

The Many-Worlds Interpretation and Quantum Computation

Armond Duwell^{†‡}

David Deutsch and others have suggested that the Many-Worlds Interpretation of quantum mechanics is the *only* interpretation capable of explaining the special efficiency quantum computers seem to enjoy over classical ones. I argue that this view is not tenable. Using a toy algorithm I show that the Many-Worlds Interpretation must crucially use the ontological status of the universal state vector to explain quantum computational efficiency, as opposed to the particular ontology of the MWI, that is, the computational histories of worlds. As such, any other interpretation that treats the state vector as representing real ontological features of a system can explain quantum speedup too.

1. Introduction. Why are quantum computers faster than classical computers? Deutsch (1997) and others (e.g., Jozsa 2000) have suggested that the reason quantum computers are more efficient than classical computers is that they can compute multiple values of a function in a single computational step. Call this the *quantum parallelism thesis* (QPT).¹ Deutsch (1997) has also suggested that the only interpretation of quantum mechanics that can support the QPT is the many worlds interpretation (MWI). The quantum parallelism process (QPP) is often pointed to in order to justify these claims.

Consider a system of $n + m$ qubits. Let n qubits be the input register, system i , for our quantum computer and let m qubits be the output register, system o . Each qubit can represent a single binary number, 0 or 1, where 0 corresponds to z-spin up, and 1 z-spin down. The n qubit input register can represent any integer from 0 to $2^n - 1$. Further, the input register can be put into a superposition of all quantum states representing the integers

[†]To contact the author, please write to: Department of Philosophy, Liberal Arts 101, University of Montana, Missoula, MT 59812; e-mail: armond.duwell@umontana.edu.

[‡]Thanks to Soazig Le Bihan for her critical comments on this paper.

1. This seems like a reasonable account of quantum computational efficiency if classical computers are not capable of such behavior.

Philosophy of Science, 74 (December 2007) pp. 1007–1018. 0031-8248/2007/7405-0038\$10.00
Copyright 2007 by the Philosophy of Science Association. All rights reserved.

from 0 to $2^n - 1$. For convenience, adopt the following notation: $|x\rangle$ will stand for the n qubit state representing the number x in binary. Suppose that $n = 2$. The number three would then be represented by the state $|3\rangle = |1\rangle|1\rangle$. With this notation, an n qubit system that is a superposition of all numbers from 0 to $2^n - 1$ is given by

$$\sum_{x=0}^{2^n-1} |x\rangle,$$

where normalization coefficients are suppressed throughout this paper.

Consider a unitary gate, U_f , that performs the following evolution:

$$|x\rangle|y\rangle \rightarrow |x\rangle|f(x) \oplus y\rangle, \quad (1)$$

where $f: \{0, 1\}^n \rightarrow \{0, 1\}^m$ and \oplus is addition mod 2. If the input register is put into a superposition of all points in the domain of the function, the system evolves as

$$\sum_{x=0}^{2^n-1} |x\rangle|y\rangle \xrightarrow{U_f} \sum_{x=0}^{2^n-1} |x\rangle|f(x) \oplus y\rangle. \quad (2)$$

(2) is the quantum parallelism process.

An argument needs to be made that the QPP can be appealed to show that the QPT is true. For the purposes of this paper, a functionalist criterion for computation is adopted (COMP FUNC): Given some fixed interpretation of input and output states of a computer, the functional behavior between input and output states determines which function a computer computes, and when a computation takes place (see Duwell [2004], Section 7, for details regarding this functional view of computation). If the additional assumption is made that the different pure states of quantum systems are taken to be correlated to different ontological features of a system, an assumption which will be called benign realism (BR), there is no doubt the computation was successful.² The computer delivered the functional behavior requisite for the evaluation of a function at multiple points. So, COMP FUNC + BR + QPP \rightarrow QPT is true.

Given that COMP FUNC + BR + QPP \rightarrow QPT is true, it appears that any interpretation of quantum mechanics in which BR is true can support the QPT. So, there must be something special about how the MWI secures the truth of QPT that differentiates it from other interpretations.

