

Chapter 1

A Compiler for a Functional Quantum Programming Language

Jonathan Grattage and Thorsten Altenkirch

Abstract: We introduce a compiler for the functional quantum programming language QML [1], developed in Haskell. The compiler takes QML expressions as input and outputs a representation of quantum circuits (via the category **QFC** of finite quantum computations) which can be simulated by the simulator presented here, or by using a standard simulator for quantum gates. We discuss the structure of the compiler and how the semantic rules are compiled.

1.1 INTRODUCTION

Quantum programming is now a firmly established field; see recent text books as an example [2, 4]. However, quantum programs are usually presented as low level quantum gates. There are a number of proposals as to how best to integrate quantum effects in a high level language, [5, 6, 7, 9], but none of them introduce high level structures for quantum control and quantum data. In our previous paper we introduced the language QML [1] which realizes quantum control structures acting on quantum data, and which can, in principle, be translated in quantum circuits. This paper explains in detail how a compiler, which translates QML programs into quantum circuits, is implemented. The compiler itself is written in Haskell and its source code is available from <http://www.cs.nott.ac.uk/~jjg/qml>.

We begin by describing the structure of QML programs and the compiler. The compilation of QML terms is then discussed in detail, with emphasis on the differences between the strict and non-strict case constructs, and orthogonality judgements. This leads to the definition of a compiler following the denotational semantics and rules of QML.

1.2 SYNTAX OF QML

QML programs are typed expressions defined by the syntax and grammar of QML. The BNF grammar for QML types and expressions is:

$$\begin{aligned} \sigma &= 1 \mid \sigma \otimes \tau \mid \sigma \oplus \tau \\ t &= x \quad \mid \mathbf{let} \ x = t \ \mathbf{in} \ u \\ &\quad \mid x \uparrow y s \mid () \\ &\quad \mid (t, u) \mid \mathbf{let} \ (x, y) = t \ \mathbf{in} \ u \\ &\quad \mid \mathbf{qinl} \ t \mid \mathbf{qinr} \ u \\ &\quad \mid \mathbf{case} \ t \ \mathbf{of} \ \{ \mathbf{qinl} \ x \Rightarrow u \mid \mathbf{qinr} \ y \Rightarrow u' \} \\ &\quad \mid \mathbf{case}^\circ \ t \ \mathbf{of} \ \{ \mathbf{qinl} \ x \Rightarrow u \mid \mathbf{qinr} \ y \Rightarrow u' \} \\ &\quad \mid \{ (\kappa) t \mid (\iota) u \} \end{aligned}$$

As an illustration of a program, a simple example is the Hadamard operation on a single qbit, given by:

$$\begin{aligned} \mathit{had} &: \mathbf{Q}_2 \multimap \mathbf{Q}_2 \\ \mathit{had} \ x &= \mathbf{if}^\circ \ x \\ &\quad \mathbf{then} \ \{ \mathit{qfalse} \mid (-1) \ \mathit{qtrue} \} \\ &\quad \mathbf{else} \ \{ \mathit{qfalse} \mid \mathit{qtrue} \} \end{aligned}$$

This program can be read as an operation which, depending on the input qbit x , returns one of two possible superpositions of a qbit. Note that \mathbf{if}° is a special case of \mathbf{case}° . We can also easily calculate that applying the program twice returns the qbit to the original state, cancelling the amplitudes. Note that qtrue and qfalse are syntactic-sugar for $\mathbf{qinl} \ ()$ and $\mathbf{qinr} \ ()$, respectively. The type of a single qbit, denoted \mathbf{Q}_2 , is $1 \oplus 1$.

The compiler must have some way of interpreting these typed terms, and this is achieved by implementing the grammar in the compiler as Haskell datatypes: Ty representing the QML expression types and Tm representing the QML expressions themselves.

$$\begin{aligned} \mathbf{data} \ Ty &= TI \mid Ty \oplus Ty \mid Ty \otimes Ty \\ \mathbf{data} \ Tm &= Var \ String \quad \mid Let \ String \ Tm \ Tm \ (Maybe \ Ty) \\ &\quad \mid Weak \ Tm \ [String] \mid Void \\ &\quad \mid Tm \otimes Tm \quad \mid LetP \ String \ String \ Tm \ Tm \ (Maybe \ Ty) \\ &\quad \mid Inl \ Tm \quad \mid Inr \ Tm \\ &\quad \mid Case \ Tm \ (String, Tm) \ (String, Tm) \\ &\quad \mid Case^\circ \ Tm \ (String, Tm) \ (String, Tm) \ Orth \\ &\quad \mid Sup \ (\mathbb{C}, Tm) \ (\mathbb{C}, Tm) \\ &\quad \mid Embed \ Tm \end{aligned}$$

The final construct above in Tm , $Embed$, is not syntactic, but is a type level construct. It embeds a strict computation into a non-strict computation, see [1].

Returning to the example program had , passing it to the QML parser would rewrite it according to the internal representation, producing the term:

$$\begin{aligned} Case^\circ \ (Var \ "x") \ (& _1, Sup \ (1, \mathit{qfalse}) \ (-1, \mathit{qtrue})) \\ & \ (_x, Sup \ (1, \mathit{qfalse}) \ (1, \mathit{qtrue})) \ OSup \end{aligned}$$

where $OSup$ is an orthogonality judgement, which is required whenever the strict morphism $Case^\circ$ is used; see section 1.5.4.

