir}(\alpha @ P, A : I) \\ * \text{subdir}(\beta @ P', D[C \wedge \text{can\_create}(B)]) \end{array} \right\} \\
\left\{ \begin{array}{l} \mathbf{r} := \text{link}(\text{existing}, \text{new}) \\ \text{var}(\mathbf{r}, 0) * \mathcal{E} * \text{subdir}(\alpha @ P, A : I) \\ * \text{subdir}(\beta @ P', D[C + B : I]) \end{array} \right\} \\
\left\{ \begin{array}{l} \text{expr}(\text{existing}, P/D/A) \\ \wedge \text{expr}(\text{new}, P/D/B) \wedge \text{var}(\mathbf{r}, -) * \mathcal{E} \\ * \text{subdir}(\alpha @ P, D[(C + A : I) \wedge \text{can\_create}(B)]) \end{array} \right\} \\
\left\{ \begin{array}{l} \mathbf{r} := \text{link}(\text{existing}, \text{new}) \\ \text{var}(\mathbf{r}, 0) * \mathcal{E} * \text{subdir}(\alpha @ P, D[C + A : I + B : I]) \end{array} \right\} \\
\left\{ \begin{array}{l} \text{expr}(\text{path}, P/A) \wedge \text{var}(\mathbf{r}, -) \\ * \mathcal{E} * \text{subdir}(\alpha @ P, A : I) \end{array} \right\} \\
\left\{ \begin{array}{l} \mathbf{r} := \text{unlink}(\text{path}) \\ \text{var}(\mathbf{r}, 0) * \mathcal{E} * \text{subdir}(\alpha @ P, \emptyset) \end{array} \right\} \\
\left\{ \begin{array}{l} \text{expr}(\text{old}, P/A) \wedge \text{expr}(\text{new}, P'/D/B) \\ \wedge \text{var}(\mathbf{r}, -) * \mathcal{E} * \text{subdir}(\alpha @ P, A[C \wedge \text{complete}]) \\ * \text{subdir}(\beta @ P', D[C' \wedge \text{can\_create}(B)]) \end{array} \right\} \\
\left\{ \begin{array}{l} \mathbf{r} := \text{rename}(\text{old}, \text{new}) \\ \text{var}(\mathbf{r}, 0) * \mathcal{E} * \text{subdir}(\alpha @ P, \emptyset) \\ * \text{subdir}(\beta @ P', D[C' + B[C]]) \end{array} \right\} \\
\left\{ \begin{array}{l} \text{expr}(\text{old}, P/D/A) \wedge \text{expr}(\text{new}, P/D/B) \\ \wedge \text{var}(\mathbf{r}, -) * \mathcal{E} \\ * \text{subdir}(\alpha @ P, D[(C + A[C' \wedge \text{complete}]) \\ \wedge \text{can\_create}(B)]) \end{array} \right\} \\
\left\{ \begin{array}{l} \mathbf{r} := \text{rename}(\text{old}, \text{new}) \\ \text{var}(\mathbf{r}, 0) * \mathcal{E} * \text{subdir}(\alpha @ P, D[C + B[C']]) \end{array} \right\} \\
\left\{ \begin{array}{l} \text{expr}(\text{old}, P/A) \wedge \text{expr}(\text{new}, P'/B) \\ \wedge \text{var}(\mathbf{r}, -) * \mathcal{E} \\ * \text{subdir}(\alpha @ P, A[C \wedge \text{complete}]) \\ * \text{subdir}(\beta @ P', B[\emptyset]) \end{array} \right\} \\
\left\{ \begin{array}{l} \mathbf{r} := \text{rename}(\text{old}, \text{new}) \\ \text{var}(\mathbf{r}, 0) * \mathcal{E} * \text{subdir}(\alpha @ P, \emptyset) \\ * \text{subdir}(\beta @ P', B[C]) \end{array} \right\} \\
\left\{ \begin{array}{l} \text{expr}(\text{old}, P/A) \wedge \text{expr}(\text{new}, P'/B) \\ \wedge \text{var}(\mathbf{r}, -) * \mathcal{E} * \text{subdir}(\alpha @ P, A : I) \\ * \text{subdir}(\beta @ P', B : I') \end{array} \right\} \\
