Procrastination in Quantum Coding and Computation

Paul Fitzpatrick

AI Lab, MIT, Cambridge, USA
paulfitz@ai.mit.edu

Abstract

When storing, communicating, or processing information, we are often forced to make decisions earlier than we would like. Perhaps we can only afford to store a subset of our data, and must choose what to discard long before we know which parts are truly important. Or perhaps we only have time to compute responses for a certain number of scenarios, and must decide which to prepare before the situation we face becomes into focus. This paper examines whether quantum resources can help us to delay those decisions until they can be made in an informed manner.

Quantum Procrastination

The possible states of a quantum system form a continuum. We could theoretically encode as many classical bits as we like in such a system, the catch being that when we measure it we can extract just a small part of that data. But there are many situations of practical interest where we care about just a small part of a large body of information. In particular, when we prepare for some future event, we need to cover many contingencies – despite the fact that once we reach that event, only one set of preparations will be relevant. In a situation like this it might be possible to prepare a compact quantum store covering all the contingencies, such that when the time comes to read the store we can choose our measurement carefully to extract the information relevant to the contingency we care about. Of course, the price paid for this compactness is likely to be that the store can only be used once – reading it may destroy its quantum state and any potential information about other contingencies.

The first part of this paper shows that the procedure just described is in fact feasible. For example, it is possible to code two classical bits in a single qubit, later decide which bit we care about, and read it out with an accuracy of 85%. And in a variant of the teleportation process, it turns that that this qubit can be passed to a party with whom we share an Einstein-Podolsky-Rosen pair by transmitting a single classical bit.

Another situation in which procrastination would be useful occurs for computation. Suppose we expect some input in the future that we wish to react to as quickly as possible, but the computation needed to respond is very long and involved. Is there any way to go ahead and do the computation and decide later what input we care about?

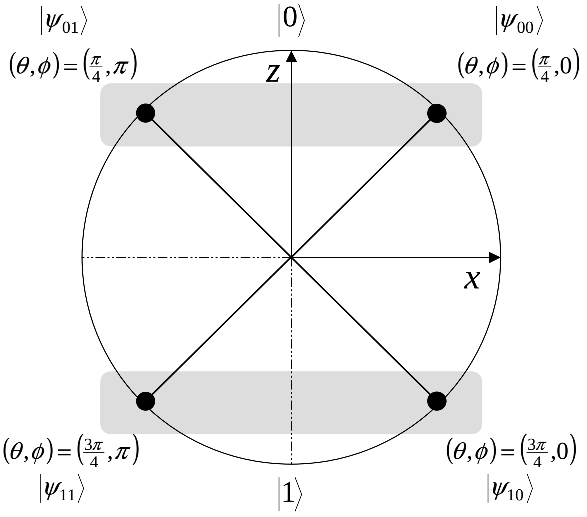
$|\psi_{01}\rangle$
$(\theta,\phi)=\left(\frac{\pi}{4},\pi\right)$
$|0\rangle$
z
$|\psi_{00}\rangle$
$(\theta,\phi)=\left(\frac{\pi}{4},0\right)$
x
$(\theta,\phi)=\left(\frac{3\pi}{4},\pi\right)$
$|\psi_{11}\rangle$
$|1\rangle$
$(\theta,\phi)=\left(\frac{3\pi}{4},0\right)$
$|\psi_{10}\rangle$

Fig. 1. Bloch sphere viewed along the y axis (with the positive sense of the axis pointing into the page). The gray shading shows pairs of states that have the same behavior when the observable Z is measured. When the observable X is measured instead, different pairs form, now grouped vertically. We can therefore choose a state to give whatever behavior we want for X and Z independently.

The second part of this paper reviews the work of Brukner, Pan, Simon, Weihs & Zeilinger (2001), in which the authors propose a mechanism for doing just this with some (small) probability.

Two Bits per Qubit

Suppose we have two classical bits of information that we might need to consult at some future time. We know that only one of the bits will actually be important, but we don't yet know which one. Can we trade the fact that we will only want to measure one bit for some reduction in the storage capacity required? This section shows that we can in fact store the two bits in a single qubit, if we are willing to accept some possibility of error (15%).

