

NISQ computing for decision making under uncertainty

H.W.L. Naus*

Quantum Technology, TNO
PO Box 155, 2600 AD Delft, Netherlands

November 15, 2019

Abstract

Noisy Intermediate-Scale Quantum computers are expected to be available this year. It is proposed to exploit such a device for decision making under uncertainty. The probabilistic character of quantum mechanics reflects this uncertainty. Concomitantly, the noise may add to it. The approach is standard in the sense that Bayes decision rule is used to decide on the basis of maximum expected reward. The novelty is to model the various action profiles and the development of ‘nature’ as unitary transformations on a set of qubits. Measurement eventually yields samples of classical binary random variables in which the reward function has to be expressed. In order to achieve sufficiently low variances for reliable decision making more runs of such a quantum algorithm are necessary. Some simple examples have been worked out to elucidate the idea. Here the calculations are still analytically feasible. Presently lacking an operating quantum device, the QX simulator of Quantum Inspire has been used to generate the necessary samples for comparison and demonstration. First obtained results are promising and point at a possible useful application for noisy intermediate-scale quantum computers.

1 Introduction

Large-scale universal quantum computers [1] allowing fault tolerant quantum computing [2] are not foreseen in the near future. However, Noisy Intermediate-Scale Quantum (NISQ) devices will be available in the next years [3]. Such quantum computers have a limited number of qubits, say up to 100 and suffer from noise which reduces the fidelity of the quantum gates. A surface code to correct errors [2] cannot be implemented. The question therefore arises how to usefully apply NISQ technology. In the seminal paper [3] such opportunities, especially for computing, have been discussed.

This note proposes another application for NISQ computers, that is decision making in the face of uncertainty [4]. Classical probabilities and, more generally, probability theory are of course the essential guidelines in the process of decision making if uncertainties are involved. Probability density functions and their possible updates have been used in practice, see, e.g., [5]. In our proposal we do not choose probability distributions, but exploit quantum mechanical probabilities. In essence, we model the problem in terms of the state¹ of nature, typically with a uncertain evolution, and in terms of different actions (alternatives). The latter, in combination with the occurring state of nature, may give different rewards (payoffs). These actions and the evolution of the state of nature are modeled quantum mechanically by means a set of unitary transformations on a limited number of qubits. The latter are supposed to be initialized in their ground states as is commonly assumed in quantum algorithms. Eventually, we measure the product of all σ_z operators of the qubits; usually but somewhat sloppy this is called ‘measuring the qubits’. The measurement yields a number of classical bits, corresponding to one sample. Repeating the calculation, including the measurements, N times then yields N samples which are used to estimate the expectation value of the reward (payoff). The Bayesian decision rule [4] maximizes this expected reward.

*E-mail: rik.naus@tno.nl

¹This is, of course, not a quantum state

Of course, we need to motivate our choice for a NISQ device in decision making. First, we expect that even for complex decision making problems the necessary number of qubits is well below the expected limit of 100 qubits. Although the algorithms are deterministic, the outcome of the measurements is inherently probabilistic. In this way, the uncertainties are simulated in a natural way. The noise in a NISQ computer obviously also reflects these uncertainties and is therefore no nuisance like in other applications. Updating probabilities and probability density functions on the basis of previous experimentation [4] and/or new acquired information [5] is conveniently implemented by means of controlled quantum operations like a CNOT gate. Note that all two-qubit gates may introduce entanglement in the system, in this way modeling correlated random variables. We emphasize, however, that the merits of this method need eventually to be demonstrated in practice. Two advantages can be envisaged. Since the quantum technique is principally different from the classical one, it may be more reliable and robust. The involved quantum sampling may also speed up the process of classical decision making where time consuming Monte Carlo simulations are often necessary.

At present, a NISQ computer is not yet available. Quantum simulators, however, are already available. In our examples, we have used the QX simulator of Quantum Inspire [8]. As of yet, the simulator supports only one error model which may be used to include the imperfections of a real device. Quantum Inspire actually aims for a spin-qubit quantum computer in the very nearby future.

There is an obvious relation with quantum sampling where a quantum computer is exploited to generate samples of a probability density function. Here we also generate samples to obtain classical bits in which the reward function is expressed. No guiding classical probability distribution is however used; the guiding principle is the combination of the uncertain state of nature and the various actions. This needs to be modeled in by unitary transformations on the multi-qubit state. The method has also some resemblance with quantum random walks [7].

The outline of this note is as follows. First, we very briefly introduce the necessary concepts of decision making. Next we fix our notation for quantum computing. In section 3 the general concept of exploiting quantum computing in decision making is presented. It is followed by sections working out three simple examples with increasing complexity. Finally, some conclusions and an outlook are presented.

2 Preliminaries

In this section we start by shortly introducing the few aspects of decision analysis and theory which we actually want to use. The more interested reader may consult the standard textbook [4] and/or the recent review [9]. Our work is also based on an approach with classical probability density functions and Bayesian updates [10].

2.1 Decision making

The simplest example of decision making under uncertainty may be to decide whether or not to take an umbrella while going out. Such a decision is typically based on a weather forecast. It contains the essential ingredients which are several actions or alternatives to choose from, an eventual ‘state of nature’ and a resulting reward or payoff. In the example the possible eventual states of nature are rain or no rain. The reward is the difference between value, dry versus soaked clothes, and cost, carrying the umbrella. These concepts need to be quantified and often monetary values are used². Possibly unsubstantiated and subjective probabilities or probability distributions are invoked for predicting the state of nature. The decision maker chooses the optimal action based on some criterion. In [4], three such criteria are given:

- Maximin payoff

For each action find the minimum reward over all states. Next, find the maximum of these minimum payoffs. Choose the action whose minimum payoff yields this maximum. Obviously, it is a pessimistic view and is very cautious.

²The concept utility goes beyond this.

- Maximum likelihood

Identify the most likely state of nature, that is with the highest prior probability. For this state choose the decision with the highest reward. This criterion excludes low-probability high-payoff gambling.

