

Encoding Purely Functional Languages in an RTL-based Compiler

Amelia Dobis*

Gongqi Huang*

Mae Milano

Princeton University
Department of Computer Science
Princeton, NJ, USA

Abstract

Unifying compiler infrastructure for high-level hardware languages has long been the goal of the CIRCT project. However, by following trends in hardware language design, CIRCT has specialized itself in HLS and RTL-like paradigms, while requiring elements from other paradigms to be handled mostly in their respective frontends. In this work, we present 3 approaches to support purely functional languages natively in CIRCT. We demonstrate these by creating a CIRCT-based compiler for Clash, a Haskell-as-hardware language. These methods range from simple conversions leveraging existing frontend transformations, to the introduction of full lambda calculus support in CIRCT via the `lc` dialect. By retaining the original functional structure of a Clash design, the `lc` dialect captures the source’s high-level structure, enabling optimizations and higher-quality emission by backends. This paves the way for non-RTL paradigms to leverage the power of a unified compiler.

1 Introduction

Non-HLS (High-Level Synthesis) high-level hardware languages have evolved through distinct “waves”, each moving further from gate-level netlists. Initially, behavioral descriptions of hardware were introduced via VHDL and Verilog. While popular, the limitations of these languages were already apparent early on and a **first wave** of new high-level hardware languages, adopting components from functional programming languages, followed with most notably Lava [1], and BlueSpec [2]. This was succeeded by a **second wave** of embedded Domain Specific Languages (eDSLs), such as Chisel (Scala) [3], pyMTL (Python) [4], and HardCaml (OCaml) [5], which embed structural components in a host language to generate hardware. Finally, a **third wave** has emerged, focusing on custom languages with advanced type systems, such as Spade [6], Filament [7], or Anvil [8].

Despite the distinct philosophies of these waves, their paradigms remain close to the Register Transfer Level (RTL), often utilizing the initial two behavioral languages as their compilation targets. This motivates a unified, MLIR-based [9] infrastructure, namely CIRCT [10], to consolidate the

compilation of these diverse hardware languages via three central RTL-based core dialects. While CIRCT has had success in unifying the compilation of RTL-like paradigms, it is primarily optimized for structural abstractions with explicit stateful elements. As a result, CIRCT struggles to natively encode purely functional languages, where hardware is defined through function composition and Algebraic Data Types (ADTs), without requiring high-level functional constructs to be elaborated away by frontends.

Clash. Being neither an eDSL, a custom language, nor a form of HLS, Clash [11] leverages the power of Haskell by reinterpreting System FC [12], the internal representation of the Glasgow Haskell Compiler (GHC) [13], to produce hardware. As a result, designers are able to describe hardware through purely functional abstractions, such as higher-order functions, polymorphism, and ADTs, without sacrificing the explicit structural nature of hardware languages. To bridge the semantic gap, Clash employs a normalization process that transforms System FC into a flat, “normalized” representation, while only eliminating recursive constructs that lack a direct hardware equivalent. This form is then finally lowered to a synthesizable target in VHDL or Verilog.

In this work, we propose 3 approaches for encoding purely functional hardware languages in a unified RTL-based compiler such as CIRCT. We demonstrate these approaches by creating a CIRCT-based Clash compiler.

2 Building A CIRCT-Based Clash Compiler

In porting a compiler to CIRCT, we start by identifying potential entry-points, i.e. stages at which we can transfer computation from our frontend (Clash) to our compiler (CIRCT). For Clash, we identify 3 such entry-points. These trade-off utilizing existing lowering methods found in our frontend with preserving the source’s high-level structure in backends. Figure 1 illustrates how an example description of control logic, i.e. f_{sm} , is compiled using the existing Clash compiler, and then integrated with CIRCT through entry-points $\{\text{fsm}\}$. In particular, f_{sm} is a higher-order function that takes a integer op , an ADT $state$ representing its state, and returns a tuple of the next state and output. bar is a function that performs some computation. We omit Clash-specific operations to generate the top-level module to focus on the functional description of a stateful circuit.

