

Scheme Pearl: Quantum Continuations

VIKRAMAN CHOUDHURY, University of Glasgow, UK
BORISLAV AGAPIEV and AMR SABRY, Indiana University, USA

We advance the thesis that the simulation of quantum circuits is fundamentally about the efficient management of a large (potentially exponential) number of delimited continuations. The family of Scheme languages, with its efficient implementations of first-class continuations and with its imperative constructs, provides an elegant host for modeling and simulating quantum circuits.

1 INTRODUCTION

The non-intuitive properties and widely-heralded computational advantage of quantum computing are often associated with concepts like superposition, entanglement, and complementarity [Khrennikov 2021]. For the programming language aficionados, we instead present a different perspective: the power of quantum computing can be attributed to an extraordinarily efficient management of an exponential number of continuations.

We demonstrate that quantum computing is just one step away from well-established techniques with roots in the Scheme community. Specifically, quantum computing can be reduced to the use of continuations to implement backtracking search and non-determinism [Hinze 2000; Kiselyov et al. 2021, 2005; Wand and Vaillancourt 2004], augmented with a global mechanism to manage these continuations. Thus, in addition to its expository role, this paper opens up a new perspective where clever implementation techniques of continuations [Hieb et al. 1990], especially those relying on speculative evaluation, can enable more efficient classical simulations of quantum algorithms in particular cases.

Outline. We begin the next section with a review of the aspects of quantum computing that are necessary for our exposition. We focus on a minimal model of quantum circuits that is still rich enough to express general quantum computations up to arbitrary small errors, expressed inside Racket.

To make the connections to continuations apparent, in Section 3, we express the semantics of these quantum circuits using a slight generalization of search trees where the edges are decorated with probability amplitudes. These amplitudes may be positive or negative to express constructive or destructive interference patterns.

In Section 4, we give an interpreter that uses delimited continuations to evaluate the quantum circuits. The interpreter is parameterized by a function that collects and manages the generated continuations. A management strategy that invokes all the continuations and appends their results produces full information about the probabilistic behavior of the quantum computation. Although instructive, this information requires an exponential cost to generate and outputs information that is physically unobservable.

An adjustment to the previous strategy, given in Section 5, that produces physically observable results is to calculate interference patterns among the continuations allowing for some branches to cancel each other, and then sample from the remaining branches. The straightforward implementation of this latter strategy would still be exponential, but it might be possible in some cases to

Authors' addresses: Vikraman Choudhury, School of Computing Science, University of Glasgow, Glasgow, G12 8QQ, UK, vikraman.choudhury@glasgow.ac.uk; Borislav Agapie, borislav.agapie@gmail.com; Amr Sabry, Department of Computer Science, Indiana University, Bloomington, 47408, USA, sabry@indiana.edu.

2022. XXXX-XXXX/2022/8-ART \$15.00
<https://doi.org/10.1145/nnnnnnn.nnnnnn>

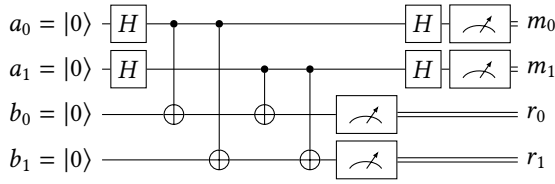


Fig. 1. Circuit for a small instance of Simon’s problem with $n = 2$ and $a = 3$. Given a 2-1 function $f : \mathbb{B}^2 \rightarrow \mathbb{B}^2$ with the property that there exists an a such that $f(x) = f(x \text{ XOR } a)$ for all x , the problem is to determine a .

discover more ingenious implementation strategies, which is an idea we leave for consideration in future work, as explained in the concluding section.

2 QUANTUM CIRCUITS

The circuit model is a well-established universal model of quantum computing [Nielsen and Chuang 2010]. We introduce the main idea using an extended example and then define a small circuit language whose semantics we model using continuations.

2.1 An Example: Simon’s Problem

As shown in Fig. 1, a typical quantum circuit consists of an initialization phase (on the left), followed by sequential and parallel compositions of elementary gates, and ending with measurements. Each horizontal line in the circuit corresponds to a quantum bit (qubit). The initial state, and the result of measuring each qubit, is a classical bit (a boolean value). As we explain below, between initialization and measurement, the state of the qubit can be in a *superposition* of boolean values, or even be *entangled* with other qubits.