Figure 1 shows one way to achieve this. Two classical bits together can take on four possible values. We will associate one bit with the x axis and one with the z axis of the Bloch sphere. There are four states that are within 45° of both axes, as shown in the figure. If the Z observable is measured, the states segregate horizontally into the pairs shown in gray, in terms of measurement statistics. If the X observable is measured, the states segregate vertically. So we can encode the two classical bits by picking whichever of the four states will give appropriate X and Z measurements. Describing the four states in the computational basis gives:

$$\begin{aligned}|\psi_{00}\rangle &= \mathit{\Delta}|0\rangle + \delta|1\rangle\\ |\psi_{01}\rangle &= \mathit{\Delta}|0\rangle - \delta|1\rangle\\ |\psi_{10}\rangle &= \delta|0\rangle + \mathit{\Delta}|1\rangle\\ |\psi_{11}\rangle &= \delta|0\rangle - \mathit{\Delta}|1\rangle\end{aligned}$$

where

$$\begin{aligned}|\psi\rangle &= \cos\frac{\theta}{2}|0\rangle + e^{i\phi}\sin\frac{\theta}{2}|1\rangle\\ \delta &= \sin\frac{\pi}{8} = \frac{1}{\sqrt{(1+\sqrt{2})^2+1}} = 0.38\\ \mathit{\Delta} &= \cos\frac{\pi}{8} = \sqrt{1-\delta^2} = 0.92\end{aligned}$$

Measuring $|\psi_{00}\rangle$ or $|\psi_{01}\rangle$ in the computational basis will result in a $|0\rangle$ with probability $\mathit{\Delta}^2 = 0.85$, while measuring $|\psi_{10}\rangle$ or $|\psi_{11}\rangle$ in this basis will result in a $|1\rangle$ with probability 0.85. Similarly, measuring $|\psi_{00}\rangle$ or $|\psi_{10}\rangle$ in the basis formed by the eigenvectors of X will result in one basis vector with probability 0.85, while measuring $|\psi_{01}\rangle$ or $|\psi_{11}\rangle$ will result in the other basis vector with the same probability. Notice that the pairs of states with the same behavior with respect to the X observable are different from those that behave the same with respect to the Z observable.

Putting it all together, the procedure is as follows. Suppose Alice has two classical bits, A and B. First she generates a qubit in the state $|\psi_{AB}\rangle$ – for example by preparing the state $|0\rangle$ and then applying the appropriate rotation around the y axis, $R_y((1+2A)(1-2B)\frac{\pi}{4})$. She then sends this qubit to Bob (or waits to read it herself at a later point – the situations are equivalent). Bob decides if he wants to know about the value of A or the value of B. If he wants A, he immediately measures the qubit in the computational basis and gets the value of A with probability 0.85. If he wants B, he measures the observable Y (or equivalently he first rotates the qubit by $R_y(-\frac{\pi}{2})$ and then measures in the computational basis), and gets the value of B with probability 0.85. In either case, once he makes one measurement he cannot make the other.

If we were unwilling to accept some chance of error in our measurement, then storing two bits would clearly be impossible. To guarantee the outcome of a measurement with probability one, the qubit would need to be in one of the two eigenstates of the corresponding observable, leaving no room to encode information about a second bit. But if some possibility of error is acceptable, that opens up enough wiggle-room for storing information about the two bits. A later section shows that if we can tolerate a somewhat higher probability of error, even more wiggle-room opens up and we can squeeze a third bit into our long-suffering qubit.

