

CLASSICAL SIMULATION OF QUANTUM COMPUTATION, THE GOTTESMAN-KNILL THEOREM, AND SLIGHTLY BEYOND

MAARTEN VAN DEN NEST

*Max-Planck-Institut für Quantenoptik, Hans-Kopfermann-Str. 1
D-85748 Garching, Germany*

Received December 18, 2008

Revised November 9, 2009

We study classical simulation of quantum computation, taking the Gottesman-Knill theorem as a starting point. We show how each Clifford circuit can be reduced to an equivalent, manifestly simulatable circuit (normal form). This provides a simple proof of the Gottesman-Knill theorem without resorting to stabilizer techniques. The normal form highlights why Clifford circuits have such limited computational power in spite of their high entangling power. At the same time, the normal form shows how the classical simulation of Clifford circuits fits into the standard way of embedding classical computation into the quantum circuit model. This leads to simple extensions of Clifford circuits which are classically simulatable. These circuits can be efficiently simulated by classical sampling (“weak simulation”) even though the problem of exactly computing the outcomes of measurements for these circuits (“strong simulation”) is proved to be $\#P$ -complete—thus showing that there is a separation between weak and strong classical simulation of quantum computation.

Keywords: Quantum computation, classical simulation

Communicated by: B Kane & H Zbinden

1. Introduction

The Gottesman-Knill theorem [1] is a well-known result in quantum information theory which states that a certain class of non-trivial quantum circuits, called “Clifford circuits”, can be simulated efficiently on a classical computer, and can hence not provide any (exponential) speed-up w.r.t. classical computation. A Clifford circuit is any quantum circuit which is composed of Hadamard, PHASE and CNOT gates. The Gottesman-Knill theorem asserts that each (uniform family of) Clifford circuit(s), when acting on the computational basis state $|0\rangle \equiv |0\rangle^N$, and when followed by a computational basis measurement, can be simulated efficiently on a classical computer.

While in fact not so hard to prove, this result exhibits some rather remarkable and sometimes puzzling features, not all of which are fully understood. For example, even though they are efficiently classically simulatable, Clifford circuits can generate a high degree of entanglement [2]; the highly entangled cluster states [3] can e.g. be generated by Clifford circuits. This very feature raises doubts about the often-recited mantra that “entanglement is responsible for the quantum computational speed-up”. In particular, it highlights that, while the presence of (certain types of) entanglement in a quantum computation (QC) is provably necessary to disallow efficient classical simulation (see e.g. [4, 5, 6, 7, 8, 9, 10]), it is certainly not sufficient.

Further, it is known that Clifford circuits can efficiently be simulated classically by a rather restricted classical computer, namely a circuit model computer which only uses NOT and CNOT gates [11]. In other words, it is not necessary to invoke the full power of classical computers to efficiently simulate arbitrary Clifford circuits. Hence, Clifford circuits are most likely not even universal for *classical* computation. In complexity theoretic terms, the power of classical computation with NOT and CNOT gates—and hence of Clifford circuits—is captured by the complexity class $\oplus\mathbf{L}$ (“parity-L”) [12, 13]. This class is known to be contained in \mathbf{P} but not expected to be equal to it (although this is an unproven conjecture). Nevertheless, supplementing Clifford operations with essentially any non-Clifford gate immediately yields the full QC model [14, 15]. This yields an immediate “jump” in computational power from $\oplus\mathbf{L}$ to \mathbf{BQP} rather than a “smooth” transition $\oplus\mathbf{L} \rightarrow \mathbf{BPP} \rightarrow \mathbf{BQP}$. In particular, this property makes it hard to extend Clifford circuits to a class of efficiently simulatable quantum circuits which has the same computational power as full classical computation.

It is the aim of this note to obtain some insight in the above list of features. First we will show that each Clifford circuit \mathcal{C} can be (efficiently) reduced to an equivalent circuit \mathcal{C}' which yields the same output. The “normal form” \mathcal{C}' is also a Clifford circuit but has a very simple structure, as displayed in Fig. 2; this normal form is based on an earlier result [16], for which we also provide a simple proof. We will see that \mathcal{C}' is manifestly efficiently classically simulatable. We argue that the normal form also sheds some light on why the high degree of entanglement and interference generated by Clifford operations does not result in any (exponential) quantum computational speed-up. Finally, the normal form shows how the simulation of Clifford operations fits within the standard embedding of probabilistic classical computation into the quantum circuit model. The latter is related to the notion of HT circuits, as discussed in section 4.

Along the way, we will make some general remarks regarding classical simulation of QC. One of them regards the different possible definitions of the notion of “classical simulation”. When a QC is to be simulated classically, the aim may be to either (i) *compute* the probabilities of the output measurement efficiently classically with high accuracy (“strong simulation”) or (ii) *sample* from this distribution efficiently using a classical computer (“weak simulation”). Both variants constitute valid classical simulation techniques; however, a priori it is not clear whether there is a clear separation between these two notions. The vast majority of all works (see e.g. [1, 6, 7, 9, 10, 18, 19, 20, 21, 22, 23, 24, 25]) regarding classical simulation of QC considers quantum circuits which can be simulated efficiently classically in the strong sense, whereas not much is known about weak simulation.

We will find that the simulatable quantum circuits considered in the present paper are *not* amenable to strong simulation—in fact, it is easy to show that strong simulation of these circuits constitutes a $\#\mathbf{P}$ -complete problem. Nevertheless, weak classical simulation of the same circuits is efficiently possible by means of a very simple sampling technique. This shows that strong and weak simulation of QC are fundamentally different notions. Moreover, these examples highlight that any serious attempt to understand the relationship between classical and quantum computation should not rely on the notion of strong simulation.

2. Classical simulation of quantum computation

Here we state more precisely what we mean by “efficient classical simulation of quantum

computation”. We essentially follow the definitions from [4, 21].