- Bayes decision rule

Use the best available estimates of the probabilities (or probability distributions) of the respective states and calculate the expected reward for the various actions. Choose the one with the maximum reward.

In this study we restrict ourselves to the latter criterion, that is Bayes decision rule. It obviously requires more computations (simulations) than the other two.

Up to now, we have described decision making without experimentation. It can be extended to include experimentation [4] or Bayesian(-like) updates based on the already acquired information. The aim is to improve the estimated prior probabilities (or density functions). The improvements are called posterior probabilities (or updated density functions [10]). For simple examples we refer to [4, 10].

2.2 Notation quantum mechanics

In this section we define the notation to describe the necessary NISQ quantum mechanical part. For a single qubit we use as standard or computational basis

$$|0\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, |1\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}. \quad (1)$$

It is straightforwardly extended to more qubits; for two qubits we explicitly have

$$|00\rangle = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}, |01\rangle = \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \end{pmatrix}, |10\rangle = \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix}, |11\rangle = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix}. \quad (2)$$

The initial state for n qubits will always be chosen as $|0\rangle^{\otimes n}$. Note that the computational states are the eigenstates of the Pauli σ_z (or Z) operators

$$\sigma_z |0\rangle = |0\rangle, \quad \sigma_z |1\rangle = -|1\rangle. \quad (3)$$

The eigenvalues $1, -1$ are the possible results of a measurement of Z . They correspond to classical bits $0, 1$, thereby referring to the eigenstate. The following NISQ algorithms are always terminated by a simultaneous measurement of $Z^{[1]}, Z^{[2]} \dots$ and $Z^{[n]}$. In this way a classical bit vector \vec{s} with $s_k = 0, 1; k = 1, \dots, n$ is generated. Because of the inherent probabilistic character of quantum physics, it is produced with some probability depending on the state just before the measurements. This deterministic state follows from applying unitary operations on the initial state. Repeating the whole procedure, that is including the final measurement, then yields samples of the classical, random binary vector \vec{s} .

3 Concept

This leads us to the concept of decision making under uncertainty by means of a NISQ device. For each possible action A_j , the reward function has to be expressed in the classical random binary variables \vec{s} :

$$r_j(\vec{s}) = v_j(\vec{s}) - c_j(\vec{s}) \quad (4)$$

with value function $v(\vec{s})$ and cost function $c(\vec{s})$. These functions may be different for the various actions, but are not necessarily so. Furthermore, they may depend on other deterministic variables or parameters. The uncertainty is simulated by obtaining the random samples \vec{s} by means of the NISQ

computer. Apart from the quantum probabilities, the latter actually adds some uncertainty because of noise.

For each action, a quantum algorithm in terms of unitary operations U_k has to be implemented. Usually also the evolution of the state of nature is mimicked by an unitary transformation $U(\tau)$, where τ is a evolution parameter like a dimensionless time. For example, we get the final state

$$|\psi\rangle = U\mathcal{U}(\tau)|0\rangle^{\otimes n}. \quad (5)$$

The resulting algorithms have to be performed a number, say N , of times. In this way, N samples of the classical bits are obtained for each action. Inserting the samples into the reward function eventually yields samples of various rewards and therefore also their averages as sample means. These are used as decision criterion, the action with the highest expected reward is selected. Concomitantly, it is possible to estimate the variances of the rewards function by their sample variance. It of course yields an impression how sensitive a decision may be for the uncertainties. Specifally, we obtain for each action the variates $\vec{s}_k, k = 1, \dots, N$ and thus the samples $r(\vec{s}_k)$. The estimated expected reward³ E follows as

$$E[r] = \frac{1}{N} \sum_{k=1}^N r(\vec{s}_k). \quad (6)$$

The output is sometimes provided as a list of found different samples $\vec{s}(i), i = 1, L$ and their fractions n_i/N which is an estimate of their probabilities p_i . In that case, we obtain as average

$$E[r] = \sum_{i=1}^L \frac{n_i}{N} r(\vec{s}(i)) = \sum_{i=1}^L p_i r(\vec{s}(i)). \quad (7)$$

This is especially convenient if the number of qubits is small and, consequently, the number of different variates is not that large. Recall that we do not yet have a NISQ device at our disposal and therefore use the QX simulator of Quantum Inspire.

Although this concept appears to be simple, some thinking is required if a practical decision making needs to be solved. First, the set of possible actions has to be identified. Next, one needs to define appropriate reward functions in terms of the binary random variables \vec{s} and other parameters. Obviously, this resembles the approach in classical decision making. Note that there been published a vast amount of literature on this subject. Exploiting a quantum computer requires that the evolution of the state of nature and the consequences of the various actions need to be expressed as unitary transformations. For the case of decision making ‘with experimentation’ we propose to apply controlled unitary two-qubit gates, for example the CNOT. This leads to a generalized form of eq.(5)

$$|\psi\rangle = U_m \mathcal{U}(\tau_m, \tau_{m-1}) U_{m-1} \mathcal{U}(\tau_{m-1}, \tau_{m-2}) \cdots U_1 \mathcal{U}(\tau_1, \tau_0) |0\rangle^{\otimes n}, \quad (8)$$

where ‘time’ is ordered: $\tau_m > \tau_{m-1} > \cdots > \tau_0$. The unitary operations $U_j, j = 1, \dots, m$ can be controlled gates.

4 Taking an umbrella

The decision to be made is (not) taking an umbrella, given a certain probability p_r of rain. The umbrella, however, is heavy – so carrying it has a price. The possible states of nature can be coded in two classical bits: $s_0 = 0, 1$ corresponding to rain, no rain and $s_1 = 0, 1$ corresponding to travelling light, heavy. The quantum implementation therefore needs two qubits, *i.e.*,

$$\begin{array}{ll} \text{qubit}[0] : & |0\rangle \quad \text{no rain,} & |1\rangle \quad \text{rain,} \\ \text{qubit}[1] : & |0\rangle \quad \text{light,} & |1\rangle \quad \text{heavy.} \end{array} \quad (9)$$

The reward functions therefore depend on the two binary variables s_0, s_1 .