Consider the story a MWI advocate might tell about the QPP. Each term in the initial superposition in (2) corresponds to a different world. Those terms are the state vectors describing the quantum mechanics of

2. BR does not entail a commitment to definite values or any particular kind of ontological feature or entity.

each world. The initial states of these worlds correspond to points at which values of a function are to be computed. In each world, the quantum system representing the point of the function to be computed is then submitted to a quantum gate that computes each value of the function. After the quantum system passes through the gate, in each world, the value of the function will have been computed. Put another way, the functional behavior of each world secures the computation of a value of the function. All values of the function are computed in virtue of all individual values being computed in different worlds. So, the truth of the QPT supervenes on the computational histories of individual worlds, not only on the computational history of the whole universe. This view seems especially compelling because in each world, the computer computes the value of a function at a particular point just as a classical circuit computer would.

Distinguish two ways in which multiple values of a function might be computed:

Local computation: The functional behavior of the computational device secures the computation of each value individually by satisfying COMP FUNC for each value of the function.

Global computation: The functional behavior of the computational device secures the computation of all values by satisfying COMP FUNC, but not for each value.

In order to claim a privileged status regarding the account of quantum computational efficiency, proponents of the MWI would have to claim that only “local” accounts of the truth of the QPT are satisfactory, “global” accounts are not. So, interpretations of quantum theories can account for quantum computational efficiency only if they are such that the ontological features of the computational device evolve in such a way that each of the values of the function are computed individually. In terms of the MWI, the computational histories of the worlds are what accounts for the truth of QPT, and the computational history of the entire system only derivatively.

Regardless of whether there is a local account of the truth of the QPT or not, Steane (2003) argues that an instance of the QPP does not constitute the calculation of multiple values of a function, that is, the QPT is false. Steane points out that for a system that undergoes the quantum parallelism process, there is no way to access the points of a function that were purportedly evaluated. If access to the values is necessary to qualify the quantum parallelism process as a computation of multiple values in a single step, then there is no evidence that the QPT is true.

Steane’s challenge to the truth of the QPT can be met. One can simply

make an appeal to the fact that $\text{COMP FUNC} + \text{BR} + \text{QPP} \rightarrow \text{QPT}$ is true. Ontologically the computer has done its job, regardless of whether the computational results are accessible. Though Steane's objection does not successfully challenge the truth of the QPT, it does challenge the sufficiency of any account of quantum computational efficiency that simply points to the truth of the QPT. The mere fact that a quantum computer can compute multiple values in a single step shows nothing when at most one of those results is accessible. One must indicate how the computation of those inaccessible values can be used to perform computational work efficiently.

The MWI account of quantum computational efficiency sketched above needs to be extended to give an account of how one utilizes the results of the QPT to do useful computational work. It will be argued that in order to do this, an advocate of the MWI must go global. The advocate must appeal to computational history of the universe, not simply the computational histories of each individual world. In order to account for quantum computational efficiency, the advocate of the MWI must appeal to the ontological features of the computational device that are crucial to account for computational efficiency, but irrelevant to the truth of the QPT. In doing so, advocates of the MWI must utilize the type of ontological features of the computational device that they reject in other accounts of the truth of the QPT. They must appeal to BR and the evolution of the universal state vector throughout the computational process to account for quantum computational efficiency. Put another way, what advocates of the MWI reject in other accounts of quantum computational efficiency in order to claim superiority for their account is exactly what they need to account for quantum computational efficiency. Hence, it will be argued that MWI accounts of quantum computational efficiency are no better off than others in which BR is true. Note that it will not be argued that the MWI account is deficient or that other accounts are better.

Section 2 describes an algorithm for determining if a binary function is constant or balanced. This simple algorithm makes it explicit how the results of the QPP can be utilized to do useful computational work. Section 3 shows that the success of the algorithm cannot supervene on the individual computational histories of each world. Indeed, it must supervene on the computation history of the entire universe, and hence make appeal to the ontological status of the universal state vector.

2. A Problematic Algorithm. In this section a toy algorithm is examined to decide whether a function is constant or balanced. A function is balanced if exactly half of the values in the range of the function are 0 and the other half are 1. A constant function takes a single value for every point in its domain. This algorithm seems to compute multiple values of

a function in a single computational step and is a good test case for the claim that the MWI interpretation is a superior explanation of quantum efficiency.

