

Qimaera: Type-safe (Variational) Quantum Programming in Idris

LILIANE-JOY DANDY, Ecole Polytechnique/EPFL, France/Switzerland

EMMANUEL JEANDEL, Université de Lorraine, LORIA, France

VLADIMIR ZAMDZHIEV, Inria, France

Variational Quantum Algorithms are hybrid classical-quantum algorithms where classical and quantum computation work in tandem to solve computational problems. These algorithms create interesting challenges for the design of suitable programming languages. In this paper we introduce Qimaera, which is a set of libraries for the Idris 2 programming language that enable the programmer to implement (variational) quantum algorithms where the full power of the elegant Idris language works in synchrony with quantum programming primitives that we introduce. The two key ingredients of Idris that make this possible are (1) dependent types which allow us to implement unitary (i.e. reversible and controllable) quantum operations; and (2) linearity which allows us to enforce fine-grained control over the execution of quantum operations that ensures compliance with the laws of quantum mechanics. We demonstrate that Qimaera is suitable for variational quantum programming by providing implementations of the two most prominent variational quantum algorithms – QAOA and VQE. To the best of our knowledge, this is the first implementation of these algorithms that has been achieved in a *type-safe* framework.

Additional Key Words and Phrases: variational quantum programming, dependent types, linear types

1 INTRODUCTION

Variational Quantum Algorithms [8, 15, 20] are becoming increasingly important for quantum computation. The main idea behind this computational paradigm is to use hybrid classical-quantum algorithms that work in tandem to solve computational problems. The classical part of the algorithm is performed by a classical processor and the quantum part of the algorithm is executed on a quantum device. During the computation process, intermediary results produced by the quantum device are passed onto the classical device which performs further computation on them that is used to tune the parameters of the quantum part of the algorithm, which therefore has an effect on the quantum dynamics. The hybrid classical-quantum back and forth process repeats until a desired termination condition is satisfied.

This hybrid classical-quantum computational paradigm opens up interesting and important challenges for the design of suitable programming languages. It is clear that if we wish to program within such computational scenarios, we need to develop a language that correctly models the manipulation of *quantum resources*. In particular, quantum measurements give rise to *probabilistic computational effects* that are inherited by the classical side of the language. Another issue is that quantum information behaves very differently compared to classical information. As an example, quantum information cannot be copied in a uniform way [28], unlike classical information, which may be freely copied without restriction. Therefore, if we wish to avoid runtime errors, the quantum fragment of the language needs to be equipped with features for fine-grained control, such as for example, having a *substructural typing discipline* [3–5, 10, 14] where contraction (i.e. copying) is restricted. On the other hand, when doing classical computation, such restrictions are unnecessary and often inconvenient. One solution to this problem is to design a language with a classical (non-linear) fragment together with a quantum (linear) one, both of which interact nicely with each other. In fact, this can be achieved within an existing language that has a sufficiently advanced type system, as we show in this paper.

Our Contributions. In this paper, we describe *Qimaera* (named after the hybrid creature Chimaera from Greek mythology), which is a set of libraries for the Idris 2 language [7] that allow the programmer to implement (variational) quantum algorithms in a *type-safe* way. Idris 2 is an elegant functional programming language that is equipped with an advanced type system based on Quantitative Type Theory [3, 14] that brings many useful features to the programmer, most notably *dependent types* and *linearity*. These two features of Idris are crucial for the development of *Qimaera* and, in fact, are the reason we chose Idris in the first place. Dependent types are used throughout our entire development in order to correctly represent and formalise the compositional nature of quantum operations. Linearity is used in order to enforce the proper consumption of quantum resources (during execution) in a way that is admissible with respect to the laws of quantum mechanics. The combination of dependent types and linearity allows us to *statically* detect and reject erroneous quantum programs and this ensures the type safety of our approach to variational quantum programming.

In our intended computational scenario, we have access to both a classical computer and a quantum computer. Since we cannot directly observe quantum information, we directly interact with the classical computer which sends instructions to, and receives data from, the quantum device via a suitable interface that makes use of the IO monad. In our view, this is an adequate representation of a realistic computational environment for variational quantum programming. We design a suitable (abstract) interface that allows us to model this situation accurately and which makes use of the IO monad. However, since the authors do not personally have any quantum hardware, we provide only one concrete implementation of our interface that simulates the relevant quantum operations on our classical computers by using the proper linear-algebraic formalism, but while still using the IO monad as prescribed by the abstract interface. From a high-level programming perspective, both the abstract interface and its concrete implementation (via linear-algebraic simulation) address all of the programming challenges induced by the classical-quantum device computational scenario.

We emphasise that we can achieve type-safe (variational) quantum programming in an *existing* programming language by implementing suitable libraries. This is important for *variational* quantum programming, because in most variational quantum algorithms, the classical part of the algorithm is considerably larger, more complicated and more difficult to implement, compared to the quantum part of the algorithm. Therefore, it is important for the programming language to have first-class support for classical programming features. Our chosen language, Idris, is definitely such a language. The advanced type system of Idris allows us to elegantly mix quantum and classical programming primitives and therefore allows us to get the best of both worlds. We demonstrate that *Qimaera* is suitable for variational quantum programming by providing implementations of the two most prominent variational quantum algorithms – QAOA and VQE. To the best of our knowledge, this is the first implementation of these algorithms that has been achieved in a *type-safe* framework. In particular, this means that common quantum programming errors (e.g. copying of qubits, applying a CNOT operation with the same source and target, etc.) are *statically* detected and rejected by the Idris type checker.

Overview and summary of contributions. We begin by providing some background on quantum computation (§2.1) and the Idris programming language (§2.2). We then explain how we represent unitary (i.e. reversible and controllable) quantum operations in Idris and we provide some important and non-trivial examples (§3). In §4 we describe how we represent arbitrary (non-unitary, effectful) quantum operations and we present some examples of effectful quantum programs and algorithms. The linear features of Idris are crucial for achieving this. We discuss why *Qimaera* is suitable for variational quantum programming and we provide implementations of the VQE and

QAOA algorithms in §5. Finally, we discuss related work and future work in §6. The Idris source code for Qimaera is available at <https://github.com/zamdzhiev/Qimaera> under the MIT license. In this paper, we provide some excerpts of the source code, but we recommend consulting the code itself, because it contains many useful comments.

2 BACKGROUND

In this section we introduce the relevant background so that readers may follow our subsequent development.

2.1 Quantum Computation

Readers interested in a detailed introduction to quantum computing may consult [16]. In this section we summarise the basic notions that are relevant for our development.

2.1.1 Qubits. The simplest non-trivial quantum system is the *quantum bit*, often abbreviated as *qubit*. Qubits may be thought of as the quantum counterparts of the bit from classical computation. A qubit $|\psi\rangle$ is represented as a normalised vector in \mathbb{C}^2 . The *computational basis* is given by the pair of vectors $|0\rangle \stackrel{\text{def}}{=} \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and $|1\rangle \stackrel{\text{def}}{=} \begin{pmatrix} 0 \\ 1 \end{pmatrix}$, which may be seen as representing the classical bits 0 and 1. An arbitrary qubit is described by $|\psi\rangle = a|0\rangle + b|1\rangle$ where $a, b \in \mathbb{C}$ and $|a|^2 + |b|^2 = 1$.

2.1.2 Superposition. A qubit may be in (uncountably) many different states, whereas a classical bit is either 0 or 1. When the linear combination $|\psi\rangle = a|0\rangle + b|1\rangle$ is non-trivial, then we say that $|\psi\rangle$ is in *superposition* of $|0\rangle$ and $|1\rangle$. Superposition is a very important quantum resource which is used by many quantum algorithms.

2.1.3 Composite Systems. The state space that describes a system of n qubits is the Hilbert space \mathbb{C}^{2^n} . If $|\psi\rangle$ and $|\phi\rangle$ are two states of n and m qubits respectively, then the composite $n + m$ qubit state $|\psi\phi\rangle \stackrel{\text{def}}{=} |\psi\rangle \otimes |\phi\rangle$ is described by the Kronecker product \otimes of the original states.

