

# CMSC838L - Final report

Arjun Vedantham  
Yusuf Bham

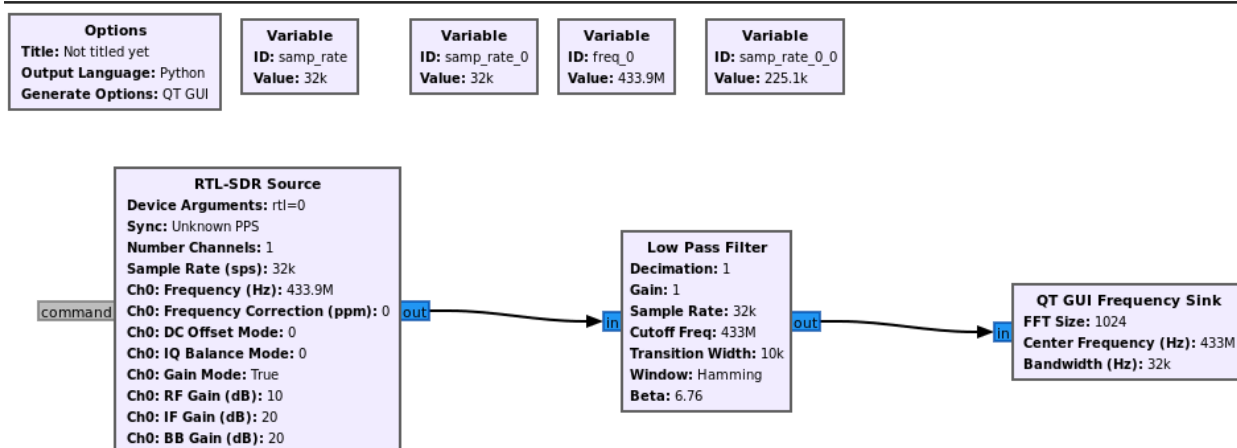
May 2024

## 1 Introduction and Motivation

Digital signal processing techniques are extremely important in telecommunications, computer vision, and a number of other related fields. In particular, digital signal processing techniques form a key component behind the idea of software defined radio (SDR), which refers to analyzing digital samples that represent radio signals with software, instead of discrete hardware components that operate over signals in analog formats. Software defined radio presents a notable improvement in flexibility for radio engineers, and removes the need for specialized hardware components - instead, new signal demodulation techniques or formats can be deployed just through a simple software update.

SDR users typically define signal processing pipelines using platforms like GNURadio. GNURadio presents a graphical format for creating these pipelines, with discrete blocks representing a signal source, signal sink, or an intermediate step in the processing pipeline. As an example, see the "flowgraph" (as called in GNURadio's documentation) below. This starts by instantiating a signal source from an RTL-SDR (a type of hobbyist SDR that can be used over USB with consumer PCs), and sets both a sampling rate (32000 Hz) and a listening frequency (signals at 433 MHz). From here, these samples are fed into a low pass filter block (which smooths out high frequency noise in the signal), and the resulting signal is transformed into a stream that is passed to a GUI block that graphs the signal, using a Fast Fourier Transform (FFT) to move the signal into the frequency domain.

This simple flowgraph is stored in an XML format, and is used by GNURadio's backend to generate a Python script that actually runs the defined processing pipeline. In addition, there are also C/C++ modules loaded into the runtime system for low level tasks - for instance, a USB driver for the SDR.



## 1.1 Problems Identified with State of the Art

There are number of problems with this current approach. First, GNURadio's practice of emitting Python scripts means that parallelism is limited on two fronts - first, because of language design choices (Python's infamous "global interpreter lock", which essentially forbids runtime concurrency), and also because we are ultimately running this script on standard PC hardware, which uses an inherently sequential von Neumann architecture. Additionally, as previously mentioned, there is also an extensive library of C/C++ modules, and the implementation of these modules are opaque to the Python-level code generated by GNURadio's flowgraph compiler.

As such, we identified two research questions that we aimed to answer in this project:

- Could we use hardware acceleration to get better performance and greater parallelism for DSP applications?
- Could we design a language that is more conducive to defining correct DSP pipelines?

## 2 Literature Review

We started by conducting a literature review of existing languages designed for DSP problems.

### 2.1 Ziria

One of the first papers we considered as Ziria [8], a domain specific language that was designed to aid development in implementations of the physical layer of wireless protocols. Ziria presented a functional language design syntax, and was specifically intended for wireless protocol implementations on IOT hardware. As such, it contained many primitives that we thought would be important to add to our language - for instance, an "FFT" primitive function. However Ziria had a key limitation - it was designed for standard CPUs, which meant that there were still parallelism limitations.

## 2.2 Calyx

One technique generally used for hardware acceleration for specialized applications like this is to deploy them to FPGAs. One example of this was the Catapult paper that we read in class, where Microsoft deployed FPGAs to accelerate running PageRank as part of their Bing server infrastructure. However, FPGAs are notoriously difficult to program, and generally require extremely fine-grained circuit configurations written in a hardware description language like Verilog.

Calyx [7] is an intermediate representation for compilers developed by the CAPRA research group at Cornell. It defines circuits in three distinct parts - a collection of memories (consisting of combinational memories/ flip-flops and registers), wire groups (which denote assignments between different memory components in the circuit), and a statically defined control schedule that orders wire assignments. Calyx has already been used in the Filament HDL, another domain specific language project from the CAPRA group that incorporated signal timing into the language's type system.

