Quantum Property Testing[‡]

Harry Buhrman[§] Lance Fortnow[¶] Ilan Newman[‖] Hein Röhrig[§]

Suppose we have a large data set, for example, a large chunk of the world-wide web or a genomic sequence. We would like to test whether the data has a certain property, but we may not have the time to even look at the entire data set or even a large portion of it.

To handle these types of problems, Rubinfeld and Sudan and Goldreich, Goldwasser and Ron have developed the notion of property testing. Property testing comes in many varieties including graph properties, algebraic properties of functions and regular languages.

In this model, the property tester has random-access to the n input bits similar to the black-box oracle model. The tester can query only a small, usually some fixed constant, randomly-chosen set of bits of the input. Clearly we cannot determine from these small number of bits whether the input sits in some language L. However, for many languages we can distinguish the cases as to whether the input is in L from whether the input differs from all inputs in L by some constant fraction of input bits.

Since we have seen many examples where quantum computation gives us an advantage over classical computation one may naturally ask whether using quantum computation may lead to better property testers. By using the quantum oracle-query model developed by Beals, Buhrman, Cleve, Mosca and de Wolf, we can easily extend the definitions of property testing to the quantum setting.

Beals, et. al. have shown that for all total functions we have a polynomial relation between the number of queries required by quantum machine and that needed by a deterministic machine. For separations one needs to require a promise in the input and the known examples, such as those due to Simon and Bernstein-Vazirani, require considerable structure in the promise. In property testing you have a promise of either being in the language or far from any input in the language. This promise would seem to have too little structure to give a separation but in fact one can prove that quantum property testing does improve on classical testing.

We show that any subset of Hadamard codes have a quantum property tester and most subsets would require $\Theta(\log n)$ queries to test with a probabilistic tester. This shows that indeed quantum property testers are more powerful than classical testers.

Perhaps every language has a quantum property tester. We show that this is not the case. We prove that for most properties of a certain size, any quantum algorithm requires $\Omega(n)$ queries. We then show that a natural property, namely, the range of a d-wise independent pseudorandom generator cannot be quantumly tested with $d/2$ queries for any odd $d = O(n/\log n)$.

[‡]Research done while all authors were visiting the NEC Research Institute.

[§]CWI and University of Amsterdam. Partially supported by the European Union fifth framework project QAIP, IST-1999-11234.

[¶]NEC Research Institute

[‖]Haifa University