1.3 QUANTUM MACHINE CODE

The low level machine code that our compiler generates is an **FQC** morphism that is represented as a quantum circuit. This is included along with some other information in the computation type, *Comp*:

data $Comp = Comp\{uCon \in Snoc\ String, ty \in Ty, fqc \in FQC\}$

This datatype is constructed by three functions. The first, *uCon* (used context), returns a list which contains the names of all variables that are used by the computation. This is for type correctness: a computation is not correct unless every variable passed to it by the context has been used. The second contains the returned type of the function, and the last contains the **FQC** morphism, represented by the type *FQC*:

data $FQC = FQC\{a, h, b, g \in Int, \phi \in Unitary\}$

This datatype is made up from five destructor functions: *a*, *h*, *b*, and *g*, which all return integers, and ϕ , which returns an object of type *Unitary*. *a* represents the ‘size’ of the input to this **FQC** (i.e, the size of the input type; the number of qbits required to represent it), *b* represents the size of the output, while *h* and *g* represent the size of the heap and garbage required, respectively. There is the condition that $a + h = b + g$, but this does not need to be enforced as the compiler never builds *FQC* objects where this is not true. $\phi \in A \otimes H \xrightarrow{\circ_{unitary}} B \otimes G$ represents a reversible quantum computation as a unitary operator on Hilbert spaces — for details see [1]


We represent unitary operators as circuits which can be constructed using the following combinators:

data $Unitary = \otimes [Unitary] \mid \odot [Unitary]$
 $\quad \mid Perm [Int] \quad \mid Cond\ Unitary\ Unitary$
 $\quad \mid Rotate\ (\mathbb{C}, \mathbb{C})\ (\mathbb{C}, \mathbb{C})$

\otimes takes a list of *Unitary*, and represents these operations acting in parallel (hence the use of the tensor symbol). Similarly, \odot represents the *Unitary* operations passed to it acting in series (the choice of symbol represents composition). *Perm*.*xs* is simply a permutation operator. *Rotate* represents any 1-qbit unitary rotation to the state passed, which is used to generate any given superposition - we can represent negation as a special rotation $unot = Rotate\ (0, 1)\ (1, 0)$. It follows from the Kitaev-Solovay theorem, see [4], pp. 616-624, that we can represent all unitary operators this way.

Finally, *Cond* represents a conditional operation, and is slightly unusual. If the control qbit is true then the first operation is performed on the remaining qbits, while if it is false the second operation is performed. If the control qbit is in a superposition, then both operations are applied to generate the correct superposition of outputs. This is different to the standard *controlled - U* operation, where the identity is applied if the control qbit is false, but *Cond* can be generated by simple application of two *controlled - U* operations and two *Not* operations, as shown

in the following circuit, for which *Cond* is a useful shorthand:



Two further functions are often used when defining *Unitary* operators, to aid understanding and to simplify what could otherwise become cumbersome definitions. These are *idu* $x \in \text{Int} \rightarrow \text{Unitary}$ and *swapN* $o \ l \ n \in \text{Int} \rightarrow \text{Int} \rightarrow \text{Int} \rightarrow \text{Unitary}$. The function *idu* simply creates an identity permutation of the required size. *swapN* creates another permutation operation, based on the three input integers, *o*, *l* and *n*. The permutation generated by this operation swaps a ‘block’ of qbits from anywhere in a circuit to the top, putting what was originally first after the moved block, followed by the remainder. *o* gives the size of the ‘offset’ to the block we wish to move, with *l* giving the size of this block. *n* is simply the size of the entire segment, for purposes of plumbing.

All objects of *Unitary* also have an associated arity; the number of qbits required to implement them. These are given by the following function:

```

arity ∈ Unitary → Maybe Int
arity (⊗ []) = Just 0
arity (⊗ (x:xs)) = do m ← arity x
                      n ← arity (⊗ xs)
                      return (m + n)
...
arity (Cond φ ψ) = do m ← arity φ
                      n ← arity ψ
                      guard (m ≡ n)
                      return (1 + m)

```

We define a function *eval* which calculates the matrix representation of a unitary circuit — here are some examples of how this mapping is achieved, with *matMult* being matrix multiplication, and *tensor* calculating the matrix tensor product:

```

type Row = [C]
type Matrix = [Row]
eval ∈ Unitary → Matrix
eval (⊗ xs) = mPar (map eval xs)
eval (⊙ xs) = mSeq (map eval xs)
...
mPar ∈ [Matrix] → Matrix
mPar [] = [[1]]
mPar (x:xs) = tensor x (mPar xs)
mSeq ∈ [Matrix] → Matrix
mSeq [x] = x
mSeq (x:xs) = matMult x (mSeq xs)

```

The compiler also allows a direct mapping from programs to matrices using

the *compute* rather than *compile* function. Additionally, the compiler also optimises *Unitary* types to give an optimised representation of the circuit, without changing the meaning.

1.4 STRUCTURE OF THE COMPILER

While the general structure of the compiler is quite standard, it is unusual because our language doesn't allow implicit weakenings. Hence we have to implement some book keeping of used variables. It takes as input parsed QML programs, and then performs a translation into its object language, in this case **FQC** morphisms. These morphisms, explained in [1], fully describe the computation, giving the size of input required, how much heap is required, how much garbage is produced, and a unitary operation that performs the computation. This unitary operation is essentially a description of a low-level quantum circuit (our 'machine' code). The **FQC** object can then be compiled further into the matrix that describes the computation (the lowest level 'machine' code).