## 3 Technical Contribution

We decided to use the Calyx IR and design a more general, hardware acceleratable language for DSP tasks, called Zinnia. <sup>1</sup> We also considered emitting circuits using CIRCT, the LLVM framework that allows compilers to generate MLIR that is subsequently lowered to a Verilog hardware description, however we quickly found that CIRCT lacked the project maturity needed to develop even a basic language around it. This included no binaries to link against, thus requiring us to compile a large subset of the LLVM project, which was not practical. Calyx also had better documentation for its IR, and a small standard library of memory cells that could be easily integrated into circuits synthesized when the compiler lowers from the IR to a Verilog hardware description.

### 3.1 Parsing and Typing

Zinnia's parsing is a fairly standard approach. We used the Rust parser-combinator library `chumsky` for both the tokenization and subsequent parsing steps.

Zinnia is backed by a bidirectional type system, heavily inspired by Complete and Easy Bidirectional Typing for Higher-Rank Polymorphism and Purescript [2]. Also similarly to Purescript, while the system is based on CAEBDT, our system doesn't use the "main" takeaway of the paper – an ordered context, and stays with a more simple one. This was largely due to familiarity, and in the future we would like to rewrite it to use ordered contexts. The choice of a bidirectional type system stems from its ease of extensibility compared to HM, which is the more common choice for a language like this. This extensibility allows us to more easily add features like linear typing in the future, which is important in this space as it means we can keep the *pure* and *functional* aspect, while still preserving performance and memory constraints. Currently types are restricted to functions and primitives, just

---

<sup>1</sup>Named "Zinnia" because "calyx" refers to the petals of a flower and zinnias are a type of flower. Plus it sounds like Ziria.

due to time constraints. General universal quantification is supported, but access to it is currently limited to the `scan` primitive. While generics over types are generally supported, being generic over a *value* is currently restricted to the vector primitive.

Full semantics and syntax for our language can be found at the end of our report.

### 3.1.1 Ease of Use

One key consideration when designing the typechecker and parser was support for rich error reports. The parser has support for error backtracking, allowing for better reports on syntax. Similarly, the typechecker preserves as much location info as possible in order to give full-featured errors.

```

Error: Mismatched operand types!
[test_files/basic4.z:1:1]
1 let main: i8 = (let ((x: i8, (+ 1u8 4i8))) (+ x 3));
                                |
                                | this had type u8
                                |
                                | while this had type i8
                                |
                                | while checking expression had type i8

```

Figure 1: An error report for mismatched operands

## 3.2 Code Generation

The compiler that we wrote for Zinnia traverses the AST recursively. Since there is no concept of dynamic memory allocation for circuits, we focused on generating the required memories for the circuit first.

At the top level, "let" bindings are used to generate register memories (for integers) and combinational memories (for vectors). For constant values, we use Calyx's "structure" macro to instantiate a "constant" register with the required value at compile time, which is subsequently removed in the process of optimization and lowering to Verilog. Variable bindings are stored in a hashmap and are accessed using either the component name (which is guaranteed to be unique through the Calyx frontend API) or the user-defined variable name.

For binary operations or if expressions, we need to use comparator primitives and wire them correctly - to do this, we call the memory generation function on each of the parameters of the operation, and then use a "wire" generation function which generates assignments between the inputs and outputs of the memories and logical primitives. After this is complete, we generate a small schedule of operations depending on the subexpression that we are compiling, using the Calyx "seq" block to sequentially order operations (e.g. ensuring that a register memory gets its constant value before the addition actually takes place). Finally, the top most register memory's value is copied into a specially instantiated combinational memory. This is because register memories in Calyx do not have latching/persistence guar-

antees, however, values in combinational/vector memories can be examined as part of the circuit simulation.