Below is an example, taken from Mermin (2000), of a quantum algorithm used to determine if a function is constant or balanced. Consider a function $f: \{0, 1\} \rightarrow \{0, 1\}$. Now, let the initial input system and output system be in the states $(|0\rangle - |1\rangle)$. The initial state of the system is

$$|\Psi_{ini}\rangle = |0\rangle_i|0\rangle_o - |1\rangle_i|0\rangle_o - |0\rangle_i|1\rangle_o + |1\rangle_i|1\rangle_o. \quad (3)$$

Allow these qubits to pass through quantum gate U_f as described by (1). The state, ignoring normalization factors, becomes

$$|\Psi_f\rangle = |0\rangle_i|f(0)\rangle_o - |1\rangle_i|f(1)\rangle_o - |0\rangle_i|\bar{f}(0)\rangle_o + |1\rangle_i|\bar{f}(1)\rangle_o, \quad (4)$$

where $\bar{f} = 1 - f$. If the function is constant, $f(0) = f(1)$, then $\bar{f}(0) = \bar{f}(1)$, and the state is

$$(|0\rangle_i - |1\rangle_i)(|f(0)\rangle_o - |\bar{f}(0)\rangle_o), \quad (5)$$

and if the function is balanced, $f(0) = \bar{f}(1)$, and $f(1) = \bar{f}(0)$, the state is

$$(|0\rangle_i + |1\rangle_i)(|f(0)\rangle_o - |\bar{f}(0)\rangle_o). \quad (6)$$

It is plain to see from (5) and (6) that a measurement on the input state will determine whether the function is constant or balanced. The states $(|0\rangle_i - |1\rangle_i)$ and $(|0\rangle_i + |1\rangle_i)$ are eigenstates of x-spin down and up, respectively. Furthermore, the two possible final states for i are orthogonal and can always be distinguished by a single measurement.

3. MWI Account of the Toy Algorithm. The challenge for the MWI in accounting for the success of this algorithm is twofold. The MWI must secure the truth of the QPT, but also secure the evaluation of the function as constant or balanced, as that is what is required to account for the efficiency of the algorithm. Recall that in order to differentiate between the MWI account of quantum computational efficiency and other interpretations' accounts, one must demand that the QPT be secured locally. An advocate of the MWI must produce an account of quantum computational efficiency that secures the evaluation of the function as constant or balanced locally as well. That is, the sure-fire disposition to indicate if a function was constant or balanced must supervene on the computational histories of each world individually. In terms more specific to the situation, the evaluation cannot depend on phase relations between the worlds that are not relevant to the computations performed in each in-

dividual world.³ If the MWI does not have the resources to do this, then the claimed superiority vanishes, as the resources required to account for quantum computational efficiency will be exactly what advocates of the MWI find faulty with other accounts. It will be argued that this superiority cannot be maintained.

There appear to be four exhaustive options for an advocate of the MWI to account for the truth of the QPT and the evaluation of the function as constant or balanced (EVAL), using the ontology of worlds: Option 1: the history of quantum states in each world account for the truth of the QPT and EVAL; Option 2: the history of definite values associated with the worlds account for the truth of the QPT and EVAL; Option 3: the history of quantum states associated with each world accounts for the truth of the QPT, while the history of definite values associated with the worlds account for the EVAL; and Option 4: the history of definite values associated with each world accounts for the truth of the QPT, while the EVAL is secured by the quantum states associated with each world. Options 2 and 4 assume that measurement devices are responsive to the possessed definite values in worlds as opposed to the quantum state of a world. This assumption is not as bad as it seems, as the possessed values attributed to worlds correspond to eigenstates of the global state assigned to the universe.⁴

It is rather quick and easy to dismiss Options 1 and 2. For simplicity, adopt the following notation. To describe a world at a given time, two features will be considered: the quantum state of the world and the definite values that it possesses. The state that a system possesses will occur in kets, and definite values will be without kets. A description of a world will be given by a pair (\cdot, \cdot) , where the first slot describes the feature of the world relevant for the truth of the QPT, while the second slot describes the feature of the world relevant for the EVAL. So, $(|00\rangle, C)$ would be a world with a quantum state given by $|00\rangle$ which is relevant for the truth of the QPT, but which possesses a definite value that underwrites the EVAL as Constant.

