

VyZX: A Vision for Verifying the ZX Calculus

Adrian Lehmann*

University of Chicago

adrianlehmann@uchicago.edu

Ben Caldwell*

University of Chicago

caldwellb@uchicago.edu

Robert Rand

University of Chicago

rand@uchicago.edu

Optimizing quantum circuits is a key challenge for quantum computing. The PyZX compiler broke new ground by optimizing circuits via the ZX calculus, a powerful graphical alternative to the quantum circuit model. Still, it carries no guarantee of its correctness. To address this, we developed VyZX, a verified ZX-calculus in the Coq proof assistant. VyZX provides two distinct representations of ZX diagrams for ease of programming and proof: A graph-based representation for writing high-level functions on diagrams and a block-based representation for proving ZX diagrams equivalent. Through these two different views, VyZX provides the tools necessary to verify properties and transformations of ZX diagrams. This paper explores the proofs and design choices underlying VyZX and its application and the challenges of verifying a graphical programming language.

1 Introduction

As quantum computers transition from fiction to a feature of our daily lives, there has been a surge of interest in quantum optimizers [21, 23, 25, 16, 1, 29]. The goal of a quantum optimizer is to reduce the number of bottlenecks in a quantum circuit, whether those be two-qubit gates in the near term or T gates in the longer term. Many of these optimizers do some form of model checking [23, 29] or translation validation [25, 16] to gain confidence that their optimizations are correct, out of awareness that bugs in quantum optimizers are both common and costly [16]. Of particular note, the VOQC compiler [14] is fully verified in the Coq proof assistant, guaranteeing that its optimizations preserve the semantics of the original quantum circuit.

Unfortunately, the quantum circuit model has many weaknesses, particularly from an optimization perspective. Quantum circuits come equipped with a variety of different gate sets: A good optimizer for the Clifford+T gate set is not guaranteed to perform well on IBM’s gates or Google’s. They are also rigid: They consist of a large sequence of vertically and horizontally ordered gates, whereas an optimizer only cares about the connections between gates. In this spirit, Kissinger and van de Wetering developed PyZX [16], an optimizer for the ZX calculus [5], a graphical language for quantum computing in which *only connectivity matters* [8]. Like prior tools PyZX checks for correctness by translation validation, either converting diagrams to their underlying linear maps in NumPy and checking if all elements of the linear map are equal up to a global nonzero scalar, or “optimizing” a circuit concatenated with its adjoint, and checking to see if it returns the identity. Unfortunately, these methods are slow and are not guaranteed to succeed as showing circuit equivalence is known to be QMA-complete in the general case [15].

Drawing inspiration from VOQC and PyZX and the verified classical compiler CompCert [19], we present VyZX, a formalization of the ZX calculus in the Coq [9] proof assistant. VyZX is intended to be a fully verified implementation of the PyZX compiler and a platform for mechanized reasoning about the ZX calculus and related graphical calculi. Given the versatility of the ZX-calculus, VyZX should allow us to tackle correctness issues in a range of domains, including lattice surgery [4], circuit simulation [18], and natural language processing [7].

* Equal contribution

Unfortunately, while “only connectivity matters” is an excellent slogan for a graphical language, it poses significant challenges for formal verification. Computers do not talk in pictures: Internally, they impose a rigidity akin to that of the circuit model. Even the standard representations of graphs, like adjacency lists and adjacency matrices, are ill-suited to inductive reasoning of the sort Coq excels in. To address this, we formalize two views of ZX diagrams: *block representation*, in which wires and nodes are composed horizontally and vertically, and *graph representation*, which is more faithful to the standard representation of ZX diagrams.

We explore the motivation for creating VyZX in [Section 2](#). We layout the design decisions underlying VyZX and discuss their potential in [Section 3](#). There we also cover the inductive definition for ZX diagrams, how we apply semantics to them, and how we prove equivalence of two diagrams. We discuss how to convert from standard quantum circuits to our inductive diagrams ([Section 3.5](#)) and how we can view our inductive diagrams as graphs ([Section 4](#)). We sketch out a path from our current formalization of the ZX calculus to a full-fledged quantum optimizer that is integrated with the VOQC compiler and conclude with the many potential use-cases for VyZX in [Section 5](#). All code referenced in this paper can be found at <https://github.com/inQWIRE/VyZX>.

2 Verified Optimization and the ZX-Calculus

2.1 Verified Optimization

Quantum circuit optimizers take varied approaches to verify that their optimizations are well behaved. Some compilers, such as the one created by Nam et al. [21], solely rely on unit testing to ensure correctness, though unit testing can only show the presence – not absence – of bugs. Quantinuum’s $t|ket\rangle$ [24] goes a step further and uses a Hoare logic system to check that certain postconditions, such as “the circuit contains no wire swaps”, hold for various optimization passes. This is a useful system for a quantum optimizer to have, but checking that the postcondition holds does not guarantee that the optimization has returned an equivalent circuit to the input circuit. This gap in how we can verify quantum circuit optimizations has been filled by a few other compilers.

The quantum compilers CertiQ [23] and Quartz [29] attempt to alleviate this concern by adding systems to check for circuit equivalence. They use these systems throughout their development process to check optimizations. To check equivalence, they generate some proof obligations as SMT formulas and pass them along to Z3. This steps beyond compilers like $t|ket\rangle$ as it actually attempts to validate the optimization. The key feature of compilers like Quartz and CertiQ is that they automate this validation of optimization passes. While this is a valuable way to make it easy to engineer new optimization passes, it is incomplete. Not every optimization can be verified by an SMT solver, and compilers using such solvers hence include certain optimization without validation. This undermines the validation itself if the entire thing is not validated. If we want to verify the optimizer completely, we need a stronger system that allows us to prove the correctness, not just pass it to an SMT solver.

The flaws of existing compilers inspired the development of the VOQC verified compiler [14]. VOQC uses SQIR [13] and QuantumLib [26] to provide a full-stack verification pipeline. All three libraries above are written in the Coq proofs assistant, providing them with strong correctness guarantees. VOQC ingests SQIR circuits from QASM [10], then applies optimization passes that are proven correct at compile time in Coq, and hence do not require any automation as we have seen with CertiQ. Upon completion SQIR circuits are converted back into QASM. SQIR can also handle multiple different gate sets, which are related to a base gate set that is used for proof.