2.1.4 Unitary Quantum Operations. A quantum state $|\psi\rangle \in \mathbb{C}^{2^n}$ may undergo a *unitary evolution* described by a unitary matrix $U \in \mathbb{C}^{2^n \times 2^n}$ in which case the new state of the system is described by the vector $U|\psi\rangle$. Unitary operations (and matrices) are closed under sequential composition (described by matrix multiplication \circ) and under parallel composition (described by Kronecker product \otimes). Sequential composition of unitary operations is used to describe the temporal evolution of quantum systems, whereas the parallel composition is used to describe their spatial structure.

The unitary quantum operations are also often called *unitary gates*. One typically chooses a *universal gate set* which is a small set of unitary operations that suffices to express all other unitary operations via (parallel and sequential) composition. The universal gate set that we choose for our development is standard and we specify these unitary operations next by giving their action on the computational basis (which uniquely determines the operations).

The *Hadamard Gate*, denoted H , is the 1-qubit unitary map whose action on the computational basis is given by

$$H|0\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \quad H|1\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$$

and its primary purpose is to generate superposition. The *Phase Shift Gate*, denoted $P(\alpha)$, for $\alpha \in \mathbb{R}$, is a 1-qubit unitary map whose action on the computational basis is given by:

$$P(\alpha)|0\rangle = |0\rangle \quad P(\alpha)|1\rangle = e^{i\alpha}|1\rangle$$

and its primary purpose is to modify the phase of a quantum state. The family of Phase Shift Gates is parameterised by the choice of $\alpha \in \mathbb{R}$ and important special cases include the unitary gates




Fig. 1. The Hadamard, Phase Shift and CNOT gates.




Fig. 2. Preparation of the Bell state using atomic gates.

$T \stackrel{\text{def}}{=} P(\pi/4)$ and $Z \stackrel{\text{def}}{=} P(\pi)$. The *Controlled-Not Gate* (CNOT), is a 2-qubit unitary map whose action on the computational basis is given by

$$\begin{aligned} \text{CNOT } |00\rangle &= |00\rangle & \text{CNOT } |01\rangle &= |01\rangle \\ \text{CNOT } |10\rangle &= |11\rangle & \text{CNOT } |11\rangle &= |10\rangle \end{aligned}$$


and this unitary map may be used to generate quantum entanglement (see §2.1.7).

Unitary gates admit a diagrammatic representation as *quantum circuits*. The atomic unitary gates we described above are shown in Figure 1. Composite unitary gates may also be described as circuits (see Figure 2): sequential composition amounts to plugging wires of subdiagrams and parallel composition amounts to juxtaposition.

2.1.5 Controlled Unitary Operations. The CNOT gate is the simplest example of a *controlled unitary gate*. Given a unitary gate $U: \mathbb{C}^{2^n} \rightarrow \mathbb{C}^{2^n}$, the controlled- U unitary gate is the unitary gate $CU: \mathbb{C}^{2^{n+1}} \rightarrow \mathbb{C}^{2^{n+1}}$ whose action is determined by the assignments

$$CU(|0\rangle \otimes |\psi\rangle) = |0\rangle \otimes |\psi\rangle \quad CU(|1\rangle \otimes |\psi\rangle) = |1\rangle \otimes (U|\psi\rangle).$$

Controlled unitary operations are ubiquitous in quantum computing and they are graphically depicted as



using a similar notation to that of the CNOT gate.

2.1.6 Inverse Unitary Operations. Every unitary operation U is *reversible* with the inverse operation given by the conjugate transpose, denoted U^\dagger , which is again a unitary matrix. Applying the inverse operation (also known as the *adjoint*) of a given unitary is also ubiquitous.

2.1.7 Quantum Entanglement. A quantum state $|\psi\rangle \in \mathbb{C}^{2^n}$, with $n > 1$, is said to be *entangled* when there exists no non-trivial decomposition $|\psi\rangle = |\phi\rangle \otimes |\tau\rangle$. Quantum entanglement is a very important resource in quantum computation which is exhibited by many quantum algorithms. Because of the possibility of entanglement, we cannot, in general, break down quantum systems into smaller components and we are often forced to reason about such systems in their entirety. The most important example of an entangled state is the *Bell state* given by $|\text{Bell}\rangle \stackrel{\text{def}}{=} \frac{|00\rangle + |11\rangle}{\sqrt{2}}$.

2.1.8 Preparation of Quantum States. Preparing a new qubit in state $|0\rangle$ is an admissible physical operation. This, together with application of unitary gates as part of the computation, allows us to

prepare arbitrary quantum states, e.g., the Bell state can be prepared by taking $|\text{Bell}\rangle = (\text{CNOT} \circ (H \otimes I)) |00\rangle$. See Figure 2 for the corresponding circuit.

2.1.9 Measurements. Quantum information cannot be directly observed without affecting the state of the underlying system. In order to extract information from quantum systems, we need to perform a *quantum measurement* on (parts of) our systems. For example, when performing a quantum measurement on a qubit in the state $|\psi\rangle = a|0\rangle + b|1\rangle$, there are two possible outcomes: either the quantum system will collapse to state $|0\rangle$ and we obtain the classical bit 0 as evidence of this event, or, the quantum system will collapse to state $|1\rangle$ and we obtain the classical bit 1 as evidence of this event. The first outcome (corresponding to bit 0) occurs with probability $|a|^2$ and the second outcome (corresponding to bit 1) occurs with probability $1 - |a|^2 = |b|^2$. In general, when we measure n qubits simultaneously, we obtain a bit string of length n which determines the event that occurred and the quantum system collapses to a corresponding state with some probability, both of which are determined via the Born rule of quantum mechanics. Therefore, quantum measurements induce evolutions which are *probabilistic* and *irreversible* (or *destructive*), which distinguishes them from unitary evolutions, which are *deterministic* and *reversible*.

2.1.10 No-Cloning Theorem. Unlike classical information, quantum information cannot be uniformly copied. This is made precise by the *no-cloning* theorem of quantum mechanics [28]: there exists no unitary operation $U : \mathbb{C}^4 \rightarrow \mathbb{C}^4$, such that for every qubit $|\psi\rangle : U(|\psi\rangle \otimes |0\rangle) = |\psi\rangle \otimes |\psi\rangle$. This means that copying of quantum information is a *physically inadmissible* operation. Ideally, quantum programming languages should be designed so that these kinds of errors are detected during type checking and not at runtime.

2.2 The Idris 2 Language

In this section, we give a short overview of the Idris 2 language and its main features that are relevant for the development of Qimaera. Idris 2 is a pure functional language with a syntax influenced by that of Haskell. The main additional features, compared to Haskell, are dependent types and linearity, both of which are crucial for Qimaera. Its type system is based on Quantitative Type Theory [3, 14], which specifies how dependent types and linearity are combined.

2.2.1 Dependent Types. In Idris, types are first-class primitives and they may be manipulated like any other construct of the language. This allows us to formulate more expressive types that can depend on values, and hence it enables us to make some properties and program invariants explicit.

Example 2.2.1. The type of vectors is a simple and useful example of a dependent type. A vector is a list with a fixed length that is part of the type. It can be defined as follows, where S is the successor function for natural numbers, and a is a polymorphic type:

```
data Vect : Nat -> Type -> Type where
  Nil    : Vect 0 a
  (:::)  : a -> Vect k a -> Vect (S k) a
```

The type `Vect` has two constructors (i.e., introduction rules). The first one constructs the empty vector, of length zero. The second one is used to introduce non-empty vectors: a vector with $k+1$ elements of type a is constructed by combining an element of type a and a vector of size k .

Type dependency allows us to specify useful program properties and type checking ensures that they hold. For instance, we can define an `append` function that concatenates two vectors. Then, the size of the output vector is the sum of the sizes of the input vectors and this is specified by its type.

```
append : Vect n a -> Vect m a -> Vect (n + m) a
```

This information allows the language to detect a larger class of programming errors. Note that type dependency information is not available for the analogous function on lists. Type dependency may also be used to express constraints on the inputs of a function, e.g. we can define a *total* function, called `pop`, that cannot be applied to an empty vector.

```
pop : Vect (S k) a -> Vect k a
pop (x :: xs) = xs
```

Writing "`pop []`" is now an error which is detected statically, rather than dynamically, and we note that the same cannot be achieved if we were to replace vectors with lists.