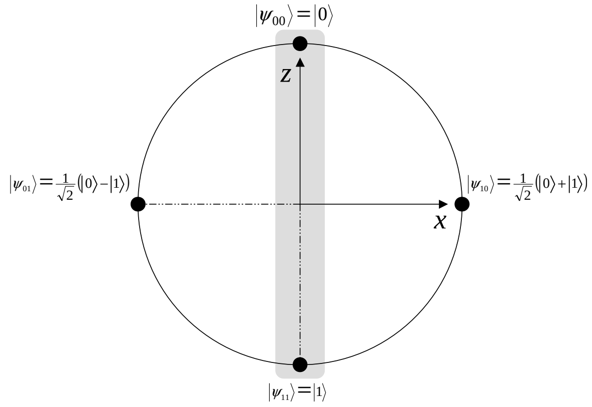
$|\psi_{00}\rangle = |0\rangle$
z
$|\psi_{01}\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$
$|\psi_{10}\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$
x
$|\psi_{11}\rangle = |1\rangle$

Fig. 2. Bloch sphere with the y axis pointing into the page. These states are similar to those in Figure 1, but rotated to make one pair impervious to bit flips and the other pair impervious to phase flips.

Two Bits per Bit (offer subject to terms and conditions)

If we combine the procedure from the previous section with teleportation, it turns out that we can send two bits from Alice to Bob using a single classical bit and a shared EPR pair – subject to the important constraints that Bob must only care about one of the bits and be willing to tolerate some probability of error. The key observation is that while quantum teleportation in general requires two classical bits to be sent, for a qubit in a known state drawn from one of four states separated by 90° on some plane through the Bloch sphere as in Figure 1, we can make do with a single bit.

If Alice and Bob share a pair of qubits in a Bell state, Alice can pass a qubit $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ to Bob by interacting $|\psi\rangle$ with her side of the Bell state pair. In the normal procedure, Bob will end up with a bit-flipped and/or phase-flipped variant of ψ, $|\psi'\rangle = \alpha|0\rangle \pm \beta|1\rangle$ or $|\psi'\rangle = \beta|0\rangle \pm \alpha|1\rangle$. During the teleportation process, Alice learns which of these four variants Bob will have. By sending Bob this information encoded in two classical bits, Alice gives Bob everything he needs to transform $|\psi'\rangle$ into the desired state $|\psi\rangle$.

Suppose we have a qubit prepared as described in the previous section. If we rotate the state of the qubit by 45° around the y axis, we get one of the four states shown in Figure 2. Bob can easily undo this rotation if Alice can transmit the qubit state correctly. When $|\psi\rangle$ is in one of these four states, Alice need only send Bob a single bit of information, telling him whether the $|\psi'\rangle$ he has access to after teleportation is the same as $|\psi\rangle$, or a variant. She does not need to tell him which variant it is. Each of the states in Figure 2 is left unchanged by one type of variation, either bit-flips or phase-flips. Since Alice constructed $|\psi\rangle$, she can easily tell whether $|\psi'\rangle$ will in fact be identical to $|\psi\rangle$ based on which kinds of flips occurred. If Bob is told that $|\psi'\rangle$ is a variant, he can recover $|\psi\rangle$ by simply applying the operator XZ to the qubit. Whatever the state is, it will be invariant to either the bit-flip X or the phase-flip Z, so Bob doesn't need to know which flip occurred in order to correct it. Hence one bit of classical information is all Alice need send to give Bob access to a state prepared as described in the previous section.

So with the expenditure of a shared pair of qubits in a Bell state, Alice can send Bob a single bit of classical information that allows him to answer one of two binary questions with 85% accuracy, where the choice of which question to ask is his and not Alice's.

It is worth comparing this process to superdense coding (Bennett & Wiesner 1992). In both cases, Alice has two classical bits she would like to send to Bob. They share an EPR pair of qubits. Alice performs some manipulation on her qubit to encode the two bits – in the case of superdense coding by directly applying one of four operators, and in our case by applying one of four operators to an intermediate qubit and then using the teleportation machinery to transfer the state. In superdense coding, Alice sends Bob her qubit, allowing him to recover the two bits exactly by determining which of four operators Alice applied to it. In our case, Alice only needs to send Bob a classical bit. The price for this is that Bob can only recover (his choice of) one of the bits, and even that has some probability of error – so this coding scheme is by no means superdense, but it is nevertheless moderately well packed.