Consider a uniform family of quantum circuits* $\mathcal{U} \equiv \mathcal{U}_N$ acting on the N -qubit input state $|\mathbf{0}\rangle \equiv |0\rangle^{\otimes N}$, and followed by a measurement of, say, the first qubit in the computational basis. Then this quantum computation yields as an outcome a bit $\alpha \in \{0, 1\}$. The probability that the outcome α occurs is given by $\pi(\alpha) = \langle \mathbf{0} | \mathcal{U}^\dagger [|\alpha\rangle\langle\alpha| \otimes \mathbf{I}] \mathcal{U} | \mathbf{0}\rangle$. We say that the above quantum computation can be efficiently simulated classically in the *strong* sense if it is possible to evaluate $\pi(0)$ up to M digits in $\text{poly}(N, M)$ time on a classical computer. Furthermore, we say that the quantum computation can be efficiently simulated classically in the *weak* sense if it is possible to sample once from the probability distribution $\{\pi(\alpha)\}$ in $\text{poly}(N)$ time on a classical computer.†

We further point out that to date the vast majority of results on simulation of quantum computation regard *strong* classical simulation. This regards e.g. the Gottesman-Knill theorem, the simulation of “matchgates” [18, 19, 20, 21], the simulation of the quantum Fourier transform [22, 23, 24], simulation results involving tensor contracting techniques [6, 7, 9, 10], etc. However, in order to study the difference between classical and quantum computation, the notion of weak classical simulation is much more appropriate. In fact, the strong simulation of general quantum circuits can easily be shown to be a $\#\mathbf{P}$ -hard problem (this can e.g. been showed by considering the the family of circuits U_f as in section 4). As \mathbf{BQP} is believed to be much smaller than $\#\mathbf{P}$, this indicates that strong classical simulation inevitably seems too strong a requirement to study the relation between classical and quantum computation. However, not much is known about the difference between strong and weak simulation of quantum computation. In section 4 we will give an example of a class of quantum circuits which is efficiently simulatable in the weak sense but for which simulation in the strong sense is intractable (unless \mathbf{P} is equal to $\#\mathbf{P}$).

3. The Gottesman-Knill theorem

In this section we recall the Gottesman-Knill theorem and the relation between Clifford operations and the complexity class $\oplus\mathbf{L}$. A Clifford circuit is any quantum circuit composed of Hadamard (H), controlled-NOT (CNOT) and PHASE (P) gates, where H and CNOT are defined as usual and $P = \text{diag}(1, i)$. (A variant of) the Gottesman-Knill theorem can be formulated as follows [1]:

Theorem 1 (Gottesman-Knill) *Every (uniform family of) Clifford circuit(s), when applied to the input state $|\mathbf{0}\rangle \equiv |0\rangle^{\otimes N}$ and when followed by a Z measurement of the first qubit, can be efficiently simulated classically in the strong sense.*

The standard technique to prove this result involves a connection between Clifford operations and groups of commuting Pauli operations called “stabilizer groups”. We omit this proof here and refer to e.g. [1, 26]. Below we will provide an alternative proof of the Gottesman-Knill theorem which does not use stabilizer techniques.

*In the following, when we refer to a “quantum circuit” we will always mean a uniform family of quantum circuits.

†More precisely, one should require to sample from a probability distribution which is not necessarily exactly equal to $\{\pi(\alpha)\}$, but sufficiently close to it (see e.g. [4]). In the examples in this paper, however, it will always be possible to perform an exact sampling, such that we omit such accuracy issues for simplicity.

It turns out that a rather restricted classical computer suffices to efficiently simulate arbitrary Clifford circuits. Indeed, every Clifford circuit can be efficiently simulated by a classical circuit model computer which uses NOT and CNOT gates only, applied to the input string $\mathbf{0} = (0, \dots, 0)$ [11]. More precisely, in [11] the problem GOTTESMAN-KNILL was defined as follows: the input is an n -qubit Clifford circuit \mathcal{C} ; the problem is to decide whether the first qubit will be in the state $|1\rangle$ with certainty after \mathcal{C} has been applied to the input $|0\rangle^n$. It was showed in [11] that this problem can be mapped, under a logarithmic-space reduction, to a problem of simulating a classical poly-size CNOT-NOT circuit acting on the all-zeroes input state. The complexity class of problems that are log-space reducible to the simulation of CNOT-NOT circuits, is called $\oplus\mathbf{L}$ (“Parity- \mathbf{L} ”) ‡[12, 13]. Hence, it was showed in [11] that GOTTESMAN-KNILL is in $\oplus\mathbf{L}$. As CNOT and NOT gates are Clifford operations, it is clear that every $\oplus\mathbf{L}$ problem can be reduced to an instance of GOTTESMAN-KNILL. In conclusion, the problem GOTTESMAN-KNILL is $\oplus\mathbf{L}$ -complete [11].

The class of problems in $\oplus\mathbf{L}$ are centered around linear algebra over the finite field \mathbf{Z}_2 . It has been shown that the problems of solving linear equations over \mathbf{Z}_2 , finding the inverse of a nonsingular \mathbf{Z}_2 -matrix, multiplying matrices over \mathbf{Z}_2 etc., are $\oplus\mathbf{L}$ problems [13].

4. Embedding classical in quantum computation

Here we briefly review how classical (probabilistic) computation can be regarded as being a “part” of quantum computation. All material in this section is standard (see e.g. [26]), except for an observation about the difference between strong and weak simulation of quantum computation, which will be made at the end of this section.

Probabilistic classical computation is classical computation (e.g. considered in the circuit model) supplemented with the possibility of deciding, in each step of the computation, which gate to apply based on the random outcome of a coin toss. It is well understood how probabilistic classical computation can be embedded in the full pure-state circuit model of quantum computation. The subclass of quantum circuits that corresponds to probabilistic classical computation is the following: (ROUND 1) apply Hadamard gates to an arbitrary subset of qubits; (ROUND 2) apply a (uniform family of) circuit(s) consisting of classical[§] gates only (e.g. NOT, CNOT, Toffoli gates). A quantum circuit with the above structure will be called an HT circuit (short for Hadamard-Toffoli); see also Fig 1. It can easily be shown that every probabilistic classical computation can be accounted for by a suitable HT circuit acting on the all-zeroes input state and followed by a single-qubit Z measurement (say, on the first qubit). Conversely, it is also straightforward to see that each HT circuit (with input and measurement as above) can be simulated efficiently—in the weak sense—by a probabilistic classical computer: if the quantum register has N qubits, if m denotes the number of Hadamard gates applied in ROUND 1, and if the classical gates in ROUND 2 compute an efficiently computable (invertible) function $f : \{0, 1\}^N \rightarrow \{0, 1\}^N$, then the state of the system after the first two rounds has the form $\sum_{x \in \{0, 1\}^m} |f(x, 0, \dots, 0)\rangle$, where there are $N - m$ zeroes in the argument of f . We will use the abbreviation $F(x) \equiv f(x, 0, \dots, 0)$.

