

The Specification of POSIX File Systems

Gian Ntzik, Pedro da Rocha Pinto and Philippa Gardner

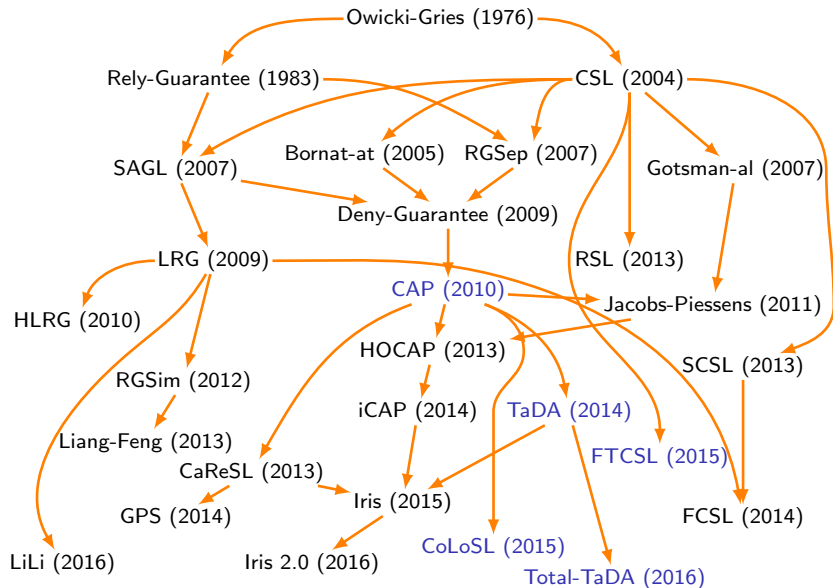


Imperial College London

`gian.ntzik08@imperial.ac.uk`
`pedro.da-rocha-pinto09@imperial.ac.uk`
`p.gardner@imperial.ac.uk`

INVEST 2017

Logics for Specifying Concurrent Programs




Thanks, Ilya Sergey.

POSIX File Systems


- The POSIX standard is written in English
- Several formal sequential specifications have been given:
 - not good for client reasoning if based on first-order logic;
 - good for client reasoning if based on separation logic.
- The English specification of the concurrent behaviour is poor.
- We have the first concurrent specification of POSIX file systems.

A POSIX operation: `unlink(/usr/bin/git)`

`unlink(/usr/bin/git)` takes a sequence of atomic actions



unlink(/usr/bin/git): a Sequence of Atomic Actions



Concurrent Specification of `unlink(path)`

```
unlink(path)  $\sqsubseteq$  let  $p = \text{dirname}(path)$ ;  
  let  $a = \text{basename}(path)$ ;  
  let  $r = \text{resolve}(p, \iota_0)$ ;  
  if  $\neg \text{iserr}(r)$  then  
    return  $\text{link\_delete}(r, a)$   
       $\sqcup$   $\text{link\_delete\_notdir}(r, a)$   
  else return } r fi
```

- $\text{link_delete}(r, a)$ atomically removes link a , even if it links a directory
- $\text{link_delete_notdir}(r, a)$ atomically removes link a , only if it does not link a directory
- \sqcup : non-deterministic *angelic choice*

Concurrent Specification of path resolution

```
letrec resolve(path,  $\iota$ )  $\triangleq$   
  if path = null then return  $\iota$  else  
    let a = head(path);  
    let p = tail(path);  
    let r = link_lookup( $\iota$ , a);  
    if iserr(r) then return r  
    else return resolve(p, r) fi  
fi
```

- `link_lookup(ι , a)` **atomically** lookup the link named a in the directory with inode ι .

Concurrent Specification of `unlink(path)`

```
unlink(path)  $\sqsubseteq$  let p = dirname(path);  
                let a = basename(path);  
                let r = resolve(p,  $\iota_0$ );  
                if  $\neg$ iserr(r) then  
                    return link_delete(r, a)  
                         $\sqsubseteq$  link_delete_notdir(r, a)  
                else return r fi
```


Concurrent Specification of `unlink(path)`

```
unlink(path)  $\sqsubseteq$  let  $p = \text{dirname}(path)$ ;  
                  let  $a = \text{basename}(path)$ ;  
                  let  $r = \text{resolve}(p, \iota_0)$ ;  
                  if  $\neg \text{iserr}(r)$  then  
                    return link_delete( $r, a$ )  
                             $\sqcup$  link_delete_notdir( $r, a$ )  
                  else return  $r$  fi
```

link_delete(ι, a) \triangleq

$\forall FS. \langle \text{fs}(FS) \wedge \text{isdir}(FS(\iota)), a \in FS(\iota) \implies \text{fs}(FS[\iota \mapsto FS(\iota) \setminus \{a\}]) * \text{ret} = 0 \rangle$

\sqcap return **enoent**(ι, a)

\sqcap return **enotdir**(ι)

enoent(ι, a) and **enotdir**(ι) describe error cases. \sqcap is non-deterministic *demonic* choice.

Concurrent specification of `link(source, target)`


```
link(source, target)
  ⊑ let  $p_s = \text{dirname}(source)$ ; let  $a = \text{basename}(source)$ ;
    let  $p_t = \text{dirname}(target)$ ; let  $b = \text{basename}(target)$ ;
    let  $r_s, r_t = \text{resolve}(p_s, \iota_0) \parallel \text{resolve}(p_t, \iota_0)$ ;
    if  $\neg \text{iserr}(r_s) \wedge \neg \text{iserr}(r_t)$  then
      return link_insert( $r_s, a, r_t, b$ )
        ⊑ link_insert_notdir( $r_s, a, r_t, b$ )
    else if  $\text{iserr}(r_s) \wedge \neg \text{iserr}(r_t)$  then return  $r_s$ 
    else if  $\neg \text{iserr}(r_s) \wedge \text{iserr}(r_t)$  then return  $r_t$ 
    else if  $\text{iserr}(r_s) \wedge \text{iserr}(r_t)$  then return  $r_s$  ⊑ return  $r_t$  fi
```

Lock-file Client

- ▶ `lock(path)`: atomically create a non-existing lock file at `path`
- ▶ `unlock(path)`: remove the lock file identified by `path`
- ▶ Implemented similarly to spin locks
 - ▶ `open(path, O_CREAT|O_EXCL)` to try to lock
 - ▶ `unlink` to unlock

Lock-file Protocol Agreement

We want the module and the environment to agree on a boundary



Lock-file Specification

$LF(path) \vdash \text{lock}(path) \sqsubseteq \forall v \in \{0, 1\}. \langle \text{Lock}(path, v) , \text{Lock}(path, 1) * v = 0 \rangle$

$LF(path) \vdash \text{unlock}(path) \sqsubseteq \langle \text{Lock}(path, 1) , \text{Lock}(path, 0) \rangle$

- ▶ $LF(path)$ is a *context invariant* denoting what the implementation and the environment can and cannot do to the path:
 - ▶ The path-prefix of $path$ is not changed (everyone can only read)
 - ▶ Only the module create the lock file at $path$ (locks the lock)
 - ▶ Only the module removes the lock file at $path$ (unlocks the lock)

Conclusions

- Concurrent specification of the POSIX file system (fragment)
- Combination of the refinement calculus and a modern concurrent separation logic called TaDA
- Client reasoning: lock files, named pipes, email server
- Future: executable specification which we can test against implementations
- Future: now what about distribution....