Ribbon Proofs for Separation Logic

John Wickerson
University of Cambridge
john.wickerson@cl.cam.ac.uk

Mike Dodds
University of Cambridge
mike.dodds@cl.cam.ac.uk

Matthew Parkinson
Microsoft Research Cambridge
mattpark@microsoft.com

A program proof should not merely certify *that* a program is correct; it should explain *why* it is correct. A proof should be more than 'true': it should be informative, and it should be intelligible. Extending work by Bean [1], we introduce a system that produces readable program proofs that are highly scalable and easily modified.

The de facto standard for presenting program proofs in Hoare logic [2] is the *proof outline*, in which the program's instructions are interspersed with 'enough' assertions to allow the reader to reconstruct the derivation tree. As an example, Fig. 1a presents a proof outline for a program that performs in-place list reversal. A key asset of the proof outline is what we shall call *instruction locality*: that one can verify each instruction in isolation (by confirming that the assertions immediately above and below it form a valid Hoare triple) and immediately deduce that the entire proof is correct.

The proof outline suffers several drawbacks, however. First, there is much repetition: '$\mathit{list}\,\alpha\,\mathrm{x}$' appears redundantly in six consecutive assertions before it is used on line 25. Second, there is no distinction between those parts of an assertion that are affected by an instruction and those that are merely in what separation logic calls the *frame*. For instance, line 19 affects only the second and fourth conjuncts of its preceding assertion, but it is difficult to deduce its effect because two unchanged conjuncts are interspersed. (Had we followed common practice and reduced the size of the proof outline by combining this line with the assignment on line 17, the effect would be even harder to deduce.) Third, the use of logical variables is unclear. For instance, spotting that the β in line 20 differs from the one in line 22 requires careful examination, or else, as we have done, an explicit textual comment. These minor problems in our illustrative example quickly become devastating when scaled to large programs.

Separation logic [3], [4] provides a mechanism for handling a second dimension of locality: *resource locality*. One can use separation logic's *small axioms* to reason about an instruction operating only on the resources (i.e. memory cells) that it needs, and immediately deduce its effect on the entire state using the *frame rule*. To depict this mechanism in a proof outline, one must show applications of the frame rule explicitly. But this is tedious; moreover, it is difficult to know when and what to frame. Meanwhile, the ribbon proof inherently supports resource locality. Its primitive steps correspond exactly to the small axioms. It is thus an ideal representation for exploiting both forms of locality that separation logic provides.

Figure 1b recasts our proof as a ribbon proof. The state is distributed across several *ribbons* (thick borders). Horizontally separated ribbons describe disjoint parts of the state. The instructions are in grey bars, and the scope of each logical variable is delimited by an *existential box* (thin borders). We are free to stretch ribbons as required by the layout, and, because $*$ is commutative, we can 'twist' them too. A temporarily inactive ribbon slides discreetly down the side of the proof. This effect is achieved by invoking the frame rule at each instruction; but crucially in a ribbon proof, these invocations are implicit and do not increase the diagram's complexity. Observe that the repetition has disappeared, and that each instruction's effect is clear: it affects exactly those assertions directly above and below it, while framed assertions (which must not mention variables written by the instruction) bypass to the left or right. Existential boxes extend vertically to indicate the range of steps over which the same witness is used, thus making the usage of logical variables visually clear.

In our full paper [5]:

- we present an Isabelle-checked graph-based formalisation of our proof system;
- we showcase, with a ribbon proof of the memory manager from Version 7 Unix, the ability of our diagrams to present readable proofs of large, complex programs; and
- we describe a prototype tool for mechanically checking ribbon proofs in Isabelle. Provided with a small proof script for each primitive step, our tool assembles a script that verifies the entire diagram. The tool handles tediums such as the associativity and commutativity of $*$ automatically, leaving the user to concentrate on the interesting parts of the proof.

This work lays the foundations for a new way to use logic to understand programs. Where a proof outline essentially flattens a proof to a list of assertions, our system produces geometric objects that illuminate the structure of proofs, and which can be navigated, modified and simplified by leveraging human visual intuition.

References

[1] J. Bean, "Ribbon proofs," in *MFPS*, 2003.
[2] C. Hoare, "An axiomatic basis for computer programming," *Communications of the ACM*, vol. 12, no. 10, October 1969.
[3] J. C. Reynolds, "Separation logic: A logic for shared mutable data structures," in *LICS*, 2002.
[4] S. Ishtiaq and P. W. O'Hearn, "BI as an assertion language for mutable data structures," in *POPL*, 2001.
[5] J. Wickerson, M. Dodds, and M. J. Parkinson, "Ribbon proofs for separation logic," May 2012, http://www.cl.cam.ac.uk/~jpw48/ribbons.html.