2.2 ZX calculus

In contrast to the circuit model used by VOQC and other optimizers, ZX diagrams are a graphical representation of quantum operations. The ZX calculus [6] uses such diagrams together with a set of rewrite rules to manipulate quantum operations. Fundamentally, ZX diagrams are graphs with green and red nodes¹, called Z and X spiders, with *in* inputs and *out* outputs, along with a rotation angle $\alpha \in [0, 2\pi)$. If the rotation angle is 0 it can be omitted. The semantics of Z spiders and X spiders are shown in Figure 1.

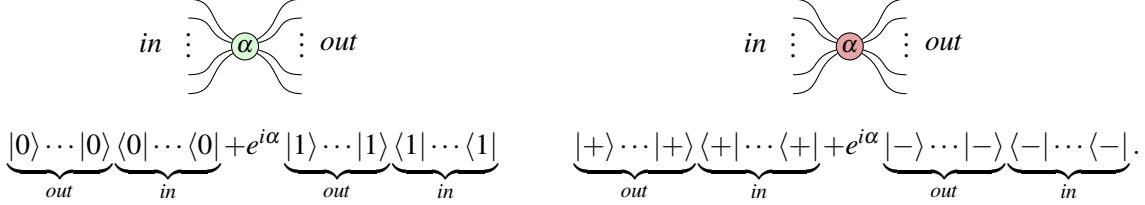


Figure 1: Z and X spiders with their standard bra-ket semantics.

These spiders are connected through edges. Edges can either be regular edges or Hadamard edges, which are represented as dotted lines in diagrams and implicitly add a Hadamard gate on their path. These Hadamard edges can be treated as syntax for regular edges with three nodes, see Figure 5. For more on the ZX-calculus, we recommend John van de Wetering's excellent survey of the topic [27].

2.3 ZX Calculus Optimization

Our work is inspired by PyZX optimizations, which is based on work by Duncan et al. [11] and Kissinger and van de Wetering [17]. Duncan et al.'s describe how to optimize ZX diagrams with graph-theoretic rules. They begin by outlining a restriction on ZX diagrams, called *graph-like form*. Such diagrams are restricted to having only one kind of node (the Z spider), one kind of edge (Hadamard edge), no self-loops, and the condition that each node has a unique parent (i.e., either only one input or output). They show that one can freely convert between unrestricted ZX diagrams and graph-like ZX diagrams with equivalent semantics. Using graph-like form, they use graph operations, such as local complementation and pivoting, to reduce the number of Clifford gates (corresponding to nodes containing multiples of $\pi/2$) in their circuits by combining nodes.

Since the algorithm above does not reduce the number of non-Clifford gates in a circuit, Kissinger and van de Wetering [17] devised an optimization strategy for ZX calculus to reduce the T-gate count, T gates being the most expensive operations for error-corrected quantum computers. To achieve this, they remove any non-Clifford angles from spiders by splitting them into so-called *phase gadgets*, as shown in Figure 2. Then, the resulting Clifford gates are optimized as previously described. Finally, phase gadgets are fused back into the diagram resulting in an optimized diagram. In merging multiple phase gadgets spiders, the phase gadgets themselves can be merged, creating angles that are multiples of $\frac{\pi}{2}$ that can be fused into the Clifford part of the diagram, reducing the number of non-Clifford gates. With these optimizations, the ZX-based optimizer PyZX achieves state-of-the-art performance [16] on Maslov's reversible benchmark suite [20].

VOQC and PyZX both stand out as significant quantum optimizers for their use of formal verification and the ZX-calculus, respectively. With VyZX, we set out to combine these two ideas into one quantum circuit optimizer.

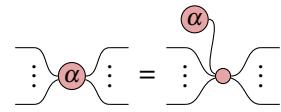


Figure 2: A phase gadget

¹We chose accessible shades of green and red for this paper (see <https://zxcalculus.com/accessibility.html>)

3 VyZX

In designing VyZX, we wanted to make it easy to write recursive or inductive functions over diagrams. In a proof assistant like Coq, inductive structures allow for inductive proofs. Having the ability to use inductive proofs was a core goal for our definition as it would greatly simplify proofs. For inspiration, we looked at diagrams for symmetric monoidal categories as described by Selinger [22]. We reduced the basic requirements for our string diagrams to:

1. The unit object, which is the empty diagram,
2. The single wire,
3. Morphisms, which take n inputs to m outputs,
4. Braids, which swap two wires,
5. Sequential composition, which composes two diagrams in sequence , and
6. Tensor products, which arrange two diagrams in parallel.

These core objects give us our base language, consisting of a set of base morphisms, an empty diagram, a way to swap wires around, and the ability to compose diagrams in sequence and parallel. When we wish to apply this to the ZX calculus, we must decide what our language's *signature* will be. This signature constitutes the morphisms of our string diagram. For the ZX calculus, a simple signature could include just the Z and X spiders. We also include caps and cups in our signature to make diagrams easier to write, which are standard morphisms for string diagrams. The building blocks for creating string diagrams inductively are given in [Figure 3](#), which we expand upon to build different representations for ZX diagrams.

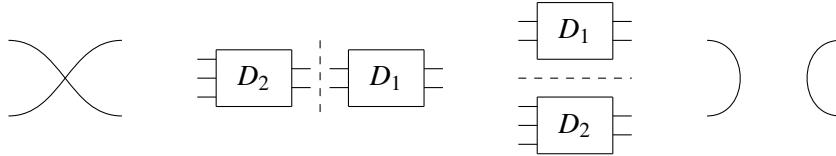


Figure 3: From left to right, the braid, sequential composition, tensor product, cap, and cup for symmetric monoidal string diagrams.