With the addition of dependent types, we sometimes have to prove the equivalence between types. For example, the types `Vect (n + k) Nat` and `Vect (k + n) Nat` are equivalent, but Idris cannot (at present) infer this, so the programmer has to sometimes provide a proof that addition is commutative (which is easy to do). All proof obligations needed for our development are simple and concise.

2.2.2 Linearity. The type system of Idris 2 is based on Quantitative Type Theory, where every function argument is associated with a multiplicity that states the number of times the variable is used at runtime. This multiplicity can be 0, 1 or ω . An argument with multiplicity 0 is only used at compile time (to determine type dependency information) and is erased at runtime. A *linear* argument has multiplicity 1 and it is used exactly once at runtime. Finally, ω represents the unrestricted multiplicity, which is also default, where the function argument may be used any number of times.

Example 2.2.2. Consider the `pop` function which we just discussed. The (implicitly bound) variables `k` and `a` have multiplicity 0, because they are not explicitly specified as separate arguments, and they are *not* accessible at runtime in the function. The variables `x` and `xs`, which are explicitly bound, have the default (unrestricted) multiplicity.

Example 2.2.3. An important type which we define in Qimaera is the type of linear vectors, which we write as `LVect`. The only difference, compared to the standard vectors in Idris, is that the `(::)` constructor for `LVect` is a linear function in all of its arguments. Linearity in Idris 2 is specified by writing the multiplicity 1 in front of each argument.

```
data LVect : Nat -> Type -> Type where
  Nil : LVect 0 a
  (::) : (1 _ : a) -> (1 _ : LVect k a) ->
        LVect (S k) a
```

We also use linear pairs that are already defined in Idris 2.

```
data LPair : Type -> Type -> Type
  (#) : (1 _ : a) -> (1 _ : b) -> LPair a b
```

Linearity allows us to specify and enforce constraints on function arguments, e.g. it prevents us from duplicating data, so the function definition below leads to an error:

```
copy : (1 _ : a) -> LPair a a
copy x = x # x
```

Error: While processing right hand side of `copy`. There are 2 uses of linear name `x`.

```

data Unitary : Nat -> Type where
  IdGate : Unitary n
  H      : (j : Nat) ->
    {auto prf : (j < n) = True} ->
    Unitary n -> Unitary n
  P      : (p : Double) -> (j : Nat) ->
    {auto prf : (j < n) = True} ->
    Unitary n -> Unitary n
  CNOT   : (c : Nat) -> (t : Nat) ->
    {auto prf1 : (c < n) = True} ->
    {auto prf2 : (t < n) = True} ->
    {auto prf3 : (c /= t) = True} ->
    Unitary n -> Unitary n

```

Fig. 3. The Unitary data type (file: Unitary.idr).

Linearity is prominently used in Qimaera. In particular, when manipulating quantum information, linearity is enforced in order to properly consume quantum resources and comply with the laws of quantum mechanics.

3 UNITARY OPERATIONS IN QIMAERA

As we saw in §2.1, unitary transformations have a special role in quantum computation. In most non-variational quantum algorithms, the vast majority of the programming effort consists in implementing the required unitary operations. In this section, we describe our representation of unitary transformations in Qimaera as an algebraic data type called `Unitary`. Every value of this type is, by design, an *algebraic decomposition* of a unitary operation in terms of the atomic unitary gates that we identified in §2.1.4.

The `Unitary` data type is very useful, because it allows us to adopt a high-level *algebraic* and *scalable* approach towards the reversible fragment of quantum computation. This provides the programmer with many benefits as we show in this section. However, using the `Unitary` data type is actually entirely optional. Users who are interested in effectful quantum programming do *not* have to use it (see §4) and they may still recover the full power of (variational) quantum programming, but at the cost of losing the algebraic decomposition of unitary operations. To utilise the full potential of Qimaera, programmers should make use of all the constructions described in this section and the next one.

3.1 The Unitary Data Type

Quantum unitary operations admit a compositional and algebraic representation based on the atomic gates from the universal gate set in §2.1.4. Our idea for the representation of unitary operations is based on this, or equivalently, on how unitary operations may be expressed in terms of unitary quantum circuit diagrams. Because of these reasons, linearity is not required for our formalisation of unitary operations. The code for the `Unitary` data type is listed in Figure 3 and we now describe our representation in greater detail.

Given a natural number $n : \text{Nat}$, the type of unitary operations on n qubits is given by `Unitary n`. Note that `Unitary` is an algebraic data type with a simple type dependency on the arity of the desired operation. The `Unitary` type has four different introduction rules which we describe next.

The first constructor, `IdGate`, represents the identity unitary operation on n qubits. Diagrammatically, we can see this as constructing a circuit of n wires, without applying any other gates on any of the wires. It has a unique argument, n , which is implicit – it can be omitted when calling the `IdGate` constructor and it will often be inferred by `Idris`.

The second constructor, `H`, should be understood as applying the Hadamard gate H to the j -th qubit of some previously constructed unitary circuit which is specified as the last argument. The first implicit argument, n , is simply the arity of the resulting unitary operation. The second implicit argument, `prf`, is a proof obligation that j is smaller than n . This ensures that the argument j identifies an existing wire of the previously constructed unitary circuit (last argument) and therefore the overall definition is algebraically and physically sound. We note that the implicit argument `prf` may be removed from our implementation if we change the type of j to `Fin n`, the type of integers less than n . However, in our experience, `Idris` has better support for `Nat` than for `Fin` and for this reason we chose to keep the `prf` argument.

The third constructor, `P`, should be viewed as applying the $P(p)$ gate, where the real number $p \in \mathbb{R}$ is approximated by the term `p : Double`.¹ The remaining arguments serve the same purpose as those for `H`.

The final constructor, `CNOT`, should be understood as applying the CNOT gate, where c identifies the wire used for the control (the small black dot in Figure 1), t identifies the wire of the target (the crossed circle in Figure 1) and the last (unnamed) argument is the previously constructed unitary circuit on which we are applying CNOT. The remaining arguments are implicit and often do not have to be provided by the users: the argument n is the arity of the unitary; `prf1` and `prf2` ensure that c and t identify valid wires of the unitary circuit; `prf3` ensures that the control and target wires are *distinct* and therefore the overall application of CNOT is physically and algebraically admissible.

In our representation of quantum unitary operations, we make use of type dependency to impose proof obligations on some of our constructors in order to guarantee that the representation makes sense in physical and algebraic terms. On first glance, this might seem like a burden for the users of the library. However, in our experience, `Idris` can often automatically infer the required proofs (without any assistance from the user) and we had to do very little manual theorem proving. This is discussed in detail in the next subsection.

3.2 Constructing Unitary Transformations

The four basic introduction rules of the `Unitary` type allow us to define *high-level functions* in `Idris` that can be used to construct complex unitary circuits out of simpler ones. We discuss this here and we show that the proof obligations from Figure 3 are not severe and can be easily ameliorated and often completely sidestepped using our high-level functions.

First, we point out that auto-implicit arguments may often be inferred by `Idris` via suitable search. For example, if all the arguments are known statically, the required proofs will be discovered by `Idris` and the users do not have to manually provide them.

Example 3.2.1. The unitary gate depicted in the circuit from Figure 2 may be constructed in the following way:

```
toBellBasis : Unitary 2
toBellBasis = CNOT 0 1 (H 0 IdGate)
```

¹This approximation is not a big limitation – in fault-tolerant quantum computing one usually replaces the $P(p)$ gate family with a single $T = P(\pi/4)$ gate and the resulting gate set suffices to approximate any unitary with *arbitrary* precision. So we can easily replace `P` with a `T` constructor.

