



Verified Timing Transformations in Synchronous Circuits with $\lambda\pi$ -Ware

João Paulo Pizani Flor^(✉) and Wouter Swierstra

Utrecht University, Utrecht, The Netherlands
{J.P.PizaniFlor,W.S.Swierstra}@uu.nl

Abstract. We define a DSL for hardware description, called $\lambda\pi$ -Ware, embedded in the dependently-typed language Agda, which makes the DSL well-scoped and well-typed by construction. Other advantages of dependent types are that circuit models can be simulated and verified in the same language, and properties can be proven not only of specific circuits, but of circuit generators describing (infinite) families of circuits. This paper focuses on the relations between circuits computing the same values, but with different levels of statefulness. We define common recursion schemes, in combinational and sequential versions, and express known circuits using these recursion patterns. Finally, we define a notion of convertibility between circuits with different levels of statefulness, and prove the core convertibility property between the combinational and sequential versions of our vector iteration primitive. Circuits defined using the recursion schemes can thus have different architectures with a guarantee of functional equivalence up to timing.

1 Introduction

Modelling electronic circuits has been a fertile ground for functional programming (Sheeran 2005) and theorem proving (Hanna and Daeche 1992). There have been numerous efforts to describe, simulate, and verify circuits using functional languages such as MuFP (Sheeran 1984) and more recently C λ aSH (Baaij 2015) and ForSyDe (Sander and Jantsch 2004).

Functional languages have also been used to *host* an Embedded Domain-Specific Language (EDSL) for hardware description. Some of these EDSLs, such as Wired (Axelsson et al. 2005), capture low-level information about the layout of a circuit; others aim to use the host language to provide a higher-level of abstraction to describe the circuit's intended behaviour. A notable example of the latter approach is Lava (Bjesse et al. 1999) and its several variants (Gill et al. 2009; Singh 2004).

Also interactive theorem proving and programming with dependent types have been fruitfully used to support hardware verification efforts, with some based on HOL (Melham 1993; Boulton et al. 1992), some on Coq (Braibant 2011; Braibant and Chlipala 2013) and some on Martin-Löf Type Theory (Brady et al. 2007)

Following this line of research, we utilize a dependently-typed programming language (Agda) as the *host* of our hardware EDSL, for its proving capabilities and convenience of embedding.

In particular, this paper focuses on verification related to *timing*, that is, the behaviour of a circuit in terms of its inputs *over time*. When designing hardware, a compromise must be made between the *area* occupied by a circuit and the *number of clock cycles* it takes to produce its results.

A *combinational* (stateless) architecture better harnesses potential parallelism but might negatively influence other constraints such as frequency and power consumption. A more *sequential* circuit (stateful), on the other hand, will occupy less area but might be a bottleneck in computational throughput and impact other parts of the design that depend on its outputs.

There are many different ways to implement any specific functional behaviour, and it can be difficult to find the right spot in the design space upfront. Timing-related circuit transformations are quite invasive and error-prone – making it difficult to correct bad design decisions *a posteriori*. With this paper, we attenuate some of these issues by defining a language for circuit description that facilitates the exploration of different points in the timing design space. More concretely, this paper makes the following contributions:

- We show how to embed a typed hardware DSL, $\lambda\pi$ -Ware, in the general purpose dependently typed programming language Agda (Sect. 3), together with an executable semantics based on state transitions (Sect. 4).
- Next, we define common recursion patterns to build circuits in both combinational and sequential architectures (Sect. 5). We show how some well-known circuits can be expressed in terms of these recursion patterns.
- Finally, we define a precise relation between the combinational and sequential versions of circuits that *exhibit equivalent behaviour* (Sect. 5.1). By proving that different versions of our recursion schemes are convertible, we allow hardware designers to enable different levels of parallelism while being certain that semantics are being preserved up to timing.

Altogether, these contributions help to *separate the concerns* between the *values* a circuit must produce and the *timing* with which they are produced. In this way, timing decisions can more easily be modified later in the design process.

The codebase in which the ideas exposed in this paper are developed is available online.¹ For the sake of presentation, code excerpts in this paper may differ slightly from the corresponding ones in the repository.

2 Overview

We begin by shortly demonstrating the usage of $\lambda\pi$ -Ware. Although inspired by our previous work (*II* - Ware (Pizani Flor et al. 2016)), $\lambda\pi$ -Ware uses variable binding for sharing and loops, instead of pointfree combinators. Furthermore,

¹ <https://gitlab.com/joaopizani/lambda1-hdl/tree/paper-2017-comb-seq>.

$\lambda\pi$ -Ware has a universe of (simply-)structured types, whereas the types of Π -Ware were vectors only. In this section, we illustrate the language by means of two variations on a simple circuit. Later sections cover the syntax and semantics of $\lambda\pi$ -Ware in greater detail.

Example: Horner’s Method. We look at two circuits for calculating the value of a polynomial at a given point, one with a combinational architecture and another sequential, both based on Horner’s method.

For any coefficients a_0, \dots, a_n in \mathbb{N} , we can define a polynomial as follows:

$$p(x) = \sum_{i=0}^n a_i x^i = a_0 + a_1 x + a_2 x^2 + a_3 x^3 + \dots + a_n x^n,$$

In order to compute the value of the polynomial at a specific point x_0 of its domain, Horner’s method proceeds by using the following sequence of values:

$$\begin{aligned} b_n &:= a_n \\ b_{n-1} &:= a_{n-1} + b_n x_0 \\ &\vdots \\ b_0 &:= a_0 + b_1 x_0. \end{aligned}$$

Then b_0 is the value of our polynomial at x_0 , that is, $p(x_0)$. By iteratively expanding definitions for each of the b_i in the equations above, one arrives at a factorized form of the polynomial clearly equivalent to the usual series of powers.

Combinational Version. Horner’s method is easily expressed as a *fold*, and in $\lambda\pi$ -Ware we can build a combinational (stateless) circuit to compute this fold, for any given degree n . When reading the signature of the `horner-comb` definition below, one must note that *only* the parameters with the type former $\lambda\mathbb{H}$ are circuit inputs, and the others are synthesis parameters.

```
horner-comb : ∀ n (x₀ : λH ℕ) (aₙ : λH ℕ) (as : Vec (λH ℕ) n) → λH ℕ
horner-comb x₀ = foldl-comb (λ s a → a :+: x₀ :* s)
```

This circuit computes the value of a polynomial of degree n at a given point. It has three inputs: the point at which to evaluate the polynomial (x_0), the coefficient of highest degree (a_n) and the remaining coefficients (`as`). Later in Sect. 4 we present the detailed semantics of circuits, but for now we can say that `horner-comb n` behaves similarly to `foldl` from Agda’s standard library.

