

# Refinement of Parallel Algorithms down to LLVM

Peter Lammich  

University of Twente, Netherlands

## Abstract

We present a stepwise refinement approach to develop verified parallel algorithms, down to efficient LLVM code. The resulting algorithms' performance is competitive with their counterparts implemented in C/C++. Our approach is backwards compatible with the Isabelle Refinement Framework, such that existing sequential formalizations can easily be adapted or re-used. As case study, we verify a parallel quicksort algorithm, and show that it performs on par with its C++ implementation, and is competitive to state-of-the-art parallel sorting algorithms.

**2012 ACM Subject Classification** Software and its engineering → Formal software verification; Theory of computation → Semantics and reasoning; Computing methodologies → Parallel algorithms

**Keywords and phrases** Isabelle, Concurrent Separation Logic, Parallel Sorting, LLVM

**Digital Object Identifier** 10.4230/LIPIcs.ITP.2022.12

**Supplementary Material** *Software (Isabelle Formalization)*: [https://www21.in.tum.de/~lammich/isabelle\\_llvm\\_par/](https://www21.in.tum.de/~lammich/isabelle_llvm_par/)

## 1 Introduction

We present a stepwise refinement approach to develop verified and efficient parallel algorithms. Our method can verify total correctness down to LLVM intermediate code. The resulting verified implementations are competitive with state-of-the-art unverified implementations. Our approach is backwards compatible to the Isabelle Refinement Framework (IRF), a powerful tool to verify efficient sequential software, such as model checkers [10, 7, 38], SAT solvers [24, 25, 11], or graph algorithms [22, 28, 29]. This paper adds parallel execution to the IRF's toolbox, without invalidating the existing formalizations, which can now be used as sequential building blocks for parallel algorithms, or be modified to add parallelization.

As a case study, we verify total correctness of a parallel quicksort algorithm, re-using an existing verification of state-of-the-art sequential sorting algorithms [27]. Our verified parallel sorting algorithm is competitive to state-of-the-art parallel sorting algorithms.

### 1.1 Overview

This paper is based on the Isabelle Refinement Framework, a continuing effort to verify efficient implementations of complex algorithms, using stepwise refinement techniques. Figure 1 displays the components of the Isabelle Refinement Framework.

The back end layer handles the translation from Isabelle/HOL to the actual target language. The instructions of the target language are shallowly embedded into Isabelle/HOL, using a state-error (SE) monad. An instruction with undefined behaviour, or behaviour outside our supported fragment, raises an error. The state of the monad is the memory, represented via a memory model. The code generator translates the instructions to actual code. These components form the trusted code base, while all the remaining components of the Isabelle Refinement Framework generate proofs. In the back-end, the preprocessor transforms expressions to the syntactically restricted format required by the code generator, proving semantic equality of the original and transformed expression. While there exist back ends for purely functional code [30, 21], and sequential imperative code [23, 26], this paper describes a back end for parallel imperative LLVM code (Section 2).



© Peter Lammich;

licensed under Creative Commons License CC-BY 4.0

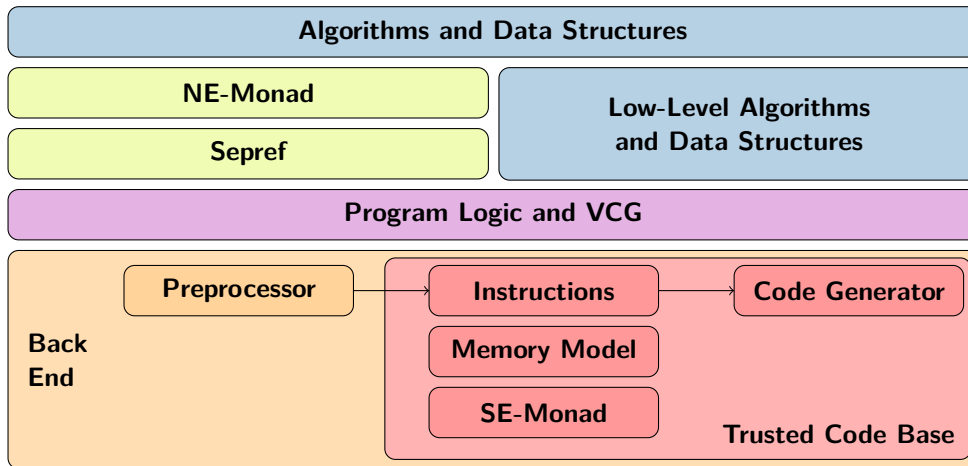
13th International Conference on Interactive Theorem Proving (ITP 2022).

Editors: June Andronick and Leonardo de Moura; Article No. 12; pp. 12:1–12:20

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



■ **Figure 1** Components of the Isabelle Refinement Framework, with focus on the back end.

44 On top of the back-end, a program logic is used to prove programs correct. It uses  
 45 separation logic, and provides automation like a verification condition generator (VCG). In  
 46 Section 3, we describe our formalization of concurrent separation logic [33], and our VCG.

47 At the level of the program logic and VCG, our framework can already be used to  
 48 verify simple low-level algorithms and data structures, like dynamic arrays and linked lists.  
 49 More complex developments typically use a stepwise refinement approach, starting at purely  
 50 functional programs modelled in a nondeterminism-error (NE) monad [30]. A semi-automatic  
 51 refinement procedure (Sepref [23, 26]) translates from the purely functional code to imperative  
 52 code, refining abstract functional data types to concrete imperative ones. In Section 4, we  
 53 describe our extensions to support refinement to parallel executions, and a fine-grained  
 54 tracking of pointer equalities, required to parallelize computations that work on disjoint  
 55 parts of the same array.

56 Using our approach, complex algorithms and data structures can be developed and refined  
 57 to optimized efficient code. The stepwise refinement ensures a separation of concerns between  
 58 high-level algorithmic ideas and low-level optimizations. We have used this approach to  
 59 verify a wide range of practically efficient algorithms [10, 7, 38, 24, 25, 11, 22, 28, 29, 27].  
 60 In Section 5, we use our techniques to verify a parallel sorting algorithm, with competitive  
 61 performance wrt. unverified state-of-the-art algorithms.

62 Section 6 concludes the paper and discusses related and future work.

## 63 **2 A Back End for LLVM with Parallel Execution**

64 We formalize a semantics for parallel execution, shallowly embedded into Isabelle/HOL. As  
 65 for the existing sequential back ends [23, 26], the shallow embedding is key to the flexibility  
 66 and feasibility of the approach. The main idea is to make an execution report the memory  
 67 that it accesses, and use this information to raise an error when joining executions that would  
 68 have exhibited a data race. We use this to model an instruction that calls two functions in  
 69 parallel, and waits until both have returned.

## 2.1 State-Nondeterminism-Error Monad with Access Reports

We define the underlying monad in two steps. We start with a nondeterminism-error monad, and then lift it to a state monad and add access reports. Defining a nondeterminism-error monad is straightforward in Isabelle/HOL:

```

74 'a neM ≡ spec ('a ⇒ bool) | fail
75 return x ≡ spec (λr. r=x)
76 bind fail f ≡ fail
77 bind (spec P) f ≡ if ∃x. P x ∧ f x = fail then fail
78                   else spec (λr. ∃x Q. P x ∧ f x = spec Q ∧ Q r)
79
80

```

A program either fails, or yields a possible set of results (`spec P`), described by its characteristic function  $P$ . The `return` operation yields exactly one result, and `bind` combines all possible results, failing if there is a possibility to fail.

Now assume that we have a state (memory) type  $'\mu$ , and an access report type  $'\rho$ , which forms a monoid  $(0, +)$ . With this, we define our state-nondeterminism-error monad with access reports, just called  $M$  for brevity:

```

87 'x M ≡ '\mu ⇒ ('x × '\rho × '\mu) neM
88 return_M x μ ≡ return_ne (x, 0, μ)
89 bind_M m f μ ≡ (x1, r1, μ) ← m μ; (x2, r2, μ) ← f x1 μ; return_ne (x2, r1 + r2, μ)
90
91

```

Here, `return` does not change the state, and reports no accesses ( $0$ ), and `bind` sequentially composes the executions, threading through the state  $\mu$ , and adding up the access reports  $r_1$  and  $r_2$ .

