Universal Quantum Compression with Minimal Prior Knowledge

Stuart Presnell, Richard Jozsa*

Abstract

A Universal Compression scheme is presented, to compress sequences of quantum information from unknown quantum sources (i.e. described by unknown density matrix ρ) asymptotically to $S(\rho)$, with fidelity bounded toward unity. The introduction of a "$\mathcal{B}$-diagonalisation" process allows us to treat the input as if it were diagonal in a known basis $\mathcal{B}$. Applying a version of the Lempel-Ziv algorithm in this basis "condenses" the sequence, leaving a tail of unentangled qubits which are then "truncated". The rate attained depends upon the basis chosen; but we then show that the process may be repeated iteratively, to search for an optimal computational basis, and thereby compress to $S(\rho) + \delta$ qubits per signal, for any $\delta > 0$.

Keywords: quantum information theory, compression, Lempel-Ziv algorithm.

1 Introduction

If and when quantum computing becomes a viable technology, the efficient storage and transmission of quantum information will be an important technological requirement. In the present work we study the limits of achievable compression of quantum information under the constraint of zero prior knowledge of the source statistics. In this respect our approach differs from most previous studies of quantum compresion [1][2][3][4] which require the encoder to know something about the source - its eigenvalues, or a bound on its entropy, for example.

We demonstrate that long sequences from an unknown i.i.d. source can be compressed to the Schumacher limit whilst maintaining high fidelity. The assumption that the quantum source is i.i.d. is not necessary for the argument, since the Lempel-Ziv algorithm on which the present work is based successfully compresses non-i.i.d. sources (e.g. the English language). The extension to more general sources therefore seems promising, but is left to further work.

A quantum source [1] may be described by a density matrix $\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|$ where $|\psi_i\rangle$ are pure states emitted by the source, and $\{p_i\}$ is a probability distribution. We assume that neither the eigenvalues nor eigenvectors of ρ are known to the encoder. The process of compressing a sequence of N quantum symbols will in general consist of some unitary manipulations of the state $\rho^{\otimes N}$, and some operations which are not unitary. Thus the whole procedure can be naturally divided into two separate stages - first a unitary "condensation" stage, to be performed by a quantum computer, and then a non-unitary "truncation" stage, in which the number of qubits used is reduced. This separation simplifies the analysis, and help us to bring in ideas from classical compression.

2 Condensation

The classical Lempel-Ziv compression scheme [5] asymptotically compresses the output of an i.i.d. source with unknown probability distribution to H bits per signal, where H is the Shannon entropy of the distribution. It relies upon the fact that at any time in the decoding process, there is a significant quantity of data that is known to both sender and receiver. Using this shared resource, by transmitting references to previous sections of data, the sender can more efficiently encode further signals from the same source. The Lempel-Ziv compression is lossless; its compression rate is therefore optimal only asymptotically.

We can construct a reversible [6] version of the Lempel-Ziv algorithm, which can be implemented on a quantum computer. The resulting quantum algorithm acts upon the orthonormal eigenstates of the computational basis as if they were classical signals, and treats the input sequence as a superposition of "pseudo-classical" sequences, each of which is condensed independently.

We demonstrate that this produces the required condensation of the whole sequence, by introducing a process called "$\mathcal{B}$-diagonalisation" which allows us to treat the input as if it were diagonal in a known basis $\mathcal{B}$. This process effectively commutes with the reversible Lempel-Ziv algorithm, and can therefore be viewed as merely a mathematical convenience, rather than a procedure we need physically to carry out.

The introduction of this procedure leads to a simple formula for the rate of condensation achieved for an arbitrary

*Merchant Venturers' Building, Department of Computer Science, University of Bristol, BRISTOL, UK

source density matrix ρ and arbitrary basis $\mathcal{B}$:

$$R_c(\rho, \mathcal{B}) = -\sum_{|e_j\rangle \in \mathcal{B}} \langle e_j|\rho|e_j\rangle \log\langle e_j|\rho|e_j\rangle \tag{1}$$

In the special case where the computational basis is the same as the eigenbasis of the source, the formula reduces to $S(\rho)$, as expected.