Figure 1 shows the architecture of `horner-comb`, where we can clearly see that the circuit contains no loops nor memory cells and that the *body* of the `foldl` is *replicated* n times. In the `horner-comb` model, area is linearly proportional to the degree of the polynomial, and if we want to reduce area occupation, we need to introduce *state* into the picture somehow.




Fig. 1. Block diagram of the horner-comb circuit.

Sequential Version. Next, we describe a fully *sequential* circuit to do the same calculation, using internal state to produce a *sequence of outputs*. With this architecture the area is constant (independent of the degree of the polynomial). The output value of the circuit at clock cycle i corresponds to the sum of all polynomial terms with degree smaller than or equal to i , evaluated at point x_0 .

$$\text{horner-seq} : \forall (x_0 : \lambda H \mathbb{N}) (a : \lambda H \mathbb{N}) \rightarrow \lambda H \mathbb{N}$$

$$\text{horner-seq } x_0 = \text{foldl-seq } (\lambda s a \rightarrow a +: x_0 .*: s)$$

The circuit takes two inputs: x_0 , the point at which we desire to evaluate the polynomial; and a , a *single* input containing the $n - (i+1)$ -th coefficient at the i -th clock cycle. The circuit is defined using the `foldl-seq` combinator, that iterates its argument function. This function corresponds to the loop body, mapping the current approximation, s , and the current value of the input a to a new approximation. As we shall see, to execute this *sequential* circuit, we will need to provide an initial value for the state, s .




Fig. 2. Block diagram of the horner-seq circuit.

Figure 2 shows the architecture of `horner-seq`, where we see that the body of the `foldl` is the same as in the combinational version. But now instead of n instances of the body we have a single instance, with one of its outputs *tied back* in a loop with a memory cell (shift register).

We have seen that the combinational and sequential definitions are syntactically similar, but have very different timing behaviour and generate very different architectures. First of all, the coefficients input of `horner-comb` is a vector (a *bus* in hardware parlance), while the corresponding input of `horner-seq` is a single number. Also, all the coefficients are consumed by `horner-comb` in a single clock

cycle, while `horner-seq` consumes the sequence of coefficients over n clock cycles. It is only after these n cycles that the results of the two circuits will coincide.

3 $\lambda\pi$ -Ware

We begin by fixing the universe of types, U , for the elements that circuits may produce or consume. This type is parameterized by the type of data carried over the circuit's wires (B). A typical choice of B would be bits or booleans, with other choices possible when modelling a higher-level circuit, such as integers or a datatype representing assembly instructions for a microprocessor.

```
data U (B : Set) : Set where
  unit   : U B
   $\iota$     : U B
   $\_ \Rightarrow \_$  : ( $\sigma \tau$  : U B)  $\rightarrow$  U B
   $\_ \otimes \_$  : ( $\sigma \tau$  : U B)  $\rightarrow$  U B
   $\_ \oplus \_$  : ( $\sigma \tau$  : U B)  $\rightarrow$  U B
  vec    : ( $\tau$  : U B) (n :  $\mathbb{N}$ )  $\rightarrow$  U B
```

The collection of type codes consists of a unit ($\mathbb{1}$) and base (ι) types, closed under function space (\Rightarrow), products (\otimes), coproducts (\oplus) and homogeneous arrays of fixed size (`vec`). Each element of $U B$ is mapped to the corresponding Agda type, in particular the code ι is mapped to B , the base type in our type universe.

Core Datatype. As mentioned before, our language is a deep-embedding in Agda, and circuits are elements of the λB datatype. Let us start by discussing the most fundamental constructors of λB , shown below. Additional constructors are discussed further ahead.

```
data  $\lambda B$  : ( $\Gamma$  : Ctxt B) ( $\tau$  : U B)  $\rightarrow$  Set where
   $\langle \_ \rangle$  : (g : Gate  $\tau$ )  $\rightarrow$   $\lambda B$   $\Gamma$   $\tau$ 
  var : (i :  $\Gamma \ni \tau$ )  $\rightarrow$   $\lambda B$   $\Gamma$   $\tau$ 
   $\_ \$ \_$  : (f :  $\lambda B$   $\Gamma$  ( $\sigma \Rightarrow \tau$ )) (x :  $\lambda B$   $\Gamma$   $\sigma$ )  $\rightarrow$   $\lambda B$   $\Gamma$   $\tau$ 
  let' : (x :  $\lambda B$   $\Gamma$   $\sigma$ ) (b :  $\lambda B$  ( $\sigma :: \Gamma$ )  $\tau$ )  $\rightarrow$   $\lambda B$   $\Gamma$   $\tau$ 
  loop : (c :  $\lambda B$  ( $\sigma :: \Gamma$ ) ( $\sigma \otimes \tau$ ))  $\rightarrow$   $\lambda B$   $\Gamma$   $\tau$ 
```

We use *typed De Bruijn* indices for variable binding, however, there is a convenience layer on top of λB , called λH , as seen in the overview section. Definitions using λH are essentially a shallow embedding of circuits into Agda (using Higher-Order Abstract Syntax (HOAS)), offering a more convenient programming interface by having named variables. The *unembedding* technique (Atkey et al. 2009) guarantees that it is always possible to go from a circuit definition using λH to an equivalent one using λB .

Returning to the $\lambda\mathbb{B}$ datatype itself, it is indexed by a context ($\Gamma : \text{Ctxt } \mathbb{B}$) representing the arguments to the circuit or any free variables currently in scope. The datatype is also indexed by the circuit's output type, ($\tau : \text{U } \mathbb{B}$).

The whole development is parameterized by a type of primitive gates, $\text{Gate} : \text{U } \mathbb{B} \rightarrow \text{Set}$, and the $\langle _ \rangle$ constructor creates a circuit from such a fundamental gate. One example of such type of gates is the usual triple ($\{\text{NOT}, \text{AND}, \text{OR}\}$) with Bool as the chosen base type; circuit designers, however, are free to choose the fundamental gates that best fit their domain.

Our language does have an eliminator ($_ \$ _$) for arrow types, but no introduction form. Arrow types can only be introduced by using gates, and this is by design, as we target synthesizability and circuits must be first-order to be synthesized. Using arrow types for gates allows for convenient partial application, while for general abstraction we use host language definitions as metaprograms.

