Quantum trajectories

A quantum system undergoing decoherence can be described as evolving by a sequence of completely positive maps:

$$\rho \to \sum_k \hat{A}_k \rho \hat{A}_k^\dagger$$

$$\to \sum_{k_1,\ldots,k_n} \hat{A}_{k_n} \cdots \hat{A}_{k_1} \rho \hat{A}_{k_1}^\dagger \cdots \hat{A}_{k_n}^\dagger,$$

where the set of operators $\{\hat{A}_k\}$ satisfy

$$\sum_k \hat{A}_k^\dagger \hat{A}_k = \hat{I}.$$

This condition is just like the condition for a generalized measurement. What would happen if this actually *were* a measurement?

Suppose the state is initially pure, $\rho = |\psi\rangle\langle\psi|$. Then after a measurement the state would become

$$|\psi\rangle \to |\psi_k\rangle = \hat{A}_k|\psi\rangle/\sqrt{p_k}$$

with probability

$$p_k = \mathrm{Tr}\{\hat{A}_k|\psi\rangle\langle\psi|\hat{A}_k^\dagger\}.$$

If we consider the set of all outcome states with their probabilities $\{p_k, |\psi_k\rangle\}$ as an ensemble, it corresponds to the density matrix

$$\rho_1 = \sum_k p_k|\psi_k\rangle\langle\psi_k| = \sum_k \hat{A}_k|\psi\rangle\langle\psi|\hat{A}_k^\dagger,$$

which will in general be a mixed state.

Suppose the outcome of the first measurement is k_1, and we then perform a second measurement. Outcome k_2 has probability

$$p_{k_2|k_1} = \mathrm{Tr}\{\hat{A}_{k_2}|\psi_{k_1}\rangle\langle\psi_{k_1}|\hat{A}^\dagger_{k_2}\}$$

$$= \frac{1}{p_{k_1}}\mathrm{Tr}\{\hat{A}_{k_2}\hat{A}_{k_1}|\psi\rangle\langle\psi|\hat{A}^\dagger_{k_1}\hat{A}^\dagger_{k_2}\}.$$

We can define the joint probability to be

$$p_{k_2k_1} = p_{k_1}p_{k_2|k_1}$$

$$= \mathrm{Tr}\{\hat{A}_{k_2}\hat{A}_{k_1}|\psi\rangle\langle\psi|\hat{A}^\dagger_{k_1}\hat{A}^\dagger_{k_2}\}.$$

Then after two measurements with outcomes k_1 and k_2, the state is

$$|\psi_{k_2k_1}\rangle = \hat{A}_{k_2}\hat{A}_{k_1}|\psi\rangle/\sqrt{p_{k_2k_1}}.$$

Treating all these states as forming an ensemble, the corresponding density matrix is

$$\rho_2 = \sum_{k_1,k_2} \hat{A}_{k_2}\hat{A}_{k_1}|\psi\rangle\langle\psi|\hat{A}^\dagger_{k_1}\hat{A}^\dagger_{k_2}.$$

Similarly, after n steps of this process, we have n outcomes $k_1,\ldots,k_n$; the probability of the outcomes is

$$p_{k_n\ldots k_1} = \mathrm{Tr}\{\hat{A}_{k_n}\cdots\hat{A}_{k_1}|\psi\rangle\langle\psi|\hat{A}^\dagger_{k_1}\cdots\hat{A}^\dagger_{k_n}\},$$

and the resulting state is

$$|\psi_{k_n\ldots k_1}\rangle = \hat{A}_{k_n}\cdots\hat{A}_{k_1}|\psi\rangle/\sqrt{p_{k_n\ldots k_1}}.$$

This ensemble of states corresponds to a density matrix

$$\rho_n = \sum_{k_1,\ldots,k_n} \hat{A}_{k_n}\cdots\hat{A}_{k_1}|\psi\rangle\langle\psi|\hat{A}^\dagger_{k_1}\cdots\hat{A}^\dagger_{k_2}.$$

This means we can replace our density matrix evolution with a *pure state*, undergoing a random evolution resulting from repeated generalized measurements. By averaging over all possible outcomes with appropriate probabilities, the density matrix is reconstructed.