3.1 Block Representation ZX Diagrams

Our first goal with VyZX is to create diagrams that can be used for proof in Coq. We accomplish this by giving an inductive definition for ZX diagrams. We refer to this representation as *block representation ZX diagrams* in reference to how they may be stacked together and how stacks can line up with one another. Each ZX diagram holds information about how many inputs and outputs it has, allowing us to define composition in a way that matches the outputs and inputs of two diagrams through a `Compose` constructor. We also have a `Stack` operation that places one diagram on top of another. Our base constructors are the `Z_Spider`, `X_Spider`, `Cap`, `Cup`, `Swap`, and `Empty` diagrams. The type of a ZX diagram is given by $\text{ZX} : \mathbb{N} \rightarrow \mathbb{N} \rightarrow \text{Type}$ and the constructors are given by [Figure 4](#). These eight constructors allow us to write simple recursive functions and inductive proofs over ZX diagrams while allowing us to describe arbitrary diagrams. Graphically these constructors correspond to the diagrams seen in [Figure 3](#) with the addition of the Z and X spiders shown in [Figure 1](#).

$\frac{\text{in out} : \mathbb{N} \quad \alpha : \mathbb{R}}{\text{Z_Spider in out } \alpha : \text{ZX in out}}$	$\frac{\text{in out} : \mathbb{N} \quad \alpha : \mathbb{R}}{\text{X_Spider in out } \alpha : \text{ZX in out}}$
$\frac{}{\text{Cap} : \text{ZX 0 2}}$	$\frac{}{\text{Swap} : \text{ZX 2 2}}$
$\frac{}{\text{zx1} : \text{ZX in mid}}$	$\frac{}{\text{Empty} : \text{ZX 0 0}}$
$\frac{\text{zx2} : \text{ZX mid out}}{\text{Compose zx1 zx2 : ZX in out}}$	$\frac{\text{zx1} : \text{ZX in1 out1} \quad \text{zx2} : \text{ZX in2 out2}}{\text{Stack zx1 zx2 : ZX (in1 + in2) (out1 + out2)}}$

Figure 4: The inductive constructors for block representation ZX diagrams

We can view the same ZX diagrams that we describe using block representation not based on their building blocks but instead on nodes' adjacency. We refer to this view as the *graph representation*. In the following, we will see how graph representation can be useful for optimization and for proving certain rules. In [Section 4](#) we will discuss how we can convert block representation to graph representation.

3.2 Semantics of diagrams

Given our inductive definition for ZX diagrams, we can write a simple function ([Algorithm 1](#)) for computing the semantics of a given ZX diagram. We use the matrix definition given in [QuantumLib](#) [26] to compute the semantics. In [QuantumLib](#) a matrix is simply a function with type $\mathbb{N} \rightarrow \mathbb{N} \rightarrow \mathbb{C}$ that takes in a row and column index and returns the associated complex number. First, we define the Z and X spider semantics, as stated in [Section 2.2](#), letting \times be matrix multiplication², $I_{m \times n}$ be the m by n identity matrix, H be the Hadamard matrix, \otimes the Kronecker product, and $H^{\otimes n}$ be the n th power of H with the Kronecker product. With spider semantics complete, we define our other base constructors, stacks and composes using the Kronecker and matrix products.

3.3 Proportionality of diagrams

Trivially two syntactically equal diagrams are to be considered equal. For making useful statements, however, we require a notion of semantic equivalence. Intuitively, one might define that as equivalence of matrices produced by `ZX_semantics` (as shown in [Algorithm 1](#)). In the ZX calculus, though, we only care about equivalence up to multiplication by a constant factor, as rules will introduce constant factors and we are able to rebuild any constant factor if necessary using ZX constructions [27].

Within [VyzX](#) we define a relation called *proportional* and give it the notation \propto . We say $zx_1 \propto zx_2$ if there is a non-zero complex number c such that $\text{ZX_semantics } zx_1 = c * \text{ZX_semantics } zx_2$. We prove that \propto is an equivalence relation as we might expect. We then prove that our composition operators respect proportionality: That is, if $zx_1 \propto zx'_1$ and $zx_2 \propto zx'_2$ then `Compose zx1 zx2` \propto `Compose zx'_1 zx'_2` and `Stack zx1 zx2` \propto `Stack zx'_1 zx'_2`. We add this fact to Coq as a *parametric morphism*, allowing us to rewrite using our equivalences even inside a broader diagram. With proportionality defined, we proceed to verify different rewrite rules within the ZX calculus by proving their diagrams to be proportional.

²We chose to align our notation with [QuantumLib](#) rather than mathematical convention

Algorithm 1 ZX Diagram Semantics

```

function Z_SPIDER_SEMATICS(in, out,  $\alpha$ )
    return  $\begin{bmatrix} 1 & 0 & \cdots & 0 & 0 \\ 0 & 0 & \cdots & 0 & 0 \\ \vdots & \ddots & & \vdots & \\ 0 & 0 & \cdots & 0 & 0 \\ 0 & 0 & \cdots & 0 & e^{i\alpha} \end{bmatrix}$ 
function X_SPIDER_SEMATICS(in, out,  $\alpha$ )
    return  $H^{\otimes \text{out}} \times Z_{\text{SPIDER\_SEMANTICS}}(\text{in}, \text{out}, \alpha) \times H^{\otimes \text{in}}$ 
function ZX_SEMATICS(zx : ZX in out) :  $\mathbb{C}^{\text{in} \times \text{out}}$ 
    switch zx do
        case Empty
            return  $I_{1 \times 1}$ 
        case Swap
            return  $\begin{bmatrix} 1, 0, 0, 0 \\ 0, 0, 1, 0 \\ 0, 1, 0, 0 \\ 0, 0, 0, 1 \end{bmatrix}$ 
        case Cap
            return  $[1, 0, 0, 1]^T$ 
        case Cup
            return  $[1, 0, 0, 1]$ 
        case Z_Spider in out  $\alpha$ 
            return  $Z_{\text{SPIDER\_SEMANTICS}}(\text{in}, \text{out}, \alpha)$ 
        case X_Spider in, out  $\alpha$ 
            return  $X_{\text{SPIDER\_SEMANTICS}}(\text{in}, \text{out}, \alpha)$ 
        case Stack zx1 zx2
            return  $ZX_{\text{SEMANTICS}}(\text{zx1}) \otimes ZX_{\text{SEMANTICS}}(\text{zx2})$ 
        case Compose zx1 zx2
            return  $ZX_{\text{SEMANTICS}}(\text{zx2}) \times ZX_{\text{SEMANTICS}}(\text{zx1})$ 

