Factoring semi-primes with (quantum) SAT-solvers

Michele Mosca*[1] and Sebastian R. Verschoor†[2]

[1,2]Institute for Quantum Computing, University of Waterloo, Canada
[1]Department of Combinatorics & Optimization, University of Waterloo, Canada
[2]David R. Cheriton School of Computer Science, University of Waterloo, Canada
[1]Perimeter Institute for Theoretical Physics, Waterloo, Canada
[1]Canadian Institute for Advanced Research, Toronto, Canada
[1]evolutionQ Inc., Waterloo, Canada

October 22, 2019

Abstract

The assumed computationally difficulty of factoring large integers forms the basis of security for RSA public-key cryptography, which specifically relies on products of two large primes or *semi-primes*. The best-known factoring algorithms for classical computers run in sub-exponential time. Since integer factorization is in NP, one can reduce this problem to any NP-hard problem, such as Boolean Satisfiability (SAT). While reducing factoring to SAT has proved to be useful for studying SAT solvers, attempting to factor large integers via such a reduction has not been found to be successful.

Shor's quantum factoring algorithm factors any integer in polynomial time. Large-scale fault-tolerant quantum computers capable of implementing Shor's algorithm are not yet available, so relevant benchmarking experiments for factoring via Shor's algorithm are not yet possible. In recent years, however, several authors have attempted factorizations with the help of quantum processors via reductions to NP-hard problems. While this approach may shed some light on some algorithmic approaches for quantum solutions to NP-hard problems, in this paper we study and question the practical effectiveness of this approach for factoring large numbers. We find no evidence that this is a viable path toward factoring large numbers, even for scalable fault-tolerant quantum computers, as well as for various quantum annealing or other special purpose quantum hardware.

1 Introduction

In this work we focus on the problem of factoring semi-primes with SAT-solvers. A semi-prime N is a composite of two primes p and q which are roughly of equal size. These particular composites are conjectured to be hard to factor, in the sense that no (classical) algorithm or heuristic is known to factor semi-primes using only polynomially many resources. This problem has great relevance for the RSA cryptosystem [49], a widely-deployed public-key cryptosystem. The RSA cryptosystem is founded upon the difficulty of factoring integers: the existence of an efficient factoring algorithm would completely break its security.

*michele.mosca@uwaterloo.ca
†srverschoor@uwaterloo.ca

Some authors have proposed an alternative approach they refer to as quantum factoring, and it is occasionally even cited in benchmarks for factoring [61]. In this paper, we explain why these approaches, while potentially helpful for studying quantum SAT-solving, are not likely a viable approach to integer factorization and, very importantly, are not a meaningful benchmark for people interested in quantum cryptanalysis of cryptosystems based on the integer factorization problem.

We attempt to generously extrapolate the kinds of speed-ups one might expect for a range of quantum solvers, and find no evidence that this is a viable path toward factoring large numbers, even for scalable fault-tolerant quantum computers, as well as for various quantum annealing or other special purpose quantum hardware.

Some researchers only implement quantum factoring for the purposes of benchmarking the experimental apparatus. There are several more relevant algorithms to implement for the purposes of benchmarking, such as work on randomized benchmarking [24] or implementations of quantum error correction. Framing the experiments as implementations of quantum factoring can easily be misinterpreted as a meaningful benchmark toward large-scale integer factorization, and we explain in this article why they are not.

For many years cryptographers have tracked and benchmarked progress in classical factorization and attempted extrapolations with an interest in estimating when RSA schemes with moduli of a given length may be broken using the number field sieve [34, 2]. The extrapolations take into account estimates of computing power increase and algorithmic improvements.

This paper highlights why none of the current literature on experimental implementations of quantum factoring serves the same purpose. In the absence of a breakthrough that demonstrates factoring can be meaningfully sped up without a fault-tolerant quantum computer, this sort of tracking of the size of numbers quantumly factored will only be meaningful after the implementation of several logical qubits.

One caveat and challenge with tracking and extrapolating is that once fault-tolerant quantum computers start factoring small numbers, a constant factor increase in available quantum resources brings a constant *factor* increase in the size of the number that can be factored (i.e. we go from being able to factor n-bit numbers to being able to factor (cn)-bit numbers for some $c > 1$ that depends on the factor of increase in time and memory) because Shor's algorithm runs in polynomial time. On the other hand, a constant factor increase in classical computing resources only implies being able to factor numbers that are a few bits larger using the number field sieve (i.e. we go from being able to factor n-bit numbers to being able to factor $(n+o(n^{2/3}))$-bit numbers). Given these quantum scalings, it will be much harder to reliably extrapolate the size of numbers that can be quantumly factored, and a relatively small change in computing resources or a relatively small algorithmic improvement can have a significant impact on the size of the number that can be quantumly factored. This is one reason why it is valuable to have post-quantum cryptography ready for wide-scale deployment before fault-tolerant quantum computers are available.

The Boolean satisfiability problem (SAT) asks whether there exists an assignment to the Boolean variables of a given propositional logic formula such that the formula evaluates to TRUE. This problem was the first that was proven to be NP-complete [16, 36]. Since no algorithms with polynomial runtime for NP-hard problems are known, solving NP-hard problems has long been considered to be intractable for real-world computers. Despite this result, coming from asymptotic analysis, modern SAT-solvers perform very well on solving large SAT instances originating from industry and academics, with formulas that have up to a million clauses [6]. At the moment of writing there exists no good general method or metric to predict if a given SAT instance is hard to solve. For practical applications it therefore makes sense to assess the performance of the solvers on the investigated instances by careful benchmarking instead of doing asymptotic analysis.

The original goal of this project was to encode the RSA factoring challenges [29] to SAT instances and see how well modern SAT solvers would perform on those instances. The smallest semi-prime of these challenges is RSA-100: a 100-digit or 330-bit number. This number was factored in a few days almost immediately after the challenge was posted [21] in 1991, whereas the current record for

factoring stands at factoring RSA-768: a 768-bit semi-prime [33]. The intention was to compare current state-of-the-art SAT solvers against the numerical results from 1991, but it turns out that even the smallest RSA semi-prime poses too big of a challenge for these solvers.

1.1 Contributions

This work provides a numerical analysis on the hardness of factoring numbers by solving the corresponding satisfiability problem, thereby confirming the folklore that factoring numbers does indeed give "hard" SAT instances. This is done by measuring the speed of the currently fastest SAT solver. We justify the choice of numerical analysis over theoretical asymptotic analysis by applying some common analysis tools from modern SAT solving theory and the observation that the tools provide no good prediction for the actual runtime. We extrapolate the numerical results to investigate the asymptotic behavior of the solver and compare the results with the asymptotics of factoring with numerical algorithms. Finally, the results are used to estimate an upper bound on the speedup that can be achieved on this specific problem using currently known quantum algorithms.

As a minor contribution, we developed a tool that can create smaller SAT instances for factoring[1] than any other publicly available tool. This tool and scripts for generating semi-primes and reproducing the results of this paper have been made available online [60].

2 SAT instances

An instance of the SAT problem is a formula in Boolean propositional logic: every *variable* (x) can take the value TRUE or FALSE as specified by the respective *literals* x and $\bar{x}$. This work considers the equivalent [32] CNF-SAT where all formulas are in conjugate normal form (CNF): each formula must be a conjunction of disjunctions of literals.[2] The disjunctions are often called *clauses*. A satisfying assignment gives a value to each variable such that at least one literal evaluates to TRUE in every clause. All tools we used for generating and solving SAT instances work with the DIMACS format which specifies formulas in CNF form.

Another (equally hard) formulation of the problem is called CircuitSAT: given a Boolean circuit with a single output, is there an input such that the output is TRUE? One can translate any Boolean circuit into a Boolean formula: assign a variable to each wire and let the clauses describe the gates. For example the Turing complete NAND-gate with input wires x, y and output wire z has the corresponding formula $(x \vee z) \wedge (y \vee z) \wedge (\bar{x} \vee \bar{y} \vee \bar{z})$. Simulating gate execution is done by fixing a value on the input wires: for example by adding the clauses $x \wedge \bar{y}$. A SAT-solver can examine those five clauses and find that the only satisfying assignment sets $z =$ TRUE. Combining gates to make a circuit is done by reusing output variables of earlier gates as input variables in later gates.

More interesting is to fix a value on the output variables of a circuit and ask the SAT-solver to find a satisfying assignment. For example adding the clause z to the NAND-gate gives three satisfying assignments: $x \wedge \bar{y}$, $\bar{x} \wedge y$, and $\bar{x} \wedge \bar{y}$. In general a circuit might have zero or more satisfying assignments. Effectively the SAT-solver is finding preimages to the function described by the circuit. An immediate cryptanalytic application that springs to mind is finding preimages to secure hash functions: indeed this has been done with varying results [41, 42, 23]. More general cryptanalytic applications can be found throughout literature [40] and occur in modern benchmarks [6], although asymmetrical cryptographic primitives are rarely targeted.

This work examines circuits that encode the multiplication of two integers p and q. We fix the multiplication output bits of the circuit to the bit-values of the semi-prime N and ask the SAT-solver to find a satisfying assignment. Only two exist[3]: those representing $N = pq$ and $N = qp$, so from the assignment one can read the factorization of N. For the remainder of this paper n represents the size

[1] using long multiplication

[2] Further restricting each clause to exactly three literals would give the equivalent 3SAT problem.

[3] The specific encoding ensures the trivial solutions $N = 1N$ and $N = N1$ do not give satisfying assignments.

of N in bits. We limit p and q similar to how the RSA cryptosystem limits its parameters: both need to be equally sized primes. We interpreted this last requirement to mean that their most significant bit may differ by at most one position.

2.1 Encoding

Despite the asymptotic worst-case exponential runtime associated with SAT instances, it turns out to be non-trivial to generate "hard" SAT instances: instances where the solver runtime grows exponentially in the number of variables. For many instantiations of the SAT problem, it turns out that the average case can be solved relatively efficient with modern SAT solvers. Specialized tools such as ToughSat [64] exist that can generate SAT instances that are hard on average, based on problems such as integer factorization.