In this example, Idris is able to infer all the implicit arguments and there is no need to provide any proofs. If we do not satisfy one of the constraints, e.g. if we write `CNOT 1 1` above (which does not make physical sense), then we get the following error during type checking:

```
Error : While processing right hand side of
toBellBasis. Can't find an implementation for
not (== 1 1) = True.
```

An error is also reported if we provide a wire number larger than 1. It is also useful to define *standalone* unitary gates for the H , $P(r)$ and CNOT gates as follows:

```
HGate : Unitary 1
HGate = H 0 IdGate

PGate : Double -> Unitary 1
PGate r = P r 0 IdGate

CNOTGate : Unitary 2
CNOTGate = CNOT 0 1 IdGate
```

3.2.1 Composing Unitary Circuits. Our libraries provide functions for sequential composition (`compose`) and parallel composition (`tensor`) of unitary operations:

```
compose : Unitary n -> Unitary n -> Unitary n
tensor : {n : Nat} -> {p : Nat} -> Unitary n -> Unitary p -> Unitary (n + p)
```

Notice that both functions do not impose any proof obligations on the user. This means that the *primary* algebraic way for composing unitary operations may be done without *any* need for theorem proving. The use of these functions is ubiquitous in practice and we introduce the infix synonyms `(.)` and `(#)` for `compose` and `tensor`, respectively.

Example 3.2.2. The `toBellBasis` gate from Example 3.2.1 may be equivalently expressed in the following way:

```
toBellBasis : Unitary 2
toBellBasis = CNOTGate . (HGate # IdGate)
```

Qimaera provides another, more general, form of composition via the function `apply` whose type is as follows:

```
apply : {i : Nat} -> {n : Nat} ->
        Unitary i -> Unitary n ->
        (v : Vect i Nat) ->
        {auto _ : isInjective n v = True} ->
        Unitary n
```

The `apply` function is used to apply a smaller unitary circuit of size i to a bigger one of size n , giving the vector v of wire indices on which we wish to apply the smaller circuit. It needs one auto-implicit proof which enforces the consistency requirement that all indices of the wires specified by v are pairwise distinct and smaller than n . In fact, the `apply` function implements the most general notion of composition possible and both sequential and parallel composition can be realised as special cases using it. The importance of the vector v is that it determines how to apply the smaller unitary circuit of arity i to *any selection* of i wires of the larger unitary circuit, and moreover, it also allows us to *permute* the inputs/outputs of the smaller unitary circuit while doing so. More specifically, if the k -th entry of the vector v is the natural number p , then the k -th

apply toBellBasis U [0,1]	
apply toBellBasis U [0,2]	
apply toBellBasis U [2,0]	
apply toBellBasis U [2,1]	

Table 1. Examples illustrating the apply function.

input/output of the smaller unitary circuit will be applied to the p -th wire of the larger unitary circuit. This is best understood by example.

Example 3.2.3. Consider the following code sample:

```
U : Unitary 3
U = HGate # IdGate {n = 1} # (PGate pi)
```

```
apply_example : Unitary 3
apply_example = apply toBellBasis U v
```

where v is a vector of length two. Here toBellBasis is given in Example 3.2.1 and represents the circuit given below left; U represents the circuit given below right:



Table 1 shows what unitary circuit is specified under different values of v . Here, Idris can automatically infer the required proofs and the user does not have to provide them.

Remark 3.2.4. Instead of using `apply`, there is another possible approach, in the spirit of *symmetric monoidal categories* [13, §XI], where we could add one extra introduction rule to the `Unitary` type for representing *permutations* of wires. However, in our view, this approach is somewhat awkward, because one does not usually think of permutations (induced by the symmetric monoidal structure) as physical gates.

3.2.2 Adjoints of Unitary Circuits. Qimaera also provides a function

```
adjoint : Unitary n -> Unitary n
```

which computes the adjoint (i.e., inverse) of a given unitary circuit. As explained previously, one often has to apply the inverse of a given unitary circuit, so having a high-level method such as this

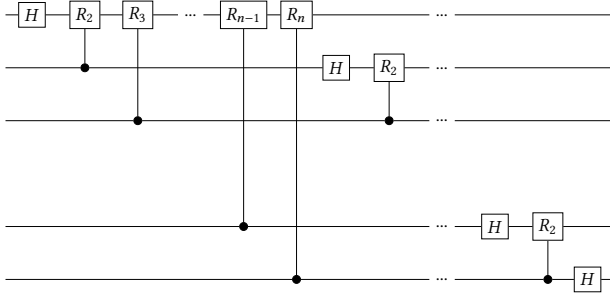


Fig. 4. The QFT unitary circuit on n qubits.

is useful. Our implementation uses the obvious algorithm for synthesising the adjoint. This may be used, for example, to automatically uncompute operations that we perform on ancilla qubits, which is often required.

3.2.3 Controlled Unitary Circuits. We also implement a function

```
controlled : {n : Nat} -> Unitary n -> Unitary (S n)
```

which given a unitary circuit U constructs the corresponding controlled unitary circuit CU . Our implementation uses the obvious algorithm for doing this, but more efficient algorithms may also be implemented in the future.

3.2.4 Analysis of Unitary Circuits. Unitary circuits are represented in a scalable and compositional way in Qimaera, thus we can use Idris to optimise them. The function:

```
optimise : Unitary n -> Unitary n
```

may be used to optimise a given (very large) unitary circuit. So far, this function provides some basic optimisations, but more sophisticated ones may be added in the future. The point we wish to make is that unitary circuits in Qimaera may be analysed and manipulated like any other algebraic data type using the full capabilities of Idris. In fact, the file `Unitary.idr` provides many functions that do this, e.g. we provide functions for calculating the circuit depth, calculating the number of specific atomic gates used by a circuit, drawing circuits in the terminal, exporting circuits to Qiskit (so that users may then use external analysis tools), etc.

3.3 Example: The Quantum Fourier Transform

The Quantum Fourier Transform (QFT) is a very important unitary operator that is used in Shor’s polynomial-time algorithm for integer factorisation [24]. The unitary circuit which realises QFT on n qubits is shown in Figure 4, where $R_n \stackrel{\text{def}}{=} P\left(\frac{2\pi}{2^n}\right)$. The Qimaera code which implements this unitary circuit is shown in Figure 5. Notice that we make use of the `controlled` function from §3.2.3 in the function `cRm`, so that we can automatically implement the many controlled R_n gates that are required. In this example, all the parameters are universally quantified, so we need a few very short and simple proofs in the code: one for using the `apply` function and one for correctly unifying the size of the circuit. Currently Idris cannot automatically discover these proofs, but we hope in the future its capabilities for proof search would improve to the point where it could. If this were to happen, then we would not have to manually provide these proofs.

```

Rm : Nat -> Unitary 1
Rm m = PGate (2 * pi / (pow 2 (cast m)))

cRm : Nat -> Unitary 2
cRm m = controlled (Rm m)

qftRec : (n : Nat) -> Unitary n
qftRec 0 = IdGate
qftRec 1 = HGate
qftRec (S (S k)) =
  let t = (qftRec (S k)) # IdGate
  in rewrite sym $ lemmaplusOneRight k
  in apply (cRm (S k)) t [S k, 0]
      {prf = lemmaInj1 k}

qft : (n : Nat) -> Unitary n
qft 0 = IdGate
qft (S k) =
  let g = qftRec (S k)
      h = (IdGate {n = 1}) # (qft k)
  in h . g

```

Fig. 5. Qimaera code for QFT (file: QFT.idr).

4 EFFECTFUL QUANTUM COMPUTATION

In the previous section we showed how unitary circuits can be represented in Qimaera in a compositional, algebraic and scalable way. This suffices to capture the pure, deterministic and reversible fragment of quantum computation. However, as we explained in §2.1, we need to also consider effectful and probabilistic quantum processes (e.g. measurements) in order to recover the full power of quantum computation. In this section we show how this can be done in a type-safe way by using monads, linearity and dependent types.

4.1 Representation of Quantum Effects

We now explain how the quantum program dynamics are represented in Qimaera in a type-safe way. We are (roughly) inspired by representing the notion of a *quantum configuration* as it appears in [12, 19, 22], which is in turn used to formally describe the operational semantics of quantum type systems.

