Quantum error-correcting codes and quantum cryptography

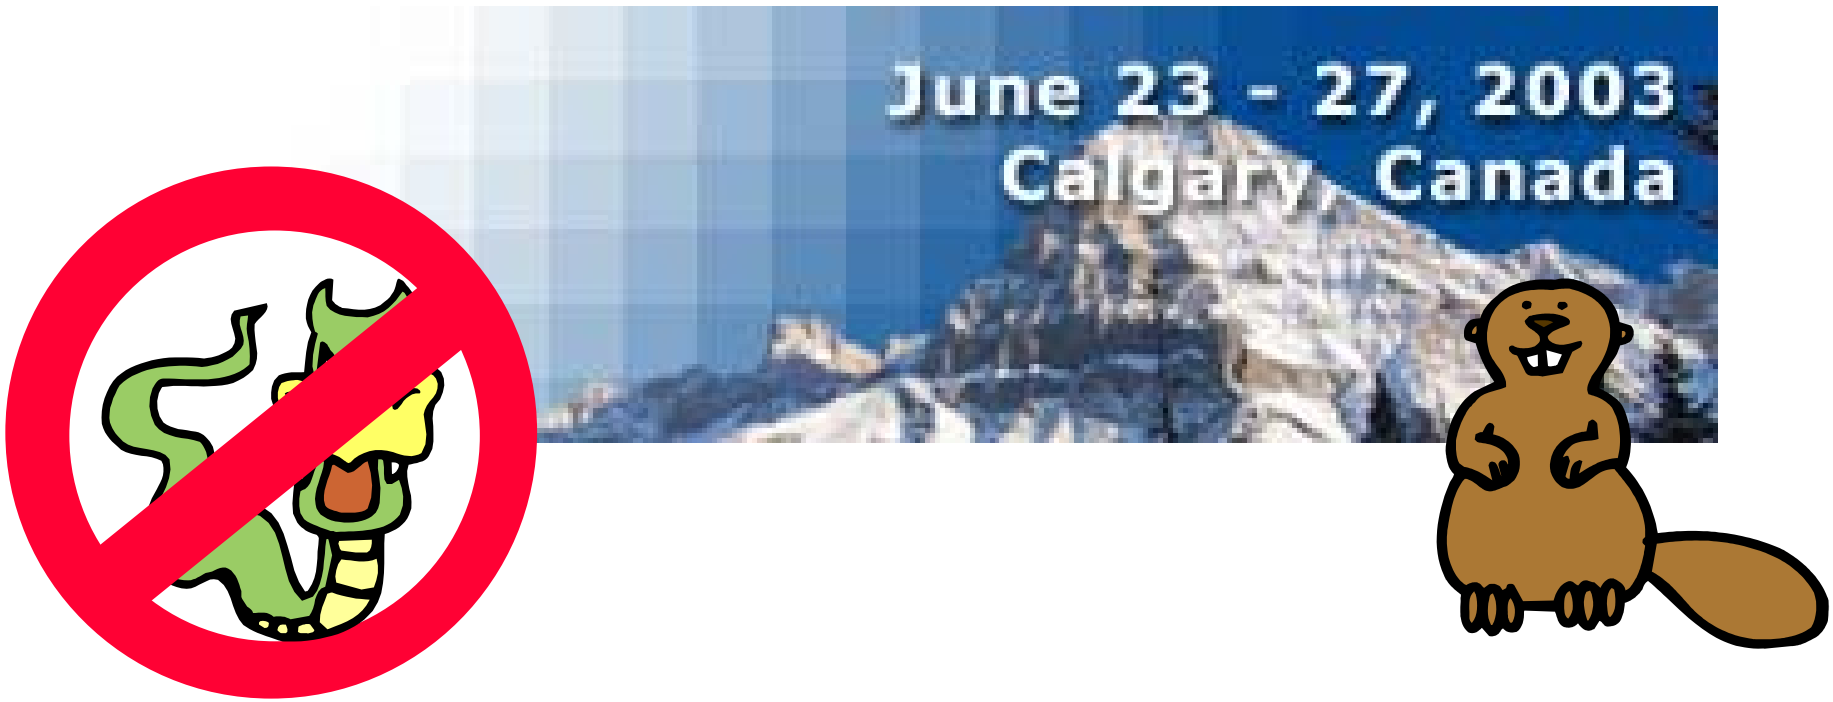

Institute for Quantum Information

John Preskill, Caltech

26 June 2003

http://www.iqi.caltech.edu/

Theoretical Quantum Information Science

is driven by ...

Three ***Great*** Ideas:

1) Quantum Computation
2) Quantum Error Correction
3) Quantum Cryptography

Quantum Error Correction

Shor ‘95

Steane ‘95

Quantum information can be protected, and processed fault-tolerantly.

Shor ‘95

Steane ‘95

Quantum error correction

Can large-scale quantum computers really be built and operated? Surely there are daunting technical challenges to be overcome. But are there obstacles *in principle* that might prevent us from ever attacking hard computational problems with quantum computers?

What comes to mind is the problem of *errors.* Quantum computers will be far more susceptible to error than conventional digital computers. A particular challenge is to prevent *decoherence* due to interactions of the computer with the environment. Even aside from decoherence, the unitary quantum gates will not be perfect, and small imperfections will accumulate over time...

Quantum factoring

For example, suppose we would like to factor a 200 digit number (which can't be done with today's classical computers). The quantum factoring algorithm requires a few thousand qubits and a few billion quantum gates. Suppose that in each gate, there is a probability p of a serious error due to an interaction with the environment. Then for the algorithm to have a good probability of success, we require

$$p \cdot 10^3 \cdot 10^9 \leq O(1) \quad \Rightarrow \quad p \leq 10^{-12}.$$

This is a very severe limitation!

Our confidence that large-scale quantum computations will someday be possible has been bolstered by the development of quantum error correction --- much larger error probabilities can be tolerated.

Quantum error correction

1. Error models and error correction
2. Quantum error-correcting codes
3. Stabilizer codes
4. 5-qubit code and 7-qubit code
5. Fault-tolerant quantum computation
6. Accuracy threshold

Quantum cryptography

1. Cryptography and security
2. Quantum key distribution
3. The BB84 (four-state) protocol
4. Security proof using QECC
5. Quantum coin flipping

Errors

The most general type of error acting on n qubits can be expressed as a unitary transformation acting on the qubits and their environment:

$$U : |\psi\rangle \otimes |0\rangle_E \rightarrow \sum_a E_a |\psi\rangle \otimes |a\rangle_E$$

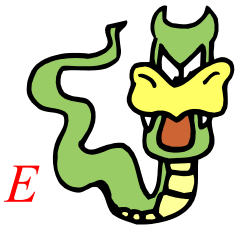

The states $|a\rangle_E$ of the environment are neither normalized nor mutually orthogonal. The operators $\{E_a\}$ are a basis for operators acting on n qubits, conveniently chosen to be "Pauli operators": $\{I, X, Y, Z\}^{\otimes n}$,

where

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

The errors could be "unitary errors" if $|a\rangle_E = C_a |0\rangle_E$ or decoherence errors if the states of the environment are mutually orthogonal.

Errors

$$U : |\psi\rangle \otimes |0\rangle_E \rightarrow \sum_a E_a |\psi\rangle \otimes |a\rangle_E$$

Our objective is to recover the (unknown) state $|\psi\rangle$ of the quantum computer. We can't expect to succeed for arbitrary errors, but we might succeed if the errors are of a restricted type. In fact, since the interactions with the environment are *local*, it is reasonable to expect that the errors are not too strongly correlated.

Define the "weight" w of a Pauli operator to be the number of qubits on which it acts nontrivially; that is $X, Y,$ or Z is applied to w of the qubits, and I is applied to n-w qubits. If errors are weakly correlated (and rare), then Pauli operators E_a with large weight have small amplitude $\| |a\rangle_E \|$.

Error recovery

We would like to devise a recovery procedure that acts on the data and an *ancilla*:

$$V : E_a \,|\psi\rangle \otimes |0\rangle_A \rightarrow |\psi\rangle \otimes |a\rangle_A$$

which works for any $E_a \in \{\text{Pauli operators of weight} \le t\}$.

Then we say that we can "correct t errors" in the block of n qubits. Information about the error that occurred gets transferred to the ancilla and can be discarded:

$$|\psi\rangle \otimes |0\rangle_E \otimes |0\rangle_A \xrightarrow{\text{error}} \sum_a E_a \,|\psi\rangle \otimes |a\rangle_E \otimes |0\rangle_A$$

$$\xrightarrow{\text{recover}} \sum_a |\psi\rangle \otimes |a\rangle_A \otimes |a\rangle_E = |\psi\rangle \otimes |\varphi\rangle_{EA}$$

Error recovery

$$|\psi\rangle\otimes|0\rangle_E\otimes|0\rangle_A \xrightarrow{\text{error}} \sum_a E_a|\psi\rangle\otimes|a\rangle_E\otimes|0\rangle_A$$

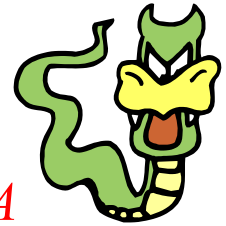

$$\xrightarrow{\text{recover}} \sum_a |\psi\rangle\otimes|a\rangle_A\otimes|a\rangle_E = |\psi\rangle\otimes|\varphi\rangle_{EA}$$

Entropy

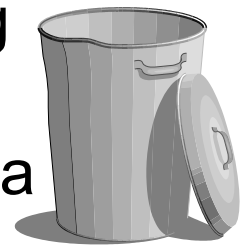

Errors entangle the data with the environment, producing *decoherence*. Recovery transforms entanglement of the data with the environment into entanglement of the ancilla with the environment, "purifying" the data. Decoherence is thus reversed. Entropy introduced in the data is transferred to the ancilla and can be discarded --- we "refrigerate" the data at the expense of "heating" the ancilla. If we wish to erase the ancilla (cool it to $T \approx 0,$ so that we can use it again) we need to pay a power bill.

Quantum error-correcting code

We won't be able to correct all errors of weight up to t for arbitrary states $|\psi\rangle \in \mathfrak{H}_{n \text{ qubits}}$. But perhaps we can succeed for states contained in a *code subspace* of the full Hilbert space,

$$\mathfrak{H}_{\text{code}} \in \mathfrak{H}_{n \text{ qubits}}.$$

If the code subspace has dimension 2^k, then we say that k encoded qubits are embedded in the block of n qubits.

How can such a code be constructed? It will *suffice* if

$$\left\{E_a \mathfrak{H}_{\text{code}}, \quad E_a \in \left\{\text{Pauli operators of weight} \le t\right\}\right\}$$

are mutually orthogonal.

If so, then it is possible in principle to perform an (incomplete) orthogonal measurement that determines the error E_a (without revealing any information about the encoded state). We recover by applying the unitary transformation E_a.

2-qubit "code"

The key concept in quantum coding theory is the *stabilizer* of a state. E.g., $|\phi^+\rangle = \left(|00\rangle_{AB} + |11\rangle_{AB}\right)/\sqrt{2}$ is the simultaneous eigenstate with eigenvalue 1 of two commuting Pauli operators:

$$M_X = X \otimes X, \quad M_Z = Z \otimes Z.$$

These two conditions on two qubits determine a one-dimensional subspace, i.e., a unique state.

Suppose that Bob's qubit is protected, but Alice's qubit might have been damaged. Can we diagnose the damage?

The space of possible errors is spanned by $\mathcal{E} = \{I, X, Y, Z\} \otimes I$

By measuring the 2 generators of the code's *stabilizer group*, we can distinguish all possible errors acting on Alice's qubit.

2-qubit "code"

$$|\phi^+\rangle = \frac{1}{\sqrt{2}}\left(|00\rangle_{AB} + |11\rangle_{AB}\right)$$

$$M_X = X \otimes X$$
$$M_Z = Z \otimes Z$$

All of the errors in $\mathcal{E} = \{I, X, Y, Z\} \otimes I$ anticommute with the stabilizer generators in distinguishable ways.

	M_X	M_Z
$I \otimes I$	+	+
$X \otimes I$	+	-
$Y \otimes I$	-	-
$Z \otimes I$	-	+

By measuring the 2 generators of the code's *stabilizer group*, we can distinguish all possible errors acting on Alice's qubit.

4-qubit code

$M_X = X\,X\,X\,X$
$M_Z = Z\,Z\,Z\,Z$

There are $n - k = 4 - 2 = 2$ encoded qubits.

	M_X	M_Z
$IIII$	+	+
$XIII$	+	-
$YIII$	-	-
$ZIII$	-	+

The 4-dimensional code space is spanned by:

$$|0000\rangle + |1111\rangle$$
$$|0011\rangle + |1100\rangle$$
$$|0101\rangle + |1010\rangle$$
$$|0110\rangle + |1001\rangle$$

(An X changes the parity, a Z changes the relative phase, a Y does both..)

Suppose that one qubit out of the four is damaged, and we know which one, but we don't know the nature of the damage. By measuring the two stabilizer generators, we can distinguish among I,X,Y,Z acting on (e.g.) the first qubit.

General stabilizer codes

Operators $M_1, M_2, \ldots, M_{n-k}$ are mutually commuting Pauli operators, $M_i^2 = I$, which generate an abelian subgroup S of the "Pauli group." S is the code's stabilizer group.

A vector $|\psi\rangle$ in the *n*-qubit Hilbert space is in the code subspace iff $M|\psi\rangle = |\psi\rangle$ for all M in S .

The dimension of the code space is: (There are *k* encoded qubits.)

$$2^n \left(\frac{1}{2^{n-k}} \right) = 2^k.$$

The code can correct t errors if each (nontrivial) Pauli operator of weight up to $2t$ *anticommutes* with at least one of the stabilizer generators:

$$\langle \psi | E_a^\dagger E_b | \psi \rangle = \langle \psi | E_a^\dagger E_b M_i | \psi \rangle = -\langle \psi | M_i E_a^\dagger E_b | \psi \rangle$$

$$= -\langle \psi | E_a^\dagger E_b | \psi \rangle \Rightarrow E_a \mathcal{H}_{\text{code}} \perp E_b \mathcal{H}_{\text{code}} \text{ for weight}(E_a) \le t$$

All errors up to weight t are distinguishable (and correctable).

5-qubit code

Suppose we would like to encode $k=1$ protected qubits in a block of n qubits, and be able to correct all weight-1 errors. How large must n be?

There are two mutually orthogonal "codewords" $|\overline{0}\rangle, |\overline{1}\rangle$ that span the code subspace. Furthermore all $E_a|\overline{0}\rangle, E_b|\overline{1}\rangle$ should be mutually orthogonal.

There are $3\times5+1=16$ Pauli operators of weight ≤ 1, and the Hilbert space of 5 qubits has dimension $2^5=32$. Therefore, for $n=5$, there is just barely enough room: $16\times2 \leq 2^5=32$.

To see that the code really exists, we can construct it explicitly.

5-qubit code

The code is the simultaneous eigenspace with eigenvalue 1 of 4 commuting *check operators (stabilizer generators):*

$$M_1 = X\,Z\,Z\,X\,I = +1$$
$$M_2 = I\,X\,Z\,Z\,X = +1$$
$$M_3 = X\,I\,X\,Z\,Z = +1$$
$$M_4 = Z\,X\,I\,X\,Z = +1$$

All of these stabilizer generators square to I; they are mutually commuting because there are two collisions between X and Z.

The other three generators are obtained from the first by cyclic permutations. (Note that $M_5 = Z\,Z\,X\,I\,X = M_1 M_2 M_3 M_4$ is not independent.) Therefore, the code is cyclic (cyclic permutations of the qubits preserve the code space).

Claim: no Pauli operator E of weight 1 or 2 commutes with all of the check operators. Weight 1: each column contains an X and a Z. Weight 2: Because the code is cyclic, it suffices to consider $??I\,I\,I$ and $?I\,?I\,I$ ….

5-qubit code

-- *k=1* protected qubit

-- corrects *t=1* error

The code is the simultaneous eigenspace with eigenvalue 1 of 4 commuting *check operators:*

$$M_1 = X\,Z\,Z\,X\,I = +1$$
$$M_2 = I\,X\,Z\,Z\,X = +1$$
$$M_3 = X\,I\,X\,Z\,Z = +1$$
$$M_4 = Z\,X\,I\,X\,Z = +1$$

By these operators, we can distinguish all possible weight-one errors. Each “syndrome” points to a unique Pauli operator of weight 0 or 1.

	M_1	M_2	M_3	M_4
X_1	+	+	+	-
Y_1	-	+	-	-
Z_1	-	+	-	+
X_2	-	+	+	+
Y_2	-	-	+	-
Z_2	+	-	+	-
X_3	-	-	+	+
Y_3	-	-	-	+
Z_3	+	+	-	+
X_4	+	-	-	+
Y_4	-	-	-	-
Z_4	-	+	+	-
X_5	+	+	-	-
Y_5	+	-	-	-
Z_5	+	-	+	+
I	+	+	+	+

5-qubit code

$$M_1 = X\,Z\,Z\,X\,I$$
$$M_2 = I\,X\,Z\,Z\,X$$
$$M_3 = X\,I\,X\,Z\,Z$$
$$M_4 = Z\,X\,I\,X\,Z$$

How do we measure the stabilizer generators without destroying the encoded state?

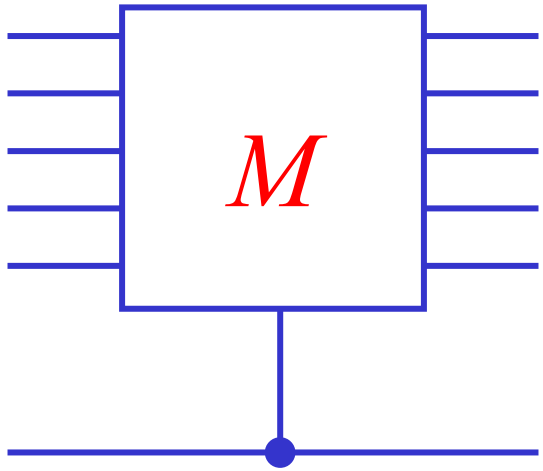

Apply M conditioned on value of an ancilla qubit.

X=1 Eigenstate: $|0\rangle_A + |1\rangle_A$ $\rightarrow$ $|0\rangle_A + M\,|1\rangle_A$

Measure X

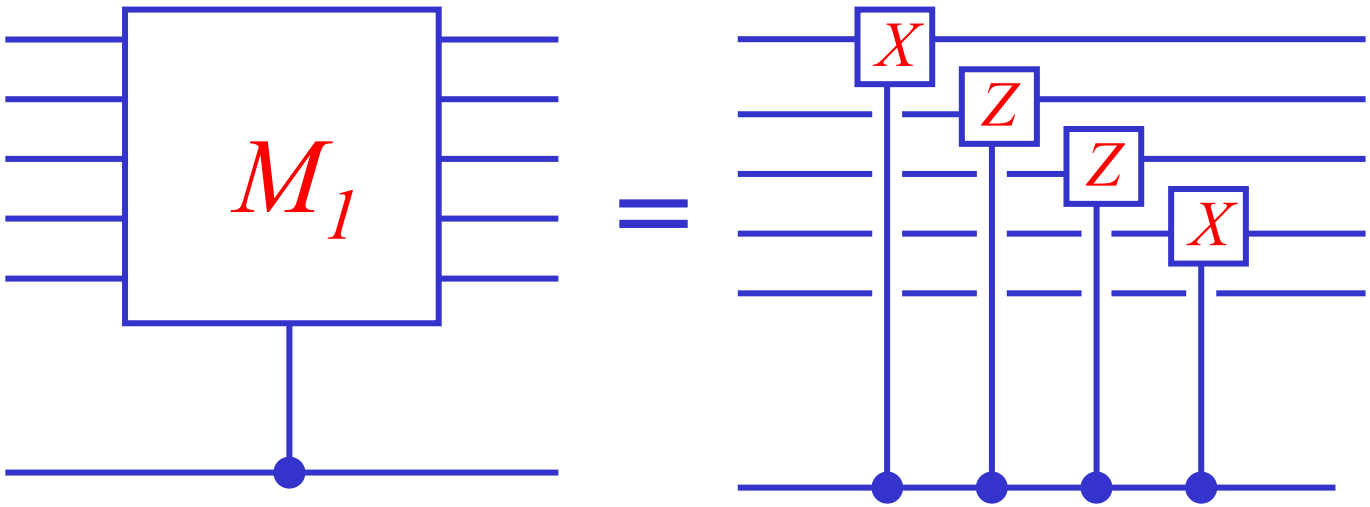

5-qubit code

What are the *encoded operations* that preserve the code space and act nontrivially within it?

$$M_1 = X\,Z\,Z\,X\,I$$
$$M_2 = I\,X\,Z\,Z\,X$$
$$M_3 = X\,I\,X\,Z\,Z$$
$$M_4 = Z\,X\,I\,X\,Z$$

We may choose them to be:

$$\bar{Z} = Z\,Z\,Z\,Z\,Z$$
$$\bar{X} = X\,X\,X\,X\,X$$

(which anticommute with one another and commute with the stabilizer).

General stabilizer code

The code stabilizer S is an abelian subgroup of order 2^{n-k} of the the n-qubit Pauli group. Its *dual* or *normalizer* $S^{\perp}$ is the subgroup of the Pauli group containing all Pauli operators that commute with S (thus $S^{\perp}$ contains S). The encoded operations are the coset space $S^{\perp} / S$. The minimum weight of $S^{\perp} \setminus S$ is the *distance* of the code. A code with distance $d=2t+1$ can correct t errors.

7-qubit code

$$
\begin{aligned}
M_{Z,1} &= Z\ I\ Z\ I\ Z\ I\ Z \\
M_{Z,2} &= Z\ Z\ I\ I\ Z\ Z\ I \\
M_{Z,3} &= Z\ Z\ Z\ Z\ I\ I\ I
\end{aligned}
$$

Corrects the bit-flip (X) errors. The three-bit string $(M_{Z,1}, M_{Z,2}, M_{Z,3})$ (if nonzero) points to the position of the error.

$$
\begin{aligned}
M_{X,1} &= X\ I\ X\ I\ X\ I\ X \\
M_{X,2} &= X\ X\ I\ I\ X\ X\ I \\
M_{X,3} &= X\ X\ X\ X\ I\ I\ I
\end{aligned}
$$

Corrects the phase (Z) errors. The three-bit string $(M_{X,1}, M_{X,2}, M_{X,3})$ (if nonzero) points to the position of the error.

The $\boldsymbol{M_Z}$'s commute with the $\boldsymbol{M_X}$'s , because each row of the $\boldsymbol{M_Z}$ matrix has an even number of "collisions" with each row of the $\boldsymbol{M_X}$ matrix; i.e., the rows are orthogonal in the sense of linear algebra over the field $\mathbb{Z}_2$. Any two matrices with this property define a quantum code, which is said to be of the "CSS" (Calderbank-Shor-Steane) type. With CSS codes, the bit-flip and phase error correction can be executed separately. The encoded operations can be chosen to be $\overline{Z} = I\ I\ I\ I\ Z\ Z\ Z$, $\overline{X} = I\ I\ I\ I\ X\ X\ X$ which commute with the code stabilizer and are not contained in it.

7-qubit code generalized: *CSS codes*

$$
\begin{aligned}
M_{Z,1} &= Z\ I\ Z\ I\ Z\ I\ Z \\
M_{Z,2} &= Z\ Z\ I\ I\ Z\ Z\ I \\
M_{Z,3} &= Z\ Z\ Z\ Z\ I\ I\ I
\end{aligned}
$$

The matrix $\boldsymbol{M_Z}$ is the *parity check matrix* of a classical code $\boldsymbol{C_Z}$: its codewords are binary strings annihilated by $\boldsymbol{M_Z}$.

$$
\begin{aligned}
M_{X,1} &= X\ I\ X\ I\ X\ I\ X \\
M_{X,2} &= X\ X\ I\ I\ X\ X\ I \\
M_{X,3} &= X\ X\ X\ X\ I\ I\ I
\end{aligned}
$$

The matrix $\boldsymbol{M_X}$ is the *generator matrix* of a classical code $\boldsymbol{C_X}^{\perp}$: its codewords are linear combinations of the rows of $\boldsymbol{M_X}$.