While the constructors shown above form the heart of the $\lambda\mathbb{B}$ datatype, there are also constructors for products, coproducts and vectors:

$$\begin{array}{ll}
 _', _ & : \lambda\mathbb{B} \Gamma \tau_1 \rightarrow \lambda\mathbb{B} \Gamma \tau_2 \rightarrow \lambda\mathbb{B} \Gamma (\tau_1 \otimes \tau_2) \\
 \text{case}\otimes_ \text{of } _ & : \lambda\mathbb{B} \Gamma (\sigma_1 \otimes \sigma_2) \rightarrow \lambda\mathbb{B} (\sigma_1 :: \sigma_2 :: \Gamma) \tau \rightarrow \lambda\mathbb{B} \Gamma \tau \\
 \text{inl} & : \lambda\mathbb{B} \Gamma \tau_1 \rightarrow \lambda\mathbb{B} \Gamma (\tau_1 \oplus \tau_2) \\
 \text{inr} & : \lambda\mathbb{B} \Gamma \tau_2 \rightarrow \lambda\mathbb{B} \Gamma (\tau_1 \oplus \tau_2) \\
 \text{case}\oplus_ \text{either } _ \text{ or } _ & : \lambda\mathbb{B} \Gamma (\sigma_1 \oplus \sigma_2) \rightarrow \lambda\mathbb{B} (\sigma_1 :: \Gamma) \tau \rightarrow \lambda\mathbb{B} (\sigma_2 :: \Gamma) \tau \\
 & \rightarrow \lambda\mathbb{B} \Gamma \tau \\
 \text{nil} & : \lambda\mathbb{B} \Gamma (\text{vec } \tau \text{ zero}) \\
 \text{cons} & : \lambda\mathbb{B} \Gamma \tau \rightarrow \lambda\mathbb{B} \Gamma (\text{vec } \tau \text{ n}) \rightarrow \lambda\mathbb{B} \Gamma (\text{vec } \tau (\text{suc } \text{n})) \\
 \text{mapAccumL-comb} & : \lambda\mathbb{B} (\sigma :: \rho :: \Gamma) (\sigma \otimes \tau) \rightarrow \lambda\mathbb{B} \Gamma \sigma \rightarrow \lambda\mathbb{B} \Gamma (\text{vec } \rho \text{ n}) \\
 & \rightarrow \lambda\mathbb{B} \Gamma (\sigma \otimes \text{vec } \tau \text{ n})
 \end{array}$$

We give the elimination forms for both products and coproducts uniformly as case constructs, instead of projections that matches on its argument and introduces newly bound variables to the context. For vectors, $\lambda\mathbb{B}$ has the two usual introduction forms: one to produce an empty vector of any type (nil) and to extend an existing vector with a new element (cons). Finally, the *accumulating map*, mapAccumL-comb , performs a combination of map and fold! : The input vector with elements of type ρ is pointwise transformed into one with elements of type τ , all the while threading an accumulating parameter of type σ from left to right.

This eliminator is less general than the usual type theoretic elimination principle for vectors; embedding this more general eliminator would require dependent types and higher-order functions in our circuit language. To keep our object language simple, however, we chose a more simple elimination principle capable of expressing the most common hardware constructs.

4 Semantics and Properties

Where the previous section defined the *syntax* of our circuit language, we now turn our attention to its *semantics*. Although there are many different interpretations that we could assign to our circuits, for the purpose of this paper we will focus on describing a circuit's input/output behaviour.

State Transition Semantics. Circuits defined in $\lambda\mathbf{B}$ can be classified in two ways. *Combinational* circuits do not have any loops; *sequential* circuits may contain loops. To define the semantics of sequential circuits, we will need to define the type of state associated with a particular circuit. To do so, we define the inductive family $\lambda\mathbf{s}$:

```

data  $\lambda\mathbf{s} : (c : \lambda\mathbf{B} \Gamma \tau) \rightarrow \mathbf{Set}$  where
   $\_s, \_ : (sx : \lambda\mathbf{s} \ x) (sy : \lambda\mathbf{s} \ y) \rightarrow \lambda\mathbf{s} \ (x, y)$ 
   $\mathbf{sLoop} : \{c : \lambda\mathbf{B} \ (\sigma :: \Gamma) \ (\sigma \otimes \tau)\} \rightarrow (si : \mathbf{El} \ \sigma) \rightarrow (sc : \lambda\mathbf{s} \ c) \rightarrow \lambda\mathbf{s} \ (\mathbf{loop} \ c)$ 
  ...

```

This family has a constructor for each constructor of $\lambda\mathbf{B}$. Most of these constructors either contain no significant information, or simply follow the structure of the circuit, like in the clause for pairs, $_s, _$, shown above. The most interesting case is \mathbf{sLoop} , in which the state required to simulate a circuit of the form $\mathbf{loop} \ c$ consists of a value of type $\mathbf{El} \ \sigma$ — where σ is the type of the state that the circuit produces — together with any additional state that may arise from the loop body.

One other constructor of $\lambda\mathbf{s}$ deserves special attention: $\mathbf{sMapAccumL-comb}$. A circuit built with $\mathbf{mapAccumL-comb}$ consists of n *copies* of a subcircuit f connected in a row. Hence, the state of such a circuit consists of a *vector* of states, one for each of the copies of f . Correspondingly, we define the state associated with such an accumulating map as follows:

$$\mathbf{sMapAccumL-comb} : (\mathbf{sf} : \mathbf{Vec} \ (\lambda\mathbf{s} \ f) \ n) (\mathbf{se} : \lambda\mathbf{s} \ e) (\mathbf{sxs} : \lambda\mathbf{s} \ xs) \rightarrow \lambda\mathbf{s} \ (\mathbf{mapAccumL-comb} \ f \ e \ xs)$$

With this definition of state in place, we turn our attention to the semantics of our circuits. We will sketch the definition of our single step semantics, $\llbracket _ \rrbracket\mathbf{s}$, mapping a circuit, initial state and environment to a new state and the value produced by the circuit.

$$\llbracket _ \rrbracket\mathbf{s} : (c : \lambda\mathbf{B} \ \Gamma \ \tau) (m : \lambda\mathbf{s} \ c) (\gamma : \mathbf{Env} \ \mathbf{El} \ \Gamma) \rightarrow \lambda\mathbf{s} \ c \times \mathbf{El} \ \tau$$

The environment γ assigns values to any free variables in our circuit definition. The base cases for our semantics are as follows:

$$\begin{aligned} \llbracket \langle g \rangle \rrbracket\mathbf{s} \ m \ \gamma &= m, (\llbracket - \rrbracket\mathbf{s} \ g \ g) \\ \llbracket \mathbf{var} \ i \rrbracket\mathbf{s} \ m \ \gamma &= m, \mathbf{lookup} \ i \ \gamma \end{aligned}$$