Consider Option 1. Assume that the computational basis is the z-basis, and worlds correspond to terms in the state vector written in that basis,

3. I take phase relations between worlds to be relational properties between worlds that do not supervene on the individual properties of worlds, hence they are global, but irreducible properties. The justification comes from the fact that relative phases between worlds can be shifted by the simple multiplication of a global phase factor that has no observational consequences. The price of rejecting this interpretation is an unnecessarily inflated ontology with absolutely no observational consequences.

4. This way of structuring the problem is inspired by Clifton (1996). Some of the arguments against these options parallel Clifton's (1996) arguments against the MWI generally.

and those terms give the state of the world they individuate.⁵ Hence, the descriptions of the worlds are initially given by $(|00\rangle, \cdot)$, $(|10\rangle, \cdot)$, $(|01\rangle, \cdot)$, and $(|11\rangle, \cdot)$. The right slots in the description of the worlds are left blank because there is nothing at the initial start of the computation that can or ought to secure the EVAL. The states for each world must evolve to $(|0f(0)\rangle, |E_1\rangle)$, $(|1f(1)\rangle, |E_2\rangle)$, $(|0\bar{f}(0)\rangle, |E_3\rangle)$, and $(|1\bar{f}(1)\rangle, |E_4\rangle)$, in order for quantum states to secure the truth of the QPT where $|E_j\rangle$ is the quantum state in world j which is supposed to secure the EVAL. Now, clearly there can be only one final state of a world after the computation. So, the states in each of the slots that compose the description of a world must be equal—for example, $|E_1\rangle$ must be equal to $|0f(0)\rangle$, and so on. Now, recall that a measurement of spin on the input state in the x-direction reveals whether f is constant or balanced. Hence in order to secure the EVAL, $|E_j\rangle$, whatever it is, would be in an eigenstate of x-spin of the input qubit. Of course, the states assigned to the worlds that can secure the truth of the QPT are z-spin eigenstates, which are incompatible with x-spin eigenstates. The only obvious move to make within the constraints of Option 1 is to suppose that the final state of the worlds is given by the final global state $|\Psi_f\rangle$, (4). That state will secure the EVAL. The price of this move is to give up the local account of the truth of the QPT. The initial states are all different, but all evolve to the same final state:

$$(|00\rangle, \cdot) \rightarrow (|\Psi_f\rangle, |\Psi_f\rangle)$$

$$(|10\rangle, \cdot) \rightarrow (|\Psi_f\rangle, |\Psi_f\rangle)$$

$$(|01\rangle, \cdot) \rightarrow (|\Psi_f\rangle, |\Psi_f\rangle)$$

$$(|11\rangle, \cdot) \rightarrow (|\Psi_f\rangle, |\Psi_f\rangle).$$

Thus, the functional behavior of the worlds cannot underwrite the truth of the QPT. So, Option 1 is unacceptable.

Consider Option 2. Local definite values secure the truth of the QPT and the EVAL. For local definite values to secure the truth of the QPT, there must be definite values for z-spin corresponding to the terms that individuate worlds in the global state vector. Moreover, the change in these

5. Clifton's (1996) paper indicates the problems associated with this last assumption, but if the truth of the QPT is to be secured with local quantum states, there appears to be no other reasonable quantum state assignments to those worlds given the story advocates of the MWI want to give about computations of individual values in those worlds. Similar justifications can be given for the definite value assignments that underwrite the truth of the QPT in other options.

values in each world must parallel the evolution of the terms in the state vector through the computational process. So, the evolution must be

$$(00, \cdot) \rightarrow (0f(0), E_1)$$

$$(10, \cdot) \rightarrow (1f(1), E_2)$$

$$(01, \cdot) \rightarrow (0\bar{f}(0), E_3)$$

$$(11, \cdot) \rightarrow (1\bar{f}(1), E_4),$$

where E_j is the definite value that secures the EVAL in world j . This is required to get the appropriate functional behavior of definite values to secure the truth of the QPT. According to Option 2, definite values must also secure the EVAL. Definite x-spin values are required in each world to secure the EVAL, but these are incompatible with the definite z-spin values that secure the truth of the QPT. Now, one might attempt to resolve this problem by suggesting that z-spin values need not be the values that underwrite the truth of the QPT. But in order to secure the EVAL, all worlds must end up with the same definite value for x-spin at the end of the computational process. Hence, similar to the objection raised in the discussion of Option 1, there is no possible way for the right functional behavior to be instituted in each world using definite values that secure the truth of the QPT and the EVAL at the same time. More sophisticated attempts to secure the truth of the QPT and the EVAL are considered below.