Multiplying larger integers requires larger circuits, which leads to instances with more variables and clauses, which leads to longer solving times. However, there are many choices to make when computing multiplication in a circuit and each choice will lead to different encodings of the SAT instance and a different solver runtime. For SAT solvers in general it turns out that the details of the encoding of a problem (beyond metrics such as number of variables and clauses) can have a significant impact on the solver runtime. The first choice is to consider different multiplication algorithms: a simple one and a more complex encoding that in theory leads to smaller instances.

Long multiplication (or schoolbook multiplication) is computed by multiplying p by each digit (bit) of q and adding the shifted results. For multiplying two m-bit numbers (where $m = n/2$) this requires $\Theta(m^2)$ bitwise multiplications and additions. The exact number of operations depends mainly on the circuit used for addition: our tool for generating instances [60] minimizes the number of both variables and clauses by maximizing the number of full-adders used in the circuit. Counting the variables in the generated instances and applying regression reveals that the number of variables grows approximately as $0.750n^2 + 0.496n - 2.05$ and similarly the number of clauses grows as $4.25n^2 - 4.01n - 9.87$ with on average 3.31 literals per clause.

Karatsuba multiplication [31] asymptotically improves upon long multiplication by a divide-and-conquer strategy and requires only $\Theta(m^{\log_2 3})$ multiplications at the cost of requiring more additions. The instances we tested were generated by the ToughSat application [64] and contain approximately $2.59n^{\log_2 3} - 7.57n + 8.75$ variables and $61.5n^{\log_2 3} - 170n - 386$ clauses with on average 6.77 literals per clause. Inspection of the generated instances reveals that the Karatsuba circuits were built from more complex gates, which explains why there are more literals per clause. It is likely that building the Karatsuba circuit with a similar gate set would increase the number of variables and clauses by another (constant) factor.

Asymptotically the Karatsuba algorithm is not the best known algorithm and is outperformed by for example Toom-Cook or FFT-multiplication. Given that these methods introduce additional overhead for small instances and given the minor difference in the runtime of long multiplication and Karatsuba (see Section 3), it appears that the cross-over point where these algorithms are faster vastly exceeds a feasible instance.

Hardware design provides alternative multiplication algorithms, which are often optimized to minimize latency and for various other physical constraints. There is no indication that these optimizations are related to optimizations that lead to smaller and/or easier SAT instances. In fact our adder encoded in the SAT instances minimizes the number of half-adders required, which gives the smallest number of variables and clauses and results in the fastest SAT solver times, but the resulting clauses encode a circuit that would give extremely high latency if built from physical components.

Since the multiplication circuit is the same for each semi-prime of the same bitlength there is an alternative strategy we can apply when we want to factor only one of many semi-primes. We encode the multiplication circuit once and then "fanout" the resulting wires to circuits that check if the output equals a semi-prime. Those results are combined with a large OR-gate, so that the entire instance evaluates to TRUE if the multiplication outcome is equal to any of the semi-primes. By inspecting which values were assigned on the circuit input wires by the solver we learn which of the semi-primes

it actually factored. The idea behind this encoding is that if there is an easy semi-prime somewhere in the input, then the solver itself may detect this and focus on solving that instance. As long as we encode only polynomially many semi-primes in the instance, the total instance size will remain polynomial.

An alternative solution for factoring numbers with SAT is to encode the integer division circuit $N/p = q + r$ and fixing the input value N and output remainder $r = 0$. The rationale for this encoding is that the solver would only have to assign values to the bits of p and can then deterministically evaluate the entire circuit and check if the remainder is zero. However, in practice this encoding leads to substantially larger SAT-instances and tests with various solvers indicate that solving such instances is significantly slower, so we did not investigate this encoding any further.

A more promising approach is to reduce some subroutine of the NFS to SAT where there is little or no increase in complexity by mapping to SAT, analogous to the approach taken in [8]. In this case, even a small quantum speed-up will lead to a faster integer factorization algorithm. This approach is studied in detail in [44].

3 Classical Solvers

Modern SAT solvers come in two classes. Conflict-Driven Clause Learning (CDCL) [19, 18] combines conflict analysis with branch heuristics to systematically backtrack the search-space of an instance. Stochastic local search approaches such as employed by WalkSAT [51] or simulated annealing combine randomized assignments with probabilistic updates to find assignments that minimize the number of clauses violated. We found that for the semi-prime instances CDCL solvers outperformed the local search solvers by an order of magnitude. The scope of this project is limited to the black-box analysis of publicly available SAT solvers. This means we will not investigate the internals of the solvers for analysis of the runtime, nor do we allow domain-specific knowledge to speed up solver times.

We tested the MapleCOMSPS [38] SAT solver for the simple reason that at the time of running the benchmarks this was the fastest solver according to the SAT Competition 2016 [27]. We compiled and ran the solver with default settings, except for the random seed which was fixed for each call to the solver to ensure reproducibility of the results.

Another solver that we tested is CryptoMiniSat 5 [55], because it has "Automatic detection of cryptographic [...] instances" [54]. One might consider this to be cheating by using domain-specific knowledge and therefore it should not be included in the benchmarks. CryptoMiniSat appears to focus on symmetric cryptography and appears to provide no speedup on public cryptography instances, which we confirmed during an initial round of benchmarking. We inspected the (partial) results and found that CryptoMiniSat 5 was consistently being outperformed by MapleCOMSPS. For this reason we did not further analyze this solver, but the results can be found in Appendix A.

All measurements were performed on a ThinkPad laptop with a 64-bit Intel Core i5–4200M (Haswell) CPU running at 2.50GHz. All measurements were executed sequentially and on a single core. Where applicable we use regression to fit a line to the data and the goodness-of-fit is quantified by the r^2 parameter.

3.1 Results

Usually when analyzing the runtime of a randomized algorithm we are interested in the expected runtime: the mean computed over the random bits. We do this by factoring the same number multiple times using a different PRNG-seed for the solver and average the runtime to compute the expected runtime numerically. We are interested in the asymptotics: the growth of the runtime as a function of the size of its input, so we group the semi-primes by their bitlength n (100 semi-primes per bitlength) and plot the mean runtime of solving five times. The results are given in Figure 1a and are showing an exponential trend. The green line is fitted against the median runtime of all semi-primes of the same bitlength.

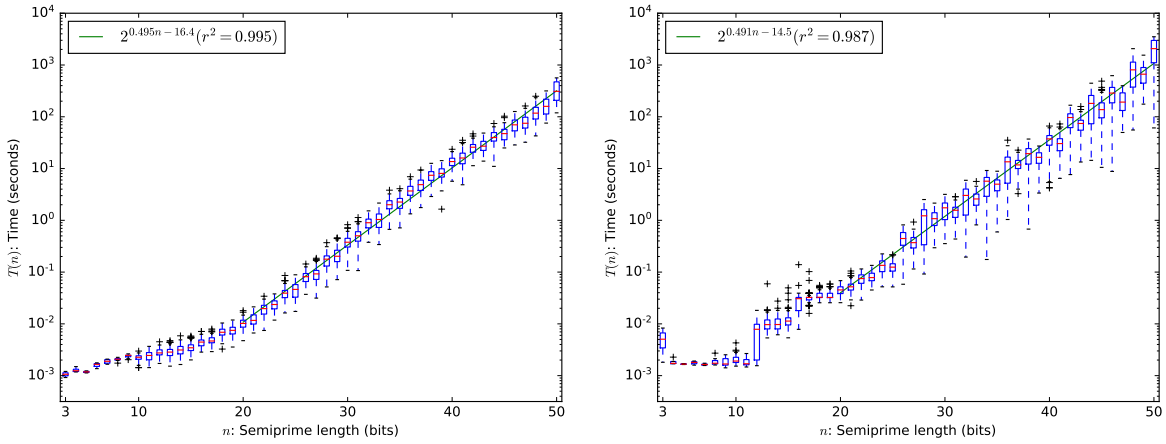

(a) schoolbook multiplication

(b) Karatsuba multiplication

Figure 1: Runtime of MapleCOMSPS on factoring semi-primes.

We repeated the same experiment for multiplication with the Karatsuba algorithm. The results are given in Figure 1b: note that asymptotic runtime has improved somewhat over schoolbook multiplication at the cost of a larger constant. We conclude that changing the multiplication algorithm does not make factoring with SAT solvers efficient. Since the larger constant dominates the runtime at this small scale, we will consider schoolbook multiplication for the remainder of our experiments.

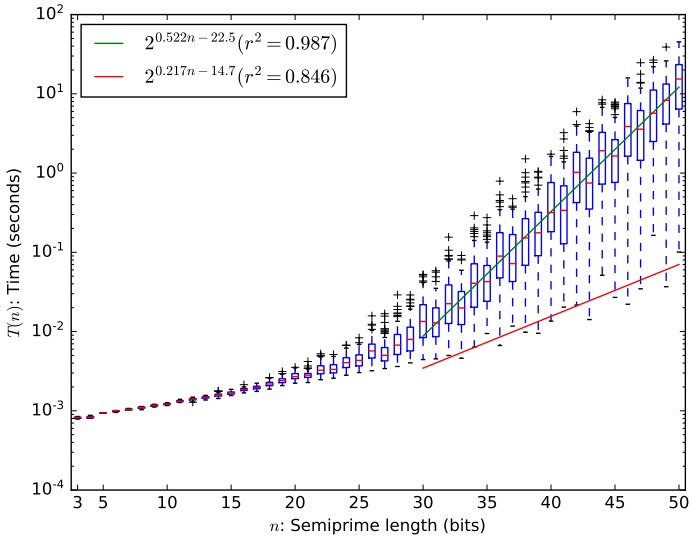

Figure 2: Minimum runtime of MapleCOMSPS on factoring semi-primes using schoolbook multiplication.

An alternative strategy for factoring is to run several solvers in parallel and wait for the first one to return a solution. We simulate this strategy by taking the minimum solver time of solving the same instance with the solver initialized with 100 different random seeds for 100 semi-primes per bitlength: the results are given in Figure 2. Asymptotically the runtime became worse by employing this strategy. Note that this strategy does push down the constant by approximately $2^{6.1}$. Since this is smaller than 100 it does not lead to a lower expected runtime on this small scale when we consider

the total runtime of all parallel solvers.