1 $\left\{ list\ \delta\ \mathrm{x} \right\}$

2 `y:=0;`

3 $\left\{ list\ \delta\ \mathrm{x} * list\ \epsilon\ \mathrm{y} \right\}$

4 `// Choose` $\alpha := \delta$ `and` $\beta := \epsilon$

5 `while` $\left\{ \exists \alpha, \beta.\ list\ \alpha\ \mathrm{x} * list\ \beta\ \mathrm{y} * \delta \doteq \beta^\dagger \cdot \alpha \right\}$

6 `(x!=0) {`

7 $\left\{ \mathrm{x} \neq \dot{0} * (\exists \alpha, \beta.\ list\ \alpha\ \mathrm{x} * list\ \beta\ \mathrm{y} * \delta \doteq \beta^\dagger \cdot \alpha) \right\}$

8 $\left\{ \exists \alpha, \beta.\ \mathrm{x} \neq \dot{0} * list\ \alpha\ \mathrm{x} * list\ \beta\ \mathrm{y} * \delta \doteq \beta^\dagger \cdot \alpha \right\}$

9 `// Unfold` *list* `def`

10 $\left\{ \begin{array}{l} \exists \alpha, \beta.\ (\exists \alpha', i, Z.\ \mathrm{x} \mapsto i, Z * list\ \alpha'\ \mathrm{z} * \alpha \doteq i \cdot \alpha') \\ * list\ \beta\ \mathrm{y} * \delta \doteq \beta^\dagger \cdot \alpha \end{array} \right\}$

11 `// Choose` $\alpha := \alpha'$

12 $\left\{ \begin{array}{l} \exists \alpha, \beta, i, Z.\ \mathrm{x} \mapsto i, Z * list\ \alpha\ Z * \delta \doteq \beta^\dagger \cdot (i \cdot \alpha) \\ * list\ \beta\ \mathrm{y} \end{array} \right\}$

13 `z:=[x+1];`

14 $\left\{ \exists \alpha, \beta, i.\ list\ \alpha\ \mathrm{z} * \mathrm{x} \mapsto i, \mathrm{z} * \delta \doteq \beta^\dagger \cdot (i \cdot \alpha) * list\ \beta\ \mathrm{y} \right\}$

15 `// Reassociate` i

16 $\left\{ \exists \alpha, \beta, i.\ list\ \alpha\ \mathrm{z} * \mathrm{x} \mapsto i, \mathrm{z} * \delta \doteq (i \cdot \beta)^\dagger \cdot \alpha * list\ \beta\ \mathrm{y} \right\}$

17 `[x+1]:=y;`

18 $\left\{ \exists \alpha, \beta, i.\ list\ \alpha\ \mathrm{z} * \mathrm{x} \mapsto i, \mathrm{y} * \delta \doteq (i \cdot \beta)^\dagger \cdot \alpha * list\ \beta\ \mathrm{y} \right\}$

19 `// Fold` *list* `def`

20 $\left\{ \exists \alpha, \beta, i.\ list\ \alpha\ \mathrm{z} * list\ (i \cdot \beta)\ \mathrm{x} * \delta \doteq (i \cdot \beta)^\dagger \cdot \alpha \right\}$

21 `// Choose` $\beta := (i \cdot \beta)$

22 $\left\{ \exists \alpha, \beta.\ list\ \alpha\ \mathrm{z} * list\ \beta\ \mathrm{x} * \delta \doteq \beta^\dagger \cdot \alpha \right\}$

23 `y:=x;`

24 $\left\{ \exists \alpha, \beta.\ list\ \alpha\ \mathrm{z} * list\ \beta\ \mathrm{y} * \delta \doteq \beta^\dagger \cdot \alpha \right\}$

25 `x:=z;`

26 $\left\{ \exists \alpha, \beta.\ list\ \alpha\ \mathrm{x} * list\ \beta\ \mathrm{y} * \delta \doteq \beta^\dagger \cdot \alpha \right\}$

27 `}`

28 $\left\{ \mathrm{x} \doteq 0 * (\exists \alpha, \beta.\ list\ \alpha\ \mathrm{x} * list\ \beta\ \mathrm{y} * \delta \doteq \beta^\dagger \cdot \alpha) \right\}$

29 $\left\{ \exists \alpha, \beta.\ \mathrm{x} \doteq 0 * list\ \alpha\ \mathrm{x} * list\ \beta\ \mathrm{y} * \delta \doteq \beta^\dagger \cdot \alpha \right\}$

30 `// Unfold` *list* `def`

31 $\left\{ \exists \alpha, \beta.\ \alpha \doteq \epsilon * list\ \beta\ \mathrm{y} * \delta \doteq \beta^\dagger \cdot \alpha \right\}$

