

# Local Reasoning about POSIX File Systems

Philippa Gardner, Gian Ntzik, and Adam Wright

Imperial College London  
p.gardner@imperial.ac.uk  
gian.ntzik08@imperial.ac.uk  
adam.wright07@imperial.ac.uk

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

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

## 1 Introduction

Local reasoning was originally introduced in separation logic [20], where the goal was to provide scalable, modular reasoning about low-level programs for manipulating heaps defined as partial functions from addresses to values. Since then, there has been much work on abstraction. We have separation logic with abstract predicates for reasoning locally about e.g. linked lists [19], concurrent abstract predicates for more abstract reasoning about e.g. sequential and concurrent sets [7, 22], context logic for reasoning abstractly about complex structured data such as the W3C DOM library for XML update [13, 5, 21], and structural separation logic for reasoning about concurrent trees [23, 12].

Despite these recent advances in abstract local reasoning, we still have insufficient technology for reasoning *abstractly* about programs that manipulate complex data by *globally* traversing the data structure to a place where they *locally* update. Such programs are common place. For example, the POSIX file system has an English specification which naturally describes commands which globally follow directory paths to locally update files or directories<sup>1</sup>. We provide a formal specification of a core subset of the POSIX file system, with the aim to be as close to the POSIX English specification as we can. We achieve this by extending structural separation logic with *path promises* that require stable information about where a directory exists within a complete file system. It enables us to provide integrated reasoning about the directory structure and the standard heap within the same logic. We demonstrate this integrated reasoning by verifying safety properties of a client software installer.

File systems are essential for manipulating persistent storage. The most commonly used file-system abstraction standard is POSIX [2]. We spent considerable effort identifying a core subset of sequential POSIX, which is both faithful to the

---

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

standard and a natural subset with which to introduce our reasoning. Various structures of the file-system state are naturally modelled as heaps: file heaps mapping file identifiers (inodes) to bytes; file-descriptor heaps mapping file descriptors to input/output related data. Separation logic can reason about these heap structures. Our challenge is to reason about the directory structure, which we regard as a tree-shaped hierarchy<sup>2</sup>, where directory tree nodes are identified by global path resolution and their updates are naturally local.

We reason about this directory tree hierarchy, extending structural separation logic with *path promises*. Structural separation logic (SSL) is a program logic for reasoning locally about complex structured data. With SSL, we can for example reason about DOM commands which have direct access to tree nodes using DOM identifiers [23, 12]. SSL allows fine-grained reasoning about tree fragments using *abstract heaps* which provide instrumentation for separating and joining tree fragments. SSL is a substantial improvement over previous work on context-logic reasoning about DOM. However, basic SSL cannot reason about global paths. We extend SSL with *promises* which declare stable global properties about data structures. The general theory is in Wright’s thesis [23]. Here we introduce the extension by studying POSIX, to ground our theory on real application. We provide small axioms which naturally correspond to the POSIX English standard. We provide integrated reasoning about the directory tree structure and the standard heap structures. Although we concentrate on sequential POSIX, our results immediately extend to a disjoint concurrent POSIX since our reasoning is underpinned by the views framework. In future, we will explore shared-memory concurrent POSIX as defined by the standard. We use our reasoning to verify that a client software installer behaves sensibly: if the installer fails, the file system will be unchanged; if it succeeds, the program will be successfully installed; and there will be no other possible outcomes.

**Related Work** There has been substantial work on the formal specification of file systems [14, 17, 9, 4], even leading to a verification challenge by Joshi and Holzmann [16, 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 [18]. Here, we concentrate on demonstrating the advantages of our local directory tree reasoning approach compared to reasoning about global directory trees and heap structures with paths as addresses.

A natural question is whether we can use first-order reasoning about a global tree directory, as in [8]. However, this leads to scalability problems. Consider the incorrect specification of the `rename(p/a, p’/b)` command:

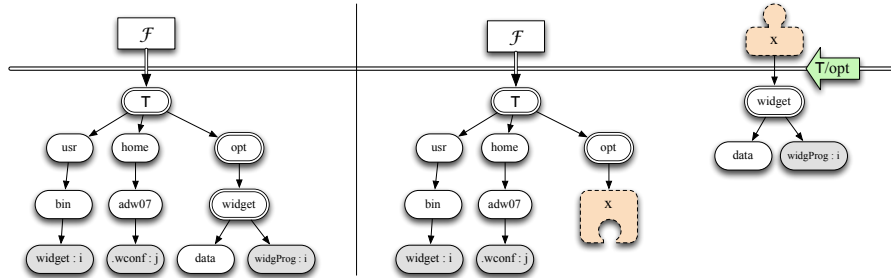
$$\frac{\{ \text{resolve}(p, d[t + a[t']]) \wedge \text{resolve}(p', d'[t'' \wedge \neg \text{exists}(b)]) \}}{\text{rename}(p/a, p'/b)} \{ \text{resolve}(p, d[t]) \wedge \text{resolve}(p', d'[t'' + b[t']]) \}$$

---

<sup>2</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.

$resolve(p, d[t + a[t']])$  states that path  $p$  resolves to the directory  $d$  containing an arbitrary forest  $t$  and the directory  $a$ ; also,  $resolve(p', d'[t'' \wedge \neg exists(b)])$  states that  $p'$  resolves to directory  $d'$  with no  $b$  entry. This precondition is incorrect because POSIX specifies an error when  $p'$  is a descendant of  $p/a$ : that is,  $p'$  is  $p/a/p''$  for some path  $p''$ . To correct the specification we require a syntactic path check to identify this error case. Now consider the command `rename(p/a, p'/b) ; rename(p''/c, p'''/d)`. In this case, we need syntactic checks that  $p'$  is not a descendant of  $p/a$ ,  $p'''$  is not a descendant of  $p''/c$ , and also  $p''/c$  is not a descendant of  $p/a$ . It is evident that this style of reasoning does not scale. Those familiar with separation logic might recognise that this example is analogous to Reynolds' list example for justifying separation logic.

A completely different approach, used in much of the work on the formal specification of file systems using Z [17] and other methods as [14], is to treat paths as heap addresses. Define the set of heaps as  $PATHS \stackrel{fin}{\mapsto} BYTES \cup \mathcal{P}(FNAMES)$ , 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, for every path in the domain mapping to a directory, paths to every descendant must also be in the domain and must be consistent with each other. This makes the specification of `rename(p/a, p'/b)` even more complicated, as it not only involves updating the heap cell of path  $p/a$ , but also the heap cells of all the descendants to ensure the validity of their path addresses.



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

## 2 Example Specifications

A substantial part of the POSIX standard describes the file system. We focus on a core fragment of 16 commands in this paper. Although small, it includes most primitive commands that manipulate the structure and perform input-output (IO). In fact, a large proportion of the file system commands can be implemented using our core commands. This means that we can derive specifications of these other commands without additional machinery, as we shall demonstrate.

First consider the English description of the `rmdir` command<sup>3</sup>:

<sup>3</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

$[r := \text{rmdir}(\text{path})]$  Remove the directory identified by  $\text{path}$  and set  $r$  to 0. The directory must be empty.

Intuitively, this command acts *globally* by traversing the global path  $\text{path}$  to identify the location of the update, and it acts *locally* by removing the empty directory whilst leaving the rest of the file system unchanged. We capture this combination of global and local behaviour using structural separation logic extended with *promises*. Consider figure 1. The left hand side illustrates part of a structured heap, consisting of a heap cell  $\mathcal{F}$  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 cell address*  $x$  whose value is a pair consisting of a complete subdirectory and a *path promise* ‘ $\top/\text{opt}$ ’ providing the stability condition that body address  $x$  must be at the end of ‘ $\top/\text{opt}$ ’. 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. When we move from the left to the right hand side view of the state we apply *abstract allocation* in that we allocate a new abstract heap cell, while the converse is *abstract deallocation*.

With SSL and path promises, we can reason about such abstract heap cells. Our assertion  $\alpha^P \mapsto A[\emptyset]$  describes the ownership of the abstract heap cell denoted by logical variable  $\alpha$  containing an empty directory  $A[\emptyset]$  and associated with path promise  $P$ . Using it, we are able to provide an axiomatic specification of the `rmdir` command:

$$\{\text{path} \mapsto P/A \wedge r \Rightarrow - * \mathcal{E} * \alpha^P \mapsto A[\emptyset]\} r := \text{rmdir}(\text{path}) \{r \Rightarrow 0 * \mathcal{E} * \alpha^P \mapsto \emptyset\}$$

As well as the assertion for the abstract heap cell at  $\alpha$ , the precondition contains assertions about the store, derived from the standard variables as resource approach [3]. The pure *expression assertion*  $\text{path} \mapsto P/A$  states that expression  $\text{path}$  has logical expression  $P/A$  as its value. This logical expression describes an arbitrary path  $P$  followed by the directory or file name  $A$ . Similarly, the *variable assertion*  $r \Rightarrow -$  states that the program variable  $r$  has some arbitrary value. The postcondition states that variable  $r$  now has value 0, whilst abstract heap cell  $\alpha$  is empty with the path promise  $P$ . The additional variable resource predicate  $\mathcal{E}$  captures any additional variable resource needed to evaluate the expression  $\text{path}$ , and is unchanged between the pre and postconditions.

This specification 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  $\text{path}$  given by  $\mathcal{E}$ ; and the abstract cell address  $\alpha$  with the directory being updated. It also describes the global information that only directories satisfying path promise  $P$  associated with  $\alpha$  can be framed on. To illustrate this, consider the directory  $\mathcal{F} \mapsto \top[C+D[A[\emptyset]]]$ , the path  $P = \top/D/A$  and the proof derivation:

---

result is to assign -1 to  $r$  and set the global variable `errno` to `ENOENT`. We discuss error specifications in the appendix.

```

{path ⇒ τ/D/A * r ⇒ - * ε * F ↦ τ[C + D[A[∅]]]}
// abstract allocation
{path ⇒ τ/D/A * ε * r ⇒ - * ∃α. (F ↦ τ[C + D[α]] * ατ/D ↦ A[∅])}
// existential elimination and frame rule
// and apply the axiom
{path ⇒ τ/D/A * r ⇒ - * ε * ατ/D ↦ A[∅]}
r := rmdir(path)
{path ⇒ τ/D/A * r ⇒ 0 * ε * ατ/D ↦ ∅}
// existential, frame rule reapplication
{path ⇒ τ/D/A * r ⇒ 0 * ε * ∃α. (F ↦ τ[C + D[α]] * ατ/D ↦ ∅)}
// abstract deallocation
{r ⇒ - * ε * F ↦ τ[C + D[∅]]}

```

The initial precondition contains a directory hierarchy beginning at the file system root  $\tau$  with arbitrary contents captured by the logical variable  $C$  and a directory named  $D$  that contains just the empty directory  $A$ . This 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. 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 what `rmdir` does not need. We now match `rmdir`'s precondition, where  $\tau/D$  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:

`[r := unlink(path)]` Remove the link to the file identified by `path`.

Again, using SSL we can formalise the English specification with the following small axiom:

$$\{\text{path} \Rightarrow P/A \wedge r \Rightarrow - * \varepsilon * \alpha^P \mapsto A : I\} r := \text{unlink}(\text{path}) \{\varepsilon * r \Rightarrow 0 * \alpha^P \mapsto \emptyset\}$$

In the precondition,  $\alpha^P \mapsto A : I$  states that a file named  $A$  is found at abstract cell address  $\alpha$  at the end of path  $P$ . Note that the contents of the file are not included in the precondition. When the last link to a file is removed, the file will no longer be accessible by any path, and we assume a garbage collection step will remove any associated file data. 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 match with a body address at the appropriate path promise.

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:

$$\{\text{path} \Rightarrow P/A \wedge t \Rightarrow - * \varepsilon * \alpha^P \mapsto A[\beta]\} t := \text{stat}(\text{path}) \{t \Rightarrow D * \varepsilon * \alpha^P \mapsto A[\beta]\}$$

Notice that the specification uses *body address*  $\beta$  in  $\alpha^P \mapsto A[\beta]$  to specify that the content of  $A$  is not changed by the command. It does not need more de-

tailed 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 in this section are enough to derive another POSIX command:  $r := \text{remove}(\text{path})$ . 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 using the `stat`, `unlink` and `rmdir` commands discussed earlier. The specification that we derive is:

$$\begin{aligned} & \{ \text{path} \Rightarrow P/A \wedge r \Rightarrow - * \mathcal{E} * \alpha^P \mapsto (A : I \vee A[\emptyset]) \} \\ & \quad r := \text{remove}(\text{path}) \\ & \{ r \Rightarrow 0 * \mathcal{E} * \alpha^P \mapsto \emptyset \} \end{aligned}$$

Note that this matches exactly the English description obtained from the POSIX standard. Following the same process, we can use the core fragment of this paper to “discover” formal specifications of many more complex commands of POSIX.

```

{ path ⇒ P/A ∧ r ⇒ - * ℰ * αP ↦ (A : I ∨ A[∅]) }
r := remove(path) ≜ local t {
  t := stat(path);
  { ∃T. path ⇒ P/A ∧ r ⇒ - * t ⇒ T * ℰ * αP ↦ (A : I ∧ T = F) ∨ (A[∅] ∧ T = D) }
  if t = F
  { path ⇒ P/A ∧ r ⇒ - * ℰ * αP ↦ A : I }
  r := unlink(path);
  { path ⇒ P/A ∧ r ⇒ 0 * ℰ * αP ↦ ∅ }
  else if t = D
  { path ⇒ P/A ∧ r ⇒ - * ℰ * αP ↦ A[∅] }
  r := rmdir(path);
  { path ⇒ P/A ∧ r ⇒ 0 * ℰ * αP ↦ ∅ }
  else r := -1;
  { r ⇒ 0 * t ⇒ - * ℰ * αP ↦ ∅ }
}
{ path ⇒ P/A * r ⇒ 0 * αP ↦ ∅ }

```

Fig. 2. Implementation of `remove` and proof that it meets its specification.