The classical code $\boldsymbol{C_X}^{\perp}$ is a subcode of $\boldsymbol{C_Z}$. Expressed in the ***Z*-basis,** a basis for the codewords of the CSS quantum code is:

$$|\bar{w}\rangle \propto \sum_{u \in C_X^{\perp}} |w+u\rangle\,, \quad w \in C_Z$$

There is a codeword associated with each *coset* of $\boldsymbol{C_X}^{\perp}$ in $\boldsymbol{C_Z}$. We use $\boldsymbol{C_Z}$ to diagnose the bit flip errors and $\boldsymbol{C_X}$ to diagnose the phase errors (in the conjugate basis).

Fault tolerance

- The measured error syndrome might be inaccurate.
- Errors might propagate during syndrome measurement.
- We need to implement a universal set of quantum gates that act on encoded quantum states, without unacceptable error propagation.
- We need codes that can correct many errors in the code block.

Error propagation in error correction

Error propagation in error correction

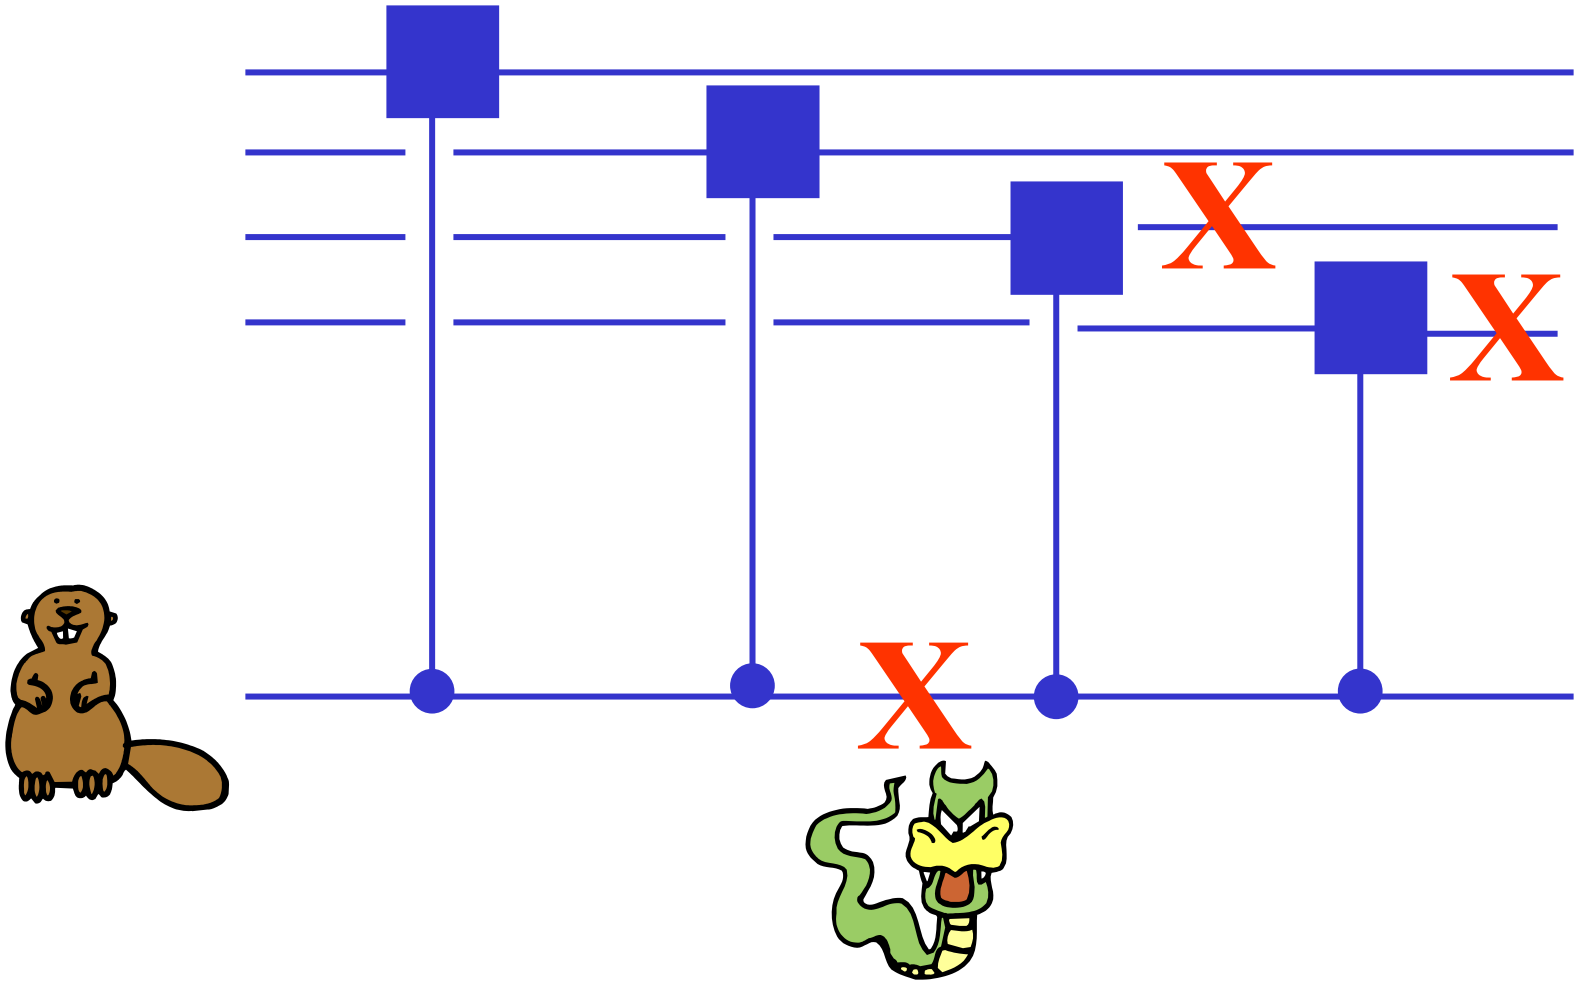

A single error due to the environment causes multiple errors in the data.

Preparation and consumption of *quantum software*

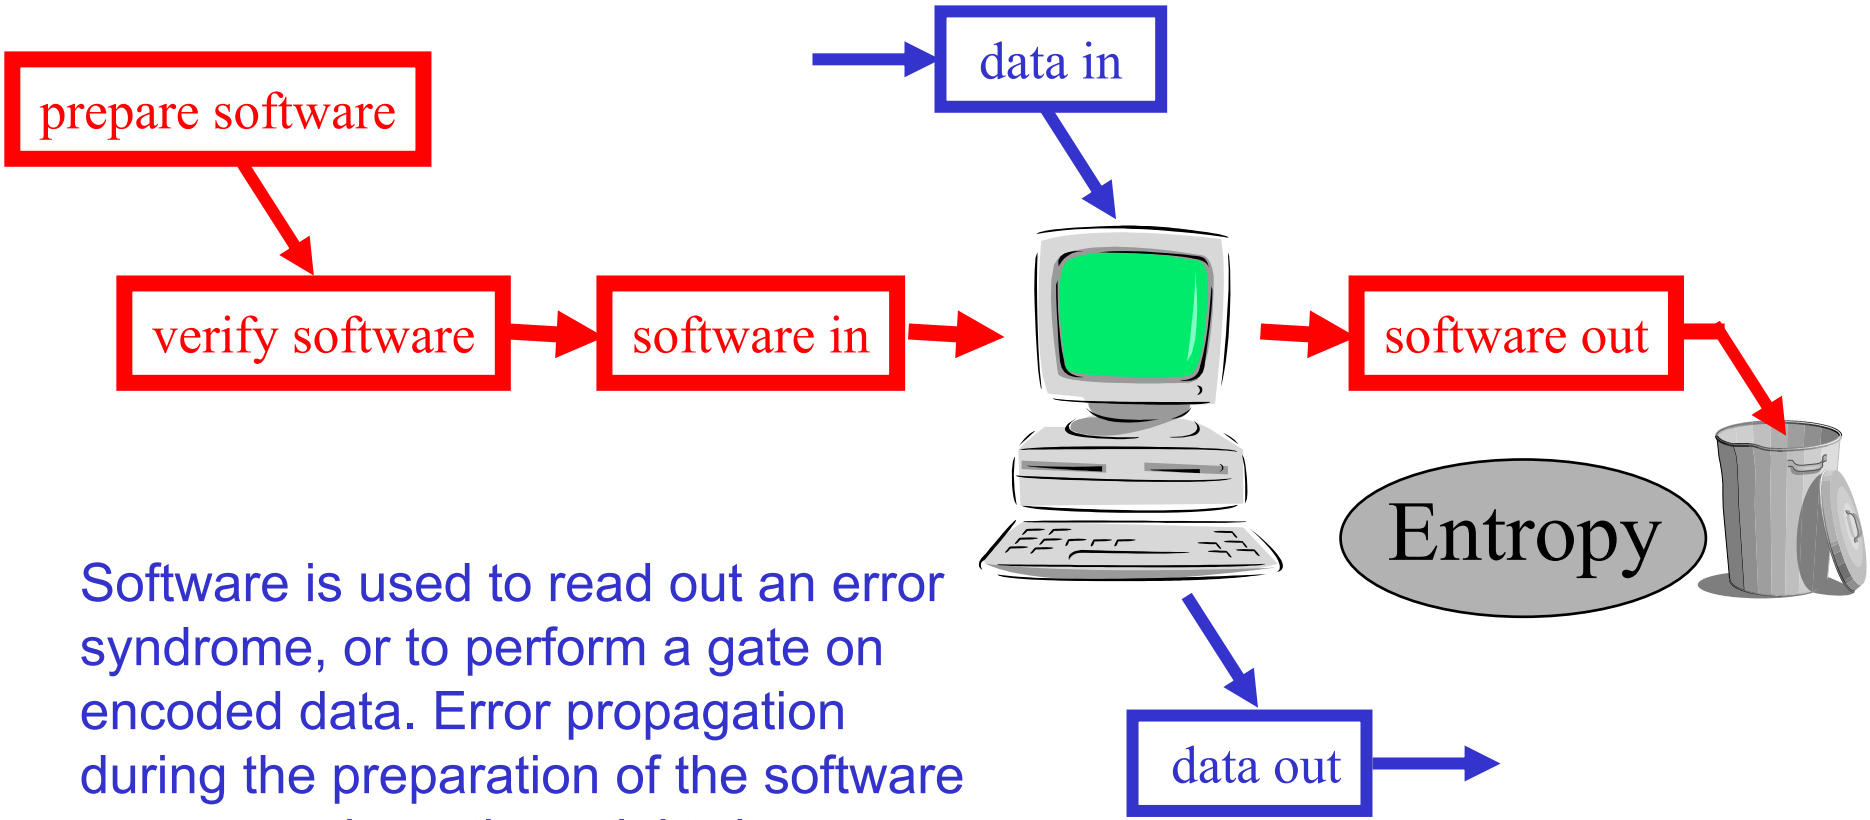

Software is used to read out an error syndrome, or to perform a gate on encoded data. Error propagation during the preparation of the software can cause *bugs* that might damage the data when the software is used.

Therefore the software must be checked and purified. After a single use, the software is irreparably damaged and must be discarded.

Consumption of *quantum software*

Gottesman & Chuang
quant-ph/9908010

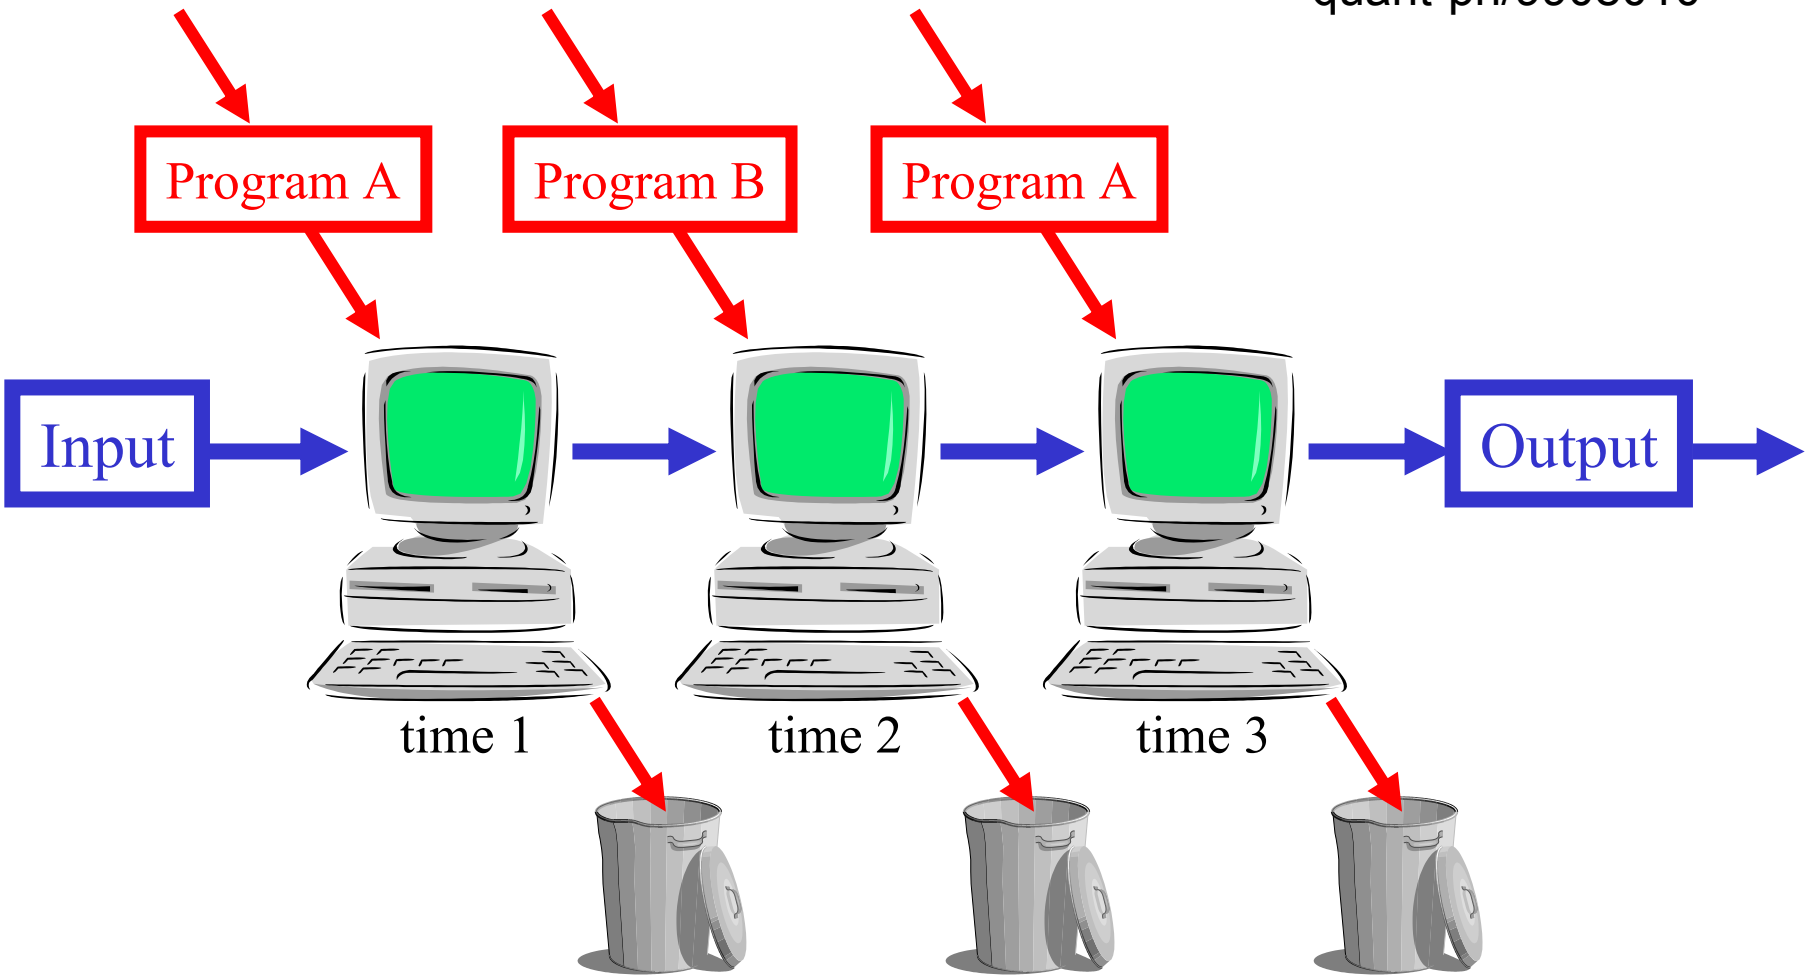

After a single use, the software is irreparably damaged and must be discarded. To execute a quantum algorithm, the user downloads and consumes a particular program many times.