4.1.1 Qubits in Qimaera. Because of the possibility of quantum entanglement (see §2.1.7), we cannot describe the state of an individual qubit which is part of a larger composite system. On the other hand, we wish to be able to refer to *parts* of the whole system by identifying specific qubit positions. In Qimaera, we introduce the following type declaration:

```

data Qubit : Type where
  MkQubit : (n : Nat) -> Qubit

```

The argument of type `Nat` is used as a *unique identifier* for the constructed qubit. The constructor `MkQubit` is *private* and users of our libraries cannot access it. Instead, our libraries provide functions (discussed later) that ensure that terms of type `Qubit` are always created with a fresh (i.e. unique) natural number that serves as its identifier. In fact, these functions are the only way users can access or manipulate qubits and, moreover, our users cannot access these unique identifiers. This allows us to formulate a representation where terms of type `Qubit` unambiguously refer to the relevant parts of larger composite systems. Therefore, a term of type `Qubit` should be understood as a pointer, or as a unique identifier, of a 1-qubit subsystem of some larger quantum state. Terms of type `Qubit` should *not* be understood as representing any sort of state, because they do not carry such information.

4.1.2 Probabilistic Effects. As we previously discussed in §2.1.9, quantum measurements induce probabilistic computational effects which are inherited by the classical side of the computation in variational quantum algorithms. Furthermore, in our intended computational scenario, the classical computer (on which Idris is running) sends instructions to, and receives data from, the quantum device. In order to correctly model all of this, it is clear that we have to use the IO monad in order to encapsulate these effects.

However, when representing quantum program dynamics, we also need to *enforce linearity*, but all the functions provided by the IO monad (e.g. `pure` which introduces pure values to monadic types) are *not* linear in any of their arguments. This creates a problem which may be solved by using the LIO library, which extends the IO monad with linearity. For brevity, we define `R` to be our linear IO monad:

```
R : Type -> Type
R = L IO {use = Linear }
```

Then, by using `R` we can combine IO effects (and thus also probabilistic effects) and linearity in a suitable way.

4.1.3 Quantum State Transformer. Quantum computation is *effectful*, and moreover, quantum information *cannot* be observed by the classical computer (on which Idris is running): it only receives classical information through communication with the quantum device. Because of this, we have to adopt a more abstract view on the hybrid classical-quantum computational process. In order to do this, we define an (abstract) *quantum state transformer* by combining several different concepts: *indexed state monads* [2]², linearity and IO (and thus also probabilistic) effects. Our representation of these ideas in Qimaera is shown in Figure 6, where we omit the function definitions for brevity.

The type `QStateT` is parameterised by a choice of three (arbitrary) types, so it is fairly abstract. Soon, we will see that it is very useful for our purposes. The intended interpretation of this type is the following: any value of type

$$\text{QStateT } \text{initialType } \text{finalType } \text{returnType}$$

represents a stateful (quantum) computation starting from a (quantum) state of type `initialType` and ending in a (quantum) state of type `finalType` which produces a user-accessible result of type `returnType` during the computation. For example, a value of type

$$\text{QStateT } (\text{LPair Qubit Qubit}) \text{ Qubit Bool}$$

should be understood as a quantum process that transforms a two-qubit state into a single-qubit state and returns a single (classical) value of type `Bool` to the user. The functions presented in Figure 6 allow us to adopt a *monadic programming discipline* when working with `QStateT` and we

²See [23] for a Haskell implementation of this idea.

```

data QStateT :
  Type -> Type -> Type -> Type where
  MkQST :
    (1 _ : (1 _ : initialType) ->
      R (LPair finalType returnType)) ->
    QStateT initialType finalType returnType

runQStateT : (1 _ : initialType) ->
  (1 _ : QStateT initialType
    finalType returnType) ->
  R (LPair finalType returnType)

pure : (1 _ : a) -> QStateT t t a

(>>=) : (1 _ : QStateT i m a) ->
  (1 _ : ((1 _ : a) -> QStateT m o b)) ->
  QStateT i o b

```

Fig. 6. Quantum state transformer (file: QStateT.idr).

```

interface QuantumOp (0 t : Nat -> Type) where
  newQubits : (p : Nat) ->
    QStateT (t n) (t (n+p)) (LVect p Qubit)

  applyUnitary : {n : Nat} -> {i : Nat} ->
    (1 _ : LVect i Qubit) -> Unitary i ->
    QStateT (t n) (t n) (LVect i Qubit)

  measure : {n : Nat} -> {i : Nat} ->
    (1 _ : LVect i Qubit) ->
    QStateT (t (i + n)) (t n) (Vect i Bool)

  run : QStateT (t 0) (t 0) (Vect n Bool) ->
    IO (Vect n Bool)

```

Fig. 7. The QuantumOp interface (file: QuantumOp.idr).

do so henceforth. We remark `QStateT` makes use of the monad `R` which encapsulates the IO (and probabilistic) effects. Note also that linearity is enforced when working with `QStateT`.

4.1.4 Effectful Quantum Programming. The `QStateT` monad can be used to define a suitable *abstract interface* for quantum programming. In Figure 7, we present an excerpt of the `QuantumOp` interface which allows us to easily write quantum programs and execute them in a type-safe way. All of the (variational) quantum algorithms we present are implemented using this interface that we describe next.

The function `newQubits` is used to prepare `p` new qubits in state $|0\rangle$ and the function returns a linear vector of length `p` with the qubit identifiers of the newly created qubits.

The function `applyUnitary` is used to apply a unitary operation of arity `i` to the qubits specified by the argument `LVect` (which also determines the order of application) and the operation returns an `LVect` which serves the same purpose – it identifies the qubits which were just modified by the unitary operator. The file `QuantumOp.idr` also provides functions `applyH`, `applyP` and `applyCNOT`

which can be seen as special cases of `applyUnitary`. However, these three functions do not depend on the `Unitary` type.

The measure function is used to measure `i` qubits identified by the `LVect` argument and it returns a value of type `Vect i Bool` that represents the result of the measurement. During this process, the `i` qubits are destroyed, as one can see from the provided type information.

Finally, the function `run` is used to *execute* quantum algorithms on the quantum device and obtain the classical information returned from it. Notice that `run` can be used to execute effectful quantum processes which start from the trivial quantum state (on zero qubits) and which terminate in the same trivial quantum state, but which also produce some number of classical bits as a user-accessible return result. This may be used to run quantum algorithms: in a typical situation, we start with the trivial quantum state (on zero qubits), we prepare n qubits in state $|0\rangle$, we apply some unitary operations on them, and we finally measure all the qubits, thereby destroying all the qubits and producing n bits of classical information. This quantum algorithm may then be represented as a value of type `QStateT (t \emptyset) (t \emptyset) (Vect n Bool)`. Running it, however, produces a classical value of type `IO (Vect n Bool)`, because the execution is probabilistic and because our classical computer (on which we are running Idris) has to perform IO actions to communicate with the quantum device.

In fact, *all* of the above operations modify the quantum state on the quantum device and may cause IO effects, because of the need to communicate with the quantum device. This is indeed reflected by our implementation. Observe, that our interface is defined using the `QStateT` monad transformer which does incorporate IO effects (via the `R` monad we discussed previously).

The best way to understand how this interface may be used for quantum programming is to look at examples.

Example 4.1.1. A fair coin toss may be implemented using quantum resources. The process is simple: (1) prepare the state $|0\rangle$; (2) apply the H gate to it; (3) measure the qubit and return this as output. We implement this as follows:

```
coin : QuantumOp t => IO Bool
coin = do
  [b] <- run (do
    q <- newQubit {t = t}
    q <- applyH q
    r <- measure [q]
    pure r
  )
  pure b
```

The top-level `do` block simply realises monadic sequencing for the standard IO monad. The `do` block within the `run` environment is more interesting and crucial for our development. It performs monadic sequencing for the `QStateT` monad and it represents the simple three-step algorithm we just described. The call to the `run` function executes this algorithm and users obtain the produced classical information by storing it in the variable `b` of type `Bool`. We emphasise that linearity is *enforced* within the `run` environment and this is what ensures the type safety of our approach, e.g. all of the following errors are statically detected and rejected by Idris: passing the qubit `q` to a non-linear function, copying the qubit `q`, forgetting to measure the qubit `q` so that it is implicitly discarded, etc. For example, if in the above code we replace the last two statements in the `run` environment with `"pure True"`, then Idris statically detects this error.