### 3 File System Specification

We provide an axiomatic specification of our POSIX commands, based on structural separation logic extended with naturally stable promises for describing update on abstract heaps. An abstract program state comprises: an abstract *file-system heap*, which represents the directory tree and associated stored files, as might intuitively reside on a hard disk; a *process heap*, which represents the contents of computer memory during execution; and a *variable store*, which represents the values of program variables.

#### 3.1 File-system Heap

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  $\mathbf{x}, \mathbf{y}, \mathbf{z}, \dots$ , 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:*

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

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  $\mathbf{x} \in \text{ABSADDRS}$  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  $\text{DETTYPES} \triangleq \{\text{F}, \text{D}\}$ : the F denotes a hard link to a file, and D a directory.

Each body address can be replaced with a directory via *context application*.

**Definition 2 (Context application).** *The addresses function,  $\text{addrs} : \text{DIRS} \rightarrow \mathcal{P}(\text{ABSADDRS})$  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:*

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

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

Many POSIX commands refer to entries in the file system tree by *absolute (or rooted) paths* through the directory tree. In general, POSIX paths are complex with e.g. backwards paths ( $\dots$ ) and symbolic links. In this paper, we work with simple *linear paths*, just called paths for this paper, where the path structure matches the inductive directory structure. In fact, general POSIX paths will be very interesting for our overall agenda of local reasoning about complex data structures, since they walk right across the structure. We should be able to handle general paths using a combination of promises and *obligations* discussed in the conclusions: the abstract address  $\mathbf{x}$  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  $\text{ABSPATHS} = \{p/\mathbf{x} \mid p \in \text{ABPATHS}, \mathbf{x} \in \text{ABSADDRS}\}$ . The set of paths,  $p \in \text{PATHS}$ , is  $\text{PATHS} \triangleq \text{RELPATHS} \cup \text{ABPATHS} \cup \text{ABSPATHS}$ .

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

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

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

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

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  $\text{inodes}(rd) \subseteq \text{dom}(pfs)^4$ .

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

**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  $\mathbf{addr} \in \text{ABSADDRS} \cup \{\mathcal{F}\}$  and unique  $\mathbf{y} \in \text{ABSADDRS}$  such that:

1.  $pfs_1(\mathbf{addr}) = (p, d)$  and  $pfs_1(\mathbf{y}) = (p_{\mathbf{y}}, d_{\mathbf{y}})$ ;

<sup>4</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.



2.  $\mathbf{y} \in \text{addrs}(d)$ ;
3. *there is some  $q \in \text{PATHS}$  such that  $p_{\mathbf{y}} = p/q$ ;*
4.  $\text{resolve}(q, d) = \mathbf{y}$ ;
5.  $pfs_2 = pfs_1[\mathbf{addr} \mapsto (p, d \circ_{\mathbf{y}} d_{\mathbf{y}})]/\mathbf{y}$ <sup>5</sup>.

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 [23] 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 define file-system heaps.

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

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

### 3.2 Process Heap

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. `readdir` 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. We adopt a *snapshot semantics*. A copy of the set of entries is taken when `opendir` is used. Calls to `readdir` return elements of that set in a non-deterministic order, but will observe no future changes.

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

<sup>5</sup> That is,  $pfs_2$  is equal to the function obtained from  $pfs_1$  by removing  $\mathbf{y}$  from the domain, mapping `addr` to  $(p, d \circ_{\mathbf{y}} d_{\mathbf{y}})$ , and leaving the other mappings the same.

### 3.3 Programming Language

To write programs that use our POSIX subset, we define a standard imperative sequential WHILE language. 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:

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

*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 three file-system specific expressions: path concatenation,  $\text{Expr}/\text{Expr}$ ; base name extraction  $\mathbf{base}(\text{Expr})$ ; and directory name extraction  $\mathbf{dirExpr}$ . Expression evaluation  $\llbracket \cdot \rrbracket_{\sigma} : \text{EXPR} \rightarrow \Sigma \rightarrow \text{VALUES}$  is mostly standard, with the file system specific evaluations being:

$$\begin{aligned} \llbracket \text{Expr}/\text{Expr}' \rrbracket_{\sigma} &\triangleq \llbracket \text{Expr} \rrbracket_{\sigma} / \llbracket \text{Expr}' \rrbracket_{\sigma} && \text{iff } \llbracket \text{Expr}' \rrbracket_{\sigma} \notin \text{ABPATHS} \\ \llbracket \mathbf{base}(\text{Expr}) \rrbracket_{\sigma} &\triangleq a && \text{iff } \llbracket \text{Expr} \rrbracket_{\sigma} \equiv \top / rp / a \wedge a \notin \epsilon \\ \llbracket \mathbf{dir}(\text{Expr}) \rrbracket_{\sigma} &\triangleq rp && \text{iff } \llbracket \text{Expr} \rrbracket_{\sigma} \equiv \top / rp / a \wedge a \notin \epsilon \\ \llbracket \mathbf{base}(\text{Expr}) \rrbracket_{\sigma} &\triangleq \top && \text{iff } \llbracket \text{Expr} \rrbracket_{\sigma} \equiv \top \\ \llbracket \mathbf{dir}(\text{Expr}) \rrbracket_{\sigma} &\triangleq \top && \text{iff } \llbracket \text{Expr} \rrbracket_{\sigma} \equiv \top \end{aligned}$$

Together, file system heaps, process heaps and variable stores form the *abstract program states* of file system programs.

**Definition 8 (Abstract Program States).** *The set of abstract program states,  $as \in \text{ASTATES}$ , is defined as:  $\text{ASTATES} \triangleq \text{FS} \times \text{PH} \times \Sigma$*

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  $\mathbb{C}_{Str} \in \text{COMM}_{Str}$ , IO commands  $\mathbb{C}_{IO} \in \text{COMM}_{IO}$ , and state commands  $\mathbb{C}_{Stat} \in \text{COMM}_{Stat}$ :*

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

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

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

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.

### 3.4 Program Logic

We describe our *program logic* for reasoning about POSIX programs, extending structural separation logic [23, 12] with naturally stable promises for describing update on abstract heaps.<sup>6</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  $\alpha, \beta, \dots$

Abstract Directory Cell	$\alpha^E \mapsto \phi$	Empty Entry	$\emptyset$
File Cell	$I \xrightarrow{F} E$	File Type Entry	$E : I$
Open File Descriptor Cell	$X \xrightarrow{\text{OF}} (I, E)$	Directory Type Entry	$E[\phi]$
Directory Stream Cell	$X \xrightarrow{\text{DS}} E$	File System Root	$\top[\phi]$
Normal Heap Cell	$E \xrightarrow{H} E$	Logical Expression	$E$
Variable Value	$\text{var} \Rightarrow E$	Entry List	$\phi + \phi$
Expression Value	$\text{Expr} \Rightarrow E$	Context Application	$\phi \circ_\alpha \phi$
		Path Resolution	$@E$

**Fig. 4.** Assertion language

Assertions,  $P, Q \in \text{ASRTS}$ , are constructed from: standard first order logic connectives and quantifiers; the *separating conjunction* of separation logic,  $P \star Q$ , and its unit,  $\text{emp}$ ; and the specific assertions of figure 4 describing file-system heaps, process heaps and variable stores. The key assertion is the directory cell assertion,  $\alpha^E \mapsto \phi$ , 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 assertion axioms, which we discuss after introducing the remaining assertions.

The file cell assertion,  $I \xrightarrow{F} E$ , 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  $\text{var} \Rightarrow E$ , as described previously in section 2, maps the program variable  $\text{var}$  to the value given by the logical expression  $E$ . Our core program commands accept parameters given by program expressions. The assertion  $\text{Expr} \Rightarrow E$  states that the program expression  $\text{Expr}$  evaluates to

<sup>6</sup> We have been asked whether ramified separation logic for reasoning about dags might be worth exploring [15]. 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.

the value of the logical expression  $E$ . This expression evaluation requires that we own all the variables used in the expression. Since in an arbitrary expression the variables are unknown, we will typically use this assertion in conjunction with an *exact* assertion  $\mathcal{E}$ , leading to assertions of the form  $\text{Expr} \Rightarrow E \wedge \mathcal{E}$ , where  $\mathcal{E}$  captures all the variable resource required to evaluate  $\text{Expr}$ .

Directory assertions,  $\phi, \psi \in \text{DIRASRTS}$ , are constructed from standard first-order connectives and quantifiers, and the 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). The *context application* assertion,  $\phi \circ_\alpha \psi$ , taken from context logic, describes a directory tree that can be separated into an partial directory satisfying  $\phi$  and containing abstract body address  $\alpha$ , 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:*

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

The assertion  $\Diamond\phi$  is read “somewhere  $\phi$ ”, and describes directories containing some directory satisfying  $\phi$ . The assertion  $\Diamond\phi$  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.

A crucial part of the reasoning is the *abstract allocation and deallocation* axioms. They 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 axioms to be used whenever required.

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

$$\begin{aligned}
&\{\alpha^P \mapsto ((\phi_1 \wedge @Q/\beta) \circ_\beta \phi_2)\} \text{id} \{ \exists\gamma. (\alpha^P \mapsto (\phi_1 \wedge @Q/\beta) \circ_\beta \gamma * \gamma^{P/Q} \mapsto \phi_2) \} \\
&\{ \exists\gamma. (\alpha^P \mapsto (\phi_1 \wedge @Q/\beta) \circ_\beta \gamma * \gamma^{P/Q} \mapsto \phi_2) \} \text{id} \{ \alpha^P \mapsto ((\phi_1 \wedge @Q/\beta) \circ_\beta \phi_2) \}
\end{aligned}$$

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 two separate parts: the context directory described by  $\phi_1$  which contains a relative path  $Q$  ending in body address  $\beta$ ; and the subdirectory described by  $\phi_2$ . The postcondition states that directory really can be separated into its two subparts: the subdirectory satisfying  $\phi_2$  can be “allocated” into its own abstract heap cell  $\gamma$  whose corresponding body address is at rooted path  $P/Q$ ; and the context directory satisfying  $\phi_1$  still at  $\alpha$ . *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.

### 3.5 Axiomatic Specification

Figure 5 provides the axioms for the commands used in our software installer example, plus the axioms for `rename` as it is the most challenging command. The rest are IO commands for regular files and their axioms are given in the appendix. 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 pre-conditions contain no abstract body addresses.

Consider the `mkdir(path)` command. According to its POSIX description, it creates a new empty directory identified by the `path` argument. An existing entry with the same name must not already exist. In our specification’s precondition, the `path` argument evaluates to a path of the form  $P/B/A$ , and the abstract heap cell  $\alpha^P \mapsto B[C \wedge \text{can\_create}(A)]$  states that the directory  $B$  must be at abstract address  $\alpha$  found at the end of path  $P$  with contents  $C$  where the predicate  $\text{can\_create}(A)$  (definition 10) states that it is safe to create a new entry  $A$ . Indeed, in the postcondition, a new empty directory named  $A$  is 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  $\tau$ .

Next, 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` except that it involves two paths with one abstract heap cell for each path. The  $\alpha$  heap cell simply states that the `existing` path  $P/A$  identifies a file link named  $A$  to the file with inode address  $I$ . The  $\beta$  heap cell states as in `mkdir` that an entry with the name we want to create does not already exist. In the postcondition this new entry is created with another file link to the same file. Note that, in the first `link` axiom, the update takes place between two different directories, while in the second they take place within the same directory.

The `rename` command moves and/or renames the entry identified by the path `old` to that identified by `new`. Consider the first axiom case, where `old`

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

Fig. 5. Some of the POSIX axiomatic specifications.

is a directory and `new` does not exist. This captured by the two abstract heap cells  $\alpha$  and  $\beta$  in the precondition. Note that in the  $\alpha$  heap cell, we require the directory `A` to be complete: we capture the entire subtree. This means that the target of the operation is not the same as or a descendant of this directory. In the postcondition, the entire contents of the directory, captured by `C` have been moved to the target location. As in `link` we distinguish the cases of different and same directory updates with distinct axioms.

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 we require the contents of the directory to be complete at the top level with the  $top(C)$  predicate (definition 10). In the allocated directory stream,  $H \xrightarrow{DS} A$ , the set of the directory's entry names `A` is given by the  $names(A)$  predicate. 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 that in POSIX the order of directory entries is implementation defined. The `closedir` command simply deallocates the directory stream given as argument.

