

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Splitting Recursion Schemes into Reversible and Classical Interacting Threads

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1794961> since 2021-11-03T22:24:16Z

Publisher:

Springer International Publishing

Published version:

DOI:10.1007/978-3-030-79837-6_12

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

Splitting recursion schemes into reversible and classical interacting threads

Armando B. Matos¹, Luca Paolini²^[0000-0002-4126-0170], and Luca Roversi²^[0000-0002-1871-6109]

¹ Universidade do Porto, Departamento de Ciência de Computadores – Portugal
armandobcm@yahoo.com

² Università degli Studi di Torino, Dipartimento di Informatica – Italy
{luca.paolini,luca.roversi}@unito.it

Abstract. Given a simple recursive function, we show how to extract from it a reversible and an classical iterative part. Those parts can synchronously cooperate under a Producer/Consumer pattern in order to implement the original recursive function. The reversible producer is meant to run on reversible hardware. We also discuss how to extend the extraction to a more general compilation scheme.

1 Introduction

Our goal is to compile a class of recursive functions in a way that parts of the object code produced can leverage the promised green foot-print of truly reversible hardware. This work illustrates preliminary steps towards that goal. We focus on a basic class of recursive functions in order to demonstrate its feasibility.

Contributions. Let $\text{recF}[p,b,h]$ be a recursive function defined in some programming formalism, where p is a *predecessor* function, h a *step* function, and b a *base* function. We show how to compile $\text{recF}[p,b,h]$ into $\text{itFClS}[b,h]$ and $\text{itFRev}[p,pInv]$ such that:

$$\text{recF}[p,b,h] \simeq \text{itFClS}[b,h] \parallel \text{itFRev}[p,pInv] \quad , \quad (1)$$

where: (i) “ \simeq ” stands for “*equivalent to*”; (ii) $\text{itFClS}[b,h]$ is a classical **for**-loop that, starting from a value produced by b , iteratively applies h ; (iii) $\text{itFRev}[p,pInv]$ is a reversible code with two **for**-loops in it one iterating p , the other its inverse $pInv$; (iv) “ \parallel ” is interpreted as an *interaction* between $\text{itFClS}[b,h]$ and $\text{itFRev}[p,pInv]$, according to a Producer/Consumer pattern, where $\text{itFRev}[p,pInv]$ produces the values that $\text{itFClS}[b,h]$ consumes to implement the initially given recursion $\text{recF}[p,b,h]$. In principle, $\text{itFRev}[p,pInv]$ can drive a real reversible hardware to exploit its low energy consumption features.

In this work we limit the compilation scheme (1) to use: (i) a predecessor p such that the value $p(x)-x$ is any *constant* Δ_p equal to, or smaller than, -1 ; (ii) recursion functions $\text{recF}[p,b,h]$ whose *condition* identifying the base case

is $x \leq 0$ instead than the more standard $x == 0$; this means that more than one base *non positive* value for $\text{recF}[p, b, h]$ exists in the interval $[\Delta_p + 1, 0]$. This slight generalization will require a careful management of the reversible behavior of $\text{itFRev}[p, p\text{Inv}]$ and its interaction with $\text{itFCls}[b, h]$ in order to reconstruct $\text{recF}[p, b, h]$.

Contents. Section 2 sets the stage to develop the main ideas about (1), restricting $\text{recF}[p, b, h]$ to a recursive function that identifies its base case by means of the standard condition $x == 0$; this ease the description of how $\text{itFRev}[p, p\text{Inv}]$ and $\text{itFCls}[b, h]$ interact. Section 3 extends (1) to deal with $\text{recF}[p, b, h]$ having $x \leq 0$, and not $x == 0$, to identify its base case(s); this impacts on how $\text{itFRev}[p, p\text{Inv}]$ must work. In both cases, the programming syntax we use can be interpreted into the reversible languages SRL [3,4] and RPP [5,6,4], up to minor syntactic details. Section 4 addresses future work.

```

1  Fix recF(x) {
2    if (c(x)) { b(x); }
3    else { h(x, recF(p(x))); } }