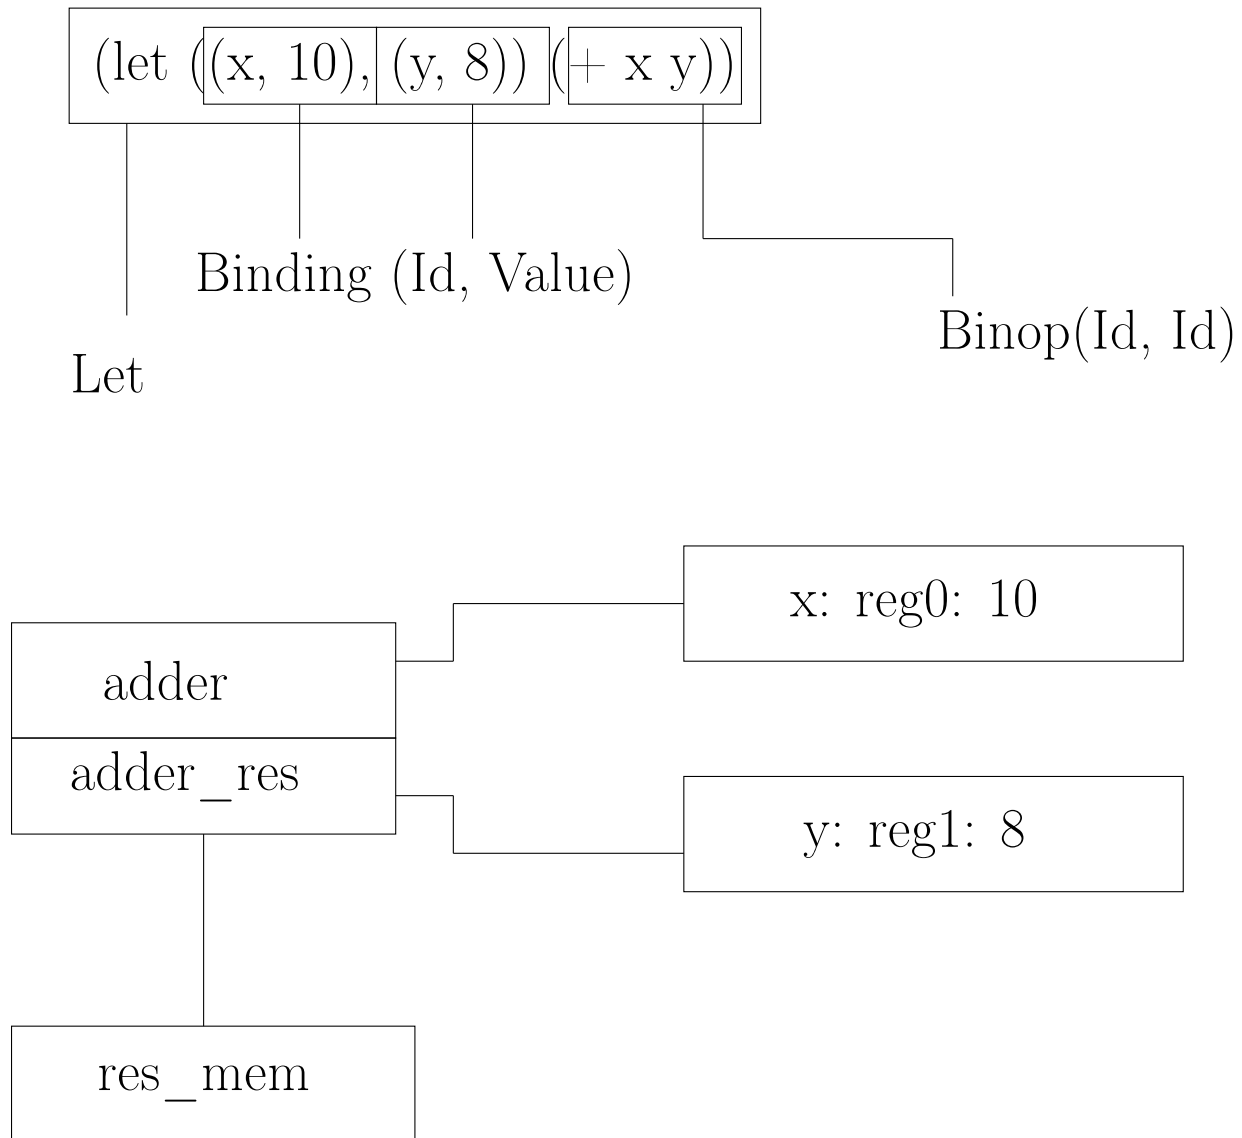


Figure 2: Simplified compilation/memory instantiation flow for a Zinnia program.

### 3.3 Supporting Parallelization in Compilation

One of the main goals of our DSL was to improve parallelism. While Calyx has a "par" primitive for parallelizing parts of the circuit, Zinnia currently only generates circuits using the "seq" ("sequential") control flow keyword for easier debugging and testing. Even though parallelism at the user level is limited, we still hoped to expose some of the potential parallel speedups through the use of a function primitive that could be used to get implicit parallelism.

### 3.3.1 Prefix Sums Scan

To expose some parallelism without explicitly supporting compiling using the "par" control flow block, we implemented a version of parallel scan using the prefix sums algorithm [1]. This primitive takes a reference to a vector with 8 integers and performs the addition operation, returning a modified array where the latter half of the array contains the scan result.

This function was supposed to be invoked as a primitive within the language, with the compiler inlining a separate Calyx component whose input memories would be populated with a reference to the vector supplied by the caller. However, while we were able to test the scan component as a separate circuit generated from our hand-written IR code, we were not able to fully resolve linking issues that would allow us to pass memories by reference into subcomponents.

```
arjun@legion5:~/coursecode/cnsc8381/verilog-test$ fud exec
{
  "cycles": 71,
  "memories": {
    "input_arr": [
      3,
      1,
      7,
      0,
      4,
      1,
      6,
      3
    ],
    "output_final": [
      0,
      0,
      11,
      0,
      4,
      11,
      16,
      0,
      3,
      4,
      11,
      11,
      15,
      16,
      22,
      25
    ]
  }
}
arjun@legion5:~/coursecode/cnsc8381/verilog-test$
```

Figure 3: Screenshot of the final state of the prefix sums array - a separate vector memory is instantiated in the IR implementation to track the intermediate values as they flow up and down the prefix-sums tree.

## 4 Evaluation

We decided to evaluate our language in three different ways: correctness testing, scalability testing, and hardware testing.

### 4.1 Correctness Testing

To verify the correctness of the language, we attempted to use the Cocotb RTL testing framework and the Icarus Verilog simulator. This allows us to run a simulated version of the circuit's clock and write unit tests around elements of the circuits that we generate. Although we did not have time to implement extensive tests, we were able to informally verify correct




Figure 4: Example timing diagram for the program shown in Figure 1. The operands (represented in hex form here) are loaded into `reg0` and `reg1`, before being loaded into the "left" and "right" inputs of the combinational adder. The adder register's result (5th line down) is then moved into the "output" combinational memory using its "write\_data" line. Finally, the last line displays the clock signal for the whole circuit.

execution using unit tests for the binary operations supported by the language, as well as correct outputs values for the scan primitive (which was tested separately since we could not link this into the language easily) through Icarus Verilog. We were also able to generate timing diagrams that we could view with GTKWave, and this showed that values used during the circuit's evaluation were being loaded into registers correctly. However, we would have wanted to use randomized testing with cocotb to be more confident in our compiler's correctness.