Typically, the access report will contain read and written addresses, such that data races can be detected. Moreover, if parallel executions can allocate memory, we must detect those executions where the memory manager allocated the same block in both parallel strands. As we assume a thread safe memory manager, those *infeasible* executions can safely be ignored. Let  $norace :: '\rho \Rightarrow '\rho \Rightarrow bool$  and  $feasible :: '\rho \Rightarrow '\rho \Rightarrow bool$  be symmetric predicates, and let  $combine :: ('\rho \times '\mu) \Rightarrow ('\rho \times '\mu) \Rightarrow ('\rho \times '\mu)$  be a commutative operator to compose two pairs of access reports and states. Then, we define a parallel composition operator for  $M$ :

```

102 (m1 || m2) μ ≡
103   (x1, r1, μ1) ← m1 μ; (x2, r2, μ2) ← m2 μ;           — execute both strands
104   assume feasible ρ1 ρ2;                                       — ignore infeasible combinations
105   assert norace ρ1 ρ2;                                           — fail on data race
106   return_ne ((x1, x2), combine (ρ1, μ1) (ρ2, μ2))         — combine results
107
108
109 assume P ≡ if P then return () else spec (λ_. False)
110 assert P ≡ if P then return () else fail
111

```

Here, we use `assume` to ignore infeasible executions, and `assert` to fail on data races. Note that, if one parallel strand fails, and the other parallel strand has no possible results `spec (λ_. False)`, the behaviour of the parallel composition is not clear. For this reason, we fix an invariant  $invar_M :: ('\mu \Rightarrow ('x \times '\rho \times '\mu) neM) \Rightarrow bool$ , which implies that every non-failing execution has at least one possible result. We define the actual type  $M$  as the subtype satisfying  $invar_M$ . Thus, we have to prove that every combinator and instruction of our semantics preserves the invariant, which is an important sanity check. As additional sanity check, we prove symmetry of parallel composition:

## 12:4 Refinement of Parallel Algorithms down to LLVM

```

120
121  $m_1 \parallel m_2 = mswap (m_2 \parallel m_1)$    where    $mswap\ m \equiv (x_1, x_2) \leftarrow m; \mathbf{return}\ (x_2, x_1)$ 
122

```

### 123 2.2 Memory Model

124 Our memory model supports blocks of values, where values can be integers, structures, or  
125 pointers into a block:

```

126 datatype  $addr \equiv ADDR\ (bidx: nat)\ (idx: nat)$ 
127 datatype  $ptr \equiv PTR\_NULL \mid PTR\_ADDR\ (the\_addr: addr)$ 
128 datatype  $val \equiv LL\_INT\ lint \mid LL\_STRUCT\ val\ list \mid LL\_PTR\ ptr$ 
129
130
131 datatype  $block \equiv FRESH \mid FREED \mid is\_alloc: ALLOC\ (vals: val\ list)$ 
132 typedef  $memory \equiv \{ \mu :: nat \Rightarrow block.\ finite\ \{b.\ \mu\ b \neq FRESH\} \}$ 
133

```

134 A block is either fresh, freed, or allocated, and a memory is a mapping from block indexes  
135 to blocks, such that only finitely many blocks are not fresh. Every block's state transitions  
136 from fresh to allocated to freed. This avoids ever reusing the same block, and thus allows  
137 us to semantically detect use after free errors. Every program execution can only allocate  
138 finitely many blocks, such that we will never run out of fresh blocks<sup>1</sup>. An allocated block  
139 contains an array of values, modelled as a list. Thus, an address consists of a block number,  
140 and an index into the array.

141 To access and modify memory, we define the functions *valid*, *get*, and *put*:

```

142  $valid\ \mu\ (ADDR\ b\ i) \equiv is\_alloc\ (\mu\ b) \wedge i < |vals\ (\mu\ b)|$ 
143  $get\ \mu\ (ADDR\ b\ i) \equiv vals\ (\mu\ b)\ !\ i$ 
144  $put\ \mu\ (ADDR\ b\ i)\ x \equiv \mu(b := ALLOC\ ((vals\ (\mu\ b))[i:=x]))$ 
145

```

147 where  $|xs|$  is the length of list  $xs$ ,  $xs[i]$  returns the  $i$ th element of list  $xs$ , and  $xs[i:=x]$  replaces  
148 the  $i$ th element of  $xs$  by  $x$ .

149 Note that our LLVM semantics does not support conversion of pointers to integers, nor  
150 comparison or difference of pointers to different blocks. This way, a program cannot see the  
151 internal representation of a pointer, and we can choose a simple abstract representation,  
152 while being faithful wrt. any actual memory manager implementation.

### 153 2.3 Access Reports

154 We now fix the state of the M-monad to be memory, and the access reports to be sets of  
155 read and written addresses, as well as sets of allocated and freed blocks:

```

156  $acc \equiv ( r :: addr\ set; w :: addr\ set; a :: nat\ set; f :: nat\ set )$ 
157  $\emptyset \equiv ( \{\}, \{\}, \{\}, \{\} )$ 
158  $(r_1, w_1, a_1, f_1) + (r_2, w_2, a_2, f_2) \equiv ( r_1 \cup r_2, w_1 \cup w_2, a_1 \cup a_2, f_1 \cup f_2 )$ 
159

```

161 Two parallel executions are feasible if they did not allocate the same block, and they  
162 have a data race if one strand accesses addresses or blocks modified by the other strand:

```

163  $feasible\ (r_1, w_1, a_1, f_1)\ (r_2, w_2, a_2, f_2) \equiv a_1 \cap a_2 = \{\}$ 
164

```

<sup>1</sup> If the actual system does run out of memory, we will terminate the program in a defined way.

```

166 norace ( $r_1, w_1, a_1, f_1$ ) ( $r_2, w_2, a_2, f_2$ )  $\equiv$ 
167   let  $m_1 = w_1 \cup \{ ADDR\ b\ i.\ b \in a_1 \cup f_1 \}$  in
168   let  $m_2 = w_2 \cup \{ ADDR\ b\ i.\ b \in a_2 \cup f_2 \}$  in
169      $(r_1 \cup m_1) \cap m_2 = \{\}$   $\wedge$   $m_1 \cap (r_2 \cup m_2) = \{\}$ 
170

```

171 The invariant for  $M$  states that blocks transition only from fresh to allocated to free, allocated  
 172 blocks never change their size, and the access report matches the observable state change  
 173 (*consistent*). It also states, that for each finite set of blocks  $B$ , there is an execution that  
 174 does not allocate blocks from  $B$ . The latter is required to show that we always find feasible  
 175 parallel executions:

```

176 invarM  $c \equiv \forall \mu.\ P.\ c\ \mu = \text{spec } P \implies$ 
177    $(\forall x\ \rho\ \mu'.\ P(x, \rho, \mu') \implies \text{consistent } \mu\ \rho\ \mu')$ 
178    $\wedge (\forall B.\ \text{finite } B \implies (\exists x\ \rho\ \mu'.\ P(x, \rho, \mu') \wedge \rho.a \cap B = \{\}))$ 
179

```

181 The combine function joins the access reports and memories, preferring allocated over fresh,  
 182 and freed over allocated memory. When joining two allocated blocks, the written addresses  
 183 from the access report are used to join the blocks. We skip the rather technical definition of  
 184 combine, and just state the relevant properties: Let  $\rho_1 = (r_1, w_1, a_1, f_1)$  and  $\rho_2 = (r_2, w_2, a_2, f_2)$  be  
 185 feasible and race free access reports, and  $\mu_1, \mu_2$  be memories that have evolved from a common  
 186 memory  $\mu$ , consistently with the access reports  $\rho_1, \rho_2$ . Let  $(\rho', \mu') = \text{combine}(\rho_1, \mu_1)(\rho_2, \mu_2)$ ,  
 187 and *addr* a valid address in  $\mu'$ . Then

```

188
189 (1)  $\mu' b = \text{FRESH} \iff \mu b = \text{FRESH} \wedge b \notin a_1 \cup a_2$ 
190 (2)  $\text{is\_alloc}(\mu' b) \iff (\text{is\_alloc}(\mu b) \vee b \in a_1 \cup a_2) \wedge b \notin f_1 \cup f_2$ 
191 (3)  $\mu' b = \text{FREED} \iff \mu b = \text{FREED} \vee b \in f_1 \cup f_2$ 
192
193 (4)  $a \in w_1 \vee b \in a_1 \implies \text{get\_addr } \mu' a = \text{get\_addr } \mu_1 a$ 
194 (5)  $a \in w_2 \vee b \in a_2 \implies \text{get\_addr } \mu' a = \text{get\_addr } \mu_2 a$ 
195 (6)  $a \notin w_1 \cup w_2 \vee b \notin a_1 \cup a_2 \implies \text{get\_addr } \mu' a = \text{get\_addr } \mu a$ 
196

```

197 The properties (1)–(3) define the state of blocks in the combined memory: a fresh block in  
 198  $\mu'$  was fresh already in  $\mu$ , and has not been allocated (1); an allocated block was already  
 199 allocated or has been allocated, but has not been freed (2); and a freed block was already  
 200 freed, or has been freed (3). The properties (4)–(6) define the content: addresses written or  
 201 allocated in the first or second execution get their content from  $\mu_1$  (4) or  $\mu_2$  (5) respectively.  
 202 Addresses not written or allocated at all keep their original content (6).