2.1 The Low-Hanging Fruits

The most straightforward integration path of the 3 entry-points is $\{\text{fsm}\}$, where the fully lowered netlist, generated by

*Equal contribution.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
LATTE '26, March 23, 2026, Pittsburgh, PA, USA
© 2026 Copyright held by the owner/author(s).

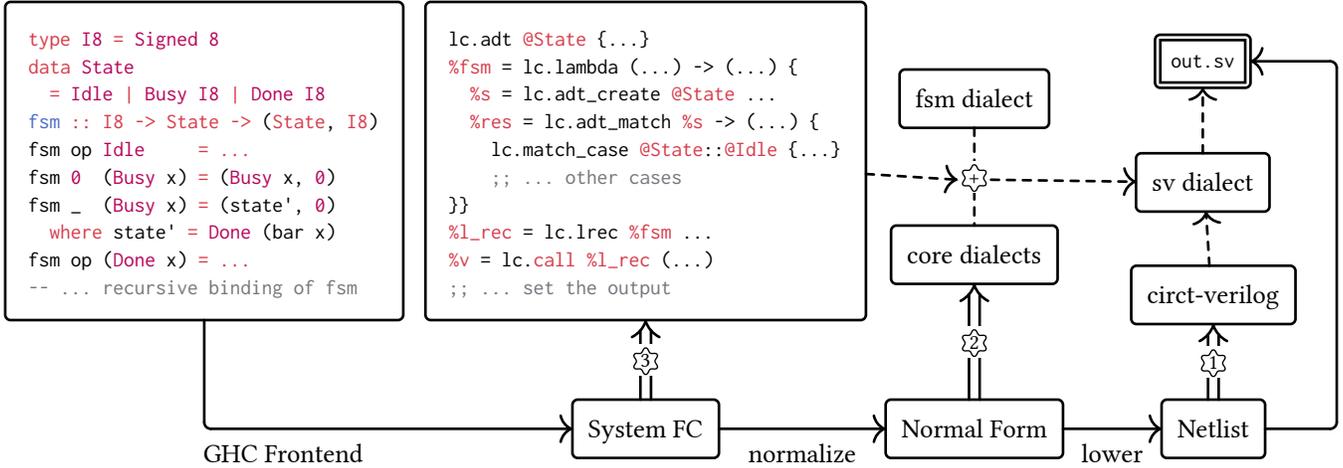


Figure 1: The Clash compilation pipeline integrated with CIRCT. Starting from an example described using pattern matching and ADTs in Clash, point $\text{\textcircled{X}}$ shows how the high-level structure is preserved in CIRCT using the proposed `lc` dialect.

Clash, is fed into CIRCT through its Verilog frontend. While this isn't a full integration, we still benefit from all of the extra optimizations provided by CIRCT's LLHD/Moore compiler [14] used in the Verilog frontend, and its industry-strength SystemVerilog (SV) backend. The result is a more consistent and compact output, yet it retains the unintuitive signal names and design structure generated by Clash.

Rather than relying on the existing backend in Clash, another approach, marked as $\text{\textcircled{X}}$, is to directly convert the normalized System FC into CIRCT's core dialects. The abstraction level of the normalized form is close to that of the core dialects, allowing the translation between the two to be mostly syntactic. This is due to the normalization ensuring that all terms are representable in hardware, where, e.g., recursive data types are fully elaborated. With $\text{\textcircled{X}}$, we are no longer tied to backend-specific passes, and now benefit from the sophisticated global optimizations, and all of the other backends, provided for core dialects. While this is an improvement over $\text{\textcircled{X}}$, it is still a shallow integration of our frontend, as none of CIRCT's backends or optimizations have access to the original functional structure of the design. While the original structure directly captures the design's intent, normalization flattens the design making many optimizations difficult without further analyses.

2.2 Preserving High-Level Structure

In our two previous approaches, the source description was structurally flattened before reaching CIRCT, obscuring high-level intent and limiting the scope of downstream optimizations. We thus propose a final approach, which integrates Clash into CIRCT directly from the pre-normalized System FC, i.e. $\text{\textcircled{X}}$. This form preserves the original structure of the design, including all recursive data types and higher-order functions, which can enable otherwise difficult analyses when encoded in CIRCT.

