Simultaneous Communication Protocols with Quantum and Classical Messages

Oded Regev* Ronald de Wolf†

July 17, 2008

Abstract

We study the simultaneous message passing model of communication complexity, for the case where one party is quantum and the other is classical. We show that in a protocol where the first party sends q qubits and the second party sends c classical bits, the quantum message can be replaced by a *randomized* message of $O(qc)$ classical bits, as well as by a *deterministic* message of $O(qc \log q)$ classical bits. In particular, our results imply that quantum-classical protocols need to send $\Omega(\sqrt{n/\log n})$ bits/qubits to compute equality, and hence are not significantly better than classical-classical protocols (and are much worse than quantum-quantum protocols such as quantum fingerprinting). This essentially answers a recent question of Wim van Dam [7]. Our proofs rely heavily on earlier results due to Scott Aaronson [1, 2].

1 Introduction

We consider the simultaneous message passing (SMP) model of communication complexity. Here Alice receives input x and Bob receives input y. They each send one message to a third party, called the "referee." Given the two messages, the referee outputs a value which should equal the function value $f(x, y)$ with probability at least, say, 2/3.

We are interested in comparing classical and quantum SMP protocols. Consider for instance the equality function: $x, y \in \{0,1\}^n$, and $f(x, y) = 1$ iff $x = y$. If Alice and Bob do not share randomness, then this function exhibits an exponential quantum-classical gap: there is an SMP protocol for f where Alice and Bob each send $O(\log n)$ *quantum* bits to the referee [6], using a technique called "quantum fingerprinting." On the other hand, if the messages are classical, then $\Theta(\sqrt{n})$-bit messages are necessary and sufficient [3, 13, 4], as we will explain in Section 3.

Here we consider a question recently asked by Wim van Dam [7]: what happens if one of the messages (say Alice's) is quantum, while the other is restricted to be classical? We call such protocols *quantum-classical* SMP protocols. For instance, one may ask whether some variant of quantum fingerprinting still works if one of the two messages is classical.

*School of Computer Science, Tel-Aviv University, Tel-Aviv 69978, Israel. Supported by the Binational Science Foundation, by the Israel Science Foundation, by the European Commission under the Integrated Project QAP funded by the IST directorate as Contract Number 015848, and by a European Research Council (ERC) Starting Grant.

†CWI Amsterdam, The Netherlands. Partially supported by a Veni grant from the Netherlands Organization for Scientific Research (NWO), and by the European Commission under the Integrated Project Qubit Applications (QAP) funded by the IST directorate as Contract Number 015848.

Our main results here say that the quantum message can be "simulated" by a classical message that is only moderately larger. More specifically, we show that the following hold for protocols computing a Boolean function:

- In Theorem 4 we show that a private-coin (resp. public-coin) quantum-classical protocol where Alice sends q qubits and Bob sends c bits, gives a private-coin (resp. public-coin) classical-classical protocol where Alice sends a randomized $O(qc)$-bit message and Bob sends a randomized c-bit message.

- In Corollary 6 we show that a private-coin quantum-classical protocol where Alice sends q qubits and Bob sends c bits, gives a private-coin classical-classical protocol where Alice sends a *deterministic* $O(qc\log(q))$-bit message and Bob sends a randomized c-bit message.

Our proofs rely heavily on earlier results in the one-way communication model by Aaronson [1, 2].

The latter result implies that quantum-classical private-coin protocols for equality need to send $\Omega(\sqrt{n/\log n})$ bits and/or qubits. This is not significantly better than the $\Theta(\sqrt{n})$ bits that are necessary and sufficient for classical protocols.

The results above are rather unusual in communication complexity, where typically introducing quantum elements gives exponentially more power. The exponential improvements given by the quantum-quantum fingerprinting protocol is a prime example of this. This is also the case with one-way communication complexity, where we know Boolean functions whose quantum complexity is exponentially smaller than their classical complexity [9]. As another example, Gavinsky et al. [9] exhibit a function with an exponential gap between classical-classical SMP protocols with shared entanglement and those with only shared randomness. Yet another example is obtained when considering *relations* instead of *functions* in our model [5]. We elaborate on this in Section 7.

Remark: After finishing a first version of this paper, we learned that Dmitry Gavinsky had independently (and much earlier) obtained essentially the same results, but did not publish them. After some friendly discussion, he declined our offer to co-author this paper. The observation of functions vs relations given in Section 7 is due to him, and included here with his permission.