The initial state of each of the four qubits in Fig. 1 is $|0\rangle$, which is essentially equivalent to a classical boolean 0. In preparation for possible entanglement among the four qubits, we collect the individual values into an aggregate $|0000\rangle$, where, by convention, the most significant position refers to the top wire. The first step in the execution is to apply the Hadamard (H) gate. The gate maps $|0\rangle$ to $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ and $|1\rangle$ to $\frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$. Both results of applying H represent states that are in equal superpositions of $|0\rangle$ and $|1\rangle$, with the sign expressing whether the values will lead to constructive or destructive interference. The aggregate state is now:

$$\frac{1}{2} ((|0\rangle + |1\rangle) (|0\rangle + |1\rangle) |00\rangle) = \frac{1}{2} (|0000\rangle + |0100\rangle + |1000\rangle + |1100\rangle)$$

The next four gates are all *controlled-not* gates which negate the target bit (marked with \oplus) if the control bit is true. The execution therefore continues with the following states:

$$\begin{aligned} \frac{1}{2} (|0000\rangle + |0100\rangle + |1000\rangle + |1100\rangle) &\rightarrow \frac{1}{2} (|0000\rangle + |0100\rangle + |1010\rangle + |1110\rangle) \\ &\rightarrow \frac{1}{2} (|0000\rangle + |0100\rangle + |1011\rangle + |1111\rangle) \\ &\rightarrow \frac{1}{2} (|0000\rangle + |0110\rangle + |1011\rangle + |1101\rangle) \\ &\rightarrow \frac{1}{2} (|0000\rangle + |0111\rangle + |1011\rangle + |1100\rangle) \end{aligned}$$

At this point, the two qubits in the least significant positions are measured. By inspection, we see that there is an equal probability of these two qubits being 00 or 11. Without any loss of generality,

we assume these qubits are measured as 11. Then, the entire state collapses to:

$$\frac{1}{\sqrt{2}} (|0111\rangle + |1011\rangle)$$

and ignoring the now irrelevant measured qubits, the top two qubits are in the state:

$$\frac{1}{\sqrt{2}} (|01\rangle + |10\rangle)$$

Applying the two H gates, the state before measurement is:

$$\frac{1}{2\sqrt{2}} (|00\rangle - |01\rangle + |10\rangle - |11\rangle + |00\rangle + |01\rangle - |10\rangle - |11\rangle)$$

The values with the same sign interfere constructively and the values with opposite signs interfere destructively leading to the simplified state:

$$\frac{1}{2\sqrt{2}} (2|00\rangle - 2|11\rangle) = \frac{1}{\sqrt{2}} (|00\rangle - |11\rangle)$$

A measurement of the two qubits is equally likely to produce 00 or 11 as the value of the hidden a . In the former case, the algorithm will have produced the vacuous identity that $f(x) = f(x)$ and the circuit needs to be re-executed. In the latter case, the algorithm successfully terminates with the correct decimal value 3 for a .

2.2 An Approximately Universal Circuit Language

For theoretical studies of quantum computing, it is particularly appealing to exploit two facts which justify focus on a tiny foundational core. First, by the principle of *deferred measurement* [Nielsen and Chuang 2010], it is always possible to defer all measurements to the last step of a quantum computation. This allows one to isolate what is called “pure quantum computing” which consists of sequential and parallel compositions of elementary gates not involving measurements. Second, if one is willing to tolerate arbitrary small errors, then all of pure quantum computing just needs two elementary gates [Aharonov 2003]: the Hadamard gate introduced above and the *controlled-controlled-not* (or Toffoli) gate which negates its target qubit when *both* control qubits are true.