Three Bits per Qubit

What if we want to pack three bits into a qubit? This is possible, although at the cost of lower accuracy (79%). Imagine growing cones around each direction along the three axes of a Bloch sphere, as shown in Figure 3. The intersection of these cones with the sphere represents the set of states that will give a particular result for the observable associated with an axis, with a certain probability that approaches chance as the cone becomes wider. We need to expand the cones until we get intersections between triplets of cones representing all the axes:

$$\begin{aligned} x^2 + y^2 &= \sin^2\theta \\ x^2 + z^2 &= \sin^2\theta \\ y^2 + z^2 &= \sin^2\theta \end{aligned}$$

The angle θ controls the size of the cones. This set of equations has eight solutions:

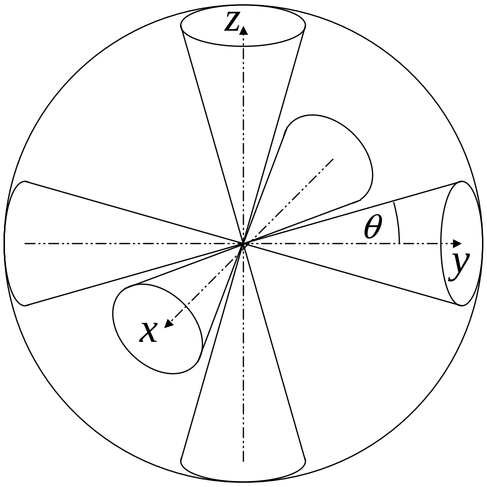

Fig. 3. Deriving a three-bit code. We need to find eight states to encode our three bits. This figure shows what can be coded for a given acceptable error rate (determined by the angle θ). We need to increase θ until x, y, and z cones intersect with each other – otherwise there will be some states that we cannot represent. Note that as θ grows, pairs of cones intersect before first, giving three two-bit codes equivalent to that described in the previous section.

$$(x, y, z) = \frac{1}{\sqrt{3}}(\pm 1, \pm 1, \pm 1)$$

Letting (s_x, s_y, s_z) be the signs of (x, y, z), chosen to correspond to three classical bits we wish to encode, we can translate this result into the following states:

$$\begin{aligned} |\psi\rangle &= 0.89|0\rangle + 0.32(s_x + s_y i)|1\rangle \quad \text{if } s_z = +1 \\ |\psi\rangle &= 0.46|0\rangle + 0.63(s_x + s_y i)|1\rangle \quad \text{if } s_z = -1 \end{aligned}$$

The expected value of the observable X works out to be $0.58 s_x$. Similarly $\langle Y\rangle$ is $0.58 s_y$ and $\langle Z\rangle$ is $0.58 s_z$. These results show that for whichever one of the three observables we choose to measure, we get the value of the corresponding classical bit with probability 79%.

More Bits?

If we try to pack four bits into a single qubit using a similar procedure, we run into trouble. It is easiest to see why by looking at how the three qubit case would fail if we were constrained to two dimensions. We require an arrangement of three axes, such that there is a point within 90° of any combination of poles of those axes. If this is not possible, then we will be unable to represent all the settings of the three bits. With three axes in two dimensions, for a solution that is symmetric in the bits we need the six poles

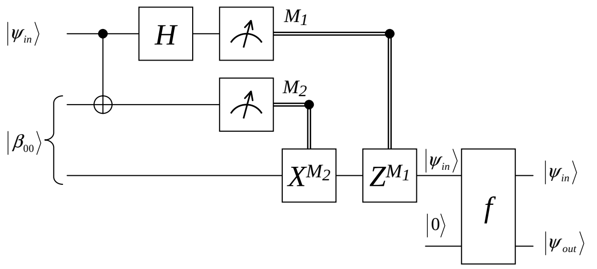

Fig. 4. A circuit for quantum teleportation is used to supply the input to a function. This circuit is adapted from Nielsen & Chuang (2001). $|\beta_{00}\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ is an Einstein-Podolsky-Rosen pair. With some probability the measurements M_1 and M_2 will reveal that neither of the bit flip or phase flip corrections implemented by X and Z are necessary, and we are free to rearrange the circuit as shown in Figure 5.