2 Preliminaries

We assume familiarity with quantum computing [14] and with the basic notions of classical and quantum communication complexity [11, 15]. Informally, the setting of communication complexity is as follows. Alice receives some input $x \in X$, Bob receives some input $y \in Y$, and together they want to compute some function $f(x,y)$, with error probability at most $1/3$ for all (x,y) in some domain $\mathcal{D} \subseteq X \times Y$. If $\mathcal{D} = X \times Y$ then the problem is called *total*, otherwise it is called *partial*. A *Boolean* function f has range $\{0,1\}$. The *communication matrix* M_f corresponding to f is the $|X| \times |Y|$ matrix whose (x,y)-entry equals $f(x,y)$ if $(x,y) \in \mathcal{D}$, and equals '*' otherwise.

In the simultaneous message passing (*SMP*) model, Alice and Bob each send a message to a third party (called the *referee*), who then computes the output. In the *one-way* model Alice sends one message to Bob (no referee here), while in the *two-way* model they can interact arbitrarily. The *cost* of a communication protocol is its total communication on the worst-case input. The (bounded-error) communication complexity of f (in one of the above models) is the minimal cost among all protocols that compute f with probability of error at most $1/3$ for each input.

Classical protocols may be deterministic or randomized. When dealing with randomness, we have to distinguish between *private coins* and *public coins*.[1] The former are visible only to individual parties, while the latter are shared among all parties and may help them coordinate their actions. In one-way and two-way models, the difference between public-coin and private-coin communication complexity is at most an additive $O(\log n)$ bits [12], but in the SMP model it can make a big difference. In particular, while the private-coin SMP complexity of equality is $\Theta(\sqrt{n})$ bits, its public-coin complexity is constant: the protocol picks a shared random n-bit string r, Alice and Bob send the inner product of their input with r (mod 2) to the referee, who checks whether the two received bits are equal.

The following theorem can be derived from the proof of [2, Theorem 1.4]. Here by a measurement operator E we mean a positive semidefinite matrix with eigenvalues in $[0,1]$. The *acceptance probability* of the two-outcome measurement with operators E and $I-E$ on density matrix ρ is $p = \mathrm{Tr}(E\rho)$.

Theorem 1 (Aaronson)**.** *For all $\delta > 0$ the following holds. If Alice has a q-qubit density matrix ρ and Bob has a measurement operator $E \in \{E_b\}_{b\in\{0,1\}^c}$, then Alice can send Bob an $O(qc)$-bit randomized message to enable him to output a value that is within $\pm\delta$ of $\mathrm{Tr}(E\rho)$ with probability at least $1-\delta$ (the constant in the $O(\cdot)$ depends on δ, and no public coin is used).*

3 Warmup: replacing a randomized message by deterministic

Our goal in this paper is to replace a quantum message by a randomized or deterministic one that is not much bigger. Let us first consider a known case: replacing a *randomized* message by a *deterministic* one. Babai and Kimmel [4] showed the following. If there is a bounded-error private-coin SMP protocol for a Boolean function f, where Alice sends c_A bits and Bob sends c_B bits, then there exists another bounded-error private-coin SMP protocol for f where Alice *deterministically* sends $O(c_A c_B)$ bits, and Bob (who is still randomized) sends c_B bits.

As an application of this result, note that in any bounded-error protocol for equality where Alice is deterministic, she needs to send a different message for each of the 2^n inputs x: if she sends the same message for x and for x', then the referee will have the same acceptance probability on inputs (x,x) and (x',x) and hence on at least one of those pairs he will err with probability at least $1/2$. This implies that in any bounded-error private-coin SMP protocol for equality, $n \leq O(c_A c_B)$, and hence we obtain the bound $c_A + c_B \geq \Omega(\sqrt{n})$ mentioned in the introduction.

We sketch a simple proof of the Babai-Kimmel result. Consider any Boolean function f, partial or total, and a bounded-error private-coin SMP protocol for f. Let $r(a,b) \in [0,1]$ denote the referee's acceptance probability if he receives message $a \in \{0,1\}^{c_A}$ from Alice and $b \in \{0,1\}^{c_B}$ from Bob. We use A_x for the random variable which is Alice's message (its distribution depends on her input x), and similarly B_y for Bob's message. Note that for every input x, y where f is defined, the acceptance probability $\mathbb{E}_{a\sim A_x, b\sim B_y}[r(a,b)]$ approximates the function value: $|f(x,y) - \mathbb{E}_{a\sim A_x, b\sim B_y}[r(a,b)]| \leq 1/3$.