In the case for gates, we apply the semantics of our atomic gates, described by the auxiliary function $\llbracket - \rrbracket g$; in the case for variables, we lookup the corresponding value from the environment. Both these cases do not refer to the circuit's state. This state becomes important when simulating loops. In the clauses for application, **let'** and **loop**, shown in Listing 1, we do need to consider the circuit's state.

$$\begin{aligned}
\llbracket f \$ x \rrbracket s (mf \$ mx) \gamma &= \mathbf{let} (mx', rx) = \llbracket x \rrbracket s mx \gamma \\
&\quad (mf', rf) = \llbracket f \rrbracket s mf \gamma \\
&\quad \mathbf{in} ((mf' \$ mx'), (rf rx)) \\
\llbracket \mathbf{let}' x b \rrbracket s (sLet mx mb) \gamma &= \mathbf{let} (mx', rx) = \llbracket x \rrbracket s mx \gamma \\
&\quad (mb', rb) = \llbracket b \rrbracket s mb (rx :: \gamma) \\
&\quad \mathbf{in} ((sLet mx' mb'), rb) \\
\llbracket \mathbf{loop} f \rrbracket s (sLoop ml mf) \gamma &= \mathbf{let} (mf', (ml', rl)) = \llbracket f \rrbracket s mf (ml :: \gamma) \\
&\quad \mathbf{in} ((sLoop ml' mf'), rl)
\end{aligned}$$

Listing 1: State-combining clauses of the single-step state transition semantics.

In the cases of application and **let**, each subcircuit simply “takes a step” independently and the next state of the whole circuit is a combination of the next states of each subcircuit. The case for **loop** is slightly more interesting: the loop body, f , takes an additional input, namely the current state given by the ml parameter of **sLoop** constructor.

The further clauses of the transition function handle the introduction and elimination forms of products, coproducts and vectors. They are all defined simply by recursive evaluation of the subcircuits, and are straightforward enough to omit from the presentation here. For example, the clause for coproduct elimination is shown below:

$$\begin{aligned}
\llbracket \mathbf{case} \oplus x \vee y \mathbf{either} f \mathbf{or} g \rrbracket s (sCase \oplus mxy mf mg) \gamma &= \\
\mathbf{let} (mxy', rx \vee ry) &= \llbracket x \vee y \rrbracket s mxy \gamma \\
\mathbf{in} [\mathbf{map} \times (\mathbf{flip} (sCase \oplus mxy') mg) \mathbf{id} \circ (\llbracket f \rrbracket s1 mf \gamma) \\
&\quad , \mathbf{map} \times ((sCase \oplus mxy') mf) \mathbf{id} \circ (\llbracket g \rrbracket s1 mg \gamma) \\
&\quad] rx \vee ry
\end{aligned}$$

First the coproduct value ($x \vee y$) is evaluated, computing a result value and its next state. The result of the evaluation ($rx \vee ry$) is then fed to Agda's coproduct eliminator ($[_,_]$); the functions that process the left and right injections proceed accordingly. In either case, the value is fed into evaluation of the appropriate body (either f or g), and the result is then used as the result of the whole coproduct evaluation.

Similarly our elimination principle for vectors, **mapAccumL-comb**, is worth highlighting:

$$\begin{aligned} \llbracket \text{mapAccumL-comb } f \ e \ xs \rrbracket s \ (s \text{MapAccumL-comb } mfs \ me \ mxs) \ \gamma &= \\ \text{let } (me', re) &= \llbracket e \rrbracket s \ me \ \gamma \\ (mxs', rxs) &= \llbracket xs \rrbracket s \ mxs \ \gamma \\ (rz, mfs', rys) &= \text{mapAccumL2 } (\text{transformF } \llbracket f \rrbracket s2 \ \gamma) \ re \ mfs \ rxs \\ \text{in } (s \text{MapAccumL-comb } mfs' \ me' \ mxs', (rz, rys)) \end{aligned}$$

The above clause is key in the relation that we later establish (Sect. 5.1) between combinational and sequential versions of circuits. The three key sub-steps involved in this clause are: evaluation of the left identity element (e), the evaluation of the row of inputs (xs) and the row of step function copies (f).

The first two steps are as expected: both the identity and row of inputs take a step, and we thus obtain the next state and result values of each. The core step is then evaluating the row of copies of f , and its semantics are given using the auxiliary function `mapAccumL2`.

The `mapAccumL2` function is simply a two-input version of an accumulating map, which works by simply zipping the pair of input vectors and calling the `mapAccumL` function from Agda's standard library.

$$\begin{aligned} \text{mapAccumL} &: (\sigma \rightarrow \alpha \rightarrow (\sigma \times \beta)) \rightarrow \sigma \rightarrow \text{Vec } \alpha \ n \rightarrow \sigma \times \text{Vec } \beta \ n \\ \text{mapAccumL } f \ s \ \llbracket &= s, \llbracket \\ \text{mapAccumL } f \ s \ (x :: xs) &= \text{let } s', y = f \ s \ x \\ & \quad s'', ys = \text{mapAccumL } f \ s' \ xs \\ & \quad \text{in } s'', (y :: ys) \\ \text{mapAccumL2} &: (\sigma \rightarrow \alpha \rightarrow \gamma \rightarrow (\sigma \times \beta \times \delta)) \rightarrow \sigma \rightarrow \text{Vec } \alpha \ n \rightarrow \text{Vec } \gamma \ n \\ & \quad \rightarrow \sigma \times \text{Vec } \beta \ n \times \text{Vec } \delta \ n \\ \text{mapAccumL2 } f \ s \ xs \ ys &= \text{map} \times \text{id} \ \text{unzip} \ \$ \ \text{mapAccumL } (\text{uncurry} \circ f) \ s \ (\text{zip } xs \ ys) \end{aligned}$$

In the semantics of `mapAccumL-comb`, we apply `mapAccumL2` to the vector with the result of xs (called rxs) as well as the vector with states for the copies of f (called mfs). Then, as the result of the application we obtain the final accumulator value and vector of result values, *together with the vector of next state values* (mfs').

Multi-step Semantics. To describe the behaviour of a circuit *over time*, we need to define another semantics. More specifically, in this work we consider only discrete-time synchronous circuits, and thus we will show how to use $\llbracket _ \rrbracket s$ to define a multi-step state-transition semantics.

$$\begin{aligned} \llbracket _ \rrbracket n &: (c : \lambda B \ \Gamma \ \tau) \ (m : \lambda s \ c) \ n \ (\gamma : \text{Vec } (\text{Env } \text{El } \Gamma) \ n) \rightarrow \lambda s \ c \times \text{Vec } (\text{El } \tau) \ n \\ \llbracket _ \rrbracket n \ c \ m \ n &= \text{mapAccumL } \llbracket c \rrbracket s \ m \end{aligned}$$

When simulating a circuit for n cycles, we need to take not *one* input environment but n , and instead of producing a single value, the simulation returns a vector of n values. Just as we saw for `mapAccumL-comb`, we ensure that the newly computed state is threaded from one simulation cycle to the next.