32 `// Concatenate empty sequence`

33 $\left\{ \exists \beta.\ list\ \beta\ \mathrm{y} * \delta \doteq \beta^\dagger \right\}$

34 `// Fold` *list* `def`

35 $\left\{ list\ \delta^\dagger\ \mathrm{y} \right\}$

(a) A proof outline

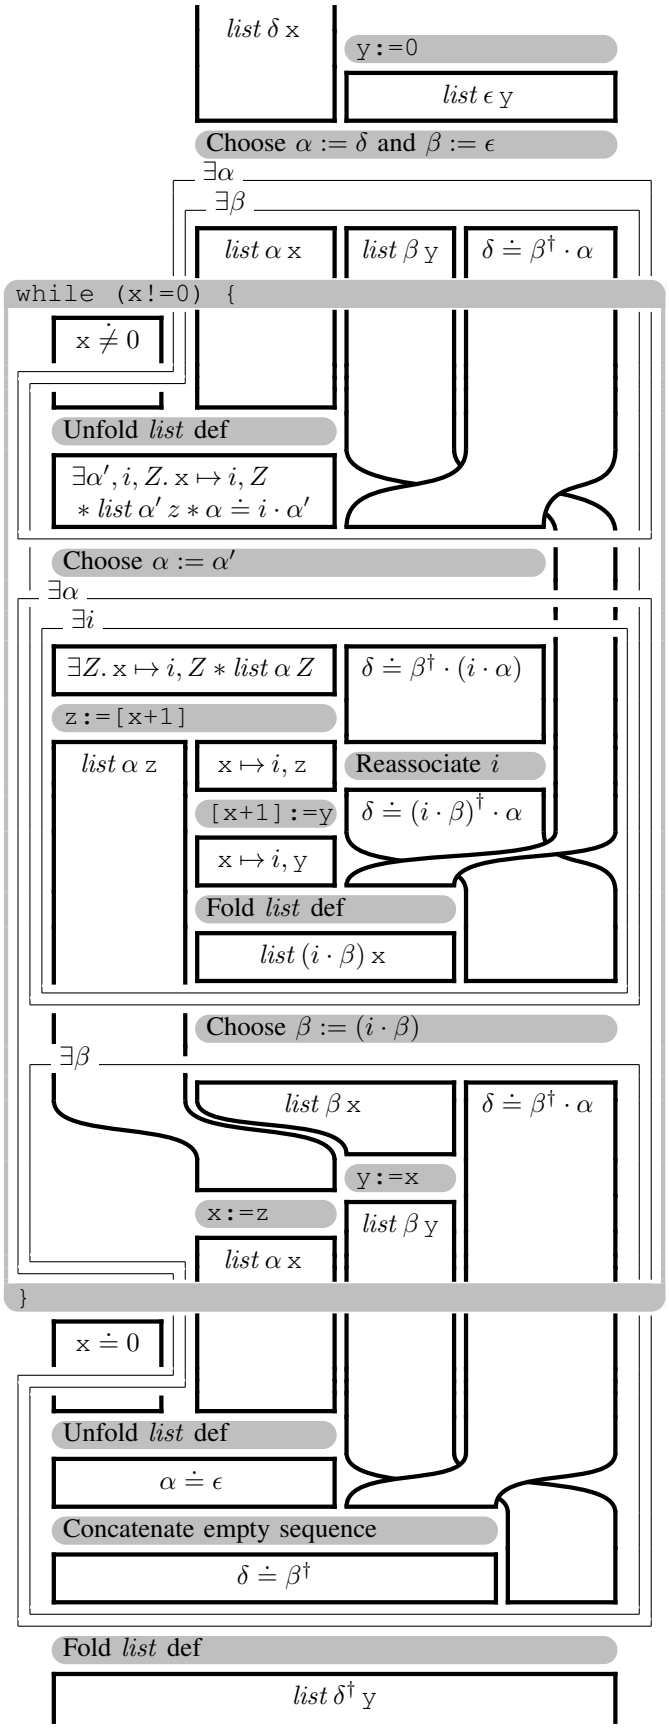

(b) A ribbon proof

Fig. 1. Two proofs of list reverse. For a binary relation r, we write $x\ \dot{r}\ y$ for $x\ r\ y \wedge emp$. We write $\cdot$ for sequence concatenation, $(-)^\dagger$ for sequence reversal and ϵ for the empty sequence, and define *list* as the smallest predicate satisfying $list\ \alpha\ x \Leftrightarrow x \doteq 0 * \alpha \doteq \epsilon \vee x \neq 0 * \exists \alpha', i, x'.\ x \mapsto i, x' * \alpha \doteq i \cdot \alpha' * list\ \alpha'\ x'$.