We now modify the protocol as follows. Let Alice send her (probabilistic) message $s \cdot c_B$ times, for some integer s to be determined later, at a total communication cost $s c_A c_B$. For a fixed message b from Bob, the referee can obtain an approximation $\widetilde{p}_b$ of the quantity $p_b = \mathbb{E}_{a\sim A_x}[r(a,b)]$, such that $\Pr[|\widetilde{p}_b - p_b| > 1/100] \ll 2^{-c_B}$. Here $\widetilde{p}_b$ is just the average of $r(a_i,b)$ over all messages a_i

[1] We will not consider models with shared entanglement here.

received from Alice. Choosing s a sufficiently large constant, the Chernoff bound implies that $\widetilde{p}_b$ is within $1/100$ of its expectation p_b, except with probability $\ll 2^{-c_B}$. Hence by a union bound, the referee can obtain, with probability close to 1, approximations $\widetilde{p}_b$ for all $b \in \{0,1\}^{c_B}$ that are all simultaneously $1/100$-close to their true value. But now the referee can compute $f(x,y)$ for each y of his choice, since he can compute $\mathbb{E}_{b \sim B_y}[\widetilde{p}_b]$ exactly, and we have

$$|f(x,y) - \mathbb{E}_{a \sim A_x, b \sim B_y}[r(a,b)]| \leq 1/3 \quad \text{and} \quad |\mathbb{E}_{a \sim A_x, b \sim B_y}[r(a,b)] - \mathbb{E}_{b \sim B_y}[\widetilde{p}_b]| \leq 1/100.$$

Thus the referee obtains (with probability very close to 1) complete information about the row of the communication matrix corresponding to x. But Alice has given him this information by sending a message of only $O(c_A c_B)$ bits. This implies that Alice can actually *deterministically* tell the referee which row she has using $O(c_A c_B)$ bits: because the referee's behavior as described above is deterministic, Alice can actually compute which of her messages will give the referee correct results, and just send him that message. We have reproved:

Theorem 2 (Babai & Kimmel)**.** *Let f be a (possibly partial) Boolean function. If there is a bounded-error private-coin SMP protocol for f where Alice sends c_A bits and Bob sends c_B bits, then there is a bounded-error private-coin SMP protocol for f where Alice* deterministically *sends $O(c_A c_B)$ bits, and Bob (who is still randomized) sends c_B bits.*

Actually, Babai and Kimmel showed something slightly stronger, namely that both Alice's and Bob's randomized message can simultaneously be replaced by deterministic messages of length $O(c_A c_B)$.

Theorem 3 (Babai & Kimmel)**.** *Let f be a (possibly partial) Boolean function. If there is a bounded-error private-coin SMP protocol for f where Alice sends c_A bits and Bob sends c_B bits, then there is a deterministic SMP protocol for f where Alice and Bob send $O(c_A c_B)$ bits.*

This theorem is essentially tight for the equality problem, where the deterministic communication complexity is $2n$. Consider the following bounded-error private-coin protocol, adapted from [3]. Alice and Bob fix, beforehand, a good error-correcting code $C : \{0,1\}^n \to \{0,1\}^m$ with $m = O(n)$. For $m = c_A c_B$, they can view the codewords as $c_A \times c_B$ Boolean matrices. Alice sends a random column of $C(x)$ together with its index, sending $c_A + \log c_B$ bits in total. Bob sends a random row with its index, sending $c_B + \log c_A$ bits. The referee checks whether the row and column agree in the one point where they intersect. If $x = y$ then the two bits are the same, otherwise they differ with constant probability. Repeating a constant number of times, we obtain a bounded-error private-coin SMP protocol where Alice sends $O(c_A + \log c_B)$ bits and Bob sends $O(c_B + \log c_A)$ bits.

Note that Theorem 2 fails spectacularly for *public-coin* SMP protocols: the bounded-error public-coin SMP complexity of equality is constant, while a deterministic player needs to send n bits, no matter what the other player sends.

4 Replacing a quantum message by a randomized one

We now prove an analogue of the Babai-Kimmel theorem for quantum-classical SMP protocols: the "quantum leg" of the protocol can be replaced by a randomized message.

Theorem 4. *Let f be a (possibly partial) Boolean function. If there is a private-coin (resp. public-coin) bounded-error quantum-classical SMP protocol for f where Alice sends q_A qubits and Bob*

sends c_B classical bits, then there is a private-coin (resp. public-coin) bounded-error SMP protocol for f where Alice sends $O(q_A c_B)$ classical bits and Bob sends c_B classical bits.