Such a random pure state evolution is called a *quantum trajectory*. If we re-write the completely positive evolution as an average over trajectories, we say that we have *unraveled* the evolution.

In fact, in principle it is possible to convert a deterministic, completely positive evolution into a random pure state evolution by doing a complete measurement on the environment. We can see how this works by looking at our simple q-bit model of decoherence.

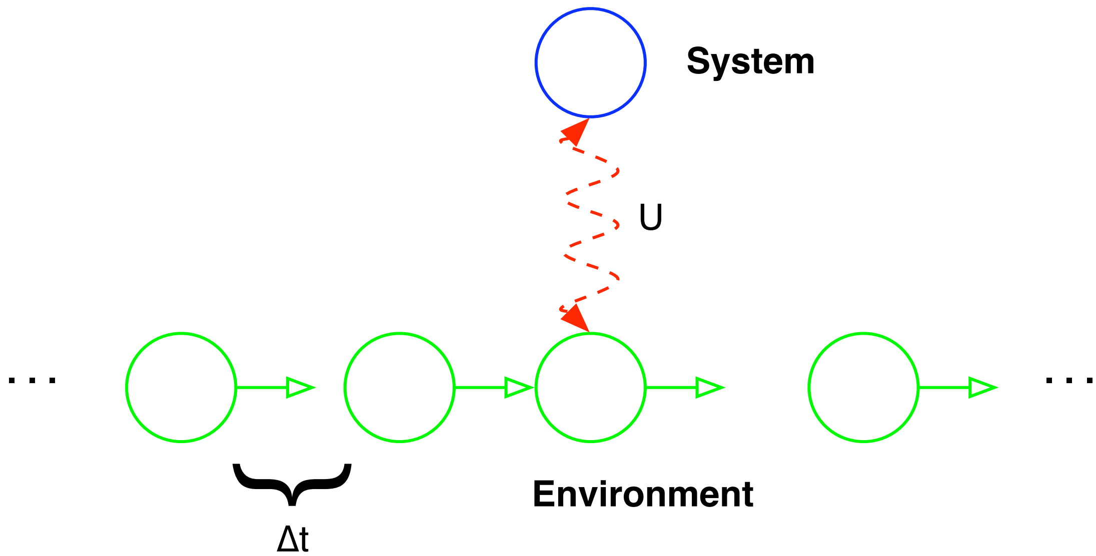

Suppose the environment bits are in the state $\rho = (1-p)|0\rangle\langle 0| + p|1\rangle\langle 1|$, and the interaction $\hat{U}$ is a CNOT from the environment onto the system. If the system is in state $|\psi\rangle$, then after the interaction the reduced density matrix for the system is

$$|\psi\rangle\langle\psi| \rightarrow (1-p)|\psi\rangle\langle\psi| + p\hat{X}|\psi\rangle\langle\psi|\hat{X}.$$

If, however, we *measure* the environment bit after the interaction in the $|0\rangle, |1\rangle$ basis, then the system is left in the state $|\psi\rangle$ or $\hat{X}|\psi\rangle$ with probabilities $1-p$ or p.

This change, however, depends on the choice of measurement. Suppose that instead of measuring the environment bit in the $|0\rangle, |1\rangle$ basis, we measure it in the basis $(|0\rangle \pm |1\rangle)/\sqrt{2}$. In this case, we learn *nothing* about whether or not the system bit was flipped. The transformation remains

$$|\psi\rangle\langle\psi| \to (1-p)|\psi\rangle\langle\psi| + p\hat{X}|\psi\rangle\langle\psi|\hat{X}.$$

The important point is that we have to choose a measurement of the environment which gives information about the system. In some cases (for instance, if the initial state of the environment bits is pure), there may be many different measurements which give information about the system, which will in general result in different unravelings.

Noisy gates

Environmental decoherence can lead to errors in QIP; but it is not the *only* source of error. Another problem is imprecision in carrying out quantum gates.

Here is a simple example. Suppose we wish to carry out the 1-bit gate $\hat{R}_z(\theta)$. We could do this (for example) by turning on a magnetic field in the Z direction with a particular strength for a particular length of time.