This is exactly the behaviour of an accumulating map, thus the use of `mapAccumL` here. The use of `mapAccumL` here is the key to the connection between the multi-cycle of circuits using `loop` and the single-cycle behavior of circuits using `mapAccumL-comb`.

5 Combinational and Sequential Combinators

With $\lambda\pi$ -Ware we intend to give a hardware developer more freedom to explore the trade-offs between area, frequency and number of cycles that a circuit might take to complete a computation. This freedom comes from the proven guarantees of convertibility between combinational and sequential versions of circuits.

To make it easier to explore this design space, we provide some *circuit combinators* for common patterns. Each of these patterns comes in a pair of sequential and combinational versions, with a lemma relating the two. If a circuit is defined using one of these combinators, changing between architectures is as easy as changing the combinator version used. The associated lemma guarantees the relation between the functional behaviour of the versions.

All combinators in this section are derived from the two primitive constructors `loop` and `mapAccumL-comb`. By appropriate partial application and the use of “wrappers” to create the loop body, all sequential combinators are derived from `loop`. Similarly, using the same wrappers but with `mapAccumL-comb`, we derive all combinational combinators.

Of notice is also the fact that, in this section, we present the combinators in *De Bruijn* style, as this is the most useful representation to use when evaluating circuit (generators), which is covered in 5.1.

The map Combinators. For example, we might want to easily build circuits that *map* a certain function over its inputs. We will define both the sequential and combinational `map` combinators in terms of a third circuit, `mapper`. The sequential version is given by `map-seq`:

$$\begin{aligned} \text{mapper} &: (f : \lambda B (\rho :: \Gamma) \tau) \rightarrow \lambda B (\sigma :: \rho :: \Gamma) (\sigma \otimes \tau) \\ \text{mapper } f &= \#_0, K_1 f \\ \text{map-seq} &: (f : \lambda B (\rho :: \Gamma) \tau) \rightarrow \lambda B (\rho :: \Gamma) \tau \\ \text{map-seq } f &= \text{loop } \{\sigma = \mathbb{1}\} (\text{mapper } f) \end{aligned}$$

We define `map-seq` by applying `loop` to the `mapper f` circuit. In `mapper`, the next state (first projection of the pair) is a copy of its first input ($\#_0$), whereas the second projection is made by the *weakened* `f`, which discards its first input.

The combinational version of the same combinator (`map-comb`) is defined in terms of `mapAccumL-comb` and `mapper`:

$$\begin{aligned} \text{map-comb} &: (f : \lambda B (\rho :: \Gamma) \tau) (xs : \lambda B \Gamma (\text{vec } \rho \ n)) \rightarrow \lambda B \Gamma (\text{vec } \tau \ n) \\ \text{map-comb } f \ xs &= \text{snd } (\text{mapAccumL-comb } (\text{mapper } f) \ \text{unit } \ xs) \end{aligned}$$

In the above definition we note that we are free to choose the type of the “initial element” (2nd argument), but we use $\mathbb{1}$ (value unit), as units can always be used regardless of the base type chosen in the development. Furthermore, we use `snd` to extract only the second element of the pair (the output vector), and discard the “final element” outputted.

The foldl-scanl Combinators. Perhaps even more useful than mapping is *scanning* and *folding* over a vector of inputs. To obtain the sequential and combinational versions of such combinators, we again apply the `loop` and `mapAccumL-comb` primitives to a special body which wraps the binary operation (`f`) of the scan/fold.

$$\begin{aligned} \text{folder} &: (f : \lambda B (\sigma :: \rho :: \Gamma) \sigma) \rightarrow \lambda B (\sigma :: \rho :: \Gamma) (\sigma \otimes \sigma) \\ \text{folder } f &= \#_0, f \\ \text{foldl-scanl-seq} &: (f : \lambda B (\sigma :: \rho :: \Gamma) \sigma) \rightarrow \lambda B (\rho :: \Gamma) \sigma \\ \text{foldl-scanl-seq } f &= \text{loop } (\text{folder } f) \end{aligned}$$

The wrapper called `folder` makes the next state equal to the first input of the binary operator, and the output be the result of applying the binary operator. In the above definition of `foldl-scanl-seq`, we get the behaviour of `scanl` and `foldl` *combined*: The circuit outputs from clock cycle 0 to `n` form the result of the `scanl` operation, and the last one at cycle `n+1` is the value of the `foldl`.

The combinational version also has such a combined behaviour:

$$\begin{aligned} \text{foldl-scanl-comb} &: (f : \lambda B (\sigma :: \rho :: \Gamma) \sigma) (e : \lambda B \Gamma \sigma) (xs : \lambda B \Gamma (\text{vec } \rho \ n)) \\ &\rightarrow \lambda B \Gamma (\sigma \otimes \text{vec } \sigma \ n) \\ \text{foldl-scanl-comb } f \ e \ xs &= \text{mapAccumL-comb } (\text{folder } f) \ e \ xs \end{aligned}$$

In `foldl-scanl-comb`, we obtain a pair as output, of which the first element is the `foldl` component, and the second element is the `scanl` (vector) component. Thus by simply applying the `fst` and `snd` functions we can obtain the usual `foldl` and `scanl`.

Whereas these combinators capture some common patterns in hardware design, their usefulness also depends on lemmas relating their combinational and sequential versions.

5.1 Convertibility of Combinational and Sequential Versions

In this section we make precise the relation between circuits with different levels of statefulness. For conciseness, only the extreme cases are handled: completely stateless (combinational) versus completely sequential. However, nothing in the following treatment precludes it from being used for *partial unrolling*.

We will show that when two circuits are deemed “convertible up to timing”, they can be substituted for one another with minor interface changes in the surrounding context but no alteration of the values ultimately produced.