Option 3 is that the quantum state of each world secures the truth of the QPT while a definite value secures the EVAL. To secure the EVAL, add a definite value of x-spin to the final state of each world. Now, there is no problem with the quantum state of a world being a z-eigenstate, but possessing a definite value for x-spin, so the objections that were raised for Options 1 and 2 just don't apply. So, it seems that the evolution could be

$$(|00\rangle, \cdot) \rightarrow (|0f(0)\rangle, E)$$

$$(|10\rangle, \cdot) \rightarrow (|1f(1)\rangle, E)$$

$$(|01\rangle, \cdot) \rightarrow (|0\bar{f}(0)\rangle, E)$$

$$(|11\rangle, \cdot) \rightarrow (|1\bar{f}(1)\rangle, E),$$

where E is the definite value of x -spin that secures the EVAL. We can assume that it is the same for all worlds because of the independence of states and definite values. (Note that E is a definite value corresponding to the global state.)

If this path is taken, one has to say what made the worlds take on the particular definite value of x -spin that they did. The only plausible option is to suggest that the possessed value that secures the evaluation supervenes on the computational history of all of the individual worlds. This does not violate the constraint on accounts of quantum computational efficiency because the definite value supervenes on the computational history of the individual worlds, and not the computational history of the entire universe. Recall that the computational history of the entire universe involves phase relations between worlds, which cannot be part of an acceptable account of quantum computational efficiency if the MWI is to retain its superiority.

There is a problem with this suggestion. An argument can be made that the definite values that the worlds acquire cannot supervene on the computational histories of the individual worlds alone. Consider a different initial state:

$$|\Phi_{ini}\rangle = |0\rangle_i|0\rangle_o + |1\rangle_i|0\rangle_o + |0\rangle_i|1\rangle_o + |1\rangle_i|1\rangle_o. \quad (7)$$

The only differences between $|\Psi_{ini}\rangle$ and $|\Phi_{ini}\rangle$ are the phase relations between worlds, which should be irrelevant to an account of quantum computational efficiency if the MWI is to retain its superiority.

So, the evolution of the worlds if the initial state is $|\Phi_{ini}\rangle$ should be

$$(|00\rangle, \cdot) \rightarrow (|0f(0)\rangle, E)$$

$$(|10\rangle, \cdot) \rightarrow (|1f(1)\rangle, E)$$

$$(|01\rangle, \cdot) \rightarrow (|0\bar{f}(0)\rangle, E)$$

$$(|11\rangle, \cdot) \rightarrow (|1\bar{f}(1)\rangle, E).$$

There should be a definite value that is taken on in the course of the computation, that underwrites the EVAL, but that isn't what happens. If $|\Phi_{ini}\rangle$ is the initial state, no matter what function is computed, the final state is always equal to $|\Phi_{ini}\rangle$. $|\Phi_{ini}\rangle$ has the disposition to indicate that the function is balanced (see (6)), even if the function computed was constant. So, the definite value that each world is to acquire either does not exist, does not dispose measurement devices to indicate whether the function is constant or balanced, or supervenes on more than the com-

putational histories of the individual worlds. Either way, Option 3 is unacceptable.

Option 4 is that the truth of the QPT is underwritten by definite values, but the EVAL is underwritten by a quantum state. As with Option 3, there is no obvious incompatibility between the ontological features that secure the truth of the QPT and the EVAL. The evolution in this case would be