## 4.2 Scalability Testing

Second, we wanted to test the scalability of our language. For this, we focused on evaluating the "scan" primitive, since it is the main way end-users would be able to obtain parallelism through the language (at least, as it is currently implemented). We found that in simulation, our parallel scan implementation always took 71 cycles to generate outputs, given a set of 8 integers to iterate over. While parallel scan theoretically provides work bounds of  $O(n)$  and span bounds of  $O(\lg(n))$ , in practice the work and span bounds of our implementation are worse. This is because we were not sure if it was possible to instantiate an arbitrarily high number of registers, so we focused on using the combinational memory primitives to hold intermediate values (particularly between the "sweep-up" and "sweep-down" phases of parallel scan). Combinational memory primitives in the Calyx IR are exclusive read/write memory cells, and take at least one cycle to set the data address and read the stored value out of the memory into a register. This restriction essentially limited the parallelism that we could achieve to only a few steps of the algorithm - specifically, summing values between nodes at each level of the prefix sums tree, since the additions were purely combinational and depended only on values in registers that were ready to go at the last cycle.

## 4.3 Physical Testing

We hoped to eventually deploy circuits that were synthesized using Zinnia to a Lattice Icestick FPGA. This FPGA can be used with consumer devices over a USB connection, and has a large library of open source tooling. This includes the Yosys synthesis toolkit (which handles transforming high level Verilog hardware descriptions into lower level RTL circuits), the nextpnr place and route tool (which handles connections between logic units on the FPGA), and iceprog, an open source programming tool that could deploy circuits

with complete routing to the flash memory on-board the Icestick.

Circuits synthesized by Calyx include three signals by default - a "go" line, which when activated, activates the circuit, a "clk" (clock) line, which handles cycle-level signalling for the sequential logic elements of the circuit, and a "reset" line, which can be used to reset the circuit while active.

In order to actually use the synthesized circuits, the "go" line must be pulled to logical high, and while the reset lines needed to be set to logical low, and the clock line must be tied to the 12 MHz oscillator built into the Icestick. Values for the "scan" primitive could then be loaded in from the Icestick's block RAM module.

To support dynamically loading values into block RAM, we wrote a small serial communication driver for the Icestick. Specifically, this was a form of one way (half-duplex) SPI, with an Arduino serving as the primary device on the SPI bus, and the Icestick serving as the secondary device. We chose SPI because it is relatively easy to implement on the receiving side (in fact, a shift register is enough to receive data transmitted over the data line). While testing the SPI driver, we wanted to use an external clock source to control the speed of data transmission, and ensure that the Icestick was sampling the data line in phase with the transmitting device's clock (in this case an Arduino). Unfortunately, setting the clock line for the circuit loaded onto the FPGA also sets the clock line for the USB programmer chip, which requires a 12MHz signal in order to function correctly. This misconfiguration prevented us from loading new circuits onto the FPGA, and required us to use simulation as our primary evaluation method instead.




Figure 5: Half duplex SPI waveform, where an Arduino would have controlled the SCK and SDO lines and the FPGA would bit shift the received values in block RAM.

[6]

## 5 Conclusion

In the end, we were able to implement a version of the scan primitive in the Calyx IR, implement a bidirectional typing system, add some basic language features into our language,



and use external circuit verification and simulation tools to show that our circuits performed as the high level program specified.

However, there is still room for improvement, both in terms of feature support and performance - we do not have support for looping structures, even though the Calyx IR supports it. Additionally, our parallelism benefits are limited by the fact that we can only really exploit the parallelism present in the hardware through the scan primitive, even though Calyx supports parallelized control schedules. While part of this is due to limitations in the Calyx standard library, future libraries could be used to enable concurrent read/exclusive write accesses to these memories instead of the current exclusive read/write access pattern.

We were also not able to show our circuits running on real hardware, which was the ultimate goal of the project. Our current version of Zinnia would not be usable with real software defined radio/signal processing workloads since our hardware interfaces are not finished yet. However, we feel that there is enough expressivity in the language (e.g. customizable bit widths, a clean and functional syntax) that could eventually make it practical for creating correct signal processing pipelines.

One obstacle that we did not foresee was the difficulty of working with the Calyx IR from Rust. While LLVM's CIRCT project has a robust Rust API, it seems that the authors of the Calyx IR primarily want language authors to use their IR from a Python builder module or the CIRCT toolchain, even though the Filament HDL (which was authored by the same research group and uses Calyx as its IR) is also written in Rust. For some issues, like linking against the combinational memory primitives or inlining other components, the Rust library documentation was so poor that we ended up looking through Filament's implementation and using that as a guide for building our compiler.

Altogether, this was a very interesting exploration into building a domain specific language that interacted so closely with hardware, and was a great way to get experience with compilers, parallel algorithms, type systems, and using FPGAs. We hope to continue working on Zinnia as a hobby research project over the summer.