The relation of convertibility relies on the fact that any sequential circuit will have an occurrence of the `loop` constructor. As such, a less stateful variant of such a circuit can be obtained by substituting the occurrence of `loop` with one of `mapAccumL-comb`, thereby unrolling the loop. The fundamental relation between `loop` and `mapAccumL-comb` is what we now establish. First, recall the types of the single- and multi-step semantic functions:

$$\begin{aligned} \llbracket _ \rrbracket_s &: (c : \lambda B \Gamma \tau) (m : \lambda s c) (\gamma : \text{Env El } \Gamma) \rightarrow \lambda s c \times \text{El } \tau \\ \llbracket _ \rrbracket_n &: (c : \lambda B \Gamma \tau) (m : \lambda s c) n (\gamma : \text{Vec (Env El } \Gamma) n) \rightarrow \lambda s c \times \text{Vec (El } \tau) n \end{aligned}$$

Now, to establish the desired relation, we apply both the single and multi-cycle semantics. The *step function* subcircuit (called `f`) is equal in both cases, and the `mapAccumL-comb` case takes 2 extra parameters besides `f`.

$$\begin{aligned} \llbracket \text{mapAccumL-comb } f \text{ (val } m) \text{ xs} \rrbracket_s & \text{ (sMapAccumL-comb (replicate } mf)) \gamma : \text{Tpar} \\ \llbracket \text{loop } f \rrbracket_n & \text{ (sLoop } m \text{ mf) } n \text{ (map } (_ :: \gamma) \text{ xs)} : \text{Tloop} \\ \text{where } \text{Tpar} &= \lambda s (\dots) \times \text{Tp } \sigma \times \text{Vec (Tp } \tau) n \\ \text{Tloop} &= \lambda s (\dots) \times \text{Vec (Tp } \tau) n \end{aligned}$$

The second parameter of `mapAccumL-comb` must be a circuit whose value is the same as the first parameter of `sLoop`, and we use here the simplest possible such circuit: `(val m)`. The third parameter (`xs`) is the input vector of size `n`, and is used to build the vector of environments used by the multi-cycle semantics (`map (_ :: \gamma) xs`).

Finally, the state of the `mapAccumL-comb` case is built by simply replicating one state of `f` by `n` times. Stating the convertibility property in this way makes it be valid only for a state-independent `f`, that is, when the input/output semantics of `f` is independent of the state.

$$\begin{aligned} \text{state-independent} &: \forall (c : \lambda B \Gamma \tau) \rightarrow \text{Set} \\ \text{state-independent } c &= \forall sa sb \gamma \rightarrow \llbracket c \rrbracket_s sa \gamma \equiv \llbracket c \rrbracket_s sb \gamma \end{aligned}$$

This restriction on `f` could be somewhat further loosened (as is discussed in Sect. 6.2), but we work here with `state-independent` loop bodies to simplify the presentation.

As we have seen, the results from applying each semantic function have different types (`Tpar` and `Tloop`), so the relation comparing these results is more subtle than just equality. We define this relation, called `_=*_`, as follows:

$$\begin{aligned} _ =^* _ &: \text{Tpar} \rightarrow \text{Tloop} \rightarrow \text{Set} \\ (_ , s', xs') =^* (sm'' , xs'') &= (s' \equiv \text{gets}_0 sm'') \times (xs' \equiv xs'') \end{aligned}$$

Both sides of `(_ =^* _)` consist of a pair of next state and circuit outputs. In the `mapAccumL-comb` case, the next state can be ignored in the comparison, but in the `loop` case, the *value* stored in the loop state (obtained by `gets0`) must be

equal to the first output of evaluating `mapAccumL-comb`. With the comparison function defined, we can finally completely express the relation we desire:

$$\llbracket \text{mapAccumL-comb } f \text{ (val } m \text{) } xs \rrbracket s \text{ (sMapAccumL-comb (replicate } mf) \text{) } \gamma \\ \equiv^* \llbracket \text{loop } f \rrbracket n \text{ (sLoop } m \text{ } mf) \text{ } n \text{ (map } (_ :: \gamma) \text{) } xs$$

Proof of the Basic Relation. The proof of the basic convertibility relation between `mapAccumL-comb` and `loop` proceeds by induction on the input vector `xs`. Due to the deliberate choice of semantics for both constructors involved, and the choice of the right parameters for the application of each, a considerable part of the proof is achieved by just the built-in reduction behaviour of the proof assistant (Agda).

The only key lemma involved is shown below. Namely, the state-independence principle is shown to hold for a whole vector, assuming that it holds for the body circuit `f`.

$$\text{state-independent-vec} : \forall (mfas mfbs : \text{Vec } (\lambda s \text{ } f) \text{ } n) (p : \text{state-independent } f) \\ \rightarrow \text{mapAccumL2 (transformF } \llbracket f \rrbracket s2 \text{ } \gamma) \text{ } e \text{ } mfas \text{ } xs \\ \equiv \text{mapAccumL2 (transformF } \llbracket f \rrbracket s2 \text{ } \gamma) \text{ } e \text{ } mfbs \text{ } xs$$

This lemma is useful because both left-hand side and right-hand side of the convertibility relation can be transformed into applications of `mapAccumL2` simply by reduction, but with different state vector parameters. Thus the lemma is used to bring the sub-goals to a state where they can be closed by using the induction hypothesis.

$$\text{mapAccumL-comb-seq} : \\ \llbracket \text{mapAccumL-comb } f \text{ (val } m \text{) } xs \rrbracket s \text{ (sMapAccumL-comb (replicate } mf) \text{) } \gamma \\ \equiv^* \llbracket \text{loop } f \rrbracket n \text{ (sLoop } m \text{ } mf) \text{ } n \text{ (map } (_ :: \gamma) \text{) } xs \\ \text{mapAccumL-comb-seq } f \text{ } mf \text{ } m \text{ (} x :: xs \text{) } \gamma \text{ } p = g:m , g:ys \textbf{ where} \\ m' , mf' , y = (\text{transformF } \llbracket f \rrbracket s2 \text{ } \gamma) \text{ } m \text{ } mf \text{ } x \text{ -- take one step} \\ ih:m , ih:ys = \text{mapAccumL-comb-seq } f \text{ } mf' \text{ } m' \text{ } xs \text{ } \gamma \text{ } p \text{ -- ind. hyp.} \\ \text{lemma} = \text{state-independent-vec } f \text{ } xs \text{ (replicate } mf) \text{ (replicate } mf') \text{ } p \\ g:m : p_1 (p_2 (\llbracket \text{mapAccumL-comb } f \text{ } \dots \rrbracket s \text{ } \dots)) \equiv \text{gets}_0 (p_1 (\llbracket \text{loop } f \rrbracket n \text{ } \dots)) \\ g:ys : p_2 (p_2 (\llbracket \text{mapAccumL-comb } f \text{ } \dots \rrbracket s \text{ } \dots)) \equiv p_2 (\llbracket \text{loop } f \rrbracket n \text{ } \dots) \\ g:m = (\text{cong } \dots \text{ lemma}) \langle \text{trans} \rangle ih:m \\ g:ys = \text{cong}_2 \dots ((\text{cong } \dots \text{ lemma}) \langle \text{trans} \rangle ih:ys)$$

Convertibility of Derived Combinators. When building circuits using the derived combinators (`map`, `foldl-scanl`, etc.), the convertibility between different (more or less stateful) variants of such circuits rely on the convertibility between the different variants of the combinators themselves.