\left\{ \begin{array}{l} \mathbf{r} := \text{rename}(\text{old}, \text{new}) \\ \text{var}(\mathbf{r}, 0) * \mathcal{E} * \text{subdir}(\alpha @ P, \emptyset) \\ * \text{subdir}(\beta @ P', B : I) \end{array} \right\} \\
\left\{ \begin{array}{l} \text{expr}(\text{old}, P/D/A) \wedge \text{expr}(\text{new}, P/D/B) \\ \wedge \text{var}(\mathbf{r}, -) * \mathcal{E} \\ * \text{subdir}(\alpha @ P, D[(C + A : I) \\ \wedge \text{can\_create}(B)]) \end{array} \right\} \\
\left\{ \begin{array}{l} \mathbf{r} := \text{rename}(\text{old}, \text{new}) \\ \text{var}(\mathbf{r}, 0) * \mathcal{E} * \text{subdir}(\alpha @ P, \emptyset) \\ * \text{subdir}(\beta @ P', D[C + B : I]) \end{array} \right\} \\
\left\{ \begin{array}{l} \text{expr}(\text{old}, P/A) \wedge \text{expr}(\text{new}, P/A) \\ \wedge \text{var}(\mathbf{r}, -) * \mathcal{E} \\ * \text{subdir}(\alpha @ P, C \wedge \text{entry}(A)) \end{array} \right\} \\
\left\{ \begin{array}{l} \mathbf{r} := \text{rename}(\text{old}, \text{new}) \\ \text{var}(\mathbf{r}, 0) * \mathcal{E} * \text{subdir}(\alpha @ P, C) \end{array} \right\} \\
\left\{ \begin{array}{l} \text{expr}(\text{path}, P/A) \wedge \text{var}(\mathbf{t}, -) * \mathcal{E} \\ * \text{subdir}(\alpha @ P, A[\beta]) \end{array} \right\} \\
\left\{ \begin{array}{l} \mathbf{t} := \text{stat}(\text{path}) \\ \text{var}(\mathbf{t}, D) * \mathcal{E} * \text{subdir}(\alpha @ P, A[\beta]) \end{array} \right\} \\
\left\{ \begin{array}{l} \text{expr}(\text{path}, P/A) \wedge \text{var}(\mathbf{t}, -) \\ * \mathcal{E} * \text{subdir}(\alpha @ P, A : I) \end{array} \right\} \\
\left\{ \begin{array}{l} \mathbf{t} := \text{stat}(\text{path}) \\ \text{var}(\mathbf{t}, F) * \mathcal{E} * \text{subdir}(\alpha @ P, A : I) \end{array} \right\} \\
\left\{ \begin{array}{l} \text{expr}(\text{path}, P/D) \wedge \text{var}(\text{dir}, -) * \mathcal{E} \\ * \text{subdir}(\alpha @ P, D[\text{top}(C)]) \end{array} \right\} \\
\left\{ \begin{array}{l} \text{dir} := \text{opendir}(\text{path}) \\ \exists H. \text{var}(\text{dir}, H) * \mathcal{E} \\ * \text{subdir}(\alpha @ P, D[C \wedge \text{names}(A)]) \\ * \text{ds}(H, A) \end{array} \right\} \\
\left\{ \begin{array}{l} \text{var}(\text{dir}, H) * \text{var}(\text{fn}, -) \\ * \text{ds}(H, A) \wedge A \neq \{\} \end{array} \right\} \\
\left\{ \begin{array}{l} \text{fn} := \text{readdir}(\text{dir}) \\ \text{var}(\text{fn}, B) * \text{var}(\text{dir}, H) \\ * \text{ds}(H, (A \setminus \{ B \})) \wedge B \in A \end{array} \right\} \\
\left\{ \begin{array}{l} \text{var}(\text{dir}, H) * \text{var}(\text{fn}, -) \\ * \text{ds}(H, \{\}) \end{array} \right\} \\
\left\{ \begin{array}{l} \text{fn} := \text{readdir}(\text{dir}) \\ \text{var}(\text{dir}, H) * \text{var}(\text{fn}, \epsilon) \\ * \text{ds}(H, \{\}) \end{array} \right\} \\
\left\{ \begin{array}{l} \text{var}(\text{dir}, H) * \text{ds}(H, A) \\ \text{closedir}(\text{dir}) \\ \text{var}(\text{dir}, H) \end{array} \right\}
\end{array}$$