³Below we include a subindex to indicate the action, *i.e.*, E_j denotes the expected reward for action A_j .

4.1 Simplest case

In the simplest case, we only consider two actions, action 1 is *not* taking the umbrella whereas action 2 is taking the umbrella. Let us define the value function by noting that is some value v to remain dry which does not happen if $s_0 = 1$ and $s_1 = 0$; so we take

$$v(s_0, s_1) = v(1 - s_0(1 - s_1)). \quad (10)$$

Because the umbrella is heavy the cost is chosen as

$$c(s_0, s_1) = cs_1, \quad (11)$$

which is independent of s_0 . These functions apply for both decisions, so we also have only one reward function in this case

$$r(s_0, s_1) = v(1 - s_0(1 - s_1)) - cs_1. \quad (12)$$

The weather development is described with the unitary one-qubit gate

$$R_y(\tau) |0\rangle = \cos\left(\frac{\tau}{2}\right) |0\rangle - \sin\left(\frac{\tau}{2}\right) |1\rangle, \quad (13)$$

that is a y -rotation around τ . It yields a probability of rain

$$p_r = \sin^2\left(\frac{\tau}{2}\right), \quad (14)$$

so τ can be chosen to match the prediction. Despite the subjective experience, rain versus no rain does not depend on the decision⁴. The complete unitary operation for action 1 is given by the rotation on the $q[0]$ and the identity on $q[1]$, resulting in the state

$$|\Psi_1\rangle = (R_y^{[0]}(\tau) \otimes \mathcal{I}^{[1]}) |00\rangle. \quad (15)$$

Taking the umbrella corresponds to an X -operation on $q[1]$, yielding

$$|\Psi_2\rangle = (R_y^{[0]}(\tau) \otimes \sigma_x^{[1]}) |00\rangle. \quad (16)$$

The corresponding complete circuits, *i.e.*, including the measurements are shown in Figure (1).

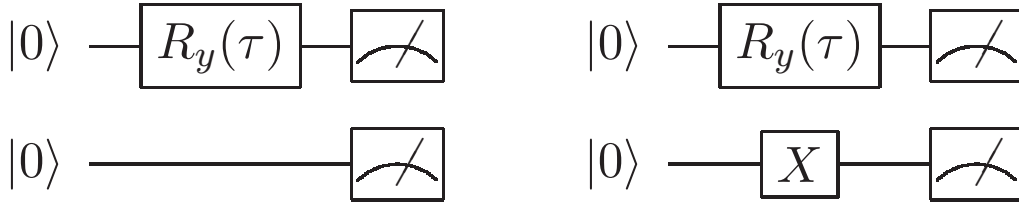


Figure 1: Corresponding quantum circuits, l.h.s: action 1; r.h.s: action 2

This example is that simple that we can proceed analytically. Standard quantum mechanical calculations yields for action 1

$$\begin{aligned} p(s_0 = 0, s_1 = 0) &= \cos^2\left(\frac{\tau}{2}\right), & p(s_0 = 1, s_1 = 0) &= \sin^2\left(\frac{\tau}{2}\right), \\ p(s_0 = 0, s_1 = 1) &= p(s_0 = 1, s_1 = 1) = 0, \end{aligned} \quad (17)$$

and eventually the expected reward

$$E_1[r] = v \cos^2\left(\frac{\tau}{2}\right). \quad (18)$$

The analogous computations for action 2 give

$$\begin{aligned} p(s_0 = 0, s_1 = 1) &= \cos^2\left(\frac{\tau}{2}\right), & p(s_1 = 1, s_1 = 1) &= \sin^2\left(\frac{\tau}{2}\right), \\ p(s_0 = 0, s_1 = 0) &= p(s_0 = 1, s_1 = 0) = 0, \end{aligned} \quad (19)$$

⁴Excluding ‘butterfly effects’.

with expected reward

$$E_2[r] = v - c. \quad (20)$$

Given the probability of rain p_r (14), one can now take the ‘best’ decision. It is taking the umbrella for $p_r > c/v$.

We nevertheless have implemented this example in the QX simulator to generate samples of \vec{s} . For each action 1024 runs have been done, *i.e.*, $N = 1024$. The resulting expected rewards are consistent with the analytical results.

4.2 With experimentation - entanglement

We extend this example by adding an third action. It consists out of waiting a certain time and then taking the decision also on the basis of the weather development thus far. An obvious criterion is rain at that moment⁵. A quantum mechanical implementation of this action is

$$|\Psi_3\rangle = (R_y^{[0]}(\tau - \tau_0) \otimes \mathcal{I}^{[1]}) C_N (R_y^{[0]}(\tau_0) \otimes \mathcal{I}^{[1]}) |00\rangle, \quad (21)$$

where C_N denotes the entangling CNOT gate. It corresponds to the quantum circuit shown in Figure (2). It is furthermore assumed that waiting has its price and for this decision the cost function is