```

3.4 Proving the Correctness of the ZX-Calculus

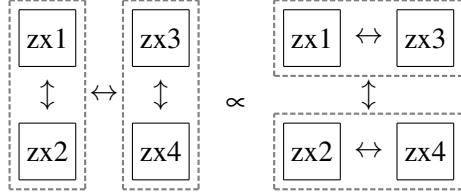
We can now show the rules to manipulate and simplify VyZX diagrams. For readability, we show most rules in graph representation.

Common gates We translated common gates from the circuit model to the ZX-calculus and proved their semantic correctness, shown in [Figure 5](#).

$$\begin{array}{lllll}
 -\textcircled{\pi}- & \equiv & -[\text{Z}]- & -\textcircled{\pi}- & \equiv -[\text{X}]- \\
 -\textcircled{\pi}\textcircled{\pi}- & \equiv & -[\text{Y}]- & -\textcircled{\frac{\pi}{4}}- & \equiv -[\text{T}]- \\
 \begin{array}{c} \textcircled{\pi} \\ \textcircled{\pi} \end{array} & \equiv & \text{CNOT} & \begin{array}{c} \textcircled{\pi} \\ \textcircled{\pi} \\ \textcircled{\pi} \end{array} & \equiv -\square-
 \end{array}$$

Figure 5: Quantum gates represented in the ZX calculus. Note that the H node omits its label and is common to ZX diagrams as syntactic sugar for the rotations above.

Stack & Compose distribute Sequential composition and stacking distribute as long as the individual diagrams have compatible dimensions by the rules stated in [Section 3.1](#). [Figure 6](#) shows this property visually. This fact is central to proving statements in block representation as it enables the diagram's structure to be changed while keeping the components the same.

Figure 6: The distribution of stack (\ominus) and compose (\leftrightarrow).

Bi-Hadamard color changing We define a *color-swapped ZX diagram* as a ZX diagram with the same structure but with all Z spiders being replaced by X spiders (while keeping the angle) and vice versa; all other constructions such as Cap, Cup, and Swap do not change [8]. Henceforth, we denote the color-swapped version of a ZX diagram zx as $\odot zx$. For a given spider, one can swap the spider's color while keeping the angle by composing a stack of Hadamards to the in and outputs. This “bi-Hadamard” construction is shown in Figure 7a. Further, we see this holds for all other non-compositional ZX diagram components (SWAPs, Caps, and Cups) since they do not have color and do not cause rotation. We go on to show that, in fact, the bi-Hadamard rule is true for all color swapped ZX diagrams, as shown in Figure 7b.

Color swapping Using the previous fact about bi-Hadamard constructions, we prove that if a rule can be applied to a ZX diagram zx_1 transforming it into zx_2 , then it can be applied to the color swapped diagram of zx_1 transforming it into the color swapped diagram of zx_2 . With this fact in mind, we only show one color configuration for any rule, understanding that it applies to the color-swapped version. In practice, this allows us to prove any rule only for one color configuration and get the color-swapped lemma for free. Since many proofs are computationally expensive, this greatly speeds up verification.

Bi-algebra rule The bialgebra rule, while not intuitive, is crucial for many ZX proofs as it allows for the rearranging of edges between nodes [27]. Unlike the rules we have so far, it requires proof at matrix level and hence has a computationally expensive proof. Though, once proven, it can be repeatedly applied, and given it rearranges edges can be used to prove many future facts.

Hopf rule The Hopf rule, like the bialgebra rule, deals with the interaction between Z and X spiders. It says that two edges between an X and a Z spider can be removed. Intuitively, this tells us that no matter how much “information” we know about the X basis, given by our input, we do not get any information about the orthogonal Z basis [27]. In practice, the Hopf rule allows us to disconnect specific nodes in the diagram instead of changing their connection. As with the bi-algebra rule, this rule is proven directly on matrices and requires computation, though since the intermediate matrices are smaller this proof is less computationally expensive.

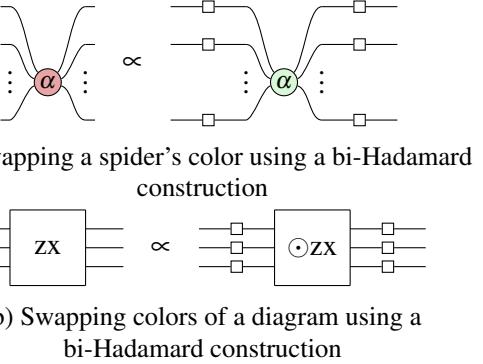
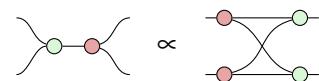


Figure 7: Color swapping ZX diagrams



Bi- π rule Any X spider is equal to itself, surrounded by Z rotations by π . Intuitively this rule is true due to the orthogonal nature of the X and Z basis and the fact that we are performing in total a full rotation. A corollary of this rule is the π -copy rule [27] as shown and derived in the figure on the right. The π -copy rule is not yet implemented, given it semantically builds on adjacency rather than a block like construction, since any input or output could be the one with the π spider on it. Generally speaking, rules that modify an arbitrary subset of the inputs and outputs are hard to represent in block representation and are left to be proven in for graph representation. Since the π -copy rule is what is usually presented in the literature, the bi- π is an interesting case study of a rule tailored to block representation that is equivalent to a more traditional rule.