Fig. 5. Axioms for some POSIX commands

for each path. The assertion  $\text{subdir}(\alpha@P, A : I)$  states that the **existing** path  $P/A$  identifies a file link named  $A$  to the file with inode address  $I$ . The assertion  $\text{subdir}(\beta@P', D[C \wedge \text{can\_create}(B)])$  states, as in `mkdir`, that an entry with the name we want to create does not already exist. In the postcondition, the assertion  $\text{subdir}(\beta@P', D[C + B : I])$  states that this new entry has been created with another file link to the same file. Note that, in this first axiom, the update takes place between two different directories; in the second, they take place within the same directory.

The `rename(old, new)` command moves and/or renames the entry identified by the path `old` to that identified by `new`. Consider the first axiom where `old` is a directory and `new` does not exist. In the precondition, the assertion  $\text{subdir}(\alpha@P, A[C \wedge \text{complete}])$  states that the subdirectory  $A$  must be complete and the assertion  $\text{subdir}(\beta@P', D[C' \wedge \text{can\_create}(B)])$  states that it is possible to create the new directory  $B$  under  $D$ . The only path constraints are that the global paths  $P$  and  $P'$  must exist in the underlying global directory, thus restricting the application of the frame rule. In contrast to the first-order `rename` axiom discussed in related work, we do not require any additional path constraints to ensure that  $P'$  is not a descendant of  $P/A$ . It comes automatically from the separating conjunction.

Finally, the `dir := opendir(path)` command allocates a new directory stream for the directory identified by `path` and assigns its address to `dir`. In the precondition, the assertion  $\text{top}(C)$  (definition 10) states that the entries of the identified subdirectory  $D$  are complete at the top level. In the postcondition, the assertion  $\text{subdir}(\alpha@P, D[C \wedge \text{names}(A)])$  declares the set  $A$  of all the entries of the directory  $D$ , and uses it in the assertion  $\text{ds}(H, A)$  to describe the allocated directory stream at  $H$ . Elements of the directory stream are obtained via the `readdir` command, for which we have two cases: one when the directory stream is not empty; and one where it is. Note that `readdir` non-deterministically selects which entry name to return and remove from the  $A$  set. This mirrors the fact the order of directory entries is implementation defined in POSIX. The `closedir(dir)` command simply deallocates the directory stream given by `dir`.

**Sequential Soundness.** We believe it is enough to justify our axiomatic specification by comparing it with the POSIX English standard, since the descriptions are naturally close. However, this is perhaps a controversial point. In [11], we give a standard soundness result for sequential programs (no external processes modifying the file system), providing an operational semantics and proving soundness in the style of the views framework [6].

## 4 Software Installer

We now demonstrate our reasoning by considering a *software installer*. New software is typically provided as a bundle, either downloaded onto the users file system or provided on some media containing a file system. The goal is to place the bundle's contents correctly in the users' file system which may involve other tasks such as removing any previous installations and dealing with incompatible

user files. Installers are a common class of client programs that perform complex manipulation of file system structure.

Here, we develop an installer for the fictional software “Widget v2”. It supersedes “Widget v1”, but is incompatible with any v1 user configuration files. Widget v2 consists of a program executable, ‘*widgProg*’ and a data file, ‘*widgData*’. Following common conventions [1], we place the files in ‘ $\tau/opt/widget/$ ’ and create a link from ‘ $\tau/usr/bin/widget$ ’ to ‘ $\tau/opt/widget/widgProg$ ’. An example situation the installer may encounter is that in figure 1, where v1 exists and the user ‘*adw07*’ has a configuration file.