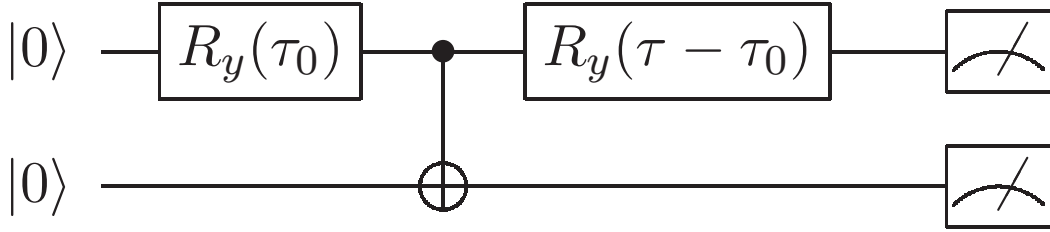


Figure 2: Circuits

therefore taken as

$$c_3(s_0, s_1) = cs_1 + d. \quad (22)$$

Since the cost increases for longer waiting we take $d \propto \tau_0/\tau$. Although the calculations are somewhat more involved, one can still proceed analytically. The resulting probabilities follow as

$$\begin{aligned} p(s_0 = 0, s_1 = 0) &= \cos^2\left(\frac{\tau - \tau_0}{2}\right) \cos^2\left(\frac{\tau_0}{2}\right), \\ p(s_0 = 0, s_1 = 1) &= \sin^2\left(\frac{\tau - \tau_0}{2}\right) \sin^2\left(\frac{\tau_0}{2}\right), \\ p(s_0 = 1, s_1 = 0) &= \sin^2\left(\frac{\tau - \tau_0}{2}\right) \cos^2\left(\frac{\tau_0}{2}\right), \\ p(s_0 = 1, s_1 = 1) &= \cos^2\left(\frac{\tau - \tau_0}{2}\right) \sin^2\left(\frac{\tau_0}{2}\right), \end{aligned} \quad (23)$$

eventually yielding the expected reward

$$\begin{aligned} E_3[r] &= v \left\{ \left(\cos^2\left(\frac{\tau - \tau_0}{2}\right) \cos^2\left(\frac{\tau_0}{2}\right) + \sin^2\left(\frac{\tau_0}{2}\right) \right) - c \sin^2\left(\frac{\tau_0}{2}\right) - d \right\} \\ &= v \left\{ 1 - \sin^2\left(\frac{\tau - \tau_0}{2}\right) \cos^2\left(\frac{\tau_0}{2}\right) \right\} - c \sin^2\left(\frac{\tau_0}{2}\right) - d. \end{aligned} \quad (24)$$

Once again, we also have performed QX simulations and have approximately reproduced the analytical probabilities for the third action as well.

If all parameters in the reward functions and in the unitary transformations are fixed, the decision is taken by selecting the one with maximal reward. It may illustrative to show the expected reward as a function of one the parameters. A natural choice is the probability of rain p_r , which determines the ‘time’ (rotation angle) τ . We fix the additional cost as $d = \frac{\tau_0}{2\tau}$ and take as parameters $v = 1.0, c = 0.8$. Figure (3) depicts the rewards for the three possible actions and two waiting times.

⁵Alternatives like an updated weather forecast are also reasonable.

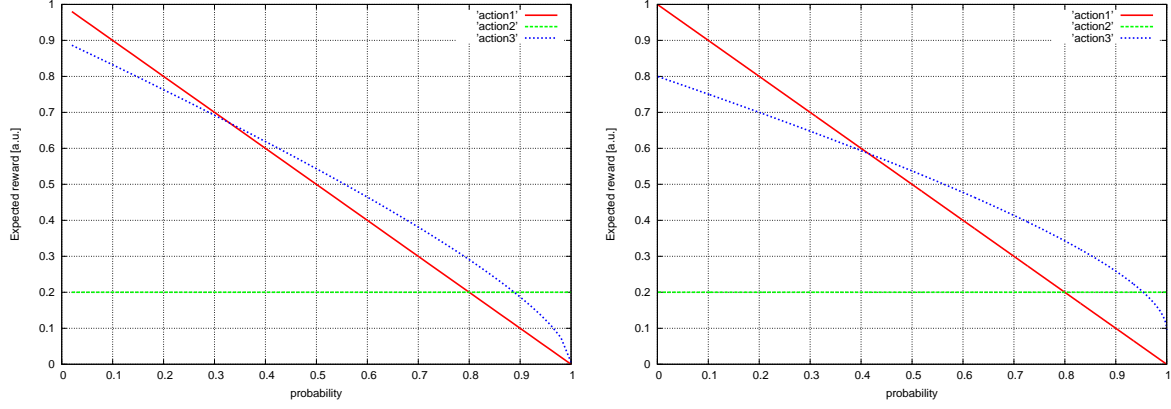


Figure 3: Expected rewards for the three different actions as a function of the rain probability p_r . L.h.s. $\tau_0 = 0.2\tau, d = 0.1$; r.h.s $\tau_0 = 0.4\tau, d = 0.2$.

4.3 Including variances

It is also possible to analytically derive the variances⁶ of the expected rewards for the various actions. In case of only sampling, either by means of the QX simulator or a NISQ computer, these quantities are replaced by sample variances. For action 1 we get

$$E_1[r^2] = v^2 \cos^2\left(\frac{\tau}{2}\right) \Rightarrow \sigma_1 = \frac{1}{2}v \sin \tau. \quad (25)$$

The variance of the expected reward for action vanishes identically. Sampling with a NISQ device will yield a finite but hopefully small variance. The expected squared reward for action 3 is obtained as

$$\begin{aligned} E_3[r^2] &= (v - d)^2 \cos^2\left(\frac{\tau - \tau_0}{2}\right) \cos^2\left(\frac{\tau_0}{2}\right) \\ &+ (v - c - d)^2 \sin^2\left(\frac{\tau_0}{2}\right) + d^2 \sin^2\left(\frac{\tau - \tau_0}{2}\right) \cos^2\left(\frac{\tau_0}{2}\right). \end{aligned} \quad (26)$$

The variance follows from $\sigma_3^2 = E_3[r^2] - (E_3[r])^2$, cf. (24), but we omit this rather lengthy expression.