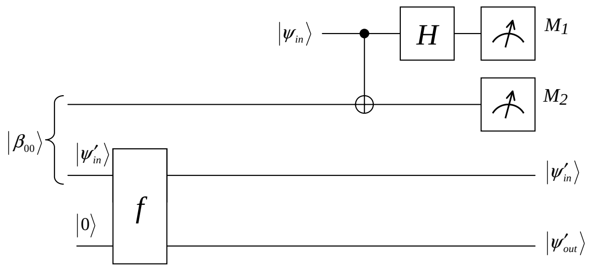

Fig. 5. With some probability the measurements M_1 and M_2 will reveal that no transformation need be applied to the unmeasured partner in the EPR pair in order for it to become $|\psi_{in}\rangle$. The central idea of Brukner et al. (2001) is that in this case, we can begin computing on that qubit as early as we like. In fact, we are free to compute with it long before $|\psi_{in}\rangle$ even becomes available, and need only wait for M_1 and M_2 to reveal that we can trust the result.

of the three axes to be 60° apart from their nearest neighbors. If we choose three poles that are 120° apart from each other, then it is clear that we cannot find a state that is within 90°of all of them. Hence we cannot represent at least one setting of the three bits. If we switch around the senses of the axes, or move them from the equi-angular arrangement, this fact will remain true. Another dimension is needed. The same is true of the four bit case in the three-dimensional Bloch sphere; we simply run out of dimensions, and can't represent all the possible settings of the four bits within the hyper-cone constraints.

We might also consider what happens when the number of qubits is increased. Suppose we have n qubits, from which we will later want to extract just a single bit with reasonable accuracy. How many bits can be encoded? I don't yet have any result better than simply coding two or three bits in the individual qubits and then choosing which to measure. Intuitively, it seems that this should be suboptimal, since it permits multiple of the original bits to be read rather than just one – we are retaining more information than is required. If it is possible to do better it will clearly require moving away from simple product states. Since the number of dimensions required to specify the state of an unentangled n qubit system is just $3n$, we can encode at most $3n$ bits this way.

Procrastination and Computation

The previous sections have shown that if we are willing to take a moderate gamble, it is possible to procrastinate on a choice that would otherwise need to be made up front: namely, which of a set of bits should be encoded for future use. Another choice we would often like to delay is which input to feed to a computation. Might it be possible to perform a computation before we have decided what the actual input should be? This could be useful if the computation is very long and we want to be prepared to make a fast response to an event.

In Brukner et al. (2001), the authors show that it is in fact possible to perform a quantum computation before its input has been specified, with some (unfortunately small) probability. To simplify their argument a little, let us consider a computation performed with just a single qubit as input. Suppose that qubit is supplied via the machinery of quantum teleportation, as shown in Figure 4. With teleportation, there is one chance in four that the qubit is transmitted unpermuted. So if we start computing with it immediately (Figure 5), before the input has actually been defined, there is some chance that when the teleportation procedure is finally applied we will end up with the correct result without any further processing required.

If there are n input qubits, the odds of success are just 1 in 4^n. This result is not interesting when applied to classical inputs, since we can do just as well here by simply picking a random input, computing for it, then accepting the result if we eventually find out we guessed the right input (our odds are 1 in 2^n, which is also what we would get in Figure 5 if we know that $|\psi\rangle$ is $|0\rangle$ or $|1\rangle$ and hence immune to phase flips). It is potentially useful for quantum inputs, but the low odds of success limit any obvious applications.

Suppose we are willing to give a mistaken result with some probability for a computation on a classical input. Can we do any better than simply computing for a random input and then using the result if it turns out that we chose the right input, or just guessing at the answer otherwise? This gives the right answer with probability $\frac{2^m+2^n-1}{2^{m+n}}$, where m and n are the number of inputs and outputs of the computation. This probability is tiny if the number of inputs and outputs are large.