```

Fig. 1. The recursive function recF .

```

1  /** Assumption: the initial value of x is 3 */
2  x = p(x) // ==2
3  x = p(x) // ==1
4  x = p(x) // ==0
5  y = b(x) // ==b(p(p(p(3))))
6  y = h(x,y) // ==h(p(p(p(3))), b(p(p(p(3))))
7  x = pInv(x) // ==pInv(p(p(p(3))))==p(3)
8  y = h(x,y) // ==h(p(p(3)), h(p(p(p(3))), b(p(p(p(3))))))
9  x = pInv(x) // ==pInv(p(p(3)))==p(3)
10 y = h(x,y) // ==h(p(3), h(p(p(3)))
11 // , h(p(p(p(3))), b(p(p(p(3))))))
12 x = pInv(x) // ==pInv(p(3))==3
13 y = h(x,y) // ==h(3, h(p(3), h(p(p(3)))
14 // , h(p(p(p(3))), b(p(p(p(3))))))

```

Fig. 2. Iterative unfolding $\text{recF}(3)$: the bottom-up part.

2 The driving idea

Let `recF[p,b,h]` in (1) have a structure as in **Fig. 1** where `b(x)` is the *base* function, `h(x,y)` the *step* function, `p(x)` the *predecessor* `x-1`, and `c(x)` the *condition* `x==0` to identify a unique base case.

Fig. 2 details out `h(3,h(p(3),h(p(p(3))), h(p(p(p(3))),b(p(p(p(3))))))`, unfolding of `recF(3)`. Every comment asserts a property of the values that `x` or `y` stores. Lines 2–4 unfold an iteration that computes `p(p(p(3)))`, which eventually sets the value of `x` to 0. Line 5 starts the construction of the final value of `recF(3)` by applying the base case of `recF`, i.e. `b(x)`. By definition, let `pInv` denote the inverse of `p`, i.e. `pInv(p(z))==p(pInv(z))==z`, for any `z`. Clearly, in our running example, the function `pInv(x)` is `x+1`. Lines 6–13 alternate `h(x,y)`, whose result `y`, step by step, gets closer to the final value `recF(3)`, and `pInv(x)`, which produces a new value for `x`.

```

1  s = 0, e = 0, g = 0, w = 0
2  w = w + x;
3  for (i = 0; i<=w; i++)      {
4      if      (x> 0) { g++; }
5      else if (x==0) { e++; }
6      else          { s++; }
7      x = p(x);              }
8
9  for (i = 0; i<=w; i++)      {
10     x = pInv(x);
11     if      (x> 0) { g--; y = h(x,y); }
12     else if (x==0) { e--; y = b(x);   }
13     else          { s--;              } }
14  w = w - x;

```

Fig. 3. Iterative `itF` equivalent to `recF`.

Let us call `itF` the code in **Fig. 3**. It implements `recF` by means of finite iterations only. Continuing with our running example, if we run `itF` here above starting with `x==3`, then `x==0` holds at line 8, just after the first `for`-loop; after the second `for`-loop `y==recF(3)` holds at line 14.

The code of `itF` has two parts. Through lines 2–7 the variable `g` counts how many times `x` remains positive, the variable `e` how many it stays equal to 0, and the variable `s` how many it becomes negative. In this running example we notice that `x` never becomes negative, for the iteration at lines 3–7 is driven by the value of `x` which, initially, we can assume non negative, and which `p(x)` decreases of a single unity. We shall clarify the role of `s` later. Lines 9–13 undo what lines 2–7 do by executing `pInv(x)`, `g--`, `e--`, `s--`, i.e. the inverses, in reversed order, of `p(x)`, `g++`, `e++`, `s++`. So the correct values of `x` are available at lines 12, and

11, ready to be used as arguments of $b(x)$ and $h(x,y)$ to update y as in **Fig. 3**, according to the results we obtain by the recursive calls to `recF`.

```

1  s = 0, e = 0, g = 0, w = 0
2  w = w + x;
3  for (i=0; i<=w; i++)      {
4      if      (x> 0) { g++; } //number of times x is 'g'reater than 0
5      else if (x==0) { e++; } //number of times x is 'e'qual to 0
6      else      { s++; } //number of times x is 's'maller than 0
7      x = p(x);              }
8
9  for (i=0; i<=w; i++)      {
10     x = pInv(x);
11     if      (x> 0) { g--; /* Value of x for h availabe here */ }
12     else if (x==0) { e--; /* Value of x for b availabe here */ }
13     else      { s--; }
14     w = w - x;

```

Fig. 4. Reversible side of `itF`.