It is clear, however, that the resulting standard deviation is of the same order of magnitude as the expected reward. Consequently, it is impossible to make a decision based on one experiment/run. Performing N runs with the QX simulator, and eventually on the NISQ device as well, reduces the standard deviation by a factor of $1/\sqrt{N}$. Figure (4) shows the expected reward including standard deviation calculated for $N = 1024$, at present the maximum number of runs for the QX simulator. We see that sensible decision making is still possible. The variance for the results generated with a

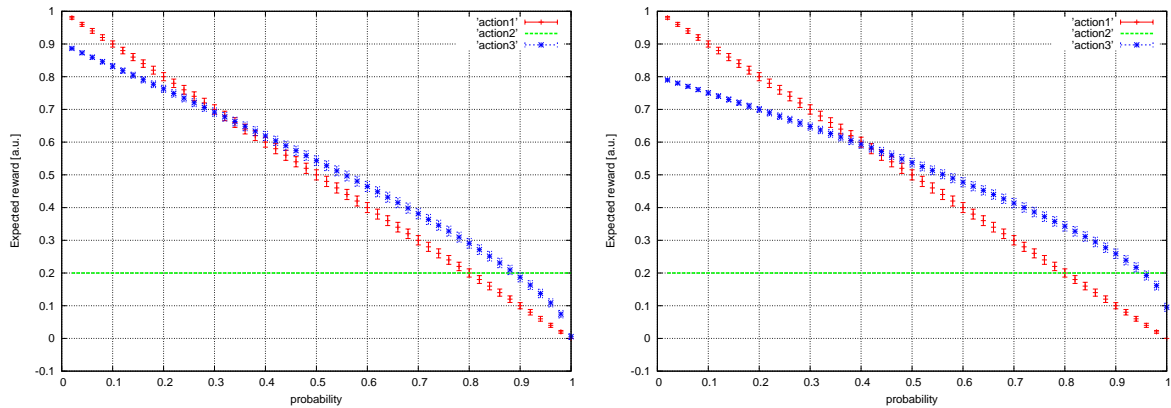


Figure 4: Expected rewards and variances for the three different actions as a function of the rain probability p_r ; $N = 1024$. L.h.s. $\tau_0 = 0.2\tau, d = 0.1$; r.h.s $\tau_0 = 0.4\tau, d = 0.2$.

NISQ computer may be larger.

⁶As is well known the standard deviation σ is the square root of the variance.

5 Example: entangling nature

5.1 Scenario and model

In this section we extend and modify the previous example. Apart from considering the possibility of rain the temperature is considered as well. If the temperature is high then one can leave a jacket at home, which is convenient. However, in case of low temperatures a jacket is necessary. Unfortunately, it is inconvenient to carry. Hence one needs to decide about taking an umbrella *and* carrying the jacket. The possible states of nature are now coded in four classical bits: $s_0 = 0, 1$ corresponding to rain and no rain, $s_1 = 0, 1$ corresponding to warm and cold, $s_2 = 0, 1$ corresponding to travelling light and heavy and $s_3 = 0, 1$ corresponding to travelling conveniently and inconveniently. The quantum implementation therefore needs four qubits, *i.e.*,

$$\begin{array}{llll} \text{qubit}[0] : & |0\rangle & \text{no rain,} & |1\rangle \text{ rain,} \\ \text{qubit}[1] : & |0\rangle & \text{warm,} & |1\rangle \text{ cold,} \\ \text{qubit}[2] : & |0\rangle & \text{light,} & |1\rangle \text{ heavy,} \\ \text{qubit}[3] : & |0\rangle & \text{convenient,} & |1\rangle \text{ inconvenient.} \end{array} \quad (27)$$

The reward functions therefore depend on the four binary variables s_0, s_1, s_2, s_3 . Analogously to the previous example, we define the value function

$$v(s_0, s_1, s_2, s_3) = v_0(1 - s_0(1 - s_2)) + v_1(1 - s_1(1 - s_3)). \quad (28)$$

Since inconvenience has its price, the cost is extended as

$$c(s_0, s_1, s_2, s_3) = c_0 s_2 + c_1 s_3. \quad (29)$$

The reward remains value minus cost. Just as above the cost function has to be adapted for one action; we will get back to this point.

First, we further modify the scenario by assuming that the decision cannot be based on the weather forecast but only on experience. That tells us that the conditions are either “warm and dry” *or* “cold and rainy”, with about 50% probability. Such a (weather) state of nature corresponds with the following Bell state⁷ of qubits [0] and [1]

$$|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle). \quad (30)$$

The unitary operation $C_N H^{[0]}$ transforms $|00\rangle$ to $|\Phi^+\rangle$ and, consequently, we define the fixed⁸ ‘weather evolution operator’ as

$$U_W = (\mathcal{I}^{[3]} \otimes \mathcal{I}^{[2]} \otimes C_N^{[01]}) (\mathcal{I}^{[3]} \otimes \mathcal{I}^{[2]} \otimes \mathcal{I}^{[1]} \otimes H^{[0]}) = U_W^{[01]} = C_N^{[01]} H^{[0]}, \quad (31)$$

where H is the Hadamard gate. Note the introduction of a shorter notation by only indicating the non-trivial operations and thus omitting identity gates.

We continue by defining the four obvious actions and their implementation. Action 1 is merely not taking anything. The resulting state therefore is

$$|\psi_1\rangle = U_W(|00\rangle \otimes |00\rangle) = U_W^{[01]} |00\rangle \otimes |00\rangle. \quad (32)$$

Carrying a jacket but no umbrella defines action 2. Hence the state is transformed as

$$|\psi_2\rangle = U_W \sigma_x^{[3]}(|00\rangle \otimes |00\rangle) = U_W^{[01]} |00\rangle \otimes \sigma_x^{[3]} |00\rangle. \quad (33)$$

Action 3 is defined as taking an umbrella and no jacket; the final state follows as

$$|\psi_3\rangle = U_W \sigma_x^{[2]}(|00\rangle \otimes |00\rangle) = U_W^{[01]} |00\rangle \otimes \sigma_x^{[2]} |00\rangle. \quad (34)$$

⁷Of course, the Bell state with the relative $-$ sign also does the job.

⁸No ‘time’ parameter is required here.