Proof: We prove the theorem for private-coin protocols. The proof for public-coin protocols is essentially the same, just fix the shared randomness at the start of the argument and average over it at the end. In general the three-party protocol has the following form: Alice sends the referee a q_A-qubit density matrix ρ_x, while Bob sends a classical message $b \in \{0,1\}^{c_B}$, whose distribution depends on his input y. The referee then measures ρ_x with a measurement operator E_b, and outputs 1 if the measurement accepts (which happens with probability $p_b = \mathrm{Tr}(E_b \rho_x)$).

The SMP protocol promised by the theorem is as follows: Bob sends a c_B-bit message b, exactly as in the original quantum-classical protocol. Using Theorem 1, Alice sends to the referee a randomized message of $O(q_A c_B)$ bits to enable him to obtain with probability at least $1-\delta$ an approximation $\widetilde{p}_b$ to $p_b = \mathrm{Tr}(E_b \rho_x)$ to within $\pm\delta$, where $\delta = 1/10$. Finally, the referee outputs 1 with probability $\widetilde{p}_b$, and 0 otherwise. The overall error probability is at most 2δ worse than in the original protocol. ■

Note the difference between the above two proofs. The proof of Theorem 2 obtains approximations p_b for all $b \in \{0,1\}^{c_B}$ simultaneously, which enables the referee to learn $f(x,y)$ for each y (i.e., learn the whole row of the communication matrix that corresponds to Alice's input x). On the other hand, the proof of Theorem 4 only obtains an approximation p_b for the specific b that the referee received from Bob, which enables the referee to predict $f(x,y)$ for the specific input y that Bob holds.

5 Replacing a quantum message by a deterministic one

By combining Theorem 4 with the Babai-Kimmel theorem (Theorem 2), we see that in every quantum-classical private-coin SMP protocol with q_A qubits and c_B bits, we can replace Alice's message by a *deterministic* message of $O(q_A c_B^2)$ bits. However, we can obtain something that is usually stronger, namely a deterministic message of $O(q_A c_B \log q_A)$ bits. The crucial tool is the following result, which is an extension of a result of Aaronson [1, Theorem 3.4]:

Theorem 5. *Suppose Alice has the classical description of an arbitrary q-qubit density matrix ρ, and Bob has 2^c measurements operators $\{E_b\}_{b \in \{0,1\}^c}$. There is a deterministic message of $O(qc \log q)$ bits from Alice that allows Bob to approximate $p_b = \mathrm{Tr}(E_b \rho)$ to within $\pm\delta$, simultaneously for all $b \in \{0,1\}^c$.*

It is interesting to compare this with Theorem 1. While Theorem 1 allows us to approximate one p_b to within $\pm\delta$ (with some small probability of error) using an $O(qc)$-bit message, Theorem 5 allows us to approximate *all* p_b to within $\pm\delta$ (*without* probability of error) at the expense of increasing the message length by a factor $\log q$. Theorem 5 generalizes Aaronson's [1, Theorem 3.4], which proves the special case where $\mathrm{Tr}(E_b \rho)$ is close to 0 or 1 for all b. Our proof is a modification of his. We conjecture that the $\log q$ factor is not needed.

Proof: Suppose Alice sends $r = O(\log q)$ many copies of her state. Let $\rho' = \rho^{\otimes r}$ be the state she

sends, and $K = rq = O(q \log q)$ its total number of qubits. Define the observable[2]

$$F_b = \frac{1}{r} \sum_{j=1}^{r} E_b^{(j)},$$

where $E_b^{(j)}$ applies E_b to the jth copy. This measures the fraction of successes if you separately measure each copy of ρ with E_b. By a Chernoff bound, the outcome of this measurement applied to ρ' will be at most $\delta/2$ away from its expectation $\mathrm{Tr}(E_b\rho)$, except with probability $1/\mathrm{poly}(q)$.

Alice's classical message. Consider all $b = 1, \ldots, 2^c$ in order. We will sequentially build a sequence of K-qubit density matrices ρ_b, one for each E_b. Call b *good* if $|\mathrm{Tr}(F_b\rho_b) - p_b| \leq \delta$; call b *bad* otherwise. Note that if Bob has a classical description of a good ρ_b, then he can approximate p_b to within $\pm\delta$ (since he knows what F_b is). We start with the completely mixed state: $\rho_1 = I/2^K$ and define the subsequent ρ_b one by one, as follows. If b is good, then define ρ_{b+1} to be equal to ρ_b. If b is bad, Alice appends the pair $(b, \widetilde{p}_b)$ to her message, where $\widetilde{p}_b$ is the $\log(1/\delta) + O(1)$ most significant bits of p_b, so $|\widetilde{p}_b - p_b| \ll \delta$. In this case, let M_b be the projector on the subspace spanned by the eigenvectors of F_b with eigenvalues in the interval $[\widetilde{p}_b - \delta/2, \widetilde{p}_b + \delta/2]$, and let ρ_{b+1} be the renormalized projection of ρ_b on this subspace.[3] Continuing all the way to $b = 2^c$, we obtain a message $(b_1, \widetilde{p}_{b_1}), \ldots, (b_T, \widetilde{p}_{b_T})$ for some T. We need to show two things: (1) this message enables Bob to approximate all p_b to within $\pm\delta$, and (2) $T = O(K)$, which implies that the message length is $O(qc \log q)$ bits.