Fault-Tolerant Error Correction

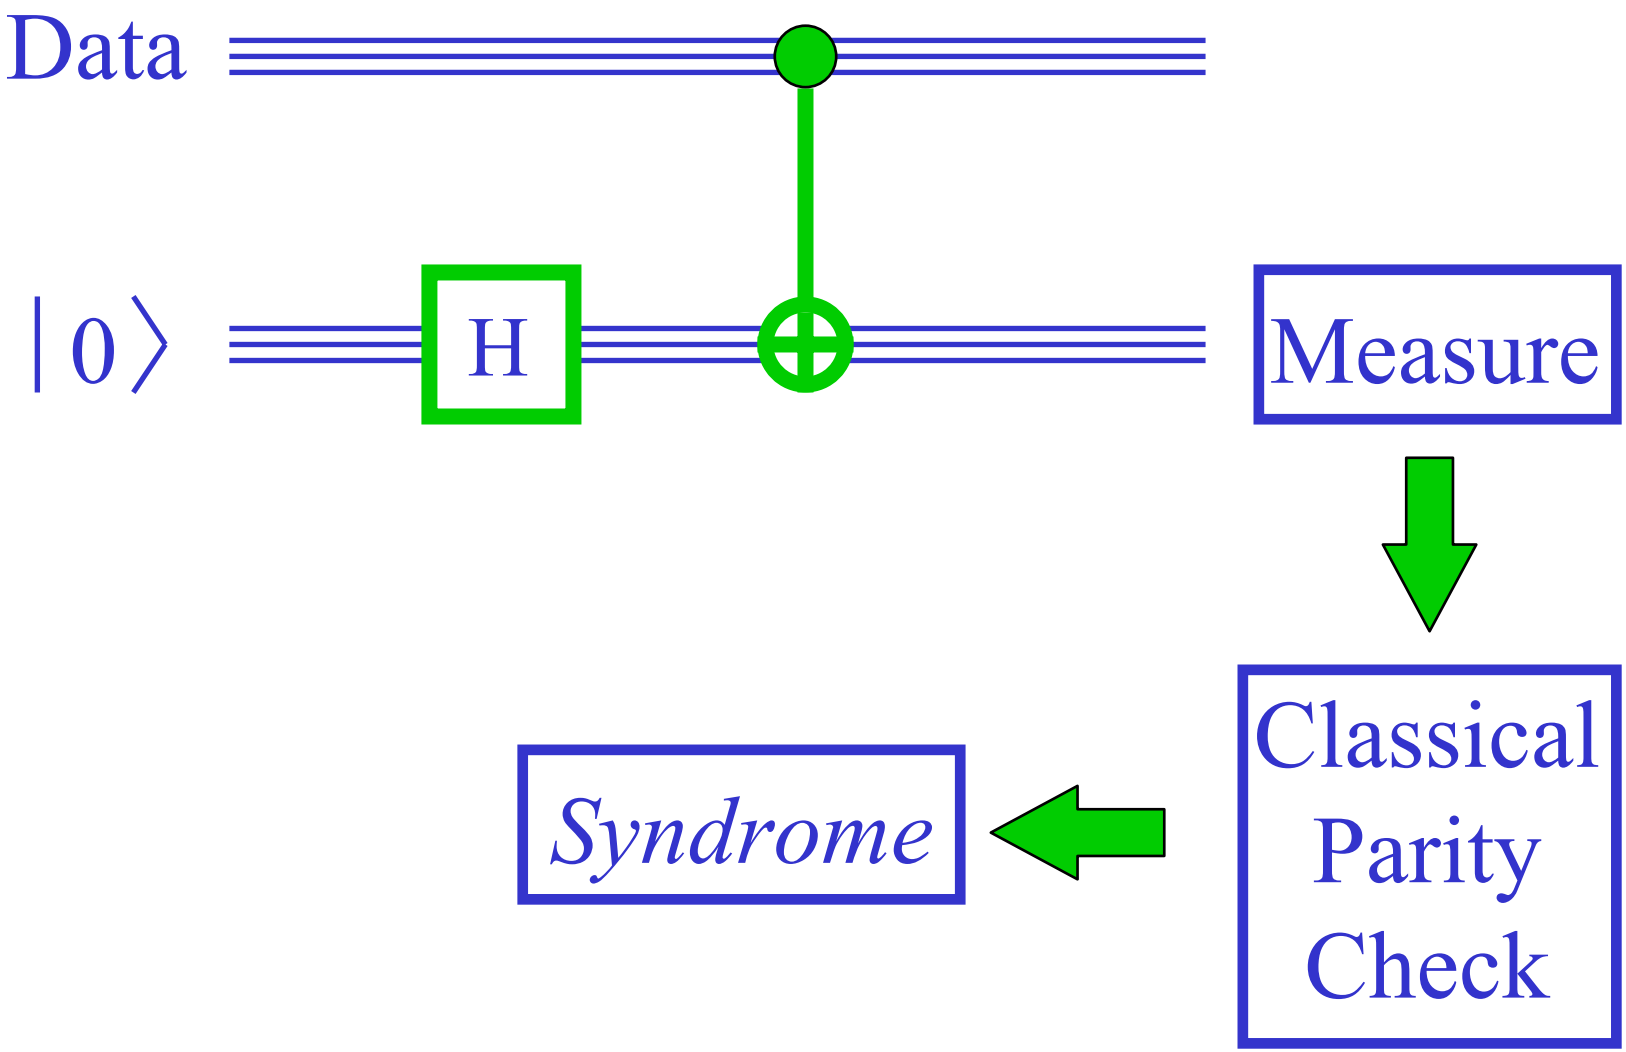

Fault-Tolerant Error Correction

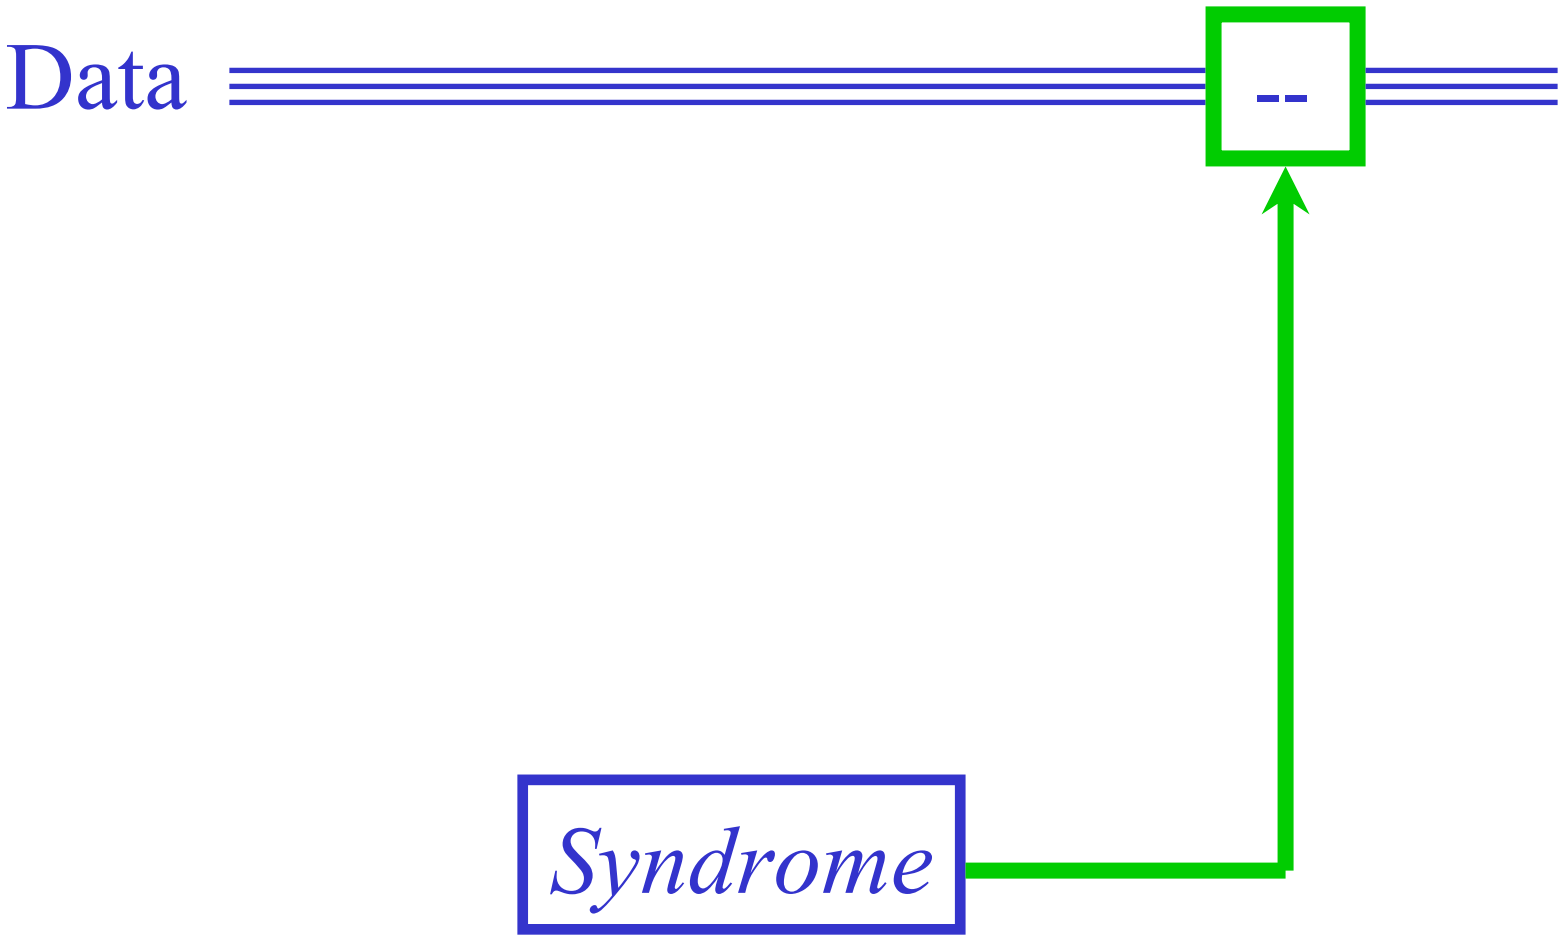

Fault-Tolerant Error Correction

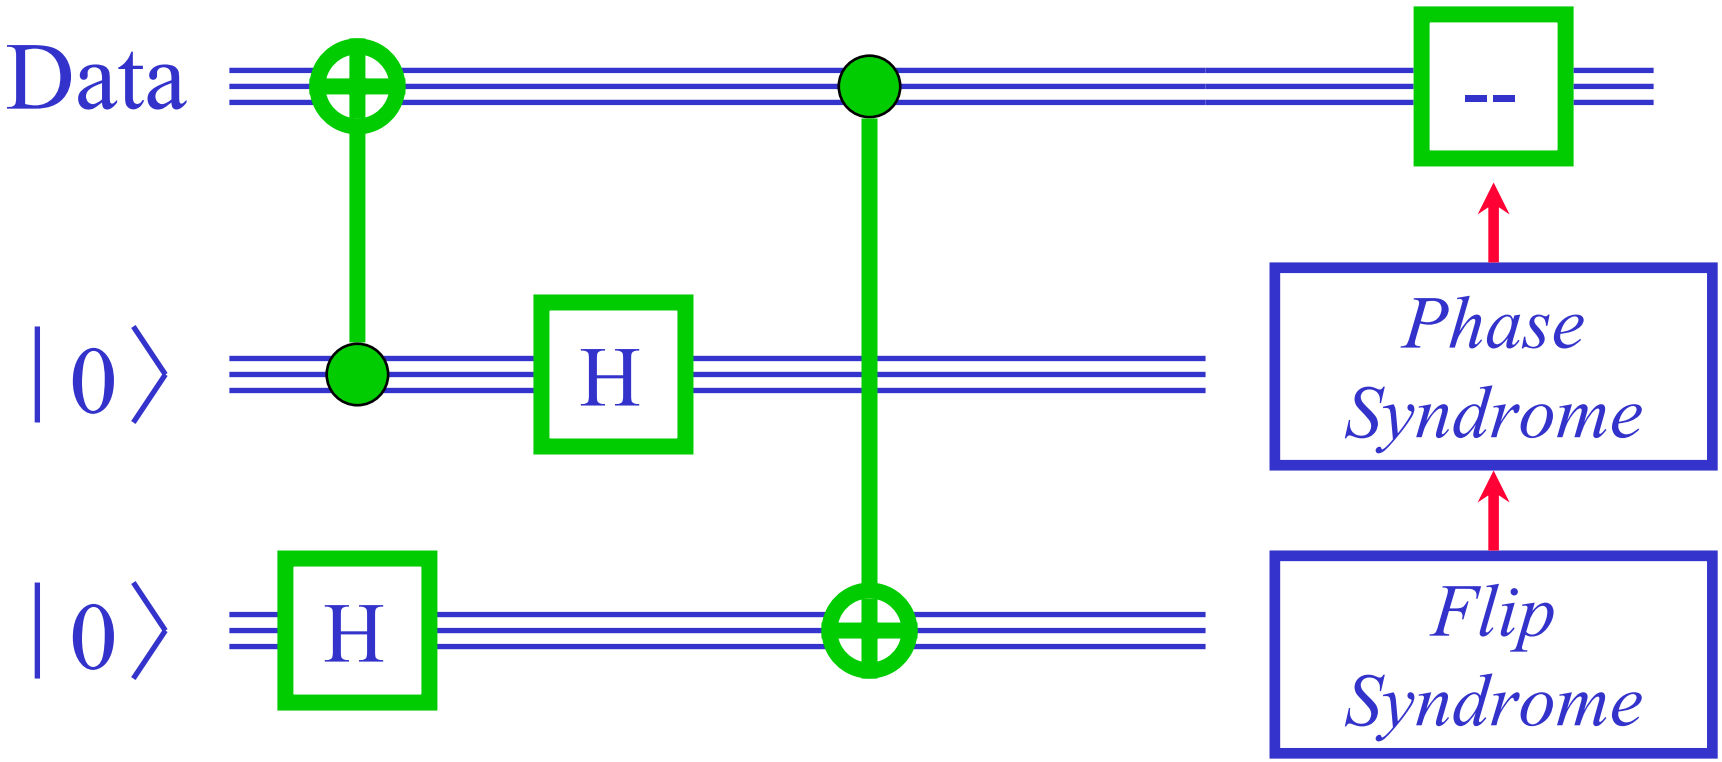

Fault-Tolerant Error Correction

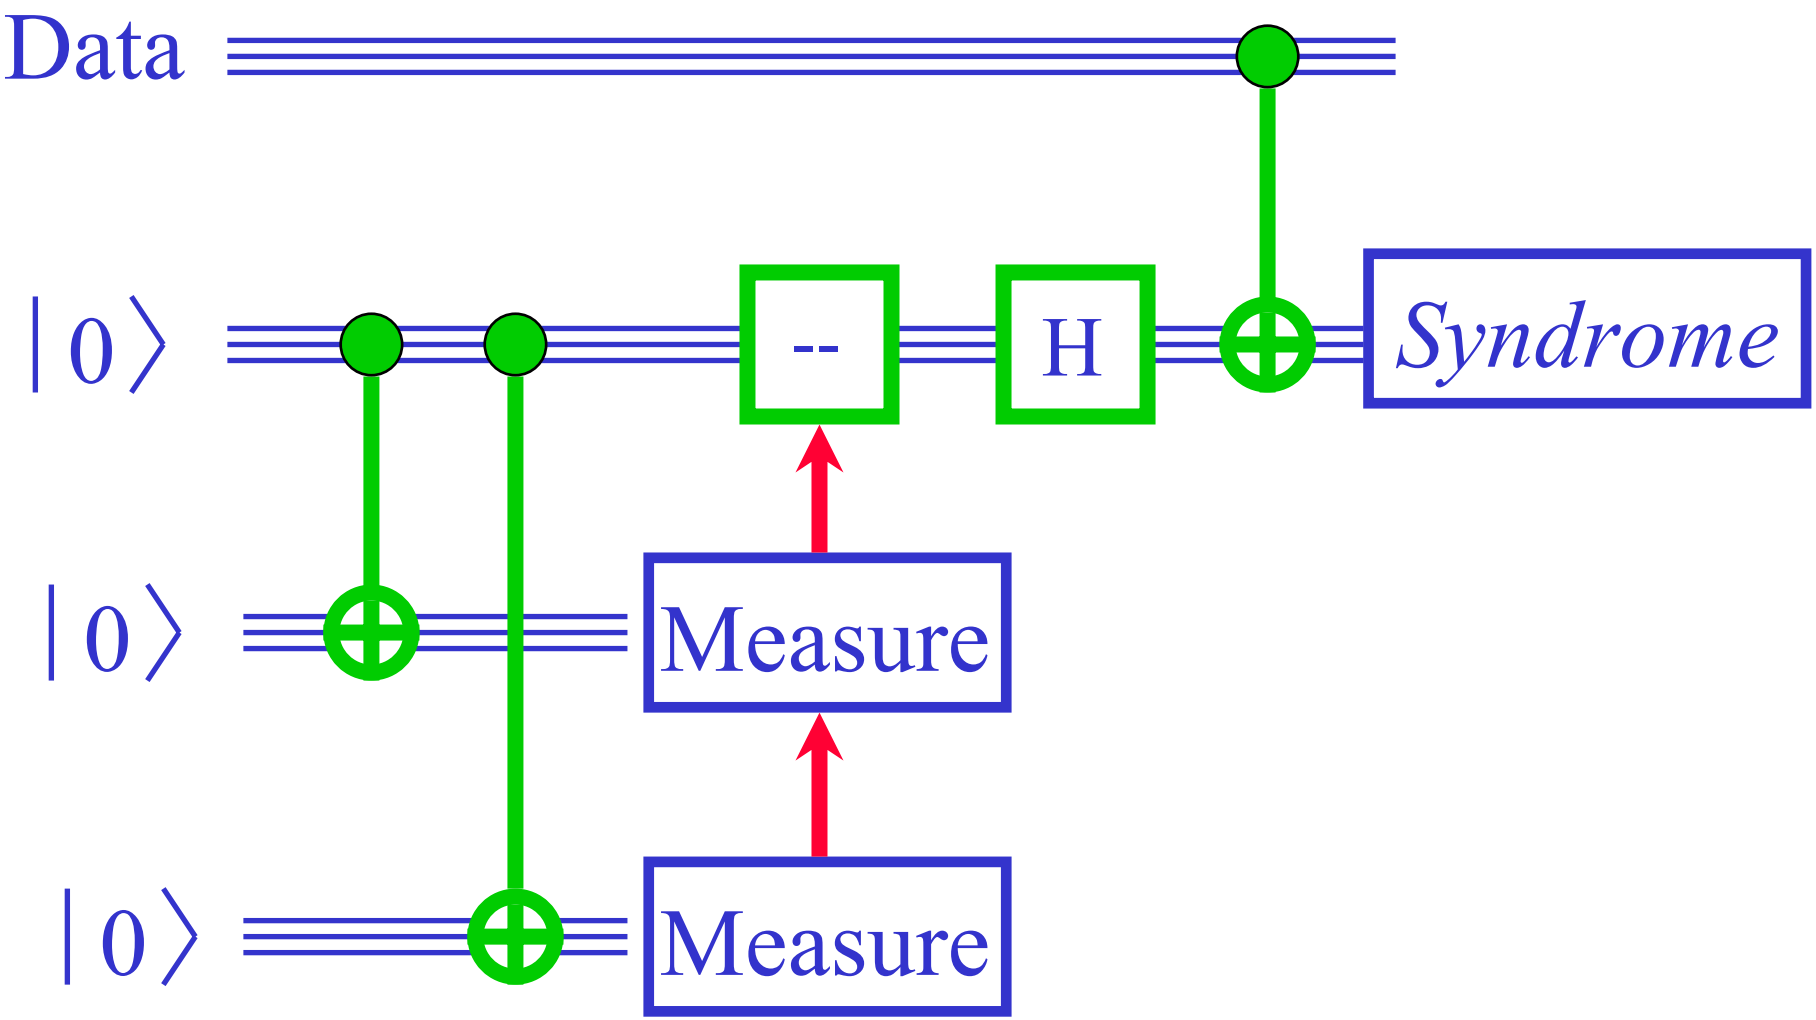

Fault-Tolerant Error Correction

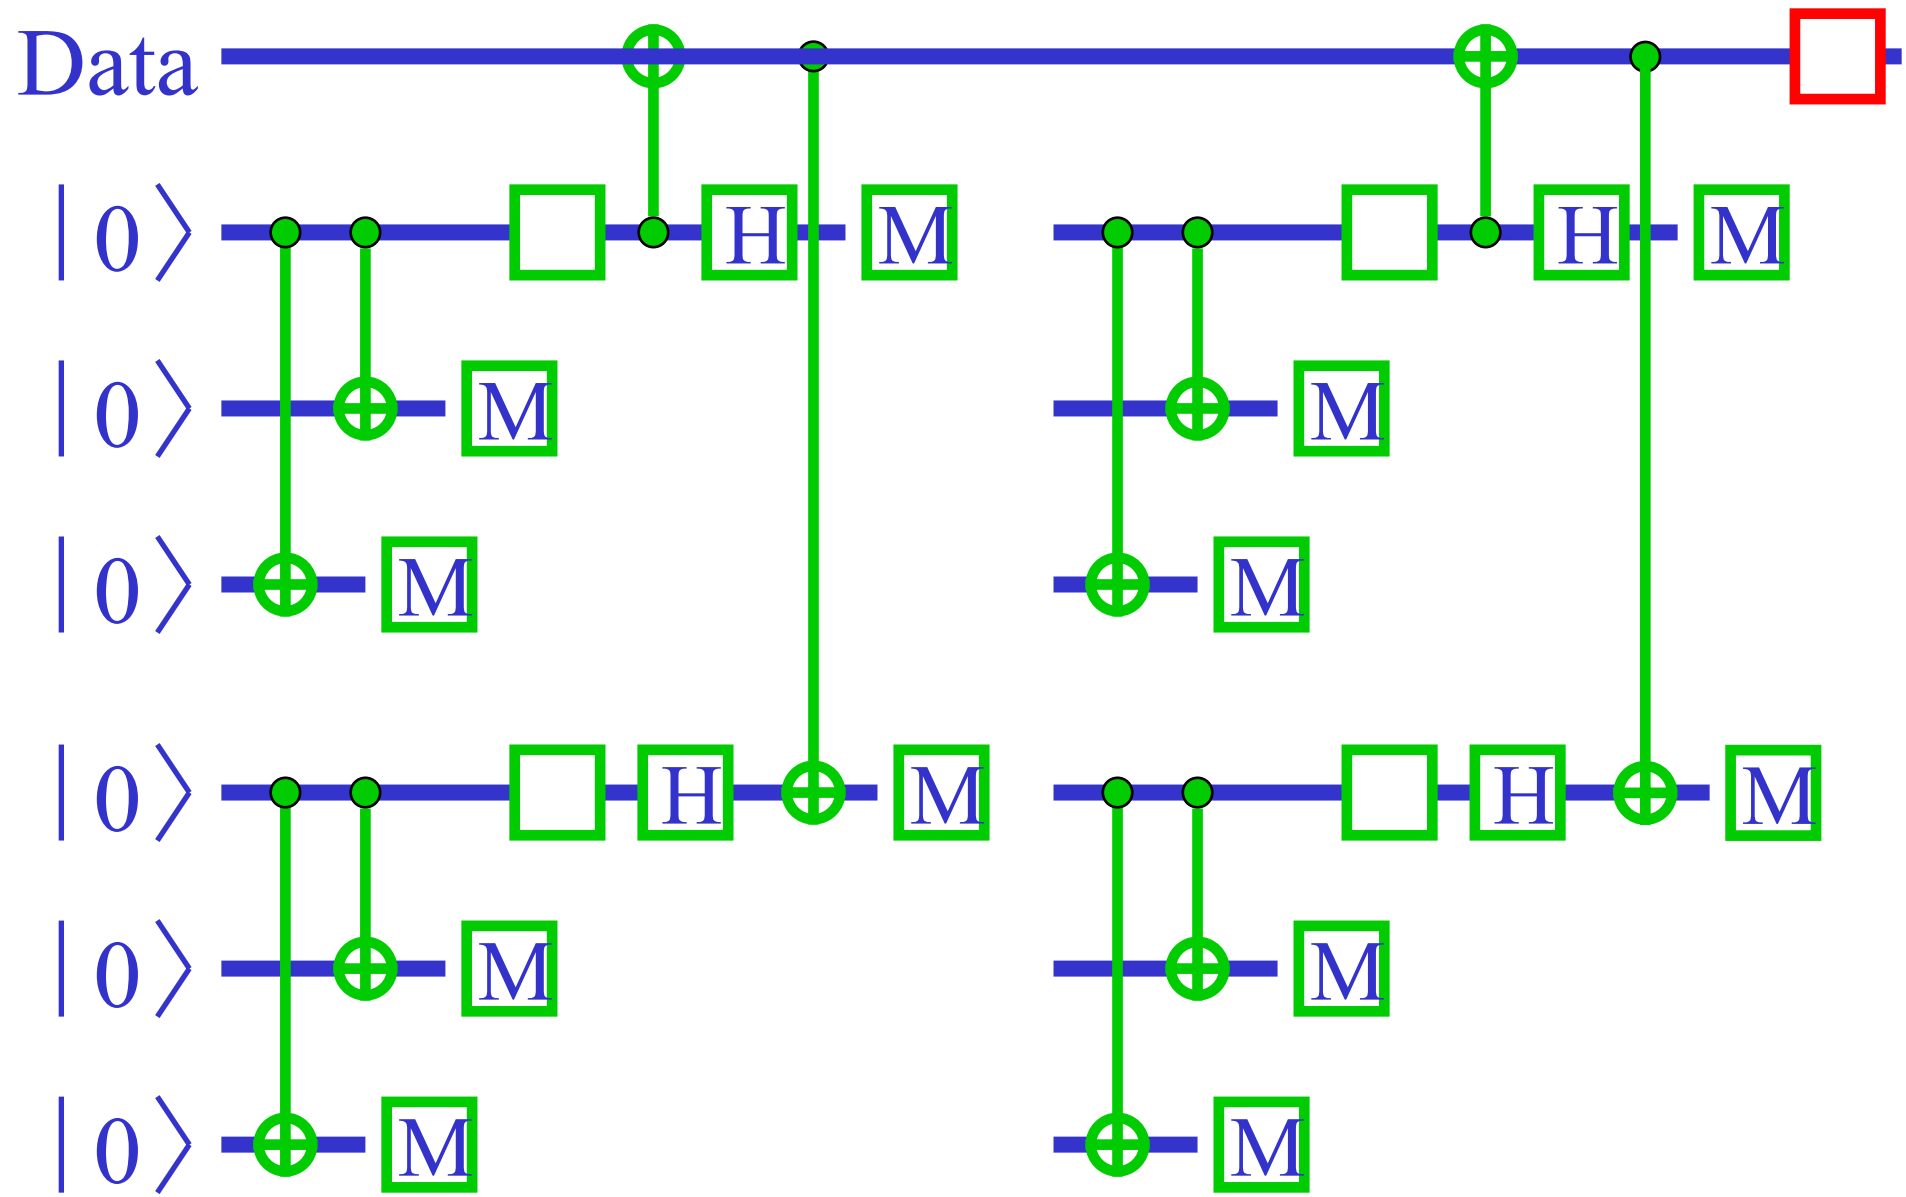

Transversal gates are fault tolerant:

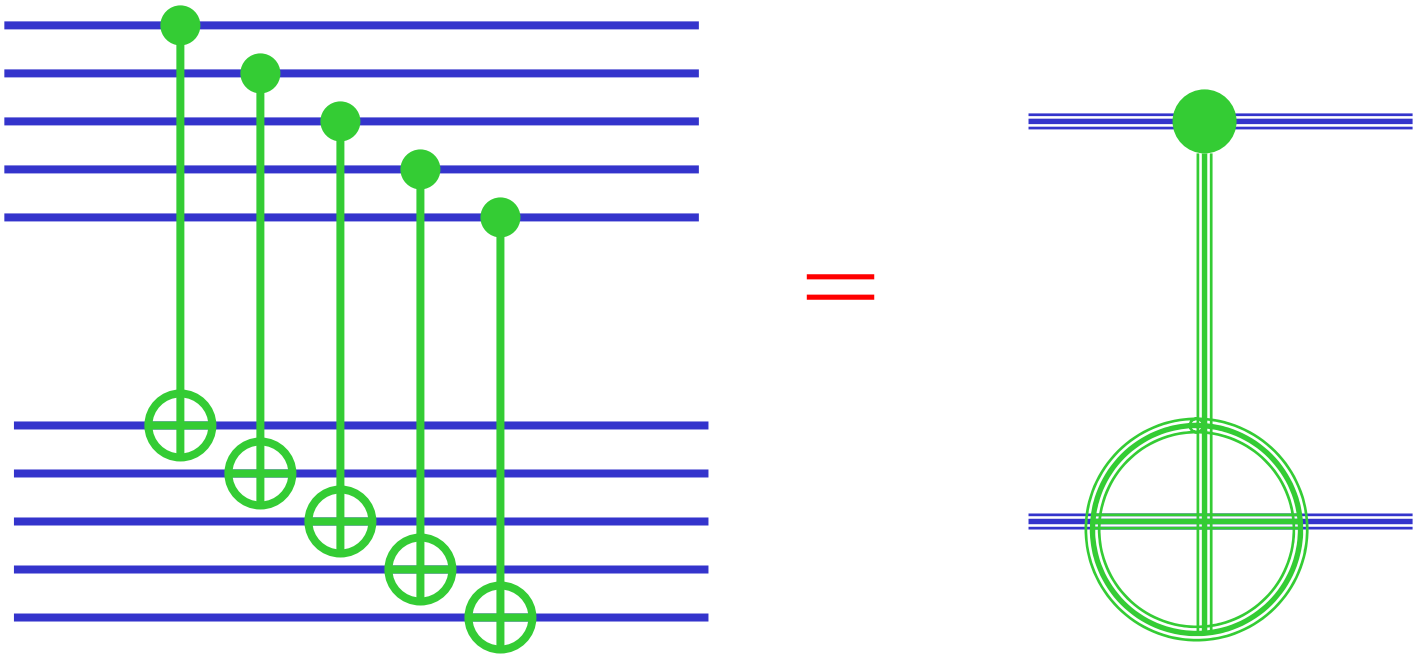

For codes of the CSS type, an encoded CNOT can be implemented by executing CNOTs between the corresponding qubits in the encoded blocks. This gate may propagate errors from one block to another, but not from one qubit to another in the same block.