Now, let us focus on the main difference between **Fig. 4** and **Fig. 3**.

Both $x=b(x)$ and $y=h(x,y)$ at lines 12, and 11 of **Fig. 3** are missing from lines 12, and 11 of **Fig. 4**. Dropping them let **Fig. 4** be the *reversible side* of `itF`; calling $b(x)$ and $h(x,y)$ in it generates y , which is the result we need, so preventing the possibility to reset the value of every variable dealt with in **Fig. 4** to their initial value. This is why we also need a *classical side* of `itF` that generates y in collaboration with the *reversible side* in order to implement the initial `recF` correctly.

```

1  /** Assumption. The value of the input x is available here */
2  /* Inject the current x at line 2 of itFRev to let it start */
3  iterations = /* Probe line 9 of itFRev to get the
4               number of iterations to execute */
5  y = b(/* Probe line 14 of itFRev to get the argument */);
6  for (i = 0; i<iterations; i++)      {
7      y = h(/* Probe line 12 itFRev to get
8            the first argument of h    */ , y); }

```

Fig. 5. Classical side of `itF`: the consumer `itFCls`.

The previous observations lead to **Fig. 5** which defines the *classical side* `itFCls` of `recF`, and to **Fig. 6** which defines the *reversible side* `itFRev` of `recF`.

```

1  s = 0, e = 0, g = 0, w = 0;
2  x = /* Inject here the value of x at line 2 of itFClS */
3  w = w + x;
4  for (i = 0; i<=w; i++)      {
5      if      (x> 0) { g++; }
6      else if (x==0) { e++; }
7      else           { s++; }
8      x = p(x);           }
9  /* itFClS probes here g which has the number of iterations */
10 for (i = 0; i<=w; i++)      {
11     x = pInv(x);
12     if      (x> 0) { g--; /* itFClS probes here the
13                          first argument value of h */ }
14     else if (x==0) { e--; /* itFClS probes here the
15                          argument value of b      */ }
16     else           { s--;           } }
17     w = w - x;

```

Fig. 6. Reversible side of `itF` updated to be the producer `itFRev` of the values that the consumer `itFClS` needs.

So, here below we can illustrate how `itFClS` and `itFRev` synchronously interact, `itFRev` producing values, `itFClS` consuming them as arguments of `b(x)` and `h(x,y)`.

Line 2 of `itFClS` is the starting point of the synchronous interaction between `itFClS` and `itFRev`; its comment:

```
/* Inject the current x at line 2 of itFRev to let it start */
```

describes what, in a fully implemented version of `itFClS`, we expect in that line of code. The comment says that `itFClS` injects (sends, puts) its input value `x` to line 2 of the *reversible side* `itFRev` (cf. **Fig. 6**). Once `itFRev` obtains that value at line 2, as outlined by:

```
/* Inject here the value of x from line 2 of itFClS */
```

its `for`-loop at lines 4–8 executes.

After line 2, `itFClS` stops at line 3. It waits for `itFRev` to produce the number of times that `itFClS` has to iterate line 7. Accordingly to:

```
/* Probe line 9 of itFRev to get the number of iterations to execute */
```

`itFRev` makes that value available in its variable `g` at line 9:

```
/* itFClS probes here g which has the number of iterations */ .
```

Once gotten the value in `iterations`, `itFClS` proceeds to line 5 and stops, waiting for `itFRev` to produce the argument of `b` which is eventually available for probing at line 14 of `itFRev`.

Once the argument becomes available `b` is applied, and `itFClS` enters its `for`-loop, stopping at line 7 at every iteration. The reason is that `itFClS` waits for line 12 in `itFRev` to produce the value of the first argument of `h(x,y)`. This interleaved dialog between line 7 of `itFClS` and line 12 of `itFRev` lasts `iterations` times.

```

1  Fix recG(x)                                {
2    if (x<=0) { b(x);                        }
3    else     { h(x,recG(p(x))); } }

```

Fig. 7. The generic structure of `recG`.