The compiler starts by taking parsed QML programs, which we assume to be originally written as typed expressions, but have been converted by the parser into its internal representation. We have already stated that QML programs are typed expressions, and given the type for terms, *Tm*, and for QML types, *Ty*. These types are combined to give the final type for QML programs, implemented in the following way:

```

type Prog = Env FDef
data Env a = Env { unEnv ∈ Snoc (String, a) }
data FDef = FDef FSig Tm
data FSig = FSig Con Ty
type Con = Env Ty

```

These constructors tell us that a QML program (*Prog*) is an environment (*Env*) of function definitions (*FDef*). Function definitions are in turn defined as having a function signature, *FSig*, which contains the required context (*Con*) with the expected return type, and the QML term, *Tm*.

To illustrate the **FQC** morphism, the morphism generated for the quantum equivalent of the negation operation is described, which in QML syntax could be written:

```

qnot : Q2 → Q2
qnot x = ifo x
         then qfalse
         else qtrue

```

From this the parser generates:

```

qnot ∈ Prog
qnot = Env (SNil :< ("qnot", FDef (FSig (Env (SNil :< ("x", qb))) qb)
                  (Caseo (Var "x") ("_1", qfalse) ("_x", qtrue) (OInrl))))

```

The compiler now applies the compilation function to produce an **FQC** morphism

that performs the action of the program:

$$FQC\{a = 1, h = 0, b = 1, g = 0, \phi = Not\}$$

This is packaged up inside the *Comp* type, which also contains information about the variables used to compute the program, and the output type of the program:

$$Comp\{uCon = SNil:< "x", ty = \mathbf{Q}_2, FQC = \dots\}$$

The morphism ϕ in the *FQC* object describes the simple circuit that negates a single qbit, just the Not gate, while using the *eval* function described in section 1.3, we could further compile this into the standard matrix representation.

The actual mechanics of how compilation is achieved are described in the next section, 1.5.

1.5 COMPILING QML

The compiler for QML is defined as a single Haskell function, *compile*. This function compiles each subterm of the expression that represents the QML program into *Comp* objects, which contain the **FQC** morphism, by calling the recursive function *compileTm*.

$$compile \in Prog \rightarrow Error (Env Comp)$$

As can be seen from the type, *compile* takes as input a QML program and returns, if there are no errors, a list (environment) of computations. At present we only consider programs consisting of a single function, and hence only return a single computation. The *compile* function is mainly concerned with error checking, using the Error monad, and passes the work of compilation onto the *compileTm* function. If there are no type, arity, or context errors then it returns the computation, via *cleanComp*, which simply ‘cleans’ the computation by optimising the underlying unitary morphism that *compileTm* produced.

The function that performs the majority of the compilation, *compileTm*, has the type

$$compileTm \in (Env Comp) \rightarrow Con \rightarrow Tm \rightarrow Maybe Ty \rightarrow Error Comp$$

The type *(Env Comp)* simply contains the functions that have already been compiled, and are available for use by the compiler. *compileTm* also takes as input the current context, the term to be compiled, and the expected type (if known, hence the use of the *Maybe*) of the term. It returns either an error or a *Comp* object. The function is defined using pattern matching over the type of terms, *Tm*, so there is a separate function for each of the twelve term forms given by the grammar for QML, given by recursion over the term syntax. We begin by looking at how the variable rule is compiled, as it is the first rule – and is very simple.

$$compileTm _ \Gamma (Var\ x) mTy =$$

$$\mathbf{do} \ \sigma \leftarrow (elookup\ x\ \Gamma)$$

$$\mathbf{let} \ |\sigma| = size\ \sigma$$


$$greturn\ mTy\ (Comp\{uCon = (SNil:<x), ty = \sigma, \\ fqc = FQC\{a = |\sigma|, h = 0, b = |\sigma|, g = 0, \\ \phi = idu\ |\sigma|\}\})$$

To compile a variable, we first look up the type of the variable *x* in the context Γ , and call it σ . The *Env* lookup function, *elookup*, would return an Error type

if the variable did not exist. The size of the type is then calculated, and this information is used to generate the **FQC** morphism. The context used in compiling this term is simply the variable, hence *uCon* (*used-Context*) only contains x . The type of the computation is the type of the variable, σ , and the *FQC* is the identity morphism of the appropriate size. *greturn* is a small extension of the usual *return* function that confirms the type correctness of the returned computation, or else gives an informative error.

1.5.1 Context Sharing

Several of the more complicated circuit diagrams make use of a strict **FQC** morphism labelled ϕ_C .




This morphism is used in the compilation of several terms, and allows the interpretation of the operator \otimes on contexts. It makes use of δ_2 , which shares a single qbit using a simple controlled-not operation [1]. We define the ϕ_C morphism as $C_{\Gamma, \Delta} \in \mathbf{FQC}^\circ [[\Gamma \otimes \Delta]] ([[\Gamma]] \otimes [[\Delta]])$ by induction over the definition of $\Gamma \otimes \Delta$. If a variable $x : \sigma$ appears in both contexts we must use $\delta_\sigma \in \mathbf{FQC}^\circ [[\sigma]] ([[\sigma]] \otimes [[\sigma]])$ which generalises δ_2 by applying it in parallel to all qbits. All the other cases can be dealt with by applying monoidal isomorphisms. The implementation works by taking the context, $\Gamma \otimes \Delta$, and the two lists of the variables required by the operations; one of which requiring Γ as its context, and the other Δ . From this we can calculate which variables need to be shared and the appropriate permutations to output a **FQC** with the appropriate behaviour. The first $|\Gamma|$ qbits contain Γ , the remaining $|\Delta|$ contain Δ , and the list of variables in these two contexts (for verification purposes). This is arguably the most complicated function in the compiler.