Of course one can take an umbrella as well as a jacket - defining action 4. Obviously, it corresponds to the final state

$$|\psi_4\rangle = U_W \sigma_x^{[2]} \sigma_x^{[3]} (|00\rangle \otimes |00\rangle) = U_W^{[01]} |00\rangle \otimes (\sigma_x^{[2]} \otimes \sigma_x^{[3]}) |00\rangle. \quad (35)$$

It is clear that for actions 1-4 the weather two-qubit state is entangled but that the remaining two qubits are separable. The quantum circuits are shown in Figures (5, 6).

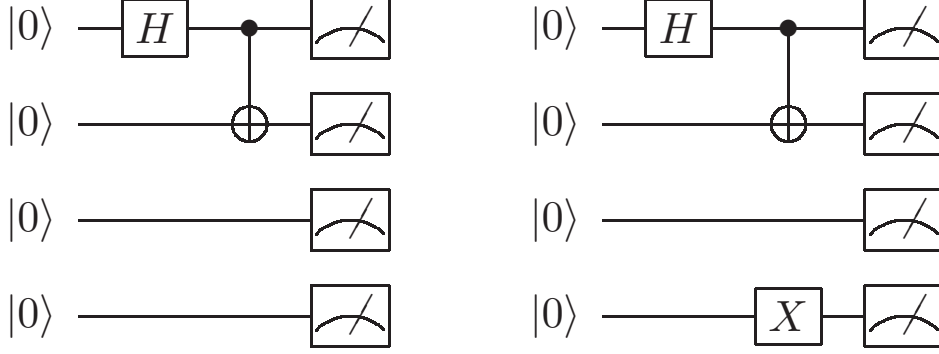


Figure 5: Corresponding quantum circuits, l.h.s: action 1; r.h.s: action 2

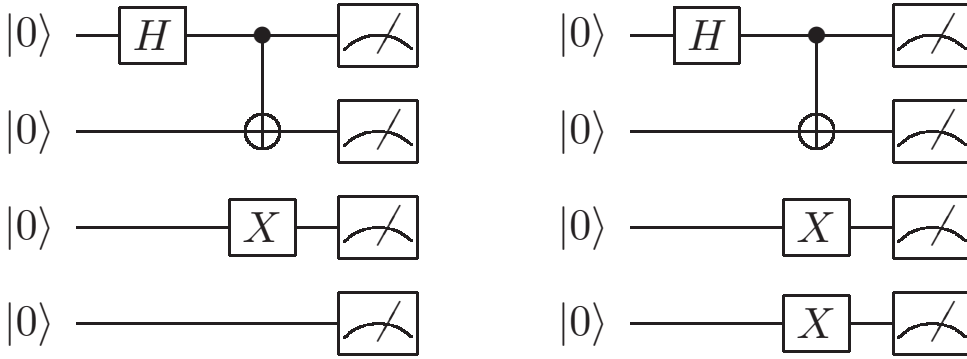


Figure 6: Corresponding quantum circuits, l.h.s: action 3; r.h.s: action 4

Once again, an alternative action 5 is obtained by modeling an intermediate decision moment. It depends on the actual weather condition and is implemented by a CNOT operation. Explicitly, we then obtain

$$|\psi_5\rangle = U_W C_N^{[13]} C_N^{[02]} U_W (|00\rangle \otimes |00\rangle). \quad (36)$$

The quantum circuit is depicted in Figure (7).

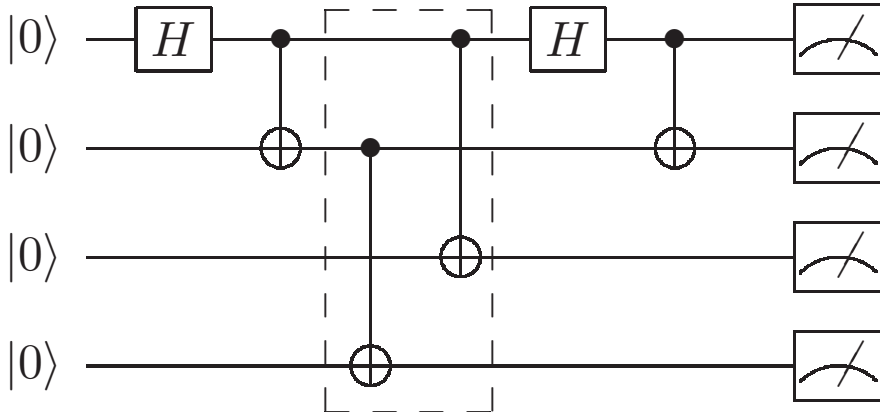


Figure 7: Quantum circuit for action 5, the CNOTs in the box can be done simultaneously

For action 5 there is an additional cost d leading to

$$c_5(s_0, s_1) = c_0 s_2 + c_1 s_3 + d. \quad (37)$$

This action and its implementation may be called ‘with experimentation’.

5.2 Results

The quantum mechanical calculations in this extended example are still analytically feasible. As above, we supplement and confirm these by the use of the QX simulator. For more complicated use cases, analytical predictions are impossible and one has to rely on the QX simulator only. Eventually, the aim is to get samples from the NISQ computer.

We present the results for each action in terms of probabilities for obtaining the classical bits \vec{s} and the eventual expected value for the reward. Note that we only list the non-zero probabilities

- Action 1

$$p(0, 0, 0, 0) = p(1, 1, 0, 0) = \frac{1}{2}, \quad E_1[r] = \frac{1}{2}(v_0 + v_1). \quad (38)$$

- Action 2

$$p(0, 0, 0, 1) = p(1, 1, 0, 1) = \frac{1}{2}, \quad E_2[r] = \frac{1}{2}v_0 + v_1 - c_1. \quad (39)$$

- Action 3

$$p(0, 0, 1, 0) = p(1, 1, 1, 0) = \frac{1}{2}, \quad E_3[r] = \frac{1}{2}v_1 + v_0 - c_0. \quad (40)$$

- Action 4

$$p(0, 0, 1, 1) = p(1, 1, 1, 1) = \frac{1}{2}, \quad E_4[r] = v_1 + v_0 - c_0 - c_1. \quad (41)$$

- Action 5

$$\begin{aligned} p(0, 0, 0, 0) &= p(1, 1, 0, 0) = p(0, 1, 1, 1) = p(1, 0, 1, 1) = \frac{1}{4}, \\ E_5[r] &= \frac{3}{4}(v_0 + v_1) - \frac{1}{2}(c_0 + c_1) - d. \end{aligned} \quad (42)$$

In principle, there are no free parameters and one just has to compare the expected rewards in order to take the optimal decision. It can be shown that for $d > 0$, action 5 is never the best action.