To enable this approach, we introduce the `lc` (lambda calculus) speciality dialect which encodes ADTs, pattern

matching, and higher-order functions. Unlike other purely functional MLIR dialects, such as `rise` [15], `lc` captures the recursive and nested structure inherent in functional descriptions, in a manner that interoperates and lowers to existing hardware dialects, in particular the `fsm` dialect.

The `lc` dialect provides native support for ADTs via `lc.adt` and `lc.adt_case`. Concrete values are instantiated using `lc.adt_create` or `lc.adt_constant`. Individual ADT cases can include block arguments, which function as parameters that can be used during pattern matching, e.g.

```

lc.adt @State {
  lc.adt_case @Idle           %s = lc.adt_create @State
                               init @State::@Idle
  lc.adt_case @Busy(x: i8)    %i = lc.adt_constant
  lc.adt_case @Done(x: i8)    @State::@Idle
}

```

Pattern matching is supported via `lc.adt_match` and `lc.match_case`, as illustrated in Figure 1. Pattern matches can contain arbitrary logic, including nested pattern matches, using the case's block argument alongside other hardware dialects to produce a result with `lc.yield`. Lambdas can be defined using `lc.lambda`. Unlike `func.func`, these support functions as arguments and recursive value bindings using `lc.rec` and `lc.call`. As a result, the structure of `lc` allows us to identify high-level constructs such as FSMs and encode them using speciality dialects, e.g. the `fsm` dialect, which contain specific optimizations and have direct lowerings to several backends allowing for high-level structure to be retained in the outputs. More specifically, for the example in Figure 1, the FSM structure is identified by the recursive value binding of `% fsm` using `lc.rec`. The `%s` value in `% fsm` tracks the current state represented by each case of the pattern match. These cases yield two results, the next state of the FSM and their return value, whose logics become a transition guard and a state output respectively in the `fsm` dialect. A full example of how we encode a clash design using the `lc` dialect, can be found in Appendix §A.

```

1 %bar = lc.lambda (x: i8) -> (i8) {...}
2 lc.adt_match %s -> i8 {
3   lc.match_case @State::@Idle {
4     %r1 = lc.call %bar %x : i8
5     %r2 = lc.call %bar %y : i8
6     %r = comb.add bin %r1, %r2 : i8
7     lc.yield %r : i8 }
8   lc.match_case @State::@Busy %x {
9     %r = lc.call %bar %z : i8
10    lc.yield %r : i8 } ... }

```

Listing 1: Multiple calls to the same function.

3 High-level Optimizations

Method [3](#) highlights the importance of maintaining the high-level structure of our source design throughout compilation, as doing so enables otherwise difficult optimizations. We now describe 3 such optimizations.

Function Sharing. In Clash, each function application typically results in a distinct instantiation of that circuit [11], meaning that a function’s logic is aggressively inlined at its call site. Nevertheless, benefiting from the functional structure of the hardware, Clash can reuse identical functions via its Common Sub-expression Elimination pass, and those with different inputs across mutually exclusive match cases through its de-duplication pass. The `lc` dialect is capable of achieving the same level of function sharing found in Clash, as it preserves the original function call hierarchy. More importantly, this structural awareness allows such optimization to be more aggressive, all while avoiding a complicated de-duplication pass.

Concretely, to convert a function into a shared module, we start by performing a use-def analysis to retrieve the calling hierarchy. If the number of users is greater than a given threshold (defined as a pass parameter), we create a module where the inputs are the function parameters and the outputs are the function’s results. This new module is given a latency insensitive interface, i.e. using ready/valid signals, in order to better enable resource sharing. To share the module across several call sites, its inputs are paired with an identifier and guarded with a round-robin arbiter, enabled by the module’s input ready signal, which arbitrates the arguments from every call site. The caller then only considers the output from the function to be valid when the module’s output identifier matches the caller’s input identifier. Finally, we can handle higher-order functions by checking if the parent block of the caller is the same as that of the callee’s definition, if it is then no additional work is needed, if not, we need to add inputs and outputs to the caller’s parent module to allow for arguments, results, and identifiers to be connected to the callee.