Even though our installer is fictional, it follows a common workflow found in real practice. In our example, this workflow translates to the following steps:

- ① Test if entries already exist at the locations we wish to place Widget v2 files. If they exist, we expect ‘ $\tau/usr/bin/widget$ ’ to be a file and ‘ $\tau/opt/widget$ ’ to be a directory. If this is the case, we remove them. If it is not, the installer aborts without modifying the system to avoid damaging other components.
- ② Check for v1 configuration files in home directories, and remove them where they exist as they are assumed to be incompatible.
- ③ Copy Widget v2 files to the target location on the file system.
- ④ Make a link to the Widget v2 executable, so the user can run it.

Before implementing the installer we need to consider errors. So far, our specifications describe only when commands succeed. However, commands can also fail with an error result. Our installer relies on the `stat` command returning an error when a path does not exist. We consider error specifications for the entire subset in [11]. Here, we discuss only the `ENOENT` error for `stat`, triggered when a path does not resolve to a file or directory. To describe a file system in which a path cannot resolve, we define the following:

$$\text{ENOENT}(P) \triangleq P \equiv \epsilon \vee (\exists P', A, B, P''. P \equiv P'/A/B/P''. \text{subdir}(\alpha @ P', A[\text{can\_create}(B)]))$$

This predicate states that the path `P` has a prefix which can be resolved, but a suffix which *cannot*. All paths which do not resolve will satisfy this specification and with it, we can give the following error axiom for `stat`:

$$\begin{aligned} & \{ \text{expr}(\text{path}, P) \wedge \text{var}(\mathbf{t}, -) * \text{var}(\mathbf{errno}, -) * \mathcal{E} * S \wedge \text{ENOENT}(P) \} \\ & \quad \mathbf{t} := \text{stat}(\text{path}) \\ & \{ \text{var}(\mathbf{t}, -1) * \text{var}(\mathbf{errno}, \text{ENOENT}) * \mathcal{E} * S \} \end{aligned}$$

In the precondition we use the predicate on the value of `path` to assert that we are in the error case. Note that we capture the state satisfying the predicate in the logical variable `S`. In the postcondition, this state is preserved and the global variable `errno` is assigned the error value, for which we use the same name as the predicate for convenience.

To remove an existing Widget installation (point ②) we need to be able to remove non-empty directories, but `rmdir` only removes empty directories. We can implement a program `rmdirRec` that recursively removes all the directories entries before removing the directory itself. The specification is:

$$\begin{aligned} & \{ \text{expr}(\text{path}, P/A) \wedge \text{var}(\mathbf{r}, -) * \mathcal{E} * \text{subdir}(\alpha @ P, A[\text{complete}]) \} \\ & \quad \mathbf{r} := \text{rmdirRec}(\text{path}) \\ & \{ \text{var}(\mathbf{r}, 0) * \mathcal{E} * \text{subdir}(\alpha @ P, \emptyset) \} \end{aligned}$$



Finally, to copy files (point ④), we can implement the program `fileCopy` with the following specification:

$$\left\{ \begin{array}{l} \text{expr}(\text{source}, P/A) \wedge \text{expr}(\text{target}, P'/D) \wedge \text{var}(\mathbf{r}, -) * \mathcal{E} \\ * \text{subdir}(\alpha_{\text{@P}}, A : I) * \text{subdir}(\beta_{\text{@P}'}, D[C \wedge \text{can\_create}(A)]) * \text{file}(I, SD) \end{array} \right\}$$

$$\mathbf{r} := \text{fileCopy}(\text{source}, \text{target})$$

$$\left\{ \begin{array}{l} \exists I'. \text{var}(\mathbf{r}, 0) * \mathcal{E} * \text{subdir}(\alpha_{\text{@P}}, A : I) * \text{subdir}(\beta_{\text{@P}'}, D[C + A : I']) \\ * \text{file}(I, SD) * \text{file}(I', SD) \end{array} \right\}$$

We have implemented both `rmdirRec` and `fileCopy` and derived their specifications in [11].