1.5.2 Compiling Let

The QML *Let* statement is the standard construction, and has the usual interpretation. *Let* is, however, interesting, as it demonstrates some of the subtleties of context handling in QML. The *Let* rule and diagram are given as:

$$\frac{\Gamma \vdash^a t : \sigma \quad \Delta, x : \sigma \vdash^b u : \tau}{\Gamma \otimes \Delta \vdash^{a \uparrow b} \text{let } x = t \text{ in } u : \tau} \text{let}$$



The function begins with a recursive call to compute the subterm t , and then extracts some data:

$$\begin{aligned}
& \text{compileTm } c \Gamma (\text{Let } x \text{ } t \text{ } u \text{ } tTy) \text{ } mTy = \\
& \quad \mathbf{do} \ t^c \quad \leftarrow \text{compileTm } c \Gamma \ t \ tTy \\
& \quad \mathbf{let} \ nt \quad = \text{uCon } t^c \\
& \quad \quad \sigma \quad = \text{ty } t^c \\
& \quad \quad t^F \quad = \text{fqc } t^c
\end{aligned}$$

Next, we extend the context Γ with the typed variable $x : \sigma$. We call the extended context Γ' , and this is the context that will be passed to functions within the scope of the *Let* statement. This is achieved by calling the auxiliary function *extEnv*, which simply adds the pair (x, σ) to the context environment Γ :

$$\mathbf{let} \ \Gamma' = \text{extEnv } \Gamma \ (x, \sigma)$$

The compilation of the subterm over which our new variable has scope, u , can now be performed. This is done in the usual way; but passing Γ' as the context.

$$\begin{aligned}
& \quad u^c \quad \leftarrow \text{compileTm } c \Gamma' \ u \ mTy \\
& \quad \mathbf{let} \ nu \quad = \text{uCon } u^c \\
& \quad \quad \tau \quad = \text{ty } u^c \\
& \quad \quad u^F \quad = \text{fqc } u^c
\end{aligned}$$

We must now verify the subterm u did indeed use the new variable x . This is done by checking that the name of x appears in the list of variables used by the subterm, $nu = \text{uCon } u^c$. If it is there then it is removed to give us nu' , which gives us the list of variables in the context Γ that are used the subterm (giving us Γ from $\Gamma \otimes \Delta$, in the diagram). If x does not appear in nu then an *Error* is returned and propagated by the Error monad. This is done by the simple auxiliary function *findrm* (find and remove):

$$nu' \quad \leftarrow \text{findrm } x \ \sigma \ nu$$

Using the used context lists from the two subterms we can now calculate ϕ_C , using the function described before. From the diagram, the size of Γ is equal to $a \ t^F$, and the size of Δ in the diagram is given by $a \ c^F - a \ t^F$:

$$\begin{aligned}
& \mathbf{let} \ (ngd, c^F) = \phi_C \ \Gamma \ nt \ nu' \\
& \quad |\Delta| \quad = a \ c^F - a \ t^F \\
& \quad \text{greturn } mTy \ (\text{Comp} \{ \text{uCon} = ngd, \text{ty} = \tau,
\end{aligned}$$

Finally, we compute the *FQC* morphism using this data, with ϕ following the construction in the diagram exactly, which is returned after *greturn* verifies the inferred and expected types match:

$$\begin{aligned}
& \text{fqc} = \text{FQC} \{ a = a \ c^F, h = h \ c^F + h \ t^F + h \ u^F, b = b \ u^F, g = g \ t^F + g \ u^F, \\
& \quad \phi = \odot \ [\otimes \ [\phi \ c^F, \text{idu} \ (h \ t^F + h \ u^F)], \\
& \quad \quad \text{swapN} \ (a \ t^F) \ |\Delta| \ (a \ t^F + |\Delta| + h \ t^F + h \ u^F), \\
& \quad \quad \otimes \ [\text{idu} \ |\Delta|, \phi \ t^F, \text{idu} \ (h \ u^F)], \\
& \quad \quad \otimes \ [\text{idu} \ (a \ u^F), \text{swapN} \ (h \ t^F) \ (h \ u^F) \ (h \ t^F + h \ u^F)], \\
& \quad \quad \otimes \ [\phi \ u^F, \text{idu} \ (g \ t^F)], \\
& \quad \quad \otimes \ [\text{idu} \ (b \ u^F), \text{swapN} \ (g \ u^F) \ (g \ t^F) \ (g \ u^F + g \ t^F)] \} \}
\end{aligned}$$

Note how the construction of ϕ follows the construction of the circuit diagram, with each line representing a column of the diagram. The computation of *uCon* is simply the shared context, and the type follows from the computation. The *FQC* constructors a , h , b , and g are easily calculated.

The implementation of the diagram for \otimes introduction, also denoted $t \otimes u$ (and

akin to pairing), is very similar. We simply compile the two subterms, generate the appropriate context for each using ϕ_C , and then construct the *FQC* object following the circuit diagram given in [1].