The codebase for Zinnia is available on GitHub: <https://github.com/javathunderman/zinnia>

## 6 Syntax

$\langle decl \rangle \quad ::= \text{let } \langle ident \rangle : \langle ty \rangle = \langle expr \rangle ;$

$\langle ident \rangle \quad ::= (\langle letter \rangle \mid \_)\{\langle letter \rangle \mid 0..9 \mid \_ \}$

$\langle ty \rangle \quad ::= ()$   
 $\quad \quad \mid \text{bool}$   
 $\quad \quad \mid \langle num\text{-}ty \rangle$   
 $\quad \quad \mid \text{Vec} \langle \langle num\text{-}ty \rangle, \{0..9\}^+ \rangle$

$\langle num\text{-}ty \rangle \quad ::= (\text{u} \mid \text{i}) \{0..9\}^+$

$\langle expr \rangle$  ::=  $()$   
|  $\langle bool-lit \rangle$   
|  $\langle int-lit \rangle$   
|  $\langle vec-lit \rangle$   
|  $(\{\langle expr \rangle\}^+)$   
|  $(\langle expr \rangle \langle bop \rangle \langle expr \rangle)$   
|  $(\mathbf{let} (\langle binders \rangle) \langle expr \rangle)$   
|  $(\mathbf{if} \langle expr \rangle \langle expr \rangle \langle expr \rangle)$

$\langle binders \rangle$  ::=  $\langle binder \rangle$   
|  $\langle binder \rangle, \langle binders \rangle$

$\langle binder \rangle$  ::=  $(\langle ident \rangle, \langle expr \rangle)$   
|  $(\langle ident \rangle: \langle ty \rangle, \langle expr \rangle)$

$\langle num-cmp-op \rangle$  ::=  $>$   
|  $>=$   
|  $<$   
|  $<=$

$\langle num-op \rangle$  ::=  $+$   
|  $-$   
|  $*$   
|  $/$

$\langle cmp-op \rangle$  ::=  $==$   
|  $!=$

$\langle bop \rangle$  ::=  $\langle num-op \rangle$   
|  $\langle num-cmp-op \rangle$   
|  $\langle cmp-op \rangle$

$\langle int-lit \rangle$  ::=  $[-]\{0..9\}^+$   
|  $[-]\{0..9\}^+ (\mathbf{u} | \mathbf{i}) \{0..9\}^+$

$\langle bool-lit \rangle$  ::=  $\mathbf{true}$   
|  $\mathbf{false}$

$\langle vec-lit \rangle$  ::=  $[\langle vec-elems \rangle]$

$$\langle \text{vec-elems} \rangle ::= \langle \text{int-lit} \rangle$$

$$| \langle \text{int-lit} \rangle, \langle \text{vec-elems} \rangle$$

## 7 Typing

Types	$A, B, C ::= ()$	$  \alpha   \hat{\alpha}$	$  \forall \alpha. A$	$  A_1, \dots, A_n \rightarrow B$
Monotypes	$\tau, \sigma, \pi ::= ()$	$  \alpha   \hat{\alpha}$	$  \tau_1, \dots, \tau_n \rightarrow \sigma$	$  \vartheta   \vartheta_{i,s}   \hat{\vartheta}   \hat{\vartheta}_{i,s}$
			$  \text{Vec}(\tau, c)   \widehat{\text{Vec}}$	$  \text{Bool}$
Contexts	$\Gamma, \Delta, \Theta ::= \cdot$	$  \Gamma, \alpha   \Gamma, (x : A)$		

Figure 6: Syntax of types, monotypes, and contexts.

$\boxed{\Gamma \vdash A}$  Under context  $\Gamma$ , type  $A$  is well-formed.

$\boxed{\Gamma \vdash_{i=1}^n A}$  Under context  $\Gamma$ , types  $\{A_i\}_1^n$  are well-formed.

$\frac{}{\Gamma \vdash ()}$ UnitWF	$\frac{}{\Gamma \vdash \text{Bool}}$ BoolWF	$\frac{n > 0}{\Gamma \vdash \vartheta_n^i}$ SNumWF	$\frac{n > 0}{\Gamma \vdash \vartheta_n^u}$ UNumWF
$\frac{\Gamma \vdash \tau \quad \text{num}(\tau) \quad c > 0}{\Gamma \vdash \text{Vec}(\tau, c)}$ VecWF		$\frac{\Gamma \vdash_{i=1}^n A_i \quad \Gamma \vdash B}{\Gamma \vdash A_1, \dots, A_n \rightarrow B}$ ArrowWF	
$\frac{\Gamma, \alpha \vdash A}{\Gamma \vdash \forall \alpha. A}$ ForAllWF		$\frac{\Gamma \vdash A_1 \quad \vdots \quad \Gamma \vdash A_n}{\Gamma \vdash_{i=1}^n A}$ WFs	

Figure 7: Well-formedness of types

$\boxed{\Gamma \vdash A \sqcap B \Rightarrow C \dashv \Delta}$  Under context  $\Gamma$ , types  $A$  and  $B$  unify to  $C$  with output context  $\Delta$ .