**Definition 12 (Hoare Logic).** *The Hoare judgements  $\vdash \{P\} C \{Q\}$ , are constructed by the POSIX command axioms and the standard inference rules of separation logic (including the frame rule)<sup>7</sup>.*

**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 the appendix, we give a theoretical soundness result, providing an operational semantics and proving soundness in the style of the views framework [6]. More interestingly, with Tom Ridge, we are currently exploring a more practical soundness result. Ridge has an ML implementation of the POSIX file system. Together, we are working on a mechanised HOL soundness result with respect to his implementation, plus testing and comparison of his implementation with a real-world implementation.

## 4 Software Installer

We now demonstrate our reasoning by considering a *software installer*. Newly obtained 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 take this bundle and place the contents at the correct points in the users' file system such that the software can run. This may involve other tasks such as removing any previous installations and dealing with incompatible user files. Software installers are a common class of client file system programs that perform complex manipulation of file system structure.

<sup>7</sup> The angelic behaviour of abstract allocation means we do not have the rule of conjunction. This is typical of angelic behaviours, as is well-known from concurrency theory.

Here, we develop an installer for the fictional software “Widget Version 2”. It supersedes “Widget Version 1”, but is incompatible with any V1 user configuration files. Widget V2 consists of a program executable, ‘*widgProg*’ and a data file, ‘*widgData*’. We follow common conventions for placing these files in the file system [1]. The two files will be placed in the directory ‘ $\tau/opt/widget/$ ’. The program will be made usable by creating a hard link from ‘ $\tau/usr/bin/widget$ ’ to the file ‘ $\tau/opt/widget/widgProg$ ’. An example situation the installer may encounter is that in figure 1, where Widget 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 the path does not exist. We discuss error specifications for all of the commands in the appendix. Here, we discuss only the `ENOENT` error for `stat`, triggered when a path argument does not resolve to a file or directory. To describe a file system tree in which a path cannot resolve, we define the following:

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

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`:

$$\begin{aligned} \{ \text{path} \mapsto P \wedge t \Rightarrow - * \text{errno} \Rightarrow - * \mathcal{E} * S \wedge \text{ENOENT}(P) \} \\ t := \text{stat}(\text{path}) \\ \{ t \Rightarrow -1 * \text{errno} \Rightarrow \text{ENOENT} * \mathcal{E} * S \} \end{aligned}$$

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. The command `rmdir` only removes empty directories. To remove a non-empty directory we can implement a program `rmdirRec` that recursively removes all of a directories entries before removing the directory itself. The specification is:



$$\left\{ \begin{array}{l} \text{path} \Rightarrow P/A \wedge r \Rightarrow - * \mathcal{E} * \alpha^P \mapsto A[\text{complete}] \\ r := \text{rmdirRec}(\text{path}) \\ r \Rightarrow 0 * \mathcal{E} * \alpha^P \mapsto \emptyset \end{array} \right\}$$

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

$$\left\{ \begin{array}{l} \text{source} \Rightarrow P/A \wedge \text{target} \Rightarrow P'/D \wedge r \Rightarrow - * \mathcal{E} \\ * \alpha^P \mapsto A : I * \beta^{P'} \mapsto D[C \wedge \text{can\_create}(A)] * I \xrightarrow{F} \text{SD} \\ r := \text{fileCopy}(\text{source}, \text{target}) \\ \exists I'. r \Rightarrow 0 * \mathcal{E} * \alpha^P \mapsto A : I * \beta^{P'} \mapsto D[C + A : I'] * I \xrightarrow{F} \text{SD} * I' \xrightarrow{F} \text{SD} \end{array} \right\}$$

We have implemented both `rmdirRec` and `fileCopy` and derived their specifications in the appendix.

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:

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

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  $E$ .

$$\begin{aligned} \text{instSrcPre} &\triangleq \delta^{\text{IL}} \mapsto (\text{widgProg} : J + \text{widgData} : K) * J \xrightarrow{F} \text{PROG} * K \xrightarrow{F} \text{DAT} \\ \text{homePre} &\triangleq \alpha^{\text{T}} \mapsto \text{home} \left[ \vdash_{(N, C) \in \text{H}} N [\text{infile}(C, \text{wconf}) \vee \text{out}(C, \text{wconf})] \right] \\ \text{binPre} &\triangleq \gamma^{\text{T/usr}} \mapsto \text{bin} [\text{in}(B, \text{widget}) \vee \text{out}(B, \text{widget})] \\ \text{optPre} &\triangleq \beta^{\text{T}} \mapsto \text{opt} [\text{top}(T) + (\emptyset \vee \text{widget}[\text{T}_w \wedge \text{complete}])] \end{aligned}$$

Each of these describes the states that parts of the file system may be in for the Widget installer to run safely. The directory entries and file data that make up the Widget version 2 installation sources are described by  $\text{instSrcPre}$ . We require them to be in a location given by the variable `iloc`.  $\text{homePre}$  captures all the home directories of the system, along with the fact that some of them will contain a Widget V1 ‘*wconf*’ configuration file. The  $\text{binPre}$  resource captures the UNIX executables directory, that may contain a ‘*widget*’ entry. Finally,  $\text{optPre}$  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.

$$P_{\text{ins}} \triangleq \text{iloc} \Rightarrow \text{IL} * r \Rightarrow - * \text{errno} \Rightarrow - \wedge W \wedge \text{instSrcPre} * \text{homePre} * \text{binPre} * \text{optPre}$$

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:

$$\begin{aligned}
v2Files_{Post}(J, K) &\triangleq J \stackrel{F}{\mapsto} \text{PROG} \star K \stackrel{F}{\mapsto} \text{DAT} \\
home_{Post} &\triangleq \alpha^T \mapsto \text{home} [ +_{(N,C) \in H} N [ C \wedge \text{can\_create}(\cdot, \text{wconf}) ] ] \\
bin_{Post}(J) &\triangleq \gamma^{T/usr} \mapsto \text{bin} [ B + \text{widget} : J ] \\
opt_{Post}(J, K) &\triangleq \beta^T \mapsto \text{opt} [ T + \text{widget} [ \text{widgProg} : J + \text{widgDat} : K ] ]
\end{aligned}$$

The postcondition is built from these sub-assertions.

$$\begin{aligned}
&\exists R, J', K'. \text{iloc} \Rightarrow \text{IL} \star \mathbf{r} \Rightarrow R \star \text{errno} \Rightarrow - \wedge (R = -1 \Rightarrow W) \\
Q_{ins} &\triangleq \wedge R = 0 \Rightarrow \left( \begin{array}{l} \text{instSrcPre} \star \text{home}_{Post} \star \text{opt}_{Post}(J', K') \star \text{bin}_{Post}(J') \\ \star v2Files_{Post}(J', K') \end{array} \right)
\end{aligned}$$

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,  $\{P_{ins}\} \text{installWidgetV2} \{Q_{ins}\}$ , 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 as described in section 2.

## 5 Conclusions and Future Work

The POSIX file system provides an interesting challenge for local reasoning: complex abstract data update with global rooted paths for identifying the place to do local update. We have extended structural separation logic with *path promises*, to provide a natural axiomatic specification of the POSIX file system; the general theory is in [23]. We are able to verify properties of client programs such as the software installer, demonstrating integrated reasoning for the file system and the heap. Our POSIX reasoning provides an illustrative example of this combination of abstract global and local reasoning; others include reasoning about the  $i$ th element of a list [23] and the combination of JQuery and 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 ( $\cdot\cdot$  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.

```

{  $P_{ins}$  }
r := installWidgetV2 ≐
local t1, t2, hDir, user {
  {  $t1 \Rightarrow - * t2 \Rightarrow - * errno \Rightarrow - * bin_{Pre} * opt_{Pre}$  }
  // 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( $\top/opt/widget$ );
  {
     $t1 \Rightarrow T1 \wedge (T1 = F \vee D \vee -1) * t2 \Rightarrow T2 \wedge (T2 = F \vee D \vee -1)$ 
    {  $* errno \Rightarrow E \wedge ((T1 = -1 \vee T2 = -1) \Rightarrow E = ENOENT) * bin_{Pre} * opt_{Pre}$  }
  }
  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.
    {  $r \Rightarrow - * t1 \Rightarrow F \vee t1 \Rightarrow -1 * t2 \Rightarrow D \vee t2 \Rightarrow -1 * bin_{Pre} * opt_{Pre}$  }
    if t1 = F
      // Remove previous installation executable.
      r := unlink( $\top/usr/bin/widget$ );
    if t2 = D
      // Remove previous installation directory. We apply the rmdirRec specification.
      {  $t2 \Rightarrow D * r \Rightarrow - * \beta^T \mapsto opt[T + widget[T_w \wedge complete]]$  }
      r := rmdirRec( $\top/opt/widget$ );
      {  $t2 \Rightarrow D * r \Rightarrow 0 * \beta^T \mapsto opt[T]$  }
      {  $r \Rightarrow 0 * t1 \Rightarrow F \vee t1 \Rightarrow -1 * t2 \Rightarrow D \vee t2 \Rightarrow -1 * \gamma^{\top/usr} \mapsto bin[B] * \beta^T \mapsto opt[T]$  }
      // Remove any stale Widget configuration files (point ②)
      {  $hDir \Rightarrow - * user \Rightarrow - * \alpha^T \mapsto home[+(N,C) \in H N[infile(C, .wconf) \vee out(C, .wconf)]]$  }
      hDir := opendir( $\top/home$ );
      user := readdir(hDir);
      {
         $\exists HD, U, Us. hDir \Rightarrow HD * user \Rightarrow U * HD \xrightarrow{DS} Us$ 
        {  $* \alpha^T \mapsto home[+(N,C) \in H N[out(C, .wconf) + N \in Us \Rightarrow (\emptyset \vee \exists I. .wconf : I) \wedge N \notin Us \Rightarrow \emptyset]]$  }
      }
      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( $\top/home/user/.wconf$ ); user := readdir(hDir);
      closedir(hDir);
      // In the end, there are no Widget V1 configuration files.
      {  $\alpha^T \mapsto home[+(N,C) \in H N[C \wedge can\_create(.wconf)]]$  }

      // Now we create the new installation, copy the new Widget files
      // and link the executable (Points ③ and ④)
      {
         $r \Rightarrow - * iloc \Rightarrow IL * \delta^{ll} \mapsto v2Dir_{Pre} * \alpha^T \mapsto home_{Post}$ 
        {  $* \gamma^{\top/usr} \mapsto bin[B] * \beta^T \mapsto opt[T] * J \xrightarrow{F} PROG * K \xrightarrow{F} DAT$  }
      }
      r := mkdir( $\top/opt/widget$ );
      r := fileCopy(instLoc/ $\top/widgProg$ ,  $\top/opt/widget$ );
      r := fileCopy(instLoc/ $\top/widgData$ ,  $\top/opt/widget$ );
      r := link( $\top/opt/widget/widgProg$ ,  $\top/usr/bin/widget$ ); r := 0
      {
         $\exists J', K'. r \Rightarrow 0 * iloc \Rightarrow IL * instSrc_{Pre} * home_{Post}$ 
        {  $* \beta^T \mapsto opt[T + widget[widgProg : J' + widgData : K']]$  }
        {  $* \gamma^{\top/usr} \mapsto bin[B + widget : J'] * J' \xrightarrow{F} PROG * K' \xrightarrow{F} DAT$  }
      }
    }
    {
       $\exists R, J', K'. iloc \Rightarrow IL * r \Rightarrow R * errno \Rightarrow - \wedge (R = -1 \Rightarrow W)$ 
      {  $\wedge R = 0 \Rightarrow (instSrc_{Pre} * home_{Post} * opt_{Post}(J', K') * bin_{Post}(K') * v2Files_{Post}(J', K'))$  }
    }
  }
{  $Q_{ins}$  }

```

Fig. 6. Widget V2 software installer.