## 203 2.4 LLVM Instructions

204 Based on the M-monad, we define shallowly embedded LLVM instructions. For most  
 205 instructions, this is analogous to the sequential case [26]. The exceptions are memory allocation,  
 206 which nondeterministically allocates some available block (the original formalization  
 207 deterministically counted up the block indexes), and an instruction for parallel function call:

```

208 llc_par  $f\ g\ a\ b \equiv f\ a \parallel g\ b$ 
209

```

211 The code generator only accepts this, if  $f$  and  $g$  are constants (i.e., function names). It then  
 212 generates some type-casting boilerplate, and a call to an external *parallel* function, which we  
 213 implement using the Threading Building Blocks [36] library:

```

214 void parallel(void (*f1)(void*), void (*f2)(void*), void *x1, void *x2) {
215

```

## 12:6 Refinement of Parallel Algorithms down to LLVM

```
216 tbb::parallel_invoke([=]{f1(x1);}, [=]{f2(x2);}); }
```

218 I.e., the two functions  $f1(x1)$  and  $f2(x2)$  are called in parallel. The generated boilerplate code  
219 sets up  $x1$  and  $x2$  to point to both, the actual arguments and space for the results.

### 3 Parallel Separation Logic

221 In the previous section, we have defined a shallow embedding of LLVM programs into  
222 Isabelle/HOL. We now describe how to reason about these programs, using separation logic.

#### 3.1 Separation Algebra

224 In order to reason about memory with separation logic, we define an abstraction function  
225 from the memory into a separation algebra [8]. Separation algebras formalize the intuition of  
226 combining disjoint parts of memory. They come with a *zero* ( $0$ ) that describes the empty  
227 part, a *disjointness predicate*  $a \# b$  describing that the parts  $a$  and  $b$  do not overlap, and a  
228 *disjoint union*  $a + b$  that combines two disjoint parts. For the exact definition of a separation  
229 algebra, we refer to [8, 20]. We note that separation algebras naturally extend over functions  
230 and pairs, in a pointwise manner.

231 ► **Example 1.** (Trivial Separation Algebra) The type  $\alpha \text{ option} = \text{None} \mid \text{Some } \alpha$  forms a  
232 separation algebra with:

```
233  $0 \equiv \text{None} \quad a \# b \equiv a=0 \vee b=0 \quad a + 0 \equiv a \quad 0 + b \equiv b$ 
```

236 Intuitively, this separation algebra does not allow for combination of contents, except if one  
237 side is zero. While it is not very useful on its own, the trivial separation algebra is a useful  
238 building block for more complex separation algebras.

239 For our memory model, we define the following abstraction function:

```
240  $\alpha :: \text{memory} \rightarrow (\text{addr} \rightarrow \text{val option}) \times (\text{nat} \rightarrow \text{nat option})$ 
```

```
241  $\alpha \mu \equiv (\alpha_m \mu, \alpha_b \mu)$ 
```

```
242  $\alpha_m \mu \text{ addr} \equiv \text{if } \text{valid } \mu \text{ addr} \text{ then } \text{Some } (\text{get } \mu \text{ addr}) \text{ else } 0$ 
```

```
243  $\alpha_b \mu b \equiv \text{if } \text{is\_alloc } (\mu b) \text{ then } \text{Some } (|\text{vals } (\mu b)|) \text{ else } 0$ 
```

247 An abstract memory  $\alpha \mu$  consists of two parts:  $\alpha_m \mu$  is a map from addresses to the values  
248 stored there. It is used to reason about load and store operations.  $\alpha_b \mu$  is a map from  
249 block indexes to the sizes of the corresponding blocks. It is used to ensure that one owns all  
250 addresses of a block when freeing it.

251 We continue to define a separation logic: assertions are predicates over separation algebra  
252 elements. The basic connectives are defined as follows:

```
253  $\text{false } a \equiv \text{False} \quad \text{true } a \equiv \text{True} \quad \square a \equiv a=0$ 
```

```
254  $(P * Q) a \equiv \exists a_1 a_2. a_1 \# a_2 \wedge a = a_1 + a_2 \wedge P a_1 \wedge Q a_2$ 
```

257 That is, the assertion *false* never holds and the assertion *true* holds for all abstract memories.  
258 The empty assertion  $\square$  holds for the zero memory, and the separating conjunction  $P * Q$  holds  
259 if the memory can be split into two disjoint parts, such that  $P$  holds for one, and  $Q$  holds for  
260 the other part. The lifting assertion  $\uparrow \phi$  holds iff the Boolean value  $\phi$  is true:

```
261  $\uparrow \phi \equiv \text{if } \phi \text{ then } \square \text{ else } \text{false}$ 
```

264 It is used to lift plain logical statements into separation logic assertions owning no memory.  
 265 When clear from the context, we omit the  $\uparrow$ -symbol, and just mix plain statements with  
 266 separation logic assertions.

## 267 3.2 Weakest Preconditions and Hoare Triples

268 We define a *weakest precondition* predicate directly via the semantics:

$$269 \quad wp\ m\ Q\ \mu \equiv \text{case } m\ \mu \text{ of spec } Q' \Rightarrow \forall x\ \rho\ \mu'. Q'(x, \rho, \mu') \Longrightarrow Q\ x\ \rho\ \mu' \mid \text{fail} \Rightarrow \text{False}$$

272 That is,  $wp\ m\ Q\ \mu$  holds, iff program  $m$  run on memory  $\mu$  does not fail, and all possible  
 273 results (return value  $x$ , access report  $\rho$ , new memory  $\mu'$ ) satisfy the *postcondition*  $Q$ .

274 To set up a verification condition generator based on separation logic, we standardize the  
 275 postcondition: the reported memory accesses must be disjoint from some abstract memory  
 276  $amf$ , called the *frame*. We define the *weakest precondition with frame*:

$$277 \quad wpf\ amf\ c\ Q\ \mu \equiv wp\ c\ (\lambda x\ \rho\ \mu'. Q\ x\ \mu' \wedge \text{disjoint } \rho\ amf)\ \mu$$

$$280 \quad \text{disjoint } (r, w, a, f)\ (m, b) \equiv (\forall \text{addr}. m\ \text{addr} \neq 0 \Longrightarrow \text{addr} \notin r \cup w \wedge \text{addr}. \text{bidx} \notin f) \\
 281 \quad \wedge (\forall i. b\ i \neq 0 \Longrightarrow i \notin f)$$

283 that is, when executed on memory  $\mu$ , the program  $c$  does not fail, every return value  $x$  and  
 284 new memory  $\mu'$  satisfies  $Q$ , and no memory described by the frame  $amf$  is accessed.

285 Equipped with a weakest precondition with access restrictions, we define a Hoare-triple:

$$286 \quad ABS\ amf\ P\ \mu \equiv \exists am. am \# amf \wedge \alpha\ \mu = am + amf \wedge P\ am$$

$$289 \quad ht\ P\ c\ Q \equiv \forall \mu\ amf. ABS\ amf\ P\ \mu \Longrightarrow wpf\ amf\ c\ (\lambda x\ \mu'. ABS\ amf\ (Q\ x)\ \mu')\ \mu$$

291 The predicate  $ABS\ amf\ P\ \mu$  specifies that the abstract memory  $\alpha\ \mu$  can be split into a  
 292 part  $am$  and the given frame  $amf$ , such that  $am$  satisfies the precondition  $P$ . A Hoare-  
 293 triple  $ht\ P\ c\ Q$  specifies that for all memories and frames for which the precondition holds  
 294 ( $ABS\ amf\ P\ \mu$ ), the program will succeed, not using any memory of the frame, and every  
 295 result will satisfy the postcondition wrt. the original frame ( $ABS\ amf\ (Q\ x)\ \mu'$ ).