‡Obviously, $\oplus\mathbf{L} \subseteq \mathbf{P}$; the inclusion is believed to be strict, but a proof of this conjecture has not yet been found.

§We call any unitary gate U a “classical gate” if it maps computational basis states to computational basis states.




Fig. 1. The class of quantum circuits with the above structure are here called HT circuits (first a round of HADAMARD gates is applied to a subset of the qubits, followed by a round of “classical gates”, i.e. TOFFOLI, NOT, etc.). HT circuits followed by a computational basis measurement have the same computational power as probabilistic classical computation: each probabilistic classical computation can be simulated efficiently by an HT circuit, and each HT circuit can be weakly simulated efficiently on a classical computer. However, strong classical simulation of HT circuits is a $\#\mathbf{P}$ -complete problem.

Performing a $\{|0\rangle, |1\rangle\}$ measurement on the first qubit yields as an outcome a bit $\alpha \in \{0, 1\}$. The probability that α occurs is equal to

$$\pi(\alpha) = \frac{|\{x \in \{0, 1\}^m : F(x)_1 = \alpha\}|}{2^m}, \quad (1)$$

where $F(x)_1$ denotes the first bit of $F(x)$. Sampling from this probability distribution can easily be done classically: simply generate an m -bit string x uniformly at random, then compute $F(x)$ and finally set $\alpha := F(x)_1$; then α is generated with probability $\pi(\alpha)$ as desired. This shows that weak classical simulation of HT quantum circuits is efficiently possible. We can therefore conclude that the computational power of HT quantum circuits is equivalent to probabilistic classical computation.

We remark that *strong* classical simulation of HT circuits is much harder—in fact, we point out that such strong simulation is a $\#\mathbf{P}$ -complete problem. The complexity class $\#\mathbf{P}$ is concerned with *counting problems*: given an efficiently computable Boolean function $f \in \mathbf{P}$, the problem of determining the integer $\#f := |\{x : f(x) = 0\}|$ defines the complexity class $\#\mathbf{P}$ [17]. It is easy to see that the problem of computing the probabilities (1) is at least as hard as computing $\#f$ for arbitrary functions $f \in \mathbf{P}$. To show this, consider an arbitrary Boolean function $f : \{0, 1\}^m \rightarrow \{0, 1\}$ which is in \mathbf{P} . Then there exists a poly-size HT quantum circuit \mathcal{U} such that $\mathcal{U}|0\rangle \propto \sum_x |x\rangle |f(x)\rangle$. Performing a measurement in the computational basis on the last qubit yields a single bit with probability distribution $\pi(0) = \#f/2^m = 1 - \pi(1)$. Computing these probabilities with perfect accuracy is thus equivalent to computing $\#f$. This shows that strong simulation of HT circuits is at least as hard as any problem in $\#\mathbf{P}$ and is thus $\#\mathbf{P}$ -hard. Moreover, it can easily be shown that, for arbitrary HT circuits and for arbitrary computational basis measurements, the problem of computing the probabilities (1) is always in $\#\mathbf{P}$. This shows that the strong simulation of HT circuits is a $\#\mathbf{P}$ -complete problem. In conclusion, *efficient strong classical simulation of HT quantum circuits is intractable (unless \mathbf{P} is equal to $\#\mathbf{P}$), while weak simulation of these circuits is efficiently possible* (even trivially so).

5. Clifford circuits as HT circuits

In this section we show that, modulo some “redundancies” (to be explained below), each Clifford circuit can be reduced to an HT circuit, which is manifestly classically simulatable. To this aim, we provide an alternative proof of the Gottesman-Knill theorem which will allow to make the connection to HT circuits. The proof will not be centered around stabilizer groups (and hence does not require any knowledge of the stabilizer formalism) but will simply consist of tracking the coefficients ψ_x of the state $|\psi\rangle = \sum_x \psi_x |x\rangle$ of the quantum register throughout the application of an arbitrary Clifford circuit. To achieve this, it will be useful to know what the most general form is which the coefficients ψ_x may assume. The latter issue has in fact already been fully investigated and understood in [16], where the following result was proved: every N -qubit state $|\psi\rangle = \mathcal{C}|0\rangle^N$, where \mathcal{C} represents a Clifford circuit, is given by an expression of the form

$$|\psi\rangle \propto \sum_{x \in A} i^{l(x)} (-1)^{q(x)} |x\rangle, \tag{2}$$

where we have the following notations. First, $A \subseteq \mathbf{Z}_2^N$ is an affine subspace, i.e. a subset of the form $A = \{Ru + t | u \in \mathbf{Z}_2^m\}$, for some fixed (nonsingular) $N \times m$ matrix R and vector $t \in \mathbf{Z}_2^N$. Second, $l(x)$ is a linear function on \mathbf{Z}_2^N , i.e., l maps x to $l(x) = d^T x$ for some $d \in \mathbf{Z}_2^N$. The exponent $l(x)$ of i is computed modulo two. Finally, $q(x)$ is a quadratic function on \mathbf{Z}_2^N . That is, q maps N -bit strings $x = (x_1, \dots, x_N)$ to $q(x) = \sum c_{ij} x_i x_j + c_i x_i$, for some (fixed) $c_{ij}, c_i \in \{0, 1\}$. Moreover, q , l and A can be efficiently computed. Conversely, every state of the form (2) can be prepared by applying a suitable (poly-size) Clifford circuit to the input $|0\rangle$ (such states are called *stabilizer states*).