Not all gates in a universal set can be realized transversally and fault tolerantly. To complete a universal fault-tolerant gate set, we must add to the transversal gates additional gates that are implemented by consuming *quantum software*.

Quantum software for the "π/8 gate"

$P = \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}$ is easy to implement.

$V = \sqrt{P} = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{i} \end{pmatrix}$ is hard and needed for universality.

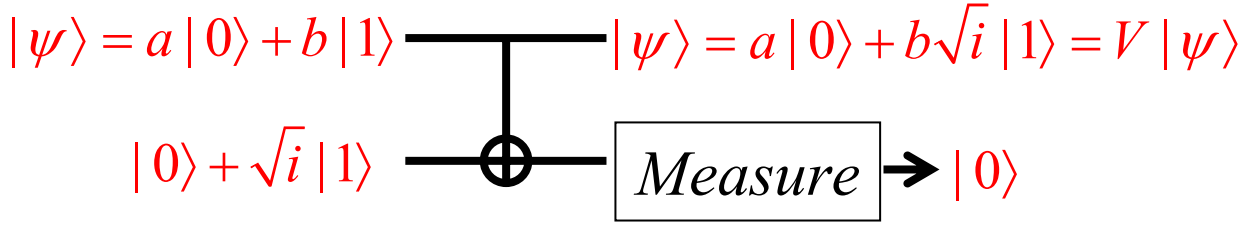

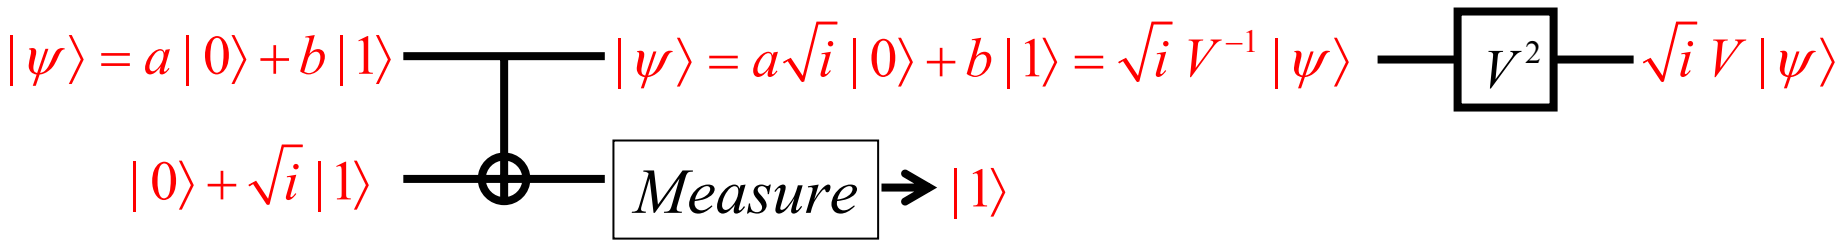

Performing the gate V is reduced to the task of preparing the software, which is achieved by measuring:

$\sqrt{-i}\,PX = \begin{pmatrix} 0 & \sqrt{-i} \\ \sqrt{i} & 0 \end{pmatrix}$

Preparing the software:

If a gate M can be *applied* transversally, then it can *measured* using a *cat state*:

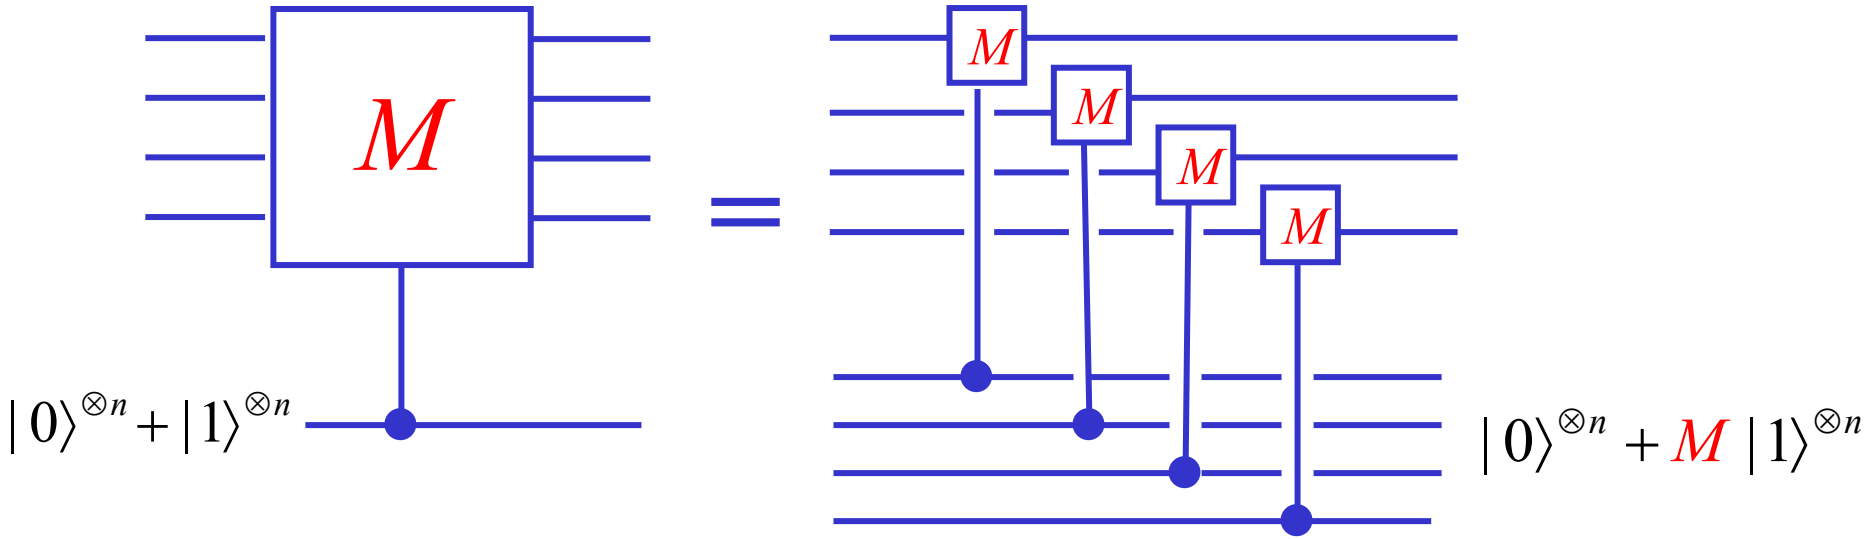

The phase of the cat is determined by measuring X of each qubit and computing the parity of the outcomes. The cat state should be verified before use, and the measurement should be repeated to improve reliability.

The teleportation trick (Gottesman-Chuang)

Suppose we teleport with a "twisted" entangled state:

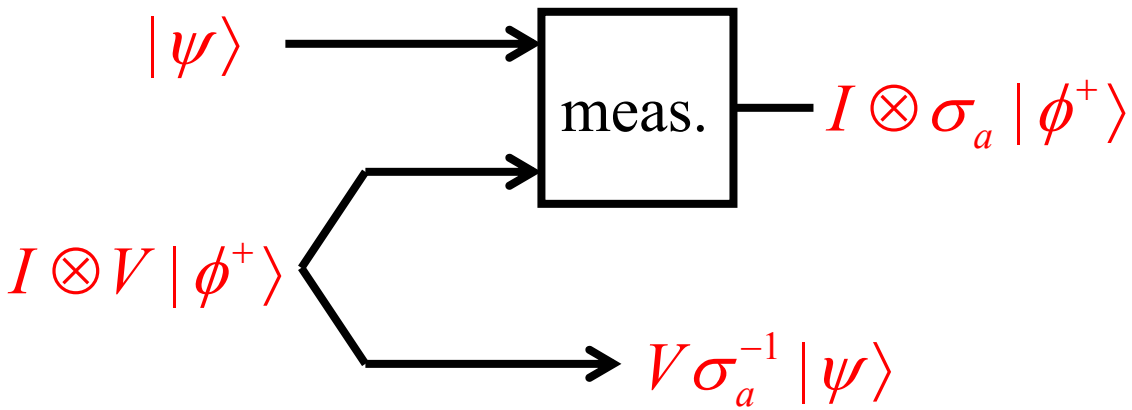

To recover the state that was destroyed by Alice's Bell measurement, Bob applies $\sigma_a V^{-1}$.

But if Bob applies $V\sigma_a V^{-1}$ instead, then he recovers the state $V|\psi\rangle$. For a given code, applying V might be hard, but Bell measurement and applying $V\sigma_a V^{-1}$ are easy. Then the problem of constructing the V gate reduces to the problem of preparing the state $I \otimes V |\phi^+\rangle$ --- the quantum software. We can ensure the quality of the software (before use) either by "distilling" a high-fidelity copy from several noisier copies, or by performing repeated measurements.

Toffoli (controlled-controlled-NOT) Gate

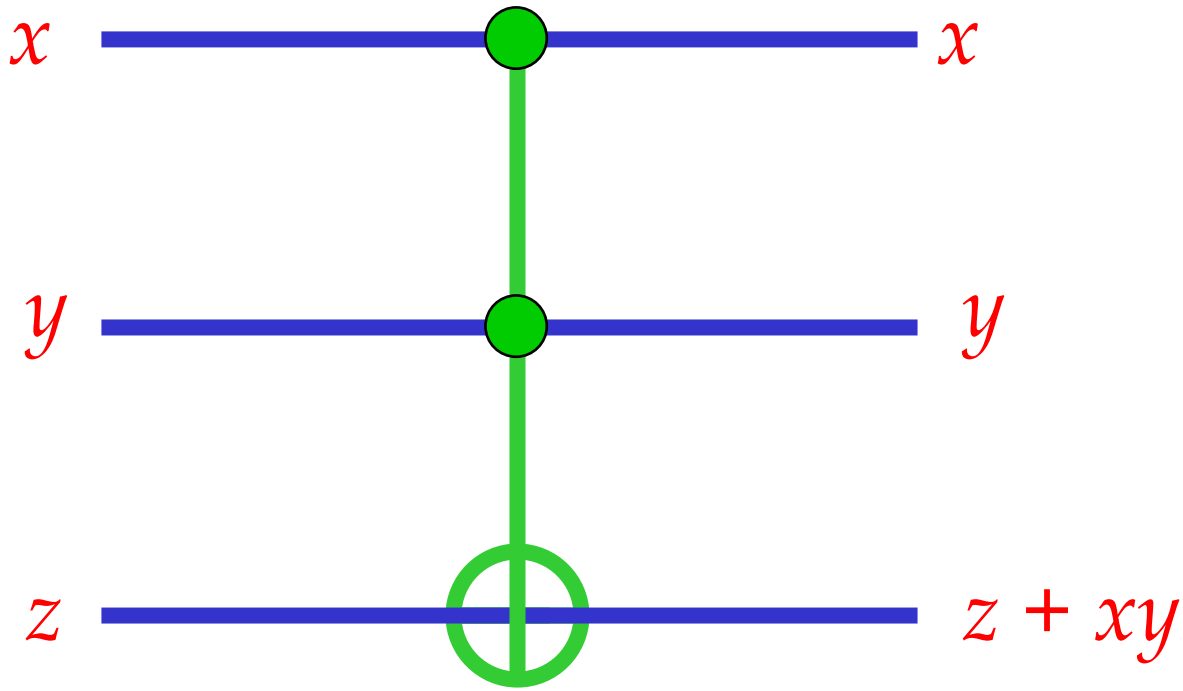

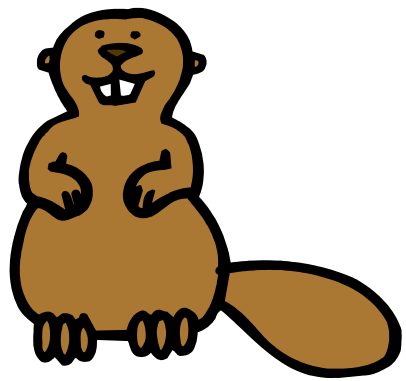

Fault-tolerant Toffoli Gate

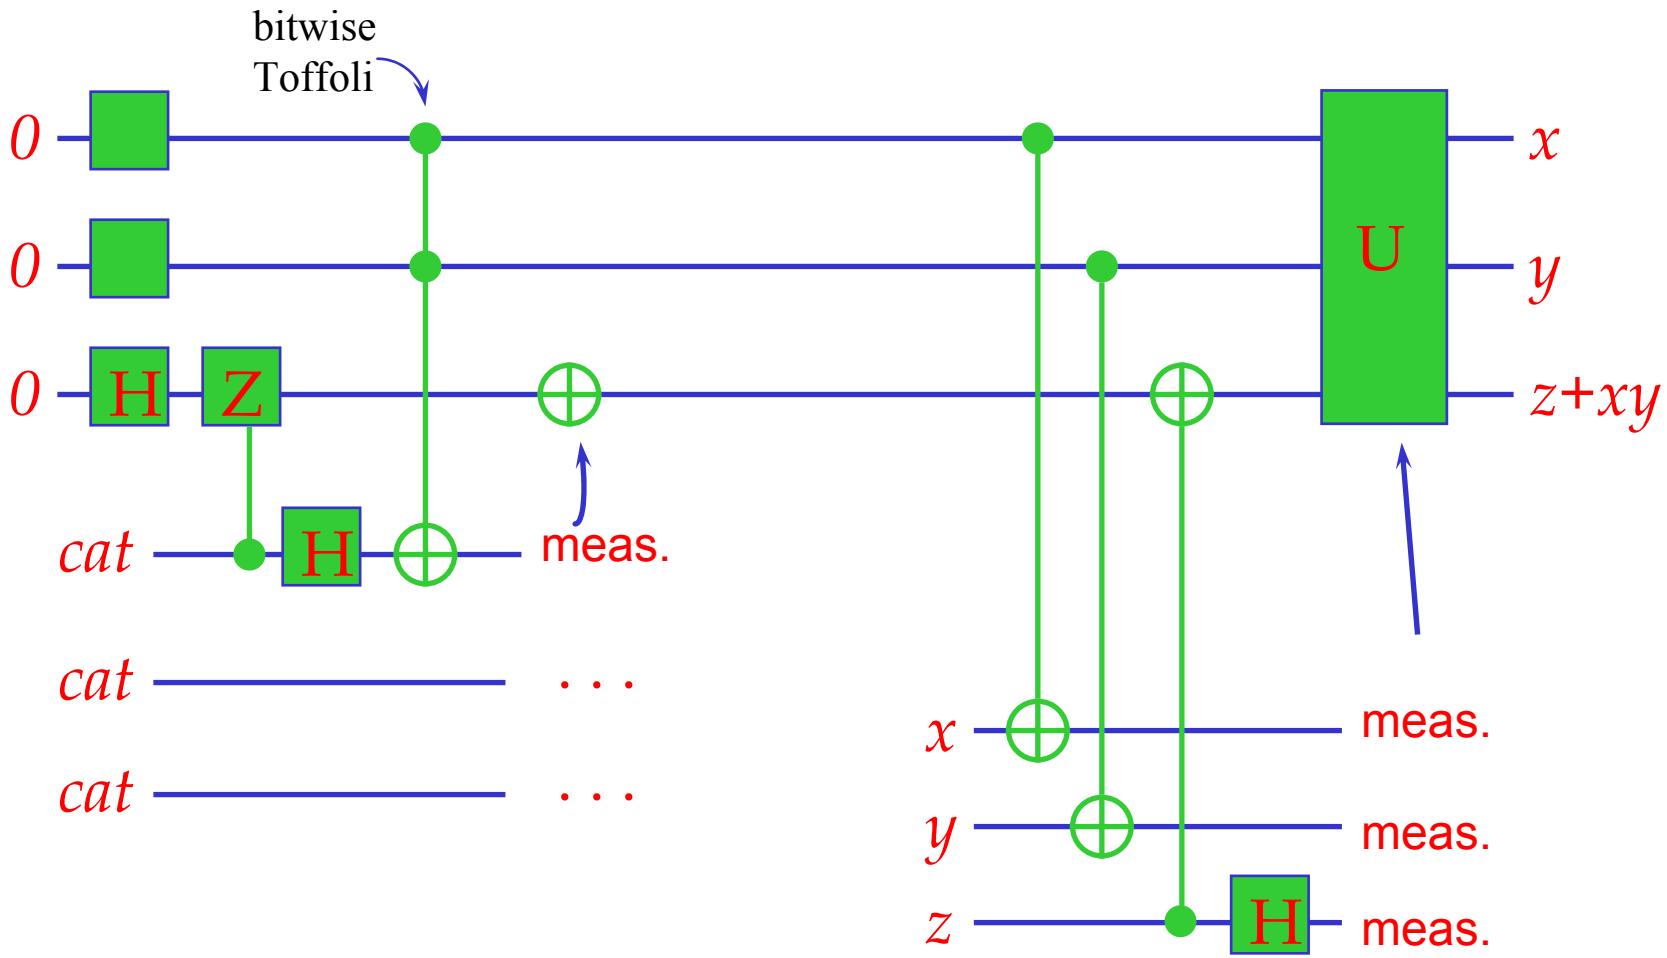

Concatenated Quantum Coding

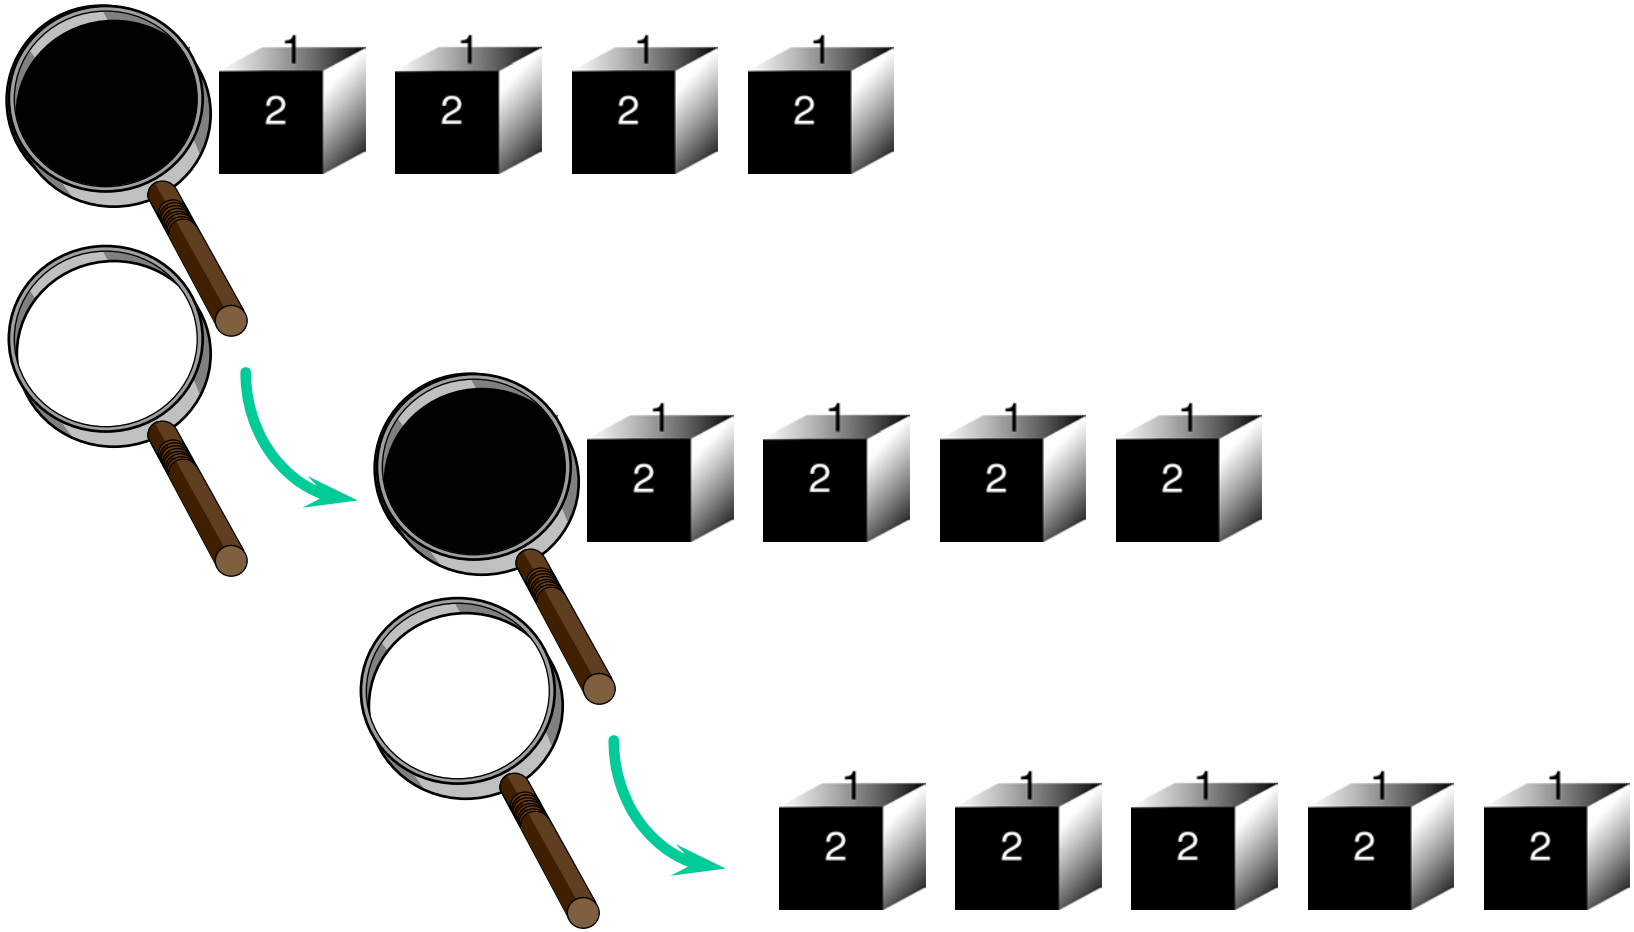

Each box, when examined with higher resolution, is itself a block of five boxes.

Higher Reliability!

Accuracy Threshold for Quantum Computation

Quantum error-correcting codes can protect quantum information. But if the probability of error per elementary quantum gate is too high, then coding does not improve the performance of a quantum computer.