## 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] K. Arkoudas, K. Zee, V. Kuncak, and M. Rinard. Verifying a File System Implementation. In *LNCS: Formal Methods and Software Engineering*. 2004.
- [5] C. Calcagno, P. Gardner, and U. Zarfaty. Context logic and tree update. *SIGPLAN Not.*, 2005.
- [6] T. Dinsdale-Young, L. Birkedal, P. Gardner, M. Parkinson, and H. Yang. Views: compositional reasoning for concurrent programs. POPL, 2013.
- [7] T. Dinsdale-Young, M. Dodds, P. Gardner, M. Parkinson, and V. Vafeiadis. Concurrent abstract predicates. In *ECOOP*. 2010.
- [8] K. Fisher, N. Foster, D. Walker, and K. Q. Zhu. Forest: a language and toolkit for programming with filestores. ICFP, 2011.
- [9] L. Freitas, Z. Fu, and J. Woodcock. POSIX file store in Z/Eves: an experiment in the verified software repository. *Engineering of Complex Computer Systems, IEEE International Conference*, 2007.
- [10] L. Freitas, J. Woodcock, and A. Butterfield. POSIX and the verification grand challenge: A roadmap. In *ICECCS*, 2008.
- [11] P. Gardner, G. Ntzik, and A. Wright. Local Reasoning about File Systems: Technical Report. Technical report, Imperial College London, 2012 <http://www.doc.ic.ac.uk/~gn408/POSIXFS/>.
- [12] P. Gardner, A. Raad, M. Wheelhouse, and A. Wright. Abstract Local Reasoning for Concurrent Libraries, Submitted, 2013.
- [13] P. Gardner, G. Smith, M. Wheelhouse, and U. Zarfaty. Local Hoare reasoning about DOM. In *PODS*, 2008.
- [14] W. H. Hesselink and M. Lali. Formalizing a hierarchical file system. *REFINE*, 2009.
- [15] A. Hobor and J. Villard. The ramifications of sharing in data structures. POPL, 2013.
- [16] R. Joshi and G. J. Holzmann. A mini challenge: build a verifiable filesystem. *Form. Asp. Comput.*, 2007.
- [17] C. Morgan and B. Sufrin. Specification of the UNIX Filing System. *Software Engineering, IEEE Transactions on*, 1984.
- [18] G. Ntzik. *Local Reasoning about File Systems*. PhD thesis, 2014 (Expected).
- [19] M. Parkinson and G. Bierman. Separation logic and abstraction. POPL, 2005.
- [20] J. C. Reynolds. Separation logic: A logic for shared mutable data structures. In *LICS*, 2002.
- [21] G. Smith. *Local Reasoning about Web Programs*. PhD thesis, 2011.
- [22] A. Turon, D. Dreyer, and L. Birkedal. Unifying refinement and hoare-style reasoning in a logic for higher-order concurrency. ICFP, 2013.
- [23] A. Wright. *Structural Separation Logic*. PhD thesis, Imperial College London, 2013.

## Appendix

### Satisfaction Relation

**Definition 13 (Satisfaction Relation).** *The satisfaction relations,  $\models_P: \text{LEnv} \rightarrow \text{ASTATES} \times \text{ASRTS}$  and  $\models_\phi: \text{LEnv} \rightarrow \text{DIRS} \times \text{DIRASRTS}$ , for the abstract state and directory assertions of figure 4, parameterised by a logical environment, are defined as:*

$$\begin{aligned}
e, (fs, ph, \sigma) \models_P \alpha^E \mapsto \phi &\iff \exists d. fs((\alpha)_e) = ((E)_e, d) \wedge e, d \models_\phi \phi \\
e, (fs, ph, \sigma) \models_P X \xrightarrow{F} E &\iff fs((X)_e) = (E)_e \\
e, (fs, ph, \sigma) \models_P X \xrightarrow{OF} (I, E) &\iff \exists f. (X)_e = f \wedge f \in \text{OFADDRES} \wedge ph(f) = ((I)_e, (E)_e) \\
e, (fs, ph, \sigma) \models_P X \xrightarrow{DS} E &\iff \exists n. (X)_e = n \wedge n \in \text{DSADDRES} \wedge ph(n) = (E)_e \\
e, (fs, ph, \sigma) \models_P E \xrightarrow{H} E' &\iff \exists n. (E)_e = n \wedge n \in \mathbb{N}^+ \wedge ph(n) = (E')_e \\
e, (fs, ph, \sigma) \models_P \text{var} \Rightarrow E &\iff \sigma(\text{var}) = (E)_e \\
e, (fs, ph, \sigma) \models_P \text{Expr} \Rightarrow E &\iff [[\text{Expr}]]_\sigma = (E)_e \\
e, d \models_\phi \emptyset &\iff d = \emptyset \\
e, d \models_\phi E_1 : E_2 &\iff d = (E_1)_e : (E_2)_e \\
e, d \models_\phi E[\phi] &\iff \exists d'. d = (E)_e[d'] \wedge e, d' \models_\phi \phi \\
e, d \models_\phi \top[\phi] &\iff \exists d'. d = \top[d'] \wedge e, d' \models_\phi \phi \\
e, d \models_\phi E &\iff d = (E)_e \\
e, d \models_\phi \phi_1 + \phi_2 &\iff \exists d', d''. d = d' + d'' \wedge e, d' \models_\phi \phi_1 \wedge e, d'' \models_\phi \phi_2 \\
e, d \models_\phi \phi_1 \circ_\alpha \phi_2 &\iff \exists d', \mathbf{x}, d''. d = d' \circ_{\mathbf{x}} d'' \wedge e[\alpha \mapsto \mathbf{x}], d' \models_\phi \phi_1 \wedge e, d'' \models_\phi \phi_2 \\
e, d \models_\phi @E &\iff \text{resolve}((E)_e, d) \text{ defined}
\end{aligned}$$

### IO Command Axioms

Here we include axioms for IO commands on regular files in figure 7. These supplement the axioms of figure 5. Together, these provide the specification for commands when they succeed. Error specifications are discussed in the next section.

### Errors

So far, in figure 5 we have given specifications of when the POSIX commands succeed. In any other case they fail. In POSIX this generally means that an *error code* is returned to explain the cause of the failure. Programmers should examine this error code, and take remedial action if the command failed. Errors are also sometimes required to perform elementary tasks. For example, to see if a path  $p$  exists in the file system, one typically performs `stat(p)` and checks if an error is returned. No error indicates that the path resolves.

We model the following errors that are relevant to our POSIX fragment:

1. ENOENT: A component of a path does not resolve to an existing file or the path is an empty string. For example, this error occurs when resolving  $\top/a/b/c/d$ , but the directory  $\top/a/b$  does not exist in the file system.

2. ENOTDIR: A component of the path prefix is not a directory. For example, this error occurs when resolving  $\tau/a/b/c/d$ , but  $\tau/a/b$  identifies a file rather than a directory.
3. EEXIST: The path resolves to an existing file or directory. For example, the `link(existing, new)` command returns this error as it expects no file to exist at the `new` path.
4. ENOTEMPTY: The path argument names a directory that is not an empty directory. This error occurs when the path identifies a directory that has entries, but the command expected an empty directory. For example, the `rmdir` command returns this error.
5. EPERM: The intended operation is not supported by the POSIX implementation. This error occurs when you attempt to use a command, but the POSIX implementation running the program does not support using the command on the resource identified by the path. For example, as we do not allow hard links to directories, our `link` command returns this error when it is applied to directories.
6. EINVAL: The arguments supplied to the intended operation are of the wrong form. For example, in order to avoid creating cycles, `rename(old, new)` returns this error if `new` is a descendant `old`.

Commands communicate their success or failure via their return value. Most commands return 0 when they succeed — `open` and `opendir` return non integer values when they succeed — and `-1` otherwise. If `-1` is returned, a constant integer value representing the error is assigned to the global variable `errno`. The types of errors each command can fail with are given in table 1.

Command	Possible errors
<code>r := mkdir(path)</code>	ENOENT, ENOTDIR, EEXIST
<code>r := rmdir(path)</code>	ENOENT, ENOTDIR*, ENOTEMPTY
<code>r := link(to, from)</code>	ENOENT*, ENOTDIR, EPERM EEXIST (only for the <code>from</code> parameter)
<code>r := unlink(path)</code>	ENOENT*, ENOTDIR, EPERM
<code>r := rename(old, new)</code>	ENOENT*, ENOTDIR*, ENOTEMPTY, EISDIR, EINVAL
<code>t := stat(path)</code>	ENOENT, ENOTDIR
<code>fd := open(path, flags)</code>	ENOENT*, ENOTDIR, EISDIR
<code>o := lseek(fd, off, wh)</code>	No errors
<code>sz := write(fd, str)</code>	No errors
<code>str := read(fd, size)</code>	No errors
<code>close(fd)</code>	No errors
<code>h := opendir(path)</code>	ENOENT, ENOTDIR
<code>s := readdir(h)</code>	No errors
<code>closedir(h)</code>	No errors

**Table 1.** Possible errors that each command can return

However, we still consider certain types of command failure to be program faults. Calling `wrire(fd)` with an `fd` that does not correspond to a valid file descriptor still faults in our specifications, despite POSIX considering this an EBADF error. We make this distinction based on an intuition of *ownership*. File descriptors are owned and managed by the programs that create them. Writing a program that could encounter EBADF is always going to be a programmer mistake. We can use our program logic to prove that a program does not fault: it would never cause EBADF to occur. Error checking for EBADF could then be elided as unnecessary.

On the other hand, the file system is a shared resource, not owned by any one program. Even in our non-concurrent system, there is nothing to stop someone mutating the file system against our expectations before our program runs. Thus, it is sensible and expected that programmers use errors as a form of “defensive programming”, to cater for unanticipated external changes to the file system structure. We can use our reasoning to determine if a program detects all possible error cases, and more importantly how it handles them.

For each command, and each error it can produce, we provide a small axiom that describes the states on which the command will produce the error. Each additional small axiom describes all the states that cause a specific error code.

The properties of each error code are independent of its use by any one command. For example, failing to resolve a path correctly is the same whether the command that failed is `open` or `mkdir`. As such, in figure 8 we provide predicates that describe the states on which resolving a given path would result in each error. We name the predicates identically to the error they represent. There are two variants of ENOENT:  $\text{ENOENT}_*(\text{path})$  is satisfied when `path` or any prefix of `path` do not resolve, and  $\text{ENOENT}(\text{path})$  is satisfied simply when any prefix of `path` does not resolve. Similarly there are two variants of ENOTDIR.

With these error predicates, we can axiomatize commands that may return errors. Here, we consider the `stat` command:

$$\begin{array}{cc}
 \left\{ \begin{array}{l} \text{path} \Rightarrow P/A \wedge \alpha^P \mapsto A : I \\ \text{t}, \text{r} := \text{stat}(\text{path}) \\ \text{t} \Rightarrow F \wedge \text{r} \Rightarrow 0 \wedge \alpha^P \mapsto A : I \end{array} \right\} & \left\{ \begin{array}{l} \text{path} \Rightarrow P/A \wedge \alpha^P \mapsto A[\beta] \\ \text{t}, \text{r} := \text{stat}(\text{path}) \\ \text{t} \Rightarrow D \wedge \text{r} \Rightarrow 0 \wedge \alpha^P \mapsto A[\beta] \end{array} \right\} \\
 \\
 \left\{ \begin{array}{l} S \wedge \text{ENOENT}_*(\text{path}) \\ \text{t}, \text{r} := \text{stat}(\text{path}) \\ \text{r} \Rightarrow -1 \wedge \text{errno} \Rightarrow \text{ENOENT} \wedge S \end{array} \right\} & \left\{ \begin{array}{l} S \wedge \text{ENOTDIR}(\text{path}) \\ \text{t}, \text{r} := \text{stat}(\text{path}) \\ \text{r} \Rightarrow -1 \wedge \text{errno} \Rightarrow \text{ENOTDIR} \wedge S \end{array} \right\}
 \end{array}$$

Error cases for the rest of the commands of table 1 are defined similarly.

## Example Proofs

Here we give the implementations and proofs for the removal of a non-empty directory and file copying.

`rmdirRec`

Recall the intuitive specification of `rmdirRec` from the software installer example.

$$\left\{ \begin{array}{l} \text{path} \Rightarrow P/A \wedge r \Rightarrow - * \mathcal{E} * \alpha^P \mapsto A[\text{complete}] \\ r := \text{rmdirRec}(\text{path}) \\ \{ r \Rightarrow 0 * \mathcal{E} * \alpha^P \mapsto \emptyset \} \end{array} \right\}$$

Our implementation of this command and a proof sketch of the specification is given in figure 9. Notice that at each loop iteration, we use abstract allocation to address the entry we are going delete. When we remove, e.g. with `unlink`, we implicitly apply the Frame Rule to isolate only that resource. Specifications of the cases where `rmdirRec` fails with an error can be easily derived.

`fileCopy`

Finally, following is the specification of `fileCopy`.

$$\left\{ \begin{array}{l} \text{source} \Rightarrow P/A \wedge \text{target} \Rightarrow P'/D \wedge r \Rightarrow - * \mathcal{E} \\ * \alpha^P \mapsto A : I * \beta^{P'} \mapsto D[C \wedge \text{can\_create}(A)] * I \xrightarrow{F} S_D \\ r := \text{fileCopy}(\text{source}, \text{target}) \\ \{ \exists I'. r \Rightarrow 0 * \mathcal{E} * \alpha^P \mapsto A : I * \beta^{P'} \mapsto D[C + A : I'] * I \xrightarrow{F} S_D * I' \xrightarrow{F} S_D \} \end{array} \right\}$$

Because file sizes can exceed available process memory, they are copied chunk by chunk. In figure 10 we show an implementation of a chunked file copy command and prove it's correct with respect to the above specification.

## Soundness

We reason about programs manipulating abstract program states (definition 8) that contain instrumentation in the form of abstract addresses and path promises. Operationally, commands operate on *machine states* where the instrumentation is erased.

**Definition 14 (Machine States).** Machine states,  $s \in \text{STATES}$ , are those abstract program states of definition 8 where the file system heap is complete.

$$\text{STATES} \triangleq \{ s \in \text{ASTATES} \mid s \downarrow_1 \text{ is complete} \}$$

We relate abstract states to machine states via *reification*. Reification accounts for the (possibly) partial file system by using completions. Process heaps are not instrumented, so are reused unchanged.