The proof that an arbitrary Clifford circuit always maps $|0\rangle$ to a state of the form (2) was first derived in [16] and is based on mappings between Clifford operations, the stabilizer formalism and arithmetic over \mathbf{Z}_2 . Here we will give a simple alternative proof of this theorem.

Proof of Eq. (2): The proof can be obtained by induction on the number of gates in the circuit \mathcal{C} . Evidently, (2) is true when \mathcal{C} is the identity. Now suppose that the result is true for every Clifford circuit consisting of K gates. We then have to prove that the result still holds for circuits with one additional gate. Consider an arbitrary circuit \mathcal{C} of size $K + 1$, and write \mathcal{C} as a product of a single Clifford gate $U \in \{H, P, \text{CNOT}\}$ and a circuit \mathcal{C}' of size K : $\mathcal{C} = U\mathcal{C}'$. By the induction step, we may assume that there exist A , l and q such that $\mathcal{C}'|0\rangle$ is given by (2). It is then simply a matter of verifying that the general form (2) is kept when a gate from the set $\{H, P, \text{CNOT}\}$ is applied. This can be shown using straightforward arithmetic, and for completeness these calculations are given in the appendix. \square

The above proof of expression (2) leads to a simple algorithm to compute A , l and q if a circuit \mathcal{C} is given: first, one initializes A , l and q to their trivial values, corresponding to the state $|0\rangle$. Then one sequentially updates A , l and q corresponding to the first, second, etc. gate in \mathcal{C} ; it can easily be seen that each of these updates is efficient. The final values for A , l and q then correspond to the state $\mathcal{C}|0\rangle$.

It is now easy to see that each Clifford computation can be simulated efficiently. If a circuit \mathcal{C} is given, the first step is to compute the triple (A, l, q) parameterizing $\mathcal{C}|0\rangle$. The goal is then to efficiently simulate $\{|0\rangle, |1\rangle\}$ measurements on this state. To do so, one argues as follows. First, the outcome probabilities of a measurement in the computational basis are *independent* of the phases $i^{l(x)}(-1)^{q(x)}$. That is, the specific values of l and q are in this

context completely redundant, such that we may set these quantities to their trivial values $l \equiv 0, q \equiv 0$. This leaves us with the state $\sum_{x \in A} |x\rangle$. Let the affine space A be given by $A = \{Ru + t : u \in \mathbf{Z}_2^m\}$, for some invertible $N \times m$ matrix R and $t \in \mathbf{Z}_2^N$. Further, suppose that the first qubit is measured, yielding as an outcome a bit α which occurs with probability

$$\pi(\alpha) = \frac{|\{x \in A : x_1 = \alpha\}|}{|A|}. \quad (3)$$

It is trivial to sample from this probability distribution: simply generate a uniformly random m -bit string u , compute $Ru + t$ and set α to be the first coefficient of $Ru + t$. This shows that each Clifford circuit can efficiently be simulated in the weak sense.

However, one can do more than this, as the probabilities $\pi(\alpha)$ can be *computed* efficiently as well. Indeed, each of these probabilities has the form $2^{-\gamma}$ for some integer γ which can be obtained by solving a system of linear equations over \mathbf{Z}_2 . Hence, strong simulation is possible as well, thus recovering the Gottesman-Knill theorem.

6. Normal form

Next we note that the expression (2) provides an alternative way—different than the circuit \mathcal{C} —to prepare the output state $\mathcal{C}|\mathbf{0}\rangle$. If the affine space A is given by $A = \{Ru + t : u \in \mathbf{Z}_2^m\}$, then this state can be prepared from $|\mathbf{0}\rangle$ as follows. First, apply m Hadamard gates such as to yield $\sum_u |u\rangle|0\rangle^{N-m}$. Second, apply appropriate CNOT gates such as to yield $\sum_u |Ru\rangle$. Third, apply the operation $X^{t_1} \otimes \dots \otimes X^{t_N}$, yielding $\sum_u |Ru + t\rangle$. Fourth, apply appropriate PHASE and CPHASE = $\text{diag}(1, 1, 1, -1)$ gates such as to finally yield $\mathcal{C}|\mathbf{0}\rangle$. The overall structure of this preparation method is illustrated in Fig. 2. We have now arrived the following result.

Theorem 2 *Let \mathcal{C} be an arbitrary poly-size Clifford circuit. Then there exists a poly-size Clifford circuit \mathcal{C}' satisfying $\mathcal{C}|\mathbf{0}\rangle = \mathcal{C}'|\mathbf{0}\rangle$ such that \mathcal{C}' can be decomposed into three “rounds”: (ROUND 1) apply Hadamard gates to an arbitrary subset of qubits; (ROUND 2) apply a poly-size circuit of NOTs and CNOTs; (ROUND 3) apply a poly-size circuit of PHASEs and CPHASEs. The circuit \mathcal{C}' can be efficiently determined.*

It is important to remark that it is generally not true that the circuits \mathcal{C} and \mathcal{C}' are equal as $2^N \times 2^N$ matrices.[¶] These circuits merely have the same effect on the input state $|\mathbf{0}\rangle$, i.e., they generate the same output state. Note that it is the output state $|\psi_{out}\rangle = \mathcal{C}|\mathbf{0}\rangle$ of the computation which is relevant for our purposes—i.e. a single column of the $2^N \times 2^N$ matrix \mathcal{C} —and not the entire circuit. We further emphasize that the choice of a standard input state $|\mathbf{0}\rangle$ as opposed to arbitrary products of X, Y, Z eigenstates as inputs, does not entail any loss of generality, as such alternative input states can always be rotated into the state $|\mathbf{0}\rangle$ using a (local) Clifford operation.