Suppose it would be possible to postpone the intermediate decision to the moment of the truth⁹. It is obvious that such an action profile gives better rewards than action 5 and plausible that it is the best of all for small enough d . In order to check these statements, we define the action 6 as action 5 without the last U_W operation. Its quantum circuit is shown in Figure (8).

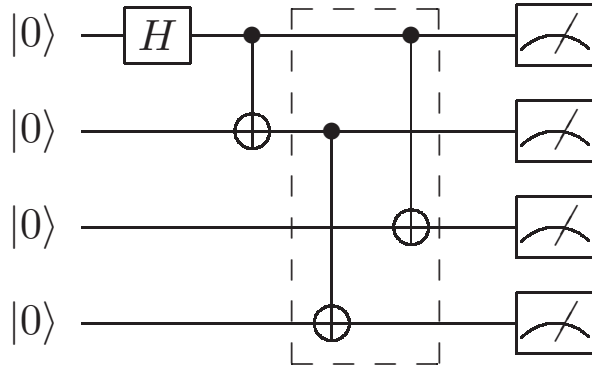


Figure 8: Quantum circuit for action 6, the CNOTs in the box can be done simultaneously

Thus we get as state

$$|\psi_6\rangle = C_N^{[13]} C_N^{[02]} U_W (|00\rangle \otimes |00\rangle). \quad (43)$$

⁹In the quantum model corresponding to the ideal measurement, that is instantaneous projection.

Of course, the additional cost d has to be included in the reward. The resulting probabilities and expected reward are obtained as

$$p(0, 0, 0, 0) = p(1, 1, 1, 1) = \frac{1}{2}, \quad E_6[r] = v_1 + v_0 - \frac{1}{2}(c_0 + c_1) - d. \quad (44)$$

We illustrate typical outcomes in this example in Figure (9). The rewards for the various actions are plotted as a function of the additional cost d in actions 5-6. The value parameters are fixed as $v_0 = 1.25, v_1 = 1.0$.

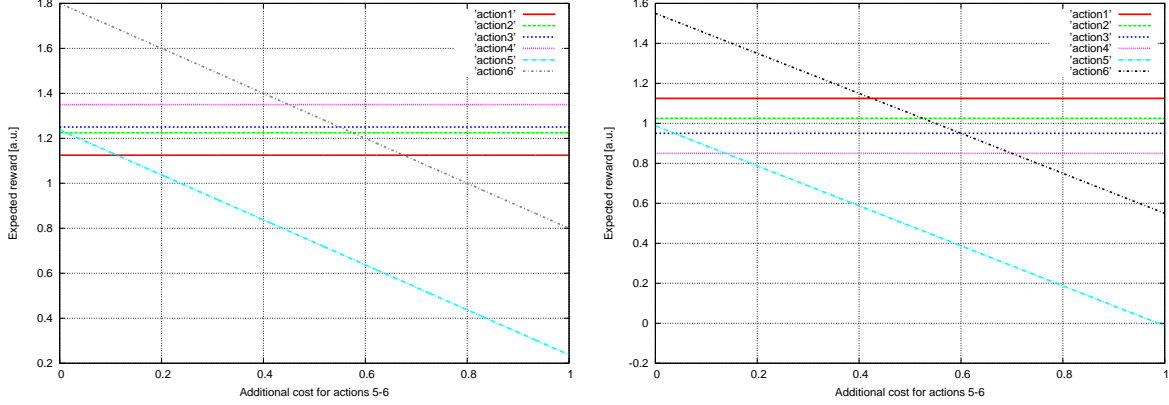


Figure 9: Expected rewards for the six different actions as a function of the additional cost d . L.h.s. $c_0 = 0.5, c_1 = 0.4$; r.h.s $c_0 = 0.8, c_1 = 0.6$.

5.3 Including variances

The variances and standard deviations of the rewards can also be straightforwardly calculated for the various actions. The following results are obtained

$$\begin{aligned} \sigma_1 &= \frac{1}{2}(v_0 + v_1), \quad \sigma_2 = \frac{1}{2}v_0, \quad \sigma_3 = \frac{1}{2}v_1, \quad \sigma_4 = 0, \\ \sigma_5 &= \frac{1}{2}\sqrt{\frac{3}{4}(v_0 + v_1)^2 + (c_0 + c_1)^2 - (v_0 + v_1)(c_0 + c_1)}, \\ \sigma_6 &= \frac{1}{2}(c_0 + c_1). \end{aligned} \quad (45)$$

Also in this example the standard deviations are of the same order of magnitude as the rewards. It prohibits decision making based on one run. As above, we assume that $N = 1024$ runs are performed which reduces the standard deviation by $1/\sqrt{N}$. The expected rewards including these standard deviations are depicted in Figure (10)