We can also see in Figure 2 that some semi-primes are significantly easier to solve than others with this strategy. Even if we only manage to factor some semi-primes that may be important to (for example) cryptography. For this method to be asymptotically efficient, it is required that the runtime is pushed down exponentially for more than just negligibly many cases. To see if it does we can inspect the distribution of the solver runtime given different seeds. Here we focus on three different semi-primes[4]: the easiest, average and hardest semi-prime from the 100 semi-primes of 35 bits, where hardness is defined by the expected (mean) solve time computed over 360 seeds.

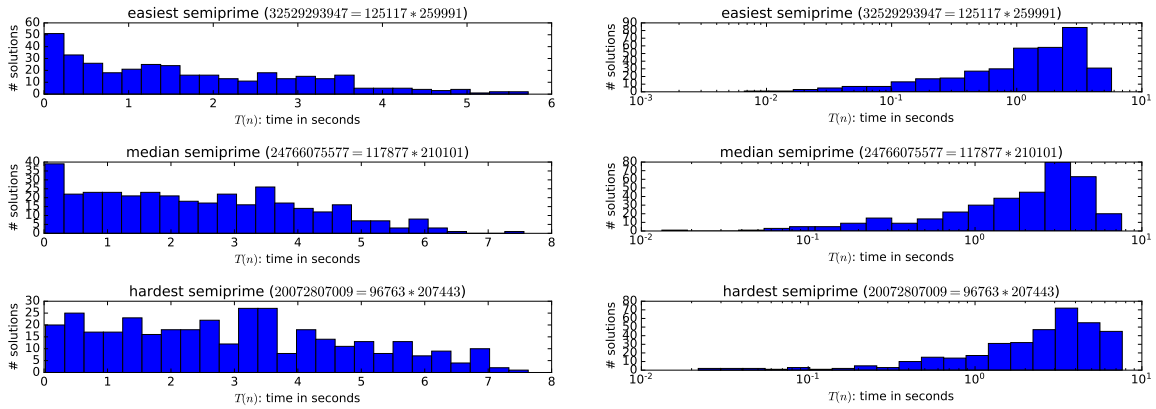

(a) Linear time-scale

(b) Logarithmic time-scale

Figure 3: Histogram of the MapleCOMSPS runtime on factoring semi-primes using schoolbook multiplication.

Although no strong conclusions should be drawn from the results in Figure 3a, the distribution does suggest that running a few parallel solvers may lower the total runtime. To see if it may be considered efficient we again inspect the distribution but this time on a logarithmic scale: see Figure 3b.

This data suggests that even if the method could push down the runtime significantly for any semi-prime, it only does so with negligible probability. Another way of interpreting this data is that employing parallel SAT solvers to factor a semi-prime does not appear to be better than employing a single solver.

The last strategy we investigate is that of encoding multiple semi-primes into a single instance. We encoded 100 semi-primes per bitlength in each instance and solved it 100 times using different seeds. The results are given in Figure 4 Note that whereas the vertical boxplots in previous plots show a distribution over different primes, here a distribution over different solver PRNG-seeds is shown. From the data we conclude that this strategy is less efficient than solving instances with a single semi-prime. From inspection of the solver solution we can see which semi-prime was factored (see [60]). This reveals that some semi-primes in the same instance are factored more often than others, suggesting that these are easier to factor by the solver, although we note that these are not "easy enough" to make the overall method efficient.

3.2 Patterns

The above results support that it is hard *on average* to factor a number with SAT solvers, but we can also observe that some numbers are easier to factor than others. It is an interesting question whether there is some structure in the semi-primes that is picked up by solver that allows it to factor more efficiently or whether the solver's heuristic choices accidentally lead to a faster solution. We try to

[4]the distribution for all other semi-primes can be generated at [60]

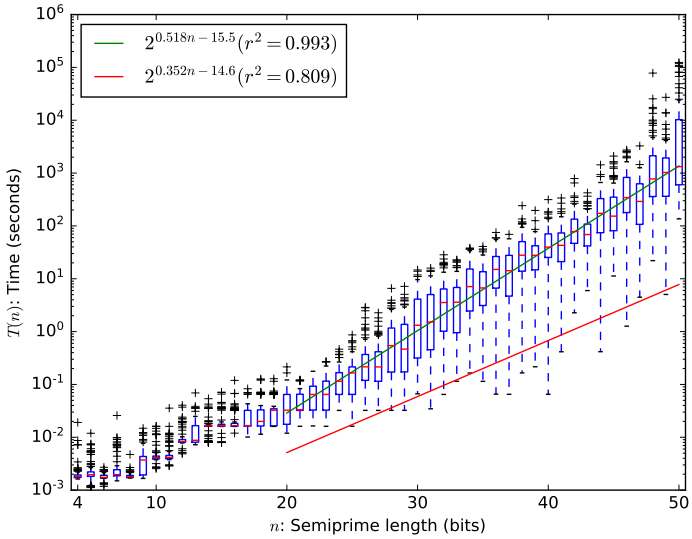

Figure 4: Runtime of MapleCOMSPS on factoring one of 100 semi-primes encoded in each instance using schoolbook multiplication.

answer this question by inspecting the instances using two analytic methods from the SAT literature (backdoors and community structure) and we do some manual inspection of the instances. Because we are interested in the fastest solver time, we focus on the minimum solver time per instance given different random seeds.

3.2.1 Backdoors

Backdoors in SAT instances were introduced by Williams, Gomes and Selman [62]. A backdoor is a subset of variables such that setting these variables to any value allows a so-called subsolver to assert if the entire formula is satisfiable in polynomial time. If the solver can find such a backdoor of size k with a subsolver that runs in time l, the entire solver can run in time $O(l2^k)$. Any CircuitSAT instance has a trivial backdoor in the form of the input wires/variables: set these and the rest of the clauses can be determined deterministically.[5] A backdoor subset for CircuitSAT therefore only becomes interesting when it is smaller than the set describing the input variables.

Every instance has n input wires, but the solver runtime suggests that a backdoor of $k \approx n/2$ variables was found. Given the structure of the problem this is not surprising (division to find the other k input bits only takes polynomial time), but it is somewhat surprising given that it is unlikely that the SAT solver was programmed to perform this division. More meaningful analysis of this observation would require inspecting the internals of the solver to look for potential subsolvers and backdoor detection capabilities. We consider this outside the scope of this project. We simply conclude that even if a backdoor of size $n/2$ is found then the runtime of the SAT solver would still be exponential and therefore would not impact the security of RSA.

3.2.2 Community structure

A SAT instance can be represented as a graph where each variable is a vertex and an edge is drawn between vertices when the variables occur in the same graph. The community structure of a graph is often characterized by a quality metric Q (also known as the modularity of the graph). According to Newsham, Ganesh, Fischmeister, Audemard and Simon [45] the community structure of a SAT

[5] This is why we also encoded a division circuit: the input to that circuit contains one prime instead of two.

instance should provide us with a good prediction on how hard it is to solve that instance: instances are harder to solve when $0.05 \leq Q \leq 0.12$.

An immediate problem that occurs when applying the above theory to the generated circuits is that all instances for semi-primes of the same bitlength have the same structure: the encoded circuit is simply an m by m bit multiplier. Therefore, we compute the community structure only on the instances after they are simplified by the solver's preprocessor. We approximate the value of Q with the greedy algorithm by Clauset, Newman and Moore [15].

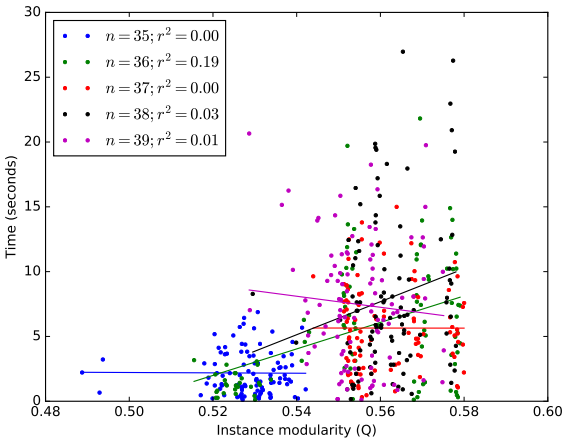

Figure 5: Solver times per community structure. We only display the results for some values of n to avoid more clutter, but similar results hold for all n.

Even after this preprocessing step the instances for long multiplication have too little variation to conclude anything about the relation between Q and solver time. For some Karatsuba instances the results are given in Figure 5. The results are grouped according to the bitlength n and per group linear regression is applied to each group. The low r^2-values suggest there is no relation between modularity and solver time for these instances.

Interestingly, the values of Q are relatively high and far outside the range $0.05 \leq Q \leq 0.12$ for which the instances were conjectured to be hard, yet the instances are still hard to solve. The above data leads to the conclusion that the community structure does not provide a good prediction for solver runtime when applied to SAT instances that encode multiplication circuits.

3.2.3 Other metrics

Besides the above metrics that can be computed on any SAT instance, one might consider if there is any correlation between metrics that apply only to this specific use case. In particular we are interested if there is any pattern in N, p and/or q that the solver is able to exploit for a faster solving time. Since SAT instances are defined over Boolean variables we considered the Hamming weight of: N, p, q, and $p \oplus q$. We also measure if the solver is able to pick up on some patterns that make a number easier to factor according to number theoretic methods (such as Pollards $p-1$ method): smoothness of $p-1$, smoothness of $q-1$, $|p-q|$, and $\log N$. We measured the correlation with the solver time (see Appendix B for details). No metric shows any significant correlation.

3.3 Comparison to number-theoretical methods

One can put the above results in context by comparing the absolute runtime to that of other number-theoretical results. Using SageMath [20] we measured the runtime of two approaches: factoring with the built-in factor function: Figure 6a and factoring by trial division: Figure 6b.