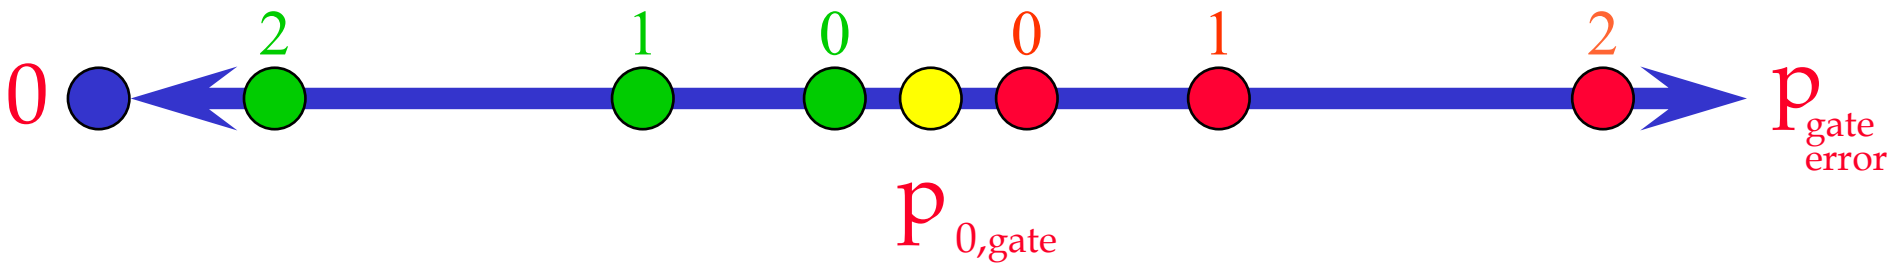

If the error probability of error per gate is below a critical value, the *accuracy threshold*, then an arbitrarily long computation *can* be executed reliably.

$$p_{gate\ error} \leq p_{0,gate} \sim 10^{-4} \rightarrow$$

Reliable Computation

Quantum vs. Classical

Very Quantum

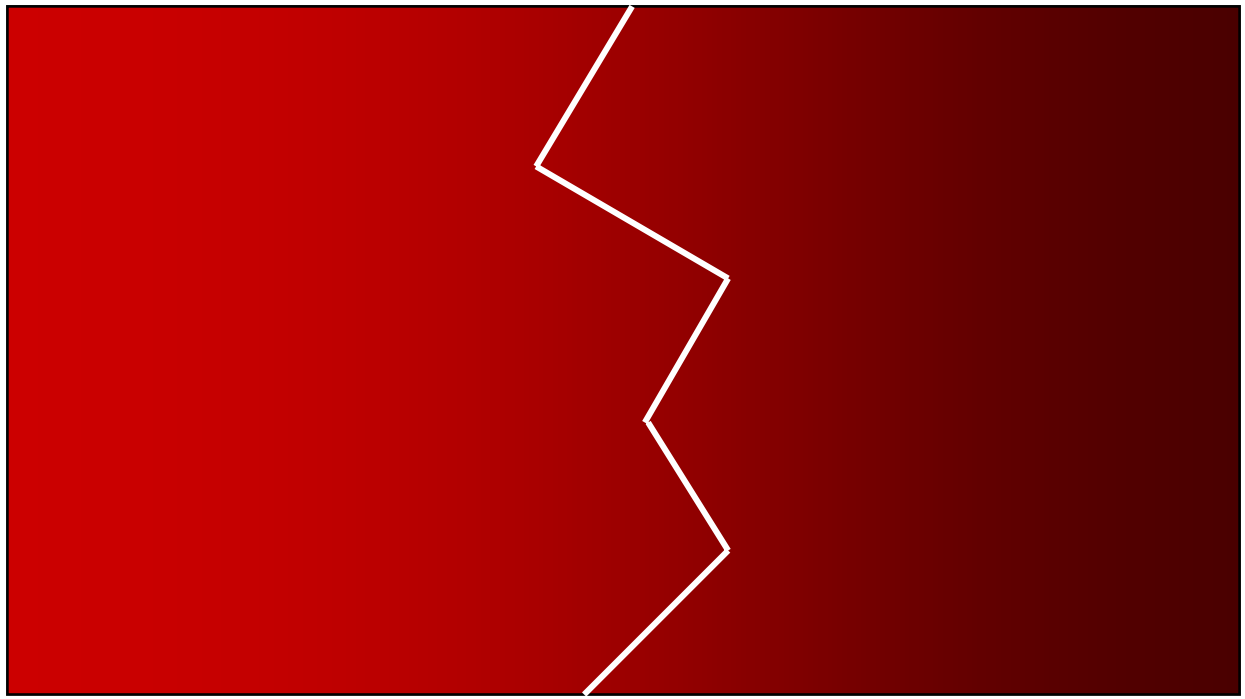

Very Classical

There is a sharp boundary between a classical phase (in which robust quantum computation is not possible) and quantum phase (not efficiently simulable by a classical computer).

Beating Decoherence

$$\frac{1}{\sqrt{2}}$$

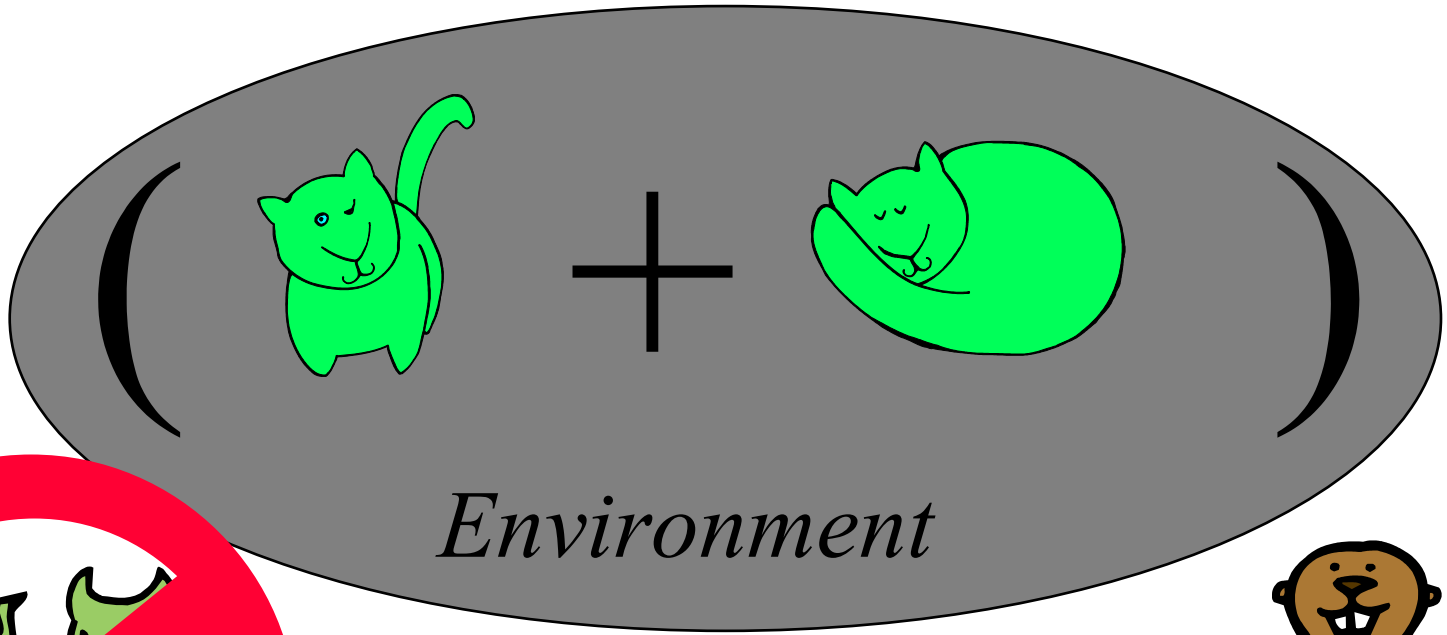

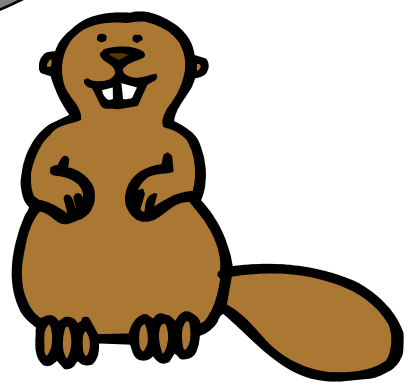

Theoretical Quantum Information Science

is driven by ...

Three ***Great*** Ideas:

1) Quantum Computation
2) Quantum Error Correction
3) Quantum Cryptography

Quantum Cryptography

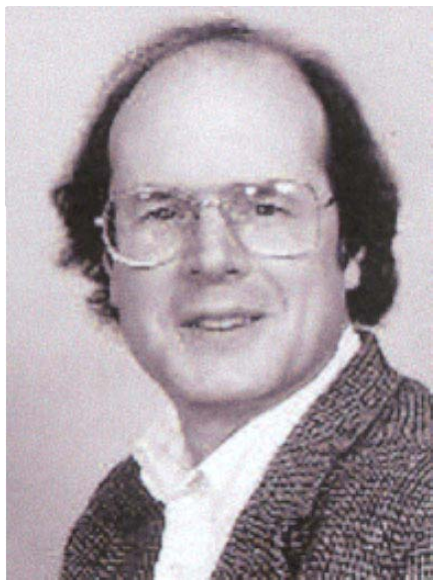

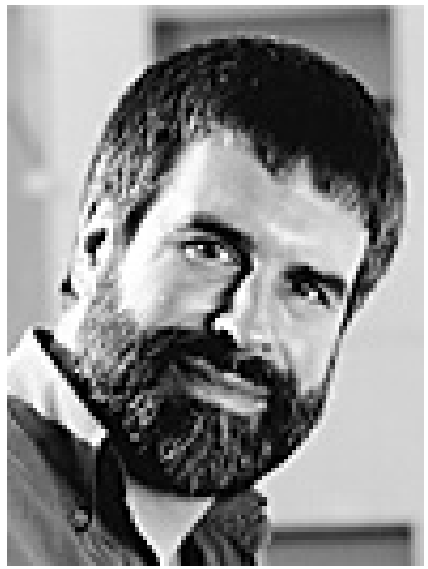

Bennett Brassard ’84

Eavesdropping on quantum information can be detected; key distribution via quantum states is *unconditionally* secure.

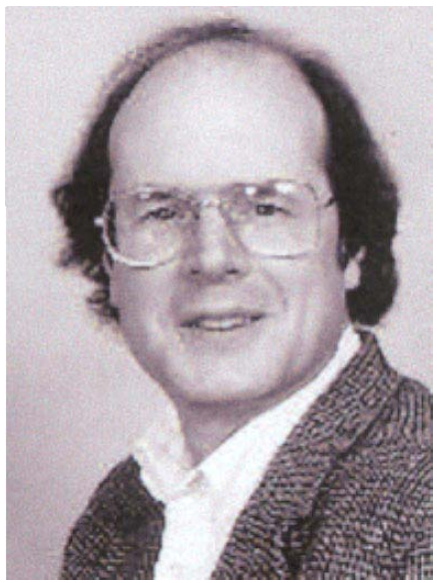

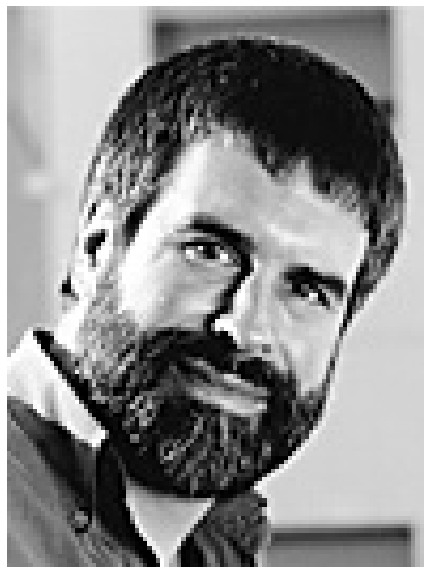

Bennett Brassard ’84

Quantum Cryptography

Alice

Eve

Bob

Privacy is founded on principles of fundamental physics, not the assumption that eavesdropping requires a difficult computation. Gathering information about a quantum state unavoidably disturbs the state.

Quantum error correction

1. Error models and error correction
2. Quantum error-correcting codes
3. Stabilizer codes
4. 5-qubit code and 7-qubit code
5. Fault-tolerant quantum computation
6. Accuracy threshold

Quantum cryptography

1. Cryptography and security
2. Quantum key distribution
3. The BB84 (four-state) protocol
4. Security proof using QECC
5. Quantum coin flipping

Cryptography

In a cryptographic protocol, two or more parties perform a task while protecting privileged information from unauthorized parties.

For example, Alice might wish to send a secret to Bob, without allowing the eavesdropper Eve to learn the secret.

Typical classical cryptographic protocols are *computationally secure.* This means that the security is founded on an (unproven) assumption that a certain computation that would break the protocol is too *hard* for the adversary to execute.

If the adversary might have a *quantum computer*, the usual assumptions about classical cryptography need to be reexamined.

Alice

Bob

One-time pad

Stronger than computational security is *information-theoretic security*. This means that even an adversary with unlimited computational power is unable to break the protocol.

A classical protocol for secret communication that is information-theoretically secure is the *one-time pad*. If Alice and Bob share a string of random bits (the “key”), then that key can be used to encipher and decipher a message. If Eve knows nothing about the key then she will not learn anything about the message by intercepting the ciphertext.

The key should be used only once (if it is used repeatedly information-theoretic security will be compromised), and then should be destroyed to ensure that Eve will not acquire a copy.

Alice

Eve

Bob

Message: HI BOB

01001000 01001001 00100000 01000010 01001111 01000010
01110100 10111001 00000101 10101001 01011100 01110100

00111100 11110001 00100101 11101011 00010011 00110110

Alice

Message: HI BOB

01001000 01001001 00100000 01000010 01001111 01000010
01110100 10111001 00000101 10101001 01011100 01110100

Eve

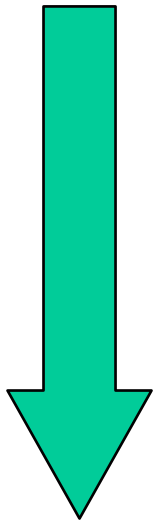

Bob

00111100 11110001 00100101 11101011 00010011 00110110
01110100 10111001 00000101 10101001 01011100 01110100

01001000 01001001 00100000 01000010 01001111 01000010

HI BOB

Alice

Eve

Bob

Message: HI BOB

01110100 10111001 00000101 10101001 01011100 01110100

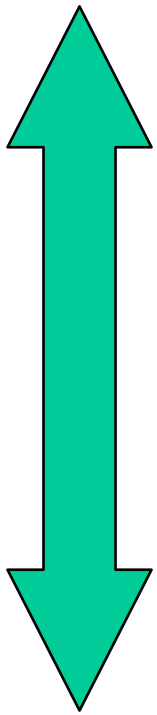

Alice and Bob can communicate privately if they share a random key that Eve doesn't know.

01110100 10111001 00000101 10101001 01011100 01110100

HI BOB

One-time pad

But what if Alice and Bob possess no shared secret random key? Perhaps they are far apart, and have never met. Or perhaps they have already consumed the key they previously shared, and do not dare to reuse it. They could ask their friend Charlie to act as an intermediary, distributing the key to Alice and Bob, but can Charlie be trusted? Perhaps Charlie is covertly in cahoots with Eve.

Alice

Bob

Public-key cryptography

Classically, this difficulty can be overcome by "public key cryptography." In public key cryptography, there are two keys, one public and one private. Everyone (including Alice and Eve) knows the public key, which is used for enciphering, but only Bob knows the private key, which is used for deciphering. Thus, anyone can send an encrypted message to Bob, but only Bob can read it! Actually, it is possible in principle to infer the private key from the known public key, but this requires a computation that is believed to be prohibitively difficult.

Public key cryptography uses a *1-way function* f, a function that is easy to compute, but hard to invert.

$$a \rightarrow b = f(a) \quad \text{encrypts}$$
$$b \rightarrow a = f^{-1}(b) \quad \text{decrypts}$$

Alice

Bob

RSA

RSA is a widely used public key crypto-system whose security is founded on the presumed difficulty of factoring large numbers.

Bob generates two prime numbers p and q (primality is easy to check), computes their product $N=pq$ and the Euler function $\varphi(N)=(p-1)(q-1)$. He choses $e < \varphi(N)$ co-prime to $\varphi(N)$, and computes $d=e^{-1} \pmod{\varphi(N)}$. Bob announces e and N, but he keeps d and $\varphi(N)$ secret.

Alice encrypts $a < N$ by computing

$$b = f(a) = a^{e} \pmod{N}$$

Bob decrypts by computing

$$a = f^{-1}(b) = b^{d} \pmod{N} = a \pmod{N}$$

It works because of Euler's theorem:

$$a^{\varphi(N)} = 1 \pmod{N}$$ (where a is co-prime to N).

RSA

Bob generates two prime numbers p and q (primality is easy to check), computes their product $N=pq$ and the Euler function $\varphi(N)=(p-1)(q-1)$. He choses $e < \varphi(N)$ co-prime to $\varphi(N)$, and computes $d=e^{-1} \pmod{\varphi(N)}$. Bob announces e and N, but he keeps d and $\varphi(N)$ secret.

A quantum computer (or any superfast factoring machine) can break RSA! If Eve can factor N she easily computes $\varphi(N)$ and d.

In fact it suffices to compute the *order* of $b=a^e \pmod{N}$ which is the same as the order of $a \pmod{N}$ and to "invert" e modulo $Order(a)$. Therefore, Eve can crack RSA if she can determine the period of a function (an abelian hidden subgroup problem).

There are other public-key schemes, but these are also vulnerable to quantum attacks...

Alice

Quantum cryptography?

Bob

Thus, if and when quantum computers become available, much of classical cryptography will become obsolete. But that won't happen for a while, so do Alice and Bob need to worry about it today? Possibly. Sometimes, it is important for a secret to remain confidential for a long time in the future. What if Alice is telling Bob about the classified design of a nuclear weapon, or the identities of covert agents who have penetrated Al Qaeda? How certain can Alice and Bob be that today's communications, intercepted and archived (but not yet decoded) by the adversary, will not be deciphered in, say, 30 years?

So quantum computing may be bad news for cryptologists. But while quantum theory taketh away, quantum theory also giveth: *quantum key distribution* is information-theoretically secure!

Alice

Quantum key distribution and the one-time pad

Bob

But what if Alice and Bob possess no shared secret random key? Perhaps they are far apart, and have never met. Or perhaps they have already consumed the key they previously shared, and do not dare to reuse it. They could ask their friend Charlie to act as an intermediary, distributing the key to Alice and Bob, but can Charlie be trusted? Perhaps Charlie is covertly in cahoots with Eve.

They can solve the problem of distributing a secure (classical) key by using quantum information. Furthermore, quantum key distribution (unlike quantum computation) is feasible with today's technology.

EPR quantum key distribution

Here is one way to accomplish QKD. Suppose that Alice and Bob share many copies of the maximally entangled (EPR, Bell) state of two qubits:

Alice

$$|\phi^+\rangle = \frac{1}{\sqrt{2}}\left(|00\rangle_{AB} + |11\rangle_{AB}\right)$$

Bob

This state can be conveniently characterized as the simultaneous eigenstate with eigenvalue one of two commuting operators: $X \otimes X = 1 = Z \otimes Z$, where

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

denote the Pauli matrices.

EPR quantum key distribution

$$|\phi^+\rangle = \frac{1}{\sqrt{2}}\left(|00\rangle_{AB} + |11\rangle_{AB}\right)$$

This state has the property that if Alice or Bob measures either X or Z, the outcome is random, e.g., $Z=1$ and $Z=-1$ occur each with probability 1/2. Furthermore, if they measure the same observable, Bob's outcome is perfectly correlated with Alice's, since $X \otimes X = 1 = Z \otimes Z$.

Consider this protocol:

1) On her half of each pair, Alice decides at random to measure either X or Z.
2) Bob does the same.
3) Through public discussion, Alice and Bob discard the results in the cases where they measured in different bases, retaining the rest.

Thus, Alice and Bob generate a shared random string.

EPR quantum key distribution

But … is it *secure?* Eve may have tampered with the pairs at some time in the past, and could have entangled them with a probe that she controls. After Alice and Bob publicly announce their bases, she might measure her probe to collect information about the key. If Eve *has* tampered with the pairs, the joint state of the pairs and Eve's probe has the form:

$$|\Phi\rangle_{ABE} = \frac{1}{\sqrt{2}}\left(\begin{array}{c} |00\rangle_{AB}|e_{00}\rangle_E + |01\rangle_{AB}|e_{01}\rangle_E \\ +|10\rangle_{AB}|e_{10}\rangle_E + |11\rangle_{AB}|e_{11}\rangle_E + \end{array} \right)$$

(where the $|e\rangle$'s need not be normalized or mutually orthogonal). *Suppose* that Alice and Bob can verify that each of their pairs satisfies $X \otimes X = 1 = Z \otimes Z$. Then...

EPR quantum key distribution

$$|\Phi\rangle_{ABE} = \frac{1}{\sqrt{2}}\Big(|00\rangle_{AB}|e_{00}\rangle_E + |01\rangle_{AB}|e_{01}\rangle_E + |10\rangle_{AB}|e_{10}\rangle_E + |11\rangle_{AB}|e_{11}\rangle_E + \Big)$$

Suppose that Alice and Bob can verify that each of their pairs satisfies $X \otimes X = 1 = Z \otimes Z$. If $Z \otimes Z = 1$, then the state must be

$$|\Phi\rangle_{ABE} = \frac{1}{\sqrt{2}}\Big(|00\rangle_{AB}|e_{00}\rangle_E + |11\rangle_{AB}|e_{11}\rangle_E + \Big)$$

And if also $X \otimes X = 1$ then it must be

$$|\Phi\rangle_{ABE} = \frac{1}{\sqrt{2}}\Big(|00\rangle_{AB} + |11\rangle_{AB} \Big)|e\rangle_E = |\phi^+\rangle_{AB}|e\rangle_E$$

Thus the pairs are unentangled with Eve, and she can't learn anything about the key by measuring her probe!

EPR quantum key distribution

$$|\phi^+\rangle_{AB} = \frac{1}{\sqrt{2}}\left(|00\rangle_{AB} + |11\rangle_{AB}\right)$$

How do Alice and Bob *verify* that their pairs are really in the state $|\phi^+\rangle$? They check that $X \otimes X = 1 = Z \otimes Z$ by conducting a statistical test. To generate an n bit key, they start out with $4n(1+\varepsilon)$ pairs. With high probability (if n is large), they measure in the same basis at least $2n$ times (otherwise, they abort the protocol). They randomly choose (say) n bits from these $2n$ bits of sifted key, and publicly compare. If all or nearly all of these bits agree, they have high statistical confidence that the remaining n bits were generated by measuring a state that was quite close to $|\phi^+\rangle^{\otimes n}$. (But how close is "close enough"? More on that later…)

BB84 quantum key distribution

$$|\phi^+\rangle_{AB} = \frac{1}{\sqrt{2}}\left(|00\rangle_{AB} + |11\rangle_{AB} \right)$$

EPR QKD illustrates that quantum entanglement is a potentially useful resource --- we can exploit it to perform a task that would otherwise be difficult (we can put the weirdness to work). But there is another way to look at the EPR protocol, such that entanglement makes no explicit appearance.

Imagine that Alice prepares all of the $|\phi^+\rangle$ pairs herself, keeping half of each pair and shipping the other half to Bob. It would not in any way reduce the efficacy of the protocol if Alice measured X or Z on her half *before* sending off the other half. In effect, then, she prepares (equiprobably) and sends to Bob one of four possible states.

BB84 quantum key distribution

Alice prepares one of four states:

Bob measures either X or Z.

$$Z = 1: \quad |0\rangle = |\uparrow\rangle$$
$$Z = -1: |1\rangle = |\downarrow\rangle$$
$$X = 1: \quad (|0\rangle + |1\rangle)/\sqrt{2} = |\rightarrow\rangle$$
$$X = -1: (|0\rangle - |1\rangle)/\sqrt{2} = |\leftarrow\rangle$$

This is called the “four-state protocol” or the “BB84 protocol” (because it was first described by Bennett and Brassard in 1984 --- the idea of quantum cryptography was first conceived by Wiesner in the early ’70’s, but he was unable to get his work published). BB84 QKD (a “prepare and measure” protocol) is no less secure than EPR QKD which uses quantum entanglement to distribute the key.