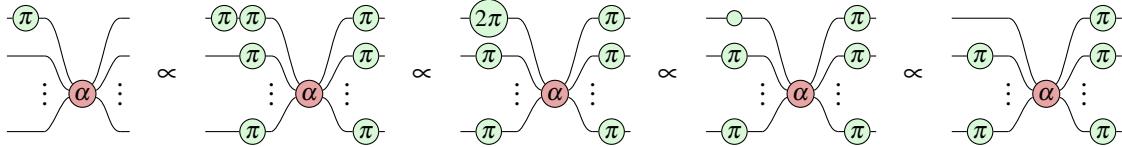


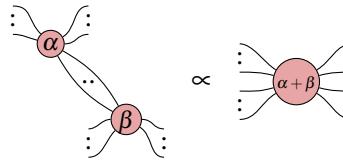
Figure 8: The π -copy rule derivation using the bi- π rule

Spider fusion/splitting One of the most important rules is that spiders connected by an arbitrary (non-zero) number of edges can be fused into a single node with the angles added [27]. This rule is shown in Figure 9a. Further, the reverse is true: any spider can be split such that the two new spiders add to the original angle. A corollary of this is that we can split phase-gadgets (as discussed in Section 2.3) off nodes.

Since this rule fundamentally works based on adjacency rather than block representation construction, we have not fully implemented it at the time of writing. We plan to follow with the general fusion rule, using the graph representation described in Section 4. We do, however, have a restricted version proven where two nodes are just connected to each other by a single wire, which will form the basis of the general spider fusion proof. Figure 9b shows this restricted version. We implemented this by using the bra-ket semantics of spiders, as shown in Figure 1, instead of the direct matrix semantics shown in Algorithm 1. This allows us to use QuantumLib’s algebraic rewrites of complex matrices to combine the angles easily and shows why having both versions of our semantics is valuable for proof. Given the algebraic rewrite this rule is computationally very efficient.

3.5 SQIR Ingestion

We want VyZX to be able to ingest arbitrary circuits. To achieve this, VyZX reads in circuits written in SQIR, an intermediate representation for the VOQC compiler embedded in Coq. By choosing SQIR we maintain interoperability with another verified compiler. SQIR represents circuits with q qubits as compositions of arbitrary x, y, z rotations acting on qubit n and CNOTs between arbitrary qubits n, m , so this transformation and its verification is not trivial.



(a) The spider fusion rule: Two connected nodes fused into a new node with added angles.



(b) A restricted version of spider fusion currently supported by VyZX

Figure 9: Spider fusion rule

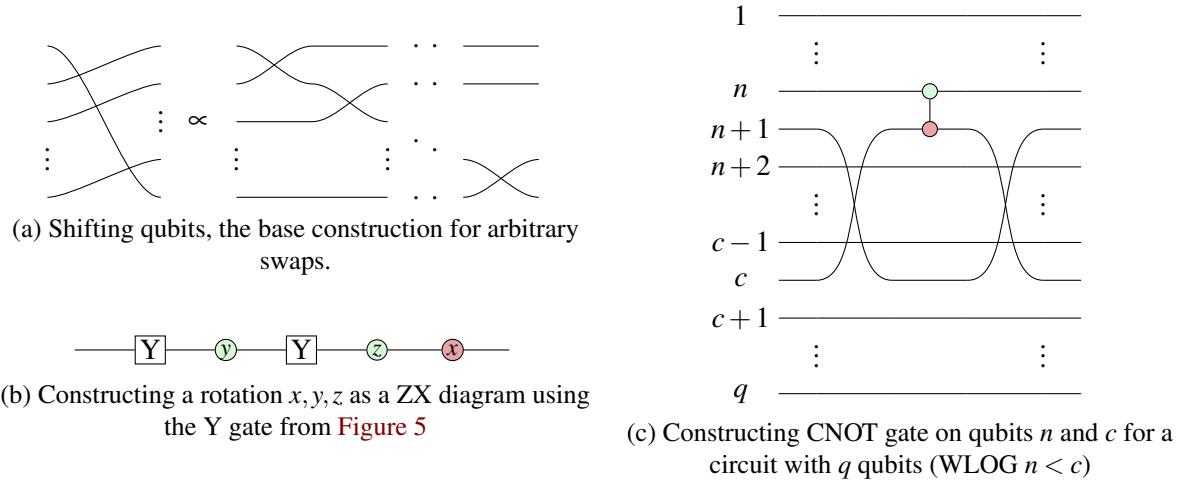


Figure 10: Construction of SQIR operations in our ZX representation

Circuit ingestion works by using arbitrary swaps, meaning instead of only having swap gates that swap two adjacent qubits, we have swap gates able to swap arbitrary two qubits³. To construct an arbitrary swap, we first build a construction to shift qubit 1 to n , whereby all other qubits are shifted up. We also build the inverse (i.e., shifting qubit n to 1). Figure 10a shows this construction. With this we can easily construct our arbitrary SWAP gate (shift 1 to n , then $n - 1$ to 1). The discussion below will be in block representation that has defined swaps (see sec Section 3.1) of two adjacent qubits as the composition of 3 CNOTs.

Using arbitrary swaps and shifts, we can now interpret any wire crossing. Hence, in block representation we can construct a CNOT acting on two qubits n, c by swapping one qubit next to the other qubit, applying the CNOT, and swapping back, as shown in Figure 10c. To then convert such a circuit, VyZX converts rotations x, y, z into the construction shown in Figure 10b. Composition of SQIR terms is also represented by composition in our block representation IR.

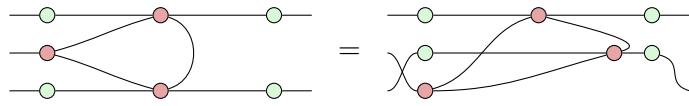


Figure 11: Two equal ZX diagrams, where the right diagram is deformed. Note that all connections, inputs/outputs, and qubit order are maintained.

4 Graph Representation

As discussed in Section 3.1, we can view diagrams in both block representation and graph representation. In graph representation, we construct ZX diagrams solely based on node adjacency. This representation brings some interesting properties that will allow us to optimize diagrams or prove further rules (as mentioned in Section 3.4). We can see that any graph representation diagram can be deformed arbitrarily, as long as inputs and outputs are kept in order. Figure 11 illustrates this: Here, we see two ZX diagrams equal up to deformation.