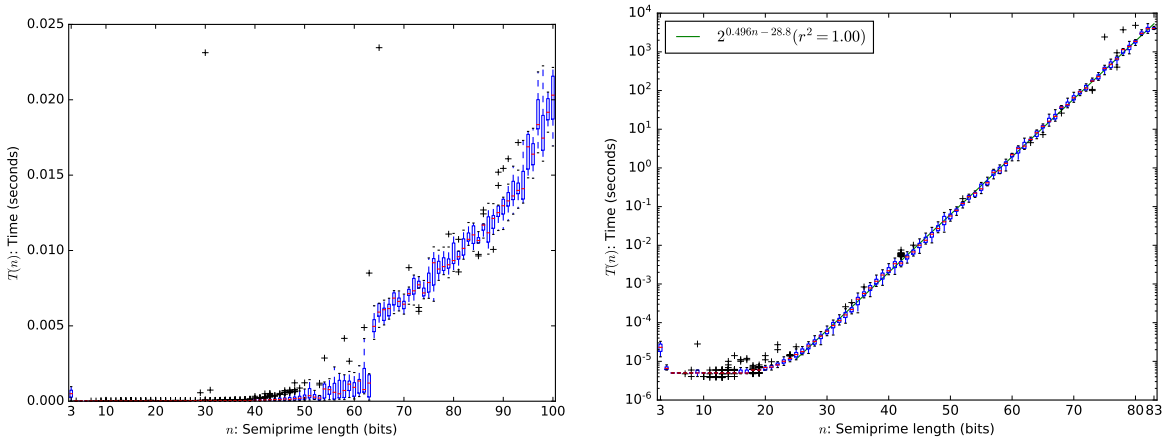

(a) SageMath `factor()`

(b) Trial division

Figure 6: Runtime of factoring using numerical methods. No randomization was applied for obtaining these results.

SageMath is able to factor almost all semi-primes up to a 100 bits in under 0.025 seconds. The tested semi-primes are so small that the asymptotic behavior of the underlying algorithm is not even visible yet, so there is no point in extrapolating these results. In fact the crossover point where the number field sieve (NFS) is faster than asymptotically slower methods such as the quadratic sieve and the elliptic-curve method is much larger than 100 bits, so that SageMath is not even using NFS to factor these small numbers. Instead, we refer to the literature to find that the best classical factoring algorithm (the general number field sieve [35]) runs in $L_N[1/3, (64/9)^{1/3}]$ and this was indeed used to factor a 768-bit RSA modulus in approximately two-thousand core-years [33], with one core being a 2.2 GHz AMD Opteron.

The timing of factoring using trial division is shown in Figure 6b. The results reveal an exponential trend and with a much smaller constant than the SAT solver. On this small scale on which measurements were performed, trial division easily outperforms the SAT solvers. The asymptotic runtime of the methods are so close together that we cannot meaningfully extrapolate the results to find a cross-over point where the SAT solvers become faster than trial division. We therefore cannot rule out that factoring with classical SAT solvers is always slower than trial division.

4 Quantum Solvers

State of the art classical factoring algorithms have super-polynomial runtime $L_N[1/3, (64/9)^{1/3}]$ [35], whereas Shor's algorithm [52] runs in polynomial time. This algorithm requires a fault-tolerant quantum computer and no scalable version has been implemented yet. Shor's algorithm has profound practical implications for currently deployed public-key cryptography such as RSA and the timing of the factoring of 1024-bit, 2048-bit or even larger semi-primes is of great practical significance for both contemporary and future security systems [43]. Mitigations for future systems and current systems requiring long-term security are being researched by the field of post-quantum cryptography [9, 13, 14].

An interesting notion of quantum computing has been proposed by Farhi et al. [26] in the form of adiabatic quantum computers. It was suggested that adiabatic quantum algorithms may be able to outperform classical computers on hard instances of NP-complete problems [25]. Since then, adiabatic quantum computation (a generalization of the adiabatic optimization explored deeply by Farhi et al.) has been proven to be polynomially equivalent to quantum computation in the standard gate model [3]. While the possibility of super-polynomial (or even just super-quadratic) quantum speed-up for NP-

hard problems remains an open question[6] it is generally believed that quantum computers (including adiabatic quantum computers) are not able to efficiently solve NP-hard problems such as SAT. Note that this assumption is implicit, e.g. in the fact that post-quantum cryptographers are working on the assumption that symmetric algorithms like AES and SHA that offers n bits of security against classical attacks offer $n/2$ bits of security against the best known quantum attacks [13].[7] In this section we consider the speedup that can be achieved by reducing the problem of factoring a semi-prime to an instance of an NP-hard problem which is then solved with a quantum computer.

When considering the runtime T of an algorithm we are most interested in the runtime as a function of the input size. In order to determine if one solver is faster than the other, we should always consider the *total* runtime. In the above analysis this is what we did by measuring the total runtime of the SAT solver *including* the runtime of the preprocessor.[8] For many solvers the total runtime can be naturally partitioned into the time spent in pre-/post-processing (T_p) and the time spent solving (T_s):

$$T(n) = T_p(n) + T_s(n), \tag{1}$$

where n is the input size of the problem.

Examples of this partitioning occur with the SAT preprocessor (T_p) and the SAT solver (T_s), the compiling (T_p) and running (T_s) of Shor's algorithm or the creation of a Hamiltonian (T_p) and the execution of the adiabatic algorithm (T_s).

In order to properly analyze the runtime of any algorithm we need to consider $T(n)$ and not just $T_s(n)$, since an unbounded amount of preprocessing can find a solution and render $T_s(n)$ to be trivial. We should also take care to set n to be the input size of the problem. Concretely this means we should let n be the size of the semi-prime and not the number of variables or clauses in our SAT instance. It is also important to analyse instance sizes larger than some lower bound ($n \geq n_0$), as the asymptotic behaviour is not visible for smaller sizes. For example the asymptotics of the MapleCOMSPS solver on integer factorization only become apparent at $n_0 = 20$ bit semi-primes.

4.1 Faster SAT solvers

One might hope that we can apply a quantum strategy that can improve on the best known classical methods. We chose SAT solvers to represent the best classical methods as their implementations are the highly optimized result of years of research. Generic quantum searching methods can achieve at most a quadratic speed-up, and we are aware of no convincing evidence that more than a quadratic speed-up can be achieved by quantum SAT-solving methods. For example, many modern SAT solvers rely on machine-learning techniques [38] and many quantum methods with a quadratic speedup are known for a variety of machine-learning algorithms [10]. See also [1] for why the exponential speedup promised in some quantum machine-learning literature is unlikely to be achieved in real-world implementations.

A quick calculation shows that even with a quadratic speedup, this strategy is not a very efficient one. We set an upper bound on the number of operations required for the classical solver based on our results. Accounting for any internal parallelism in the processor (four arithmetic ports per processor) and assuming that the CPU was fully occupied at every clock cycle this means that 10^{10} operations were being executed every second during the solving time.

Under this assumption the expected number of operations required for classical SAT solving becomes $2^{16.8} \cdot 2^{0.495n}$. With a quantum computer we might hope to reduce this to $\sqrt{2^{16.8} \cdot 2^{0.495n}} = 2^{8.41} \cdot 2^{0.247n}$ operations. To put this in perspective, consider a quantum computer that can execute 10^{40} *quantum* operations per second. Note that even a classical computer with such speeds could

[6] it is known that any such speed-up must go beyond pure "black-box" search [7] as attempted by Farhi et al. [25] and must somehow exploit additional knowledge or structure [58]

[7] Excluding some specific attacks in the "quantum superposition" attack model [30].

[8] To be even more precise we should also have included the time it took to generate the SAT instances. This generation is done in polynomial time and the runtime is negligible compared to the solver time, therefore we omitted this from our benchmarks.

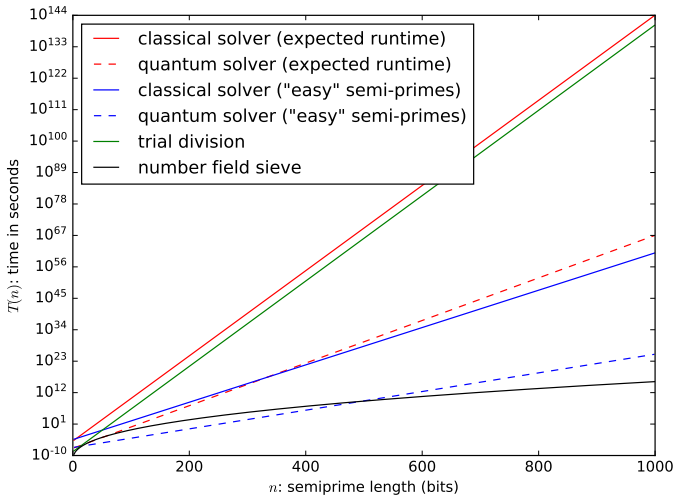

Figure 7: Comparison of efficiency of various factoring methods. The classical results are extrapolated from experimental data. The quantum results apply a quadratic speedup over the full classical computation. The number field sieve result plots $L_N[1/3, (64/9)^{1/3}]$ operations assuming the same number of operations per second.

break AES-128, SHA-256, RSA-2048 and ECC P-256 in an instant, so this is a very generous upper bound. Under these assumptions it would still take approximately a hundred times the lifetime of the universe to factor the RSA-768 number using the quantum SAT solving approach, whereas this number has been factored classically on a real machine in two-thousand core-years using number-theoretical methods. A visual comparison of these results are given in Figure 7.

Note that all estimates so far are biased towards more (classical) operations per second. The reason is that we want to compute an upper bound on the speedup that can be achieved by applying Grover's algorithm (or some alternative quadratic speedup) in order to factor numbers with SAT solving. It is almost certain that the processor executed less operations and it is very unlikely that the quadratic speedup can be applied to the full computation without any overhead of executing the algorithm. Therefore, classical algorithms will likely require less operations than reported and quantum algorithms will likely require more operations than computed.

Note also that the known speedups for quantum solvers are applied to T_s, even though the above calculation generously assumes that $T_p(n) = 0$ and the speedup can be applied to the full calculation time $T(n)$. For most classical solvers it indeed holds that $T_p(n) \ll T_s(n)$ as n grows large enough, but for many of the adiabatic factoring methods discussed below it holds that $T_p(n) \gg T_s(n)$ as n grows. This means that our calculation is a significant overestimation of the maximum speedup that can be achieved with the adiabatic factoring method.