The Clifford normal form highlights the computational “weakness” of Clifford circuits. First, the original Clifford circuit, which may contain many Hadamard gates located at different places in the circuit, causing subsequent rounds of constructive and destructive interference, is (efficiently) mapped to the normal form which does not display any interference

[¶]A counterexample is e.g. given by the single-qubit circuit $\mathcal{C} = \text{HPH}$. By exhaustive enumeration of all possibilities, it can easily be shown that no single-qubit circuit \mathcal{C}' of the above structure satisfies $\mathcal{C} = \mathcal{C}'$.

at all. Indeed, up to the irrelevant last round of diagonal gates, each Clifford circuit is reduced to a circuit of NOT and CNOT gates applied to a superposition of computational basis states $|+\rangle^K|0\rangle^L$ —in other words, nothing but a (very simple instance of an) HT circuit, which are trivially classically simulatable. Also the connection between Clifford circuits and the complexity class $\oplus\mathbf{L}$ is in this way highlighted, given the relation between the class $\oplus\mathbf{L}$ and circuits consisting of NOT and CNOT operations.

With a little extra work, theorem 2 may be used to arrive at a normal form for Clifford circuits \mathcal{C} which does regard the entire $2^N \times 2^N$ matrix. The normal form is in fact highly similar to theorem 2. The significance of this normal form is again that it contains only a single round of basis-changing operations (i.e. Hadamards). In fact, we will see that every Clifford operation can be written as a tensor product of Hadamard operations $\mathcal{H} = H^S \otimes I$ acting nontrivially on a subset S of the qubits, multiplied on the left and the right with *basis-preserving* Clifford circuits M_1 and M_2 , i.e. circuits composed of CNOT, NOT, PHASE and CPHASE gates. Similar to (2), (a variant of) theorem 3 was proved in [16] using mappings between the stabilizer formalism, Clifford operations and \mathbf{Z}_2 -arithmetic; here we provide a direct proof.

Theorem 3 *Let \mathcal{C} be an arbitrary n -qubit Clifford operation. Then there exist: (a) poly-size circuits M_1 and M_2 composed of CNOT, NOT, PHASE and CPHASE gates and (b) a tensor product of Hadamard gates and identities $\mathcal{H} = H^S \otimes I$ acting nontrivially on a subset S of the qubits, such that $\mathcal{C} \propto M_2\mathcal{H}M_1$. Moreover, M_1 , M_2 and \mathcal{H} can be determined efficiently.*

Proof: Let $a = (a_1, \dots, a_n)$ denote an arbitrary n -bit string. For every $i = 1, \dots, n$, define $\sigma_i := \mathcal{C}X_i\mathcal{C}^\dagger$ where X_i denotes the Pauli matrix X acting on qubit i . Since \mathcal{C} is a Clifford operation, each σ_i is a Pauli operator, possibly with an overall minus sign. Denoting $X(a) = X^{a_1} \otimes \dots \otimes X^{a_n}$ and $\sigma(a) := \prod_i \sigma_i^{a_i}$, we thus have $\sigma(a) := \mathcal{C}X(a)\mathcal{C}^\dagger$. Since $|a\rangle = X(a)|0\rangle^n$, we have $\mathcal{C}|a\rangle = \mathcal{C}X(a)|0\rangle^n = \sigma(a)\mathcal{C}|0\rangle^n$. We can now apply theorem 2 to the state $\mathcal{C}|0\rangle^n$: there exists a poly-size Clifford circuit M consisting of CNOT, NOT, PHASE and CPHASE gates, and an operation $\mathcal{H} = H^S \otimes I$ such that $\mathcal{C}|0\rangle^n = \gamma M\mathcal{H}|0\rangle^n$, where γ is some overall phase factor. Now define $\tau_i := \mathcal{H}M^\dagger\sigma_i M\mathcal{H}$ for every i ; also τ_i is a Pauli operator, possibly supplemented with an overall minus sign. Denoting $\tau(a) := \prod_i \tau_i^{a_i}$, we thus have $\tau(a) := \mathcal{H}M^\dagger\sigma(a)M\mathcal{H}$. This leads to

$$\mathcal{C}|a\rangle = \gamma\sigma(a)M\mathcal{H}|0\rangle^n = \gamma M\mathcal{H}\tau(a)|0\rangle^n. \tag{4}$$

Note that γ does not depend on a .

We now study the general form of the unitary mapping $|a\rangle \rightarrow \tau(a)|0\rangle^n$. Since each τ_i is an element of the Pauli group, there exist n -bit strings $u = (u_1, \dots, u_n), v = (v_1, \dots, v_n), R^i = (R_1^i, \dots, R_n^i)$ and $T^i = (T_1^i, \dots, T_n^i)$ such that $\tau_i = (-1)^{u_i v_i} X(R^i)Z(T^i)$ and hence

$$\tau(a) = \prod_i (-1)^{a_i u_i} X(a_i R^i)Z(a_i T^i). \tag{5}$$

||A Pauli operator has the form $P = P_1 \otimes \dots \otimes P_n$ where each P_i is either the identity or one of the Pauli matrices X, Y or Z .




Fig. 2. For each Clifford circuit \mathcal{C} acting on the computational basis state $|0\rangle$ (left), there exists a normal form \mathcal{C}' (right) of the above structure such that $\mathcal{C}|0\rangle = \mathcal{C}'|0\rangle$, i.e., these two circuits yield the same output state. The last round, consisting of diagonal gates, is undetected by a computational basis measurement and can therefore be completely disregarded. Moreover, the first two rounds in \mathcal{C}' constitute a (restricted) HT circuit and hence are trivially classically simulatable.