We therefore restrict our attention to circuits built from sequences of the following two elementary operations:

```
(define-syntax-rule (CCX a b c) `(ccx ,a ,b ,c))
(define-syntax-rule (H a)      `(h ,a))
```

where a , b , and c are natural numbers giving the index of the qubit in question, or boolean constants $\#t$ or $\#f$. As is obvious, CCX denotes the controlled-controlled-not gate, and H denotes the Hadamard gate. For convenience, we further introduce the following two abbreviations:

```
(define-syntax-rule (X a)      (CCX #t #t a))
(define-syntax-rule (CX a b)   (CCX #t a b))
```

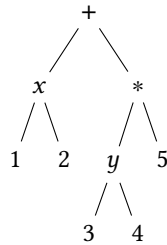
which introduce special cases of CCX where one or both control qubits are known to be true. Here X denotes the not gate, and CX denotes the controlled-not gate. For example, the pure part of the circuit in Fig. 1 can be written as the following expression:

```
(list
  (H 0))
```

```
(H 1)
(CX 0 2)
(CX 0 3)
(CX 1 2)
(CX 1 3)
(H 0)
(H 1))
```

3 CONTINUATIONS AND SEARCH TREES

Before considering quantum circuits, let us examine a toy example of an arithmetic expression. We illustrate the use of continuations to evaluate $x + (y * 5)$ where x can take the value 1 or 2, and y can take the value 3 or 4. We begin by visualizing the evaluation as a tree:



Given primitives `shift` and `reset` to manipulate continuations [Danvy and Filinski 1990], the four possible results of the expression can be collected in a list as follows. We set a continuation delimiter [Felleisen 1988] at the beginning of the evaluation (i.e., at the root of the tree), and then every time we encounter a node with a choice, we capture and *shift* two copies of the continuation to explore both branches. Each individual result is wrapped in a list and the results of invoking the continuations are appended:

```
(define (choose^ a b)
  (shift k (append (k a) (k b))))

(define expr
  (reset (let ([x (choose^ 1 2)]
               [y (choose^ 3 4)])
           (list (+ x (* 5 y))))))
```

The result of evaluating `expr` is a list.

```
> expr
'(16 21 17 22)
```

We now adapt this idea to simulate the quantum circuit in Fig. 1. We begin by visualizing the evaluation of the pure gates as a tree, similar to the arithmetic expression above. The resulting tree, shown in Fig. 2, starts with the initial state at the root. The semantics of X , CX , and CCX is straightforward, only making a local change to the state. At each occurrence of a Hadamard gate, evaluation splits into two branches. The crucial aspect of the tree is that the edges are labeled with probability amplitudes that are cumulatively multiplied as evaluation progresses. These amplitudes capture dependencies among the different paths that are highlighted with the colored connectors on the right. Each dashed red arrow connects two identical states reached by different