4.2 Adiabatic factoring

The method used for factoring with the adiabatic algorithm first reduces factorization to finding the roots in a set of integer equations in which the unknown variables are restricted to binary values, corresponding to the input bits of the prime and carry bits of the intermediate computation. This is translated to the pseudo-Boolean optimization problem by squaring all equations (so that the roots correspond to the minimum values) and summing over all equations. This reduction was first suggested by Burges as a method for generating unconstrained optimization problems whose complexity can be easily controlled [12]. The adiabatic algorithm [26] is particularly well suited for encoding optimization problems of this kind: the resulting sum describes a Hamiltonian of which the ground state encodes

the solution and every variable corresponds to a single qubit. In general it is not easy to physically initialize the system in the ground state of the Hamiltonian, so instead an easier Hamiltonian encodes the initial state which is easy to initialize in the ground state. The adiabatic theorem tells us that if we evolve the physical system from the initial Hamiltonian to the final Hamiltonian slow enough, the system will remain in its ground state. Measuring the final state will then provide the answer to the optimization problem.

To assess the power of the adiabatic algorithm it is therefore important to quantify how fast this evolution can be done. A coarse lower bound is given by the spectral width of the time-dependent Hamiltonian, but sharper bounds on the runtime so far elude us [58]. This has led some researchers to study the applicability of the adiabatic algorithm to some NP-complete problems [25]. Most evidence for a speed-up is based on noise-free simulations on small instances (for which the asymptotic behaviour might not be visible) which are chosen randomly, shedding light on typical performance for small instances. Cryptographic problems require average-case hardness in order to be practical, which is why they are so suitable for testing worst-case behaviour of algorithms that solve them, especially when the reduction to an NP-hard problem is as simple as reducing factoring to SAT as demonstrated in the previous section.

Pseudo-Boolean optimization is known to be NP-hard, meaning amongst other things that a polynomial reduction exists from the SAT problem. The objective function for factorization instances using the above method is a quartic polynomial. Real-world demonstrations of the adiabatic algorithm suffer from additional limitations (besides noise-resistance) in the number of available qubits and multi-qubit interactions. The latter limitation means that quartic terms in the objective function cannot always be realized. Using quadratization [50] each objective function can be simplified to a quadratic polynomial at the price of additional variables, giving an instance of the well-studied quadratic unconstrained binary optimization (QUBO)[9] problem. This simplification runs in polynomial time and results in only polynomially many variables overhead, so the problems are equivalent.

However for many real-world systems the extra variables (qubits) are not available, so additional simplifications are required. This is fine as long as these simplification steps do not dominate the overall runtime of the program. More precisely we can execute polynomially many simplification operations and $T_p(n)$ will remain polynomial in n, thereby not significantly increasing the runtime $T(n)$ which is dominated by the super-polynomial runtime $T_s(n)$. When the simplification process is allowed to have an exponential runtime it can absorb the hardness of the problem, leaving a weaker problem to be solved (trivially) in polynomial time.

4.2.1 Implementations

The first adiabatic factorization [47] was implemented in 2008 using nuclear magnetic resonance (NMR) to factor 21 using three qubits. The authors fit a quadratic curve against a theoretical approximation in a noiseless model, they measure the runtime as a function of the number of input qubits (not the size of the factored number) and they only consider the small domain of seven to sixteen input qubits. It is doubtful that such small instances are a good indicator of polynomial asymptotic behaviour.

Later work [63] translates the problem of factoring 143 into a pseudo-binary optimization instance, which is an NP-hard problem [11]. The authors manage this by introducing the additional assumption that both factors must be of equal bitlength with the most significant bit set to one. Combining these assumption with some simplifications in the pseudo-Boolean equations simplifies the problem so that it only concerns four input bits of the prime factors. Although the used simplifications are efficient, only an upper bound of their effectiveness is given.

Subsequent research [17] observes that a minor generalization of the previous method reduces the problem to four input qubits whenever the two primes composing the semi-prime differ only in two positions, which likely occurs for infinitely many semi-primes [48]. This provides some evidence that

[9] also known as unconstrained binary quadratic programming (UBQP)

the simplifications do not generalize and the factored number 143 was identified as a particularly easy number to factor. In other words, this example was hand-picked from an exponentially unlikely family of semi-primes that are by design easy to factor. The authors report the number 56153 as being the largest semi-prime factored quantumly and at the same time argue that the work has factored an arbitrarily large set of semi-primes (since they can be pre-processed into solving the same pseudo-Boolean equations). The reason for not reporting a bigger number appears to be the large runtime T_p of the simplification process.

Much subsequent research in the adiabatic factoring field has focussed on methods such as deduc-reduc [57], split-reduc [46] and energy landscape manipulation [56], all of which can be seen as improvements on the preprocessor runtime T_p and none of which do any improvements on T_s.

The problems with viewing these works as relevant quantum integer factorization benchmarks is highlighted even further in the more recent paper that claims to have factored 291311 with adiabatic quantum computation [37]. The authors take the above approach and reduce the problem of factoring 291311 to the integer equations

$$q_1 + q_2 - 2q_1q_2 = 1 \tag{2}$$
$$q_2 + q_5 - 2q_2q_5 = 0 \tag{3}$$
$$q_1 + q_5 - 2q_1q_5 = 1, \tag{4}$$

where the variables q_i must take on binary values and represent unknown bits in the binary representation of factor $q = 1000q_501q_2q_11$. The authors stop their simplification process at this point and fail to notice that the above equations can be further simplified to

$$q_1 = 1 - q_2 = 1 - q_5. \tag{5}$$

Both solutions $q_1 = 0$ and $q_1 = 1$ correspond with respective factors $q = 557$ and $q = 523$. In other words, the number was already factored by the simplification process and the adiabatic quantum computation was a complicated way of flipping a coin and deciding between the two factors. The above criticism of these claims to meaningful quantum factoring benchmarks was in fact already made in 2013 [53].

A method called Variational Quantum Factoring (VQF) [5] employs the same strategy for factoring, which is to reduce it to an NP-hard optimization problem. The authors are careful to ensure that preprocessing takes only polynomial time. Although the authors claim that "VQF could be competitive with Shor's algorithm even in the regime of fault-tolerant quantum computation", we find no convincing argument to support this conjecture. In particular, they do not provide convincing evidence that the solving step is efficient: no semi-prime larger than 2^{15} is considered by their work and they observe that "the mere presence of carry bits negatively affects the algorithm".

The criticism from [53] applies equally well against "compiled versions" of Shor's algorithm: both implementations require much precomputation and therefore do not scale to factoring larger numbers. The problem is that both precomputations require prior knowledge of the solution. "Compiled versions" of Shor's algorithm were never intended to scale to meaningful input sizes, as is highlighted in the abstract of the work factoring 15 with NMR: "scalability is not implied by the present work. The significance of our work lies in the demonstration of experimental and theoretical techniques" [59].

The important difference is that the runtime of Shor's algorithm is well understood and provides a super-polynomial speedup in T_s over even the best numerical methods for factoring. As fault-tolerant hardware emerges, we can simply strip away the non-scalable optimizations. On the other hand the runtime of reducing factoring to an NP-hard problem and then solving it with (quantum) solvers is not understood very well, but all evidence points in the direction that it cannot even compete with classical numerical methods for factoring.

4.3 D-Wave

The D-Wave systems work by a process called quantum annealing, which can be viewed as a noisy version of adiabatic quantum computing. It has been shown that $O(n^2)$ qubits suffice to encode factoring into a quantum annealing instance with local interactions [39]. The article "Boosting integer factoring performance via quantum annealing offsets" [4] describes a "boost" when comparing factoring on the D-Wave machine with annealing offsets against the D-Wave machine without annealing offsets. The largest factored number has 20 bits.

All semi-primes up to 200000 (18 bits) have been factored with help of the D-Wave 2X by heuristically mapping the optimization problem to the Chimera graph underlying the machine [22]. Exponential methods from computational algebraic geometry are used for preprocessing the instances without quantification of the (asymptotic or measured) runtime so that there is no indication of the efficiency of this preprocessing step. Although some statistics on the annealing process are provided for six semi-primes, not enough information is given for a meaningful assessment on the scalability of both the efficiency and effectiveness of this method.

Integer factorization has been implemented on the D-Wave 2000Q by a similar strategy [28]. Quantified experimental results are only provided for factoring 15 and 21. As the authors note, there is no evidence that quantum annealing will find factors with significant likelihood in polynomial (or even sub-exponential) time.

5 Conclusions

SAT solvers are not known or believed to be able to factor semi-primes efficiently. Overall, even the fastest solver (MapleCOMSPS) has an exponential runtime in the size of the factors. Closer inspection of the solver runtime indicates that the solver is not able to detect any pattern in the SAT formulas that encode the factorization problem. Asymptotically the solver runtime appears to be comparable to that of trial division, but this advantage is almost completely negated by the overhead in the constant term. The performance of SAT solvers does not even come close to that of number-theoretical methods.

Quantum SAT solvers are not expected to do much better. Even when calculating a very optimistic speedup to the current state-of-the-art classical solvers, these solvers are outperformed with orders of magnitude by (classical) number-theoretical factoring methods. This approach to factoring reduces factoring, a problem with an $L_N[1/3, (64/9)^{1/3}]$ algorithm, to an NP-hard problem and then running (classical or quantum) solvers that have exponential runtime in the worst-case. At the surface, this obviously does not sound like a promising idea, as the quantum SAT solver must make up the exponential ground lost by translating the problem with subexponential algorithms to one where the best known algorithms are exponential. One might hope that good SAT solving heuristics for solving SAT on random or average-case instances could nevertheless have a practical impact on integer factorization, but there is no evidence of this. Of course if it were that easy RSA would be broken regularly by SAT solvers which does not appear to be the case. Furthermore, in practice it appears that SAT instances derived from integer factorization instances are hard SAT instances. Thus it would be especially surprising if a SAT solver of any kind (quantum or classical) could solve these instances with resources comparable to that of using the classical number field sieve (i.e. subexponential complexity). Our work explores this possibility more deeply and reinforces the folklore that reducing multiplication to SAT and then applying SAT solvers, classical or quantum, is not useful for factoring numbers of sizes relevant to cryptography.