For every two $x, z \in \mathbf{Z}_2^n$, we have $X(x)Z(z) = (-1)^{x^T z} Z(z)X(x)$. Hence, we may reshuffle the factors in the product (5) at the cost of an overall phase factor of the form $(-1)^{q(a)}$, for some quadratic function $q : \mathbf{Z}_2^n \rightarrow \mathbf{Z}_2$. In particular, we may write

$$\tau(a) = (-1)^{q(a)} \left[\prod_i i^{a_i v_i} \right] \left[\prod_i X(a_i R^i) \right] \left[\prod_i Z(a_i T^i) \right]. \tag{6}$$

Letting R (T) denote the $n \times n$ matrix with columns R^i (T^i), we then have

$$\tau(a) = (-1)^{q(a)} \left[\prod_i i^{a_i v_i} \right] X(Ra)Z(Ta). \tag{7}$$

The action of $\tau(a)$ on $|0\rangle^n$ then has the following form:

$$\tau(a)|0\rangle^n = (-1)^{q(a)} \left[\prod_i i^{a_i v_i} \right] X(Ra)Z(Ta)|0\rangle^n = (-1)^{q(a)} \left[\prod_i i^{a_i v_i} \right] |Ra\rangle, \tag{8}$$

where we have used that $Z(Ta)|0\rangle^n = |0\rangle^n$ and that $X(Ra)|0\rangle^n = |Ra\rangle$. It then easily follows that there exists a Clifford circuit M' (independent of a) composed of CNOT, NOT, PHASE and CPHASE gates such that $\tau(a)|0\rangle^n = M'|a\rangle$. Together with (4), this completes the proof. \square

7. Slightly beyond Gottesman-Knill

The Clifford normal form in theorem 2 makes it easy to extend the Gottesman-Knill theorem. Next we provide a class of quantum circuits which are efficiently simulatable classically (in the weak sense), and which are extensions of Gottesman-Knill in the sense that these circuits can efficiently generate any output state which can be efficiently generated by arbitrary Clifford circuits. Moreover, contrary to Clifford circuits, the extensions encompass full probabilistic classical computation.

We consider circuits of the following structure: first, apply an arbitrary local unitary operation; afterwards, apply an arbitrary (uniform family) of quantum circuit(s) consisting of Toffoli and diagonal gates only. We show that such circuits, when applied to the input $|0\rangle$ and when followed by a computational basis measurement, can be efficiently simulated

classically in the weak sense. After the first round, the state of the quantum register is a complete product state $|\chi\rangle = |\chi_1\rangle \otimes \dots \otimes |\chi_N\rangle$. Denoting $|\chi_i\rangle = a_i|0\rangle + b_i|1\rangle$, the coefficients χ_x of $|\chi\rangle$ in the computational basis (where $x = (x_1, \dots, x_N)$ is an N -bit string) are given by $\chi_x = \prod_i a_i^{1-x_i} b_i^{x_i}$. The final state of the quantum register has the form $\sum_x \chi_x \theta_x |f(x)\rangle$, where the θ_x are complex phases (which are efficiently computable as a function of x , but this will not be relevant here) and where $f : \{0, 1\}^N \rightarrow \{0, 1\}^N$ is an invertible, efficiently computable Boolean function. Further, suppose that a subset S of the N qubits is measured, yielding as an outcome a bit string $\alpha = (\alpha_j : j \in S)$, where the bit α_j is the measurement outcome for qubit j . Then α occurs with probability $\pi(\alpha) = \sum |\chi_x|^2$, where the sum is over all bit strings x such that $f(x)_j = \alpha_j$ for all $j \in S$. To state it differently, $\pi(\alpha)$ is the total probability that an N -bit string x which is generated with probability $p(x) = |\chi_x|^2$, satisfies $f(x)_j = \alpha_j$ for all $j \in S$. Now, it is possible to efficiently sample from the probability distribution $\{p(x)\}$. As a consequence, the following procedure (i)-(ii)-(iii) allows to efficiently generate a bit string α with probability $\pi(\alpha)$ on a classical computer: (i) generate an N -bit string x with probability $p(x) = |\chi_x|^2$ (note that $\{p(x)\}$ is a simple product distribution since $|\chi\rangle$ is a product state); (ii) compute $f(x)$; (iii) set α_j to be the j -th coefficient of $f(x)$, for every $j \in S$. This shows it is possible to efficiently sample from the probability distribution $\{\pi(\alpha)\}$.

Note that the first round, consisting of a local unitary operation, can be generalized while keeping the entire circuit classically simulatable. To do so, consider a quantum circuit \mathcal{U} such that the computation $|\mathbf{0}\rangle \rightarrow \mathcal{U}|\mathbf{0}\rangle$, followed by a computational basis measurement of *all* N qubits, can be efficiently simulated classically in the weak sense. That is, it is classically possible to efficiently sample once from the probability distribution $p(x) = |\langle x|\mathcal{U}|0\rangle|^2$. Then such a circuit, followed by an arbitrary (uniform) quantum circuit consisting of classical and diagonal gates only, and followed by a computational basis measurement of an arbitrary subset of the qubits, can be simulated efficiently classically in the weak sense. This last observation e.g. implies that a ‘‘matchgate circuit’’ [18, 19, 20, 21], a circuit of ‘‘bounded tree-width’’ [6, 7, 9], a circuit which generates ‘‘bounded Schmidt-rank’’ [5, 10], or a quantum Fourier transform [22, 23, 24] (which are all known to be strongly simulatable), followed by an arbitrary poly-size Toffoli-Diagonal circuit, is still weakly simulatable.

8. Conclusion

We have studied classical simulation of quantum computation, taking the Gottesman-Knill theorem as a starting point. We have showed that each Clifford circuit \mathcal{C} can be reduced to a simple equivalent circuit \mathcal{C}' which, when applied to the input $|\mathbf{0}\rangle$, provides the same output state. Using this reduction to the normal form, we have attempted to provide a better understanding in the somewhat peculiar features of the Gottesman-Knill theorem. We have argued that the normal form provides insight in why the large amount of entanglement which can be generated by Clifford circuits fails to provide any quantum computational speed-up. Furthermore, the normal form shows how the simulation of Clifford operations fits within the standard embedding of probabilistic classical computation into the quantum circuit model, related to the notion of HT circuits.

The class of HT circuits considered here, while indeed simple, exhibits features which are quite different from many results regarding classical simulation of quantum computation which have been found so far (such as the Gottesman-Knill theorem itself, matchgates, circuits

of small tree-width, etc). For example, while such circuits are classically simulatable in the weak sense (i.e. they can be simulated using classical sampling techniques), strong classical simulation (i.e. the problem of *computing* the output probabilities of measurement outcomes with high accuracy) is $\#\mathbf{P}$ -complete and hence intractable. This is in contrast with previous results, the vast majority of which consider circuits where strong simulation is efficiently possible. Also, HT circuits can generate unbounded amounts of entanglement (as can e.g. Clifford circuits) while several previous results use techniques which allow classical simulation of certain quantum circuits only in cases where the entanglement generated by these circuits is (in a well-defined sense) “bounded” (see e.g. [5, 6, 10]). Thus, the present results provide simple examples of quantum circuits which can generate vast amounts of entanglement, which nevertheless does not result in any quantum computational speed-up.