Why this works. Note that Bob knows which $b \in [2^c]$ are bad, since those b are exactly the ones in Alice's message. Bob can in fact compute the whole sequence $\rho_1, \ldots, \rho_{2^c}$ given the message: $\rho_1 = I/2^K$; if b is good then $\rho_{b+1} = \rho_b$; if b is bad then $(b, \widetilde{p}_b)$ is part of Alice's message and ρ_{b+1} can be computed from this information. Suppose Bob wants to approximate $p_b = \mathrm{Tr}(E_b\rho)$. If b is good then by definition $|\mathrm{Tr}(F_b\rho_b) - p_b| \leq \delta$ and Bob can calculate $\mathrm{Tr}(F_b\rho_b)$. If b is bad, then the pair $(b, \widetilde{p}_b)$ is part of Alice's message, so Bob knows p_b with sufficient precision. Hence Bob can approximate all p_b up to $\pm\delta$, for all b simultaneously.

Why the message is not too long. Here we show $T = O(K)$. Define $\eta = 1 - \delta/(2 - \delta)$ and $t = \lceil (K+1)/\log(1/\eta) + 1 \rceil = O(K)$. Suppose, by way of contradiction, that $T \geq t$. We consider the sequence $b_1, \ldots, b_t$ of the first t bad b's. Let

$$p = \mathrm{Tr}\left(M_{b_t} \cdots M_{b_1} \frac{I}{2^K} M_{b_1} \cdots M_{b_t} \right)$$

be the probability that all t measurements succeed if we start with the completely mixed state and sequentially measure $M_{b_1}, \ldots, M_{b_t}$. We will derive contradicting upper and lower bounds on p.

First, the upper bound on p. If we sequentially measure $M_{b_1}, \ldots, M_{b_t}$, starting from the completely mixed state, and if all t measurements succeed, then we exactly have the sequence of density matrices $\rho_{b_1} = I/2^K, \ldots, \rho_{b_t}, \rho_{b_t+1}$. Note that if ρ_b is bad, then $\mathrm{Tr}(M_b\rho_b) \leq \eta$ by Markov's inequality.[4] Hence the probability that all t measurements succeed is $p \leq \eta^t$.

[2] An *observable* F is a Hermitian matrix that describes a measurement, as follows. By diagonalization we can write $F = \sum_i \lambda_i P_i$, where P_i is the projector on the eigenspace corresponding to eigenvalue λ_i. These eigenspaces are all orthogonal to each other and $\sum_i P_i = I$. The corresponding measurement on a density matrix ρ gives outcome λ_i with probability $\mathrm{Tr}(P_i\rho)$. Hence the *expectation* of the measurement is $\sum_i \lambda_i \mathrm{Tr}(P_i\rho) = \mathrm{Tr}(F\rho)$.

[3] The fact that this projection is nonzero (and hence can be renormalized to have trace 1) follows from the argument in the "Second, the lower bound on p" paragraph below.

[4] To see this, let X denote the random variable which is the outcome of measuring ρ_b with the observable F_b.

Second, the lower bound on p. By a Chernoff bound, we have $\mathrm{Tr}(M_b\rho') \geq 1 - 1/\mathrm{poly}(q)$. This allows us to measure ρ' with M_b while hardly disturbing the state (see for instance the "almost as good as new lemma" [1, Lemma 2.2]). If we measure each of M_b, for the first t bad b's in sequence, starting in ρ', then with probability at least $1/2$ all measurements will succeed. However, the completely mixed state can be written as $\frac{I}{2^K} = \frac{1}{2^K}\rho' + (1 - \frac{1}{2^K})\rho''$ where ρ'' is orthogonal to ρ'. Hence if we start from $I/2^K$, then the probability of all measurements succeeding is $p \geq 1/2^{K+1}$.