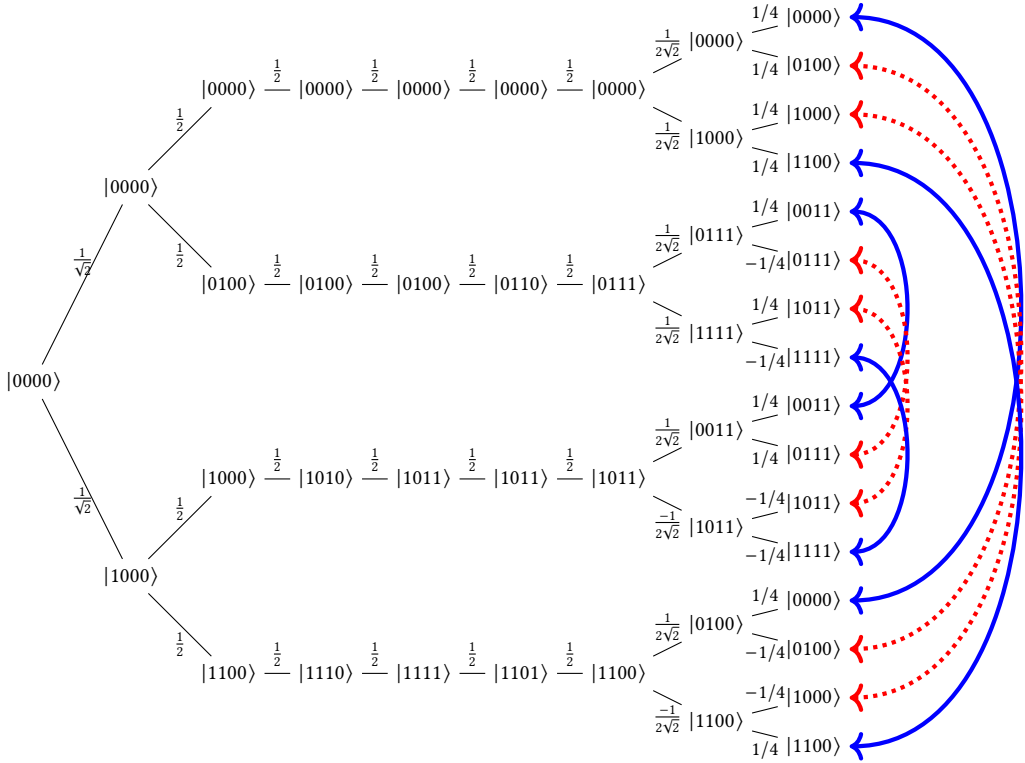


Fig. 2. Visualizing the evaluation of the pure gates in Fig. 1

execution paths with opposite probability amplitudes: they annihilate. The solid blue arrows similarly connect two identical states reached by different execution paths, but with probability amplitudes of the same sign: they reinforce each other. Thus, the final state of the execution is $\frac{1}{2}(|0000\rangle + |1100\rangle + |0011\rangle - |1111\rangle)$.

4 CONTINUATIONS AND QUANTUM CIRCUITS

We now codify the analysis of the examples in the previous section in an interpreter for the idealized quantum circuit language consisting of the ccx and H gates, described in Section 4.1. In Section 4.2, we use the interpreter to run three examples including the running example of Fig. 1.

4.1 Continuation-based Evaluator

The ccx gate is a classical gate whose semantics is just as simple as the semantics of $+$ or $*$. The Hadamard H gate is the one interesting gate that introduces choices. Specifically the application of H introduces a choice between 0 and ± 1 with the proper weights. The evaluator in Fig. 3 processes just one gate: it uses the delimited continuation primitive `shift` to generate continuations, and a parameterized way to *collect* and process these continuations.

The `evalg^` function takes two arguments: v is a tuple of the form (c, d, bs) , where d is the probability amplitude, bs is a bit-vector representing the state, and g is the quantum gate we are evaluating. The first two clauses define the semantics of the ccx gate by simply negating the target qubit when the control qubits are set. If no control qubits are set, the state is unchanged. The

```

(define (evalg^ v g)
  (match `(,v ,g)
    ;; ccx with both control bits set
    [^((,d ,bs) (ccx ,ctrl1 ,ctrl2 ,targ))
     #:when (and (is-set? bs ctrl1)
                  (is-set? bs ctrl2))
     `(,d ,(neg bs targ))]
    ;; ccx with control bit(s) unset
    [^((,d ,bs) (ccx ,ctrl1 ,ctrl2 ,targ))
     `(,d ,bs)]
    ;; hadamard with target bit set
    [^((,d ,bs) (h ,targ))
     #:when (is-set? bs targ)
     (collect^ `(,* hscale d) ,(neg bs targ))
              `(,* -1.0 hscale d) ,bs))]
    ;; hadamard with target bit unset
    [^((,d ,bs) (h ,targ))
     (collect^ `(,* hscale d) ,bs)
              `(,* hscale d) ,(neg bs targ)))]))

```

Fig. 3. Continuation-based evaluator for quantum gates

semantics of \mathbb{H} parallels the semantics of choice in the previous section: the rest of the computation is duplicated, executed once for each choice, and the choices are collected using a `collect^` function. If the target qubit is set, we negate the amplitude, multiplying it by -1.0 .

```

(define (evalc^ v c)
  (foldl (lambda (g st) (evalg^ st g)) v c))

(define (run-circ t circ st)
  (reset (match (evalc^ st circ)
               [^((,d ,bs)
                  (match t
                     [^list (list `(,bs ,d))]
                     ...))))))

```

Fig. 4. Continuation-based evaluator for complete circuits

As we represent circuits as lists of gates, to evaluate a circuit we simply fold over the circuit, as shown in Fig. 4. Finally, the entry point to run a circuit is to call `evalc^` with the circuit and the initial state, enclosing the entire evaluation using a `reset` combinator which acts as the delimiter for the captured continuation. We use an additional argument `t` in that entry point which is a tag that lets us choose which instance of `collect^` to use.