**Definition 15 (Reification).** The reification function,  $[\cdot] : \text{ASTATES} \rightarrow \mathcal{P}(\text{STATES})$ , is defined by:

$$[as] = \{ (fs, as \downarrow_2, as \downarrow_3) \mid fs \in \text{completion}(as \downarrow_1) \}$$

where  $\text{completion}(fs) \triangleq \{ fs_c \in \text{FS} \mid fs' \in \text{FS} \wedge fs_c = fs \sqcup fs' \wedge fs_c \text{ is complete} \}$  where  $\downarrow_i$  denotes the *i*th projection. We lift reification to sets of instrumented states pointwise.



The operational semantics of our programming language are directly derived from the operational semantics of the Views framework.

**Definition 16 (Program Evaluation).** *The evaluation relation,  $\Downarrow$ , relates programs and their initial states to either a terminal state or a distinguished fault state  $\Downarrow$ , where  $\Downarrow \notin \text{VALUES}$ . Command/state pairs not in the relation are deemed to be divergent.*

$$\Downarrow \subset (\text{COMM} \times \text{STATES}) \times (\text{STATES} \uplus \{\Downarrow\})$$

We derive the evaluation relation from the operational semantics framework in Views.

We use the standard partial correctness fault avoiding interpretation of Hoare triples.

**Definition 17 (Hoare Triple Interpretation).**

$$\models \{P\} \mathbb{C} \{Q\} \iff \forall e \in \text{LENV}, s \in \llbracket P \rrbracket_e. \mathbb{C}, s \Downarrow \Downarrow \wedge (\exists s'. \mathbb{C}, s \Downarrow s' \Rightarrow s' \in \llbracket Q \rrbracket_e)$$

**Theorem 1 (Soundness).** *If  $\vdash \{P\} \mathbb{C} \{Q\}$ , then  $\models \{P\} \mathbb{C} \{Q\}$ .*

We prove theorem 1 via the Views framework of [6]. Readers interested in the details of Views are referred to [6], which gives a detailed account of the parameters and properties required to prove the soundness theorem. Here we simply instantiate the parameters of the Views framework and prove the required properties.

**Definition 18 (File Systems Atomic Commands).** *The atomic commands are those of definition 9.*

**Definition 19 (File Systems View Monoid).** *The File Systems View Monoid is:*

$$\text{FSVIEW} \triangleq (\mathcal{P}(\text{ASTATES}), \circ, u)$$

where

1.  $(ifs_1, ph_1, \sigma_1) \circ' (ifs_2, ph_2, \sigma_2) = (ifs_1 \sqcup ifs_2, ph_1 \sqcup ph_2, \sigma_1 \sqcup \sigma_2)$
2.  $\circ$  is the lift of  $\circ'$  to sets
3.  $u = \{(\emptyset, \emptyset, \emptyset)\}$ .

**Definition 20 (File Systems Axiomatisation).** *The File System Axiomatisation is given in figures 5 and 7.*

**Definition 21 (File System Machine States).** *The machine states are the program states of definition 14.*

**Definition 22 (File System Interpretation of Atomic Commands).** *The File System interpretation of Atomic Commands are given in figures 11 and 12.*

**Definition 23 (File Systems View Reification).** *The File System View Reification is the lift of the Instrumented State Reification of definition 15 to sets.*

**Definition 24 (File Systems View Disjunction).** *The File Systems View Disjunction function is set union,  $\cup$ .*

We now prove each of the required properties. We begin with some additional lemmas. The Views framework assumes that the  $*$  of Views is the separating conjunction. As we work with logical environments, we have “baked”  $*$  into the assertion language. Assertions are then transformed into views by interpretation through a logical environment. However, the  $*$  of our assertion language is essentially identical to the  $*$  of our Views monoid.

**Lemma 1 ( $*$  equivalence).** *Let  $*$  be the separating conjunction of our assertion language, figure . Recall that the separating conjunction of our views monoid is  $\circ$ . For all  $p, q \in \mathcal{P}(\text{ASTATES}), e \in \text{LENV}$ , it is the case that*

$$\llbracket p * q \rrbracket_e = \llbracket p \rrbracket_e \circ \llbracket q \rrbracket_e$$

*Proof.* Calculating using definitions 13 and 19

$$\begin{aligned} \llbracket p * q \rrbracket_e &= \{as \in \text{ASTATES} \mid e, as \models p * q\} \\ &= \{(ifs_1 \bullet ifs_2, ph_1 \bullet ph_2, \sigma) \mid ifs_1, ph_1, \sigma \models p, ifs_2, ph_2, \sigma \models q\} \\ &= \{(ifs_1, ph_1, \sigma) \circ' (ifs_2, ph_2, \sigma) \mid ifs_1, ph_1, \sigma \models p, ifs_2, ph_2, \sigma \models q\} \\ &= \{P \circ' Q \mid P \in \llbracket p \rrbracket_e, Q \in \llbracket q \rrbracket_e\} \\ &= \llbracket p \rrbracket_e \circ \llbracket q \rrbracket_e \end{aligned}$$

**Lemma 2 (Atomic Soundness of Commands).**

*Proof.* We will cover only characteristic examples of commands in our fragment. Atomic soundness for the remaining command axioms is proven similarly.

First, consider the `mkdir` command axiom in figure 5. By the satisfaction relation (definition 13) we have:

$$\begin{aligned} \llbracket \text{path} \Rightarrow P/B/A \wedge \alpha^P \mapsto B[C \wedge \neg \text{top\_entry}(A) \wedge \neg \text{top\_address}] * r \rrbracket_e &= \\ &= \left\{ \begin{array}{l} (ifs, ph, \sigma) \in \text{ASTATES} \mid \\ \exists ifs_1, ifs_2, ph_1, ph_2. ifs = ifs_1 \bullet ifs_2 \wedge ph = ph_1 \bullet ph_2 \\ \wedge \llbracket \text{path} \rrbracket_\sigma = (\llbracket P \rrbracket_e / (\llbracket B \rrbracket_e) / (\llbracket A \rrbracket_e) \\ \wedge \exists id. ifs_1(\llbracket \alpha \rrbracket_e) = ((\llbracket P \rrbracket_e), (\llbracket B \rrbracket_e[id]) \vee ifs_1(\mathcal{F})) = \top[id] \wedge id = (\llbracket C \rrbracket_e) \\ \wedge \neg(\exists id', id'', \iota. id = id' + ((\llbracket A \rrbracket_e[id''] \vee (\llbracket A \rrbracket_e : \iota)) \\ \wedge \neg(\exists \mathbf{x}, id'. id = id' + \mathbf{x}) \\ \wedge e, (ifs_2, ph_2, \sigma) \models r \\ = \text{ass} \end{array} \right\} \end{aligned}$$

the set of instrumented pre-states. By reification, the pre-states are:

$$\begin{aligned} \llbracket \text{ass} \rrbracket &= \\ &= \left\{ \begin{array}{l} (fs, ph, \sigma) \mid fs \in \text{completion}(as \downarrow_1) \wedge ph = as \downarrow_2 \wedge \sigma = as \downarrow_3 \\ \wedge \llbracket \text{path} \rrbracket_\sigma = (\llbracket P \rrbracket_e / (\llbracket B \rrbracket_e) / (\llbracket A \rrbracket_e) \\ \wedge \exists id, id', \mathbf{x}. fs(\mathcal{F}) = id \circ_{\mathbf{x}} (\llbracket B \rrbracket_e[id']) \\ \wedge \text{resolve}(\llbracket P \rrbracket_e / \mathbf{x}, id) = \mathbf{x} \wedge \text{noent}(id', (\llbracket A \rrbracket_e) \end{array} \right\} \end{aligned}$$

Then, by the command interpretation

$$\begin{aligned} & \llbracket \text{mkdir}(\text{path}) \rrbracket (\llbracket \text{ass} \rrbracket) \\ & = \\ & \left\{ \begin{array}{l} (fs, \text{ph}, \sigma) \mid \text{ph} = \text{as} \downarrow_2 \wedge \sigma = \text{as} \downarrow_3 \\ \wedge \llbracket \text{path} \rrbracket_\sigma = (\text{P})_e / (\text{B})_e / (\text{A})_e \\ \wedge \exists id, id', \mathbf{x}. fs(\mathcal{F}) = id \circ_{\mathbf{x}} (\text{B})_e[id' + (\text{A})_e[\emptyset]] \\ \wedge \text{resolve}((\text{P})_e/\mathbf{x}, id) = \mathbf{x} \end{array} \right\} \end{aligned}$$

we obtain the post-states.

By the satisfaction relation we get the instrumented post-states from the command axiom's postcondition:

$$\begin{aligned} & \llbracket \alpha^P \mapsto \text{B}[C + \text{A}[\emptyset]] * r \rrbracket_e \\ & = \\ & \left\{ \begin{array}{l} (ifs, \text{ph}, \sigma) \in \text{ASTATES} \mid \\ \exists ifs_1, ifs_2, \text{ph}_1, \text{ph}_2. ifs = ifs_1 \bullet ifs_2 \wedge \text{ph} = \text{ph}_1 \bullet \text{ph}_2 \\ \wedge \llbracket \text{path} \rrbracket_\sigma = (\text{P})_e / (\text{B})_e / (\text{A})_e \\ \wedge (\exists id. ifs_1((\alpha)_e) = ((\text{P})_e, (\text{B})_e[id + (\text{A})_e[\emptyset]]) \\ \wedge e, (ifs_2, \text{ph}_2, \sigma) \models' r \\ = \text{ass}' \end{array} \right\} \end{aligned}$$

where facts about the variable store are derived from logical variable immutability. The instrumented pre and post-states have the same completing pre-instrumented file systems  $\text{pifs}_c$  in their completions. By this, it is easy to see that:

$$\llbracket \text{mkdir}(\text{path}) \rrbracket (\llbracket \text{ass} \rrbracket) = \llbracket \text{ass}' \rrbracket$$

Structural and state commands are proven in a similar way.

Now consider the third axiom of the **open** command. By the satisfaction relation we have:

$$\begin{aligned} & \llbracket \text{path} \Rightarrow \text{P}/\text{A} \wedge \text{flags} \Rightarrow \text{0\_TRUNC} \wedge \alpha^P \mapsto \text{A} : \text{I} * \text{I} \stackrel{\text{F}}{\mapsto} \text{S} * \text{emp} * r \rrbracket_e \\ & = \\ & \left\{ \begin{array}{l} (ifs, \text{ph}, \sigma) \in \text{ASTATES} \mid \\ \exists ifs_1, ifs_2, \text{ph}_1, \text{ph}_2. ifs = ifs_1 \bullet ifs_2 \wedge \text{ph} = \text{ph}_1 \bullet \text{ph}_2 \\ \wedge \llbracket \text{path} \rrbracket_\sigma = (\text{P})_e / (\text{A})_e \wedge \llbracket \text{flags} \rrbracket_\sigma = \text{0\_TRUNC} \\ \wedge ifs_1((\alpha)_e) = ((\text{P})_e, (\text{A})_e : (\text{I})_e) \wedge ifs_2((\text{I})_e) = (\text{S})_e \\ \wedge e, (ifs_2, \text{ph}_2, \sigma) \models' r \\ = \text{ass} \end{array} \right\} \end{aligned}$$

the set of instrumented pre-states. We reify these to pre-states:

$$\begin{aligned} & \llbracket \text{ass} \rrbracket \\ & = \\ & \left\{ \begin{array}{l} (fs, \text{ph}, \sigma) \mid fs \in \text{completion}(\text{as} \downarrow_1) \wedge \text{ph} = \text{as} \downarrow_2 \wedge \sigma = \text{as} \downarrow_3 \\ \wedge \llbracket \text{path} \rrbracket_\sigma = (\text{P})_e / (\text{A})_e \wedge \llbracket \text{flags} \rrbracket_\sigma = \text{0\_TRUNC} \\ \wedge \exists id. fs(\mathcal{F}) = id \circ_{\mathbf{x}} (\text{A})_e : (\text{I})_e \\ \wedge \text{resolve}((\text{P})_e/\mathbf{x}, id) = \mathbf{x} \wedge fs(\text{I}) = (\text{S})_e \end{array} \right\} \end{aligned}$$

We use the pre-states as input to the command interpretation

$$\begin{aligned} & \llbracket \text{fd} := \text{open}(\text{path}, \text{flags}) \rrbracket (\llbracket \text{ass} \rrbracket) \\ & = \\ & \left\{ \begin{array}{l} (fs', \text{ph}', \sigma') \mid \text{ph} = \text{as} \downarrow_2 \wedge \sigma = \text{as} \downarrow_3 \\ \wedge \sigma' = \sigma[\text{fd} \mapsto f] \wedge \text{ph}' = \text{ph}[f \mapsto ((\text{I})_e, 0)] \\ \wedge fs' = fs[\text{I} \mapsto \epsilon] \wedge f \notin \text{dom}(\text{ph}) \end{array} \right\} \end{aligned}$$

to obtain the post-states. Again, by the satisfaction relation we get the instrumented post-states.

