Abstract

We consider the *Online Boolean Matrix-Vector Multiplication* (OMV) problem studied by Henzinger *et al.* [STOC'15]: given an $n \times n$ Boolean matrix M, we receive n Boolean vectors $v_1, \ldots, v_n$ one at a time, and are required to output Mv_i (over the Boolean semiring) before seeing the vector v_{i+1}, for all i. Previous known algorithms for this problem are combinatorial, running in $O(n^3/\log^2 n)$ time. Henzinger *et al.* conjecture there is no $O(n^{3-\varepsilon})$ time algorithm for OMV, for all $\varepsilon > 0$; their OMV conjecture is shown to imply strong hardness results for many basic dynamic problems. We give a substantially faster method for computing OMV, running in $n^3/2^{\Omega(\sqrt{\log n})}$ randomized time. In fact, after seeing $2^{\omega(\sqrt{\log n})}$ vectors, we already achieve $n^2/2^{\Omega(\sqrt{\log n})}$ amortized time for matrix-vector multiplication. Our approach gives a way to reduce matrix-vector multiplication to solving a version of the Orthogonal Vectors problem, which in turn reduces to "small" algebraic matrix-matrix multiplication. Applications include faster independent set detection, partial match retrieval, and 2-CNF evaluation. We also show how a modification of our method gives a cell probe data structure for OMV with worst case $O(n^{7/4}/\sqrt{w})$ time per query vector, where w is the word size. This result rules out an unconditional proof of the OMV conjecture using purely information-theoretic arguments.