$$(00, \cdot) \rightarrow (0f(0), |E\rangle)$$

$$(10, \cdot) \rightarrow (1f(1), |E\rangle)$$

$$(01, \cdot) \rightarrow (0\bar{f}(0), |E\rangle)$$

$$(11, \cdot) \rightarrow (1\bar{f}(1), |E\rangle),$$

where $|E\rangle$ is the state which underwrites the EVAL. The obvious candidate for $|E\rangle$ is $|\Psi_f\rangle$. Much like the case in Option 3, something needs to be said about where this state comes from. It is tantalizing to simply make an appeal to the initial global state and the unitary evolution that instantiates the computation, but this is exactly what cannot be done. This would be an explicit appeal to the ontological status of the global state vector and unitary evolution, which must be deemed illegitimate in order to claim superiority for the MWI. Again, the only recourse seems to be that the quantum state supervenes on the computational histories of each world. But this will not work. Again, suppose the initial global state is $|\Phi_{ini}\rangle$. If $|\Phi_{ini}\rangle$ is the initial state, no matter what function is computed, the final state is always equal to $|\Phi_{ini}\rangle$, which is disposed to indicate that the function is balanced, no matter whether it is or not. So, the quantum state that each world is supposed to acquire does not exist, does not dispose measurement devices to indicate whether the function is constant or balanced, or supervenes on more than the computational histories of the individual worlds. Either way, Option 4 is unacceptable.

One might think that more options might be entertained whereby some combination of quantum states and definite values might underwrite the truth of the QPT, and then some quantum state or definite value might secure the EVAL. But the same argument in the discussion of Options 3 and 4 can be given against this possibility. The ontological feature that secures the EVAL has to supervene on the computational history that secures the truth of the QPT, but an appeal to $|\Phi_{ini}\rangle$ demonstrates that it cannot. Phase relations between worlds are crucial to accounting for the success of the toy algorithm. That said, appeal to phase relations and the

unitary evolution of that state is an appeal to the ontological significance of the global state as opposed to the local ontological features of worlds, precisely what an advocate of the MWI must find unacceptable in alternative accounts of quantum computational efficiency in order to claim superiority for the MWI.

4. Conclusion. The claim examined in this paper is that the MWI is the only interpretation that can account for why quantum computers seem faster than classical computers. This claim is based on the ability of the MWI to underwrite the truth of the QPT locally; however, if one takes the ontological status of the state vector seriously, there is little question about whether the QPT is true. In order to maintain that the MWI is superior to alternative accounts of quantum computational efficiency, proponents of the MWI must suggest that the ontological status of the universal state vector alone is not sufficient to explain quantum computational efficiency. In order for this objection to have any significance, the proponent of the MWI must show how to account for the extraction of information about a function that depends on the truth of the QPT without having to implicate the ontological status of the universal state vector.

This paper examined how various ontological features of worlds might account for the success of the toy algorithm in computing whether a function was constant or balanced. This was subject to the constraint that worlds were at least individuated by the terms in the computational basis. Four Options were considered as ways that local changes in worlds would secure the truth of the QPT and have the requisite ontological features to account for the success of the toy algorithm.

The universal validity of the Schrödinger evolution of the *entire quantum system* must be appealed to to account for the success of the efficiency of the toy algorithm. Now, those who endorse MWI surely would accept this account. The universal validity of the Schrödinger evolution is a fundamental assumption for MWIs. Moreover, it involves no introduction of definite values, which many advocates of the MWI would find objectionable. However, surely quantum computations as the one above cannot be used as evidence for the superiority of the MWI over other interpretations. Any interpretation in which BR is true and subscribes to the universal validity of the Schrödinger evolution will at least be on equal footing with the MWI regarding an account of quantum computational efficiency.

REFERENCES

- Clifton, Robert K. (1996), "On What Being a World Takes Away", *Philosophy of Science* 63: S151–S158.
Deutsch, David (1997), *The Fabric of Reality*. New York: Penguin.

- Duwell, Armond (2004), *How to Teach an Old Dog New Tricks: Quantum Information, Quantum Computing, and the Philosophy of Physics*. Ph.D. Dissertation. Pittsburgh: University of Pittsburgh.
- Jozsa, Richard (2000), "Quantum Algorithms", in D. Boumeester, A. Ekert, and A. Zeilinger (eds.), *The Physics of Quantum Information: Quantum Cryptography, Quantum Teleportation, Quantum Computation*. Germany: Springer, 104–125.
- Mermin, N. David (2000), "The Contemplation of Quantum Computation", <http://www.aip.org/pt/vol-53/iss-7/p11.html>.
- Steane, Andrew (2003), "A Quantum Computer Needs Only One Universe", *Studies in History and Philosophy of Modern Physics* 34: 469–478.