The basic convertibility principle shown above between `mapAccumL-comb` and `loop` is the most general one, and can be directly applied to the derived combinators as well, as they are all just a specialized instance of `mapAccumL-comb` or `loop`. However, for the derived combinators, some more specific properties are useful.

With regards to the `map` combinators, for example, we wish that the vectors produced by the combinational and sequential versions be equal, without any regard for initial or final states. This can be succinctly expressed as:

$$\begin{aligned} & \text{snd } (\llbracket \text{map-comb } f \text{ xs } \rrbracket_s \text{ units } \gamma) \\ \equiv & \text{snd } (\llbracket \text{map-seq } f \quad \rrbracket_n \text{ units}' (\text{map } (_:: \gamma) \text{ xs})) \end{aligned}$$

Where `units` and `units'` are simply the states (composed of units) that need to be passed to the semantic function but are irrelevant for the computed vectors.

On the other hand, when comparing `foldl-comb` to `foldl-seq`, the intermediate values produced in the output of `foldl-seq` are disregarded, and only the final state matters.

$$\begin{aligned} & \text{fst } (\llbracket \text{foldl-comb } f (\text{val } e) \text{ xs } \rrbracket_s m \quad \gamma) \\ \equiv & \text{fst } (\llbracket \text{foldl-seq } f \quad \rrbracket_n (\text{sFoldl } e m) (\text{map } (_:: \gamma) \text{ xs})) \end{aligned}$$

Both of these properties (for `map` and for `foldl`) can simply be proven by application of the general property shown above for `mapAccumL-comb` and `loop`. This is because the definition of the derived combinators is just a partial application of `mapAccumL-comb` and `loop`, along with projections.

5.2 Applications of the Combinational and Sequential Combinators

In this section we describe several variants of circuit families that compute matrix multiplication, as a commonly used application of the aforementioned techniques.

The first design choice involved in this example application is *how to represent matrices*, i.e., the choice of the *matrix type*. Traditionally in computing contexts, matrices are mostly represented in two ways: *row major* (vector of rows) and *column major* (vector of columns). As it turns out, *both* representations are useful for our purposes, so we show both here:

$$\begin{aligned} \text{RMat } \text{CMat} & : (r \ c : \mathbb{N}) \rightarrow \mathbb{U} \ \mathbb{N} \\ \text{RMat } r \ c & = \text{vec } (\text{vec } \mathbb{N} \ c) \ r \\ \text{CMat } r \ c & = \text{vec } (\text{vec } \mathbb{N} \ r) \ c \end{aligned}$$

Here, `RMat r c` and `CMat r c` both represent matrices with `r` rows and `c` columns, the difference being only whether they are row- or column-major. Going further with the example, we need to define the basic ingredient of matrix multiplication: the *dot product* of two equally-sized vectors.

```

dp : λH (vec ℕ n) → λH (vec ℕ n) → λH ℕ
dp xs ys = foldl-comb _:+:_ (val 0) (zipWith-comb _:*:_ xs ys)

```

The dot product is simply defined as element-wise multiplication of the vectors and summing up the results. We can then use the dot product m times in order to multiply a vector by a compatibly-sized matrix.

```

vec×mat-comb : λH (vec ℕ n) → λH (CMat n m) → λH (vec ℕ m)
vec×mat-comb v m = map-comb (dp v) m

```

Here an important detail resides: as the dot product is done for *each column* of the matrix, the matrix argument of `vec×mat` must be in *column-major* representation. Also, here we start having choices: we may either have the computation done combinationally as above, or sequentially as below:

```

vec×mat-seq : λH (vec ℕ n) → λH (vec ℕ n) → λH ℕ
vec×mat-seq v m = map-seq (dp v) m

```

With the multi-step semantics in mind, we know that each of the m columns of the matrix will be present on the circuit's second input, one per clock cycle, and that collecting the output values for m cycles gives the same vector of results as the one from the combinational version.

For defining the multiplication of two matrices, we simply use `vec×mat` on each row of the left matrix. If using `vec×mat-comb`, we obtain a matrix multiplication circuit with area proportional to $r * c$, whereas by using `vec×mat-seq` the area is proportional to $r * 1$.

```

mat×mat-comb : λH (RMat n m) → λH (CMat m p) → λH (RMat n p)
mat×mat-comb mr mc = map-comb (flip vec×mat-comb mc) mr
mat×mat-seq : λH (RMat n m) → λH (vec ℕ m) → λH (vec ℕ n)
mat×mat-seq mr mc = map-comb (flip vec×mat-seq mc) mr

```

In the combinational version (`mat×mat-comb`), all the rows in the resulting matrix are computed in parallel, with the column-positioned values inside each row computed also in parallel. In the sequential version, at each clock cycle one whole *column* is produced, with the row-positioned values inside each column computed in parallel.

Matrix multiplication as defined here has two nested recursion blocks, and thus four ways in which it could be sequentialized. Above we have shown two possible such choices, and the other two can simply be obtained by swapping `map-comb` for `map-seq`.

6 Discussion

6.1 Related Work

There is a rich tradition of using functional programming languages to model and verify hardware circuits, Sheeran (2005) gives a good overview – we restrict

ourselves to the most closely related languages here. Languages embedded in Haskell, such as Lava and Wired, typically rely on automated theorem provers and testing using QuickCheck for verification. In $\lambda\pi$ -Ware, however, we can perform *inductive* verification of our circuits. Existing embeddings in most theorem provers, such as Coquet (Braibant 2011) and *II* - Ware (Pizani Flor et al. 2016), have a more limited treatment of variable scoping and types. More recent work by Choi et al. (2017) is higher level, but sacrifices the ability to be simulated directly (using denotational semantics) in the theorem prover.

6.2 Future Work

Other Timing Transformations. While our language easily lets you explore possible designs, trading time and space, there are several alternative transformations, such as *pipelining* that we have not yet tried to describe in this setting.

While we have a number of combinators for transforming between combinational and sequential circuits, these are mostly aimed at linear, list-like data. Even though these structures are the most prevalent in hardware design, we would like to explore related timing transformations on tree-structured circuits. To this end, it would be interesting to look into the formalization and verification of *flattening* transformations, and of the work done in the field of *nested data parallelism*.

Relaxed Unrolling Restriction. In Sect. 5.1 we mention that the proof of semantics preservation for loop unrolling relies on the premise that the loop body is *state-independent*, that is, it has the same input/output behaviour for any given state. This premise can be relaxed somewhat, and proving that loop unrolling still preserves semantics under this relaxed premise is (near-)future work.