For a single qubit input and output, it is hard to think of a way to do any better than this classical strategy (which gives a 75% probability of success in this case). The teleportation mechanism doesn't help for classical inputs. And if we instead compute on a superposition of the possible inputs, then we need some way to amplify the input we care about at a later point. But this is at least as hard as the problem faced in quantum search algorithms, where terms of a particular sign need to be amplified (Grover 1997). In our case we are after terms with a particular value of $|\psi\rangle$, which can be tagged with a negative sign by applying a controlled NOT to an ancillary $\frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$ state. More significantly, we could change the oracle in a search problem to tag solutions by giving them a particular value of $|\psi\rangle$ rather than changing their sign – so if we could solve our problem, we could solve quantum search. This reduction of our problem to quantum search shows that we shouldn't expect to be able to pull out the term we want without reapplying the function many times, as in Grover's algorithm. But if we have to reapply the function, than the whole procedure is moot, since once the correct input is known we might as well discard any earlier computations and apply the function directly to the real input! So there seems to be disappointingly little room to maneuver when it comes to procrastinating on which input to feed a computation.

Conclusions

This paper has shown that, to a certain extent, it is possible to delay decisions about which piece of information is important and worth transmitting until after the transmission has taken place. But it seems harder to delay decisions about what should be computed until after the computation has been made, since it is difficult to pull out the component we care about from a superposition without knowing what the other components are – requiring that we repeat the computation. Perhaps such procrastination will call for more drastic and fanciful measures (Moravec 1992).

Obligatory Quantum Story

Suppose an epic intergalactic war is being fought between two alien races, the S'hor and the Grvr. The Grvr have divided into two strike forces, separated for generations, that are approaching the S'hor stronghold from different directions. They have adopted this tactic because they are not quite sure of the location of the S'hor homeworld. As they draw closer, the S'hor detect a one bit message passing between the two Grvr forces. This is the only such message these groups have been able to pass because of continual electromagnetic jamming by the S'hor (the Grvr had anticipated this and bypassed the jamming by a prior agreement to encode the bit in the destruction of one of a pair of binary stars). What can the S'hor infer about this message? They know that the Grvr were uncertain whether a third race, the Bel, would be providing the S'hor with their rare and powerful weaponry – a key piece of strategic information that would dwarf any other considerations. The S'hor had recently deployed Bel weapons against the Grvr group that sent the message, so it seems logical for them to conclude that this is what the message indicates – and in fact that it is a waste, since they had also revealed their possession of Bel weapons to the second group during an attack. But their assumption that the message was a simple single bit of information turns out to be a grave mistake. The use of Bel technology was of crucial strategic concern to the Grvr, true, but a close second was the location of the S'hor homeworld. And in fact the Grvr group sending the message had detected the homeworld on their long-distance sensors. By prior agreement, and using shared EPR pairs and the techniques of this paper, that group encoded both the use of Bel weapons and their detection of the homeworld in the one-bit message. The other group, having independently learned of the weaponry, chose to read the homeworld-detection bit, saw it was set (with 85% probability, a virtual certainty amidst the vagaries of battle), and converged on the other group's trajectory to join forces and win the day.

References

Bennett, C., Brassard, G., Crépeau, C., Jozsa, R., Peres, A. & Wootters, W. (1993). Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels, *Physical Review Letters* **70**(13): 1895–1899.

Bennett, C. & Wiesner, S. (1992). Communication via one- and two-particle operators on Einstein-Podolsky-Rosen states, *Physical Review Letters* **69**(20).

Brukner, C., Pan, J., Simon, C., Weihs, G. & Zeilinger, A. (2001). Probabilistic instantaneous quantum computation, *arXive e-print quant-ph/0109022*.

Grover, L. K. (1997). Quantum mechanics helps in searching for a needle in a haystack, *Physical Review Letters* **79**(2): 325–328.

Gruska, J. & Imai, H. (2001). Power, puzzles and properties of entanglement, *Lecture Notes in Computer Science*, Vol. 2055, Springer-Verlag, pp. 25–69.

Moravec, H. (1992). Time travel and computing, *Extropy* **9**: 15–20.

Nielsen, M. & Chuang, I. (2001). *Quantum Computation and Quantum Information*, Cambridge University Press.