3 From recursion to iteration

We now generalize what we have seen in Section 2. Inside (1) we use `recG` of Fig. 7 instead than `recF` of Fig. 1. This requires to generalize Fig. 6.

From the introduction we recall that, given a *predecessor* `p(x)`, we define $\Delta_p = p(x) - x$, which is a negative value. In this section Δ_p can be any *constant* $k < -1$, not only $k == -1$; this requires to consider the slightly more general *condition* `x <= 0` in `recG`. For example, let `p(x)` be `x-2`. The computation of `recG(3)` is `h(3,h(p(3),h(p(p(3)),b(p(p(3))))))` which looks for the least n of iterated applications of `p(x)` such that `p(...p(3)...) <= 0`; in our case we have $2 == n < 3$.

Fig. 8 introduces `itG` which generalizes `itF` in Fig. 3.

The scheme `itG` iteratively implements any recursive function whose structure can be brought back to the one of `recG`. We remark that line 1 in Fig. 8 initializes ancillae `s`, `e`, `g`, and `w`, like Fig. 3 initializes the namesake variables of `itF`, but line 2 of `itG` has new ancillae `z`, `predDivX`, and `predNotDivX`.

We also assume an initial *non negative* value for `x`. The reason is twofold. Firstly, it keeps our discussion as simple as possible, with no need to use the absolute value of `x` to set the upper limit of every index `i` in the `for`-loops that occur in the code. Second, negative values of `x` would widen our discussion about what a classical recursive function on negative values is and about what its reversible equivalent iteration has to be; we see this as a very interesting subject connected to [1], which is much more oriented than us to optimization issues of recursively defined functions.

We start observing that line 3 of `itG` sets `w` to the initial value of `x`; the reason is that every `for`-loop, but the one at lines 10–12, has to last `x+1` iterations, and `x` changes in the course of the computation; so, `w` stores the initial value of `x` and stays constant from line 4 through line 21. In fact it can change at lines 22–33. We will see why, but `w` is eventually reset to its initial value `0` at line 36.

```

1  s = 0, e = 0, g = 0, w = 0;
2  z = 0, predDivX = 0, predNotDivX = 1;
3  w = w + x; /* x is assumed to be the input */
4  for (i = 0; i <= w; i++) {
5      if (x > 0) { g++; }
6      else if (x == 0) { e++; }
7      else { s++; }
8      x = p(x);
9
10 for (i = 0; i < e; i++) {
11     predDivX = predDivX + predNotDivX;
12     predNotDivX = predDivX - predNotDivX; }
13
14 for (j = 0; j < predDivX; j++) {
15     for (i = 0; i <= w; i++) {
16         x = pInv(x);
17         if (x > 0) { g--; y = h(x,y); }
18         else if (x == 0) { e--; y = b(x); }
19         else { s--; }}}
20
21 for (j = 0; j < predNotDivX; j++) {
22     w++;
23     for (i = 0; i <= w; i++) {
24         x = pInv(x);
25         if (x > 0) { g--;
26                     x = p(x);
27                     if (z < 0) { }
28                     else if (z == 0) { y = b(x); z++; }
29                     else { y = h(x,y); }
30                     x = pInv(x); }
31         else if (x == 0) { e--; }
32         else { s--; }}
33     w--;
34     for (i = 0; i < predNotDivX; i++) {
35         z--;
36     }
37     w = w - x;
38     /* y carries the output */

```

Fig. 8. The iterative function `itG`.

With the here above assumptions, given a non negative x , and in analogy to `itF`, the `for`-loop at lines 4–8 of `itG` iterates the application of $p(x)$ as many times as $w+1$, i.e. the initial value of x plus 1. So, the value of x at line 9 is equal to $w+(w+1)*\Delta_p$ which cannot be positive. In particular, all the values that x assumes in the `for`-loop at lines 4–8 belong to the following interval:

$$I(w) \triangleq [w+(w+1)*\Delta_p, w+w*\Delta_p, \dots, w+\Delta_p, w] \quad (2)$$

from the least to the greatest; the counters g , e , s say how many elements of $I(x)$ are greater, equal or smaller than 0, respectively. Depending on 0 to belong to $I(x)$ determines the behavior of the reminder part of `itG`, i.e. lines 10–36.