The function `coin` from Example 4.1.1 is implemented using our *abstract* interface. This means we can use this function in any *concrete* implementation of the `QuantumOp` interface. Since the

```

RUS : QuantumOp t => (1 _ : Qubit) ->
      (u' : Unitary 2) -> (e : Unitary 1) ->
      QStateT (t 1) (t 1) Qubit
RUS q u' e = do
  q' <- newQubit
  [q',q] <- applyUnitary [q',q] u'
  b <- measureQubit q'
  if b then do
    [q] <- applyUnitary [q] (adjoint e)
    RUS q u' e
  else pure q

example_u' : Unitary 2
example_u' = H 0 $ T 0 $ CNOT 0 1 $ H 0 $ CNOT 0 1 $ T 0 $ H 0 IdGate

runRUS : QuantumOp t => IO Bool
runRUS = do
  [b] <- run (do
    q <- newQubit {t = t}
    q <- RUS q example_u' IdGate
    measure [q]
  )
  pure b

testRUS : IO Bool
testRUS = runRUS {t = SimulatedOp}

```

Fig. 8. Repeat-until-success algorithm (file: RUS.idr).

authors do not have any quantum hardware, we provide one concrete implementation of this interface, called `SimulatedOp`, which performs linear-algebraic simulation of all the required operations. Once quantum hardware becomes more readily available, we can provide additional concrete implementations of the interface. For example, if we wish to use the `coin` function, then the code:

```

testCoin : IO Bool
testCoin = coin {t = SimulatedOp}

```

defines a new function, called `testCoin`, which does the same as `coin`, but it specifically instructs Idris to use linear-algebraic simulation. We emphasise that all of our quantum algorithms are written w.r.t. our abstract interface, so there is no need to reimplement them for any additional concrete implementations of the interface.

4.2 Example: Repeat-Until-Success Algorithm

Repeat-until-success (RUS) [17] is an algorithm for implementing quantum unitary operators by using *quantum measurements* and *recursion*. The main advantage in using RUS over traditional techniques that synthesise unitary operators, is that RUS usually requires fewer applications of the T gate, which is expensive in terms of error correction.

In the simplest case, we wish to realise a fixed single-qubit unitary operator $U : \mathbb{C}^2 \rightarrow \mathbb{C}^2$. The RUS algorithm is as follows. Given an input qubit $|\psi\rangle$, then: (1) prepare a new qubit in state $|0\rangle$; (2) apply a two-qubit unitary operator U' (chosen in advance depending on U); (3) measure the first qubit; (4) if the measurement outcome is 0 (which occurs with probability $p > 0$), then the output state is $U|\psi\rangle$, as required, and the algorithm terminates; otherwise the current state is $E|\psi\rangle$, where

E is some other unitary operator (chosen in advance depending on U), so we apply E^\dagger to this state and we go back to step (1). The unitary operators U' and E are chosen in advance, depending on U , before the algorithm starts so that the above conditions are satisfied.

This process always terminates in state $U|\psi\rangle$ (provided $p > 0$) so RUS indeed implements the unitary operator U . Note that this is an *algorithmic* realisation of U , not an algebraic one, and so we cannot write a program of type `Unitary` that achieves this. Instead, we represent this as a quantum program in Figure 8. There, `RUS q u'` is the quantum state transformer which is the RUS algorithm as above. The function `runRUS` simply executes the RUS algorithm on a qubit in state $|0\rangle$, with the unitary operator chosen from [17, Figure 8], then measures the qubit and returns the outcome. Both of these functions are written using our abstract interface. The function `testRUS` is the same as `runRUS`, but it also instructs Idris to use linear-algebraic simulation for the execution.

Remark 4.2.1. It is possible to make the RUS function non-linear in the qubit argument and still get a valid function. However, such a function cannot be called from the `run` environment (which enforces linearity) as we do in the `runRUS` function. This leads to an error which is statically detected by Idris. Therefore, such a non-linear RUS function would be useless, because we can never actually *run* it.

5 VARIATIONAL QUANTUM PROGRAMMING

In the previous section we saw that Qimaera is suitable for writing recursive and effectful quantum programs that make use of quantum measurements. Moreover, Idris 2 is an excellent programming language with an advanced type system and first-class support for classical programming features. In order to demonstrate that Qimaera is suitable for *variational* quantum programming, we also have to show that both classical and quantum programming features may be elegantly combined. This is the purpose of this section and we achieve this through the ultimate test – implementing the two most prominent variational quantum algorithms: the Quantum Algorithm for Approximate Optimisation (QAOA) [8] and the Variational Quantum Eigensolver (VQE) [20]. To the best of our knowledge, this is the first implementation of these algorithms that has been achieved in a type-safe framework.

General Framework. Both variational algorithms presented in this section are trying to find the minimum (or maximum) eigenvalue of a Hamiltonian. A Hamiltonian is a Hermitian (i.e., self-adjoint) matrix H . Its minimum eigenvalue is the minimum (real) value λ s.t. $H|\psi\rangle = \lambda|\psi\rangle$ for some nonzero vector $|\psi\rangle$. As H is unitarily diagonalizable, this is equivalent to the minimum of $\langle\psi|H|\psi\rangle$ for all vectors $|\psi\rangle$ of norm 1, where $\langle\psi| \stackrel{\text{def}}{=} |\psi\rangle^\dagger$.

Both algorithms proceed in the same way to find this value. They start with some assumption on what the vector $|\psi\rangle$ looks like and usually $|\psi\rangle$ is prepared by a quantum circuit that depends on some real parameters $\alpha_1, \dots, \alpha_p$. In the VQE algorithm, this is called the *Ansatz*.

By measuring this state $|\psi\rangle$, one obtains some information on the value of $\langle\psi|H|\psi\rangle$. This information can then be fed to a *classical* optimizer to change the value of the parameters $\alpha_1, \dots, \alpha_p$ for subsequent execution.

This classical-quantum back and forth is repeated until some satisfactory termination condition has been satisfied (e.g. simply repeat this k times, where $k \in \mathbb{N}$ is some large constant). However, in both algorithms, there is no guarantee that we will find the minimum eigenvalue.

QAOA. QAOA is a variational algorithm [8] that approximately solves optimization problems. Let $f : \{0, 1\}^n \rightarrow \mathbb{R}$ be a function for which we want to find its minimum. We see f as a diagonal Hamiltonian over n qubits defined by $H|x\rangle = f(x)|x\rangle$ for all $x \in \{0, 1\}^n$. We are therefore searching for the minimum eigenvalue of this Hamiltonian.

For QAOA to work, the Hamiltonian H should have a special form so that a circuit for $e^{\gamma H}$ is easy to make, where $\gamma \in \mathbb{R}$ is a parameter (to be tuned) that is used by the algorithm. A well-known and important example is to compute the maximum cut of an undirected graph, i.e., to solve the MAXCUT problem.

Our implementation for QAOA on the MAXCUT problem is presented in the file `QAOA.idr` and an excerpt is shown in Figure 9. The Ansatz depends on the graph G for which we want the maximum cut, a depth parameter p , and some real parameters β_i, γ_i . The depth parameter $p \in \mathbb{N}$ is usually fixed to be small, and we have a guarantee that the results of our algorithm become better when p becomes larger. The real parameters β_i and γ_i are used to determine rotation angles for some of the unitary operators that we need for constructing the relevant circuits.

Following the general framework, in our implementation, we have a function `QAOA_Unitary`, that takes these parameters as input and produces a unitary circuit that may be used to prepare the state $|\psi\rangle$ when applied to the initial state $|0\rangle^{\otimes n}$. We then measure this state $|\psi\rangle$ and present the result (a cut of the graph in the obvious binary encoding) to an Optimiser. Our optimiser is implemented by the function `classicalOptimisation` that uses all observable information from all previous runs (which amounts to the values of the parameters β_i, γ_i and the value of the cuts that have been previously obtained through quantum measurements) to compute the subsequent rotation parameters β_i, γ_i that we will use for the next iteration. The type of this function indicates that it uses the IO monad: this is because we wish to allow the function to use probabilistic optimisation algorithm or even external tools. The simplest implementation of this function chooses the rotation parameters at random.