## 296 3.3 Verification Condition Generator

297 The verification condition generator is implemented as a proof tactic that works on subgoals  
 298 of the form:

$$299 \quad ABS\ amf\ P\ \mu \wedge \dots \Longrightarrow wpf\ amf\ c\ Q\ \mu$$

302 The tactic is guided by the syntax of the command  $c$ . Basic monad combinators are broken  
 303 down using the following rules:

$$304 \quad Q\ r\ \mu \Longrightarrow wpf\ amf\ (\text{return } r)\ Q\ \mu \\
 306 \quad wpf\ amf\ m\ (\lambda x. wpf\ amf\ (f\ x)\ Q)\ \mu \Longrightarrow wpf\ amf\ (\{x \leftarrow m; f\ x\})\ Q\ \mu$$

308 For other instructions and user defined functions, the VCG expects a Hoare-triple to be  
 309 already proved. It then uses the following rule:

$$310 \quad ht\ P\ c\ Q \wedge ABS\ amf\ P'\ \mu \quad \text{--- match Hoare triple and current state} \\
 311 \quad \wedge P' \vdash P * F \quad \text{--- infer frame} \\
 312 \quad \wedge (\wedge r\ \mu. ABS\ amf\ (Q\ r * F)\ \mu \Longrightarrow Q'\ r\ \mu) \quad \text{--- continue with postcondition} \\
 313 \quad \Longrightarrow wpf\ amf\ c\ Q'\ \mu$$



## 12:8 Refinement of Parallel Algorithms down to LLVM

316 To process a command  $c$ , the first assumption is instantiated with the Hoare-triple for  $c$ , and  
317 the second assumption with the assertion  $P'$  for the current state. Then, a simple syntactic  
318 heuristics infers a frame  $F$  and proves that the current assertion  $P'$  entails the required  
319 precondition  $P$  and the frame. Finally, verification condition generation continues with the  
320 postcondition  $Q$  and the frame as current assertion.

### 321 3.4 Hoare-Triples for Instructions

322 To use the VCG to verify LLVM programs, we have to prove Hoare triples for the LLVM  
323 instructions. For parallel calls, we prove the well-known disjoint concurrency rule [33]:

$$324 \quad ht\ P_1\ c_1\ Q_1 \wedge ht\ P_2\ c_2\ Q_2 \implies ht\ (P_1 * P_2)\ (par\ c_1\ c_2)\ (\lambda(r_1, r_2). Q_1\ r_1 * Q_2\ r_2)$$

327 That is, commands with disjoint preconditions can be executed in parallel.

328 For memory operations, we prove:

$$329 \quad \models \{n \neq 0\}\ ll\_malloc\ TYPE(\alpha)\ n\ \{\lambda p. range\ \{0..<n\}\ (\lambda_. init)\ p * b\_tag\ n\ p\}$$
$$330 \quad \models \{range\ \{0..<n\}\ xs\ p * b\_tag\ n\ p\}\ ll\_free\ p\ \{\lambda_. \square\}$$
$$331 \quad \models \{pto\ x\ p\}\ ll\_load\ p\ \{\lambda r. r = x * pto\ x\ p\}$$
$$332 \quad \models \{pto\ y\ p\}\ ll\_store\ x\ p\ \{\lambda_. pto\ x\ p\}$$

333

335 Here  $b\_tag\ n\ p$  asserts that  $p$  points to the beginning of a block of size  $n$ , and  $range\ I\ f\ p$   
336 describes that for all  $i \in I$ ,  $p + i$  points to value  $f\ i$ . Intuitively,  $ll\_malloc$  creates a block of  
337 size  $n$ , initialized with the default  $init$  value, and a tag. If one possesses both, the whole block  
338 and the tag, it can be deallocated by  $free$ . The rules for load and store are straightforward,  
339 where  $pto\ x\ p$  describes that  $p$  points to value  $x$ .

## 340 4 Refinement for Parallel Programs

341 At this point, we have described a separation logic framework for parallel programs in  
342 LLVM. It is largely backwards compatible with the framework for sequential programs  
343 described in [26], such that we could easily port the algorithms formalized there to our  
344 new framework. The next step towards verifying complex programs is to set up a stepwise  
345 refinement framework. In this section we describe the refinement infrastructure of the Isabelle  
346 Refinement Framework, focusing on our changes to support parallel algorithms.

### 347 4.1 Abstract Programs

348 Abstract programs are shallowly embedded into the nondeterminism error monad  $'a\ neM$  (cf.  
349 Section 2.1). They are purely functional, not modifying memory, or differentiating between  
350 sequential and parallel execution. We define a *refinement ordering* on  $neM$ :

$$351 \quad \text{spec } P \leq \text{spec } Q \equiv \forall x. P\ x \implies Q\ x \quad \text{fail} \not\leq \text{spec } Q \quad m \leq \text{fail}$$

352

354 Intuitively,  $m_1 \leq m_2$  means that  $m_1$  returns fewer possible results than  $m_2$ , and may only  
355 fail if  $m_2$  may fail. Note that  $\leq$  is a complete lattice, with top element  $\text{fail}$ .

356 We use refinement and assertions to specify that a program  $m$  satisfies a specification  
357 with precondition  $P$  and postcondition  $Q$ :

$$358 \quad m \leq \text{assert } P; \text{spec } x. Q\ x$$

359



361 If the precondition is false, the right hand side is `fail`, and the statement trivially holds.  
 362 Otherwise,  $m$  cannot fail, and every possible result  $x$  of  $m$  must satisfy  $Q$ .

363 For a detailed description on using the  $ne$ -monad for stepwise refinement based program  
 364 verification, we refer the reader to [30].

## 365 4.2 The Sepref Tool

366 The Sepref tool [23, 26] symbolically executes an abstract program in the  $ne$ -monad, keeping  
 367 track of refinements for every abstract variable to a concrete representation, which may  
 368 use pointers to dynamically allocated memory. During the symbolic execution, the tool  
 369 synthesizes an imperative Isabelle-LLVM program, together with a refinement proof. The  
 370 synthesis is automatic, but requires annotations to the abstract program.

371 The main concept of the Sepref tool is refinement between an abstract program  $c$  in the  
 372  $ne$ -monad, and a concrete program  $c_{\dagger}$  in the  $M$  monad, as expressed by the  $hnr$ -predicate:

$$373 \quad hnr \Gamma c_{\dagger} \Gamma' R CP c \equiv \\ 374 \quad c \neq \text{fail} \implies ht \Gamma c_{\dagger} (\lambda x_{\dagger}. \exists x. \Gamma' * R x x_{\dagger} * \uparrow(\text{return } x \leq c \wedge CP x_{\dagger}))$$

377 That is, either the abstract program  $c$  fails, or for a memory described by assertion  $\Gamma$ , the  
 378 LLVM program  $c_{\dagger}$  succeeds with  $x_{\dagger}$ , such that the new memory is described by  $\Gamma' * R x x_{\dagger}$ ,  
 379 for a possible result  $x$  of the abstract program  $c$ . Moreover, the predicate  $CP$  holds for the  
 380 concrete result. Note that  $hnr$  trivially holds for a failing abstract program. This makes  
 381 sense, as we prove that the abstract program does not fail anyway. Moreover it allows us to  
 382 assume that assertions actually hold during the refinement proof:

$$383 \quad (\phi \implies hnr \Gamma c_{\dagger} \Gamma' R CP c) \implies hnr \Gamma c_{\dagger} \Gamma' R CP (\text{assert } \phi; c)$$

386 ► **Example 2.** (Refinement of lists to arrays) We define abstract programs for indexing and  
 387 updating a list:

$$388 \quad lget \, xs \, i \equiv \text{assert } (i < |xs|); \text{return } xs[i] \quad lset \, xs \, i \, x \equiv \text{assert } (i < |xs|); \text{return } xs[i:=x]$$

391 These programs assert that the index is in bounds, and then return the accessed element  
 392 ( $xs[i]$ ) or the updated list ( $xs[i:=x]$ ) respectively. The following assertion links a pointer to a  
 393 list of elements stored at the pointed-to location:

$$394 \quad arr_A \, xs \, p = \text{range } \{0..<|xs|\} (\lambda i. xs[i]) \, p$$

397 That is, for every  $i < |xs|$ ,  $p + i$  points to the  $i$ th element of  $xs$ . On arrays, indexing and  
 398 updating of arrays is implemented by:

$$399 \quad aget \, p \, i \equiv ll\_ofs\_ptr \, p \, i; ll\_load \, p \quad aset \, p \, i \, x \equiv ll\_ofs\_ptr \, p \, i; ll\_store \, x \, p; \text{return } p$$

402 And the abstract and concrete programs are linked by the following refinement theorems:

$$403 \quad hnr (arr_A \, xs \, xs_{\dagger} * idx_A \, i \, i_{\dagger}) (aget \, xs_{\dagger} \, i_{\dagger}) (arr_A \, xs \, xs_{\dagger} * idx_A \, i \, i_{\dagger}) id_A (\lambda_. True) (lget \, xs \, i) \\ 404 \quad hnr (arr_A \, xs \, xs_{\dagger} * idx_A \, i \, i_{\dagger}) (aset \, xs_{\dagger} \, i_{\dagger} \, x) (idx_A \, i \, i_{\dagger}) arr_A (\lambda r. r=xs_{\dagger}) (lset \, xs \, i \, x)$$

407 That is, if the list  $xs$  is refined by array  $xs_{\dagger}$ , and the natural number  $i$  is refined by the  
 408 fixed-width<sup>2</sup> word  $i_{\dagger}$  ( $idx_A \, i \, i_{\dagger}$ ), the  $aget$  operation will return the same result as the  $lget$

<sup>2</sup> We use Isabelle's word library here, which encodes the actual width as a type variable, such that our functions work with any bit width. For code generation, we will fix the width to 64 bit.