Combining the bounds of the last two paragraphs together with our value of t gives a contradiction. ■

Now consider a quantum-classical private-coin SMP for a function f (total or partial). The previous theorem enables us to replace Alice's q_A-qubit message by a deterministic message of $O(q_A c_B \log q_A)$ bits. This message allows the referee to closely approximate $p_b = \mathrm{Tr}(E_b\rho_x)$. Since $f(x,y) \approx \mathbb{E}_{b\sim B_y}[p_b]$, this allows the referee to compute $f(x,y)$ for every y, so he now knows the row of the communication matrix corresponding to Alice's input.

Corollary 6. *Let f be a (possibly partial) Boolean function. If there is a bounded-error quantum-classical private-coin SMP protocol for f where Alice sends q_A qubits and Bob sends c_B bits, then there is a bounded-error private-coin SMP protocol for f where Alice* deterministically *sends $O(q_A c_B \log q_A)$ bits, and Bob (who is still randomized) sends c_B bits.*

As argued in Section 3, in a bounded-error protocol for equality where Alice is deterministic, she needs to send at least n bits. Hence we obtain the following lower bound on quantum-classical private-coin SMP protocols for equality, which is tight up to the $\sqrt{\log n}$-factor.

Corollary 7. *Every quantum-classical private-coin SMP protocol for equality has communication complexity $\Omega(\sqrt{n/\log n})$.*

To make the protocol of Corollary 6 fully deterministic, it remains to replace Bob's randomized message by a deterministic one. Since Alice's deterministic message gives the referee full information about her row of the communication matrix, the problem reduces to a deterministic one-way protocol from Bob to the referee. The minimal number of bits that Bob needs to send in such a deterministic one-way protocol is denoted by $D^{1,B\to A}(f)$. We have proved

Corollary 8. *Let f be a (possibly partial) Boolean function. If there is a bounded-error quantum-classical private-coin SMP protocol for f where Alice sends q_A qubits and Bob sends c_B bits, then there is a deterministic SMP protocol for f where Alice sends $O(q_A c_B \log q_A)$ bits and Bob sends $D^{1,B\to A}(f)$ bits.*

6 Tightness

The example of the equality function shows that Theorems 5 and Corollary 6 are essentially tight.

We do not know whether Theorem 4 is close to optimal, but at least it shows that the gap between quantum-classical and classical-classical SMP protocols is at most polynomial. The following communication problem, adapted from [5, 10, 9], presents an interesting quantum-classical

The statement that b is bad is equivalent to the statement $|\mathbb{E}[X] - p_b| > \delta$ (assume $p_b = \widetilde{p}_b$ for simplicity). Let $\tau = \mathrm{Tr}(M_b\rho_b) = \Pr[|X - p_b| \leq \delta/2]$. Assume $\mathbb{E}[X] > p_b$ (the other case is similar). Then we have $p_b + \delta < \mathbb{E}[X] \leq \tau(p_b + \delta/2) + (1-\tau)\cdot 1$. This implies $\tau < 1 - \delta/(2 - 2p_b - \delta) \leq 1 - \delta/(2-\delta) = \eta$.

public-coin protocol that uses about $n^{1/3}$ qubits. We do not know an equally efficient classical-classical public coin protocol for this problem; the best one we know sends about $\sqrt{n}$ bits. This suggests that quantum-classical SMP protocols can at least have *some* polynomial advantage over classical-classical protocols.

The problem is as follows. Let n be an even integer. Alice receives $x \in \{0,1\}^n$, while Bob receives a perfect matching M (i.e., a partition of $[n] = \{1,\ldots,n\}$ into $n/2$ disjoint pairs, called "edges"), and a string $w \in \{0,1\}^{n/2}$ whose bits are indexed by the edges in M. We can view the edges of M as rows (of weight 2) in an $n/2 \times n$ matrix M over $GF(2)$. Then the matrix-vector product Mx is the $n/2$-bit string obtained by taking, for each edge (i,j) of M in order, the XOR $x_i \oplus x_j$. The promise is that the Hamming distance between w and Mx is either at most $n/6$ or at least $n/3$, and the function value is 1 in the first case and 0 in the second.

One can show easily that the deterministic complexity of the problem is $\Omega(n)$, as follows. By the probabilistic method, there exists a set $S \subseteq \{0,1\}^n$ of size $|S| = 2^{\Omega(n)}$ such that all distinct $x, x' \in S$ are at distance around $n/2$. But then for each distinct $x, x' \in S$ we can find a matching M where the $n/2$-bit strings Mx and Mx' have distance close to $n/2$ (pick as many edges as possible that have one endpoint in a bitlocation where x and x' agree, and one endpoint where they disagree). Putting $w = Mx$, we have $f(x, M, w) = 1$ but $f(x', M, w) = 0$. Hence in a deterministic protocol, Alice will need to send a different message for each of the $x \in S$. Therefore the deterministic SMP (and even one-way) communication complexity of this function is $\Omega(n)$.