The relaxed restriction on the body f of a loop to be unrolled is as follows:

$$\begin{aligned} \text{state-input-independent} &: \forall (c : \lambda B \Gamma \tau) \rightarrow \text{Set} \\ \text{state-input-independent } c &= \text{fst } (\llbracket c \rrbracket_s \text{ sa } \gamma) \equiv \text{fst } (\llbracket c \rrbracket_s \text{ sa } \delta) \end{aligned}$$

That is, the next state (fst projection) is equal even with evaluation taking different input environments. This condition is necessary because when writing the combinational version of a loop construct we must give each copy of f its own initial state. As the desired initial state for each such copy must be known at verification time, it cannot depend on input.

Using the definitions from Sect. 5.1 along with the relaxed hypotheses above, we can show that not only total, but also *partial unrolling* preserves semantics up to timing.

7 Conclusion

There are several advantages to be gained by embedding a hardware design DSL in a host language with dependent types, such as Agda. Among these advantages are the easy enforcement of some well-formedness characteristics of circuits, the power given by the host's type system to express object language types and design constraints. The crucial advantage though, is the ability to have *modelling, simulation, synthesis and theorem proving* in the same language.

By using the host language's theorem-proving abilities, we are able not only to show properties of individual circuits, but of (infinite) classes of circuits, defined by using *circuit generators*. Particularly interesting is the ability to have *verified transformations*, preserving some semantics.

The focus of this paper lies on timing-related transformations, but we also recognize the promise of theorem proving for the formalization of other non-functional aspects of circuit design, such as power consumption, error correction, fault-tolerance and so forth. The formal study of all these aspects of circuit construction and program construction could benefit from mechanized verification.

Acknowledgments. We would like to thank the very helpful feedback gathered during the visit to Chalmers University of Technology, funded by COST Action EUTypes CA15123. Especially valuable were the meetings and discussions with Mary Sheeran, whose deep knowledge of the field oriented this work in its beginning stage. Also, we are very thankful to the feedback given during our presentation on this topic at the TFP2017 conference in Canterbury.

This work was supported by the Netherlands Organization for Scientific Research (NWO) project on *A Dependently Typed Language for Verified Hardware*.

References

- Atkey, R., Lindley, S., Yallop, J.: Unembedding domain-specific languages. In: Proceedings of the 2nd ACM SIGPLAN Symposium on Haskell, Haskell 2009, pp. 37–48. ACM, New York (2009). <https://doi.org/10.1145/1596638.1596644>. ISBN 978-1-60558-508-6
- Axelsson, E., Claessen, K., Sheeran, M.: Wired: wire-aware circuit design. In: Borrione, D., Paul, W. (eds.) CHARME 2005. LNCS, vol. 3725, pp. 5–19. Springer, Heidelberg (2005). https://doi.org/10.1007/11560548_4
- Baaij, C.P.R.: Digital circuits in CλaSH: functional specifications and type-directed synthesis. info:eu-repo/semantics/doctoralThesis. University of Twente, Enschede, January 2015. <https://doi.org/10.3990/1.9789036538039>
- Bjesse, P., Claessen, K., Sheeran, M., Singh, S.: Lava: hardware design in Haskell. ACM SIGPLAN Not. **34**(1), 174–184 (1999). <https://doi.org/10.1145/291251.289440>. ISSN 03621340
- Boulton, R.J., Gordon, A.D., Gordon, M.J.C., Harrison, J., Herbert, J., Van Tassel, J.: Experience with embedding hardware description languages in HOL. In: TPCD, vol. 10, pp. 129–156 (1992)
- Brady, E., Mckinna, J., Hammond, K.: Constructing correct circuits: verification of functional aspects of hardware specifications with dependent types. In: Trends in Functional Programming 2007 (2007)

- Braibant, T.: Coquet: a Coq library for verifying hardware. In: Jouannaud, J.-P., Shao, Z. (eds.) CPP 2011. LNCS, vol. 7086, pp. 330–345. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-25379-9_24
- Braibant, T., Chlipala, A.: Formal verification of hardware synthesis. In: Sharygina, N., Veith, H. (eds.) CAV 2013. LNCS, vol. 8044, pp. 213–228. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-39799-8_14
- Choi, J., Vijayaraghavan, M., Sherman, B., Chlipala, A., Arvind: Kami: a platform for high-level parametric hardware specification and its modular verification. Proc. ACM Program. Lang. 1(ICFP), 24:1–24:30 (2017). <https://doi.org/10.1145/3110268>. ISSN 2475-1421
- Gill, A., Bull, T., Kimmell, G., Perrins, E., Komp, E., Werling, B.: Introducing Kansas Lava. In: Morazán, M.T., Scholz, S.-B. (eds.) IFL 2009. LNCS, vol. 6041, pp. 18–35. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-16478-1_2
- Hanna, F.K., Daeche, N.: Dependent types and formal synthesis. Philos. Trans. Phys. Sci. Eng. **339**(1652), 121–135 (1992). <http://www.jstor.org/stable/54016>. ISSN 0962-8428
- Melham, T.: Higher Order Logic and Hardware Verification. Cambridge Tracts in Theoretical Computer Science, vol. 31. Cambridge University Press, Cambridge (1993). <https://doi.org/10.1017/CBO9780511569845>. <http://www.cs.ox.ac.uk/tom.melham/pub/Melham-1993-HOL.html>. ISBN 0-521-41718-X
- Pizani Flor, J.P., Sijlsing, Y., Swierstra, W.: π -ware: hardware description and verification in Agda. In: Uustalu, T. (ed.) 21th International Conference on Types for Proofs and Programs (TYPES 2015). Leibniz International Proceedings in Informatics (LIPIcs) (2016)
- Sander, I., Jantsch, A.: System modeling and transformational design refinement in ForSyDe [formal system design]. IEEE Trans. Comput. Aided Des. Integr. Circuits Syst. **23**(1), 17–32 (2004). <https://doi.org/10.1109/TCAD.2003.819898>. ISSN 0278-0070
- Sheeran, M.: Hardware design and functional programming: a perfect match (2005). http://www.jucs.org/jucs_11_7/hardware_design_and_functional/jucs_11_7_1135_1158_sheeran.pdf
- Sheeran, M.: muFP, a language for VLSI design. In: Proceedings of the 1984 ACM Symposium on LISP and Functional Programming, pp. 104–112. ACM Press (1984). <https://doi.org/10.1145/800055.802026>. ISBN 0897911423
- Singh, S.: Designing reconfigurable systems in lava. In: Proceedings 17th International Conference on VLSI Design 2004, pp. 299–306 (2004). <https://doi.org/10.1109/ICVD.2004.1260941>

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