We further note that a very moderate extension of HT circuits immediately leads to highly non-trivial quantum algorithms. For example, it is known that Shor’s factoring algorithm, using the phase estimation approach [27] (see also [28]), can be implemented by an HT circuit *supplemented with a final round of Hadamard gates*, and then followed by a measurement in the computational basis. This shows that a mere final round of local basis changes allows to go from an HT circuit, which is trivially classically simulatable, to Shor’s factoring algorithm, which is believed to provide an exponential speed-up.

Acknowledgements

I am very grateful to R. Jozsa, W. Dür, H. Briegel, A. Kay and I. Cirac for interesting discussions and suggestions on the manuscript. Work supported by the Excellence Cluster MAP.

References

1. D. Gottesman (1998), *The Heisenberg Representation of Quantum Computers*, talk at International Conference on Group Theoretic Methods in Physics, arXiv:quant-ph/9807006.
2. M. Hein, W. Dür, J. Eisert, R. Raussendorf, M. Van den Nest and H. J. Briegel (2005), *Entanglement in Graph States and its Applications*, Proceedings of the International School of Physics “Enrico Fermi” on “Quantum Computers, Algorithms and Chaos” ; see also arXiv:quant-ph/0602096.
3. H. J. Briegel and R. Raussendorf (2001), *Persistent Entanglement in Arrays of Interacting Particles*, Phys. Rev. Lett. **86**, 910–913.
4. R. Jozsa and N. Linden (2002), *On the role of entanglement in quantum computational speed-up*, arXiv:quant-ph/0201143.
5. G. Vidal (2003), *Efficient Classical Simulation of Slightly Entangled Quantum Computations*, Phys. Rev. Lett. **91**, 147902.
6. I. Markov and Y. Shi (2008), *Simulating quantum computation by contracting tensor networks*, SIAM J. Comp., **38**(3):963-981, 2008.
7. Y. Shi, L.-M. Duan and G. Vidal (2006), *Classical simulation of quantum many-body systems with a tree tensor network*, Phys. Rev. A **74**, 022320.
8. N. Yoran and A. Short (2006), *Classical simulation of limited-width cluster-state quantum computation*, quant-ph/0601178.
9. R. Jozsa (2006), *On the simulation of quantum circuits*, arXiv:quant-ph/0603163.
10. M. Van den Nest, W. Dür, G. Vidal and H. J. Briegel (2007), *Classical simulation versus universality in measurement-based quantum computation*, Phys. Rev. A **75**, 012337 (2007).

11. S. Aaronson and D. Gottesman (2004), *Improved Simulation of Stabilizer Circuits*, Phys. Rev. A **70**, 052328.
12. C. H. Papadimitriou and S. K. Zachos (1982), *Two Remarks on the Complexity of Counting*, preprint MIT/LCS/TM-228, MIT.
13. C. Damm (1990), *Problems complete for $\oplus L$* , Inform. Process. Lett. **36**, 247.
14. G. Nebe, E. M. Rains and N. J. A. Sloane (2001), *The Invariants of the Clifford Groups*, Designs, Codes and Cryptography **24** (1), 99-122.
15. R. Solovay (2000), talk at Mathematical Sciences Research Institute.
16. J. Dehaene and B. De Moor (2003), *Clifford group, stabilizer states, and linear and quadratic operations over $GF(2)$* , Phys. Rev. A **68**, 042318.
17. C. M. Papadimitriou (1994), *Computational complexity* (Addison-Wesley, Reading, MA).
18. L. G. Valiant (2002), *Quantum circuits that can be simulated classically in polynomial time*, SIAM J. Comput. **31**, No. 4, p. 1229.
19. D. DiVincenzo and B. Terhal (2002), *Classical simulation of noninteracting-fermion quantum circuits*, Phys. Rev. A **65**, 032325/1-10.
20. S. Bravyi (2009), *Contraction of matchgate tensor networks on non-planar graphs*, Cont. Math., **482**, pp. 179-211.
21. R. Jozsa and A. Miyake (2008), *Matchgates and classical simulation of quantum circuits*, Proc. R. Soc. A **464**, 3089-3106 (2008).
22. D. Aharonov, Z. Landau, J. Makowsky (2006), *The quantum fast Fourier transform can be classically simulated*, quant-ph/0611156.
23. N. Yoran and A. Short (2007), *Efficient classical simulation of the approximate quantum Fourier transform* Phys. Rev. A **76**, 042321.
24. D. E. Browne (2007), *Efficient classical simulation of the semi-classical Quantum Fourier Transform*, New J. Phys. **9** 146.
25. S. Bravyi and R. Raussendorf (2007), *On measurement-based quantum computation with the toric code states*, Phys. Rev. A **76**, 022304.
26. M. A. Nielsen and I. L. Chuang (2000), *Quantum Computation and Quantum Information* (Cambridge University Press).
27. A. Y. Kitaev (1995), *Quantum measurements and the Abelian Stabilizer Problem*, quant-ph/9511026.
28. D. Shepherd (2005), *On the Role of Hadamard Gates in Quantum Circuits*, arXiv:quant-ph/0508153.

Appendix A

Here we show the following claim: **Claim:** suppose that an N -qubit state $|\psi\rangle$ is given in terms of an expression of the form (2) for some (A, l, q) , and suppose that a gate $U \in \{H, P, \text{CNOT}\}$ is applied to $|\psi\rangle$, resulting in a state $|\psi'\rangle$. Then $|\psi'\rangle$ still has the form (2) for some updated (A, l, q) , and that these updates can be performed efficiently.