Fig. 5. Two small quantum circuits

4.2 Examples with List Collector

The interpreter is parameterized by an implementation of the function `collect^` which dictates how we combine and manage the results of the intermediate computations. We will explore two instances of `collect^`, one in this section, and one in the next. In both instances, we first use `shift` to capture the continuation `k`, and invoke it on the two arguments `x` and `y` giving us access to the underlying states `a` and `b`. In this section, we use the simplest strategy which collects all the states in a list:

```
(define (collect^ x y)
  (shift
   k
   (let ([a (k x)]
         [b (k y)])
     (cond
      ((and (list? a) (list? b)) (append a b))
      ...
    )
  )
```

As examples of running circuits using this implementation, we consider two small circuits, shown in Fig. 5. The first circuit is a `H` gate followed by a `CX` gate, operating on two qubits initialized to $|00\rangle$. We can represent and execute it as follows, using lists for state management.

```
(run-circ `list
  (list (H 0)
        (CX 0 1))
  `(1.0 ,(make-vector 2 #f)))
```

The execution happens as follows. The initial state is $|00\rangle$. At the `H` gate, the rest of the computation, i.e., the application of `CX` is duplicated and executed once with the state $\frac{1}{\sqrt{2}}|00\rangle$ and once with the state $\frac{1}{\sqrt{2}}|10\rangle$. The results are accumulated in a list producing the final result $[\frac{1}{\sqrt{2}}|00\rangle, \frac{1}{\sqrt{2}}|11\rangle]$.

The second circuit on the right in Fig. 5 is more interesting. It consists of an `H`, `X`, and `H` gate in series, on one qubit. We represent and execute it as follows.

```
(run-circ `list
  (list (H 0)
        (X 0)
        (H 0))
  `(1.0 ,(vector #f)))
```

The first occurrence of `H` duplicates the continuation consisting of the application of `x` followed by `H`. That continuation itself generates a nested choice, producing four distinct continuations in total. Tracing through the evaluation of the circuit, the final answer is the list $[\frac{1}{2}|0\rangle, -\frac{1}{2}|1\rangle, \frac{1}{2}|0\rangle, \frac{1}{2}|1\rangle]$.

However, this evaluation does not combine the probabilities, which is the only physically realizable observation. The state $|0\rangle$ has two positive probability amplitudes, which leads to constructive interference, and the state $|1\rangle$ has two opposite amplitudes, which cancel each other out, leading to

destructive interference. Our current evaluator does not observe this, but we can change that by switching to another instance of `collect^` in the next section.

At this point, our main thesis becomes evident: quantum computing efficiently (and somewhat mysteriously) manages these four distinct continuations using constructive and destructive interference ensuring that the only possible observable result is $|0\rangle$.

Before concluding this section, we confirm that evaluating the pure gates of the circuit in Fig. 1 indeed produces the leaves of the tree shown in Fig. 2:

```
> (run-simon-list)
(+0.25|0000>)
(+0.25|0100>)
(+0.25|1000>)
(+0.25|1100>)
(+0.25|0011>)
(-0.25|0111>)
(+0.25|1011>)
(-0.25|1111>)
(+0.25|0011>)
(+0.25|0111>)
(-0.25|1011>)
(-0.25|1111>)
(+0.25|0000>)
(-0.25|0100>)
(-0.25|1000>)
(+0.25|1100>)
```

5 CONTINUATIONS AND INTERFERENCE OF PROBABILITY AMPLITUDES

In the previous section, we evaluated the pure part of a quantum circuit, producing a list of states in the end. However, this evaluation does not combine the probabilities, that is, the states at the end are in superposition, but the probability amplitudes are not combined. By using lists to manage the quantum states, we simply explored all branches of the search tree.

For example, consider the evaluation of the second circuit in Fig. 5. The final state is given by the list $[\frac{1}{2}|0\rangle, -\frac{1}{2}|1\rangle, \frac{1}{2}|0\rangle, \frac{1}{2}|1\rangle]$. Here, the state $|0\rangle$ has two positive probability amplitudes, which leads to constructive interference, and the state $|1\rangle$ has two amplitudes of opposing signs, which cancel each other out, leading to destructive interference.

To allow our evaluator to combine the probabilities, we will use hash tables to represent these states along with their probability amplitudes. The hash table maps bit-vectors to probability amplitudes. All we have to do is change the implementation of our `collect^` function.

```
(define (collect^ x y)
  (shift
   k
   (let ([a (k x)]
         [b (k y)])
     (cond
      ...
      ((and (hash? a) (hash? b)) (hash-union a b #:combine +))))))
```


To combine two hash tables, we take their union. If a bit-vector is repeated, we simply add the probability amplitudes, using the `+` function to combine the two values that are mapped to by the same key. This means that two branches with the same state but opposite amplitudes will cancel each other out, and states with amplitudes of the same sign will reinforce each other.