$$\left[ \left[ \exists \text{FD}. \left( \begin{array}{l} \text{fd} \models \text{FD} \wedge \alpha^P \mapsto A : I \\ \star I \xrightarrow{F} \epsilon \star \text{FD} \xrightarrow{\text{OF}} (I, 0) \end{array} \right) \star r \right] \right]_e$$

$$= \left\{ \begin{array}{l} (\text{ifs}, \text{ph}, \sigma) \in \text{ASTATES} \mid \\ \exists \text{ifs}_1, \text{ifs}_2, \text{ph}_1, \text{ph}_2. \text{ifs} = \text{ifs}_1 \bullet \text{ifs}_2 \wedge \text{ph} = \text{ph}_1 \bullet \text{ph}_2 \\ \wedge \llbracket \text{fd} \rrbracket_\sigma = (\text{FD})_e \\ \wedge \text{ifs}_1(\langle \alpha \rangle_e) = (\langle P \rangle_e, \langle A \rangle_e : \langle I \rangle_e) \wedge \text{ifs}_1(\langle I \rangle_e) = \epsilon \\ \wedge \text{ph}_1(\langle \text{FD} \rangle_e) = (\langle I \rangle_e, 0) \wedge e, (\text{ifs}_2, \text{ph}_2, \sigma) \models' r \end{array} \right\}$$

$$= \text{ass}'$$

Again, the completions of the instrumented pre and post-states are the same and thus it is easy to see that:

$$\llbracket \text{fd} := \text{open}(\text{path}, \text{flags}) \rrbracket(\llbracket \text{ass} \rrbracket) = \llbracket \text{ass}' \rrbracket$$

**Lemma 3 (Entailment Properties).** *File System Entailment satisfies Entailment Locality.*

*Proof.* Let  $p \models q$ . By the definition of  $\models$ , we have that for all  $e \in \text{LENV}$ ,  $\llbracket p \rrbracket_e \subseteq \llbracket q \rrbracket_e$ . We must show that  $p \preceq q$ . Recall the definition of  $\preceq$ :  $p \preceq q \iff \forall r \in \mathcal{P}(\text{ASTATES}), e \in \text{LENV}, \llbracket p \star r \rrbracket_e \subseteq \llbracket q \star r \rrbracket_e$ . Thus, the lemma requires us to show that if  $\llbracket p \rrbracket_e \subseteq \llbracket q \rrbracket_e$ , then  $\llbracket p \star r \rrbracket_e \subseteq \llbracket q \star r \rrbracket_e$ , for all choices of logical environments and views  $r$ . But, considering lemma 1, this follows immediately as entailment is just subset inclusion.

**Lemma 4 (Disjunction Properties).** *File System Disjunction has Join Distributivity and is a Join morphism.*

*Proof.* Our Disjunction function is set union. Therefore, for Join Distributivity, we must show  $p, q_i \in \mathcal{P}(\text{ASTATES})$  and index set  $I$ ,  $p \star \bigcup_{i \in I} \{q_i\} = \bigcup_{i \in I} \{p \star q_i\}$ . For the Join Morphism property, we need to show  $\llbracket \bigcup_{i \in I} \{p_i\} \rrbracket_e = \bigcup_{i \in I} \llbracket p_i \rrbracket_e$ . Both properties follow from the set lifting definition of reification.

Finally, the soundness of the abstract allocation/deallocation axioms (definition 11) is justified by the fact that they only change the instrumentation. The reifications of pre- and post-condition are the same. With these results, we can conclude our reasoning is sound.

*Proof (Proof of Soundness Theorem).* By lemmas 2, 3, 4 according to the Views framework, the partial correctness result follows. The fault avoiding result follows by the fact that  $\not\downarrow \notin \text{STATES}$ .

$$\begin{aligned}
& \{ \text{path} \Rightarrow P/B/A \wedge \text{fd} \Rightarrow - * \mathcal{E} * \alpha^P \mapsto B[C \wedge \text{can\_create}(A)] \} \\
& \quad \text{fd} := \text{open}(\text{path}, \text{flags}) \\
& \quad \left\{ \exists \text{FD}, I. \left( \begin{array}{l} \text{fd} \Rightarrow \text{FD} * \mathcal{E} * \alpha^P \mapsto B[C + A : I] \\ * I \xrightarrow{F} \emptyset_b * \text{FD} \xrightarrow{\text{OF}} (I, 0) \end{array} \right) \right\} \\
& \{ \text{path} \Rightarrow P/A \wedge \text{flags} \Rightarrow 0 \wedge \text{fd} \Rightarrow - * \mathcal{E} * \alpha^P \mapsto A : I * I \xrightarrow{F} S \} \\
& \quad \text{fd} := \text{open}(\text{path}, \text{flags}) \\
& \quad \left\{ \exists \text{FD}. \left( \begin{array}{l} \text{fd} \Rightarrow \text{FD} * \mathcal{E} * \alpha^P \mapsto A : I \\ * I \xrightarrow{F} S * \text{FD} \xrightarrow{\text{OF}} (I, 0) \end{array} \right) \right\} \\
& \{ \text{path} \Rightarrow P/A \wedge \text{flags} \Rightarrow 0.\text{TRUNC} \wedge \text{fd} \Rightarrow - * \mathcal{E} * \alpha^P \mapsto A : I * I \xrightarrow{F} S \} \\
& \quad \text{fd} := \text{open}(\text{path}, \text{flags}) \\
& \quad \left\{ \exists \text{FD}. \left( \begin{array}{l} \text{fd} \Rightarrow \text{FD} * \mathcal{E} * \alpha^P \mapsto A : I \\ * I \xrightarrow{F} \emptyset_b * \text{FD} \xrightarrow{\text{OF}} (I, 0) \end{array} \right) \right\} \\
& \quad \left\{ \begin{array}{l} \text{fd} \Rightarrow \text{FD} * \text{FD} \xrightarrow{\text{OF}} (I, \text{OFFSET}) \\ \text{close}(\text{fd}) \\ \{ \text{emp} \} \end{array} \right\} \\
& \{ \text{fd} \Rightarrow \text{FD} \wedge \text{to} \Rightarrow O' \wedge \text{wh} \Rightarrow \text{SEEK\_SET} \wedge O' \geq 0 * \text{off} \Rightarrow - * \mathcal{E} * \text{FD} \xrightarrow{\text{OF}} (I, O) \} \\
& \quad \text{off} := \text{lseek}(\text{fd}, \text{to}, \text{wh}) \\
& \quad \left\{ \text{off} \Rightarrow O' * \text{fd} \Rightarrow \text{FD} * \mathcal{E} * \text{FD} \xrightarrow{\text{OF}} (I, O') \right\} \\
& \{ \text{fd} \Rightarrow \text{FD} \wedge \text{to} \Rightarrow O' \wedge \text{wh} \Rightarrow \text{SEEK\_CUR} \wedge (O + O' \geq 0) * \text{off} \Rightarrow - * \mathcal{E} * \text{FD} \xrightarrow{\text{OF}} (I, O) \} \\
& \quad \text{off} := \text{lseek}(\text{fd}, \text{to}, \text{wh}) \\
& \quad \left\{ \text{off} \Rightarrow O + O' * \text{fd} \Rightarrow \text{FD} * \mathcal{E} * \text{FD} \xrightarrow{\text{OF}} (I, O + O') \right\} \\
& \left\{ \begin{array}{l} \text{fd} \Rightarrow \text{FD} \wedge \text{to} \Rightarrow O' \wedge \text{wh} \Rightarrow \text{SEEK\_END} \wedge |S| + O' \geq 0 \\ \text{off} \Rightarrow - * \mathcal{E} * I \xrightarrow{F} S * \text{FD} \xrightarrow{\text{OF}} (I, O) \\ \text{off} := \text{lseek}(\text{fd}, \text{to}, \text{wh}) \end{array} \right\} \\
& \left\{ \text{off} \Rightarrow |S| + O' * \text{fd} \Rightarrow \text{FD} * \mathcal{E} * I \xrightarrow{F} S * \text{FD} \xrightarrow{\text{OF}} (I, |S| + O') \right\} \\
& \left\{ \begin{array}{l} \text{fd} \Rightarrow \text{FD} \wedge \text{buf} \Rightarrow S \wedge |S'| = O \wedge |S''| = |S| \\ * \text{sz} \Rightarrow - * \mathcal{E} * I \xrightarrow{F} S' \cdot S'' \cdot S''' * \text{FD} \xrightarrow{\text{OF}} (I, O) \\ \text{sz} := \text{write}(\text{fd}, \text{buf}) \end{array} \right\} \\
& \left\{ \text{sz} \Rightarrow |S| * \text{fd} \Rightarrow \text{FD} * \mathcal{E} * I \xrightarrow{F} S' \cdot S \cdot S''' * \text{FD} \xrightarrow{\text{OF}} (I, O + |S|) \right\} \\
& \left\{ \begin{array}{l} \text{fd} \Rightarrow \text{FD} \wedge \text{buf} \Rightarrow S \wedge |S'| = O \wedge |S''| < |S| \\ * \text{sz} \Rightarrow - * \mathcal{E} * I \xrightarrow{F} S' \cdot S'' * \text{FD} \xrightarrow{\text{OF}} (I, O) \\ \text{sz} := \text{write}(\text{fd}, \text{buf}) \end{array} \right\} \\
& \left\{ \text{sz} \Rightarrow |S| * \text{fd} \Rightarrow \text{FD} * \mathcal{E} * I \xrightarrow{F} S' \cdot S * \text{FD} \xrightarrow{\text{OF}} (I, O + |S|) \right\} \\
& \{ \text{fd} \Rightarrow \text{FD} \wedge \text{buf} \Rightarrow S \wedge O \geq |S'| \wedge I \xrightarrow{F} S' * \text{sz} \Rightarrow - * \mathcal{E} * \text{FD} \xrightarrow{\text{OF}} (I, O) \} \\
& \quad \text{sz} := \text{write}(\text{fd}, \text{buf}) \\
& \left\{ \text{sz} \Rightarrow |S| * \text{fd} \Rightarrow \text{FD} * \mathcal{E} * I \xrightarrow{F} S' \underset{O-|S'|}{\circlearrowleft} 0 \cdot S * \text{FD} \xrightarrow{\text{OF}} (I, O + |S|) \right\} \\
& \left\{ \begin{array}{l} \text{fd} \Rightarrow \text{FD} \wedge \text{sz} \Rightarrow \text{LEN} \wedge |S| = O \wedge |S'| = \text{LEN} \\ * \text{buf} \Rightarrow - * \mathcal{E} * I \xrightarrow{F} S \cdot S' \cdot S'' * \text{FD} \xrightarrow{\text{OF}} (I, O) \end{array} \right\} \\
& \quad \text{buf} := \text{read}(\text{fd}, \text{sz}) \\
& \left\{ \text{buf} \Rightarrow S' * \text{fd} \Rightarrow \text{FD} * \mathcal{E} * I \xrightarrow{F} S \cdot S' \cdot S'' * \text{FD} \xrightarrow{\text{OF}} (I, O + \text{LEN}) \right\} \\
& \left\{ \begin{array}{l} \text{fd} \Rightarrow \text{FD} \wedge \text{sz} \Rightarrow \text{LEN} \wedge |S| = O \wedge \text{LEN} > |S'| \\ * \text{buf} \Rightarrow - * \mathcal{E} * I \xrightarrow{F} S \cdot S' * \text{FD} \xrightarrow{\text{OF}} (I, O) \end{array} \right\} \\
& \quad \text{buf} := \text{read}(\text{fd}, \text{sz}) \\
& \left\{ \text{buf} \Rightarrow S' * \text{fd} \Rightarrow \text{FD} * \mathcal{E} * I \xrightarrow{F} S \cdot S' * \text{FD} \xrightarrow{\text{OF}} (I, O + |S'|) \right\} \\
& \{ \text{fd} \Rightarrow \text{FD} \wedge \text{sz} \Rightarrow \text{LEN} \wedge O > |S| * \text{buf} \Rightarrow - * \mathcal{E} * I \xrightarrow{F} S * \text{FD} \xrightarrow{\text{OF}} (I, O) \} \\
& \quad \text{buf} := \text{read}(\text{fd}, \text{sz}) \\
& \left\{ \text{buf} \Rightarrow \emptyset_b * \text{fd} \Rightarrow \text{FD} * \mathcal{E} * I \xrightarrow{F} S * \text{FD} \xrightarrow{\text{OF}} (I, O) \right\}
\end{aligned}$$

Fig. 7. Axiomatic specifications for IO commands.