We distinguish two cases in order to illustrate them.

First case. Let $w\% \Delta_p == 0$, i.e. the integer value Δ_p divides with no remainder the initial value of x that we find in w . So, $0 \in I(x)$, which implies the following relations hold at line 9:

$$e == 1 \quad g == -\frac{w}{\Delta_p} \quad s == (w+1)-g-e . \quad (3)$$

```

1  if      (e < 0) {
2  else if (e == 0) { predDivX = predDivX+predNotDivX;
3                      predNotDivX = predDivX - predNotDivX; }
4  else      {

```

Fig. 9. A possible replacement of lines 10–12 in **Fig. 8**.

Lines 10–12 execute exactly once, swapping `predDivX` and `predNotDivX`. As a remark, we could have well used the `if`-selection in **Fig. 9** (a construct of RPP) in place of the `for`-loop at lines 10–12, but we opt for a more compact code.

Swapping `predDivX` and `predNotDivX` sets `predDivX==1` and `predNotDivX==0`, computationally exploiting that Δ_p divides w with no remainder: the `for`-loop body at lines 15–19 becomes accessible, while lines 22–33, with `for`-loops among them, do not. Lines 15–19 are identical to lines 10–16 of `itF` in **Fig. 4** which we already know to correctly apply $b(x)$ and $h(x,y)$ in order to simulate the recursive function we start from.

As a second case. Let $w\% \Delta_p != 0$, i.e. the integer value Δ_p divides the initial value of x that w stores, *but with some remainder*. So, $0 \notin I(x)$, which imply:

$$e == 0 \quad g == -\left\lfloor \frac{w}{\Delta_p} \right\rfloor \quad s == (w+1)-g-e \quad (4)$$

hold at line 9. Lines 11–12 cannot execute, leaving `predDivX` and `predNotDivX` as they are: lines 22–33 become accessible and the `for`-loop at lines 15–19 does not. Line 22 increments `w` to balance the information loss that the rounding of `g` in (4) introduces; line 33 recovers the value of `w` when the outer `for`-loop starts. The `if`-selection at lines 25–32 identifies when to apply `b(x)`, which must be followed by the required applications of `h(x,y)`. We know that $0 \notin I(x)$, so `x==0` can never hold. Clearly, `s--` is executed until `x>0`. But the *first* time `x>0` holds true we must compute `b(p(x))`, because the *base* function `b(x)` *must be used the last time* `x` assumes a negative value, *not the first time* it gets positive; lines 26–30 implement our needs. Whenever `x>0` is true, the value of `x` is one step ahead the required one: we get one step back with line 26 and, if it is the first time we step back, i.e. `z==0` holds, then we must execute line 28. If not, i.e. `z!=0`, we must apply the *step* function at line 29. Line 30, restores the right value of `x`. Finally, the `for`-loop at line 34 sets `z` to its initial value.

At this point, in order to obtain the fully reversible version of **Fig. 8** we must think of replacing the calls to `h(x,y)` and `b(x)` at lines in 28, and 29 by means of actions that probe the value of `x`, in analogy to **Fig. 6**, lines 12 and 14. The full details are in [7] which we look as a playground with Java classes that implement **Fig. 8** and **Fig. 5** as synchronous and parallel threads, acting as a producer and a consumer.

4 Future work

We have shown that we can decompose every classical recursive function, based on a *predecessor* that decreases every of its input by a constant value, into reversible and classical components that cooperate to implement the original recursive functions under a Producer/Consumer pattern (see (1)).

Firstly, we plan to extend (1) to recursive functions `recF` based on predecessors `p` not limited to a constant Δ_p not greater than `-1`. A predecessor `p` should be at least such that:

1. Δ_p is not necessarily a constant. For example, $\Delta_p == -3$ on even arguments, and `-2` on odd ones can be useful;
2. the predecessor can be an integer division `x/k`, for some given `k>0`, like in a dichotomic search, which has `k==2`.