To illustrate our optimization, consider Listing 1, in which a function `bar` is used three times in a pattern match. After a

use-def analysis, which shows that `bar` is used several times in our pattern match, we start by hoisting all of the calls into the same region as the `%bar` definition. This allows us to create a singleton instance of `%bar`, which we convert from a `lc.lambda` to a `hw.module`. We then convert our `lc.call` into `hw.wire` and use the arguments, with additional arbitration logic, as inputs to our instance whose outputs are then connected to the `hw.wire` operations.

Hierarchical Retiming. Retiming is a common optimization used during synthesis to improve Performance, Power, and Area (PPA) metrics in the resulting design [16]. Since the `lc` dialect retains all of the hierarchical information about our source design, retiming can be performed in a hierarchical manner. Consider a chain of function calls `foo1(foo2(foo3(...)))` which would potentially lead to a very long critical path. In a traditional retiming algorithm, we would flatten this chain and work with the end-to-end latency of our design, which might be complex and lead to a more conservative retiming being applied. However, in a hierarchical retiming, we can optimize our design in a bottom-up manner, retiming the inner-most call of our chain before using that result to more aggressively retime parent calls, thus leading to a less conservative result.

Identifying FSMs. As presented in a previous section, the `lc` dialect allows for recursive value bindings, typically used in Clash to represent FSMs, to be explicitly encoded using the `lc.lrec` operation. This allows us to retain the high-level intent throughout compilation by lowering all `lc.lrec` operations into `fsm.machine` in the `fsm` dialect, which can then be directly emitted by CIRCT’s SystemVerilog backend such that the FSM is identifiable by downstream tools.

4 Discussion and Conclusion

The goal of this work is to illustrate how a unified hardware compiler infrastructure, mostly centered around RTL and HLS, can be extended to natively support other high-level paradigms while also enabling new optimizations. Our extension to CIRCT is meant to be generic enough to be applicable beyond the context of Clash. Specifically, the `lc` dialect could serve any high-level hardware language supporting ADTs or higher-order functions, e.g. Spade or even standard Haskell (although all of the non-hardware lowerings from `lc` would still need to be added). Given the easily extensible nature of MLIR, we could also imagine other paradigms, e.g. message passing style languages like Anvil, can have their own native representations in CIRCT. While a deep integration is ideal, we show that shallow integrations with CIRCT can still yield immediate benefits while minimizing the engineering cost. We hope that this work can function as a motivator and guide on how new languages could be added to the CIRCT ecosystem, thus strengthening our shared infrastructure and paving the way for full hardware compiler unification, regardless of the source paradigm.