## 12:10 Refinement of Parallel Algorithms down to LLVM

409 operation ( $id_A$ ). The resulting memory will still contain the original array. Note that there  
 410 is no explicit precondition that the array access is in bounds, as this follows already from the  
 411 assertion in the abstract *lget* operation. The *aset* operation will return a pointer to an array  
 412 that refines the updated list returned by *lset*. As the array is updated in place, the original  
 413 refinement of the array is no longer valid. Moreover, the returned pointer  $r$  will be the same  
 414 as the argument pointer  $xs_{\dagger}$ . This information is important for refining to parallel programs  
 415 on disjoint parts of an array (cf. Section 4.3).

416 Given refinement assertions for the parameters, and *hnr*-rules for all operations in a  
 417 program, the Sepref tool automatically synthesizes an LLVM program from an abstract *neM*  
 418 program. The tool tries to automatically discharge additional proof obligations, typically  
 419 arising from translating arithmetic operations from unbounded numbers to fixed width  
 420 numbers. Where automatic proof fails, the user has to add assertions to the abstract program  
 421 to help the proof. The main difference of our tool wrt. the existing Sepref tool [26] is the  
 422 additional condition (*CP*) on the concrete result, which is used to track pointer equalities.  
 423 We have added a heuristics to automatically synthesize and discharge these equalities.

### 4.3 Array Splitting

425 An important concept for parallel programs is to concurrently operate on disjoint parts of  
 426 the memory, e.g., different slices of the same array. However, abstractly, arrays are just lists.  
 427 They are updated by returning a new list, and there is no way to express that the new list is  
 428 stored at the same address as the old list. Nevertheless, in order to refine a program that  
 429 updates two disjoint slices of a list to one that updates disjoint parts of the array in place,  
 430 we need to know that the result is stored in the same array as the input. This is handled by  
 431 the *CP* argument to *hnr*. To indicate that operations shall be refined to disjoint parts of the  
 432 same array, we introduce the combinator `with_split` for abstract programs:

```
433 with_split  $i$   $xs$   $f$   $\equiv$   

  434   assert ( $i < |xs|$ );  

  435   ( $xs_1, xs_2$ )  $\leftarrow f$  ( $take\ i\ xs$ ) ( $drop\ i\ xs$ );  

  436   assert ( $|xs_1| = i \wedge |xs_2| = |xs| - i$ );  

  437   return ( $xs_1 @ xs_2$ )  

  438  

  439
```

440 Abstractly, this is an annotation that is inlined when proving the abstract program correct.  
 441 However, Sepref will translate it to the concrete combinator *awith\_split*:

```
442 awith_split  $i$   $xs_{\dagger}$   $f_{\dagger}$   $\equiv$   $ll\_ofs\_ptr\ xs_{\dagger}\ i;$   $f_{\dagger}\ xs_{\dagger}\ xs_{\dagger 2};$  return  $xs_{\dagger}$   

  443  

  444  

  445 hnr ( $arr_A\ xs_1\ xs_{\dagger 1} * arr_A\ xs_2\ xs_{\dagger 2}$ ) ( $f_{\dagger}\ xs_{\dagger 1}\ xs_{\dagger 2}$ )  $\square$   

  446   ( $arr_A \times arr_A$ ) ( $\lambda(xs'_{\dagger 1}, xs'_{\dagger 2}). xs'_{\dagger 1} = xs_{\dagger 1} \wedge xs'_{\dagger 2} = xs_{\dagger 2}$ )  

  447   ( $f\ xs_1\ xs_2$ )  

  448  $\implies$   

  449 hnr ( $arr_A\ xs\ xs_{\dagger} * idx_A\ i\ i_{\dagger}$ ) (awith_split  $i_{\dagger}\ xs_{\dagger}\ f_{\dagger}$ )  

  450   ( $idx_A\ i\ i_{\dagger}$ ) ( $\lambda xs\ xs_{\dagger}. arr_A\ xs\ xs_{\dagger}$ ) ( $\lambda xs'_{\dagger}. xs'_{\dagger} = xs_{\dagger}$ )  

  451   (with_split  $i\ xs\ f$ )  

  452
```

453 The refinement of the function  $f$  to  $f_{\dagger}$  requires an additional proof that the returned pointers  
 454 are equal to the argument pointers ( $xs'_{\dagger 1} = xs_{\dagger 1} \wedge xs'_{\dagger 2} = xs_{\dagger 2}$ ). Sepref tries to prove that  
 455 automatically, using a simple heuristics.

#### 4.4 Refinement to Parallel Execution

The purely functional abstract programs have no notion of parallel execution. To indicate that refinement to parallel execution is desired, we define an abstract annotation `npar`:

```

459 npar f g a b  $\equiv x \leftarrow f\ a; y \leftarrow g\ b; \mathbf{return}\ (x,y)$ 
460
461
462  $hnr\ Ax\ (f_{\dagger}\ x_{\dagger})\ Ax'\ Rx\ CP_1\ (f\ x) \wedge hnr\ Ay\ (g_{\dagger}\ y_{\dagger})\ Ay'\ Ry\ CP_2\ (g\ y)$ 
463  $\implies$ 
464  $hnr\ (Ax * Ay)\ (llc\_par\ f_{\dagger}\ g_{\dagger}\ x_{\dagger}\ y_{\dagger})\ (Ax' * Ay')\ (Rx \times Ry)$ 
465  $(\lambda(x'_{\dagger}, y'_{\dagger}).\ CP_1\ x'_{\dagger} \wedge CP_2\ y'_{\dagger})\ (\mathbf{npar}\ f\ g\ x\ y)$ 
466

```

This rule can be used to automatically parallelize any (independent) abstract computations. For convenience, we also define `nseq`. Abstractly, it's the same as `npar`, but Sepref translates it to sequential execution.

## 5 A Parallel Sorting Algorithm

To test the usability of our framework, we verify a parallel sorting algorithm. We start with the abstract specification of an algorithm that sorts a list:

```

474 sort_spec xs = spec xs'. mset xs' = mset xs  $\wedge$  sorted xs
475

```

That is, we return a sorted permutation of the original list. Note that this is a standard specification of sorting in Isabelle. Reusing the existing development of an abstract introsort algorithm [27], we easily prove with a few refinement steps that the following abstract algorithm implements `sort_spec`:

```

480 1 psort xs n  $\equiv \mathbf{assert}\ n=|xs|; \mathbf{if}\ n \leq 1 \mathbf{then}\ \mathbf{return}\ xs \mathbf{else}\ psort\_aux\ xs\ n\ (\log_2\ n * 2)$ 
481 2
482 3 psort_aux xs n d  $\equiv$ 
483 4  $\mathbf{assert}\ n=|xs|$ 
484 5  $\mathbf{if}\ d=0 \vee n < 100000 \mathbf{then}\ sort\_spec\ xs$ 
485 6  $\mathbf{else}$ 
486 7  $(xs,m) \leftarrow partition\_spec\ xs;$ 
487 8  $\mathbf{let}\ bad = m < n\ \mathit{div}\ 8 \vee (n-m < n\ \mathit{div}\ 8)$ 
488 9  $(_,xs) \leftarrow \mathbf{with\_split}\ m\ xs\ (\lambda xs_1\ xs_2.$ 
489 10  $\mathbf{if}\ bad \mathbf{then}\ \mathbf{nseq}\ psort\_aux\ psort\_aux\ (xs_1,m,d-1)\ (xs_2,n-m,d-1)$ 
490 11  $\mathbf{else}\ \mathbf{npar}\ psort\_aux\ psort\_aux\ (xs_1,m,d-1)\ (xs_2,n-m,d-1)$ 
491 12  $);$ 
492 13  $\mathbf{return}\ xs$ 
493 14
494 15 lemma psort xs  $|xs| \leq sort\_spec\ xs$ 
495
496

```

This algorithm is derived from the well-known quicksort and introsort algorithms [32]: like quicksort, it partitions the list (line 7), and then recursively sorts the partitions in parallel (l. 11). Like introsort, when the recursion gets too deep, or the list too short, we fall back to some (not yet specified) sequential sorting algorithm (l. 5). Similarly, when the partitioning is very unbalanced (l. 8), we sort the partitions sequentially (l. 10). These optimizations aim at not spawning threads for small sorting tasks, where the overhead of thread creation outweighs the advantages of parallel execution. A more technical aspect is the extra parameter  $n$  that

## 12:12 Refinement of Parallel Algorithms down to LLVM

504 we introduced for the list length. Thus, we can refine the list to just a pointer to an array,  
505 and still access its length<sup>3</sup>.