The interplay between the classical and the quantum part is presented in Figure 9. The function `QAOA` takes as input an integer k representing how many times the whole routine will be done, the depth p of the circuit, and the graph G on which to compute the cut. Notice that the call to the quantum processor is isolated inside the `run` function.

VQE. In the VQE algorithm [20], one often assumes that H is a tensor product of *Pauli matrices* (which we do not define here for simplicity) and then one can use *Hamiltonian averaging*, i.e., we produce a unitary operator U_H that we apply to $|\psi\rangle$, with the following property: if we measure the first bit of $U_H |\psi\rangle$ and denote by p_0 (resp. p_1) the probability of the outcome being 0 (resp. 1), then $\langle \psi | H | \psi \rangle = p_0 - p_1$. We follow the same approach for our implementation which is available in the file `VQE.idr` (we do not provide an excerpt here for lack of space).

This unitary U_H is built by the function `encodingUnitary` and the process to compute $\langle \psi | H | \psi \rangle$ is given in the function `computeEnergyPauli`. This function runs the quantum process `nSamples` times and returns the difference between the two possible outcomes (0 and 1).

More generally, any Hamiltonian H can be written as a linear combination of tensor products of Pauli matrices, $H = \sum_i \alpha_i H_i$, and one can compute $\langle \psi | H | \psi \rangle$ by computing each term $\langle \psi | H_i | \psi \rangle$ independently, see the code for details.

6 RELATED WORK AND FUTURE WORK

In this section we compare Qimaera with other existing quantum programming languages that are implemented in software. We omit comparisons with quantum type systems that do not have a software implementation, mostly for brevity, but also because we do not feel such comparisons are fair towards the type systems, which are usually much smaller and with fewer features (but also more formal).

The QWIRE language [18, 21] is a quantum circuit language that is embedded in the Coq proof assistant [26]. In that sense, QWIRE is similar to Qimaera, because it also has access to dependent types (provided by Coq), and because it also clearly separates the quantum and classical modes

```

QAOA_Unitary : {n : Nat} -> (betas : Vect p Double) -> (gammas : Vect p Double) ->
  (graph: Graph n) -> Unitary n

classicalOptimisation : {p : Nat} -> (graph : Graph n) ->
  (previous_info : Vect k (Vect p Double, Vect p Double, Cut n)) ->
  IO (Vect p Double, Vect p Double)

QAOA' : QuantumOp t =>
  {n : Nat} ->
  (k : Nat) -> (p : Nat) -> (graph : Graph n) ->
  IO (Vect k (Vect p Double, Vect p Double, Cut n))
QAOA' 0 p graph = pure []
QAOA' (S k) p graph = do
  previous_info <- QAOA' {t} k p graph
  (betas, gammas) <- classicalOptimisation graph previous_info
  let circuit = QAOA_Unitary betas gammas graph
  cut <- run (do
    qs <- newQubits {t} n
    qs <- applyUnitary qs circuit
    measureAll qs
  )
  pure $ (betas, gammas, cut) :: previous_info

QAOA : QuantumOp t => {n : Nat} -> (k : Nat) -> (p : Nat) -> Graph n -> IO (Cut n)
QAOA k p graph = do
  res <- QAOA' {t} k p graph
  let cuts = map (\(_, _, cut) => cut) res
  let (cut, size) = bestCut graph cuts
  pure cut

```

Fig. 9. Qimaera implementation (excerpt) for the QAOA algorithm solving the MAXCUT problem (file: QAOA.idr).

of programming. However, Idris 2 is *not* a proof assistant (it is a programming language), it does not have access to an interactive proof system and it is not suitable for formally verifying complex properties of quantum programs. Indeed, this is what QWIRE excels at. The focus of Qimaera is on *programming* and Idris 2 has better support for classical, quantum and effectful programming features. In particular, since Coq lacks general recursion, the RUS algorithm from §4.2 cannot be expressed in QWIRE and the same is true for many other classical and quantum algorithms that make use of general recursion or probabilistic effects.

Silq [6] is a recently described standalone quantum programming language which is also type-safe (like Qimaera) and whose main notable feature is automatic uncomputation of temporary values. We currently partially support this feature, because we have clearly identified and separated the reversible fragment of quantum computation (see the `Unitary` type) and we can synthesise the required adjoints by calling the `adjoint` function. We believe that it is possible to extend Qimaera with full automatic uncomputation as well and we leave this for future work. Compared to Silq, the main advantage of Qimaera is that Idris has much better support for classical programming features and so we believe that Qimaera is a better choice for *variational* quantum programming, where the classical part of the algorithm is more complicated and difficult to program. In addition, Silq does not support general recursion, so it cannot express quantum algorithms that rely on this (e.g. the RUS algorithm from §4.2).

The Quipper language [11] is a domain specific language (DSL) embedded in Haskell. Qimaera is itself embedded in Idris 2 and clearly Idris 2 has been strongly inspired by Haskell, so the programming styles in Qimaera and Quipper are similar. Quipper has an extensive library of useful quantum functions and we have not implemented all of them in Qimaera yet. We believe that we can implement most of them and we leave this for future work. The main advantage that Qimaera has over Quipper is that Qimaera is type-safe, because it is implemented in Idris 2, which has dependent types and linearity. Haskell does not support these features, and because of this, Quipper cannot statically detect quantum programs that are physically inadmissible, whereas Qimaera can do this, as we already demonstrated in this paper.

Another recent language includes Proto-Quipper-D [9] which is a type-safe circuit description language. This language is based on a novel type system which shows how linearity and dependent types can be combined. A fundamental difference between Proto-Quipper-D and Qimaera is that linearity is the default mode of operation in Proto-Quipper-D, whereas in Qimaera the default mode is non-linear. The focus in Proto-Quipper-D is on *circuit* description and generation and the language currently lacks what is commonly known as *dynamic lifting*, i.e., the language does not support effectful quantum measurements and probabilistic effects. Because of this it cannot be used for variational quantum programming at present.

Other languages, include Google's Cirq [27] (a set of python libraries), IBM's Qiskit [1] (a set of python libraries) and Microsoft's Q# [25] (standalone). These languages offer a wide-range of quantum functions and features, however, none of them are type-safe and so it is possible to write erroneous quantum programs which are not detected at compile time. Qimaera does not have this problem and this is indeed its main advantage over them, together with dependent types.