Before proceeding with the proof, we note the following: in (2), the functions q and l are defined on the space of N -bit strings \mathbf{Z}_2^N . If the affine space A is given by $A = \{Ru + t : u \in \mathbf{Z}_2^m\}$, we might as well write

$$|\psi\rangle \propto \sum_u (-1)^{\bar{q}(u)} i^{\bar{l}(u)} |Ru + t\rangle, \quad (\text{A.1})$$

where now \bar{q} and \bar{l} are quadratic, resp. linear functions defined on the space of m -bit strings u , such that $\bar{q}(u) = q(Ru + t)$ and $\bar{l}(u) = l(Ru + t)$. Note that it is computationally easy to determine \bar{q} (\bar{l}) from q (l) and vice versa. In the following we will be a bit sloppy and simply

identify $q(u) \equiv q(Ru + t)$ and similar for l ; it will be clear from the context what the notation means.

We will also need the following lemma:

Lemma A.1 *Let $q(x)$ and $l(x)$ be quadratic and linear functions, resp., on \mathbf{Z}_2^K , where $x = (x_1, \dots, x_K)$. Denote $\bar{x} = (x_2, \dots, x_K)$. Then there exist a quadratic function q' and two linear functions l' and l'' on \mathbf{Z}_2^{K-1} , such that*

$$\sum_{x_1} (-1)^{q(x)} i^{l(x)} \propto (-1)^{q'(\bar{x})} i^{l'(\bar{x})} \delta_{l''(\bar{x}), 0}, \quad (\text{A.2})$$

where \propto denotes equality up to a multiplicative constant independent of \bar{x} . Moreover, determining q', l' and l'' from q and l can be performed efficiently.

The proof of the lemma is straightforward and is omitted. We now proceed with the proof of the claim.

Proof of claim: First, suppose that $|\psi\rangle$ is given by (2) and that a PHASE gate is applied to, say, the first qubit. As P maps $|a\rangle$ to $i^a|a\rangle$ ($a = 0, 1$) the resulting state $|\psi'\rangle$ is

$$|\psi'\rangle \propto \sum_{x \in A} (-1)^{q(x)} i^{l(x)} i^{x_1} |x\rangle, \quad (\text{A.3})$$

where x_1 is the first component of x . Using the identity $i^{*a} i^b = (-1)^{ab} i^{a+b}$, for every $a, b \in \mathbf{Z}_2$, then shows that $|\psi'\rangle$ is again of the form (2).

Second, suppose that a CNOT gate is applied to, say, the first and second qubit of $|\psi\rangle$. Note that CNOT maps $|a, b\rangle$ to $|a, a + b\rangle$, for every $a, b = 0, 1$. That is, CNOT performs a linear transformation (over \mathbf{Z}_2) “within the ket”. This immediately implies that the state (A.1) is mapped to

$$|\psi'\rangle \propto \sum_{u \in \mathbf{Z}_2^m} (-1)^{q(u)} i^{l(u)} |R'u + t'\rangle, \quad (\text{A.4})$$

for some appropriate R' and t' which are easily determined by performing the CNOT. Thus, also in this case the form (2) is kept.

Finally, suppose that a Hadamard gate is applied to, say, the first qubit of $|\psi\rangle$. This is the most nontrivial case of the three. Recall that H maps $|a\rangle$ to $\sum_{b=0}^1 (-1)^{ab} |b\rangle$ (where $a = 0, 1$). Denoting the first row of R by r^T , letting \bar{R} be the $(N-1) \times m$ matrix obtained by removing the first row of R and denoting $\bar{t} = (t_2, \dots, t_N)$, one has

$$|\psi'\rangle \propto \sum_{v=0}^1 \sum_{u \in \mathbf{Z}_2^m} (-1)^{q(u) + v \cdot (r^T u) + vt_1} i^{l(u)} |v, \bar{R}u + \bar{t}\rangle. \quad (\text{A.5})$$

Note that, as R has full rank m , \bar{R} may have either full rank m , or rank $m - 1$. If \bar{R} has full rank, then (A.5) is of the form (A.1) and we are done. If \bar{R} has rank $m - 1$, some additional work is required. In this case, there is exactly one nontrivial linear combination of the columns of \bar{R} which yields the identity. Without loss of generality, we may assume that the first column of \bar{R} can be written as a linear combination of the other columns, and

*Here, the exponent $a + b$ is computed modulo 2.

that the last $m - 1$ columns of \bar{R} are linearly independent. Denoting the columns of \bar{R} by c^i ($i = 1, \dots, m$), we therefore conclude that there exists a bit string $y = (y_2, \dots, y_m)$ such that $c^1 = \sum_{i=2}^m y_i c^i$. Note that determining y is a $\oplus\mathbf{L}$ problem as y is the solution to a system of linear equations over \mathbf{Z}_2 . It is then easy to verify that $\bar{R}Q = [0|c^2|\dots|c^m]$, where the $m \times m$ invertible matrix Q is defined by

$$Q = \begin{bmatrix} 1 & & & & \\ y_2 & 1 & & & \\ \vdots & & \ddots & & \\ y_m & & & & 1 \end{bmatrix}. \tag{A.6}$$

Making the substitution $u = Qu'$ in (A.5) yields an expression of the following form:

$$|\psi'\rangle \propto \sum_{v=0}^1 \sum_{u' \in \mathbf{Z}_2^m} (-1)^{q'(v,u')} i^{l'(u')} |v, \bar{R}Qu' + \bar{t}\rangle, \tag{A.7}$$

for some quadratic and linear functions q' and l' which can be easily determined. To lighten notation we will drop all primes in the above expression, i.e., $u' \equiv u$, $q' \equiv q$, $l' \equiv l$. Now note that the vector $\bar{R}Qu$ does not depend on the first coefficient of u , since the first column of $\bar{R}Q$ is zero. Indeed, denoting $\bar{u} = (u_2, \dots, u_m)$, we have $\bar{R}Qu = [c^2|\dots|c^m]\bar{u}$ (note also that $[c^2|\dots|c^m]$ has full rank). Therefore, in (A.7) the variable u_1 can be fully summed out. Using lemma A.1 then shows that $|\psi'\rangle$ is again of the desired form (2). This completes the proof.