### 506 5.1 Implementation and Correctness Theorem

507 Next, we have to provide implementations for the fallback *sort\_spec*, and for *partition\_spec*.  
508 These implementations must be proved to be in-place, i.e., return a pointer to the same array.  
509 It was straightforward to amend our existing formalization of *pdqsort* [27] with the in-place  
510 proofs: once we had amended the refinement statements, and bug-fixed the pointer equality  
511 proving heuristics that we added to Sepref, the proofs were automatic.

512 Given the implementations of *sort\_spec* and *partition\_spec*, the Sepref tool generates an  
513 LLVM program *psort<sub>†</sub>* from the abstract *psort*, and proves a corresponding refinement lemma:

```
514 hnr (arrA xs xs† * idxA n n†) (psort† xs† n†) (idxA n n†) arrA ( $\lambda r. r = xs†$ ) (psort xs n)
```

517 Combining this with the correctness lemma of the abstract *psort* algorithm, and unfolding  
518 the definition of *hnr*, we prove the following Hoare-triple for our final implementation:

```
519 ht (arrA xs xs† * idxA n n† * n = |xs|)  
520 (psort† xs† n†)  
521 ( $\lambda r. r = xs† * \exists xs'. arr_A xs' xs† * sorted xs' * mset xs' = mset xs$ )
```

524 That is, for a pointer *xs<sub>†</sub>* to an array, whose contents are described by list *xs* (*arr<sub>A</sub>*), and a  
525 fixed-size word *n<sub>†</sub>* representing the natural number *n* (*idx<sub>A</sub>*), which must be the number of  
526 elements in the list *xs*, our sorting algorithm returns the original pointer *xs<sub>†</sub>*, and the array  
527 contents are now *xs'*, which is sorted and a permutation of *xs*. Note that this statement uses  
528 our semantically defined Hoare triples (cf. Section 3.2). In particular, its correctness does  
529 not depend on the refinement steps, the Sepref tool, or the VCG.

### 530 5.2 A Sampling Partitioner

531 While we could simply re-use the existing partitioning algorithm from the *pdqsort* formaliza-  
532 tion, which uses a pseudomedian of nine pivot selection, we observe that the quality of the  
533 pivot is particularly important for a balanced parallelization. Moreover, the partitioning in  
534 the *psort\_aux* procedure is only done for arrays above a quite big size threshold. Thus, we  
535 can invest a little more work to find a good pivot, which is still negligible compared to the  
536 cost of sorting the resulting partitions. We choose a sampling approach, using the median of  
537 64 equidistant samples as pivot. The highly optimized partitioning algorithms that we use  
538 swap the pivot to the front of the partition, such that we need to determine its index, rather  
539 than just its value. We simply use quicksort to find the median<sup>4</sup>:

```
540 sample xs  $\equiv is \leftarrow equidist |xs| 64; is \leftarrow sort\_wrt (\lambda i j. xs!i < xs!j) is; \mathbf{return} (is!32)$ 
```

543 Proving that this algorithm finds a valid pivot index is straightforward. More challenging is to  
544 refine it to purely imperative LLVM code, which does not support closures like  $\lambda i j. xs!i < xs!j$ .

545 We resolve such closures over the comparison function manually: using Isabelle's locale  
546 mechanism [19], we parametrize over the comparison function. Moreover, we thread through  
547 an extra parameter for the data captured by the closure:

<sup>3</sup> Alternatively, we could refine a list to a pair of array pointer and length.

<sup>4</sup> We leave verification of efficient median algorithms, e.g., quickselect, to future work. Note that the overhead of sorting 64 elements is negligible compared to the large partition that has to be sorted.

```

548
549 locale pcmp =
550   fixes lt :: 'p ⇒ 'e ⇒ 'e ⇒ bool and lt† :: 'p† ⇒ 'e† ⇒ 'e† ⇒ bool
551     and parA :: 'p ⇒ 'p† ⇒ assn and elemA :: 'e ⇒ 'e† ⇒ assn
552   assumes  $\forall p. \text{weak\_ordering } (lt\ p)$ 
553   assumes hnr (parA p pi * elemA a ai * elemA b bi) (lt† pi ai bi)
554     (parA p pi * elemA a ai * elemA b bi) (boolA) ( $\lambda.. \text{True}$ ) (lt p a b)
555

```

556 This defines a context in which we have an abstract compare function *lt* for the abstract  
 557 elements of type 'e. It takes an extra parameter of type 'p (e.g. the list *xs*), and forms a  
 558 weak ordering<sup>5</sup>. Note that the strict compare function *lt* also induces a non-strict version  
 559  $le\ p\ a\ b \equiv \neg lt\ p\ b\ a$ . Moreover, we have a concrete implementation *lt<sub>†</sub>* of the compare  
 560 function, wrt. the refinement assertions *par<sub>A</sub>* for the parameter and *elem<sub>A</sub>* for the elements.

561 Our sorting algorithm is developed and verified in the context of this locale (to avoid  
 562 confusion, our presentation has, up to now, just used  $<$ ,  $\leq$ , and *sorted* instead of *lt* *p*, *le* *p*,  
 563 and *sorted\_wrt* (*le* *p*)). To get a sorting algorithm for an actual compare function, we have  
 564 to instantiate the locale, providing an abstract and concrete compare function, along with a  
 565 proof that the abstract function is a weak ordering, and the concrete function refines the  
 566 abstract one. For our example of sorting indexes into an array, where the array elements are,  
 567 themselves, compared by a parametrized function *lt*, we get:

```

568
569 interpretation idx: pcmp lt_idx lt_idx† (parA × arrA) idxA <proof>
570
571 lt_idx (p, xs) i j ≡ lt p (xs!i) (xs!j)
572 lt_idx† (p†, xs†) i† j† ≡ x† ← aget xs† i†; y† ← aget xs† j†; lt† p† x† y†
573

```

574 this yields sorting algorithms for sorting indexes, taking an extra parameter for the array to  
 575 index into. For our sampling application, we use *idx.introsort* *xs*.

### 576 5.3 Code Generation

577 Finally, we instantiate the sorting algorithms to sort unsigned integers and strings:

```

578
579 interpretation unat: pcmp ( $\lambda.. <$ ) ( $\lambda.. ll\_icmp\_ult$ ) unatA64 <proof>
580 interpretation str: pcmp ( $\lambda.. <$ ) ( $\lambda.. strcmp$ ) strA64 <proof>
581

```

582 This yields implementations *unat.psort<sub>†</sub>* and *str.psort<sub>†</sub>*, and automatically proves instantiated  
 583 versions of the correctness theorems.

584 In a last step, we use our code generator to generate actual LLVM text, as well as a C  
 585 header file with the signatures of the generated functions<sup>6</sup>:

```

586
587 export_llvm
588   unat.psort† is uint64_t* psort(uint64_t*, int64_t)
589   str.psort† is llstring* str_psort(llstring*, int64_t)
590   defines typedef struct {int64_t size; struct {int64_t capacity; char *data;};} llstring;
591   file psort.ll
592

```

<sup>5</sup> A weak ordering is induced by a mapping of the elements into a total ordering. It is the standard prerequisite for sorting algorithms in C++ [17].

<sup>6</sup> For technical reasons, we represent the array size as non-negative signed integer, thus the C signature uses *int64\_t*. Moreover, we use a string implementation based on dynamic arrays, rather than C's zero terminated strings.

593 This checks that the specified C signatures are compatible with the actual types, and then  
 594 generates *psort.ll* and *psort.h*, which can be used in a standard C/C++ toolchain.

## 595 5.4 Benchmarks

596 We have benchmarked our verified sorting algorithm against a direct implementation of the  
 597 same algorithm in C++. The result was that both implementations have the same runtime,  
 598 up to some minor noise. This indicates that there is no systemic slowdown: algorithms  
 599 verified with our framework run as fast as their unverified counterparts implemented in C++.

600 We also benchmarked against the state-of-the-art implementations *std::sort* with execution  
 601 policy *par\_unseq* from the GNU C++ standard library [12], and *sample\_sort* from the Boost  
 602 C++ libraries [4, 5]. We have benchmarked the algorithm on two different machines, and  
 603 various input distributions. The results are shown in Figure 2. While our verified algorithm  
 604 is clearly competitive for integer sorting on the less parallel laptop machine, it’s slightly less  
 605 efficient for sorting strings on the highly parallel server machine. Nevertheless, we believe  
 606 that our verified implementation is already useful in practice, and leave further optimizations  
 607 to future work.