Alice

Bob

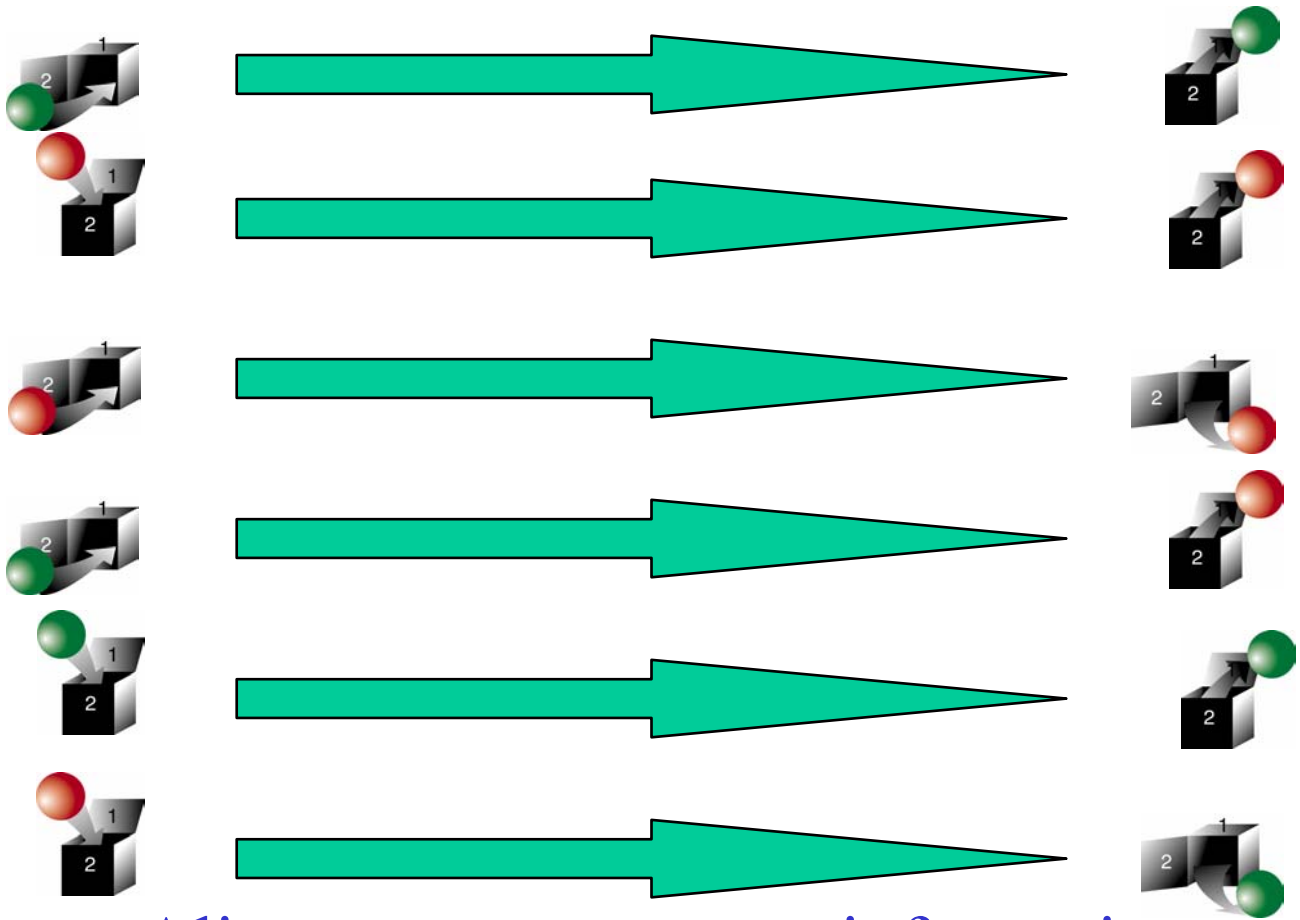

Alice can use quantum information (qubits) to send a random key to Bob.

Eve

Alice

Bob

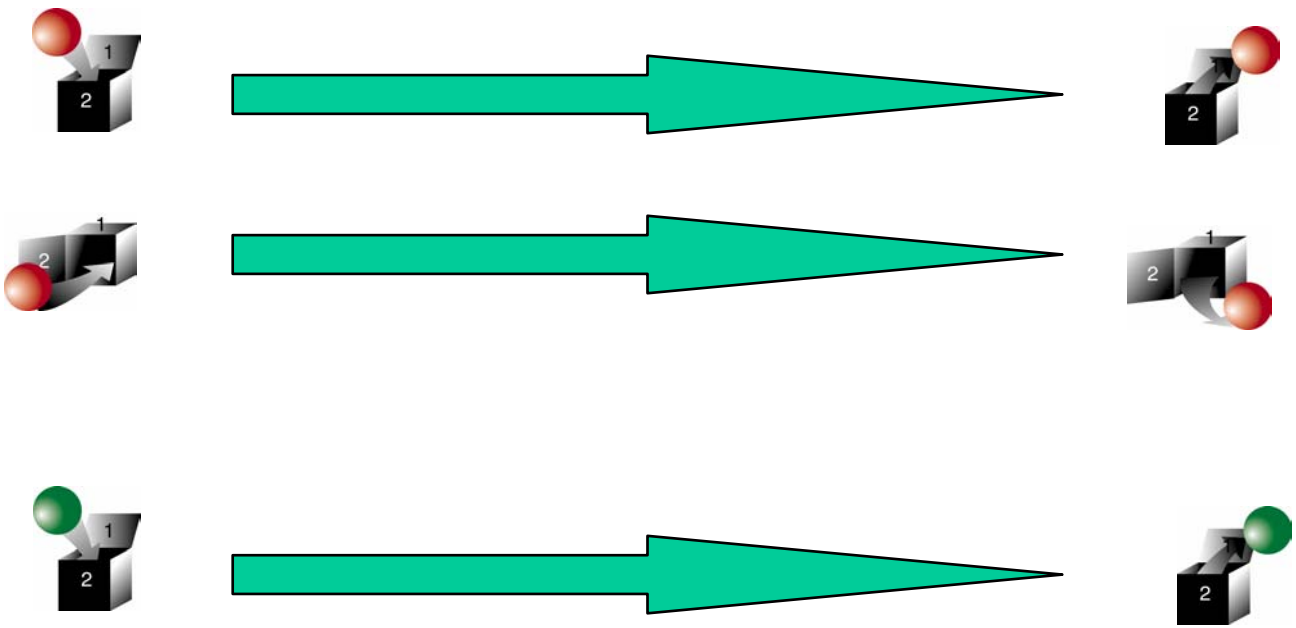

Alice can use quantum information (qubits) to send a random key to Bob.

Eve

Alice

Bob

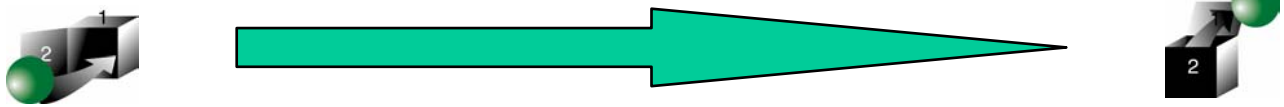

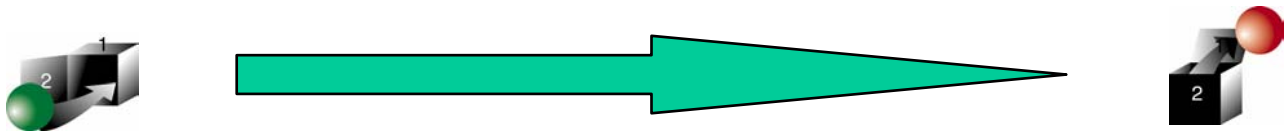

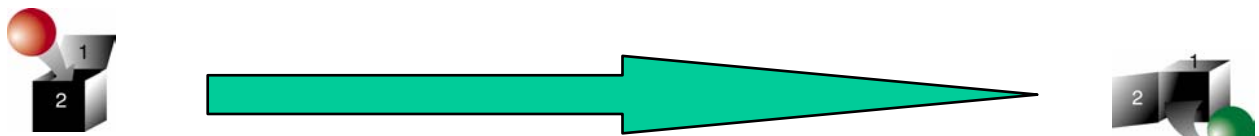

Alice can use quantum information (qubits) to send a random key to Bob.

Eve

Quantum News

"Spooky action at a distance"

January 15, 1997 Volume 1, Issue 1

Alice Announces Doors She Used!!

Alice

Lorem ipsum dolor sit amet, consectetuer adipiscing elit, sed diem nonummy nibh euismod tincidunt ut lacreet dolore magna aliguam erat volutpat. Ut wisis enim ad minim veniam, quis nostrud exerci tution ullamcorper suscipit lobortis nisl ut aliquip ex ea commodo consequat. Duis te feugifacilisi. Duis autem dolor in hendrerit in vulputate velit esse molestie consequat, vel illum dolore eu feugiat nulla facilisis at vero eros et accumsan et iusto odio dignissim qui blandit praesent luptatum zzril delenit au gue duis dolore te feugat nulla facilisi. Ut wisi enim ad minim veniam, quis nostrud exerci taion ullamcorper suscipit lobortis nisl ut aliquip ex en commodo consequat. Duis te feugifacilisi.per suscipit lobortis nisl ut aliquip ex en commodo consequat. Duis te feugifacilisi. Lorem ipsum dolor sit amet, consectetuer adipiscing elit, sed diem nonummy nibh euismod tincidunt ut lacreet dolore magna aliguam erat volutpat. Ut wisis enim ad minim veniam, quis nostrud exerci tution ullamcorper suscipit lobortis nisl ut aliquip ex ea commodo consequat. Duis te feugifacilisi. Duis autem dolor in hendrerit in vulputate velit esse molestie consequat, vel illum dolore eu feugiat nulla facilisis at vero eros et accumsan et iusto odio dignissim qui blandit praesent luptatum zzril delenit au gue duis dolore te feugat nulla facilisi.ipsum dolor sit amet, consectetuer adipiscing elit, sed diem nonummy nibh euismod tincidunt ut lacreet dolore magna aliguam erat volut-

Lorem ipsum dolor sit amet, consectetuer adipiscing elit, sed diem nonummy nibh euismod tincidunt ut lacreet dolore magna aliguam erat volutpat. Ut wisis

Lorem Ipsum

Alice

Bob

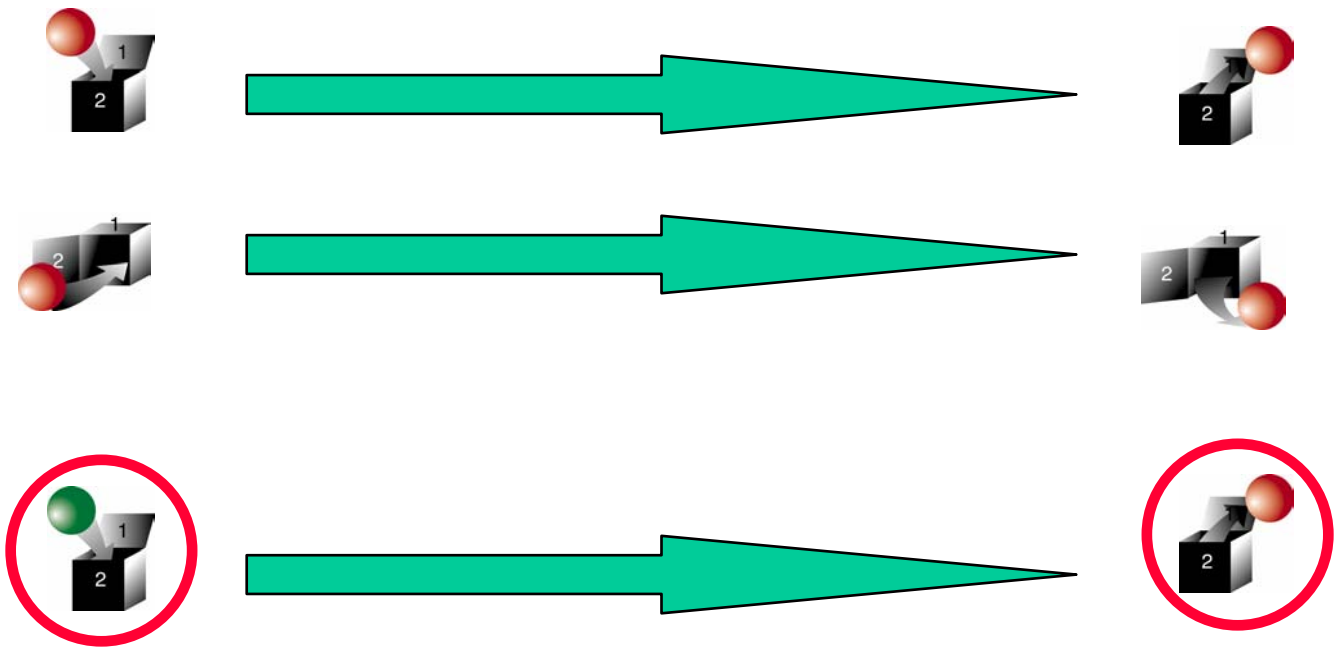

Alice can use quantum information (qubits) to send a random key to Bob.

Alice

Quantum key distribution, augmented by classical protocols that correct errors and amplify privacy, is *provably* secure against *arbitrary* eavesdropping attacks.

Bob

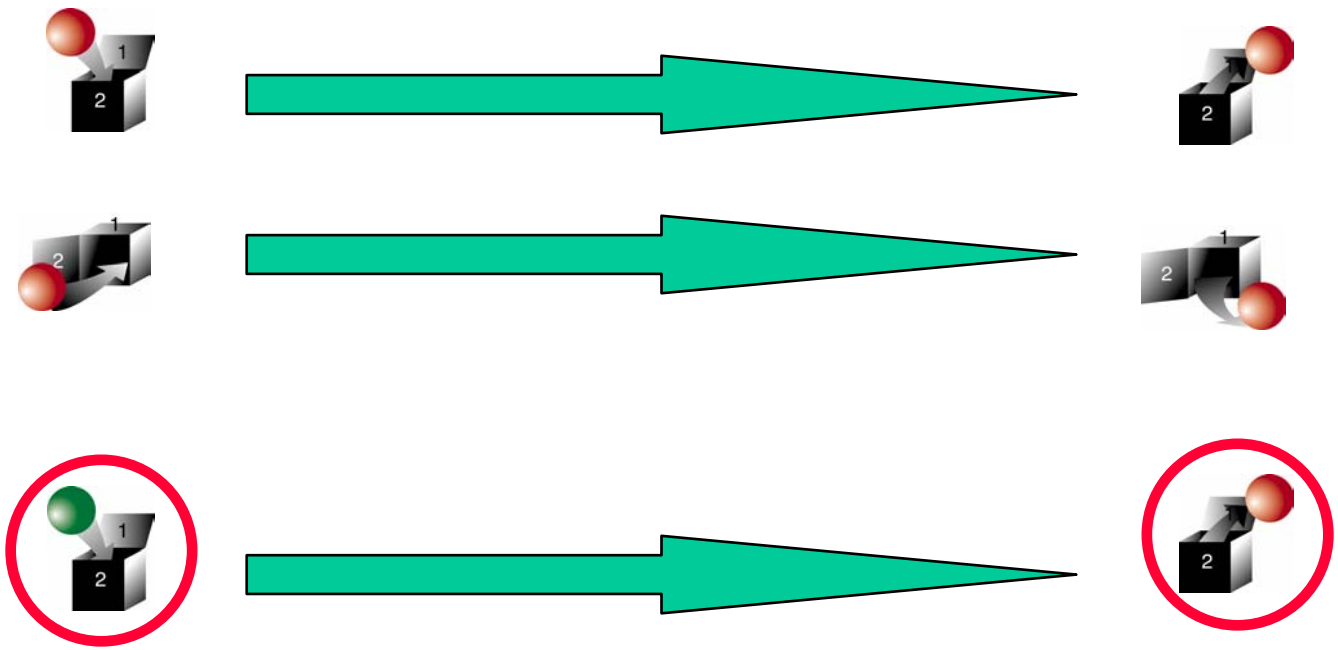

Alice can use quantum information (qubits) to send a random key to Bob.

Information vs. disturbance

Why is eavesdropping detectable? Because it is not possible to *collect information* that distinguishes among *nonorthogonal* quantum states without creating a detectable disturbance. In contrast, we can distinguish among orthogonal states (read classical information) without leaving any trace.

Consider a game in which Alice prepares either $|\uparrow\rangle$ or $|\rightarrow\rangle$ (chosen at random). Eve is supposed to guess which state Alice prepared. There is no strategy Eve can play that will win the game every time (her optimal probability of success is 85.4%), and no strategy that is better than a random guess leaves the state unmodified.

But if Alice prepares either $|\uparrow\rangle$ or $|\downarrow\rangle$, then Eve can win every time, without disturbing the state at all.

Information vs. disturbance

Suppose Alice prepares either $|\varphi\rangle$ or $|\psi\rangle$. To distinguish the two possible states, Eve performs a unitary transformation that rotates her probe while leaving Alice's state intact

$$U: \quad |\varphi\rangle_A \otimes |0\rangle_E \to |\varphi\rangle_A \otimes |e\rangle_E$$

$$|\psi\rangle_A \otimes |0\rangle_E \to |\psi\rangle_A \otimes |f\rangle_E$$

where $|e\rangle$ and $|f\rangle$ are normalized states. Since U preserves inner products,

$$\langle\psi|\varphi\rangle \cdot \langle f|e\rangle = \langle\psi|\varphi\rangle ,$$

and if $|\varphi\rangle$ and $|\psi\rangle$ are nonorthogonal, then $\langle f|e\rangle = 1$; the states of the probe are indistinguishable. Eve's measurement of the probe cannot reveal any information about whether the state is $|\varphi\rangle$ or $|\psi\rangle$. On the other hand if $|\varphi\rangle$ and $|\psi\rangle$ are orthogonal, the probe states can also be orthogonal. Eve can *copy* the info.

Information vs. disturbance

So we see that it is impossible to collect any information that distinguishes two nonorthogonal states without creating a disturbance. The same principle applies if Eve wants to distinguish the four BB84 states: $|\uparrow\rangle$, $|\rightarrow\rangle$, $|\downarrow\rangle$, $|\leftarrow\rangle$.

This is a powerful argument, but it is not quantitative. What if Eve collects just a little bit of information --- how big a disturbance must she cause? Or if she is permitted to alter the fidelity of the state slightly, how much information can she gain?

Quantum key distribution provides an excellent setting for studying the information/disturbance tradeoff, which is of fundamental interest in quantum information theory. We have well motivated ways to quantify both information gain and disturbance: what does Eve know about the key, and what error rate do Alice and Bob detect?

BB84 quantum key distribution

In the real world, communication channels (especially quantum channels) are imperfect. Therefore, Alice and Bob can expect to find some errors in their verification test even if Eve has not collected any information at all. Still, when errors occur, they (as cautious cryptologists) should pessimistically assume that the errors were caused by Eve's tampering.

Thus we must enhance the BB84 QKD protocol in two ways. First we should incorporate (classical) *error correction*, to ensure that Alice and Bob really have the same secret key. Second, we should include (classical) *privacy amplification*. After error correction, Alice and Bob agree on n bits about which Eve has only a little information. Then A. and B. both process the bits, extracting $k < n$ bits about which Eve has even less information.

Error correction and privacy amplification

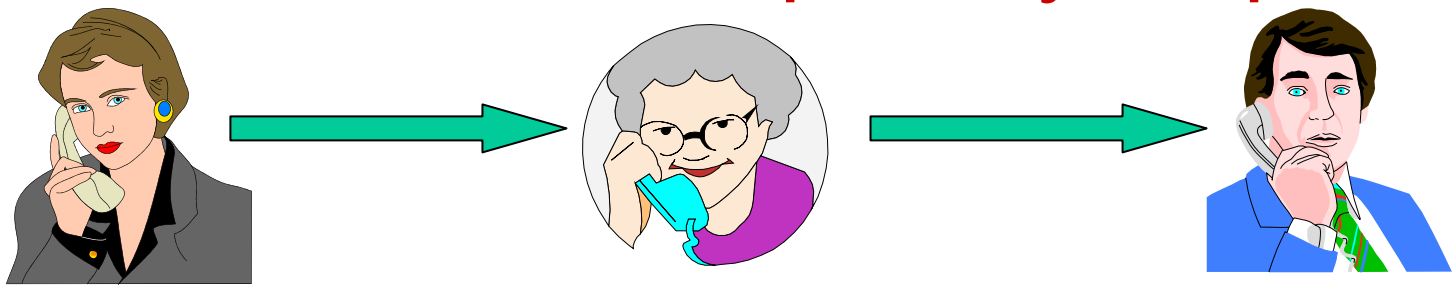

For example, to do error correction, Alice and Bob both divide their private key bits into blocks of three.

(011)(101)(001) → (111)(100)(001)

(Bob's errors are shown in red.) Then Alice announces her error syndrome: the bit (if any) in each block that differs from the other two. She flips this bit and so does Bob.

(111)(111)(000) → (011)(110)(000)

Now each of Alice's blocks is a codeword of the 3-bit repetition code. Bob decodes his block by majority voting. If there is no more than one error in a block of three, then Bob's decoded bit agrees with Alice's.

(1) (1) (0) (1) (1) (0)

Error correction and privacy amplification

After error correction, Alice and Bob are likely to share the same bits. Next they perform privacy amplification to extract bits that are more secure. For example Alice and Bob might divide their corrected key bits into blocks of three. And in each block compute the parity of the three bits.

$[(1)\,(1)\,(0)]\quad[(0)\,(1)\,(0)]\quad[(1)\,(0)\,(0)]$

$[0]\quad[1]\quad[1]$

If Eve has a little bit of information about each corrected bit, she'll know less about the parity bit of a block.

Security of BB84

To make a precise statement about the security of the BB84 protocol, we consider the asymptotic behavior for very large key length. Then:

Theorem: For *any* attack by Eve, either the verification test fails with probability exponentially close to 1, or if the test succeeds then Bob's key agrees with Alice's with probability exponentially close to 1, and Eve's information about the key is exponentially small.

"Exponentially close/small" can be taken to mean $< \exp(-Ck)$ where k is the length of the final key and C is a constant; Eve's information is the mutual information of the key and the outcome of Eve's measurement of her probe. The theorem says that either Alice and Bob are almost certain to catch Eve, or else Eve knows almost nothing.

As cautious cryptologists, we make no assumptions about Eve's technological power. In particular, she might have a quantum computer, enabling her to make collective measurements on all the qubits at once. The security is *information- theoretic*.

Security of BB84

As cautious cryptologists, we make no assumptions about Eve's technological power. In particular, she might have a quantum computer, enabling her to make collective measurements on all the qubits at once. The security is *information- theoretic*.

This information-theoretic security is sometimes called "unconditional security," meaning that Eve's attack is completely unrestricted. However there are conditions on the equipment used in the protocol --- Alice's source of BB84 states and Bob's detector that measures X or Z. For now we'll suppose that the source and the detector are perfect. More about this later...

Security of BB84

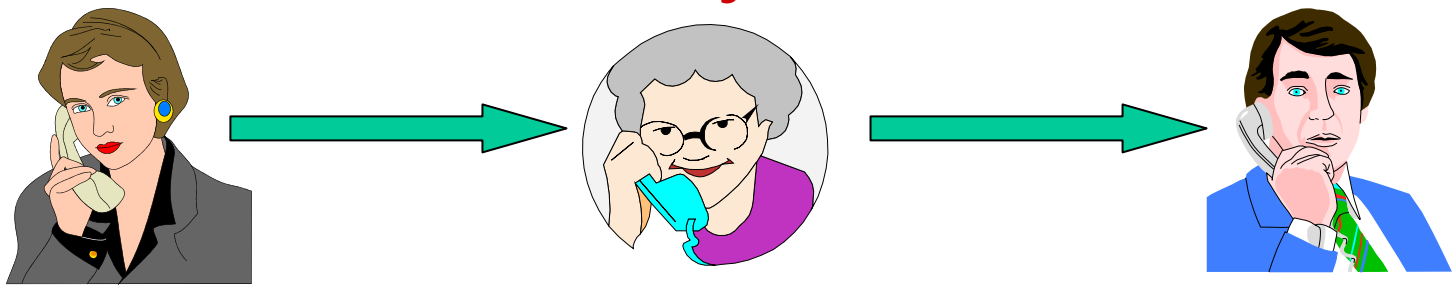

Theorem: For *any* attack by Eve, either the verification test fails with probability exponentially close to 1, or if the test succeeds then Bob's key agrees with Alice's with probability exponentially close to 1, and Eve's information about the key is exponentially small.