The runner initializes the hash table with the initial bit-vector and amplitude. To run the example with the hash implementation, we simply change the tag passed to it as an argument.

```
(define (run-circ t circ st)
  (reset (match (evalc^ st circ)
               [ `(,d ,bs)
                 (match t
                    ...
                    [ `hash (hash bs d)]))]))

(run-circ `hash
  (list (H 0)
        (X 0)
        (H 0))
  `(1.0 ,(vector #f)))
```

The final observation now is $[1 | 0]$, as expected.

We also evaluate the circuit in Fig. 1 using the hash table implementation. Following the evaluation tree in Fig. 2, the states connected by the dashed red arrows get annihilated, and the ones connected by the solid blue arrows get reinforced.

```
> (run-simon-hash)
(+0.50|0000)
(+0.50|1100)
(+0.50|0011)
(-0.50|1111)
```

6 REFLECTIONS & CONCLUSIONS

We represented quantum states as bit-vectors with their probability amplitudes, using lists and hash tables to store and combine them. Another way of representing them is to directly work with probability distributions, which can be approximately implemented using continuations. We quickly discuss a short implementation of classical probability distributions using continuations, where events are encoded as functions ($\rightarrow a p$), and probability distributions are encoded as higher-order functionals ($\rightarrow (-\rightarrow a p) p$).

We can write a `choose-p` function that takes two probability distributions `k1` and `k2`, with a sampling bias `p`, and returns a new distribution by calculating the convex combination of probabilities of events in the two distributions.

```
(define (choose-p p k1 k2)
  (lambda (f)
    (let ([p1 (* p (k1 f))]
          [p2 (* (- 1.0 p) (k2 f))])
      (+ p1 p2))))
```

To produce a constant probability distribution, given a value, for any event, we apply the value to the event. The expectation of an event is computed by simply applying the event to the underlying continuation.

```
(define (const-p a)
  (lambda (f) (f a)))

(define (expectation f k)
  (k f))
```

Using these combinators, we can write yet another implementation of the evaluator, where we use the choose-p operation (with no bias) to collect two states. For example, we can execute the example in Fig. 1.

```
> (run-simon-prob)
(+0.25|0000⟩)
(+0.25|0010⟩)
(+0.25|0001⟩)
(+0.25|0011⟩)
(+0.25|1100⟩)
(-0.25|1110⟩)
(+0.25|1101⟩)
(-0.25|1111⟩)
(+0.25|1100⟩)
(+0.25|1110⟩)
(-0.25|1101⟩)
(-0.25|1111⟩)
(+0.25|0000⟩)
(-0.25|0010⟩)
(-0.25|0001⟩)
(+0.25|0011⟩)
```

However, it is obvious that these probability amplitudes do not have cancellation, and we do not get the desired observation with interference. This encoding of probability distributions only works for classical probabilities and not quantum probabilities – to move to quantum computing, we need to move from classical probability theory to generalized probabilistic theory (GPT). We hope to investigate these ideas connecting them to this work on continuations in the future.

By using hash tables to represent probability distributions, we are sampling from the distribution at every step of the computation when we evaluate a H gate. On the other hand, by using continuations to represent probability distributions, we're avoiding sampling until the very end, when a measurement is made. Combining these techniques could lead to techniques for implementing speculative execution in quantum computers.

We have presented a continuation-based analysis of quantum computing which shows that it efficiently (and somewhat mysteriously) manages an exponential number of continuations. It would be however incorrect to conclude that quantum computing can manage arbitrary such collections of continuations as this would imply it can solve NP-complete problems efficiently which is not believed to be true. Discovering instances in which collect can be implemented without necessarily invoking each continuation might provide insights on the elusive power of quantum computing.

From a different perspective, it is known that there is a duality between values and continuations [Curien and Herbelin 2000; Filinski 1989] with values flowing from producers to consumers and continuations flowing from consumers to producers. This suggests a re-interpretation of our continuation-based semantics as a distributed system with positive offers by producers and negative counter-offers by consumers, which would be worth investigating in detail.