$\boxed{\Gamma \vdash A_i \sqcap B_i \xrightarrow[n]{\Rightarrow} C_i \dashv \Delta}$  Under context  $\Gamma$ , types  $A_i$  and  $B_i$  unify to  $C_i$  with output context  $\Delta$ .

$$\frac{\Gamma \vdash A}{\Gamma \vdash A \sqcap A \Rightarrow A \dashv \Gamma} \text{EqUni}$$

$$\frac{\Gamma \vdash \text{Vec}(\tau, c) \quad \Gamma \vdash \text{Vec}(\sigma, c) \quad \Gamma \vdash \tau \sqcap \sigma \Rightarrow \pi \dashv \Delta}{\Gamma \vdash \text{Vec}(\tau, c) \sqcap \text{Vec}(\sigma, c) \Rightarrow \text{Vec}(\pi, c) \dashv \Delta} \text{VecUni}$$

$$\frac{\begin{array}{c} \Gamma \vdash A_1, \dots, A_n \rightarrow R_A \quad \Gamma \vdash B_1, \dots, B_n \rightarrow R_B \\ \Gamma \vdash A_i \sqcap B_i \xrightarrow[n]{\Rightarrow} C_i \dashv \Delta \quad \Gamma \vdash R_A \sqcap R_B \Rightarrow R_C \dashv \Delta \end{array}}{\Gamma \vdash A_1, \dots, A_n \rightarrow R_A \sqcap B_1, \dots, B_n \rightarrow R_B \Rightarrow C_1, \dots, C_n \rightarrow R_C \dashv \Delta} \text{ArrowUni}$$

$$\frac{\Gamma \vdash \hat{\alpha} \quad \Gamma \vdash \hat{\beta}}{\Gamma \vdash \hat{\alpha} \sqcap \hat{\beta} \Rightarrow \hat{\alpha} \dashv \Delta} \text{ExAnyUni}$$

$$\frac{\Gamma \vdash \hat{\vartheta}_1 \quad \Gamma \vdash \hat{\vartheta}_2}{\Gamma \vdash \hat{\vartheta}_1 \sqcap \hat{\vartheta}_2 \Rightarrow \hat{\vartheta}_1 \dashv \Delta} \text{ExUNumUni}$$

$$\frac{\begin{array}{c} \Gamma \vdash \hat{\vartheta}_{s_1, c_1} \quad \Gamma \vdash \hat{\vartheta}_{s_2, c_2} \\ s_3 = s_1 \wedge s_2 \quad c_3 = \max(c_1, c_2) \end{array}}{\Gamma \vdash \hat{\vartheta}_{s_1, c_1} \sqcap \hat{\vartheta}_{s_2, c_2} \Rightarrow \hat{\vartheta}_{s_3, c_3} \dashv \Delta} \text{ExSSNumUni}$$

$$\begin{array}{c}
\frac{\Gamma \vdash \widehat{\vartheta}_{s_1, c_1} \quad \Gamma \vdash \widehat{\vartheta}}{\Gamma \vdash \widehat{\vartheta}_{s_1, c_1} \cap \widehat{\vartheta} \Rightarrow \widehat{\vartheta}_{s_1, c_1} \dashv \Delta} \text{ExSUNumUni} \\
\\
\frac{\Gamma \vdash \widehat{\text{Vec}}_1 \quad \Gamma \vdash \widehat{\text{Vec}}_2}{\Gamma \vdash \widehat{\text{Vec}}_1 \cap \widehat{\text{Vec}}_2 \Rightarrow \widehat{\text{Vec}}_1 \dashv \Delta} \text{ExVecUni} \\
\\
\frac{\Gamma \vdash \widehat{\text{Vec}}_\alpha \quad \Gamma \vdash \text{Vec}(\tau, c)}{\Gamma \vdash \widehat{\text{Vec}}_\alpha \cap \text{Vec}(\tau, c) \Rightarrow \text{Vec}(\tau, c) \dashv \Delta} \text{ExVecCUni} \\
\\
\frac{\Gamma \vdash \widehat{\vartheta} \quad \Gamma \vdash \vartheta_{s, c}}{\Gamma \vdash \widehat{\vartheta} \cap \vartheta_{s, c} \Rightarrow \vartheta_{s, c} \dashv \Delta} \text{ExUNumCUni} \\
\\
\frac{\Gamma \vdash \widehat{\vartheta}_{s_1, c_1} \quad \Gamma \vdash \vartheta_{s_2, c_2} \quad \neg s_1 \vee s_2 \quad c_1 \leq c_2}{\Gamma \vdash \widehat{\vartheta}_{s_1, c_1} \cap \vartheta_{s_2, c_2} \Rightarrow \vartheta_{s_2, c_2} \dashv \Delta} \text{ExSNumCUni} \\
\\
\frac{\Gamma \vdash \widehat{\alpha} \quad \Gamma \vdash \alpha}{\Gamma \vdash \widehat{\alpha} \cap \alpha \Rightarrow \alpha \dashv \Delta} \text{ExAnyCUni} \\
\\
\frac{\Gamma \vdash \forall \alpha. A \quad \Gamma \vdash B \quad \Gamma, \widehat{\alpha} \vdash A[\alpha := \widehat{\alpha}] \cap B \Rightarrow B \dashv \Delta}{\Gamma \vdash \forall \alpha. A \cap B \Rightarrow B \dashv \Delta} \text{ForAllUni} \\
\\
\frac{\Gamma \vdash A \cap B \Rightarrow C \dashv \Delta}{\Gamma \vdash B \cap A \Rightarrow C \dashv \Delta} \text{CommutUni} \\
\\
\frac{\Gamma \vdash A_1 \cap B_1 \Rightarrow C_1 \dashv \Theta_1 \quad \Theta_1 \vdash A_2 \cap B_2 \Rightarrow C_2 \dashv \Theta_2 \quad \vdots \quad \Theta_{n-1} \vdash A_n \cap B_n \Rightarrow C_n \dashv \Theta_n}{\Gamma \vdash A_i \cap B_i \xrightarrow[n]{\Rightarrow} C_i \dashv \Theta_n} \text{Unis}
\end{array}$$