On the other hand, here is a bounded-error public-coin SMP protocol where Alice sends $q_A \approx n^{1/3}$ qubits and Bob sends $c_B \approx n^{1/3}$ bits. Alice and Bob use the public coin to select a random subset $S \subseteq [n]$ of about $n^{2/3}$ elements from $[n]$. Now with high probability, $M \cap (S \times S)$ will contain $\Theta(n^{1/3})$ edges. Alice sends the referee $n^{1/3}$ copies of the uniform superposition $\frac{1}{\sqrt{|S|}} \sum_{i \in S} (-1)^{x_i} |i\rangle$. Bob sends over $\Theta(n^{1/3})$ edges in $M \cap (S \times S)$, together with the corresponding bits of w. The referee constructs two-dimensional measurement operators from the edges he received from Bob, and measures each of the quantum states with them. With probability close to 1, one of those measurements will succeed and give him a bit of the string Mx. Since the location of that bit in Mx is random, comparing that bit with the corresponding bit of w (which is part of Bob's message), gives the referee the function value with probability at least 2/3.[5]

7 Functional separations versus relational separations

As mentioned before, Theorem 4 implies that the gap between quantum-classical and classical-classical SMP protocols is at most polynomial for any Boolean function. In contrast, Bar-Yossef et al. [5] exhibited a *relational* problem for which quantum-classical SMP protocols are *exponentially* better than classical-classical SMP protocols with a public coin. In a relational problem, for each input pair (x, y) there is a *set* of valid outputs z. In the case of [5], Alice receives an arbitrary string $x \in \{0,1\}^n$ and Bob receives a perfect matching M on $[n]$ from a set of $n/2$ possible perfect matchings (assume n is even). Their goal is to output any z of the form $(i, j, x_i \oplus x_j)$ where $(i,j) \in M$. Bar-Yossef et al. exhibit a quantum-classical SMP protocol that solves this problem with success probability 1, with a $\log n$-qubit message from Alice to the referee and a $\log n$-bit message from Bob. No public coin is needed. In contrast, they proved a $\Theta(\sqrt{n})$ bound for classical-

[5]The best classical protocol that we know works similarly. It selects a set S of about $\sqrt{n}$ elements, and Alice sends the corresponding bits of x to the referee at the expense of about $\sqrt{n}$ bits of communication.

classical public-coin protocols for this relational problem.

To summarize, we see that when comparing the quantum-classical SMP model to the classical-classical SMP model, one can obtain an exponential separation for a relation but not for a Boolean function. In this section we present an observation due to Dmitry Gavinsky [8], showing that such a situation *cannot occur* for purely classical models. More precisely, we show that for models obeying the *Yao principle* [16] (defined next), any separation for a relation implies a similar separation for a function (and sometimes even for a Boolean function). Notice that the converse implications clearly holds, i.e., a separation for a functional problem is also a separation for a relational problem, since functions are a special case of relations.

Consider any computational model that has a class of deterministic algorithms (or protocols), each of a certain cost. A *randomized* algorithm in such a model is a probability distribution over deterministic algorithms. The Yao principle states the equality of two different complexities: (1) the ε-error complexity (the minimal cost of randomized algorithms whose error probability is at most ε on every input) and (2) the ε-error distributional complexity under the hardest input distribution μ (the minimal cost of deterministic protocols that have error probability at most ε under μ). Because of the minimax theorem from game theory, allowing shared randomness in one's communication model is a sufficient condition for the Yao principle to hold.

Now assume any two computational models, both obeying the Yao principle. Call these models "a" and "b", respectively. Let $R_{a,\varepsilon}$ and $R_{b,\varepsilon}$ denote the ε-error complexities in these two models. Assume we have a separation between these two models, showing that $R_{a,2\varepsilon}(P)$ is greater than $R_{b,\varepsilon}(P)$ for some relational problem P. We will next show how to construct a function f for which there is a separation between $R_{a,\varepsilon}(f)$ and $R_{b,\varepsilon}(f)$ that is at least as large.