³For the implementation we chose to have an IR with first-class arbitrary swaps that will then be translated into our base block representation IR (preserving semantics) using the following constructions.

One of our future goals is to add a graph representation for ZX diagrams, allowing us to act on diagrams. Creating a verified semantics for graph representation remains a work-in-progress since, fundamentally, this requires topologically-sorted graph traversal, which is a hard problem to implement in proof assistants. Further, converting graph representation diagrams into block representation is also challenging; if overcome, we could provide semantics to graph representation in terms of the corresponding block representation diagrams. We devised a method to convert to graph representation, which is a central problem. In the following, we shall outline this conversion.

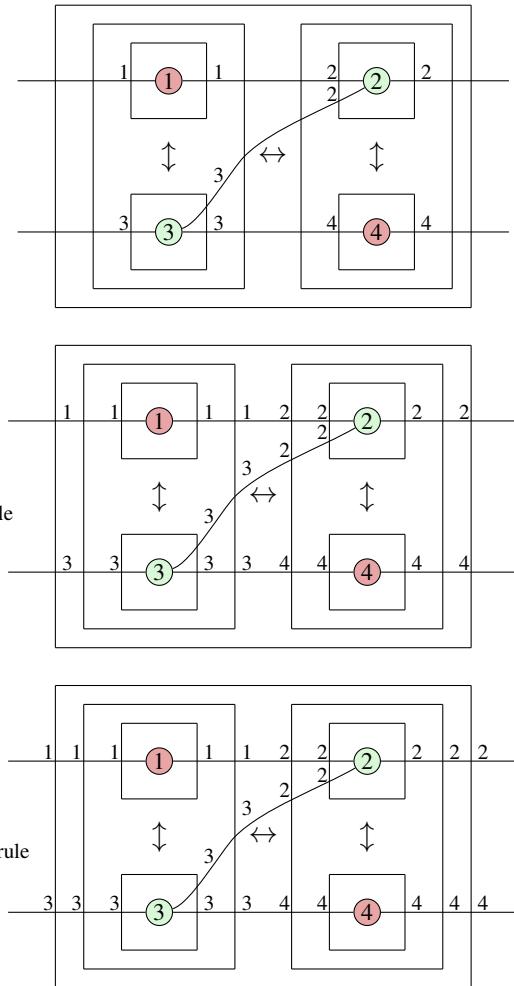
Before converting into graph representation, we convert ZX diagrams into a restricted yet just as universal form; this form is similar to Duncan et al's [11] graph-like form (as described in [Section 2.3](#)), differing only in the existence of self-loops and two or more possible links between spiders. The restrictions are as follows:

1. Only Hadamard edges
2. Only one type of spider (Z spiders)
3. Every spider has one input and zero to two outputs or vice-versa

Restrictions 1 & 2 are in place to make our graph more conventional (and akin to graph-like diagrams) by having nodes and edges. Restriction 3, however, is in place for ease of proving. To divide up the proofs into logical steps, there are separate intermediate representations that build up all three restrictions. These IRs exist to divide proofs into logical components and are mostly transparent. Developers, however, can choose to use less restricted forms.

Given our restricted form, our algorithm proceeds as follows: A procedure numbers all ZX diagram components (i.e. the constructors) with a unique integer. Then it proceeds to number edges of components as follows: Each component with n inputs and m outputs will produce a pair of lists sized n and m , where every position in the list indicates the closest fundamental component (spider or cap/cup). A non stacking/composing component with n inputs, m outputs, and node number x will return lists of size n and m , each with all entries being x .

When stacking two diagrams, the procedure concatenates respective lists, and when sequencing diagrams, it carries forward the outer lists. [Figure 12](#) shows an example of this process. It is important to note that we treat caps and cups like spiders at this stage. Once the edge numbering is complete, we



[Figure 12](#): An example of the edge annotation algorithm. Each spider is shown with a unique node number inside of it. In the first step, all wires out of the spiders are annotated with the node numbers of the spiders. Then when stacking nodes, those numbers are carried forward and applied to the stack constructions wires. Note that the order is preserved.

We see that the outer set of numbers is carried forward upon a composition in the last step. As before, \uparrow denotes stacking and \leftrightarrow composition

Algorithm 2 block representation to graph representation conversion

```

function NUMBERNODES(zx)
  assignFreshNumber(zx)
  if zx = Stack zx1 zx2 OR zx = Compose zx1 zx2 then
    NumberInnerNodes(zx1)
    NumberInnerNodes(zx2)
  function NUMBEREDGES(zx)
    if zx = Stack zx1 zx2 then
      (in1, out1) = NumberEdges(zx1)
      (in2, out2) = NumberEdges(zx2)
      return (in1 ++ in2), (out1 ++ out2)
    else if zx = Compose zx1 zx2 then
      (in1, _) = NumberEdges(zx1)
      (_, out2) = NumberEdges(zx2)
      return in1, out2
    else
      return NodeNumber(zx), NodeNumber(zx)
  function CREATEEDGES(zx)
    if zx = Compose zx1 zx2 then
      (_, out1) = GetEdgeNumbers(zx1)
      (in2, _) = GetEdgeNumbers(zx2)
      for ((in, out) ∈ (out1, in2)) do                                ▷ Note that out1 and in2 have the same length by construction
        AddEdge(in, out)
    if zx = Compose zx1 zx2 OR Stack zx1 zx2 then
      CreateEdges(zx1)
      CreateEdges(zx2)
  
```

will traverse the diagram once more, and at every Compose, we will match the output edge numbers of the left diagram with the input edge numbers of the right diagram and mark each of those as an edge. Looking at our example in [Figure 12](#), we see wires labeled (1,2), (3,2), and (3,4) bridging the main composition. We use the information from the algorithm to infer which inputs/outputs of the diagram are connected to which node by looking at the outermost annotation. So in our example, we see by looking at the outermost labels that input 1 is connected to node 1, input 2 to node 3, output 1 to node 2, and output 2 to node 4. We can then annotate the entire diagram’s inputs and outputs by looking at the outermost labels. [Algorithm 2](#) shows a pseudocode description of these processes.