The installation of a simple, two file program is a surprisingly complex task. We therefore specify our intuitions about good behaviour and prove that our installer matches them. First, we develop abstractions to assist us in the specifications. We use the following predicates to assert that entries may or may not exist within a given directory resource:

$$\begin{aligned} \text{out}(C, A) &\triangleq \text{top}(C) \wedge \neg \oplus \text{entry}(A) & \text{in}(C, A) &\triangleq C + \text{entry}(A) \\ \text{infile}(C, A) && &\triangleq C + \exists I. A : I \end{aligned}$$

The first predicate describes directory entries  $C$  in which an entry named  $A$  does not exist, whereas the second describes entries  $C$  in which it does.  $\text{infile}(C, A)$  is more specific and describes directory entries  $C$  with an  $A$  file entry.

We build a precondition for our installer out of several sub-assertions with the help of the above predicates. In these, the assertion  $\vdash_{X \in E} \phi$  is the iterated version of  $\vdash$ , interpreted as  $\phi_1 + \dots + \phi_{|E|}$  where each  $\phi_i$  has  $X$  bound to a distinct member of  $E$ .

$$\begin{aligned} \text{src}_{\text{Pre}} &\triangleq \text{subdir}(\delta_{\text{@IL}}, \text{widgProg} : J + \text{widgData} : K) * \text{file}(J, \text{PROG}) * \text{file}(K, \text{DAT}) \\ \text{home}_{\text{Pre}} &\triangleq \text{subdir}(\alpha_{\text{@T}}, \text{home} [\vdash_{(N, C) \in H} N [\text{infile}(C, \text{.wconf}) \vee \text{out}(C, \text{.wconf})]]) \\ \text{bin}_{\text{Pre}} &\triangleq \text{subdir}(\gamma_{\text{@T/usr}}, \text{bin} [\text{in}(B, \text{widget}) \vee \text{out}(B, \text{widget})]) \\ \text{opt}_{\text{Pre}} &\triangleq \text{subdir}(\beta_{\text{@T}}, \text{opt} [\text{top}(T) + (\emptyset \vee \text{widget}[\text{T}_w \wedge \text{complete}])]) \end{aligned}$$

Each of these describes the states that parts of the file system may be in for the installer to run safely. The directory entries and file data that make up the `Widget v2` installation sources are described by  $\text{src}_{\text{Pre}}$ . We require them to be in a location given by the variable `il`.  $\text{home}_{\text{Pre}}$  captures all the home directories of the system, along with the fact that some of them will contain a `v1` `.wconf` configuration file. The  $\text{bin}_{\text{Pre}}$  resource captures the UNIX executables directory, that may contain a `widget` entry. Finally,  $\text{opt}_{\text{Pre}}$  describes the target installation directory, which may already contain a previous installation, which we require to be complete, as it will be deleted.

We combine these descriptions into a precondition, where we also snapshot the initial state in the logical variable  $W$ , to show that nothing changes in the event of an error.

$$P_i \triangleq \text{var}(\mathbf{il}, \text{IL}) * \text{var}(\mathbf{r}, -) * \text{var}(\mathbf{errno}, -) \wedge W \wedge \text{src}_{\text{Pre}} * \text{home}_{\text{Pre}} * \text{bin}_{\text{Pre}} * \text{opt}_{\text{Pre}}$$