References

- [1] P. Bjesse, K. Claessen, M. Sheeran, and S. Singh, “Lava: hardware design in Haskell,” *SIGPLAN Not.*, vol. 34, no. 1, pp. 174–184, Sept. 1998, doi: [10.1145/291251.289440](https://doi.org/10.1145/291251.289440).
- [2] Arvind, “Bluespec: A language for hardware design, simulation, synthesis and verification Invited Talk,” in *Proceedings of the First ACM and IEEE International Conference on Formal Methods and Models for Co-Design*, in MEMOCODE '03. USA: IEEE Computer Society, 2003, p. 249.
- [3] J. Bachrach *et al.*, “Chisel: constructing hardware in a Scala embedded language,” in *Proceedings of the 49th Annual Design Automation Conference*, in DAC '12. San Francisco, California: Association for Computing Machinery, 2012, pp. 1216–1225. [Online]. Available: <https://doi.org/10.1145/2228360.2228584>
- [4] S. Jiang, P. Pan, Y. Ou, and C. Batten, “PyMTL3: A Python Framework for Open-Source Hardware Modeling, Generation, Simulation, and Verification,” *IEEE Micro*, vol. 40, no. 4, pp. 58–66, 2020, doi: [10.1109/MM.2020.2997638](https://doi.org/10.1109/MM.2020.2997638).
- [5] A. Ray, B. Devlin, F. Y. Quah, and R. Yesanatharao, “Hardcaml: An OCaml Hardware Domain-Specific Language for Efficient and Robust Design.” [Online]. Available: <https://arxiv.org/abs/2312.15035>
- [6] F. Skarman and O. Gustafsson, “Spade: An HDL Inspired by Modern Software Languages,” in *2022 32nd International Conference on Field-Programmable Logic and Applications (FPL)*, 2022, pp. 454–455. doi: [10.1109/FPL57034.2022.00075](https://doi.org/10.1109/FPL57034.2022.00075).
- [7] R. Nigam, P. H. Azevedo de Amorim, and A. Sampson, “Modular Hardware Design with Timeline Types,” *Proc. ACM Program. Lang.*, vol. 7, no. PLDI, June 2023, doi: [10.1145/3591234](https://doi.org/10.1145/3591234).
- [8] J. Z. Yu, A. R. Jha, U. Mathur, T. E. Carlson, and P. Saxena, “Anvil: A General-Purpose Timing-Safe Hardware Description Language.” [Online]. Available: <https://arxiv.org/abs/2503.19447>
- [9] C. Lattner *et al.*, “MLIR: scaling compiler infrastructure for domain specific computation,” in *Proceedings of the 2021 IEEE/ACM International Symposium on Code Generation and Optimization*, in CGO '21. Virtual Event, Republic of Korea: IEEE Press, 2021, pp. 2–14. doi: [10.1109/CGO51591.2021.9370308](https://doi.org/10.1109/CGO51591.2021.9370308).
- [10] S. Eldridge *et al.*, “MLIR as hardware compiler infrastructure,”
- [11] C. Baaij, “Digital circuit in ClaSH: functional specifications and type-directed synthesis,” PhD Thesis - Research UT, graduation UT, University of Twente, Netherlands, 2015. doi: [10.3990/1.9789036538039](https://doi.org/10.3990/1.9789036538039).
- [12] M. Sulzmann, M. Chakravarty, S. Peyton Jones, and K. Donnelly, “System F with type equality coercions,” 2007, pp. 53–66. doi: [10.1145/1190315.1190324](https://doi.org/10.1145/1190315.1190324).
- [13] S. P. Jones, K. Hammond, W. Partain, P. Wadler, C. B. Hall, and S. L. P. Jones, “The Glasgow Haskell Compiler: a technical overview,” 1993. [Online]. Available: <https://api.semanticscholar.org/CorpusID:56228824>
- [14] F. Schuiki, A. Kurth, T. Grosser, and L. Benini, “LLHD: a multi-level intermediate representation for hardware description languages,” in *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation*, in PLDI 2020. London, UK: Association for Computing Machinery, 2020, pp. 258–271. doi: [10.1145/3385412.3386024](https://doi.org/10.1145/3385412.3386024).
- [15] M. Lücke, M. Steuwer, and A. Smith, “Integrating a functional pattern-based IR into MLIR,” in *Proceedings of the 30th ACM SIGPLAN International Conference on Compiler Construction*, in CC 2021. Virtual, Republic of Korea: Association for Computing Machinery, 2021, pp. 12–22. doi: [10.1145/3446804.3446844](https://doi.org/10.1145/3446804.3446844).
- [16] C. E. Leiserson, F. M. Rose, and J. B. Saxe, “Optimizing Synchronous Circuitry by Retiming (Preliminary Version),” in *Third Caltech Conference on Very Large Scale Integration*, R. Bryant, Ed., Berlin, Heidelberg: Springer Berlin Heidelberg, 1983, pp. 87–116.