A more promising approach is to try to speed up the solution to some subroutine of the NFS, as is done in [8]. In particular, one could reduce some carefully chosen sub-problem solved within the number field sieve to SAT. The sub-problem should be chosen so that classically solving the SAT instance is roughly as costly as the usual approach to solving the sub-problem. In this case, any quantum speed-up for solving these SAT instances would lead to a faster implementation of the

number field sieve. This approach is explored in [44].

Of course, one cannot rule out unexpected breakthroughs in quantum SAT solving or a wide range of other quantum or classical approaches to factoring semi-primes. However, it is important to distinguish the possibility of unexpected breakthroughs (especially those that contradict conventional wisdom or lack a plausible roadmap) from tracking progress of an existing hardware platform and of an algorithm that is pertinent for cryptographically relevant semi-primes (i.e. classical computers and the NFS).

Once scalable fault-tolerant quantum computers capable of implementing Shor's algorithm are available, a similar tracking would be very meaningful (with the caveat outlined in the introduction). In the meantime, it is important to track progress toward achieving scalable fault-tolerant quantum computers.

In other words, notwithstanding other scientific merits of these works, we are not aware of any evidence that any SAT-based quantum factoring results to date, including factorization by quantum annealing, are relevant milestones toward large-scale integer factorization or the demonstration of a speed-up over the best known classical algorithms for integer factorization.

Acknowledgment

We would like to thank Vijay Ganesh and Curtis Bright for the many lessons about modern SAT solving and insightful discussions regarding this project. We also would like to thank Colin P. Williams and Kenneth Paterson for their helpful comments.

References

[1] Scott Aaronson. Read the fine print. *Nature Physics*, 11:291–293, 4 2015. doi:10.1038/nphys3272.

[2] Michel Abdalla, Tor Erling Bjørstad, Carlos Cid, Benedikt Gierlichs, Andreas Hülsing, Atul Luykx, Kenneth G. Paterson, Bart Preneel, Ahmad-Reza Sadeghi, Terence Spies, Martijn Stam, Michael Ward, Bogdan Warinschi, and Gaven Watson. Algorithms, key size and protocols report. Technical report, University of Bristol, 2 2018. URL: http://www.ecrypt.eu.org/csa/documents/D5.4-FinalAlgKeySizeProt.pdf.

[3] Dorit Aharonov, Wim van Dam, Julia Kempe, Zeph Landau, Seth Lloyd, and Oded Regev. Adiabatic quantum computation is equivalent to standard quantum computation. In *45th Annual IEEE Symposium on Foundations of Computer Science*, pages 42–51, Rome, Italy, 10 2004. IEEE Computer Society. doi:10.1109/FOCS.2004.8.

[4] Evgeny Andriyash, Zhengbing Bian, Fabian Chudak, Marshall Drew-Brook, Andrew D. King, William G. Macready, and Aidan Roy. Boosting integer factoring performance via quantum annealing offsets. Technical report, D-Wave Systems Inc., 12 2016. URL: https://www.dwavesys.com/sites/default/files/14-1002A_B_tr_Boosting_integer_factorization_via_quantum_annealing_offsets.pdf.

[5] Eric R. Anschuetz, Jonathan P. Olson, Alán Aspuru-Guzik, and Yudong Cao. Variational quantum factoring. *CoRR*, abs/1808.08927, 2018. URL: https://arxiv.org/abs/1808.08927.

[6] Tomáš Balyo, Marijn J. H. Heule, and Matti Järvisalo, editors. *Proceedings of SAT Competition 2017: Solver and Benchmark Descriptions*, Publication series B, Report B-2017-1, 9 2017. URL: http://hdl.handle.net/10138/224324.

[7] Charles H. Bennett, Ethan Bernstein, Gilles Brassard, and Umesh Vazirani. Strengths and weaknesses of quantum computing. *SIAM Journal on Computing*, 26(5):1510–1523, 1997. doi:10.1137/S0097539796300933.

[8] Daniel J. Bernstein, Jean-François Biasse, and Michele Mosca. A low-resource quantum factoring algorithm. In Tanja Lange and Tsuyoshi Takagi, editors, *Post-Quantum Cryptography*, pages 330–346, Cham, 2017. Springer International Publishing. doi:10.1007/978-3-319-59879-6_19.

[9] Daniel J. Bernstein, Johannes Buchmann, and Erik Dahmen, editors. *Post-quantum cryptography*. Springer-Verlag, Berlin, Heidelberg, 2009. doi:10.1007/978-3-540-88702-7.

[10] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. Quantum machine learning. *Nature*, 549(7671):195, 2017. doi:10.1038/nature23474.

[11] Endre Boros and Peter L. Hammer. Pseudo-boolean optimization. *Discrete Applied Mathematics*, 123(1):155 – 225, 2002. doi:10.1016/S0166-218X(01)00341-9.

[12] Christopher J. C. Burges. Factoring as optimization. Technical report, Microsoft, 8 2002. MSR-TR-2002-83. URL: https://www.microsoft.com/en-us/research/publication/factoring-as-optimization/.

[13] Lily Chen, Stephen Jordan, Yi-Kai Liu, Dustin Moody, Rene Peralta, Ray Perlner, and Daniel Smith-Tone. Report on post-quantum cryptography, 4 2016. doi:10.6028/NIST.IR.8105.

[14] Lily Chen, Dustin Moody, and Yi-Kay Liu. Post-quantum cryptography, 2018. URL: https://csrc.nist.gov/Projects/Post-Quantum-Cryptography.

[15] Aaron Clauset, Mark E. J. Newman, and Cristopher Moore. Finding community structure in very large networks. *Phys. Rev. E*, 70:066111, 12 2004. doi:10.1103/PhysRevE.70.066111.

[16] Stephen A. Cook. The complexity of theorem-proving procedures. In *Proceedings of the Third Annual ACM Symposium on Theory of Computing*, STOC '71, pages 151–158. ACM, 1971. doi:10.1145/800157.805047.

[17] Nikesh S. Dattani and Nathaniel Bryans. Quantum factorization of 56153 with only 4 qubits. *CoRR*, abs/1411.6758, 2014. URL: https://arxiv.org/abs/1411.6758.

[18] Martin Davis, George Logemann, and Donald Loveland. A machine program for theorem-proving. *Commun. ACM*, 5(7):394–397, 7 1962. doi:10.1145/368273.368557.

[19] Martin Davis and Hilary Putnam. A computing procedure for quantification theory. *J. ACM*, 7(3):201–215, 7 1960. doi:10.1145/321033.321034.

[20] The Sage Developers. *SageMath, the Sage Mathematics Software System (Version 7.5.1)*. SageMath, 2017. URL: http://www.sagemath.org.

[21] Brandon Dixon and Arjen K. Lenstra. *Factoring Integers Using SIMD Sieves*, pages 28–39. Springer Berlin Heidelberg, Berlin, Heidelberg, 1994. doi:10.1007/3-540-48285-7_3.

[22] Raouf Dridi and Hedayat Alghassi. Prime factorization using quantum annealing and computational algebraic geometry. *Scientific Reports*, 7:43048, 2 2017. doi:10.1038/srep43048.

[23] Ashutosh Dhar Dwivedi, Milos Kloucek, Pawel Morawiecki, Ivica Nikolic, Josef Pieprzyk, and Sebastian Wójtowicz. SAT-based cryptanalysis of authenticated ciphers from the CAESAR competition. *IACR Cryptology ePrint Archive*, 2016, 11 2016. URL: https://eprint.iacr.org/2016/1053.

[24] Joseph Emerson, Robert Alicki, and Karol Życzkowski. Scalable noise estimation with random unitary operators. *Journal of Optics B: Quantum and Semiclassical Optics*, 7(10):S347, 2005. doi:10.1088/1464-4266/7/10/021.

[25] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, Joshua Lapan, Andrew Lundgren, and Daniel Preda. A quantum adiabatic evolution algorithm applied to random instances of an NP-complete problem. *Science*, 292(5516):472–475, 2001. doi:10.1126/science.1057726.

[26] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Michael Sipser. Quantum computation by adiabatic evolution. *CoRR*, abs/quant-ph/0001106, 2000. URL: https://arxiv.org/abs/quant-ph/0001106.

[27] Marijn J. H. Heule, Matti Järvisalo, and Tomáš Balyo. SAT competition, 2016. Affiliated with the 19th International Conference on Theory and Applications of Satisfiability Testing. URL: https://baldur.iti.kit.edu/sat-competition-2016/index.php.

[28] Shuxian Jiang, Keith A Britt, Alexander J McCaskey, Travis S Humble, and Sabre Kais. Quantum annealing for prime factorization. *CoRR*, abs/1804.02733, 2018. URL: https://arxiv.org/abs/1804.02733.

[29] Burt Kaliski. RSA factoring challenge. http://groups.google.com/groups?selm=BURT.91Mar18092126%40chirality.rsa.com, 3 1991.

[30] Mark Kaplan, Gaëtan Leurent, Anthony Leverrier, and María Naya-Plasencia. Breaking symmetric cryptosystems using quantum period finding. In Matthew Robshaw and Jonathan Katz, editors, *Advances in Cryptology – CRYPTO 2016*, volume 9815, pages 207–237, Berlin, Heidelberg, 2016. Springer. doi:10.1007/978-3-662-53008-5_8.

[31] Anatoly A. Karatsuba and Y. Ofman. Multiplication of multidigit numbers on automata. In *Soviet Physics Doklady*, volume 7, pages 595–596, 1963.

[32] Richard M. Karp. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, The IBM Research Symposia Series, pages 85–103, Boston, MA, 1972. Springer US. doi:10.1007/978-1-4684-2001-2_9.

[33] Thorsten Kleinjung, Kazumaro Aoki, Jens Franke, Arjen K. Lenstra, Emmanuel Thomé, Joppe W. Bos, Pierrick Gaudry, Alexander Kruppa, Peter L. Montgomery, Dag Arne Osvik, Herman te Riele, Andrey Timofeev, and Paul Zimmermann. Factorization of a 768-bit RSA modulus. In Tal Rabin, editor, *Advances in Cryptology – CRYPTO 2010*, pages 333–350, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg. doi:10.1007/978-3-642-14623-7_18.