$$\begin{aligned}
\text{ENOENT}_*(\text{path}) &\triangleq (\text{path} \Rightarrow \epsilon) \vee \left( \begin{array}{c} \exists P, A, B, P'. \text{path} \Rightarrow P/A/B/P' \\ \wedge \alpha^P \mapsto A[\mathbf{true} \wedge \neg \text{top\_entry}(B) \wedge \neg \text{top\_address}] \end{array} \right) \\
\text{ENOENT}(\text{path}) &\triangleq (\text{path} \Rightarrow \epsilon) \vee \left( \begin{array}{c} \exists P, A, B, P'. \text{path} \Rightarrow P/A/B/P' \wedge P' \neq \epsilon \\ \wedge \alpha^P \mapsto A[\mathbf{true} \wedge \neg \text{top\_entry}(B) \wedge \neg \text{top\_address}] \end{array} \right) \\
\text{ENOTDIR}_*(\text{path}) &\triangleq \exists P, A, P'. \text{path} \Rightarrow P/A/P' \wedge \alpha^P \mapsto \exists I. A : I \\
\text{ENOTDIR}(\text{path}) &\triangleq \exists P, A, P'. \text{path} \Rightarrow P/A/P' \wedge \alpha^P \mapsto \exists I. A : I \\
\text{EEXIST}(\text{path}) &\triangleq \exists P, A. \text{path} \Rightarrow P/A \wedge \alpha^P \mapsto \text{entry}(A) \\
\text{ENOTEMPTY}(\text{path}) &\triangleq \exists P, A. \text{path} \Rightarrow P \wedge \alpha^P \mapsto \text{entry}(A) \\
\text{EPERM}(\text{path}) &\triangleq \exists P, A. \text{path} \Rightarrow P/A \wedge \alpha^P \mapsto A[\mathbf{true}] \\
\text{EISDIR}(\text{path}) &\triangleq \text{EPERM}(\text{path}) \\
\text{EINVAL}(\text{old}, \text{new}) &\triangleq \exists P, P', A. \text{old} \Rightarrow P/A \wedge \text{new} \Rightarrow P' \wedge \alpha^P \mapsto (A[C] \wedge \exists P''. @A/P''/\beta) * \beta^{P'} \mapsto \mathbf{true}
\end{aligned}$$

**Fig. 8.** Definitions of the error predicates used in command specifications.

```

{ path  $\Rightarrow$  P/A  $\wedge$  r  $\Rightarrow$  - *  $\mathcal{E}$  *  $\alpha^P \mapsto A[C \wedge complete]$  }
r := rmdirRec(path)  $\triangleq$  local dirH, dirEnt, t {
  dirH := opendir(path);
  if dirH  $\neq$  null
    {  $\exists H, \alpha^P \mapsto A[C \wedge names(A) \wedge complete] * H \xrightarrow{DS} A$  }
    dirEnt := readdir(dirH);
    {  $\exists H, A, B, C', \alpha^P \mapsto A \left[ \begin{array}{c} entry(B) + (C' \wedge names(A)) \\ \vee \\ B = 0 \wedge A = \emptyset \wedge \emptyset \end{array} \right] \wedge complete$  }
    while dirEnt  $\neq$   $\epsilon$ 
      {  $\exists H, A, B, C', \beta \cdot \alpha^P \mapsto A [\beta + (C' \wedge names(A) \wedge complete)]$  }
      *  $\beta^{P/A} \mapsto (entry(B) \wedge complete) * H \xrightarrow{DS} A$ 
      t := stat(path/dirEnt);
      {  $\exists H, A, B, C', \beta \cdot \alpha^P \mapsto A [\beta + (C' \wedge names(A) \wedge complete)]$  }
      *  $\beta^{P/A} \mapsto (entry(B) \wedge complete) * H \xrightarrow{DS} A$ 
      if t = F
        unlink(path/dirEnt);
      else
        {  $\exists H, A, B, C', \beta \cdot \alpha^P \mapsto A [\beta + (C' \wedge names(A) \wedge complete)]$  }
        *  $\beta^{P/A} \mapsto \exists C''. B[C'' \wedge complete] * H \xrightarrow{DS} A$ 
        rmdirRec(path/dirEnt);
        {  $\exists H, A, B, C', \beta \cdot \alpha^P \mapsto A [\beta + (C' \wedge names(A) \wedge complete)]$  }
        *  $\beta^{P/A} \mapsto \emptyset * H \xrightarrow{DS} A$ 
        dirEnt := readdir(dirH);
        {  $\exists H, path \Rightarrow P/A \wedge \mathcal{E} * r \Rightarrow - * dirH \Rightarrow H * dirEnt \Rightarrow 0 \wedge \alpha^P \mapsto A[\emptyset] * H \xrightarrow{DS} \emptyset$  }
        closedir(dirH); rmdir(path); r := 0
      else r := dirH; }
    { r  $\Rightarrow 0 \wedge \mathcal{E} * \alpha^P \mapsto \emptyset$  }

```

**Fig. 9.** Proof that `rmdirRec` meets its specification (we sketch only one branch of the if for space). We use the rule of disjunction to merge the two possible results for `stat`.

```

{ source ⇒ P/A ∧ target ⇒ P'/D ∧  $\mathcal{E} * r \Rightarrow - * \alpha^P \mapsto A : I$  }
{ *  $\beta^{P'}$   $\mapsto D[C \wedge \sim\text{top\_entry}(A) \wedge \sim\text{top\_address}] * I \xrightarrow{F} SD$  }
r := fileCopy(source, target) ≐
local s, sourceH, targetH, buffer {
  sourceH := open(source, 0);
  if sourceH ≠ -1
    targetH := open(target/base(source), 0);
  if targetH ≠ -1
    buffer := read(sourceH, 4096); // Copy in chunks of 4096 chars
    {
       $\exists SH, TH, I', SD', SB, SD''.$ 
      sourceH ⇒ SH * targetH ⇒ TH * buffer ⇒ SB
      ∧ SD = SD' · SB · SD'' ∧ ((|SB| < 4096 ∧ SD'' =  $\epsilon$ ) ∨ |SB| = 4096)
      ∧  $\alpha^P \mapsto A : I * \beta^{P'} \mapsto D[C + A : I'] * I \xrightarrow{F} SD$ 
      *  $I' \xrightarrow{F} SD' * SH \xrightarrow{OF} (I, |SD' \cdot SB|) * TH \xrightarrow{OF} (I', |SD'|)$ 
    }
    while |buffer| = 4096
      s := write(targetH, buffer);
      buffer := read(sourceH, 4096);
      // We read less than 4096, so it was the last chunk
      s := write(targetH, buffer);
      {
         $\exists SH, TH, I', SD', SB.$ 
        sourceH ⇒ SH * targetH ⇒ TH * buffer ⇒ SB ∧ s ⇒ |SB|
        ∧ SD = SD' · SB ∧  $\alpha^P \mapsto A : I * \beta^{P'} \mapsto D[C + A : I']$ 
        *  $I \xrightarrow{F} SD * I' \xrightarrow{F} SD * SH \xrightarrow{OF} (I, |SD' \cdot SB|) * TH \xrightarrow{OF} (I', |SD|)$ 
      }
      close(sourceH); close(targetH); r := 0
    else close(sourceH); r := targetH
  else r := sourceH }
{  $\exists I'. r \Rightarrow 0 * \mathcal{E} * \alpha^P \mapsto A : I * \beta^{P'} \mapsto D[C + A : I'] * I \xrightarrow{F} SD * I' \xrightarrow{F} SD$  }

```

Fig. 10. Code for fileCopy, with a proof sketch of its specification.



$$\begin{aligned}
& \text{ent}(id, a) \triangleq \exists \iota. id = a : \iota \vee \exists id'. id = a[id'] \\
& \text{noent}(id, a) \triangleq \neg \exists id', id''. id \equiv \text{ent}(id') + id'' \\
\llbracket \text{mkdir}(\text{path}) \rrbracket (fs, \text{ph}, \sigma) &= \begin{cases} \{(fs[\mathcal{F} \mapsto id \circ_x b[id' + a[\emptyset]]], \text{ph}, \sigma)\} & \text{iff } \llbracket \text{path} \rrbracket_\sigma = p/b/a \wedge fs(\mathcal{F}) \equiv id \circ_x b[id'] \wedge \text{resolve}(p/\mathbf{x}, id) = \mathbf{x} \wedge \text{noent}(id', a) \\ \{(fs[\mathcal{F} \mapsto \top[id + a[\emptyset]]], \text{ph}, \sigma)\} & \text{iff } \llbracket \text{path} \rrbracket_\sigma = \top/a \wedge fs(\mathcal{F}) \equiv \top[d] \wedge \text{noent}(d, a) \\ \{\emptyset\} & \text{otherwise} \end{cases} \\
\llbracket \text{rmdir}(\text{path}) \rrbracket (fs, \text{ph}, \sigma) &= \begin{cases} \{(fs[\mathcal{F} \mapsto id], \text{ph}, \sigma)\} & \text{iff } \llbracket \text{path} \rrbracket_\sigma = p/a \wedge fs(\mathcal{F}) \equiv id \circ_x a[\emptyset] \wedge \text{resolve}(p/\mathbf{x}, id) = \mathbf{x} \\ \{\emptyset\} & \text{otherwise} \end{cases} \\
\llbracket \text{link}(\text{existing}, \text{new}) \rrbracket (fs, \text{ph}, \sigma) &= \begin{cases} \{(fs[\mathcal{F} \mapsto (id \circ_x d[id' + b : \iota]) \circ_y a : \iota], \text{ph}, \sigma)\} & \text{iff } \llbracket \text{existing} \rrbracket_\sigma = p/a \wedge \llbracket \text{new} \rrbracket_\sigma = p'/d/b \\ & \wedge fs(\mathcal{F}) \equiv (id \circ_x d[id']) \circ_y a : \iota \wedge \text{resolve}(p'/\mathbf{x}, id) = \mathbf{x} \\ & \wedge \text{resolve}(p/\mathbf{y}, id \circ_x d[id']) = \mathbf{y} \wedge \text{noent}(id', b) \\ \{(fs[\mathcal{F} \mapsto (\top[id' + b : \iota]) \circ_x a : \iota], \text{ph}, \sigma)\} & \text{iff } \llbracket \text{existing} \rrbracket_\sigma = p/a \wedge \llbracket \text{new} \rrbracket_\sigma = \top/b \\ & \wedge fs(\mathcal{F}) \equiv \top[id'] \circ_x a : \iota \wedge \text{resolve}(p'/\mathbf{x}, \top[id]) = \mathbf{x} \\ & \wedge \text{noent}(id', b) \\ \{\emptyset\} & \text{otherwise} \end{cases} \\
\llbracket \text{unlink}(\text{path}) \rrbracket (fs, \text{ph}, \sigma) &= \begin{cases} \{(fs[\mathcal{F} \mapsto id \circ_x \emptyset], \text{ph}, \sigma)\} & \text{iff } \llbracket \text{path} \rrbracket_\sigma = p/a \wedge fs(\mathcal{F}) \equiv id \circ_x a : \iota \wedge \text{resolve}(p/\mathbf{x}, id) = \mathbf{x} \\ \{\emptyset\} & \text{otherwise} \end{cases} \\
\llbracket \text{rename}(\text{old}, \text{new}) \rrbracket (fs, \text{ph}, \sigma) &= \begin{cases} \{(fs[\mathcal{F} \mapsto (id \circ_x \emptyset) \circ_y b : \iota], \text{ph}, \sigma)\} & \text{iff } \llbracket \text{old} \rrbracket_\sigma = p/a \wedge \llbracket \text{new} \rrbracket_\sigma = p'/b \\ & \wedge fs(\mathcal{F}) \equiv (id \circ_x a : \iota) \circ_y b : \iota \wedge \text{resolve}(p/\mathbf{x}, id) = \mathbf{x} \\ & \wedge \text{resolve}(p'/\mathbf{y}, id) = \mathbf{y} \\ \{(fs[\mathcal{F} \mapsto (id \circ_x d[id' + b : \iota]) \circ_y \emptyset], \text{ph}, \sigma)\} & \text{iff } \llbracket \text{old} \rrbracket_\sigma = p/a \wedge \llbracket \text{new} \rrbracket_\sigma = p'/d/b \\ & \wedge fs(\mathcal{F}) \equiv (id \circ_x d[id']) \circ_y a : \iota \wedge \text{resolve}(p'/\mathbf{x}, id) = \mathbf{x} \\ & \wedge \text{resolve}(p/\mathbf{y}, id \circ_x d[id']) = \mathbf{y} \wedge \text{noent}(id', b) \\ \{(fs[\mathcal{F} \mapsto \top[id + b : \iota] \circ_x \emptyset], \text{ph}, \sigma)\} & \text{iff } \llbracket \text{old} \rrbracket_\sigma = p/a \wedge \llbracket \text{new} \rrbracket_\sigma = \top/b \\ & \wedge fs(\mathcal{F}) \equiv \top[id] \circ_x a : \iota \wedge \text{resolve}(p/\mathbf{x}, \top[id]) = \mathbf{x} \\ & \wedge \text{noent}(id, b) \\ \{(fs[\mathcal{F} \mapsto (id \circ_x b[id']) \circ_y \emptyset], \text{ph}, \sigma)\} & \text{iff } \llbracket \text{old} \rrbracket_\sigma = p/a \wedge \llbracket \text{new} \rrbracket_\sigma = p'/b \\ & \wedge fs(\mathcal{F}) \equiv (id \circ_x b[id']) \circ_y a[id'] \wedge \text{resolve}(p/\mathbf{x}, id) = \mathbf{x} \\ & \wedge \text{resolve}(p'/\mathbf{y}, id) = \mathbf{y} \\ \{(fs[\mathcal{F} \mapsto (id \circ_x d[id' + b[id'']] \circ_y \emptyset], \text{ph}, \sigma)\} & \text{iff } \llbracket \text{old} \rrbracket_\sigma = p/a \wedge \llbracket \text{new} \rrbracket_\sigma = p'/d/b \\ & \wedge fs(\mathcal{F}) \equiv (id \circ_x d[id']) \circ_y a[id''] \wedge \text{resolve}(p'/\mathbf{x}, id) = \mathbf{x} \\ & \wedge \text{resolve}(p/\mathbf{y}, id \circ_x d[id']) = \mathbf{y} \wedge \text{noent}(id', b) \\ \{(fs[\mathcal{F} \mapsto \top[id + b[id']] \circ_x \emptyset], \text{ph}, \sigma)\} & \text{iff } \llbracket \text{old} \rrbracket_\sigma = p/a \wedge \llbracket \text{new} \rrbracket_\sigma = \top/b \\ & \wedge fs(\mathcal{F}) \equiv \top[id] \circ_x a[id'] \wedge \text{resolve}(p/\mathbf{x}, \top[id]) = \mathbf{x} \\ & \wedge \text{noent}(id, b) \\ \{(fs, \text{ph}, \sigma)\} & \text{iff } \llbracket \text{old} \rrbracket_\sigma = p/a \wedge \llbracket \text{new} \rrbracket_\sigma = p/a \\ & \wedge fs(\mathcal{F}) \equiv id \circ_x id' \wedge \text{resolve}(p/\mathbf{x}, id) = \mathbf{x} \\ & \wedge \text{ent}(id', a) \\ \{\emptyset\} & \text{otherwise} \end{cases} \\
\llbracket \text{t} := \text{stat}(\text{path}) \rrbracket (fs, \text{ph}, \sigma) &= \begin{cases} \{(fs, \text{ph}, \sigma[\text{t} \mapsto \text{F}])\} & \text{iff } \llbracket \text{path} \rrbracket_\sigma = p/a \wedge fs(\mathcal{F}) \equiv id \circ_x a : \iota \wedge \text{resolve}(p/\mathbf{x}, id) = \mathbf{x} \\ \{(fs, \text{ph}, \sigma[\text{t} \mapsto \text{D}])\} & \text{iff } \llbracket \text{path} \rrbracket_\sigma = p/a \wedge fs(\mathcal{F}) \equiv id \circ_x a[id'] \wedge \text{resolve}(p/\mathbf{x}, id) = \mathbf{x} \\ \{(fs, \text{ph}, \sigma[\text{t} \mapsto \text{ENOENT}])\} & \text{iff } \llbracket \text{path} \rrbracket_\sigma = p/a \wedge fs(\mathcal{F}) \equiv \top[d] \wedge \text{resolve}(p/a, \top[d]) = \text{udf} \\ \{\emptyset\} & \text{otherwise} \end{cases}
\end{aligned}$$