If the installer errors, we expect the file system to be unchanged. If it succeeds, we expect `Widget v1` to be installed successfully. There should be no other outcome. We describe a successful installation with the following sub-assertions:

```

{ Pi }
r := installWidgetV2 ≐
local t1, t2, hDir, user {
  { var(t1, -) * var(t2, -) * var(errno, -) * binPre * optPre }
  // Check for preexisting files (point ①). The installer expects  $\top$ /usr/bin/widget
  // to be a file and  $\top$ /opt/widget to be a directory, if they exist.
  t1 := stat(' $\top$ /usr/bin/widget'); t2 := stat(' $\top$ /opt/widget');
  {
    var(t1, T1)  $\wedge$  (T1 = F  $\vee$  D  $\vee$  -1) * var(t2, T2)  $\wedge$  (T2 = F  $\vee$  D  $\vee$  -1)
    * var(errno, E)  $\wedge$  ((T1 = -1  $\vee$  T2 = -1)  $\Rightarrow$  E = ENOENT) * binPre * optPre
  }
  if t1 = D  $\vee$  t2 = F
    // There are preexisting entries, but not of a previous installation.
    // The installer ends here without any modifications.
    r = -1;
  else
    // Either previous entries do not exist, or they are of a previous installation.
    { var(x, -) * var(t1, F)  $\vee$  var(t1, -1) * var(t2, D)  $\vee$  var(t2, -1) * binPre * optPre }
    if t1 = F
      // Remove previous installation executable.
      r := unlink('usr/bin/widget');
    if t2 = D
      // Remove previous installation directory. We apply the rmdirRec specification.
      { var(t2, D) * var(x, -) * subdir( $\beta$ @ $\top$ , opt[T + widget[Tw  $\wedge$  complete]]) }
      r := rmdirRec('opt/widget');
      { var(t2, D) * var(x, 0) * subdir( $\beta$ @ $\top$ , opt[T]) }
      {
        var(x, 0) * var(t1, F)  $\vee$  var(t1, -1) * var(t2, D)  $\vee$  var(t2, -1)
        * subdir( $\gamma$ @ $\top$ /usr, bin[B]) * subdir( $\beta$ @ $\top$ , opt[T])
      }
      // Remove any stale Widget configuration files (point ②)
      { var(hDir, -) * var(user, -) * subdir( $\alpha$ @ $\top$ , home[+((N,C)  $\in$  H N [in file(C, .wconf)  $\vee$  out(C, .wconf)])]) }
      hDir := opendir('home');
      user := readdir(hDir);
      {
         $\exists$ Hd, U, Us. var(hDir, Hd) * var(user, U) * ds(Hd, Us)
        * subdir( $\alpha$ @ $\top$ , home[+((N,C)  $\in$  H N [out(C, .wconf) + N  $\in$  Us  $\Rightarrow$  ( $\emptyset$   $\vee$   $\exists$ I. .wconf: I)  $\wedge$  N  $\notin$  Us  $\Rightarrow$   $\emptyset$ ]]) )
      }
      while user  $\neq$   $\epsilon$ 
        // We iterate over every user's home directory and delete the file.
        // If the file does not exist, then unlink returns -1 as in stat.
        r := unlink('home/user/.wconf'); user := readdir(hDir);
      closedir(hDir);
      // In the end, there are no Widget V1 configuration files.
      { subdir( $\alpha$ @ $\top$ , home[+((N,C)  $\in$  H N [C  $\wedge$  can_create(.wconf)])]) }
      // Now we create the new installation, copy the new Widget files
      // and link the executable (Points ③ and ④)
      {
        var(x, -) * var(il, IL) * subdir( $\delta$ @IL, v2DirPre) * subdir( $\alpha$ @ $\top$ , homePost)
        * subdir( $\gamma$ @ $\top$ /usr, bin[B]) * subdir( $\beta$ @ $\top$ , opt[T]) * file(J, PROG) * file(K, DAT)
      }
      r := mkdir('opt/widget');
      r := fileCopy(il/'widgProg', ' $\top$ /opt/widget');
      r := fileCopy(il/'widgData', ' $\top$ /opt/widget');
      r := link('opt/widget/widgProg', 'usr/bin/widget'); r := 0
      {
         $\exists$ J', K'. var(x, 0) * var(il, IL) * srcPre * homePost
        * subdir( $\beta$ @ $\top$ , opt[T + widget[widgProg: J' + widgData: K']])
        * subdir( $\gamma$ @ $\top$ /usr, bin[B + widget: J]) * file(J', PROG) * file(K', DAT)
      }
    }
    {
       $\exists$ R, J', K'. var(il, IL) * var(x, R) * var(errno, -)  $\wedge$  (R = -1  $\Rightarrow$  W)
       $\wedge$  R = 0  $\Rightarrow$  (srcPre * homePost * optPost(J', K') * binPost(K') * v2FilesPost(J', K'))
    }
  }
{ Qi }