Figure 8: Unification of types

$\boxed{\Gamma \vdash e \Leftarrow A \dashv \Delta}$  Under input context  $\Gamma$ ,  $e$  checks against input type  $A$ , with output context  $\Delta$

$\boxed{\Gamma \vdash e \Rightarrow A \dashv \Delta}$  Under input context  $\Gamma$ ,  $e$  synthesizes output type  $A$ , with output context  $\Delta$

$\boxed{\Gamma \vdash e_i \xRightarrow{i=1}^n A_i \dashv \Delta}$  Under input context  $\Gamma$ ,  $\{e_i\}_1^n$  synthesizes types  $\{A_i\}_1^n$ , with output context  $\Delta$

$$\frac{\Gamma \vdash e \Rightarrow A \dashv \Theta \quad \Theta \vdash A \cap B \Rightarrow C \dashv \Delta}{\Gamma \vdash e \Leftarrow B \dashv \Delta} \text{UniToCheck}$$

$$\frac{\Gamma \vdash \Gamma[x] \Rightarrow A \dashv \Delta}{\Gamma \vdash x \Rightarrow A \dashv \Delta} \text{IdSyn}$$

$$\frac{x = \text{true} \vee x = \text{false}}{\Gamma \vdash x \Rightarrow \text{Bool} \dashv \Gamma} \text{BoolSyn}$$

$$\frac{x = \text{NuC} \quad C > 0 \quad 0 \leq N < 2^C}{\Gamma \vdash x \Rightarrow \vartheta_{u,C} \dashv \Gamma} \text{SNumLitAnnSyn}$$

$$\frac{x = \text{NiC} \quad C > 0 \quad -2^{C-1} \leq N \leq 2^{C-1} - 1}{\Gamma \vdash x \Rightarrow \vartheta_{i,C} \dashv \Gamma} \text{UNumLitAnnSyn}$$

$$\frac{x = N \quad N \in \mathbb{Z} \quad \text{sz} = \text{min-width}(N) \quad \text{sgn} = \text{sign}(N)}{\Gamma \vdash x \Rightarrow \widehat{\vartheta}_{\text{sgn}, \text{sz}} \dashv \Delta} \text{NLitSyn}$$

$$\frac{\Gamma \vdash e \Rightarrow \tau \dashv \Delta \quad \Delta \vdash \text{Vec}(\tau, 1)}{\Gamma \vdash [e] \Rightarrow \text{Vec}(\tau, 1) \dashv \Delta} \text{Vec1LitSyn}$$

$$\frac{\begin{array}{l} \Gamma \vdash [e_1, \dots, e_n] \Rightarrow \text{Vec}(\tau, i) \dashv \Theta_1 \\ \Theta_1 \vdash e_{n+1} \Rightarrow \sigma \dashv \Theta_2 \\ \Theta_2 \vdash \sigma \cap \tau \Rightarrow \pi \dashv \Delta \end{array}}{\Gamma \vdash [e_1, \dots, e_{n+1}] \Rightarrow \text{Vec}(\pi, i+1) \dashv \Delta} \text{VecLitSyn}$$

$$\frac{\begin{array}{l} \text{op} \in [+ , - , / , \cdot] \\ \Gamma \vdash l \Rightarrow \tau \dashv \Theta \quad \Gamma \vdash r \Rightarrow \sigma \dashv \Theta \\ \text{num}(\tau) \quad \text{num}(\sigma) \\ \Theta \vdash \tau \cap \sigma \Rightarrow \pi \dashv \Delta \end{array}}{\Gamma \vdash l \text{ op } r \Rightarrow \pi \dashv \Delta} \text{NumFieldOpSyn}$$

$$\begin{array}{c}
\text{op} \in [<, \leq, >, \geq] \\
\Gamma \vdash l \Rightarrow \tau \dashv \Theta \quad \Gamma \vdash r \Rightarrow \sigma \dashv \Theta \\
\text{num}(\tau) \quad \text{num}(\sigma) \\
\Theta \vdash \tau \Leftarrow \sigma \dashv \Delta \\
\hline
\Gamma \vdash l \text{ op } r \Rightarrow \text{Bool} \dashv \Delta \quad \text{NumCmpOpSyn}
\end{array}$$

$$\begin{array}{c}
\text{op} \in [=, !=] \\
\Gamma \vdash l \Rightarrow \tau \dashv \Theta \quad \Gamma \vdash r \Rightarrow \sigma \dashv \Theta \\
\Theta \vdash \tau \Leftarrow \sigma \dashv \Delta \\
\hline
\Gamma \vdash l \text{ op } r \Rightarrow \text{Bool} \dashv \Delta \quad \text{EqOpSyn}
\end{array}$$

$$\begin{array}{c}
\Gamma \vdash f \Rightarrow F \dashv \Theta_1 \\
\Theta_1 \vdash e_i \Rightarrow A_i \dashv \Theta_2 \\
\Theta_2 \vdash F \cap A_1, \dots, A_n \rightarrow R_A \Rightarrow B_1, \dots, B_n \rightarrow R_B \dashv \Delta \\
\hline
\Gamma \vdash (f \ e_1 \ \dots \ e_n) \Rightarrow R_B \dashv \Delta \quad \text{CallSyn}
\end{array}$$