Let μ be a worst-case input distribution for the distributional $R_{a,2\varepsilon}$-complexity of relation P. That is, any deterministic protocol in the "a" model solving P with distributional error at most 2ε under μ, must use at least $R_{a,2\varepsilon}(P)$ bits of communication. Next, by the Yao principle, there is a deterministic protocol in the "b" model solving P with error at most ε under distribution μ, using at most $R_{b,\varepsilon}(P)$ bits of communication. Being deterministic, this protocol necessarily computes some function with error probability 0, call it f. Hence we see that for this function, $R_{b,\varepsilon}(P) \geq R_{b,0}(f)$. Moreover, note that the probability under μ that $f(x,y)$ is not a valid answer for P, is at most ε.

To complete the argument, we now show that $R_{a,\varepsilon}(f) \geq R_{a,2\varepsilon}(P)$. Consider an optimal ε-error protocol for f in the "a" model with complexity $R_{a,\varepsilon}(f)$. By the Yao principle, there exists a deterministic protocol for f in the "a" model with distributional error at most ε under μ, and the same complexity. By the union bound, the same protocol computes P with distributional error at most 2ε under μ. The desired inequality now follows from the choice of μ. In sum, we have found a function f with $R_{a,\varepsilon}(f) \geq R_{a,2\varepsilon}(P) > R_{b,\varepsilon}(P) \geq R_{b,0}(f) \geq R_{b,\varepsilon}(f)$.

The separation we obtained above is for a functional problem, but not necessarily a *Boolean* one. We now observe that if the R_a/R_b-separation we start with is sufficiently strong, and the number of output bits of P is not too large, then one can obtain a *Boolean* function with a strong R_a/R_b-separation. Assume f has a k-bit output, and assume that $g : \{0,1\}^k \to \{0,1\}^{10k}$ is an error correcting code with constant rate and constant relative distance (which is known to exist). Let $f_1, \ldots, f_{10k}$ be the Boolean functions representing the bits of $g(f(\cdot))$. Having ε-error protocols for each of these Boolean functions, each of complexity c, implies an ε-error protocol for f of complexity $O(ck)$. Hence for at least one j we have $R_{a,\varepsilon}(f_j) = \Omega(R_{a,\varepsilon}(f)/k)$.

Acknowledgments

We thank Wim van Dam for asking the question that prompted this research, Ben Toner for useful discussions, and Dmitry Gavinsky for permission to include his observation about Boolean functions vs relations of Section 7.

References

[1] S. Aaronson. Limitations of quantum advice and one-way communication. *Theory of Computing*, 1:1–28, 2005. Earlier version in Complexity'04. quant-ph/0402095.

[2] S. Aaronson. The learnability of quantum states. *Proceedings of the Royal Society of London*, A463(2088), 2007. quant-ph/0608142.

[3] A. Ambainis. Communication complexity in a 3-computer model. *Algorithmica*, 16(3):298–301, 1996.

[4] L. Babai and P. G. Kimmel. Randomized simultaneous messages: Solution of a problem of Yao in communication complexity. In *Proceedings of the 12th IEEE Conference on Computational Complexity*, pages 239–246, 1997.

[5] Z. Bar-Yossef, T. S. Jayram, and I. Kerenidis. Exponential separation of quantum and classical one-way communication complexity. In *Proceedings of 36th ACM STOC*, pages 128–137, 2004.

[6] H. Buhrman, R. Cleve, J. Watrous, and R. de Wolf. Quantum fingerprinting. *Physical Review Letters*, 87(16), September 26, 2001. quant-ph/0102001.

[7] W. van Dam. Personal communication, May 2007.

[8] D. Gavinsky. Personal communication, December 2007.

[9] D. Gavinsky, J. Kempe, I. Kerenidis, R. Raz, and R. de Wolf. Exponential separations for one-way quantum communication complexity, with applications to cryptography. In *Proceedings of 39th ACM STOC*, pages 516–525, 2007. quant-ph/0611209.

[10] D. Gavinsky, J. Kempe, O. Regev, and R. de Wolf. Bounded-error quantum state identification and exponential separations in communication complexity. In *Proceedings of 38th ACM STOC*, pages 594–603, 2006. quant-ph/0511013.

[11] E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge University Press, 1997.

[12] I. Newman. Private vs. common random bits in communication complexity. *Information Processing Letters*, 39(2):67–71, 1991.

[13] I. Newman and M. Szegedy. Public vs. private coin flips in one round communication games. In *Proceedings of 28th ACM STOC*, pages 561–570, 1996.

[14] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.

[15] R. de Wolf. Quantum communication and complexity. *Theoretical Computer Science*, 287(1):337–353, 2002.

[16] A. C-C. Yao. Probabilistic computations: Toward a unified measure of complexity. In *Proceedings of 18th IEEE FOCS*, pages 222–227, 1977.