608 Finally, we measured the speedup that the implementations achieve for a certain number  
 609 of cores. The results are displayed in Figure 3. While the speedup on the moderately parallel  
 610 laptop is comparable to the one of the C++ standard library, our implementation achieves  
 611 lower speedups than the state-of-the-art on the highly parallel server. Again, we leave further  
 612 optimizations to future work.

## 613 6 Conclusions

614 We have presented a stepwise refinement approach to verify total correctness of efficient  
 615 parallel algorithms. Our approach targets LLVM as back end, and there is no systemic  
 616 efficiency loss in our approach when compared to unverified algorithms implemented in C++.


617 The trusted code base of our approach is relatively small: apart from Isabelle’s inference  
 618 kernel, it contains our shallow embedding of a small fragment of the LLVM semantics, and  
 619 the code generator. All other tools that we used, e.g., our Hoare logic, Sepref tool, and  
 620 Refinement Framework for abstract programs, ultimately prove a correctness theorem that  
 621 only depends on our shallowly embedded semantics.

622 As a case study, we have implemented a parallel sorting algorithm. It uses an existing  
 623 verified sequential pdqsort algorithm as a building block, and is competitive with state-of-  
 624 the-art parallel sorting algorithms, at least on moderately parallel hardware.

625 The main idea of our parallel extension is to shallowly embed the semantics of a parallel  
 626 combinator into a sequential semantics, by making the semantics report the accessed memory  
 627 locations, and fail if there is a potential data race. We only needed to change the lower  
 628 levels of our existing framework for sequential LLVM [26]. Higher-level tools like the VCG  
 629 and Sepref remained largely unchanged and backwards compatible. This greatly simplified  
 630 reusing of existing verification projects, like the sequential pdqsort algorithm [27].

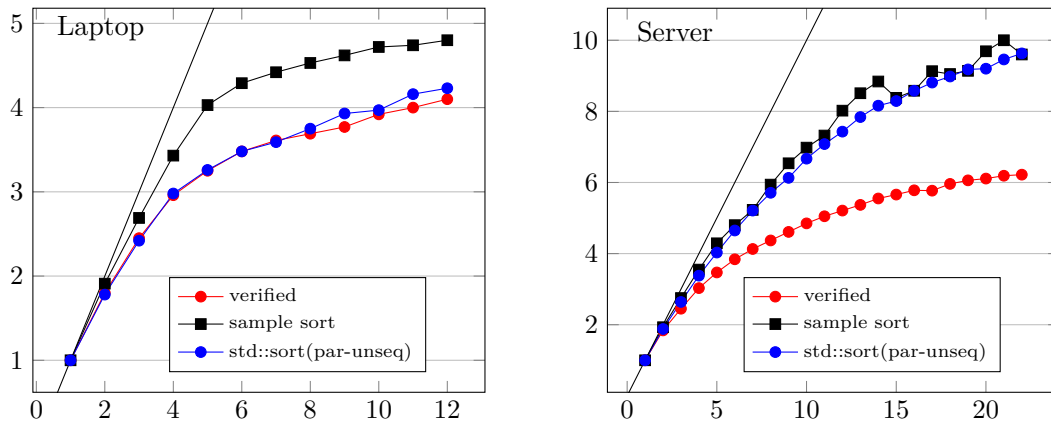
### 631 6.1 Related Work

632 While there is extensive work on parallel sorting algorithm (e.g. [9, 1]), there seems to be  
 633 almost no work on their formal verification. The only work we are aware of is a distributed  
 634 merge sort algorithm [16], for which ”no effort has been made to make it efficient”[16, Sec. 2],  
 635 nor any executable code has been generated or benchmarked. Another verification [34] uses



■ **Figure 2** Runtimes in milliseconds for sorting various distributions of unsigned 64 bit integers and strings with our verified parallel sorting algorithm, C++’s standard parallel sorting algorithm, and Boost’s parallel sample sort algorithm. The experiments were performed on a server machine with 22 AMD Opteron 6176 cores and 128GiB of RAM, and a laptop with a 6 core (12 threads) i7-10750H CPU and 32GiB of RAM.





■ **Figure 3** Speedup of the various implementations, for sorting unsigned 64 bit integers with a random distribution, on a server with 22 AMD Opteron 6176 cores and 128GiB of RAM, and a laptop with a 6 core (12 threads) i7-10750H CPU and 32GiB of RAM. The x axis ranges over the number of cores, and the y-axis gives the speedup wrt. the same implementation run on only one core. The thin black lines indicate linear speedup.

636 the VerCors deductive verifier to prove the permutation property ( $mset\ xs' = mset\ xs$ ) of  
 637 odd-even transposition sort [13], but neither the sortedness property nor termination.

638 Concurrent separation logic is used by many verification tools such as VerCors [3], and also  
 639 formalized in proof assistants, for example in the VST [37] and IRIS [18] projects for Coq [2].  
 640 These formalizations contain elaborate concepts to reason about communication between  
 641 threads via shared memory, and are typically used to verify partial correctness of subtle  
 642 concurrent algorithms (e.g. [31]). Reasoning about total correctness is more complicated  
 643 in the step-indexed separation logic provided by IRIS, and currently only supported for  
 644 sequential programs [35]. Our approach is less expressive, but naturally supports total  
 645 correctness, and is already sufficient for many practically relevant parallel algorithms like  
 646 sorting, matrix-multiplication, or parallel algorithms from the C++ STL.

## 647 6.2 Future Work

648 An obvious next step is to implement a fractional separation logic [6], to reason about parallel  
 649 threads that share read-only memory. While our semantics already supports shared read-only  
 650 memory, our separation logic does not. We believe that implementing a fractional separation  
 651 logic will be straightforward, and mainly pose technical issues for automatic frame inference.

652 Another obvious next step is to verify a state-of-the-art parallel sorting algorithm, like  
 653 Boost’s sample sort. Like our current algorithm, sample sort does not require advanced  
 654 synchronization concepts, and can be implemented only with a parallel combinator.

655 Finally, the Sepref framework has recently been extended to reason about complexity of  
 656 (sequential) LLVM programs [14, 15]. This line of work could be combined with our parallel  
 657 extension, to verify the complexity (e.g. work and span) of parallel algorithms.

658 Extending our approach towards more advanced synchronization like locks or atomic  
 659 operations may be possible: instead of accessed memory addresses, a thread could report a  
 660 set of possible traces, which are checked for race-freedom and then combined.

661 Finally, our framework currently targets multicore CPUs. Another important architecture  
 662 are general purpose GPUs. As LLVM is also available for GPUs, porting our framework to

<sup>663</sup> this architecture should be possible. We even expect that barrier synchronization, which is  
<sup>664</sup> important in the GPU context, can be integrated into our approach.