$$\begin{array}{c}
\Gamma \vdash c \Leftarrow \text{Bool} \dashv \Theta_1 \\
\Theta_1 \vdash \text{br\_t} \Rightarrow A \dashv \Theta_2 \quad \Theta_2 \vdash \text{br\_f} \Rightarrow B \dashv \Theta_3 \\
\Theta_3 \vdash A \cap B \Rightarrow C \dashv \Delta \\
\hline
\Gamma \vdash (\text{if } c \ \text{br\_t} \ \text{br\_f}) \Rightarrow C \dashv \Delta \quad \text{IfSyn}
\end{array}$$

$$\begin{array}{c}
\Gamma \vdash b \Rightarrow B \dashv \Delta \\
\hline
\Gamma \vdash (\text{let } () \ b) \Rightarrow B \dashv \Delta \quad \text{Let0Syn}
\end{array}$$

$$\begin{array}{c}
\Gamma \vdash e_1 \Rightarrow A \dashv \Theta \\
\Theta, (x:A) \vdash (\text{let } (\dots) \ b) \Rightarrow B \dashv \Delta \\
\hline
\Gamma \vdash (\text{let } ((x, e_1) \ \dots) \ b) \Rightarrow B \dashv \Delta \quad \text{LetSyn}
\end{array}$$

$$\begin{array}{c}
\Gamma \vdash e_1 \Leftarrow T \dashv \Theta \\
\Theta, (x:T) \vdash (\text{let } (\dots) \ b) \Rightarrow B \dashv \Delta \\
\hline
\Gamma \vdash (\text{let } ((x:T, e_1) \ \dots) \ b) \Rightarrow B \dashv \Delta \quad \text{LetAnnSyn}
\end{array}$$

$$\begin{array}{c}
\hline
\Gamma \vdash \text{scan} \Rightarrow \forall \text{Vec}_\alpha. \text{Vec}_\alpha \rightarrow \text{Vec}_\alpha \dashv \Delta \quad \text{ScanSyn}
\end{array}$$

Figure 9: Inference and checking of types

## 8 Zinnia Examples

```
let main: u8 = (let ((x : Vec<i8 , 8>, [1, 4, 5, 7, 1, 2, 3, 4])) 2);
```

Example of using vector types - "main" is defined as returning an 8 bit integer, x is defined as getting a vector of eight 8-bit integers, and the two is returned as the result.

```
let main: i8 = (let ((x: i8 , (+ 1u8 4i8))) (+ x 3));
```

Example of a mismatch caught by the type checker - 1 (an unsigned int) and 4 (a signed int) cannot be added together to get a signed int as a result.



## References

- [1] Guy E. Blelloch. *Prefix Sums and Their Applications*. CMU-CS-90-190. School of Computer Science, Carnegie Mellon University, Nov. 1990.
- [2] Jana Dunfield and Neelakantan R. Krishnaswami. “Complete and easy bidirectional typechecking for higher-rank polymorphism”. In: *Proceedings of the 18th ACM SIGPLAN international conference on Functional programming* (2013). URL: <https://api.semanticscholar.org/CorpusID:7586176>.
- [3] David Durst et al. “Type-directed scheduling of streaming accelerators”. In: *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation*. PLDI ’20: 41st ACM SIGPLAN International Conference on Programming Language Design and Implementation. London UK: ACM, June 11, 2020, pp. 408–422. ISBN: 978-1-4503-7613-6. DOI: 10.1145/3385412.3385983. URL: <https://dl.acm.org/doi/10.1145/3385412.3385983> (visited on 05/14/2024).
- [4] Matthew Gordon.  $\mu$ : *A Functional Programming Language for Digital Signal Processing*. Apr. 9, 2003. URL: <https://www.cs.unb.ca/tech-reports/honours-theses/Matthew.Gordon-4997.pdf>.
- [5] Andrew Lenharth and Chris Lattner. *CIRCT: Lifting hardware development out of the 20th century*. Nov. 17, 2021. URL: <https://lvm.org/devmtg/2021-11/slides/2021-CIRCT-LiftingHardwareDevOutOfThe20thCentury.pdf>.
- [6] Mathworks. *Support SPI Communication*. URL: [https://www.mathworks.com/help/simulink/supportpkg/raspberrypi\\_ug/support-spi-communication.html](https://www.mathworks.com/help/simulink/supportpkg/raspberrypi_ug/support-spi-communication.html).
- [7] Rachit Nigam et al. “A compiler infrastructure for accelerator generators”. In: *Proceedings of the 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*. ASPLOS ’21: 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems. Virtual USA: ACM, Apr. 19, 2021, pp. 804–817. ISBN: 978-1-4503-8317-2. DOI: 10.1145/3445814.3446712. URL: <https://dl.acm.org/doi/10.1145/3445814.3446712> (visited on 05/11/2024).
- [8] Gordon Stewart et al. “Ziria: A DSL for Wireless Systems Programming”. In: *Proceedings of the Twentieth International Conference on Architectural Support for Programming Languages and Operating Systems*. ASPLOS ’15: Architectural Support for Programming Languages and Operating Systems. Istanbul Turkey: ACM, Mar. 14, 2015, pp. 415–428. ISBN: 978-1-4503-2835-7. DOI: 10.1145/2694344.2694368. URL: <https://dl.acm.org/doi/10.1145/2694344.2694368> (visited on 05/11/2024).