[34] Arjen K. Lenstra. Key lengths. In Hossein Bidgoli, editor, *Handbook of Information Security*, chapter 114. Wiley, 6 2004. URL: https://infoscience.epfl.ch/record/164539/files/NPDF-32.pdf.

[35] Arjen K. Lenstra, Hendrik W. Lenstra, Mark S. Manasse, and John M. Pollard. *The number field sieve*, pages 11–42. Springer Berlin Heidelberg, Berlin, Heidelberg, 1993. doi:10.1007/BFb0091537.

[36] Leonid A. Levin. Universal Sequential Search Problems. *Problemy Peredachi Informatsii*, 9(3):115–116, 1973. URL: http://mi.mathnet.ru/eng/ppi914.

[37] Zhaokai Li, Nikesh S. Dattani, Xi Chen, Xiaomei Liu, Hengyan Wang, Richard Tanburn, Hongwei Chen, Xinhua Peng, and Jiangfeng Du. High-fidelity adiabatic quantum computation using the intrinsic hamiltonian of a spin system: Application to the experimental factorization of 291311. *CoRR*, abs/1706.08061, 2017. URL: https://arxiv.org/abs/1706.08061.

[38] Jia Hui Liang, Vijay Ganesh, Pascal Poupart, and Krzysztof Czarnecki. Learning rate based branching heuristic for SAT solvers. In Nadia Creignou and Daniel Le Berre, editors, *Theory and Applications of Satisfiability Testing – SAT 2016: 19th International Conference, Bordeaux, France, July 5-8, 2016, Proceedings*, pages 123–140, Cham, 2016. Springer International Publishing. doi:10.1007/978-3-319-40970-2_9.

[39] William G. Macready, Geordie Rose, and Peter Love. Quantum processor-based systems, methods and apparatus for solving problems as logic circuits, 10 2013. Patent No. US 8,560,282 B2, Filed August 3, 2010, Issued October 15, 2013. URL: https://patents.google.com/patent/US8560282B2/.

[40] Fabio Massacci and Laura Marraro. Logical cryptanalysis as a SAT problem. *J. Autom. Reasoning*, 24(1/2):165–203, 2000. doi:10.1023/A:1006326723002.

[41] Ilya Mironov and Lintao Zhang. Applications of SAT solvers to cryptanalysis of hash functions. In Armin Biere and Carla P. Gomes, editors, *Theory and Applications of Satisfiability Testing - SAT 2006*, pages 102–115, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg. doi:10.1007/11814948_13.

[42] Pawel Morawiecki and Marian Srebrny. A SAT-based preimage analysis of reduced keccak hash functions. *Inf. Process. Lett.*, 113(10-11):392–397, 2013. doi:10.1016/j.ipl.2013.03.004.

[43] Michele Mosca. Setting the scene for the ETSI Quantum-safe Cryptography Workshop. e-proceedings of "1st Quantum-Safe-Crypto Workshop", Sophia Antipolis, 9 2013.

[44] Michele Mosca, João Marcos Vensi Basso, and Sebastian R. Verschoor. Speeding up factoring with quantum SAT solvers. *CoRR*, abs/1910.09592, 2019. URL: https://arxiv.org/abs/1910.09592.

[45] Zack Newsham, Vijay Ganesh, Sebastian Fischmeister, Gilles Audemard, and Laurent Simon. Impact of community structure on SAT solver performance. In Carsten Sinz and Uwe Egly, editors, *Theory and Applications of Satisfiability Testing - SAT 2014 - 17th International Conference*, volume 8561 of *Lecture Notes in Computer Science*, pages 252–268, Vienna, Austria, 2014. Springer. doi:10.1007/978-3-319-09284-3_20.

[46] Emile Okada, Richard Tanburn, and Nikesh S. Dattani. Reducing multi-qubit interactions in adiabatic quantum computation without adding auxiliary qubits. part 2: The "split-reduc" method and its application to quantum determination of ramsey numbers. *CoRR*, abs/1508.07190, 2015. URL: https://arxiv.org/abs/1508.07190.

[47] Xinhua Peng, Zeyang Liao, Nanyang Xu, Gan Qin, Xianyi Zhou, Dieter Suter, and Jiangfeng Du. Quantum adiabatic algorithm for factorization and its experimental implementation. *Phys. Rev. Lett.*, 101:220405, 11 2008. doi:10.1103/PhysRevLett.101.220405.

[48] D. H. J. Polymath. Variants of the selberg sieve, and bounded intervals containing many primes. *Research in the Mathematical Sciences*, 1(1):12, 10 2014. doi:10.1186/s40687-014-0012-7.

[49] Ron L. Rivest, Adi Shamir, and Leonard Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM*, 21(2):120–126, 2 1978. doi:10.1145/359340.359342.

[50] Ivo G Rosenberg. Reduction of bivalent maximization to the quadratic case. *Cahiers du Centre d'etudes de Recherche Operationnelle*, 17:71–74, 1975.

[51] Bart Selman, Henry A. Kautz, and Bram Cohen. Local search strategies for satisfiability testing. *Cliques, coloring, and satisfiability*, 26:521–532, 1993.

[52] Peter W. Shor. Polynominal time algorithms for discrete logarithms and factoring on a quantum computer. In *ANTS*, volume 877 of *Lecture Notes in Computer Science*, page 289, Berlin, Heidelberg, 1994. Springer. doi:10.1007/3-540-58691-1_68.

[53] John A. Smolin, Graeme Smith, and Alexander Vargo. Oversimplifying quantum factoring. *Nature*, 499(7457):163–165, 2013. doi:10.1038/nature12290.

[54] Mate Soos. CryptoMiniSat 2.5.1. http://www.msoos.org/wordpress/wp-content/uploads/2010/08/cryptominisat-2.5.1.pdf, 8 2010.

[55] Mate Soos. CryptoMiniSat 5.0.1. https://github.com/msoos/cryptominisat/releases/tag/5.0.1, 9 2016.

[56] Richard Tanburn, Oliver Lunt, and Nikesh S. Dattani. Crushing runtimes in adiabatic quantum computation with energy landscape manipulation (ELM): application to quantum factoring. *CoRR*, abs/1510.07420, 2015. URL: https://arxiv.org/abs/1510.07420.

[57] Richard Tanburn, Emile Okada, and Nikesh S. Dattani. Reducing multi-qubit interactions in adiabatic quantum computation without adding auxiliary qubits. part 1: The "deduc-reduc" method and its application to quantum factorization of numbers. *CoRR*, abs/1508.04816, 2015. URL: https://arxiv.org/abs/1508.04816.

[58] Wim van Dam, Michele Mosca, and Umesh V. Vazirani. How powerful is adiabatic quantum computation? In *42nd Annual Symposium on Foundations of Computer Science, FOCS 2001*, pages 279–287, Las Vegas, Nevada, USA, 2001. IEEE Computer Society. doi:10.1109/SFCS.2001.959902.

[59] Lieven M. K. Vandersypen, Matthias Steffen, Gregory Breyta, Costantino S. Yannoni, Mark H. Sherwood, and Isaac L. Chuang. Experimental realization of Shor's quantum factoring algorithm using nuclear magnetic resonance. *Nature*, 414:883–887, 12 2001. doi:10.1038/414883a.

[60] Sebastian R. Verschoor. factoring-sat (GitHub repository), 1 2019. URL: https://github.com/sebastianv89/factoring-sat.

[61] Wikipedia. Records for efforts by quantum computers, 2018. URL: https://en.wikipedia.org/wiki/Integer_factorization_records#Records_for_efforts_by_quantum_computers.

[62] Ryan Williams, Carla P. Gomes, and Bart Selman. Backdoors to typical case complexity. In *IJCAI-03, Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence*, pages 1173–1178, Acapulco, Mexico, 2003. URL: http://ijcai.org/Proceedings/03/Papers/168.pdf.

[63] Nanyang Xu, Jing Zhu, Dawei Lu, Xianyi Zhou, Xinhua Peng, and Jiangfeng Du. Quantum factorization of 143 on a dipolar-coupling nuclear magnetic resonance system. *Phys. Rev. Lett.*, 108:130501, 3 2012. doi:10.1103/PhysRevLett.108.130501.

[64] Henry Yuen and Joseph Babel. Tough SAT Project. https://toughsat.appspot.com/, 2011.

A CryptoMiniSat 5

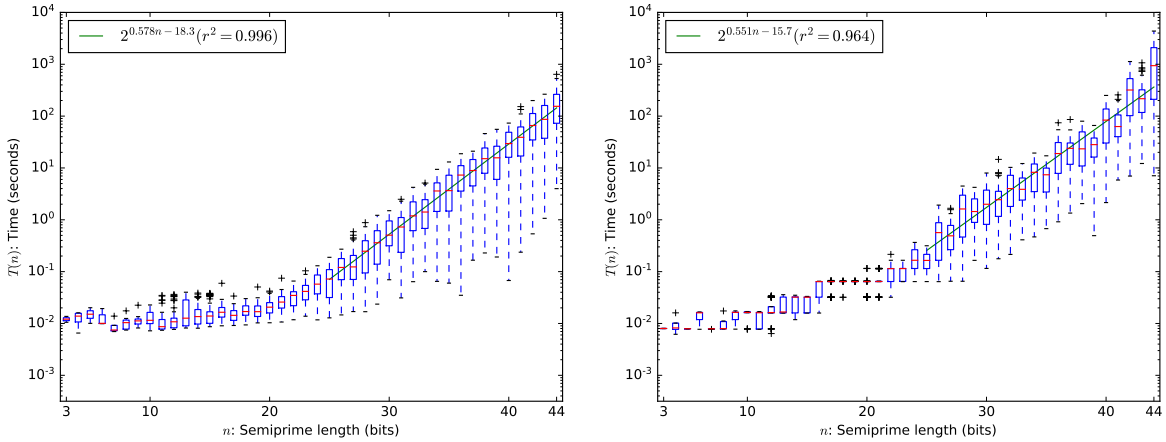

(a) schoolbook multiplication

(b) Karatsuba multiplication

Figure 8: Runtime of CryptoMiniSat 5 on factoring semi-primes. We measured 100 semi-primes per bitlength and applied no randomization.