Appendix

A Full Conversion Example

Haskell.

```
1 bar :: (Num a, Eq a, Integral a) => a -> a
2 bar x = x * x * x
3
4 type I8 = Signed 8
5 data State = Idle | Busy I8 | Done I8
6     deriving (Generic, NFDataX)
7
8 fsm :: I8 -> State -> (State, I8)
9 fsm 0 Idle     = (Idle,    0)
10 fsm op Idle   = (Busy op,  0)
11 fsm 0 (Busy x) = (Busy x,  0)
12 fsm _ (Busy x) = (Done (bar x), 0)
13 fsm 0 (Done x) = (Done x,   x)
14 fsm _ _       = (Idle,    0)
15
16 controller
17   :: (HiddenClockResetEnable dom)
18   => Signal dom (Signed 8)
19   -> Signal dom (Signed 8)
20 controller input = mealy (flip fsm) Idle input
21
22 topEntity
23   :: Clock System
24   -> Reset System
25   -> Enable System
26   -> Signal System (Signed 8)
27   -> Signal System (Signed 8)
28 topEntity = exposeClockResetEnable controller
```

LC Dialect.

```
1 hw.module @topEntity (%op: i8) -> (i8) {
2   %bar = lc.lambda (%x: i8) -> (i8) {
3     %0 = comb.mul %x, %x, %x : i8
4     lc.yield %0: i8
5   }
6
7   lc.adt @State {
8     %idle = lc.adt_case @Idle
9     %busy = lc.adt_case @Busy(x: i8)
10    %done = lc.adt_case @Done(x: i8)
11    lc.yield %idle, %busy, %done
12  }
13
14  %fsm = lc.lambda (%state: !lc.adt_type<@State>, %op: i8)
15    -> (s: !lc.adt_type<@State>, i8) {
16    %rs, %ro = lc.adt_match %state -> (!lc.adt_type<@State>, i8) {
17      lc.match_case @State::@Idle {
18        %c0_i8 = hw.constant 0: i8
19        %op_0 = comb.icmp bin eq %op, %c0_i8 : i8
20        %rs_0 = lc.adt_constant @State::@Idle: !lc.adt_type<@State>
21        %rs_1 = lc.adt_constant @State::@Busy (x: %op)
22        : !lc.adt_type<@State>
23        %mux_state = comb.mux %op_0, %rs_0, %rs_1
24        : !lc.adt_type<@State>
25      }
26      lc.yield %mux_state, %c0_i8: !lc.adt_type<@State>, i8
27    }
28  }
```

```
27
28   lc.match_case @State::@Busy (%x : i8) {
29     %c0_i8 = hw.constant 0: i8
30     %op_0 = comb.icmp bin eq %op, %c0_i8 : i8
31     %rs_0 = lc.adt_constant @State::@Busy (x: %x)
32     : !lc.adt_type<@State>
33     %bar_x = lc.call %bar (%x) : i8
34     %rs_1 = lc.adt_constant @State::@Done (x: %bar_x)
35     : !lc.adt_type<@State>
36     %mux_state = comb.mux %op_0, %rs_0, %rs_1
37     : !lc.adt_type<@State>
38     lc.yield %mux_state, %c0_i8: !lc.adt_type<@State>, i8
39   }
40
41   lc.match_case @State::@Done (%x : i8) {
42     %c0_i8 = hw.constant 0: i8
43     %op_0 = comb.icmp bin eq %op, %c0_i8 : i8
44     %ro = comb.mux %op_0, %x, %c0_i8: i8
45     %rs_0 = lc.adt_constant @State::@Done (x: %x)
46     : !lc.adt_type<@State>
47     %rs_1 = lc.adt_constant @State::@Idle: !lc.adt_type<@State>
48     %mux_state = comb.mux %op_0, %rs_0, %rs_1
49     : !lc.adt_type<@State>
50     lc.yield %mux_state, %ro: !lc.adt_type<@State>, i8
51   }
52 }
53 lc.yield %rs, %ro
54 }
55
56 %controller = lc.lambda (%input: i8) -> (i8) {
57   %l_rec = lc.lrec %fsm ("sn", %input)
58   where ("sn": s: !lc.adt_type<@State>)
59     : !lc.lambda_type<<!lc.adt_type<@State>>, <i8>>
60   %i = lc.adt_constant @State::@Idle : !lc.adt_type<@State>
61   %v = lc.call %l_rec (%i) : i8
62   lc.yield %v : i8
63 }
64
65 %out = lc.call %controller (%input) : i8
66 hw.output %out: i8
67 }
68
```