1.5.3 Compiling $Inl (\oplus)$

A third interesting example is the compilation of injections, which shows some of the complications implementing the diagrammatic form of the compilation; specifically, care is required to ensure that the subcomputations are ‘plumbed’ correctly into the rest of the computation. The rule and compilation diagram for constructing $Inl s$ are given as:

$$\frac{\Gamma \vdash^a s : \sigma}{\Gamma \vdash^a \text{inl } s : \sigma \oplus \tau} + \text{intro}_1$$

The rule tells us that from a term s of some type σ we can compute $Inl s : \sigma \oplus \tau$. The returned type cannot be inferred from the input as we know nothing about τ . For this reason, in order to compile Inl , the *compileTm* function must be passed the expected return type, mTy which should be *Just* $\sigma \oplus \tau$. *Inr*, *Case*, and *Case^o* also need this information for the same reason.

These first few lines of the definition for Inl take the type input and split it into its component types using the simple function *unPlusMaybe*. These are then extracted from the *Maybe* monad using *unMaybeTy* to give the expected return type of the computed left injection.

$$\begin{aligned} \text{compileTm } c \Gamma (Inl s) mTy = \\ \mathbf{do} (\sigma^?, \tau^?) \leftarrow \text{unPlusMaybe } mTy \\ \sigma' \quad \leftarrow \text{unMaybeTy } \sigma^? \\ \tau' \quad \leftarrow \text{unMaybeTy } \tau^? \end{aligned}$$

We next compile the subterm s by calling *compileTm* again, and passing $\sigma^?$ as the expected type of the subterm.

$$\begin{aligned} s^C \quad \leftarrow \text{compileTm } c \Gamma s \sigma^? \\ \mathbf{let} \ \sigma \quad = \text{ty } s^C \\ s^F \quad = \text{fqc } s^C \end{aligned}$$

There is no restriction that the sizes of σ and τ must be equal, so if the types are of different sizes we have to pad the smaller one to make them the same size, using the padding operator ϕ_P .

$$\begin{aligned} \max^{(s,t)} &= \max (\text{size } \sigma') (\text{size } \tau') \\ \delta &= \max^{(s,t)} - (\text{size } \sigma) \\ p^F &= \phi_P (b s^F) \delta \end{aligned}$$

Next we confirm that the inferred type for the subcomputation was indeed the type returned. If they match, then the computation is constructed, directly following the circuit diagram:

```

if  $\sigma \neq \sigma'$  then Error ("Inl type mismatch: " ++ ...)
else greturn mTy (Comp{uCon = uCon  $t^c$ , ty =  $\sigma \oplus \tau'$ ,
fqc = FQC{a =  $a s^F$ , h =  $\delta + 1$ , b =  $\max^{(s,t)} + 1$ , g =  $g s^F$ ,
 $\phi$  =  $\odot$  [  $\otimes$  [ $\phi s^F$ , idu  $\delta$ , unot],
 $\otimes$  [idu ( $b s^F$ ), swapN ( $g t^F$ )  $\delta$  ( $g s^F + \delta + 1$ )],
 $\otimes$  [ $\phi p^F$ , swapN ( $g s^F$ ) 1 ( $1 + g s^F$ )] ] })


```

1.5.4 Compiling *Case* & *Case^o*


In QML there are two versions of the traditional case construct. The first, *Case*, measures a qbit of the data over which we branch to determine which branch has to be returned. By measuring, it collapses part of the quantum state, and causes entangled states to decohere. The second, *Case^o*, is an example of the *quantum control* available in QML. It does not measure the result of executing the conditional statement, but produces a superposition of results which corresponds to the superposition of data it analyses. It is the source of quantum parallelism; intuitively both branches are run in parallel. However, there is a price to pay: we have to provide evidence that both branches are observably different, i.e. orthogonal. Only then are we able to translate our quantum case analysis into quantum circuits.

The rules and circuits for compiling *Case* and *Case^o* are:

$$\frac{\Gamma \vdash c : \sigma \oplus \tau \quad \Delta, x : \sigma \vdash t : \rho \quad \Delta, y : \tau \vdash u : \rho}{\Gamma \otimes \Delta \vdash \text{case } c \text{ of } \{\text{inl } x \Rightarrow t \mid \text{inr } y \Rightarrow u\} : \rho} \oplus\text{-elim}$$



$$\frac{\Gamma \vdash^a c : \sigma \oplus \tau \quad \Delta, x : \sigma \vdash^o t : \rho \quad \Delta, y : \tau \vdash^o u : \rho \quad t \perp u}{\Gamma \otimes \Delta \vdash^a \text{case}^o c \text{ of } \{\text{inl } x \Rightarrow t \mid \text{inr } y \Rightarrow u\} : \rho} \oplus\text{-elim}^o$$



Note both of these circuits make use of the context sharing morphism ϕ_C , defined earlier.

Looking at compiling *Case* first, the function begins by computing the *Comp* object for the conditional term c^T , from which the subtypes (via *unPlus*), context required and **FQC** morphism are extracted:

```

compileTm c  $\Gamma$  (Case  $c^T$  (x,t) (y,u)) mTy =
  do cC      ← compileTm c  $\Gamma$  cT Nothing
    ( $\sigma, \tau$ ) ← unPlus (ty cC)
  let nc      = uCon cC
      cF     = fqc cC

```

Note that the type passed when compiling c^T is *Nothing*. This is because we cannot infer the type from the information given, and therefore restricts what terms conditional statements can be composed of. Next, two extended contexts are created, $\Gamma^{x:\sigma}$ for the subterm t , and $\Gamma^{y:\tau}$ for u . These are then used in the compilation of the subterms t and u :

```

 $\Gamma^{x:\sigma}$  = extEnv  $\Gamma$  (x,  $\sigma$ )
 $\Gamma^{y:\tau}$    = extEnv  $\Gamma$  (y,  $\tau$ )
tC          ← compileTm c  $\Gamma^{x:\sigma}$  t mTy
uC          ← compileTm c  $\Gamma^{y:\tau}$  u mTy