Fig. 11. Interpretation of the *structural* and *state* commands of our subset.

$$\begin{aligned}
[[\text{fd} := \text{open}(\text{path}, \text{flags})]](f, s, \text{ph}, \sigma) &= \begin{cases} \{(f, s, \text{ph}[f \mapsto (\iota, 0)][\iota \mapsto \epsilon], \sigma[\text{fd} \mapsto f])\} & \text{iff } \llbracket \text{path} \rrbracket_\sigma = p/b/a \wedge f, s(\mathcal{F}) \equiv id \circ_{\mathbf{x}} b[id'] \wedge \text{resolve}(p/\mathbf{x}, id) = \mathbf{x} \\ & \wedge \text{noent}(id', a) \wedge f \notin \text{dom}(\text{ph}) \wedge \iota \notin \text{dom}(f, s) \\ \{(f, s, \text{ph}[f \mapsto (\iota, 0)][\iota \mapsto \epsilon], \sigma[\text{fd} \mapsto f])\} & \text{iff } \llbracket \text{path} \rrbracket_\sigma = \tau/a \wedge f, s(\mathcal{F}) \equiv \tau[d] \\ & \wedge \text{noent}(d, a) \wedge f \notin \text{dom}(\text{ph}) \wedge \iota \notin \text{dom}(f, s) \\ \{(f, s, \text{ph}[f \mapsto (\iota, 0)], \sigma[\text{fd} \mapsto f])\} & \text{iff } \llbracket \text{path} \rrbracket_\sigma = p/a \wedge \llbracket \text{flags} \rrbracket_\sigma = 0 \wedge f, s(\mathcal{F}) \equiv id \circ_{\mathbf{x}} a : \iota \\ & \wedge \text{resolve}(p/\mathbf{x}, id) = \mathbf{x} \wedge f \notin \text{dom}(\text{ph}) \wedge \iota \in \text{dom}(f, s) \\ \{(f, s, \text{ph}[f \mapsto (\iota, 0)][\iota \mapsto \epsilon], \sigma[\text{fd} \mapsto f])\} & \text{iff } \llbracket \text{path} \rrbracket_\sigma = p/a \wedge \llbracket \text{flags} \rrbracket_\sigma = \text{O\_TRUNC} \wedge f, s(\mathcal{F}) \equiv id \circ_{\mathbf{x}} a : \iota \\ & \wedge \text{resolve}(p/\mathbf{x}, id) = \mathbf{x} \wedge f \notin \text{dom}(\text{ph}) \wedge \iota \in \text{dom}(f, s) \\ \{\emptyset\} & \text{otherwise} \end{cases} \\
[[\text{off} := \text{lseek}(\text{fd}, \text{to}, \text{wh})]](f, s, \text{ph}, \sigma) &= \begin{cases} \{(f, s, \text{ph}[f \mapsto (\iota, \text{offset})], \sigma[\text{off} \mapsto \text{offset}])\} & \text{iff } \llbracket \text{fd} \rrbracket_\sigma = f \wedge \llbracket \text{to} \rrbracket_\sigma = \text{offset} \wedge \llbracket \text{wh} \rrbracket_\sigma = \text{SEEK\_SET} \\ & \wedge \text{ph}(f) = (\iota, o) \wedge \text{offset} \geq 0 \\ \{(f, s, \text{ph}[f \mapsto (\iota, o + \text{offset})], \sigma[\text{off} \mapsto o + \text{offset}])\} & \text{iff } \llbracket \text{fd} \rrbracket_\sigma = f \wedge \llbracket \text{to} \rrbracket_\sigma = \text{offset} \wedge \llbracket \text{wh} \rrbracket_\sigma = \text{SEEK\_CUR} \\ & \wedge \text{ph}(f) = (\iota, o) \wedge o + \text{offset} \geq 0 \\ \{(f, s, \text{ph}[f \mapsto (\iota, |\text{bs}| + \text{offset})], \sigma[\text{off} \mapsto |\text{bs}| + \text{offset}])\} & \text{iff } \llbracket \text{fd} \rrbracket_\sigma = f \wedge \llbracket \text{to} \rrbracket_\sigma = \text{offset} \wedge \llbracket \text{wh} \rrbracket_\sigma = \text{SEEK\_END} \\ & \wedge \text{ph}(f) = (\iota, o) \wedge f, s(\iota) = \text{bs} \wedge |\text{bs}| + \text{offset} \geq 0 \\ \{\emptyset\} & \text{otherwise} \end{cases} \\
[[\text{sz} := \text{write}(\text{fd}, \text{buf})]](f, s, \text{ph}, \sigma) &= \begin{cases} \{(f, s[\iota \mapsto \text{bs}' \cdot \text{bs} \cdot \text{bs}''], \text{ph}[f \mapsto (\iota, |\text{bs}| + o)], \sigma[\text{sz} \mapsto |\text{bs}|])\} & \text{iff } \llbracket \text{fd} \rrbracket_\sigma = f \wedge \llbracket \text{buf} \rrbracket_\sigma = \text{bs} \wedge \text{ph}(f) = (\iota, o) \\ & \wedge f, s(\iota) \equiv \text{bs}' \cdot \text{bs}'' \cdot \text{bs}''' \wedge |\text{bs}'| = o \wedge |\text{bs}'''| = |\text{bs}| \\ \{(f, s[\iota \mapsto \text{bs}' \cdot \text{bs}], \text{ph}[f \mapsto (\iota, |\text{bs}| + o)], \sigma[\text{sz} \mapsto |\text{bs}|])\} & \text{iff } \llbracket \text{fd} \rrbracket_\sigma = f \wedge \llbracket \text{buf} \rrbracket_\sigma = \text{bs} \wedge \text{ph}(f) = (\iota, o) \\ & \wedge f, s(\iota) \equiv \text{bs}' \cdot \text{bs}'' \wedge |\text{bs}'| = o \wedge |\text{bs}''| < |\text{bs}| \\ \{(f, s[\iota \mapsto \text{bs}'(\oplus_0) \cdot \text{bs}], \text{ph}[f \mapsto (\iota, |\text{bs}| + o)], \sigma[\text{sz} \mapsto |\text{bs}|])\} & \text{iff } \llbracket \text{fd} \rrbracket_\sigma = f \wedge \llbracket \text{buf} \rrbracket_\sigma = \text{bs} \wedge \text{ph}(f) = (\iota, o) \\ & \wedge f, s(\iota) \equiv \text{bs}' \wedge o \geq |\text{bs}'| \wedge z_s = o - |\text{bs}'| \\ \{\emptyset\} & \text{otherwise} \end{cases} \\
[[\text{buf} := \text{read}(\text{fd}, \text{sz})]](f, s, \text{ph}, \sigma) &= \begin{cases} \{(f, s, \text{ph}[f \mapsto (\iota, o + \text{size})], \sigma[\text{buf} \mapsto \text{bs}'])\} & \text{iff } \llbracket \text{fd} \rrbracket_\sigma = f \wedge \llbracket \text{sz} \rrbracket_\sigma = \text{size} \in \mathbb{N} \wedge \text{ph}(f) = (\iota, o) \\ & \wedge f, s(\iota) \equiv \text{bs} \cdot \text{bs}' \cdot \text{bs}'' \wedge |\text{bs}'| = o \wedge \text{size} = |\text{bs}'| \\ \{(f, s, \text{ph}[f \mapsto (\iota, o + |\text{bs}'|)], \sigma[\text{buf} \mapsto \text{bs}'])\} & \text{iff } \llbracket \text{fd} \rrbracket_\sigma = f \wedge \llbracket \text{sz} \rrbracket_\sigma = \text{size} \in \mathbb{N} \wedge \text{ph}(f) = (\iota, o) \\ & \wedge f, s(\iota) \equiv \text{bs} \cdot \text{bs}' \wedge |\text{bs}'| = o \wedge \text{size} > |\text{bs}'| \\ \{(f, s, \text{ph}, \sigma[\text{buf} \mapsto \epsilon])\} & \text{iff } \llbracket \text{fd} \rrbracket_\sigma = f \wedge \llbracket \text{sz} \rrbracket_\sigma = \text{size} \in \mathbb{N} \wedge \text{ph}(f) = (\iota, o) \\ & \wedge f, s(\iota) \equiv \text{bs} \wedge o > |\text{bs}'| \\ \{\emptyset\} & \text{otherwise} \end{cases} \\
[[\text{close}(\text{fd})]](f, s, \text{ph}, \sigma) &= \begin{cases} \{(f, s, \text{ph} \upharpoonright (\text{dom}(\text{ph}) - \{f\}), \sigma)\} & \text{iff } \llbracket \text{fd} \rrbracket_\sigma = f \wedge f \in \text{dom}(\text{ph}) \\ \{\emptyset\} & \text{otherwise} \end{cases} \\
[[\text{dir} := \text{opendir}(\text{path})]](f, s, \text{ph}, \sigma) &= \begin{cases} \{(f, s, \text{ph}[n \mapsto A], \sigma[\text{dir} \mapsto n])\} & \text{iff } \llbracket \text{path} \rrbracket_\sigma = p/b \wedge f, s(\mathcal{F}) \equiv id \circ_{\mathbf{x}} b[id'] \wedge \text{resolve}(p/\mathbf{x}, id) = \mathbf{x} \\ & \wedge n \notin \text{dom}(\text{ph}) \wedge \forall a. a \in A \iff \exists id'', id'''. id' \equiv id'' + \text{ent}(id''', a) \\ \{\emptyset\} & \text{otherwise} \end{cases} \\
[[\text{fn} := \text{readdir}(\text{dir})]](f, s, \text{ph}, \sigma) &= \begin{cases} \{(f, s, \text{ph}[n \mapsto A \setminus \{a\}], \sigma[\text{fn} \mapsto a])\} & \text{iff } \llbracket \text{dir} \rrbracket_\sigma = n \wedge \text{ph}(n) = A \wedge a \in A \\ \{(f, s, \text{ph}, \sigma[\text{fn} \mapsto 0])\} & \text{iff } \llbracket \text{dir} \rrbracket_\sigma = n \wedge \text{ph}(n) = \emptyset \\ \{\emptyset\} & \text{otherwise} \end{cases}
\end{aligned}$$

Fig. 12. Interpretation of IO commands of our subset.