Neither the field strength nor the timing will be absolutely precise. This means that the *actual* angle of rotation θ' will in general be a random variable, (hopefully) centered on the desired value of θ:

$$\theta' \equiv \theta + \delta\theta,$$

where we assume $\delta\theta$ is a Gaussian variable

$$M[\delta\theta] = 0, \quad M[\delta\theta^2] = \epsilon \ll 1.$$

In place of the desired gate $\hat{R}_z(\theta)$ we have done $\hat{R}_z(\theta+\delta\theta) = \hat{R}_z(\delta\theta)\hat{R}_z(\theta)$. We don't know the value of $\delta\theta$; so the best we can do is describe the state by a density operator

$$\int d(\delta\theta)\, p(\delta\theta)\hat{R}_z(\delta\theta)\hat{R}_z(\theta)|\psi\rangle\langle\psi|\hat{R}_z^\dagger(\theta)\hat{R}_z^\dagger(\delta\theta)$$

This may look like an infinite ensemble; but in fact, this integral can be replaced by the finite quantum operation

$$|\psi\rangle\langle\psi| \rightarrow (1-\epsilon)\hat{R}_z(\theta)|\psi\rangle\langle\psi|\hat{R}_z^\dagger(\theta)$$

$$+\epsilon\hat{Z}\hat{R}_z(\theta)|\psi\rangle\langle\psi|\hat{R}_z^\dagger(\theta)\hat{Z}.$$

We can describe the imperfect gate as a *perfect* gate followed by a *weak* (i.e., near the identity) completely positive map.

What if there are environmental interactions which go on during the gate, as well as imprecision? We can use the same trick in that case as well; instead of the desired transformation

$$|\psi\rangle\langle\psi| \to \hat{U}|\psi\rangle\langle\psi|\hat{U}^\dagger,$$

we instead get some transformation

$$|\psi\rangle\langle\psi| \to \sum_k \hat{A}_k \hat{U}|\psi\rangle\langle\psi|\hat{U}^\dagger \hat{A}_k^\dagger,$$

where the particular completely positive transformation will (in general) depend on which unitary $\hat{U}$ we were trying to perform. If the decoherence and imprecision are both small, then this completely positive transformation will be weak.

If we were attempting a series of gates

$$|\psi\rangle\langle\psi| \rightarrow \hat{U}_n \cdots \hat{U}_1 |\psi\rangle\langle\psi| \hat{U}_1^\dagger \cdots \hat{U}_n^\dagger,$$

we would instead get the state

$$\sum_{k_1,\ldots,k_n} \hat{A}_{n,k_n} \hat{U}_n \cdots \hat{A}_{1,k_1} \hat{U}_1 |\psi\rangle\langle\psi|$$

$$\times \hat{U}_1^\dagger \hat{A}_{1,k_1}^\dagger \cdots \hat{U}_n^\dagger \hat{A}_{n,k_n}^\dagger.$$

Each unitary transformation is followed by a completely positive map; the particular map depends on which unitary was performed.

Error models

We now use the correspondence between the random pure-state evolution and the density matrix evolution to define an *error model*. If the state is $|\psi\rangle$ and we attempt to do a gate $\hat{U}$, then after the gate there will be a *noise* process

$$|\psi\rangle \rightarrow \hat{A}_k \hat{U} |\psi\rangle / \sqrt{p_k}$$

with probability

$$p_k = \langle\psi|\hat{A}_k^\dagger \hat{A}_k|\psi\rangle, \quad \sum_k \hat{A}_k^\dagger \hat{A}_k = \hat{I},$$

where the set of operators $\{\hat{A}_k\}$ will usually depend on the particular gate that is done.

For practical purposes, we are usually interested in noise processes that are *weak*, i.e., that change the state little on average.

We define a weak noise process as one that is close to the identity superoperator:

$$\rho' = \sum_k \hat{A}_k \rho \hat{A}_k^\dagger,$$

$$||\rho - \rho'|| \ll 1 \, \forall \, \rho.$$

In looking at the stochastic evolution, a weak noise process does *not* imply that $\hat{A}_k|\psi\rangle/\sqrt{p_k}$ is close to $|\psi\rangle$ for all k. In fact, the change can be very large. It does, however, imply that an operator $\hat{A}_k$ which causes a large change in the state must have a low probability; and operators which have high probability must be close to the identity.