665 — **References** —

- 666 1 Mikhail Asiatici, Damian Maiorano, and Paolo Ienne. How many cpu cores is an fpga worth?  
667 lessons learned from accelerating string sorting on a cpu-fpga system. *Journal of Signal*  
668 *Processing Systems*, pages 1–13, 2021.
- 669 2 Yves Bertot and Pierre Castran. *Interactive Theorem Proving and Program Development: Coq'Art The Calculus of Inductive Constructions*. Springer Publishing Company, Incorporated,  
670 1st edition, 2010.
- 672 3 Stefan Blom, Saeed Darabi, Marieke Huisman, and Wytse Oortwijn. The vercors tool set:  
673 Verification of parallel and concurrent software. In Nadia Polikarpova and Steve Schneider,  
674 editors, *Integrated Formal Methods*, pages 102–110, Cham, 2017. Springer International  
675 Publishing.
- 676 4 Boost C++ libraries. <https://www.boost.org/>.
- 677 5 Boost C++ libraries sorting algorithms. [https://www.boost.org/doc/libs/1\\_77\\_0/libs/  
678 sort/doc/html/index.html](https://www.boost.org/doc/libs/1_77_0/libs/sort/doc/html/index.html).
- 679 6 Richard Bornat, Cristiano Calcagno, Peter O'Hearn, and Matthew Parkinson. Permission  
680 accounting in separation logic. In *Proceedings of the 32Nd ACM SIGPLAN-SIGACT Symposium*  
681 *on Principles of Programming Languages*, POPL '05, pages 259–270, New York, NY, USA,  
682 2005. ACM. URL: <http://doi.acm.org/10.1145/1040305.1040327>, doi:10.1145/1040305.  
683 1040327.
- 684 7 Julian Brunner and Peter Lammich. Formal verification of an executable LTL model  
685 checker with partial order reduction. *J. Autom. Reasoning*, 60(1):3–21, 2018. doi:  
686 10.1007/s10817-017-9418-4.
- 687 8 C. Calcagno, P.W. O'Hearn, and Hongseok Yang. Local action and abstract separation logic.  
688 In *LICS 2007*, pages 366–378, July 2007.
- 689 9 Jatin Chhugani, Anthony D Nguyen, Victor W Lee, William Macy, Mostafa Hagog, Yen-Kuang  
690 Chen, Akram Baransi, Sanjeev Kumar, and Pradeep Dubey. Efficient implementation of sorting  
691 on multi-core simd cpu architecture. *Proceedings of the VLDB Endowment*, 1(2):1313–1324,  
692 2008.
- 693 10 Javier Esparza, Peter Lammich, René Neumann, Tobias Nipkow, Alexander Schimpf, and  
694 Jan-Georg Smaus. A fully verified executable LTL model checker. In *CAV*, volume 8044 of  
695 *LNCS*, pages 463–478. Springer, 2013.
- 696 11 Mathias Fleury, Jasmin Christian Blanchette, and Peter Lammich. A verified SAT solver with  
697 watched literals using Imperative HOL. In *Proc. of CPP*, pages 158–171, 2018.
- 698 12 The GNU C++ library 3.4.28. <https://gcc.gnu.org/onlinedocs/libstdc++/>.
- 699 13 A. Nico Habermann. Parallel neighbor-sort, Jun 1972. URL: [https://kilthub.cmu.  
700 edu/articles/journal\\_contribution/Parallel\\_neighbor-sort\\_or\\_the\\_glory\\_of\\_the\\_  
701 induction\\_principle\\_/6608258/1](https://kilthub.cmu.edu/articles/journal_contribution/Parallel_neighbor-sort_or_the_glory_of_the_induction_principle_/6608258/1), doi:10.1184/R1/6608258.v1.
- 702 14 Maximilian P. L. Haslbeck and Peter Lammich. For a few dollars more - verified fine-grained  
703 algorithm analysis down to LLVM. *TOPLAS, S.I. ESOP'21*. to appear.
- 704 15 Maximilian P. L. Haslbeck and Peter Lammich. For a few dollars more - verified fine-grained  
705 algorithm analysis down to LLVM. In Nobuko Yoshida, editor, *Programming Languages*  
706 *and Systems - 30th European Symposium on Programming, ESOP 2021, Held as Part of the*  
707 *European Joint Conferences on Theory and Practice of Software, ETAPS 2021, Luxembourg*  
708 *City, Luxembourg, March 27 - April 1, 2021, Proceedings*, volume 12648 of *Lecture Notes in*  
709 *Computer Science*, pages 292–319. Springer, 2021. doi:10.1007/978-3-030-72019-3\_11.
- 710 16 Jonas Kastberg Hinrichsen, Jesper Bengtson, and Robbert Krebbers. Actris: Session-type  
711 based reasoning in separation logic. *Proc. ACM Program. Lang.*, 4(POPL), dec 2019. doi:  
712 10.1145/3371074.
- 713 17 Nicolai M. Josuttis. *The C++ Standard Library: A Tutorial and Reference*. Addison-Wesley  
714 Professional, 2nd edition, 2012.

- 715 18 Ralf Jung, Robbert Krebbers, Jacques-Henri Jourdan, Ales Bizjak, Lars Birkedal, and Derek  
716 Dreyer. Iris from the ground up: A modular foundation for higher-order concurrent separation  
717 logic. *J. Funct. Program.*, 28:e20, 2018. doi:10.1017/S0956796818000151.
- 718 19 Florian Kammüller, Markus Wenzel, and Lawrence C. Paulson. Locales a sectioning concept  
719 for isabelle. In Yves Bertot, Gilles Dowek, Laurent Théry, André Hirschowitz, and Christine  
720 Paulin, editors, *Theorem Proving in Higher Order Logics*, pages 149–165, Berlin, Heidelberg,  
721 1999. Springer Berlin Heidelberg.
- 722 20 Gerwin Klein, Rafal Kolanski, and Andrew Boyton. Mechanised separation algebra. In *ITP*,  
723 pages 332–337. Springer, Aug 2012.
- 724 21 Peter Lammich. Automatic data refinement. In *ITP*, volume 7998 of *LNCS*, pages 84–99.  
725 Springer, 2013.
- 726 22 Peter Lammich. Verified efficient implementation of gabow’s strongly connected component  
727 algorithm. In *International Conference on Interactive Theorem Proving*, pages 325–340.  
728 Springer, 2014.
- 729 23 Peter Lammich. Refinement to Imperative/HOL. In *ITP*, volume 9236 of *LNCS*, pages 253–269.  
730 Springer, 2015.
- 731 24 Peter Lammich. Efficient verified (UN)SAT certificate checking. In *Proc. of CADE*. Springer,  
732 2017.
- 733 25 Peter Lammich. The GRAT tool chain - efficient (UN)SAT certificate checking with formal  
734 correctness guarantees. In *SAT*, pages 457–463, 2017.
- 735 26 Peter Lammich. Generating Verified LLVM from Isabelle/HOL. In John Harrison, John  
736 O’Leary, and Andrew Tolmach, editors, *10th International Conference on Interactive Theorem  
737 Proving (ITP 2019)*, volume 141 of *Leibniz International Proceedings in Informatics (LIPIcs)*,  
738 pages 22:1–22:19, Dagstuhl, Germany, 2019. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.  
739 URL: <http://drops.dagstuhl.de/opus/volltexte/2019/11077>, doi:10.4230/LIPIcs.ITP.  
740 2019.22.
- 741 27 Peter Lammich. Efficient verified implementation of introsort and pdqsort. In Nicolas  
742 Peltier and Viorica Sofronie-Stokkermans, editors, *Automated Reasoning - 10th International  
743 Joint Conference, IJCAR 2020, Paris, France, July 1-4, 2020, Proceedings, Part II*, volume  
744 12167 of *Lecture Notes in Computer Science*, pages 307–323. Springer, 2020. doi:10.1007/  
745 978-3-030-51054-1\\_18.
- 746 28 Peter Lammich and S. Reza Sefidgar. Formalizing the Edmonds-Karp algorithm. In *Proc. of  
747 ITP*, pages 219–234, 2016.
- 748 29 Peter Lammich and S. Reza Sefidgar. Formalizing network flow algorithms: A refine-  
749 ment approach in Isabelle/HOL. *J. Autom. Reasoning*, 62(2):261–280, 2019. doi:10.1007/  
750 s10817-017-9442-4.
- 751 30 Peter Lammich and Thomas Tuerk. Applying data refinement for monadic programs to  
752 Hopcroft’s algorithm. In Lennart Beringer and Amy P. Felty, editors, *ITP 2012*, volume 7406  
753 of *LNCS*, pages 166–182. Springer, 2012.
- 754 31 Glen Mével and Jacques-Henri Jourdan. Formal verification of a concurrent bounded queue in  
755 a weak memory model. *Proc. ACM Program. Lang.*, 5(ICFP), August 2021. doi:10.1145/  
756 3473571.
- 757 32 DAVID R. MUSSER. Introspective sorting and selection algorithms. *Software: Practice  
758 and Experience*, 27(8):983–993, 1997. doi:10.1002/(SICI)1097-024X(199708)27:8<983::  
759 AID-SPE117>3.0.CO;2-\#.
- 760 33 Peter W. O’Hearn. Resources, concurrency and local reasoning. In Philippa Gardner and  
761 Nobuko Yoshida, editors, *CONCUR 2004 - Concurrency Theory*, pages 49–67, Berlin, Heidel-  
762 berg, 2004. Springer Berlin Heidelberg.
- 763 34 Mohsen Safari and Marieke Huisman. A generic approach to the verification of the permutation  
764 property of sequential and parallel swap-based sorting algorithms. In *International Conference  
765 on Integrated Formal Methods*, pages 257–275. Springer, 2020.

## 12:20 Refinement of Parallel Algorithms down to LLVM

- 766 **35** Simon Spies, Lennard Gäher, Daniel Gratzer, Joseph Tassarotti, Robbert Krebbers, Derek  
767 Dreyer, and Lars Birkedal. Transfinite iris: Resolving an existential dilemma of step-indexed  
768 separation logic. In *Proceedings of the 42nd ACM SIGPLAN International Conference on*  
769 *Programming Language Design and Implementation*, pages 80–95, 2021.
- 770 **36** Intel oneapi threading building blocks. <https://software.intel.com/en-us/intel-tbb>.
- 771 **37** Verified software toolchain project web page. <https://vst.cs.princeton.edu/>.
- 772 **38** Simon Wimmer and Peter Lammich. Verified model checking of timed automata. In *TACAS*  
773 *2018*, pages 61–78, 2018.