REFERENCES

- D. Aharonov. 2003. A simple proof that Toffoli and Hadamard are quantum universal. *arXiv:quant-ph/0301040* (2003).
- Pierre-Louis Curien and Hugo Herbelin. 2000. The Duality of Computation. In *Proceedings of the Fifth ACM SIGPLAN International Conference on Functional Programming (ICFP '00)*. Association for Computing Machinery, New York, NY, USA, 233–243. <https://doi.org/10.1145/351240.351262>
- Olivier Danvy and Andrzej Filinski. 1990. Abstracting Control. In *Proceedings of the 1990 ACM Conference on LISP and Functional Programming (Nice, France) (LFP '90)*. Association for Computing Machinery, New York, NY, USA, 151–160. <https://doi.org/10.1145/91556.91622>
- Mattias Felleisen. 1988. The Theory and Practice of First-Class Prompts. In *Proceedings of the 15th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (San Diego, California, USA) (POPL '88)*. Association for Computing Machinery, New York, NY, USA, 180–190. <https://doi.org/10.1145/73560.73576>
- Andrzej Filinski. 1989. Declarative Continuations and Categorical Duality.
- R. Hieb, R. Kent Dybvig, and Carl Bruggeman. 1990. Representing Control in the Presence of First-Class Continuations. In *Proceedings of the ACM SIGPLAN 1990 Conference on Programming Language Design and Implementation (White Plains, New York, USA) (PLDI '90)*. Association for Computing Machinery, New York, NY, USA, 66–77. <https://doi.org/10.1145/93542.93554>
- Ralf Hinze. 2000. Deriving Backtracking Monad Transformers. In *Proceedings of the Fifth ACM SIGPLAN International Conference on Functional Programming (ICFP '00)*. Association for Computing Machinery, New York, NY, USA, 186–197. <https://doi.org/10.1145/351240.351258>
- A. Khrennikov. 2021. Roots of quantum computing supremacy: superposition, entanglement, or complementarity? *Eur. Phys. J. Spec. Top.* 230 (2021), 1053–1057.
- Oleg Kiselyov, Shin-Cheng Mu, and Amr Sabry. 2021. Not by equations alone: Reasoning with extensible effects. *Journal of Functional Programming* 31 (2021), e2. <https://doi.org/10.1017/S0956796820000271>
- Oleg Kiselyov, Chung-chieh Shan, Daniel P. Friedman, and Amr Sabry. 2005. Backtracking, Interleaving, and Terminating Monad Transformers: (Functional Pearl). In *Proceedings of the Tenth ACM SIGPLAN International Conference on Functional Programming (Tallinn, Estonia) (ICFP '05)*. Association for Computing Machinery, New York, NY, USA, 192–203. <https://doi.org/10.1145/1086365.1086390>
- 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>
- Mitchell Wand and Dale Vaillancourt. 2004. Relating Models of Backtracking. In *Proceedings of the Ninth ACM SIGPLAN International Conference on Functional Programming (Snow Bird, UT, USA) (ICFP '04)*. Association for Computing Machinery, New York, NY, USA, 54–65. <https://doi.org/10.1145/1016850.1016861>

SUPPLEMENTARY MATERIAL

We include the complete implementation of the evaluator and the circuit examples discussed in Sections 4 and 5.

```

1  #lang racket
2
3  (require racket/control racket/hash racket/vector)
4
5  ;; vector and hash helper operations
6
7  (define (upd v i f)
8    (let ([w (vector-copy v)])
9      (vector-set! w i (f (vector-ref w i))))
10   w))