Examples

For our imprecise Z rotation gate, we saw that the completely positive map that followed it could be written

$$\rho \to (1-\epsilon)\rho + \epsilon \hat{Z}\rho\hat{Z}.$$

This gives us two *error operators*:

$$\hat{A}_0 = \sqrt{1-\epsilon}\hat{I}, \quad \hat{A}_1 = \sqrt{\epsilon}\hat{Z}.$$

In this case, the error probability is fixed: a probability ϵ after each gate of a phase-flip error. Obviously, for rotations about different axes, the type of error will be different.

We have aleady seen how our simple environment model can result in bit-flip errors:

$$\rho \to (1-p)\rho + p\hat{X}\rho\hat{X}.$$

The form is exactly the same as that for imprecise gates. So it makes sense to treat all errors the same, regardless of their source.

One might think that errors from decoherence would be independent of the gate which is performed; but this is not true, in general. If an error occurs *during* the performance of a gate, the action of the gate can transform it into a different error.

Let's look at a simple example. Suppose we perform a gate $\hat{U}$ by turning on a Hamiltonian $\hat{H}$ for a time τ. (For the moment we ignore imprecision in the gate.)

$$\hat{U} = \exp(-i\hat{H}\tau/\hbar).$$

Suppose an error $\hat{A}_k$ occurs in the middle of this gate. Then the state becomes

$$|\psi\rangle \rightarrow e^{-i\hat{H}t/\hbar}\hat{A}_k e^{-i\hat{H}(\tau-t)/\hbar}|\psi\rangle/\sqrt{p_k}$$

$$= \left(e^{-i\hat{H}t/\hbar}\hat{A}_k e^{i\hat{H}t/\hbar}\right)\hat{U}|\psi\rangle/\sqrt{p_k}.$$

Unless $\hat{A}_k$ commutes with $\hat{H}$, the error $\hat{A}_k$ has become a new error $e^{-i\hat{H}t/\hbar}\hat{A}_k e^{i\hat{H}t/\hbar}$.

Independent errors

In QIP, the usual assumption is of *independent errors*. We assume the random deviations in different gates are *not* correlated with each other; and that each q-bit interacts with a *separate* environment.

These are the usual assumptions:

1. Every gate has an error process associated with it, which is essentially the same no matter to which q-bits the gate is applied (e.g., a CNOT on bits i and j has the same error process for every value of i and j).

2. Errors in gates occur only to the bits to which the gate is applied (e.g., a gate on bits 1 and 3 will not cause an error in bit 2).

3. A q-bit not undergoing a gate has some intrinsic error process (storage errors) which is the same and independent for all bits.

A canonical kind of intrinsic noise for a q-bit is *depolarizing* noise:

$$\rho \rightarrow (1 - p_x - p_y - p_z)\rho + p_x \hat{X}\rho\hat{X}$$

$$+p_y \hat{Y}\rho\hat{Y} + p_z \hat{Z}\rho\hat{Z},$$

with error operators

$$\begin{aligned}
\hat{A}_0 =& \sqrt{1 - p_x - p_y - p_z}\,\hat{I}, \\
\hat{A}_1 =& \sqrt{p_x}\,\hat{X}, \\
\hat{A}_2 =& \sqrt{p_y}\,\hat{Y}, \\
\hat{A}_3 =& \sqrt{p_z}\,\hat{Z}.
\end{aligned}$$

At long times, all initial states will tend towards the maximally mixed state $\rho = \hat{I}/2$.

This type of noise is often used as a model for a *noisy quantum channel*. But other error processes may be more realistic for particular systems.

The independent errors assumption can be violated. For instance, if different q-bits represent different kinds of physical systems, they might be expected to have different error processes (e.g., electron spins versus nuclear spins). Some kinds of systems may have correlated errors (e.g., Pauli exchange errors for spin-1/2). Figuring out a good error model for a particular system may require very detailed experimental measurements.

Whatever the noise process may be, however, the practical question is: what do we do about it? Errors will accumulate linearly with time, which implies that quantum computations past a certain size will be impossible. Fortunately, techniques have been devised to cope with this, even when the exact error model is unknown, so long as the error process is sufficiently weak.

Next time: quantum error correction.