REFERENCES

- [1] Gadi Aleksandrowicz, Thomas Alexander, Panagiotis Barkoutsos, Luciano Bello, Yael Ben-Haim, David Bucher, Francisco Jose Cabrera-Hernández, Jorge Carballo-Franquis, Adrian Chen, Chun-Fu Chen, Jerry M. Chow, Antonio D. Córcoles-Gonzales, Abigail J. Cross, Andrew Cross, Juan Cruz-Benito, Chris Culver, Salvador De La Puente González, Enrique De La Torre, Delton Ding, Eugene Dumitrescu, Ivan Duran, Pieter Eendebak, Mark Everitt, Ismael Faro Sertage, Albert Frisch, Andreas Fuhrer, Jay Gambetta, Borja Godoy Gago, Juan Gomez-Mosquera, Donny Greenberg, Ikko Hamamura, Vojtech Havlicek, Joe Hellmers, Łukasz Herok, Hiroshi Horii, Shaohan Hu, Takashi Imamichi, Toshi-nari Itoko, Ali Javadi-Abhari, Naoki Kanazawa, Anton Karazeev, Kevin Krsulich, Peng Liu, Yang Luh, Yunho Maeng, Manoel Marques, Francisco Jose Martín-Fernández, Douglas T. McClure, David McKay, Srujan Meesala, Antonio Mezzacapo, Nikolaj Moll, Diego Moreda Rodríguez, Giacomo Nannicini, Paul Nation, Pauline Ollitrault, Lee James O’Riordan, Hanhee Paik, Jesús Pérez, Anna Phan, Marco Pistoia, Viktor Prutyantov, Max Reuter, Julia Rice, Abdón Rodríguez Davila, Raymond Harry Putra Rudy, Mingi Ryu, Ninad Sathaye, Chris Schnabel, Eddie Schoute, Kanav Setia, Yunong Shi, Adenilton Silva, Yukio Siraichi, Seyon Sivarajah, John A. Smolin, Mathias Soeken, Hitomi Takahashi, Ivano Tavernelli, Charles Taylor, Pete Taylour, Kenso Trabing, Matthew Treinish, Wes Turner, Desiree Vogt-Lee, Christophe Vuillot, Jonathan A. Wildstrom, Jessica Wilson, Erick Winston, Christopher Wood, Stephen Wood, Stefan Wörner, Ismail Yunus Akhalwaya, and Christa Zoufal. 2019. *Qiskit: An Open-source Framework for Quantum Computing*. <https://doi.org/10.5281/zenodo.2562111>
- [2] Robert Atkey. 2009. Parameterised notions of computation. *J. Funct. Program.* 19, 3-4 (2009), 335–376. <https://doi.org/10.1017/S095679680900728X>
- [3] Robert Atkey. 2018. Syntax and Semantics of Quantitative Type Theory. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018, Oxford, UK, July 09-12, 2018*, Anuj Dawar and Erich Grädel (Eds.). ACM, 56–65. <https://doi.org/10.1145/3209108.3209189>
- [4] P.N. Benton. 1995. A mixed linear and non-linear logic: Proofs, terms and models. In *Computer Science Logic: 8th Workshop, CSL ’94, Selected Papers*. <https://doi.org/10.1007/BFb0022251>
- [5] P. N. Benton and P. Wadler. 1996. Linear Logic, Monads and the Lambda Calculus. In *LICS 1996*.
- [6] Benjamin Bichsel, Maximilian Baader, Timon Gehr, and Martin T. Vechev. 2020. Silq: a high-level quantum language with safe uncomputation and intuitive semantics. In *Proceedings of the 41st ACM SIGPLAN International Conference on Programming Language Design and Implementation, PLDI 2020, London, UK, June 15-20, 2020*, Alastair F. Donaldson and Emina Torlak (Eds.). ACM, 286–300. <https://doi.org/10.1145/3385412.3386007>
- [7] Edwin C. Brady. 2021. Idris 2: Quantitative Type Theory in Practice. In *35th European Conference on Object-Oriented Programming, ECOOP 2021, July 11-17, 2021, Aarhus, Denmark (Virtual Conference) (LIPIcs, Vol. 194)*, Anders Møller and Manu Sridharan (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 9:1–9:26. <https://doi.org/10.4230/LIPIcs.ECOOP.2021.9>
- [8] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. 2014. A Quantum Approximate Optimization Algorithm. arXiv:1411.4028 [quant-ph]
- [9] Peng Fu, Kohei Kishida, Neil J. Ross, and Peter Selinger. 2020. A Tutorial Introduction to Quantum Circuit Programming in Dependently Typed Proto-Quipper. In *Reversible Computation - 12th International Conference, RC 2020, Oslo, Norway, July 9-10, 2020, Proceedings (Lecture Notes in Computer Science, Vol. 12227)*, Ivan Lanese and Mariusz Rawski (Eds.). Springer, 153–168. https://doi.org/10.1007/978-3-030-52482-1_9
- [10] J.-Y. Girard. 1987. Linear Logic. *Theoretical Computer Science* 50 (1987), 1 – 101.
- [11] A. S. Green, P. L. Lumsdaine, N. J. Ross, P. Selinger, and B. Valiron. 2013. Quipper: a scalable quantum programming language. In *PLDI*. ACM, 333–342.
- [12] Xiaodong Jia, Andre Kornell, Bert Lindenhovius, Michael W. Mislove, and Vladimir Zamdzhiev. 2022. Semantics for Variational Quantum Programming. *Proc. ACM Program. Lang.* 6, POPL (2022). arXiv:2107.13347 To appear.
- [13] Saunders Mac Lane. 1998. *Categories for the Working Mathematician (2nd ed.)*. Springer.
- [14] Conor McBride. 2016. I Got Plenty o’ Nuttin’. In *A List of Successes That Can Change the World - Essays Dedicated to Philip Wadler on the Occasion of His 60th Birthday (Lecture Notes in Computer Science, Vol. 9600)*, Sam Lindley, Conor McBride, Philip W. Trinder, and Donald Sannella (Eds.). Springer, 207–233. https://doi.org/10.1007/978-3-319-30936-1_12
- [15] Jarrod R McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik. 2016. The theory of variational hybrid quantum-classical algorithms. *New Journal of Physics* 18, 2 (2016), 023023.
- [16] Michael A. Nielsen and Isaac L. Chuang. 2010. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511976667>
- [17] Adam Paetznick and Krysta M. Svore. 2014. Repeat-until-Success: Non-Deterministic Decomposition of Single-Qubit Unitaries. *Quantum Info. Comput.* 14, 15–16 (Nov. 2014), 1277–1301.
- [18] J. Paykin, R. Rand, and S. Zdancewicz. 2017. QWIRE: a core language for quantum circuits. In *POPL*. ACM, 846–858.

- [19] Romain Péchoux, Simon Perdrix, Mathys Rennela, and Vladimir Zamdzhiev. 2020. Quantum Programming with Inductive Datatypes: Causality and Affine Type Theory. In *Foundations of Software Science and Computation Structures - 23rd International Conference, FOSSACS 2020 (Lecture Notes in Computer Science, Vol. 12077)*. Springer, 562–581. https://doi.org/10.1007/978-3-030-45231-5_29
- [20] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J Love, Alán Aspuru-Guzik, and Jeremy L O’Brien. 2014. A variational eigenvalue solver on a photonic quantum processor. *Nature communications* 5, 1 (2014), 1–7.
- [21] Robert Rand, Jennifer Paykin, Dong-Ho Lee, and Steve Zdancewic. 2018. ReQWIRE: Reasoning about Reversible Quantum Circuits. In *Proceedings 15th International Conference on Quantum Physics and Logic, QPL 2018, Halifax, Canada, 3-7th June 2018 (EPTCS, Vol. 287)*, Peter Selinger and Giulio Chiribella (Eds.). 299–312. <https://doi.org/10.4204/EPTCS.287.17>
- [22] P. Selinger and B. Valiron. 2006. A lambda calculus for quantum computation with classical control. *Mathematical Structures in Computer Science* 16, 3 (2006), 527–552.
- [23] Kwang Yul Seo. 2017. Indexed State Monad Blog Post. <https://kseo.github.io/posts/2017-01-12-indexed-monads.html>. Accessed: 13.08.2021.
- [24] Peter W. Shor. 1999. Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. *SIAM Rev.* 41, 2 (1999), 303–332. <https://doi.org/10.1137/S0036144598347011>
- [25] Krysta Svore, Alan Geller, Matthias Troyer, John Azariah, Christopher Granade, Bettina Heim, Vadym Kliuchnikov, Mariia Mykhailova, Andres Paz, and Martin Roetteler. 2018. Q#: Enabling Scalable Quantum Computing and Development with a High-Level DSL. In *Proceedings of the Real World Domain Specific Languages Workshop 2018 (Vienna, Austria) (RWDSL2018)*. Association for Computing Machinery, New York, NY, USA, Article 7, 10 pages. <https://doi.org/10.1145/3183895.3183901>
- [26] Coq Development Team. 2021. The Coq Proof Assistant Reference Manual. <https://coq.inria.fr/distrib/current/refman/>. Accessed: 19.11.2021.
- [27] Google AI Quantum Team. 2021. Cirq. <https://quantumai.google/cirq>. Accessed: 13.08.2021.
- [28] William K Wootters and Wojciech H Zurek. 1982. A single quantum cannot be cloned. *Nature* 299, 5886 (1982), 802–803.