```

Fig. 6. Widget v2 software installer

$$\begin{aligned}
v2Files_{\text{Post}}(J, K) &\triangleq \text{file}(J, \text{PROG}) \star \text{file}(K, \text{DAT}) \\
\text{home}_{\text{Post}} &\triangleq \text{subdir}(\alpha_{\text{@T}}, \text{home}[+_{(N,C)\in H} N[C \wedge \text{can\_create}(\text{.wconf})]]) \\
\text{bin}_{\text{Post}}(J) &\triangleq \text{subdir}(\gamma_{\text{@T/usr}}, \text{bin}[B + \text{widget} : J]) \\
\text{opt}_{\text{Post}}(J, K) &\triangleq \text{subdir}(\beta_{\text{@T}}, \text{opt}[T + \text{widget}[\text{widgProg} : J + \text{widgDat} : K]])
\end{aligned}$$

The postcondition is built from these sub-assertions.

$$\begin{aligned}
Q_i &\triangleq \exists R, J', K'. \text{var}(\text{i1}, \text{IL}) \star \text{var}(\mathbf{r}, R) \star \text{var}(\text{errno}, -) \wedge (R = -1 \Rightarrow W) \\
&\wedge R = 0 \Rightarrow (\text{src}_{\text{Pre}} \star \text{home}_{\text{Post}} \star \text{opt}_{\text{Post}}(J', K') \star \text{bin}_{\text{Post}}(J') \star v2Files_{\text{Post}}(J', K'))
\end{aligned}$$

Note that if the installer fails, the return variable  $\mathbf{r}$  has value -1 and the state of the file system is the same as in the precondition, captured by the logical variable  $W$ . Otherwise,  $\mathbf{r}$  is 0 and the state changes according to the sub-assertions that we have defined.

Our installer implementation, along with a proof that it meets its specification,  $\{P_i\}\text{installWidgetV2}\{Q_i\}$ , is given in figure 6. Throughout the proof we make implicit use of the frame rule to temporarily discard irrelevant state, and at the points of axiom application we implicitly use abstract allocation/deallocation.

## 5 Conclusions and Future Work

The POSIX file system provides an interesting challenge for local reasoning: complex abstract data update with global absolute paths for identifying the place to do local update. We give a natural axiomatic specification of the sequential POSIX file system using SSL; the general theory is in [25, 12]. We verify safety properties for a client software installer, demonstrating integrated reasoning for the file system and the heap. Our POSIX reasoning provides an illustrative example of reasoning about global access and local update; other natural applications include identifying the  $i$ th element of a list [25] and querying the DOM.

The promises in our POSIX reasoning are naturally stable. Wright has also explored the combination of promises and *obligations*: promises on abstract heap cells give information about what can be relied upon by the environment; obligations gives information about what data fragments must guarantee; sometimes both are needed for stability. In this paper, the only obligations are that the abstract cell and body addresses must be preserved. In general, understanding obligations is hard. A natural test example would be to extend the core POSIX fragment presented here with non-linear paths ( $\dots$  and symbolic links), where the paths can move back and forth over the structure. We also believe obligations will be useful for file-access permissions and shared-memory concurrency.

**Acknowledgements.** We acknowledge funding from an EPSRC DTA (Ntzik, Wright) and EPSRC programme grant EP/H008373/1 (Gardner, Ntzik and Wright). We also thank Pedro da Rocha Pinto, Ramana Kumar, Azalea Raad, Tom Ridge and Mark Wheelhouse for many interesting discussions.