```

Similar to the implementation of the *Let* construct, we now ensure the added variables were used and remove them from the appropriate list (or else fail). The types and *FQC* objects are then extracted:

```

nt      ← findrm x  $\sigma$  (uCon tC)
nu      ← findrm y  $\tau$  (uCon uC)
let  $\rho^t$  = ty tC
     $\rho^u$   = ty uC
    tF    = fqc tC
    uF    = fqc uC

```

We can now compute the context sharing using ϕ_C , to give us Γ and Δ from the *diagram*, and then compute the $[t|u]^F$ morphism. This is again done by an auxiliary function, *condOp*, which creates an *FQC* object using the *Cond* operator on t^F and u^F , appropriately padded so they are the same size, and plumbed correctly. Some useful mnemonics are then created to help in the definition of ϕ :

```

(–, CF) =  $\phi_C$   $\Gamma$  nc nt
[t|u]F  = condOp uF tF True
| $\Gamma$ |  = a cF
| $\Delta$ |  = a CF + h CF – | $\Gamma$ |
arity   = | $\Gamma$ | + | $\Delta$ | + h cF + h [t|u]F

```

We now begin to prepare the *Comp*, but first must ensure that that t and u make use of the same variables from the original context Γ ; otherwise there is a context error, and that ρ^t and ρ^u also agree:

```

if nt ≠ nu then Error "Context error"
  else if  $\rho^t$  ≠  $\rho^u$  then Error "Type error"
  else greturn mTy (Comp { uCon = nc ++< nt, ty =  $\rho^t$ ,

```

Finally, the construction of the *FQC* can begin. The values of a , h , b , and g can

be read from the diagram, and again the morphism is constructed using permutations and the subcomputations generated above, to match the construction of the diagram:

$$\begin{aligned}
fqc &= FQC\{a = a^{C^F}, \quad h = \quad \quad \quad h^{C^F} + h^{c^F} + h[t|u]^F, \\
&\quad b = b[t|u]^F, g = \quad \quad \quad g[t|u]^F + g^{c^F}, \\
\phi &= \odot [\otimes [\phi^{C^F}, idu(h^{c^F} + h[t|u]^F)], \\
&\quad \quad \quad swapN |\Gamma| |\Delta| arity, \\
&\quad \quad \quad \otimes [idu |\Delta|, \phi^{c^F}, idu(h[t|u]^F)], \\
&\quad \quad \quad \otimes [idu (|\Delta| + b^{c^F}), \\
&\quad \quad \quad \quad \quad \quad swapN (g^{c^F}) (h[t|u]^F) (g^{c^F} + h[t|u]^F)], \\
&\quad \quad \quad \otimes [\phi[t|u]^F, idu(g^{c^F})]]] \}
\end{aligned}$$

As the diagram for $Case^o$ is very similar to that for $Case$, so is the compilation function. Only the differences are discussed:

$$\begin{aligned}
compileTm c \Gamma (Case^o c^T (x, t) (y, u) o) mTy &= \\
\mathbf{do} \ c^C &\leftarrow compileTm c \Gamma c^T \text{ Nothing} \\
(\sigma, \tau) &\leftarrow unPlus (ty \ c^C) \\
\mathbf{let} \ nc &= uCon \ c^C \\
\quad c^F &= fqc \ c^C \\
\quad \Gamma^{x:\sigma} &= extEnv \ \Gamma \ (x, \sigma) \\
\quad \Gamma^{y:\tau} &= extEnv \ \Gamma \ (y, \tau)
\end{aligned}$$

The next step in the implementation of $Case$ is to compile the two subterms, t and u , which are passed to the $condOp$ to give $[t|u]^F$. However, note that in this case, from the diagram, we want $[f|g]^F$, where f and g are computations based on t and u , but modified using the orthogonality judgement. This is carried out by the $orthcomp$ function, see section 1.5.5. If this function fails, then we know that the terms t and u were not orthogonal. The compilation then proceeds as in $Case$:

$$\begin{aligned}
(f^C, g^C, \rho) &\leftarrow orthcomp \ c \ (t, \sigma, \Gamma^{x:\sigma}) \ (u, \tau, \Gamma^{y:\tau}) \ mTy \ o \\
nf &\leftarrow findrm \ x \ \sigma \ (uCon \ f^C) \\
ng &\leftarrow findrm \ y \ \tau \ (uCon \ g^C) \\
\mathbf{let} \ f^S &= ty \ f^C \\
\quad g^S &= ty \ g^C \\
\quad f^F &= fqc \ f^C \\
\quad g^F &= fqc \ g^C \\
(_, C^F) &= \phi_C \ \Gamma \ nc \ nf \\
[f|g]^F &= condOp \ g^F \ f^F \ False \\
|\Gamma| &= a \ c^F \\
|\Delta| &= a \ C^F + h \ C^F - |\Gamma| \\
arity &= |\Gamma| + |\Delta| + h \ c^F + h \ [f|g]^F \\
orth &= orth2unitary \ f^S \ g^S \ o
\end{aligned}$$

Now we must compile the orthogonality judgment morphism. This is done using the auxiliary function $orth2unitary$, which takes the type of f and g , and the judgment o . This function simply computes the circuit ($Unitary$) for the orthogonality judgment, and is explained in section 1.5.5.