To explain what it means for the verification test to succeed (or fail) we need to specify the maximum error rate $\delta_{\max}$ that Alice and Bob should be willing to tolerate. Furthermore, for $\delta < \delta_{\max}$, we should be able to specify the asymptotic rate $R=k/n$ at which they can extract secure final key from sifted key by choosing appropriate schemes for error correction and privacy amplification.

Intercept/resend attack

Suppose that *Eve* intercepts each qubit that Alice sends, deciding at random to measure in the X or Z basis, records the outcome, and then sends on to Bob the X or Z eigenstate found by her measurement.

This is a powerful attack: Alice cannot come to agreement with Bob on a key that Eve does not know. The reason is that, after the bases are announced, Eve can simulate Bob's outcomes with a random number generator, and Alice can't possibly know more about Bob's key than Eve does.

The intercept/resend attack generates an error rate $\delta = 1/4$; therefore, the maximum tolerable error rate satisfies $\delta_{\max} \leq .25$

Approximate cloning attack

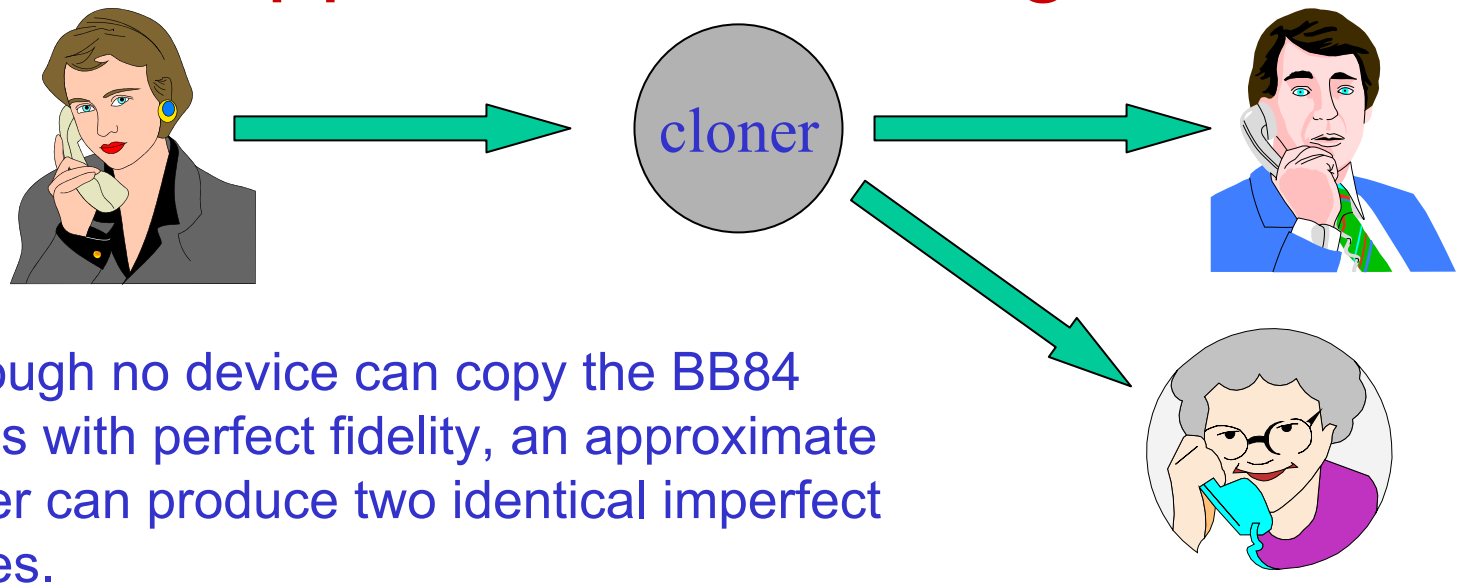

Although no device can copy the BB84 states with perfect fidelity, an approximate cloner can produce two identical imperfect copies.

The optimal approximate cloner for the BB84 ensemble achieves a fidelity of 85.4%. If Eve sends one copy for herself and sends one to Bob, then the bit error rate will be 14.6%.

If the post-processing of the sifted key is achieved with only one-way communication from Alice to Bob, then Eve knows as much as Bob, and $\delta_{\max} \leq .146$. But two-way communication between Alice and Bob breaks the symmetry between Bob and Eve, so that a higher bit-error rate could be tolerable. In fact, for this two-way case, the intercept/resent strategy provides the best currently known upper bound on $\delta_{\max}$.

Bounds on the bit-error rate

The best upper bounds on the acceptable bit-error rate are:

Optimal cloner attack: $\delta_{\text{max,one-way}} \leq .146$
Intercept/resend attack: $\delta_{\text{max, two-way}} \leq .25$

The best lower bounds found in proofs of security are:

$$\delta_{\text{max,one-way}} \geq .110$$
$$\delta_{\text{max, two-way}} \geq .189$$

We'll see where the 11% lower bound comes from...

QKD and QEC

In “prepare and measure” QKD, the error correction and privacy amplification applied during post-processing of the sifted key are *classical* protocols. Yet for proving security of e.g. BB84, the theory of *quantum* error correction is very helpful. Why?

QEC: the environment becomes entangled with our qubits. We remove the entanglement and recover a pure state by correcting both X errors and Z errors.

QKD: the eavesdropper collects information about the outcomes of our X and Z measurements by entangling her probe with the transmitted qubits.

QKD and QEC share the goal of protecting quantum states against entanglement with the outside world.

QECC's and the Security of QKD

- A quantum state encoded using a quantum error-correcting code remains pure -- the environment does not become entangled with the encoded quantum information, if the error rate is sufficiently low.

- Suppose that Alice sends to Bob qubits that are protected by a quantum error-correcting code.

- To collect information about the key bits, Eve must become entangled with the encoded state transmitted from Alice to Bob. But if the verification test shows that *quantum error correction will succeed*, then we can disentangle Eve, preventing her from acquiring information about the key. Hence, key distribution using encoded information is secure.

- In general, Alice and Bob need *quantum computers* if Alice is to encode using a QECC, and Bob is to correct the errors and decode the state.

QECC's and the Security of QKD

• But if the QECC is of the *CSS type*, then bit flip error correction and phase error correction can be separated. Bit flip error correction is needed to ensure the accuracy of the key, but phase error correction has no effect on the key; its purpose is to ensure privacy.

• Thus it is not necessary to *perform* phase error correction -- to ensure privacy, it is enough to know that it would have succeeded *if it had been done*.

• Using such reasoning based on *virtual quantum error correction*, we can prove that the BB84 quantum key distribution scheme, suitably augmented by classical error correction and privacy amplification, is secure against all possible eavesdropping strategies.

• E.g., a bit-error rate up to 11% is acceptable (if classical post-processing involves one-way communication from Alice to Bob).

7-qubit code (a *CSS code*)

$$M_{Z,1} = Z\ I\ Z\ I\ Z\ I\ Z$$
$$M_{Z,2} = Z\ Z\ I\ I\ Z\ Z\ I$$
$$M_{Z,3} = Z\ Z\ Z\ Z\ I\ I\ I$$

Corrects the bit-flip (X) errors. The three-bit string $(M_{Z,1}, M_{Z,2}, M_{Z,3})$ (if nonzero) points to the position of the error.

$$M_{X,1} = X\ I\ X\ I\ X\ I\ X$$
$$M_{X,2} = X\ X\ I\ I\ X\ X\ I$$
$$M_{X,3} = X\ X\ X\ X\ I\ I\ I$$

Corrects the phase (Z) errors. The three-bit string $(M_{X,1}, M_{X,2}, M_{X,3})$ (if nonzero) points to the position of the error.

The $\boldsymbol{M_Z}$'s commute with the $\boldsymbol{M_X}$'s , because each row of the $\boldsymbol{M_Z}$ matrix has an even number of "collisions" with each row of the $\boldsymbol{M_X}$ matrix; i.e., the rows are orthogonal in the sense of linear algebra over the field $\mathbb{Z}_2$. Any two matrices with this property define a quantum code, which is said to be of the "CSS" (Calderbank-Shor-Steane) type. With CSS codes, the bit-flip and phase error correction can be executed separately. The encoded operations can be chosen to be $\overline{Z} = I\ I\ I\ I\ Z\ Z\ Z$, $\overline{X} = I\ I\ I\ I\ X\ X\ X$ which commute with the code stabilizer and are not contained in it.

EPR QKD using the 7-qubit code

Alice and Bob share 7 EPR pairs, but the pairs are *noisy* (have imperfect fidelity). The noise could be due to tampering by Eve. Let's *assume* (for now) that the effect of Eve's tampering is first to apply X to at most one of Bob's 7 qubits and then to to apply Z to at most one (possibly the same one).

Alice measures the 3 Z stabilizer generators of the 7-qubit code: $(M_{Z,1}, M_{Z,2}, M_{Z,3})$, and if she finds a nontrivial syndrome, she applies X to the qubit identified by the syndrome. She reports her recovery operation to Bob, and he applies X to his corresponding qubit. Then Bob measures the 3 Z stabilizer generators and recovers again --- in this step he reverses the X error (if any) that Eve introduced. Alice and Bob then repeat this procedure for the 3 X stabilizer generators $(M_{X,1}, M_{X,2}, M_{X,3})$.

The state that Alice and Bob have obtained is what they would have obtained if they started with 7 perfect EPR pairs, and each had projected her/his 7 qubits onto the codespace. The state of 7 perfect pairs is an eigenstate with eigenvalue one of the two commuting encoded operations:

$$\bar{Z}_A \otimes \bar{Z}_B, \quad \bar{X}_A \otimes \bar{X}_B$$

Thus Alice and Bob share the encoded EPR pair $|\bar{\phi}\rangle_{AB}$ with perfect fidelity (Eve is unentangled with the encoded state). Alice and Bob can now each measure $\bar{Z}$ to generate a secure bit. Alice and Bob managed to "purify" their noisy entanglement, extracting perfect entanglement that could then be used to generate the key.

EPR QKD using the 7-qubit code

So far, we have considered a protocol for which Alice and Bob need quantum computers to measure the collective obervables $(M_{Z,1}, M_{Z,2}, M_{Z,3})$ and $(M_{X,1}, M_{X,2}, M_{X,3})$. This was necessary for them to be able to implement correction of both the X errors and the Z errors.

But Alice and Bob generate the key bit by measuring $\overline{Z} = I\ I\ I\ I\ Z\ Z\ Z$ ...
So there is no need to correct the Z errors --- they have no effect on the key. And if we don't correct these errors, there is no need to measure the stabilizer generators $(M_{X,1}, M_{X,2}, M_{X,3})$ that diagnose the Z errors.

What remains of our protocol? Alice and Bob measure $(M_{Z,1}, M_{Z,2}, M_{Z,3})$ and $\overline{Z}$, and Bob applies an error-correcting bit flip (if necessary) to make sure that his $\overline{Z}$ agrees with Alice's.

The reduced protocol is almost entirely classical: Alice prepares and sends bits (Z eigenstates) to Bob. Errors might occur in transit, which Bob corrects. Alice and Bob compute the parity of the last three key bits to determine one bit of their final key.

Recall that we *assumed* that at most one of the qubits would suffer an X error in the channel and that at most one would suffer a Z error. If these assumptions are justified, then Alice and Bob agree on the final key bit and Eve knows nothing about it.

"Quantum to classical reduction"

Alice and Bob share 7 EPR pairs, but the pairs are *noisy* (have imperfect fidelity). The noise could be due to tampering by Eve. Let's *assume* (for now) that the effect of Eve's tampering is first to apply X to at most one of Bob's 7 qubits and then to to apply Z to at most one (possibly the same one).

We assumed that each of the pairs is a *Bell pair*, a simultaneous eigenstate of the commuting operators $X \otimes X$ and $Z \otimes Z$, though perhaps not with the prescribed eigenvalue +1. For a general attack by Eve, this might not be the case. However, imagine that Alice and Bob are able to perform Bell measurements on their pairs right before the final measurements that determine the key. This will have no effect on the fidelity of the state with the encoded $|\bar{\phi}\rangle_{AB}$, because $|\bar{\phi}\rangle_{AB}$ is already an eigenstate of Bell measurement:

$$\langle \bar{\phi} | \Pi \, \rho \, \Pi | \bar{\phi} \rangle = \langle \bar{\phi} | \rho | \bar{\phi} \rangle$$

where ρ is the state A. and B. have at the end of the protocol and Π is the projector onto 7 $|\phi\rangle_{AB}$'s

Furthermore, we can commute this Bell measurement through the steps of the protocol, again without changing anything. The key point is that to do their error recovery Alice and Bob in effect measure stabilizer generators $M_A \otimes M_B$ that act simultaneously on Alice's system and Bob's, and these operators commute with the Bell measurements. If we imagine that they measure *only* $M_A \otimes M_B$ rather than M_A and M_B separately (which makes life no easier for Eve), then we can move the Bell measurement up to the beginning of the protocol without altering its effectiveness. This is called the "quantum to classical reduction" because we reduce a general attack by Eve to a discrete attack in which she applies X or Z to some of Bob's qubits.

7-qubit code generalized: *CSS codes*

$$
\begin{aligned}
M_{Z,1} &= Z\ I\ Z\ I\ Z\ I\ Z \\
M_{Z,2} &= Z\ Z\ I\ I\ Z\ Z\ I \\
M_{Z,3} &= Z\ Z\ Z\ Z\ I\ I\ I
\end{aligned}
$$

The matrix $\boldsymbol{M_Z}$ is the *parity check matrix* of a classical code $\boldsymbol{C_Z}$: its codewords are binary strings annihilated by $\boldsymbol{M_Z}$.

$$
\begin{aligned}
M_{X,1} &= X\ I\ X\ I\ X\ I\ X \\
M_{X,2} &= X\ X\ I\ I\ X\ X\ I \\
M_{X,3} &= X\ X\ X\ X\ I\ I\ I
\end{aligned}
$$

The matrix $\boldsymbol{M_X}$ is the *generator matrix* of a classical code $\boldsymbol{C_X}^{\perp}$: its codewords are linear combinations of the rows of $\boldsymbol{M_X}$.

The classical code $\boldsymbol{C_X}^{\perp}$ is a subcode of $\boldsymbol{C_Z}$. Expressed in the ***Z*-basis,** a basis for the codewords of the CSS quantum code is:

$$|\bar{w}\rangle \propto \sum_{u \in C_X^{\perp}} |w+u\rangle\,, \quad w \in C_Z$$

There is a codeword associated with each *coset* of $\boldsymbol{C_X}^{\perp}$ in $\boldsymbol{C_Z}$. We use $\boldsymbol{C_Z}$ to diagnose the bit flip errors and $\boldsymbol{C_X}$ to diagnose the phase errors (in the conjugate basis).

QECC's and the Security of QKD

If the QECC is of the *CSS type*, then bit flip error correction and phase error correction can be separated. Bit flip error correction is needed to ensure the accuracy of the key, but phase error correction has no effect on the key; its purpose is to ensure privacy.

When we reduce an EPR protocol to a "prepare and measure" protocol, a vestige of the QECC survives in the classical procedures we use to correct bit errors and amplify privacy. The power of the CSS code to correct X errors ensures that Alice and Bob have the same key bits. The power of the CSS code to correct Z errors ensures that Eve does not know the value of the *encoded* Z. The value of the encoded Z is found by applying the parity check matrix of $\boldsymbol{C_X}^{\perp}$ to the classical bit string, which identifies a coset in $\boldsymbol{C_Z} / \boldsymbol{C_X}^{\perp}$. This coset is the final key.

Verification

Alice and Bob share 7 EPR pairs, but the pairs are *noisy* (have imperfect fidelity). The noise could be due to tampering by Eve. Let's *assume* (for now) that the effect of Eve's tampering is first to apply X to at most one of 7 qubits and then to to apply Z to at most one (possibly the same one).

How can Alice and Bob be sure that the number of X and Z errors is small enough that error correction will work? They measure the error rate by sacrificing some of the sifted key in the verification test. From (classical) sampling theory, the joint probability that they find $n\delta$ errors in n randomly selected test bits, and the number of errors in n untested bits is greater than $n(\delta+\varepsilon)$ is

$$< \exp\left[-\varepsilon^2 n / 4\delta(1-\delta)\right]$$

In the EPR version of the protocol, we can suppose that Alice and Bob apply to randomly selected qubits a transformation H that interchanges eigenstates of X and of Z, prior to measuring Z. This imposes symmetry between the rates of X errors and of Z errors, and also ensures that when we reduce to a prepare-and-measure protocol we obtain exactly BB84, in which either X or Z is measured with equal probability. The symmetry arises because Eve doesn't know the basis used, so she can't distinguish bit and phase errors in her attack.

Bounding Eve's information

If we find error rate δ in the verification test, then we know that the error rate in the the qubits used for key generation is less than $\delta+\varepsilon$ with probability exponentially close to one. In the EPR scenario, we use a CSS code that can correct $n(\delta+\varepsilon)$ bit flip errors and $n(\delta+\varepsilon)$ phase errors in a block of n. This means that error correction succeeds with high probability, i.e., that the after error correction, the density operator ρ of the pairs is very close to the state $|\Phi^{(k)}\rangle \equiv |\phi^+\rangle^{\otimes k}$ of k encoded $|\phi^+\rangle$ states:

$$\langle \Phi^{(k)} | \rho | \Phi^{(k)} \rangle = 1-\Delta$$

where Δ is exponentially small. We pessimistically assume that the impurity of ρ is entirely due to entanglement with Eve. When Alice and Bob measure the encoded pairs, for each possible outcome some corresponding pure state is prepared for Eve. By measuring this state, Eve can acquire information about the outcome of the Alice/Bob measurement. According to *Holevo's Theorem*, this information is bounded as

$$I(AB;E) \le S(\rho) \le H_2(\Delta) + \Delta \log\left(2^{2k}-1\right)$$

where ρ is the density operator of Alice's and Bob's pairs,

$S(\rho) = -tr\left(\rho \log_2 \rho\right)$ is the Von Neuman entropy of the density operator, and

$H_2(\Delta) = -\Delta \log_2 \Delta - (1-\Delta)\log_2(1-\Delta)$ is the Shannon entropy.

Key generation rate

To determine the rate at which Alice and Bob can extract secure final key from their sifted key, we consider the asymptotic rates of CSS codes with large block size. If the block size is n there are n-k stabilizer generators and k encoded qubits. If there are $n\delta$ bit flip errors and $n\delta_p$ phase errors, what rate $R=k/n$ can be achieved, such that the probability of an encoding error becomes exponentially small for large n ?

We are entitled to imagine that Alice applies a random permutation to the qubits that is inverted by Bob (this is equivalent to randomizing our CSS code). Therefore, we may suppose that the errors occur at randomly distributed positions.

Recall Shannon's result for a binary symmetric (classical bit flip) channel. The achievable rate is:

$$R = 1 - H_2(\delta), \quad H_2(\delta) = -\delta \log \delta - (1-\delta)\log_2(1-\delta)$$

Heuristically, there must be more than enough error syndromes to point to each of the typical errors, or:

$$\binom{n}{n\delta} \approx 2^{nH_2(\delta)} < 2^{n-k}$$

We need to sacrifice a fraction $H_2(\delta)$ of the bits to correct errors. In the case of CSS codes, we need to sacrifice a fraction $H_2(\delta)$ of our qubits to correct bit flips and a fraction $H_2(\delta_p)$ to correct phase flips. The achievable rate for CSS codes is:

$$R = 1 - H_2(\delta) - H_2(\delta_p)$$

And because of the symmetrization of the bases, $\delta = \delta_p$.

Security of BB84

Theorem: For *any* attack by Eve, either the verification test fails with probability exponentially close to 1, or if the test succeeds then Bob's key agrees with Alice's with probability exponentially close to 1, and Eve's information about the key is exponentially small. Secure final key can be extracted from sifted key at the asymptotic rate: $R = \mathrm{Max}\left(1 - 2H_2(\delta), 0\right)$ where δ is the bit error rate found in the verification test.

This rate hits zero for $\delta = .1100$.

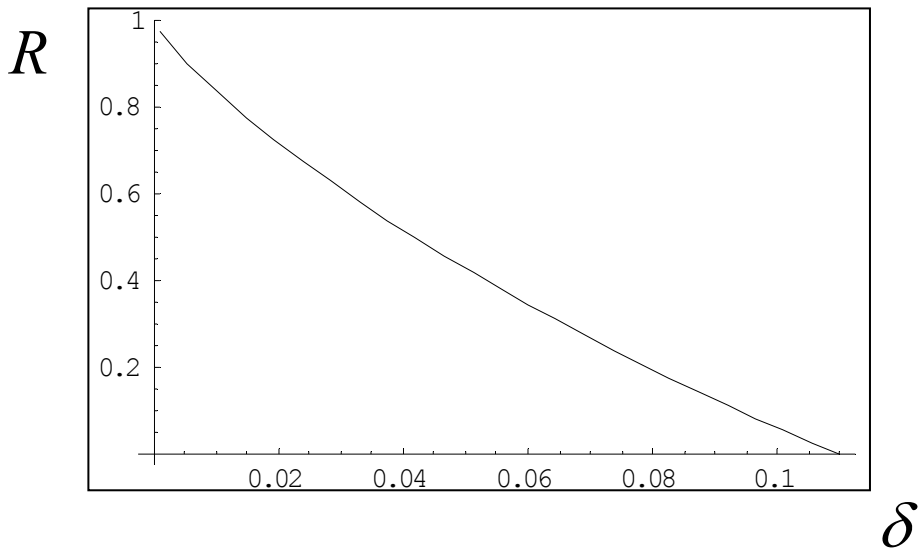

QKD for sale!

"Plug and play" quantum key distribution is *commercially available*:

Quantum Security... at last

Quantum Key Distribution System

Key distribution over optical fiber with absolute security

Main features

- First quantum cryptography system
- Security guaranteed by quantum physics
- Point-to point key distribution
- Standard optical fiber
- Distances up to 70 km
- Key rate up to 1000 bits/s
- Compact and reliable

Key distribution is a central problem in cryptography. Currently, public key cryptography is commonly used to solve it. However, these algorithms are vulnerable to increasing computer power. In addition, their security has never been formally proven.

Quantum cryptography exploits a fundamental principle of quantum physics - observation causes perturbation - to distribute cryptographic keys with absolute security.

id Quantique is introducing the first quantum key distribution system. It consists of an emitter and a receiver, which can be connected to PC's through the USB port.

id Quantique

10, rue Cingria 1205 Genève Switzerland
Tel: (+41) 022 702 69 29 Fax: (+41) 022 781 09 80
email: info@idquantique.com
web: http://www.idquantique.com

id Quantique
A Quantum Leap for Cryptography

QKD for sale!