## References

- [1] Filesystem Hierarchy Standard Group. Filesystem hierarchy standard
- [2] POSIX.1-2008, IEEE 1003.1-2008, The Open Group Base Specifications Issue 7
- [3] Variables as resource in separation logic. *Electronic Notes in Theoretical Computer Science* 155, 247–276 (2006)
- [4] Arkoudas, K., Zee, K., Kuncak, V., Rinard, M.: Verifying a File System Implementation. In: Davies, J., Schulte, W., Barnett, M. (eds.) *ICFEM 2004*. LNCS, vol. 3308, pp. 373–390. Springer, Heidelberg (2004)
- [5] Calcagno, C., Gardner, P., Zarfaty, U.: Context logic and tree update. *SIGPLAN Not.* (2005)
- [6] Dinsdale-Young, T., Birkedal, L., Gardner, P., Parkinson, M., Yang, H.: Views: compositional reasoning for concurrent programs. In: *POPL (2013)*
- [7] Dinsdale-Young, T., Dodds, M., Gardner, P., Parkinson, M.J., Vafeiadis, V.: Concurrent abstract predicates. In: D’Hondt, T. (ed.) *ECOOP 2010*. LNCS, vol. 6183, pp. 504–528. Springer, Heidelberg (2010)
- [8] Fisher, K., Foster, N., Walker, D., Zhu, K.Q.: Forest: a language and toolkit for programming with filestores. In: *ICFP (2011)*
- [9] Freitas, L., Fu, Z., Woodcock, J.: POSIX file store in Z/Eves: an experiment in the verified software repository. In: *IEEE International Conference on Engineering of Complex Computer Systems (2007)*
- [10] Freitas, L., Woodcock, J., Butterfield, A.: POSIX and the verification grand challenge: A roadmap. In: *ICECCS (2008)*
- [11] Gardner, P., Ntzik, G., Wright, D.: Local Reasoning for the POSIX File System. Technical report, Imperial College London (2014), <http://www.doc.ic.ac.uk/~gn408/POSIXFS/>
- [12] Gardner, P., Raad, A., Wheelhouse, M., Wright, A.: Abstract Local Reasoning for Concurrent Libraries. In preparation (2014)
- [13] Gardner, P., Smith, G., Wheelhouse, M., Zarfaty, U.: Local Hoare reasoning about DOM. In: *PODS (2008)*
- [14] Gardner, P., Wheelhouse, M.: Small specifications for tree update, <http://www.doc.ic.ac.uk/~pg/papers/move.pdf>
- [15] Hesselink, W.H., Lali, M.: Formalizing a hierarchical file system. In: *REFINE (2009)*
- [16] Hobor, A., Villard, J.: The ramifications of sharing in data structures. In: *POPL (2013)*
- [17] Joshi, R., Holzmann, G.J.: A mini challenge: build a verifiable filesystem. *Form. Asp. Comput.* (2007)
- [18] Morgan, C., Sufrin, B.: Specification of the UNIX Filing System. *IEEE Transactions on Software Engineering* (1984)
- [19] Ntzik, G.: Local Reasoning about File Systems. PhD thesis (expected, 2014)
- [20] Parkinson, M., Bierman, G.: Separation logic and abstraction. In: *POPL (2005)*
- [21] Reynolds, J.C.: Separation logic: A logic for shared mutable data structures. In: *LICS (2002)*
- [22] Smith, G.: Local Reasoning about Web Programs. PhD thesis (2011)
- [23] Svendsen, K., Birkedal, L.: Impredicative concurrent abstract predicates. In: Shao, Z. (ed.) *ESOP 2014*. LNCS, vol. 8410, pp. 149–168. Springer, Heidelberg (2014)
- [24] Turon, A., Dreyer, D., Birkedal, L.: Unifying refinement and hoare-style reasoning in a logic for higher-order concurrency. In: *ICFP (2013)*
- [25] Wright, A.: Structural Separation Logic. PhD thesis, Imperial College London (2013)