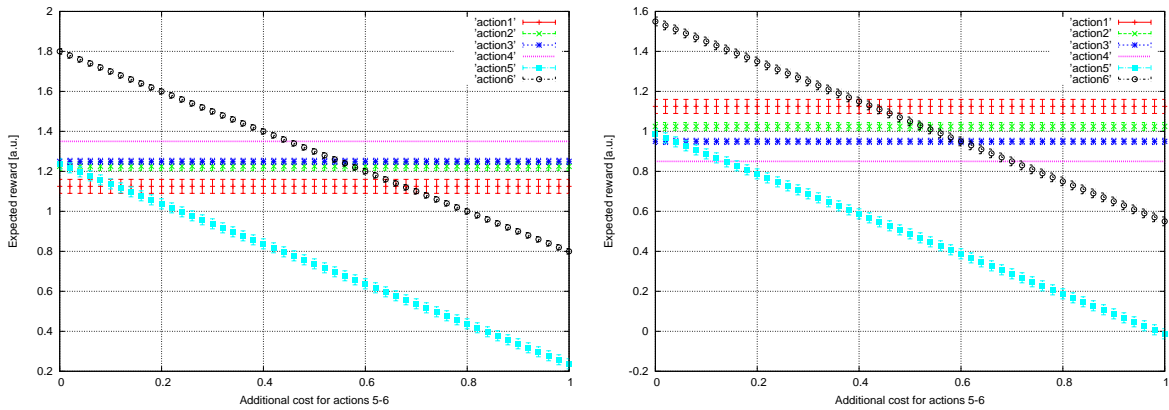


Figure 10: Expected rewards and variances for the six different actions as a function of the additional cost d . L.h.s. $c_0 = 0.5, c_1 = 0.4$; r.h.s $c_0 = 0.8, c_1 = 0.6$; $N = 1024$.

6 Including noise in the QX simulator

At present, only one error model is available in the QX simulator in Quantum Inspire, that is the “symmetric depolarizing channel” [8]. The per-operation error probability has to be set; typical values are between 0.001 and 0.01. In order to get a first assessment of the consequences, we have performed some additional simulations for action 6 of the second example. The parameters are chosen as in the l.h.s. of Figure (9); in addition we fix the additional cost as $d=0.5$. The number of runs remains $N = 1024$. The theoretical result for the expected reward and its standard deviation is $E_6[r] = 1.30 \pm 0.014$. We compare the sample mean and corresponding estimated error to theory for three error probabilities. In the noiseless case, we obtain the same result, for error probability 0.001 we get 1.33 ± 0.014 , whereas in case of error probability 0.01 we find 1.29 ± 0.024 . Of course, also these values slightly vary repeating the simulations of 1024 runs. These first results, only based on the implemented error model in the QX simulator, indicate that noise does not necessarily prohibit decision making using a NISQ computer. There will be a noise limit/threshold of course. Here we have seen that for an error probability of 0.1, results appear completely random.

7 Conclusion and outlook

An approach for decision making under uncertainty using a NISQ computer is proposed. According to Bayes decision rule [4], the decision alternative with the maximum expected reward is chosen. The reward is defined as the difference between value and cost and these are supposed to be formulated as functions of classical binary variables. Each binary variable requires a qubit in a NISQ device. Measuring these qubits in the standard basis generates a sample of the binary variables. Because of the inherent probabilistic nature of such a quantum measurement, the samples are random thereby reflecting the uncertainty in the decision problem. The noise in a NISQ computer will also add to this. The various decision actions and the development of the ‘state of nature’ are modelled as unitary transformations on the qubits. After initialization, these operations are to be performed on the qubits before measuring them. Such a quantum program has to run a sufficient number of times.

The idea is demonstrated by means of some simple examples of decision making. In fact, the calculations can still be performed analytically. Nevertheless we have cross-checked results by generating qubit measurement outcomes to obtain samples. Awaiting a NISQ device, we have actually used the QX simulator of Quantum Inspire[8]. The results indicate the feasibility of the approach. In order to reduce the variance in the expected rewards, multiple runs of the quantum algorithm are necessary. The QX simulator supports 1024 runs, which is sufficient for the problems considered.

In classical decision making in the face of uncertainty, one relies on probabilities and probability density functions. Judicious choices for such distributions and classical sampling and/or probability theory computations are necessary. In the proposed quantum approach this is replaced by unitary transformations and measurement of the qubits. Note that the formulation of the classical reward function remains in principle the same. It is constrained, however, by its dependence on random binary variables.

It is planned to test these ideas as soon as a up-and-running NISQ computer is available. A small device is already scheduled for this year [8]. If these first evaluations are succesful, then larger and more complex decision making problems may be implemented.

Acknowledgements

The author thanks Kelvin Loh for a critical reading of the manuscript and J  r  my Veltin for his support.

References

- [1] M.A. Nielsen and I.L. Chuang, *Quantum Computation and Quantum Information*, Cambridge University Press, Cambridge, UK (2000).

- [2] A.G. Fowler, M. Mariantoni, J.M. Martinis, and A.N. Cleland, *Surface codes: Towards practical large-scale quantum computation*, Phys. Rev. A **86**, 032324 (2012).
- [3] J. Preskill, *Quantum computing in the NISQ era and beyond*, arXiv:1801.00862 [quant-ph] (2018).
- [4] F.S. Hillier and G.L. Lieberman, *Introduction to Operations Research*, Eight Edition, Mc-Graw-Hill, New York (2005).
- [5] S. Giodini, A.J. Hunter, H.W.L. Naus, B. Bakker, D.J. Bekers, M. Ditzel, R. van Vossen, S. Dugelay, F. Baralli, A.L.D. Beckers, *Decision making on AUVs for adaptive minehunting surveys*, Proceedings of the European Conference on Undersea Defence Technology, UDT Europe 2014, Liverpool, UK (2014).
- [6] A.P. Lund, M.J. Brenner and T.C. Ralph, *Quantum sampling problems, BosonSampling and quantum supremacy*, NPJ Quantum Information **3**, article nr. 15 (2017).
- [7] J. Kempe, *Quantum random walks - an introductory overview*, Contemporary Physics, **44**, 307 (2003).
- [8] See www.quantum-inspire.com
- [9] C. Dimitrakakis and R. Ortner, *Decison making under uncertainty and reinforcement learning*, www.cse.chalmers.se/~chrdimi/downloads/book.pdf (2019).
- [10] H.W.L. Naus, S. Giodini, D. Bekers and B.J.H. Bakker, *On utilities, decisions and actions*, Internal TNO-note (2013); available on request.