Secondly, we aim at generalizing (1) to a compiler $\llbracket \cdot \rrbracket$:

$$\begin{aligned}
 \llbracket p \rrbracket &= \text{some implementation code} \\
 \llbracket pInv \rrbracket &= !\llbracket p \rrbracket, \text{ i.e. implementation that inverts } \llbracket p \rrbracket \\
 \llbracket \text{recF}[p,b,h] \rrbracket &= \text{itFCls}[\llbracket b \rrbracket, \llbracket h \rrbracket] \parallel \text{itFRev}[\llbracket p \rrbracket, \llbracket pInv \rrbracket] .
 \end{aligned} \tag{5}$$

The domain of $\llbracket \cdot \rrbracket$ should be a class `R` of recursive functions built by means of standard composition schemes, starting from a class of predecessors `p1`, `p2`, ... each of which must have the corresponding inverse function `p1Inv`, `p2Inv`, ...

In these lines we want to explore interpretations of \parallel more liberal than the essentially obvious synchronous Producer/Consumer that we implement in [7]. We shall very likely take advantage of parallel discrete events simulators as described in [8,9] in order to get rid of any explicit synchronization between the pairs of reversible-producer/classical-consumer that (5) would recursively generate when applied to an element in R .

We also plan to follow a more abstract line of research. The compilation scheme (5) recalls Girard’s decomposition $A \rightarrow B \simeq !A \multimap B$ of a classical computation into a linear one that can erase/duplicate computational resources. Decomposing $\text{recF}[p, b, h]$ in terms of $\text{itFClS}[b, h]$ and $\text{itFRev}[p, p\text{Inv}]$ suggests that the relation between reversible and classical computations can be formalized by a linear isomorphism $A^n \multimap B^n$ between tensor products A^n , and B^n of A , and B , in analogy to [2]. Then we can think of recovering classical computations by some functor, say γ , whose purpose is, at least, to forget, or to inject replicas, of parts of A^n , and B^n in a way that $(\gamma A^n \rightarrow \gamma A^n) \uplus (\gamma A^n \leftarrow \gamma A^n)$ can be their type. The type says that we move from a reversible computation to a classical one by choosing which is input and which is output, so recovering the freedom to manage computational resources as we are used to when writing classical programs.

References

1. E. A. Boiten. Improving recursive functions by inverting the order of evaluation. *Science of Computer Programming*, 18(2):139 – 179, 1992.
2. R. P. James and A. Sabry. Information effects. In J. Field and M. Hicks, editors, *Proceedings of the 39th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2012, Philadelphia, Pennsylvania, USA, January 22-28, 2012*, pages 73–84. ACM, 2012.
3. A. B. Matos. Linear programs in a simple reversible language. *Theor. Comput. Sci.*, 290(3):2063–2074, 2003.
4. A. B. Matos, L. Paolini, and L. Roversi. On the expressivity of total reversible programming languages. In I. Lanese and M. Rawski, editors, *Reversible Computation*, pages 128–143, Cham, 2020. Springer International Publishing.
5. L. Paolini, M. Piccolo, and L. Roversi. On a class of reversible primitive recursive functions and its turing-complete extensions. *New Generation Computing*, 36(3):233–256, Jul 2018.
6. L. Paolini, M. Piccolo, and L. Roversi. A class of recursive permutations which is primitive recursive complete. *Theor. Comput. Sci.*, 813:218–233, 2020.
7. L. Roversi, A. Matos, and L. Paolini. Eclipse java project rev2iterrev. <https://github.com/LucaRoversi/Rec2IterRev>.
8. M. Schordan, T. Opielstrup, D. R. Jefferson, and P. D. B. Jr. Generation of reversible C++ code for optimistic parallel discrete event simulation. *New Gener. Comput.*, 36(3):257–280, 2018.
9. M. Schordan, T. Opielstrup, M. K. Thomsen, and R. Glück. *Reversible Languages and Incremental State Saving in Optimistic Parallel Discrete Event Simulation*, pages 187–207. Springer International Publishing, Cham, 2020.