3 Truncation

Having condensed the sequence, we must now truncate it, removing the unentangled "tail". In doing this, we are effectively estimating the condensation rate achieved, R_c, which depends upon ρ and is therefore not known.

We define a POVM

$$\bar{\Pi}_l = \frac{1}{Y}\sum_{i=0}^{Y-1} \pi_{(l+i)} \tag{2}$$

where Y is some parameter we are free to choose, and π_l is a projector [4]:

$$\pi_l = I^{1 \cdots l} \otimes |0^{l+1 \cdots N}\rangle\langle 0^{l+1 \cdots N}| \tag{3}$$

which acts on a sequence of N qubits, projecting onto the subspace in which the last $(N-l)$ qubits are in the state $|0\rangle$. This POVM is therefore an average of Y offset projectors.

Thinking of the condensed sequence as a superposition of different sequences of varying lengths, we assume (from Shannon's Noiseless Coding Theorem) that the components with lengths significantly greater or less than R_c have exponentially small amplitude. Given this assumption we demonstrate that for any achieved condensation rate R_c, a sequence of POVMs as defined above with parameter Y will allow us to truncate to $R_c + \frac{2Y}{N}$ qubits per signal, with probability of error no greater than $\frac{1}{Y}$ and overall drop in fidelity bounded by $\frac{2}{Y}$.

4 Iteration

Having demonstrated that we can compress the original sequence while maintaining arbitrarily high fidelity and arbitrarily low probability of error, we can repeat the process with a different computational basis, iterating the compression-decompression process to search through the space of bases. Since that space is compact, we argue that for any ρ and any $\delta > 0$ there exists a compact set of orthonormal bases $\{\mathcal{B}_\delta\}$ such that $R_c(\rho, \mathcal{B}_\delta) - S(\rho) \leq \delta$. The value of the threshold δ gives the "volume" of $\{\mathcal{B}_\delta\}$ independently of $S(\rho)$, and so we can establish a mesh of $q(\delta)$ evenly-spaced points in the space of bases, choosing $q(\delta)$ large enough such that at least one point lies within $\{\mathcal{B}_\delta\}$. The value of $q(\delta)$ is therefore an upper bound on the number of iterations we will need to perform, and determines the minimum value of Y needed to achieve the specified overall fidelity.

Therefore with specified tolerances for error probability and fidelity loss and specified threshold δ, we can always choose values of Y and N large enough to ensure that those criteria are met while we compress to $S(\rho) + \delta$ qubits per signal.

5 Discussion

While the existence of universal quantum compression has already been demonstrated by Hayashi and Matsumoto [7], we feel that the present method is conceptually and technically clearer, and may be easier to implement in practice. Also, the introduction of the $\mathcal{B}$-diagonalisation step suggests that other classical deterministic algorithms could be translated into quantum algorithms in the same way, which may help us to "bridge the gap" between the two regimes of computing.

References

[1] B. Schumacher. Quantum coding. *Phys. Rev. A*, 51:2738, 1995.

[2] R. Jozsa, M. Horodecki, P. Horodecki, and R. Horodecki. Universal quantum information compression. *quant-ph/9805017*, 1998.

[3] S.L. Braunstein, C.A. Fuchs, D. Gottesman, and H.K. Lo. A quantum analog of huffman coding. *quant-ph/9805080*, 1998.

[4] B. Schumacher and M.D. Westmoreland. Indeterminate-length quantum coding. *quant-ph/0011014*, 2000.

[5] J. Ziv and A. Lempel. A universal algorithm for sequential data compression. *IEEE Trans. Inf. Theory*, IT-23:337, 1977.

[6] C.H. Bennett. Logical Reversibility of Computation. *IBM J. Res. Develop.*, 17:525, 1973.

[7] M. Hayashi and K. Matsumoto. Quantum universal variable-length source coding. *quant-ph/0202001*, 2002.