5 Future Directions

5.1 Circuit extraction

Once we have converted our block representation diagrams to graph representation diagrams, we will extract these graph representation diagrams to SQIR circuits. To accomplish this, we plan to follow Backens extraction work [3]. Our systems should allow us to define a notion of *gflow*. This graph-theoretical property is sufficient for extracting ZX diagrams into circuits, which will be a valuable tool to verify that optimizations do not break extractability. Circuit extraction will complete the core of VyZX as now we will be able to ingest circuits, write functions over them, and extract the circuits back to SQIR.

Once we have verified optimizations and extractions, we are able to pursue a couple of interesting projects, including optimization and simulation. As we are sharing base libraries with VOQC, the natural next idea is to integrate VyZX with VOQC properly.

5.2 VOQC integration

Once we complete graph representation conversion and circuit extraction, we plan to build an optimization pass using PyZX-like optimizations, as described in [Section 2.3](#). Instead of building a standalone optimizer, we plan on integrating our work into VOQC. Since we can ingest from (and later extract to) SQIR, we have a common IR that will allow us to have a wholly verified pipeline and interoperate with ease. It will be interesting to see whether non-ZX optimizations combined with ZX optimizations yield a benefit. VOQC has the advantage of a complete interface that allows for pass selection, allowing us to expose the ZX-based optimizations to users. Furthermore, integration into VOQC will allow us to benchmark our optimizer against other state-of-the-art optimizers like PyZX [16] and Quartz [28].

5.3 The ZH and the ZW Calculi

VyZX’s design is focused exclusively on the ZX calculus, but the principles described here could easily be applied to other similar calculi such as the ZW calculus [12] or ZH calculus [2], which have broad applications to quantum communication the description of quantum oracles, respectively. In fact, block construct diagrams let us easily translate between various graphical calculi. The translation can be verified by Coq by a simple inductive proof over the signature for the calculus and our shared string diagram constructions. It may be of interest to develop additional optimizations based on these different calculi, and a small extension to VyZX could allow such optimizations to be verified. With the different IRs we have right now, we are confident that extensions of the calculus are easy to integrate.

6 Conclusion

VyZX is a formal verification framework for the ZX-calculus. Creating VyZX came with several challenges that are unique to verifying a graphical language. Finding a way to easily assign semantics to a graph is inherently difficult due to traditional graph structures not being idiomatic in proof assistants. Block representation provides an inductive description of a graph that allows for easy proof while preserving the expressiveness that graphs provide. As block representation made it challenging to write diagrams, we developed a graph representation that can act as a way to implement programs over ZX diagrams. With these two views in place, we proved several core ZX diagram equivalences and added circuit ingestion from SQIR. With all these tools in place, we believe VyZX has a future as a basis for writing programs over ZX diagrams. In particular, our next step is to build a verified circuit optimizer in the style of PyZX. The core of VyZX will continue to be improved as we approach new problems in implementing programs such as a verified circuit simulator. We are confident that with VyZX’s evolution, it will provide a robust foundation for future work on the verification of graphical quantum calculi and their applications.

References

- [1] Matthew Amy (2019): *Towards Large-scale Functional Verification of Universal Quantum Circuits*. Electronic Proceedings in Theoretical Computer Science 287, pp. 1–21, doi:[10.4204/ptcs.287.1](https://doi.org/10.4204/ptcs.287.1).
- [2] Miriam Backens & Aleks Kissinger (2019): *ZH: A Complete Graphical Calculus for Quantum Computations Involving Classical Non-linearity*. Electronic Proceedings in Theoretical Computer Science 287, p. 23–42, doi:[10.4204/ptcs.287.2](https://doi.org/10.4204/ptcs.287.2).
- [3] Miriam Backens, Hector Miller-Bakewell, Giovanni de Felice, Leo Lobski & John van de Wetering (2021): *There and back again: A circuit extraction tale*. Quantum 5, p. 421, doi:[10.22331/q-2021-03-25-421](https://doi.org/10.22331/q-2021-03-25-421).