$$\mathbf{if} \ nf \neq ng \ \mathbf{then} \ Error \ "Context \ error"$$

```

else if  $f^S \neq g^S$  then Error "Type error"
else if orthcheck  $t u o \equiv \text{False}$ 
  then Error "Orthogonality Failure"
  else greturn mTy (Comp{uCon =  $nc ++ \langle nf, ty = \rho$ ,

```

The computation of the *FQC* morphism continues in the exact same way as before, but for the addition of the extra *orth* construct, as in the final column of the diagram:

$$\begin{aligned}
fqc = FQC\{ & a = a^{C^F}, \quad h = h^{C^F} + h^{c^F} + h[f|g]^F, \\
& b = b[f|g]^F, \quad g = g^{c^F}, \\
\phi = \odot [& \otimes [\phi^{C^F}, idu(h^{c^F} + h[f|g]^F)], \\
& \text{swapN } |\Gamma| \text{ } |\Delta| \text{ } \text{arity}, \\
& \otimes [idu \text{ } |\Delta|, \phi^{c^F}, idu(h[f|g]^F)], \\
& \otimes [\quad \quad \quad idu(|\Delta| + b^{c^F}), \\
& \quad \quad \quad \text{swapN}(g^{c^F})(h[f|g]^F)(g^{c^F} + h[f|g]^F), \\
& \otimes [\phi[f|g]^F, idu(g^{c^F})], \\
& \otimes [orth, idu(g^{c^F})]]] \}
\end{aligned}$$

1.5.5 Compiling *Orth* (\perp)

Here we briefly introduce the type of orthogonality judgments, *Orth*, and discuss how the *orth2unitary* function and the *orthcomp* function, used in *Case^o*, behave for some examples of the *Orth* type. The *Orth* type is defined as:

```

data Orth = OInlr   | OInrl
          | OInl Orth | OInr Orth
          | OPairl   | OPairr
          | OSup Orth Orth ( $\mathbb{C}, \mathbb{C}$ ) ( $\mathbb{C}, \mathbb{C}$ )

```

which matches the seven rules given in [1].

The idea of $t \perp u$ is that there exists a boolean observation which tells the two terms apart in every environment. Given $\Gamma \vdash t, u : \rho$, the interpretation $\llbracket t \perp u \rrbracket = (S_{t \perp u}, \phi_{t \perp u})$, where $\phi_{t \perp u} \in S_{t \perp u} \otimes \mathcal{Q}_2 \xrightarrow{\circ} \text{unitary } \llbracket \rho \rrbracket$ is defined by induction over the derivation.

OInlr & *OInrl*


$$\frac{}{\text{inl } t \perp \text{inr } u} \text{OInlr} \qquad \frac{}{\text{inr } t \perp \text{inl } u} \text{OInrl}$$

Given $\Gamma \vdash^\circ t : \sigma$ and $\Gamma \vdash^\circ u : \tau$, we define $S = \sigma \sqcup \tau$ and $\phi_{t \perp u}$ is simply the identity in the first case but negates the qbit in the second.

$$\left. \begin{array}{c} s \text{ ---} \\ \mathcal{Q}_2 \text{ ---} \end{array} \right\} \rho \qquad \left. \begin{array}{c} s \text{ ---} \\ \mathcal{Q}_2 \text{ ---} \boxed{X} \text{ ---} \end{array} \right\} \rho$$

This can be defined by the function *orth2unitary* as:
orth2unitary $\in Ty \rightarrow Ty \rightarrow Orth \rightarrow Unitary$

$orth2unitary \ t \ u \ (OInl) = \otimes [idu (\max (size \ t) (size \ u)), idu \ 1]$
 $orth2unitary \ t \ u \ (OInr) = \otimes [idu (\max (size \ t) (size \ u)), Not]$
 For both of these rules, ϕ_f is simply given by ϕ_t :



while ϕ_g is given by ϕ_u , in the same way. This is also straightforward:


$orthcomp \in Code \rightarrow (Tm, Ty, Con) \rightarrow (Tm, Ty, Con)$
 $\rightarrow Maybe \ Ty \rightarrow Orth \rightarrow Error (Comp, Comp, Ty)$
 $orthcomp \ c \ (Inl \ t, \sigma, \Gamma^{x:\sigma}) \ (Inr \ u, \tau, \Gamma^{y:\tau}) \ mTy \ (OInl) =$
do $f^c \leftarrow compileTm \ c \ \Gamma^{x:\sigma} \ t \ (Just \ \sigma)$
 $g^c \leftarrow compileTm \ c \ \Gamma^{y:\tau} \ u \ (Just \ \tau)$
 $t^c \leftarrow compileTm \ c \ \Gamma^{x:\sigma} \ (Inl \ t) \ mTy$
return $(f^c, g^c, ty \ t^c)$

OInr is identical. Note that the function also returns the type of the full subterm, for use by *Case*^o. The simplest way of doing this (though least efficient) is to compute the full subterm as well, and just extract the type.

OInl Orth

$$\frac{t \perp u}{inl \ t \perp \ inl \ u \ \ inr \ t \perp \ inr \ u} \ OInl$$

Consider the first rule, *OInl*, only: We set $S = S_{inl \ t \perp \ inl \ u} = S_{t \perp u} \otimes Q_2$, to derive $\phi_{inl \ t \perp \ inl \ u} \in (S_{t \perp u} \otimes Q_2) \otimes Q_2 \xrightarrow{\text{unitary}} [\sigma \sqcup \tau] \otimes Q_2$ from $\phi_{t \perp u} \in S_{t \perp u} \otimes Q_2 \xrightarrow{\text{unitary}} [\sigma]$ we just swap two qbits. The other case proceeds analogously.