```

```

11
12 (define (neg v i)
13   (upd v i not))
14
15 (define hscale
16   (/ 1.0 (sqrt 2.0)))
17
18 (define (is-set? v i)
19   (cond ((nonnegative-integer? i) (vector-ref v i))
20         ((boolean? i) i)
21         (else (error "is-set?: invalid index"))))
22
23 ;; syntax for gates
24
25 (define-syntax-rule (X a) (CCX #t #t a))
26 (define-syntax-rule (CX a b) (CCX #t a b))
27 (define-syntax-rule (CCX a b c) `(ccx ,a ,b ,c))
28 (define-syntax-rule (H a) `(h ,a))
29
30 ;; collect^ operation
31
32 (define (collect^ x y)
33   (shift
34     k
35     (let ([a (k x)]
36           [b (k y)])
37       (cond
38         ((and (list? a) (list? b)) (append a b))
39         ((and (hash? a) (hash? b)) (hash-union a b #:combine +))
40         (else (error "collect^: unknown state"))))))
41
42 ;; gate and circuit evaluator
43
44 (define (evalg^ v g)
45   (match `(,v ,g)
46     [ `(,d ,bs) (ccx ,ctrl1 ,ctrl2 ,targ)
47       #:when (and (is-set? bs ctrl1)
48                   (is-set? bs ctrl2))
49       `(,d ,(neg bs targ))]
50     [ `(,d ,bs) (ccx ,ctrl1 ,ctrl2 ,targ)
51       `(,d ,bs)]
52     [ `(,d ,bs) (h ,targ)
53       #:when (is-set? bs targ)
54       (collect^ `(,* hscale d) ,(neg bs targ))
55                 `(,* -1.0 hscale d) ,bs))]
56     [ `(,d ,bs) (h ,targ)
57       (collect^ `(,* hscale d) ,bs)
58                 `(,* hscale d) ,(neg bs targ))])

```

```

59     [_ (error "evalg^: invalid arguments")]))
60
61 (define (evalc^ v c)
62   (foldl (lambda (g st) (evalg^ st g)) v c))
63
64 ;; pretty printer
65
66 (define (pretty-vec bs)
67   (string-append
68     "|"
69     (foldl (lambda (b s) (string-append (if b "1" "0") s))
70           ">")
71           (vector->list bs))))
72
73 (define (pretty-prob d)
74   (~r d #:sign '+' #:precision '(= 2)))
75
76 (define (pretty-state st)
77   (cond
78     [(list? st)
79      (for-each
80        (match-lambda*
81          [^((,bs ,d)) #:when (< (abs d) 0.01) (void)]
82          [^((,bs ,d)) (printf "~a~a~n" (pretty-prob d) (pretty-vec bs))]
83          st)]
84     [(hash? st)
85      (hash-for-each
86        st
87        (match-lambda*
88          [^((,bs ,d) #:when (< (abs d) 0.01) (void)]
89          [^((,bs ,d) (printf "~a~a~n" (pretty-prob d) (pretty-vec bs)))]))]
90     [else (error "pretty-state: unknown state")])])
91
92 ;; runner
93
94 (define (run-circ t circ st)
95   (reset (match (evalc^ st circ)
96             [^((,d ,bs)
97               (match t
98                 [list (list ^((,bs ,d))]
99                 [hash (hash bs d)]
100                 [_ (error "run-circ: invalid tag")]))]))))
101
102 ;; H-CX example
103
104 (define (run-hcx-list)
105   (pretty-state (run-circ `list
106                        (list (H 0)

```

```

107             (CX 0 1))
108         `(1.0 ,(make-vector 2 #f))))))
109
110 (define (run-hcx-hash)
111   (pretty-state (run-circ `hash
112                   (list (H 0)
113                         (CX 0 1))
114                   `(1.0 ,(make-vector 2 #f))))))
115
116 ;; H-X-H example
117
118 (define (run-hxh-list)
119   (pretty-state (run-circ `list
120                   (list (H 0)
121                         (X 0)
122                         (H 0))
123                   `(1.0 ,(vector #f))))))
124
125 (define (run-hxh-hash)
126   (pretty-state (run-circ `hash
127                   (list (H 0)
128                         (X 0)
129                         (H 0))
130                   `(1.0 ,(vector #f))))))
131
132 ;; Simon example
133
134 (define (run-simon-list)
135   (pretty-state (run-circ `list
136                   (list (H 0)
137                         (H 1)
138                         (CX 0 2)
139                         (CX 0 3)
140                         (CX 1 2)
141                         (CX 1 3)
142                         (H 0)
143                         (H 1))
144                   `(1.0 ,(make-vector 4 #f))))))
145
146 (define (run-simon-hash)
147   (pretty-state (run-circ `hash
148                   (list (H 0)
149                         (H 1)
150                         (CX 0 2)
151                         (CX 0 3)
152                         (CX 1 2)
153                         (CX 1 3)
154                         (H 0))

```

155

(H 1))

156

`(1.0 ,(make-vector 4 #f))))))