Figure 8a and Figure 8b show the performance of factoring semi-primes with CryptoMiniSat 5. This solver solved each semi-prime SAT instance once, so no averaging has been applied to the shown results. In particular, one might be tempted to conclude from the longer whiskers in the depicted results that the CryptoMiniSat solver is lucky more often. However, the MapleCOMSPS solver gives similar results when only considering one solution. Closer inspection of the data reveals that CryptoMiniSat 5 is is outperformed consistently by MapleCOMSPS.

B Patterns

Figure 9 and Figure 10 examine the relation between various metrics on p, q and the solver time for long multiplication encoding and Karatsuba encoding (respectively). See also [60] for enlarged images. We examined bitwise patterns as these are most likely exploited by the SAT solver and we examined smoothness as this can determine the hardness of factoring for some number-theoretical methods.

Note that only the first two metrics ($\log_2 N$ and Hamming weight(N)) could potentially be used to predict how fast the solver will find a solution. The remaining metrics require knowledge of the value of p and q, but these metrics could be important for anyone generating primes in the RSA cryptosystem.

However, the lack of any correlation indicates that none of the investigated patterns have a significant impact on the solver time. In other words, SAT solvers do not influence the method by which a user of the RSA cryptosystem should generate primes.

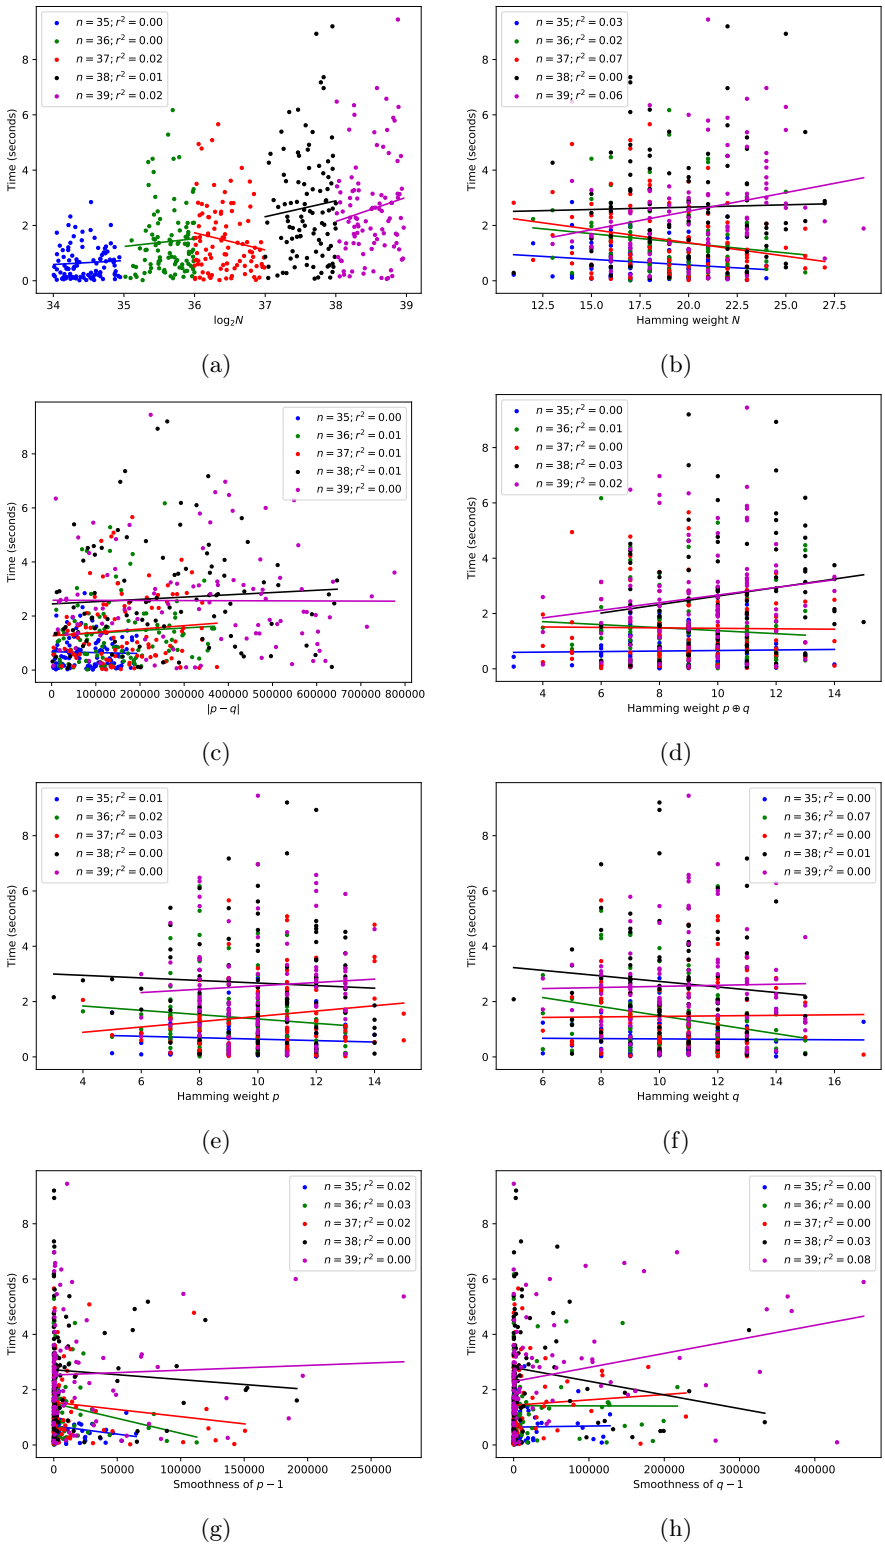

Figure 9: Solver time versus various patterns (schoolbook encoding).

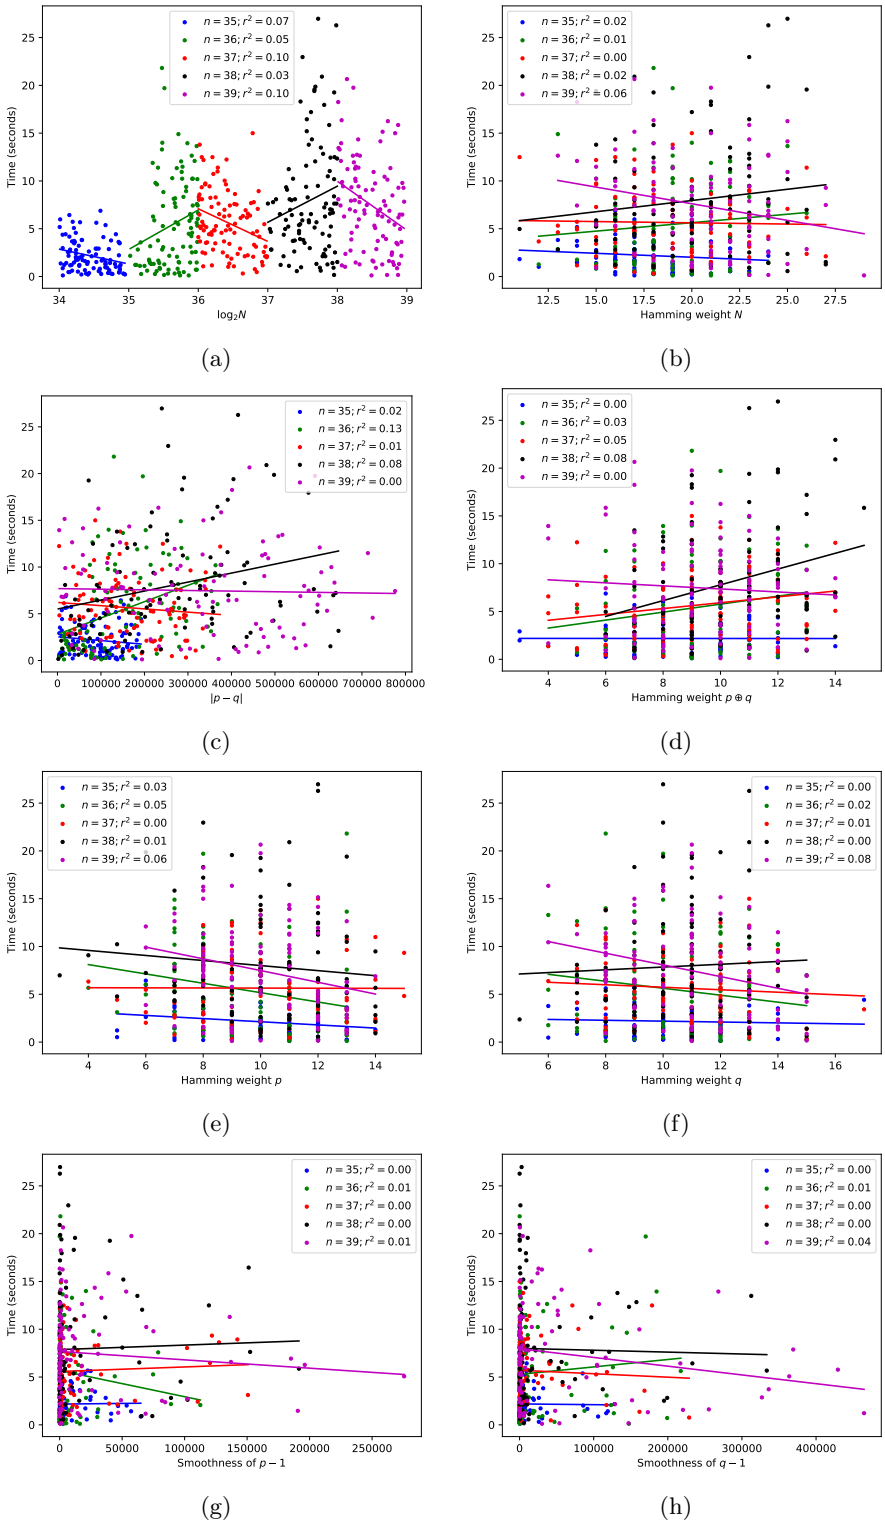

Figure 10: Solver time versus various patterns (Karatsuba encoding).