"Plug and play" quantum key distribution is *commercially available*:

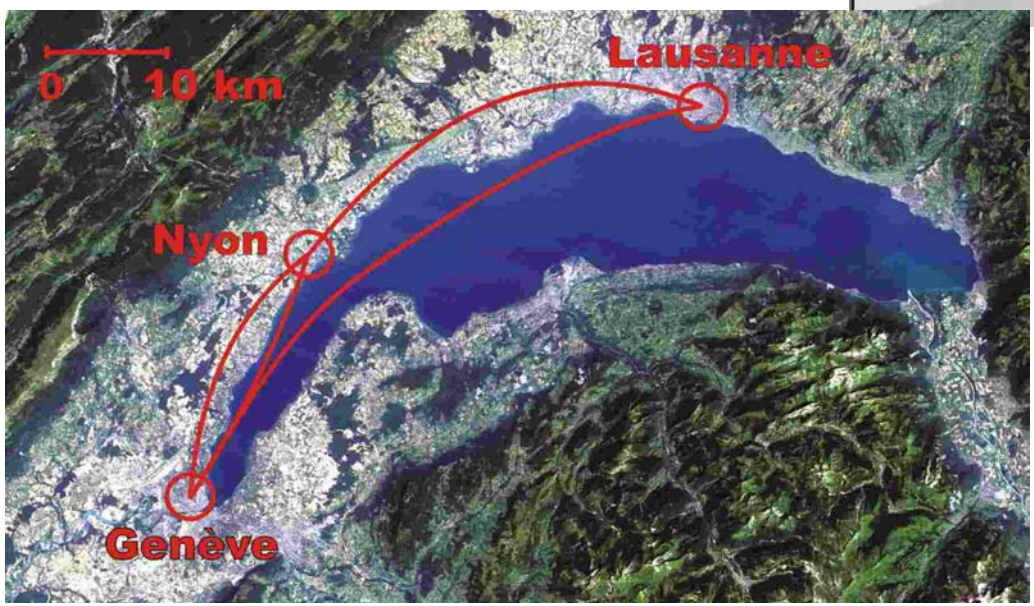

BB84 QKD has been achieved through a 67 km optical fiber under Lake Geneva.

Free-space QKD

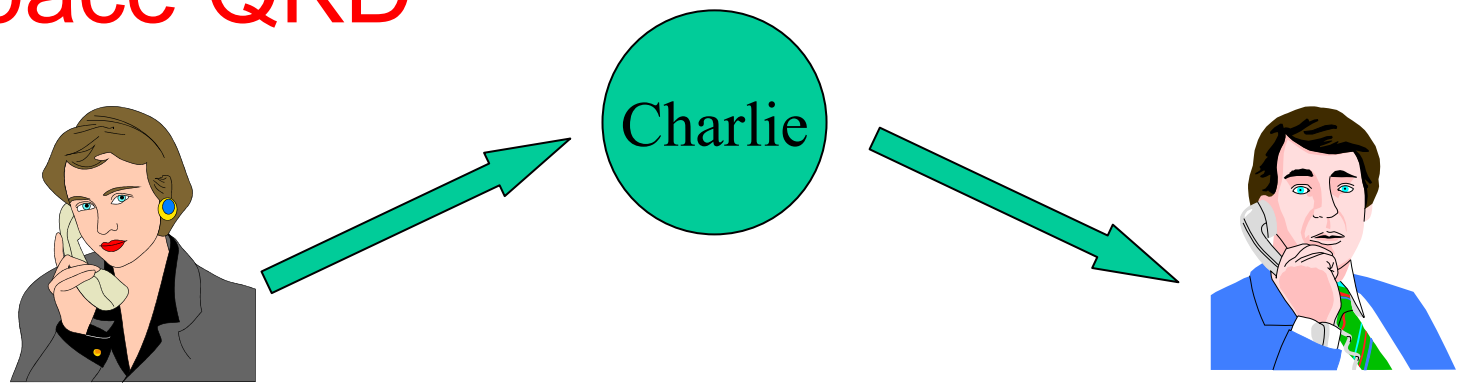

Using spectral, temporal, and spatial filtering, QKD can be executed by sending photons through the atmosphere in broad daylight! This has been achieved in Los Alamos, with sender and receiver on separate mesas, 1.6 km apart.

It should be feasible to do QKD between a party on the ground and a satellite in low earth orbit. If the satellite (Charlie) can be regarded as a trusted intermediary, it can generate key k_A shared with Alice and key k_B shared with Bob, and then announce $k_{AB} = k_A \oplus k_B$, from which Bob can recover $k_A = k_{AB} \oplus k_B$.

Security of “realistic” QKD

Sources of single photons are under development but are not readily available and are not used in current QKD systems. Instead, weak coherent states are used:

$$|\alpha\rangle = e^{-|\alpha|^2/2} e^{\alpha a^\dagger} |0\rangle \approx e^{-|\alpha|^2/2} \left(|0\rangle + \alpha |1\rangle + \left(\alpha^2 / \sqrt{2} \right) |2\rangle \right)$$

For small α, the signal is usually the vacuum, occasionally a single photon, and more rarely two or more photons.

If polarization encoding is used for key distribution, and a multiphoton state is sent, security may be compromised. Eve can skim off the extra photon(s), wait until Alice and Bob announce their bases, and then measure in the correct basis, obtaining perfect polarization information at no cost in disturbance. Our privacy amplification scheme must be sufficiently powerful (and the coherent states sufficiently weak), to nullify this advantage.

Security of "realistic" QKD

An imperfect source may leak to Eve information about what basis Alice is using. A source that emits weak coherent states instead of a single photons is one such case.

If Eve has some information about the basis, the symmetry between bit errors and phase errors is broken. The bit-error rate is measured in the verification test, but the phase error rate is inferred.

For single photons, the bit-error and phase error rates are equal. Call this rate p . For the multi-photons, pessimistically assume that the bit error rate is zero and the phase error rate is one. If the fraction of the signals that are multi-photons is Δ, then

$$\delta = (1-\Delta)p$$
$$\delta_p = (1-\Delta)p + \Delta = \delta + \Delta$$

We infer the achievable rate

$$R = 1 - H_2(\delta) - H_2(\delta + \Delta)$$

Other flaws in source and detector can be analyzed using similar methods. The protocol is secure if we have a "characterization" of the equipment --- in this case the fraction Δ of multi-photons.

Quantum repeaters?

BB84 QKD has been achieved through a 67 km optical fiber under Lake Geneva.

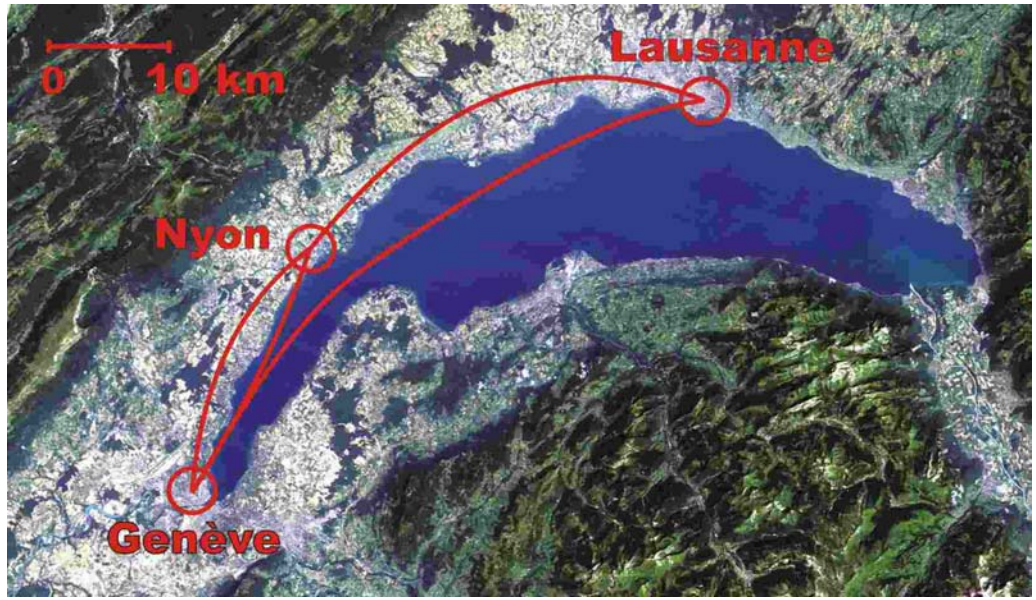

Today, the range of quantum key distribution is limited by absorption in optical fibers. Optical fiber used for classical communication comes equipped with repeaters that amplify signals, but the unknown signal states in e.g. BB84 cannot be amplified. What is needed are repeaters that use quantum error correction to protect encoded qubits from the effects of absorption. This could be a useful application for "intermediate scale" quantum computers that are powerful enough to implement quantum error recovery protocols.

Beyond key exchange: quantum coin flipping

Alice (in Calgary) and Bob (in Pasadena) want to flip a fair coin "over the telephone" --- they have just divorced, and need to decide who gets the house.

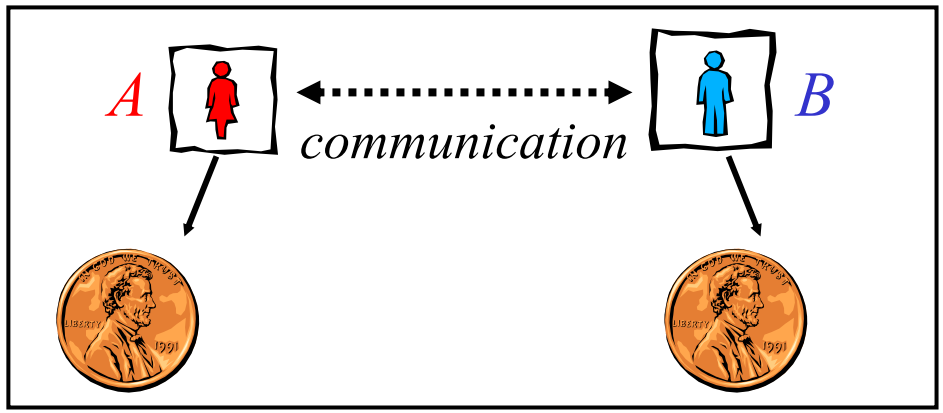

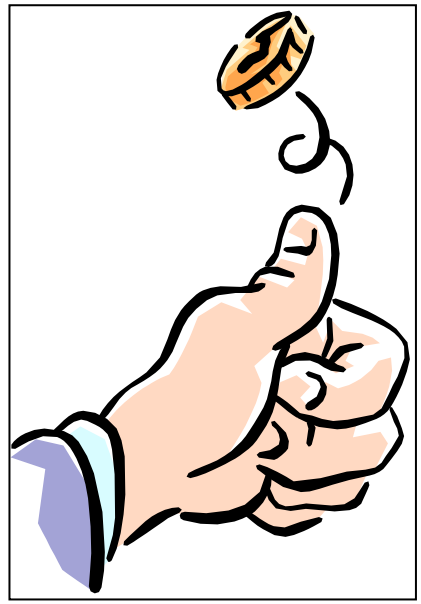

Alice could flip the coin and tell Bob the outcome, but Bob does not trust Alice. Alice could pick a random bit a, and Bob could pick a random bit b, where $a \oplus b$ is the outcome of the coin flip, but how can they ensure that neither party can cheat by delaying the selection of her/his bit until the other player's bit is known?

We'd like to devise a game in which Alice and Bob takes turns, where each player prints out the outcome of the coin flip at the end of the game. The players should agree on the outcome when they play honestly; furthermore, neither player should be able to bias the other player's outcome by cheating.

Beyond key exchange: quantum coin flipping

We'd like to devise a game in which Alice and Bob takes turns, where each player prints out the outcome of the coin flip at the end of the game. The players should agree on the outcome when they play honestly; furthermore, neither player should be able to bias the other player's outcome by cheating.

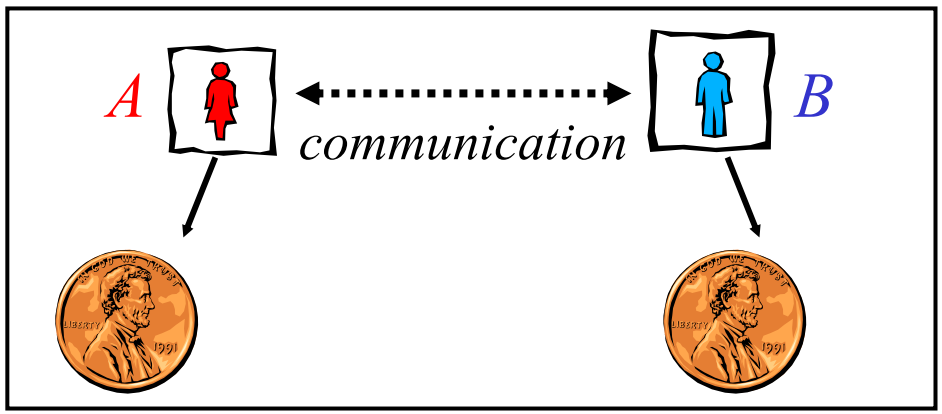

But there is no such *classical* protocol with *information-theoretic* security. Suppose Alice wins if the outcome is heads, and Bob wins if the outcome is tails. Then one player or the other has a strategy that ensures a win every time!

There are *computationally secure* classical coin-flipping protocols. For example, Alice can pick two large primes *p* and *q*, where either both *p* and *q* are congruent to +1 mod 4 ($a=0$), or both *p* and *q* are congruent to -1 mod 4 ($a=1$), and send the product *pq* to Bob. Alice reveals the prime factors only after receiving Bob's bit b. The outcome of the coin flip is $a \oplus b$.

This, and other computationally secure classical coin flipping protocols, are vulnerable to quantum attacks.

Quantum coin flipping

Ambainis, Spekkens-Rudolf



1) Alice chooses a random bit a.

For $a=0$, she sends to Bob $|0\rangle \pm |1\rangle$ (chosen equiprobably).

For $a=1$, she sends to Bob $|1\rangle \pm |2\rangle$ (chosen equiprobably).

2) Bob receives the state and stores it in his quantum memory. Bob chooses a random bit b and announces it to Alice.

3) Alice announces to Bob the value of a and the state that she sent, which Bob verifies. If the verification fails, then Alice is caught cheating and the protocol aborts.

4) If the verification succeeds, then the final outcome of the coin flip is $a \oplus b$.

Quantum coin flipping

Ambainis, Spekkens-Rudolf

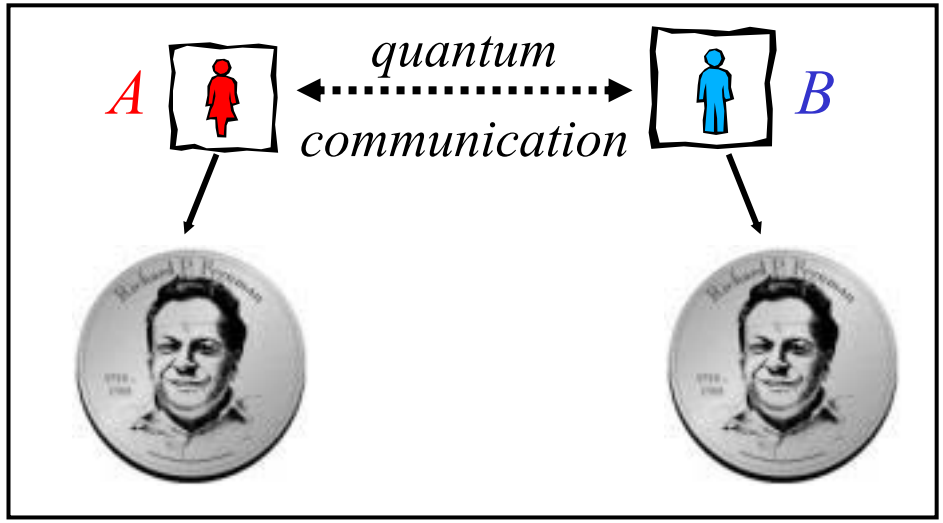

For $a=0$, A. sends $|0\rangle \pm |1\rangle$

For $a=1$, A. sends $|1\rangle \pm |2\rangle$

Bob can cheat by measuring the state sent by Alice before choosing his bit b. His best strategy is to perform an orthogonal measurement in the basis $\{|0\rangle, |1\rangle, |2\rangle\}$. The outcomes $|0\rangle, |2\rangle$ determine a unambiguously, but the outcome $|1\rangle$ reveals no information about a.

Hence, when Bob cheats he wins with probability ½ + ½ ½ =¾.

Quantum coin flipping

Ambainis, Spekkens-Rudolf

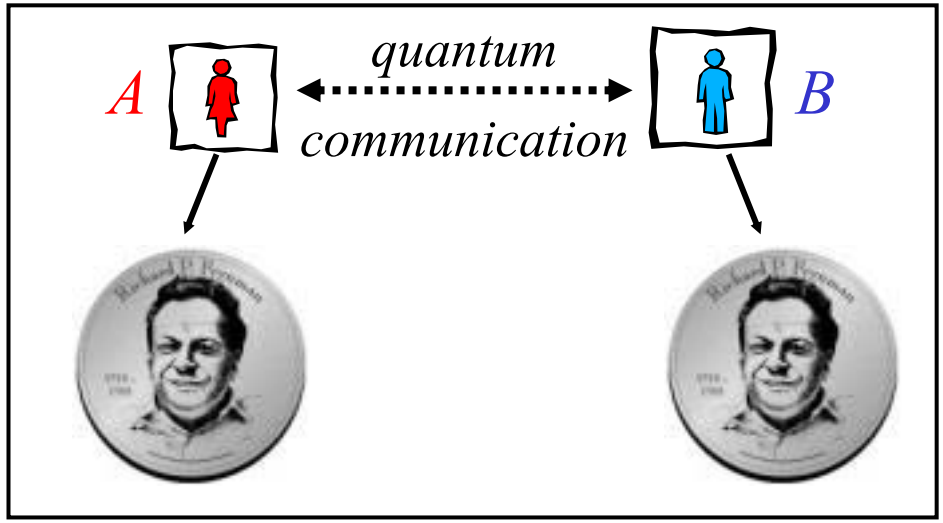

For $a=0$, A. sends $|0\rangle \pm |1\rangle$

For $a=1$, A. sends $|1\rangle \pm |2\rangle$

Alice can cheat by sending to Bob half of an entangled state, and measuring the half that she keeps after learning the value of Bob's bit b. Her best strategy is to prepare

$$\left(|00\rangle + 2|11\rangle + |22\rangle\right)/\sqrt{6}$$

and (if she wants a=0), to measure in the basis $\{|0\rangle \pm |1\rangle, |2\rangle\}$ She gets her desired outcome with probability 5/6. Bob's marginal state is $\left(|0\rangle \pm 2|1\rangle\right)/\sqrt{5}$ and so passes the verification test with probability 9/10. Altogether, then, the probability that Alice wins and her cheating escapes detection is (5/6)(9/10) = 3/4.

Quantum coin flipping

Ambainis, Spekkens-Rudolf

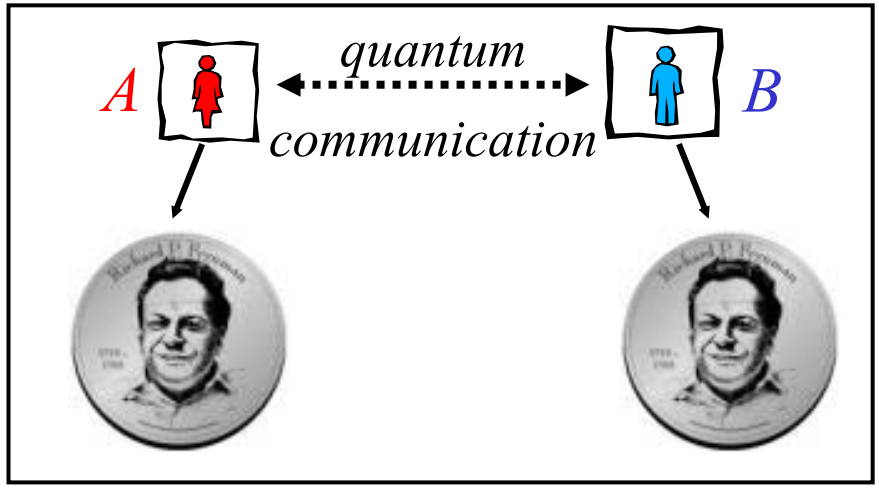

For $a=0$, A. sends $|0\rangle \pm |1\rangle$

For $a=1$, A. sends $|1\rangle \pm |2\rangle$

Bob can cheat, but the effectiveness of his cheating is limited because he cannot perfectly distinguish the nonorthogonal states that Alice might send.

Alice can cheat, but the effectiveness of her cheating is limited because operations on her half of an entangled state provide limited control over what Bob holds.

We say that the *bias* of the protocol is ε if the maximum probability of winning for a cheating player is $\frac{1}{2} + \varepsilon$. For any classical coin tossing protocol the bias is $\frac{1}{2}$. But for this quantum protocol the bias is $\frac{1}{4}$, even if the cheating player uses any strategy allowed by the laws of quantum physics.

Quantum coin flipping

Kitaev

Ambainis

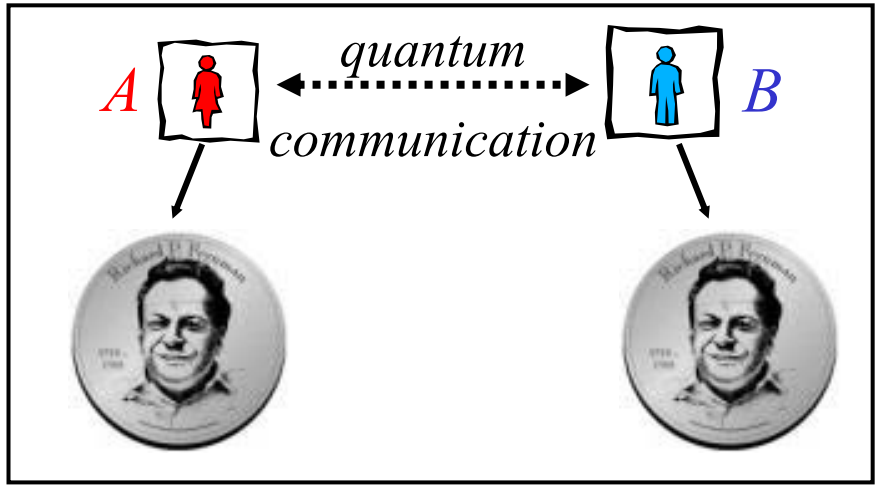

Strong coin flipping: Neither player can force *either* outcome with probability greater than ½ + ε .

Weak coin flipping: Neither player can force a *win* with probability greater than ½ + ε .

Kitaev: Strong coin tossing is *impossible* with bias

$$\varepsilon < 1/\sqrt{2} - 1/2 \cong .207.$$

Ambainis: Weak coin tossing with bias ε requires at least $\Omega\left(\log\log\left(1/\varepsilon\right)\right)$ rounds of communication.

Can the bias be arbitrarily small? An important open problem!

Weak coin flipping (Spekkens-Rudolf)

1) Alice prepares $|\psi\rangle_{AB}$ and sends half to Bob.
2) Bob performs a two-outcome POVM $\{E_0, E_1\}$ to determine the bit *b*.
3) If *b=0* (Bob wins), Bob sends *B* to Alice for verification. If *b=1* (Alice wins) Alice sends *A* to Bob for verification. If verification fails, one player is caught cheating, and the protocol aborts.
4) If verification succeeds, the outcome of the coin flip is *b*.

The best protocol of this type achieves bias $\varepsilon = 1/\sqrt{2} - 1/2 \cong .207.$

This can be realized with qubits:


Spekkens


Rudolf


Ambainis

$$|\psi\rangle_{AB} = 2^{-1/4}\,|00\rangle_{AB} + \sqrt{1-2^{-1/2}}\,|11\rangle_{AB},$$

$$E_0 = (1/\sqrt{2})\,|0\rangle\langle 0|, \quad E_1 = I - E_0.$$

Quantum error correction

1. Error models and error correction
2. Quantum error-correcting codes
3. Stabilizer codes
4. 5-qubit code and 7-qubit code
5. Fault-tolerant quantum computation
6. Accuracy threshold

Quantum cryptography

1. Cryptography and security
2. Quantum key distribution
3. The BB84 (four-state) protocol
4. Security proof using QECC
5. Quantum coin flipping