- [4] Niel de Beaudrap & Dominic Horsman (2020): *The ZX calculus is a language for surface code lattice surgery.* *Quantum* 4, p. 218, doi:[10.22331/q-2020-01-09-218](https://doi.org/10.22331/q-2020-01-09-218).
- [5] Bob Coecke & Ross Duncan (2008): *Interacting Quantum Observables*. In Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfssdóttir & Igor Walukiewicz, editors: *Automata, Languages and Programming*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 298–310, doi:[10.1007/978-3-540-70583-3_25](https://doi.org/10.1007/978-3-540-70583-3_25).
- [6] Bob Coecke & Ross Duncan (2011): *Interacting quantum observables: categorical algebra and diagrammatics.* *New Journal of Physics* 13(4), p. 043016, doi:[10.1088/1367-2630/13/4/043016](https://doi.org/10.1088/1367-2630/13/4/043016).
- [7] Bob Coecke, Giovanni de Felice, Konstantinos Meichanetzidis & Alexis Toumi (2020): *Foundations for Near-Term Quantum Natural Language Processing*, doi:[10.48550/ARXIV.2012.03755](https://arxiv.org/abs/2012.03755).
- [8] Bob Coecke & Aleks Kissinger (2017): *Picturing Quantum Processes: A First Course in Quantum Theory and Diagrammatic Reasoning*. Cambridge University Press, doi:[10.1017/9781316219317](https://doi.org/10.1017/9781316219317).
- [9] The Coq Development Team (2012): *The Coq Reference Manual, version 8.4*. Available electronically at <http://coq.inria.fr/doc>.
- [10] Andrew W. Cross, Lev S. Bishop, John A. Smolin & Jay M. Gambetta (2017): *Open Quantum Assembly Language*, doi:[10.48550/ARXIV.1707.03429](https://arxiv.org/abs/1707.03429).
- [11] Ross Duncan, Aleks Kissinger, Simon Perdrix & John van de Wetering (2020): *Graph-theoretic Simplification of Quantum Circuits with the ZX-calculus*. *Quantum* 4, p. 279, doi:[10.22331/q-2020-06-04-279](https://doi.org/10.22331/q-2020-06-04-279).
- [12] Amar Hadzihasanovic (2015): *A Diagrammatic Axiomatisation for Qubit Entanglement*, doi:[10.48550/ARXIV.1501.07082](https://arxiv.org/abs/1501.07082).
- [13] Kesha Hietala, Robert Rand, Shih-Han Hung, Liyi Li & Michael Hicks (2021): *Proving Quantum Programs Correct*. In Liron Cohen & Cezary Kaliszyk, editors: *12th International Conference on Interactive Theorem Proving (ITP 2021)*, Leibniz International Proceedings in Informatics (LIPIcs) 193, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, pp. 21:1–21:19, doi:[10.4230/LIPIcs.ITP.2021.21](https://doi.org/10.4230/LIPIcs.ITP.2021.21).
- [14] Kesha Hietala, Robert Rand, Shih-Han Hung, Xiaodi Wu & Michael Hicks (2021): *A Verified Optimizer for Quantum Circuits*. Proc. ACM Program. Lang. 5(POPL), doi:[10.1145/3434318](https://doi.org/10.1145/3434318).
- [15] Dominik Janzing, Paweł Wocjan & Thomas Beth (2003): *Identity check is QMA-complete*, doi:[10.48550/ARXIV.QUANT-PH/0305050](https://arxiv.org/abs/quant-ph/0305050).
- [16] Aleks Kissinger & John van de Wetering (2020): *PyZX: Large Scale Automated Diagrammatic Reasoning*. In Bob Coecke & Matthew Leifer, editors: *Proceedings 16th International Conference on Quantum Physics and Logic*, Chapman University, Orange, CA, USA., 10-14 June 2019, *Electronic Proceedings in Theoretical Computer Science* 318, Open Publishing Association, pp. 229–241, doi:[10.4204/EPTCS.318.14](https://doi.org/10.4204/EPTCS.318.14).
- [17] Aleks Kissinger & John van de Wetering (2020): *Reducing the number of non-Clifford gates in quantum circuits*. *Physical Review A* 102(2), doi:[10.1103/physreva.102.022406](https://doi.org/10.1103/physreva.102.022406).
- [18] Aleks Kissinger & John van de Wetering (2021): *Simulating quantum circuits with ZX-calculus reduced stabiliser decompositions*, doi:[10.48550/ARXIV.2109.01076](https://arxiv.org/abs/2109.01076).
- [19] Xavier Leroy (2009): *Formal Verification of a Realistic Compiler*. *Commun. ACM* 52(7), p. 107–115, doi:[10.1145/1538788.1538814](https://doi.org/10.1145/1538788.1538814).
- [20] Dmitri Maslov (2021): *Reversible logic synthesis benchmarks page*. Available at <https://reversiblebenchmarks.github.io/>.
- [21] Yunseong Nam, Neil J Ross, Yuan Su, Andrew M Childs & Dmitri Maslov (2018): *Automated optimisation of large quantum circuits with continuous parameters*. *npj Quantum Information* 4(1), pp. 1–12, doi:[10.1038/s41534-018-0072-4](https://doi.org/10.1038/s41534-018-0072-4).
- [22] P. Selinger (2010): *A Survey of Graphical Languages for Monoidal Categories*. In: *New Structures for Physics*, Springer Berlin Heidelberg, pp. 289–355, doi:[10.1007/978-3-642-12821-9_4](https://doi.org/10.1007/978-3-642-12821-9_4).

- [23] Yunong Shi, Runzhou Tao, Xupeng Li, Ali Javadi-Abhari, Andrew W. Cross, Frederic T. Chong & Ronghui Gu (2019): *CertiQ: A Mostly-automated Verification of a Realistic Quantum Compiler*, doi:[10.48550/ARXIV.1908.08963](https://doi.org/10.48550/ARXIV.1908.08963).
- [24] Seyon Sivarajah, Silas Dilkes, Alexander Cowtan, Will Simmons, Alec Edgington & Ross Duncan (2020): *$t|ket\rangle$: a retargetable compiler for NISQ devices.* *Quantum Science and Technology* 6(1), p. 014003, doi:[10.1088/2058-9565/ab8e92](https://doi.org/10.1088/2058-9565/ab8e92).
- [25] Kaitlin N. Smith & Mitchell A. Thornton (2019): *A Quantum Computational Compiler and Design Tool for Technology-Specific Targets.* In: *Proceedings of the 46th International Symposium on Computer Architecture*, ISCA '19, Association for Computing Machinery, New York, NY, USA, p. 579–588, doi:[10.1145/3307650.3322262](https://doi.org/10.1145/3307650.3322262).
- [26] The INQWIRE Developers (2022): *INQWIRE QuantumLib*. Available at <https://github.com/inQWIRE/QuantumLib>.
- [27] John van de Wetering (2020): *ZX-calculus for the working quantum computer scientist.* arXiv:[2012.13966](https://arxiv.org/abs/2012.13966).
- [28] Mingkuan Xu, Zikun Li, Oded Padon, Sina Lin, Jessica Pointing, Auguste Hirth, Henry Ma, Jens Palsberg, Alex Aiken, Umut A. Acar & Zhihao Jia (2022): *Quartz: Superoptimization of Quantum Circuits.* In: *Proceedings of the 43rd ACM SIGPLAN Conference on Programming Language Design and Implementation*, doi:[10.1145/3519939.3523433](https://doi.org/10.1145/3519939.3523433).
- [29] Mingkuan Xu, Zikun Li, Oded Padon, Sina Lin, Jessica Pointing, Auguste Hirth, Henry Ma, Jens Palsberg, Alex Aiken, Umut A. Acar & Zhihao Jia (2022): *Quartz: Superoptimization of Quantum Circuits (Extended Version)*, doi:[10.48550/ARXIV.2204.09033](https://doi.org/10.48550/ARXIV.2204.09033). Available at <https://arxiv.org/abs/2204.09033>.