The code for creating this using the *orth2unitary* function is simply:


$orth2unitary \ t \ u \ (OInl \ o) = \odot [\otimes [idu \ \max^{(t,u)}, swapN \ 1 \ \delta^{(t,u)} (\delta^{(t,u)} + 2)],$
 $\otimes [idu (\max^{(t,u)} + \delta^{(t,u)}), Perm [1, 0]],$
 $\otimes [idu \ \max^{(t,u)}, swapN \ 1 \ \delta^{(t,u)} (\delta^{(t,u)} + 2)],$
 $\otimes [orth2unitary \ t \ u \ o, idu (\delta^{(t,u)} + 1)],$
 $\otimes [\phi (\phi_P (\max^{(t,u)} + 1) \ \delta^{(t,u)}), idu \ 1]]$

where $\max^{(t,u)} = \max (size \ t) (size \ u)$

$\delta^{(t,u)} = \text{abs} ((size \ t) - (size \ u))$

$orth2unitary \ t \ u \ (OInr \ o) = orth2unitary \ t \ u \ (OInl \ o)$

Calculating ϕ_f and ϕ_g is more involved, as $\phi_{f_{t \perp u}}$ and $\phi_{g_{t \perp u}}$ must also be calculated. In the case where we have $inl \ t \perp \ inl \ u$, ϕ_f is given by:




Note the output, $S_{t \perp u} \otimes Q_2 \otimes H_{t-u}$, includes S . ϕ_g can be constructed in the same way, substituting $\phi_{g_t \perp u}$ for $\phi_{f_t \perp u}$. In the case where we have $\text{inr} \perp \text{inr}$, the construction of ϕ_f (and ϕ_g) is identical, except for a negation on the final wire – the control qbit. The code for *orthcomp*, which calculates these trivial circuits, is not shown here.

OPairl

$$\frac{t \perp u}{(t, v) \perp (u, w) \quad (v, t) \perp (w, u)} \text{OPairl}$$

Considering only the first rule above, given $\Gamma \vdash^\circ t, u : \sigma$ and $\Gamma \vdash^\circ v, w : \tau$. We set $S = S_{(t,v) \perp (u,w)} = S_{t \perp u} \otimes \tau$, to derive $\phi_{(t,v) \perp (u,w)} \in S_{t \perp u} \otimes \tau \otimes Q_2 \xrightarrow{\text{unitary}} [\sigma \otimes \tau]$ from $\phi_{t \perp u} \in S_{t \perp u} \otimes Q_2 \xrightarrow{\text{unitary}} [\sigma]$ requires simple rewiring.



This is computed by:

$$\text{orth2unitary } \sigma \tau \text{OPairl} = \odot [\otimes [\text{idu } |\sigma'|, \\ \text{swapN } |\tau| 1 (|\tau'| + 1)], \\ \otimes [\text{orth2unitary } \sigma' \tau' (\text{OInrl}), \\ \text{idu } |\tau'|]]$$

$$\begin{aligned} \text{where } |\sigma'| &= \text{size } s' \\ |\tau'| &= \text{size } t' \\ \text{OK } (\sigma', \tau') &= \text{unPair } \tau \end{aligned}$$

To calculate ϕ_f in this instance, we make use of ϕ_t and ϕ_v , to construct $\phi_t \otimes \phi_v$, which is an application of the \otimes - introduction rule, and can therefore be calculated simply by a call to *compileTm* ($t \otimes v$). ϕ_g is constructed analogously, replacing ϕ_t with ϕ_u and ϕ_v with ϕ_w . For the case where the second rule is applied, the same constructions are used, but with t and v (u and w for ϕ_g) swapped appropriately.

1.6 CONCLUSION

Our present work exploits the functional paradigm in two ways: first we have designed a high level quantum programming language — QML — using con-

cepts from conventional functional programming [1]; and then in the present paper we have used Haskell for a prototypical, high level implementation of QML. In related work [8] use Haskell to explore the semantic foundation of QML by implementing Super operators using the arrows [3].

ACKNOWLEDGEMENTS

Thanks are extended by the authors to Conor McBride, Peter Selinger, Amr Sabry, Juliana Vizzotto, Graham Hutton, and Janine Swarbrick.

REFERENCES

- [1] T. Altenkirch and J. Grattage. A functional quantum programming language. quant-ph/0409065, November 2004.
- [2] M. Hirvensalo. *Quantum Computing*. Springer-Verlag NewYork, Inc., 2001.
- [3] J. Hughes. Generalising monads to arrows. *Science of Computer Programming*, 37:67–111, May 2000.
- [4] M. Nielsen and I. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, Cambridge, 2000.
- [5] P. Selinger. Towards a quantum programming language. *Mathematical Structures in Computer Science*, 2004.
- [6] P. Selinger and B. Valiron. A lambda calculus for quantum computation with classical control. accepted for TLCA05, 2005.
- [7] A. van Tonder. A lambda calculus for quantum computation. quant-ph/0307150, 2003. to appear in SIAM Journal of Computing.
- [8] J. K. Vizzotto, T. Altenkirch, and A. Sabry. Structuring quantum effects: Superoperators as arrows. submitted for publication, 2004.
- [9] P. Zuliani. *Quantum Programming*. PhD thesis, Oxford University, 2001.