

# Thermodynamic Analysis of Classical and Quantum Algorithms for Preimage and Collision Search Problems

No Author Given

No Institute Given

**Abstract.** One of the earliest proposed applications of quantum computers was the quadratic speedup of classical brute force search represented by Grover's algorithm. Since then various successor algorithms have been applied to a variety of oracle problems including collision finding and claw finding. Such algorithms have typically been analyzed in terms of query complexity, which is a somewhat unphysical model of the cost of computation in the real world. Other models of these algorithms, such as the quantum circuit model and the quantum random access machine model of computation give more realistic, but sometimes conflicting answers regarding how much advantage the quantum algorithm provides over the corresponding classical algorithm.

We instead adapt Bennett's Brownian model of computation to directly estimate the cost of both classical and quantum operations in terms of time, energy, and memory. We apply this model to compare the best known quantum algorithms for collision search and preimage search to their classical counterparts. In the collision search case, our analysis agrees with previous analysis, based on the gate model of quantum computation, suggesting that quantum computation provides no improvement over the best known classical algorithm. More surprisingly, we find that a Brownian implementation of randomized classical search can achieve the same tradeoffs between time, memory and energy as Grover's algorithm (at least up to logarithmic factors.) This implementation uses thermal noise to drive a random walk within the internal state of a mostly unpowered circuit.

**Key words:** Reversible Computation, Quantum Computation, Collision Search, Preimage Search, Grover's algorithm

## 1 Introduction

## 2 Bennett's Brownian Computation Model

In contrast to ballistic models of computation (e.g. [1]), which are generally assumed to be unrealistic, Brownian computers [2, 3] are assumed to operate near thermal equilibrium at a finite temperature,  $T$ . As in the case of ballistic computers, to avoid being constrained to consume at least  $kT \ln 2$  energy per bit operation by the Landauer limit [4], the program for a Brownian computer is encoded as a reversible circuit. In the absence of any driving force, the state of the computing system at any given time may be described as a random walk on that circuit, with state transitions that undo a useful computation happening as often as those that perform it. In order for computations to proceed forward at a nonzero expected rate, a driving force, dissipating an energy of  $\epsilon$  per gate is imposed. This leads forward transitions within the circuit to occur  $e^{\frac{\epsilon}{kT}}$  times more often than backward transitions, resulting in a net forward computation rate proportional to  $\frac{\epsilon}{kT}$  for  $\epsilon$  small compared to  $kT$ .

Some additional costs are required in order for Brownian computation to be achieved fault-tolerantly. Energy barriers must be imposed to prevent transitions to physical states outside the reversible circuit, representing the prescribed computation path. In order to suppress the probability of such undesirable transitions so that a circuit of size  $G$  can be completed with high probability, the size of these energy barriers must at least be on the order of  $kT \ln(\frac{kT}{\epsilon} \cdot G)$ . Additionally, dissipating a “latching” energy of about  $kT \ln(\frac{kT}{\epsilon})$  during the computation’s final step is required to suppress backwards transitions once the computation has reached its halting state. These costs are described in detail in [3].

The above costs may, however, be assumed to be negligible in a number of important cases: In particular, the latching energy will be negligible when  $\frac{\epsilon}{kT}$  is at least logarithmically more than  $\frac{1}{G}$ . Additionally, establishing energy barriers to non-computational paths is likely to be a negligible cost when a description of the circuit can be expressed in a physically compact form, for example, using looping constructs. More precisely, if we assume that the circuit can be compressed into a program of size  $m_0$ , then the cost of imposing energy barriers should be on the order of  $m_0 \cdot kT \ln(\frac{kT}{\epsilon} \cdot G)$ . This cost is negligible as long as  $\frac{\epsilon}{kT}$  is significantly larger than  $e^{-\frac{G}{m_0}}$ .

In fact, the initialization cost may be less than this, since it may be more proper to think of the initialization process as rearranging the energy barriers already present in the available raw materials for constructing our computer. The cost is then determined by the Landauer limit and the information content of the circuit, including appropriately large energy barriers. Since the size of these barriers does not need to be precisely specified, but merely bounded above  $kT \ln(\frac{kT}{\epsilon} \cdot G)$ , the information content of the circuit may grow sublogarithmically with  $G$ . All we can say with confidence is that the information content of the circuit is at least  $m_0$ , and therefore the initialization energy is at least on the order of  $m_0 \cdot kT$ .

Thus far, we have only given asymptotic scalings for the relation between gate time and per-gate energy, assuming a fixed temperature. However, if we make the heuristic assumption that the rate at which gates are traversed due to Brownian motion is no more than  $\frac{h}{4kT}$ , following the Margolus-Levitin theorem [5], then we can specify a lower bound, independent of temperature, on the per-gate energy  $\epsilon$  required to perform  $G$  sequential operations in time  $t$ :

$$\epsilon > \frac{hG}{4t}.$$

Finally, it is worth commenting on the feasibility of Brownian computation for quantum computers. Brownian computation was originally proposed as a way to improve the thermodynamic efficiency of classical computation. It should be noted that many of the techniques that have been proposed for fault tolerance in quantum computation are thermodynamically irreversible, in particular, syndrome measurement and magic state preparation. These techniques cannot be used in a Brownian mode of computation. However, there are some proposed techniques, such as the use of Fibonacci anyons for universal quantum computation [?], that may be able to achieve fault tolerance without requiring significant thermodynamic irreversibility (although even in such cases, the cost of fault tolerance is believed to be polylogarithmic in the size of the circuit.) We will therefore optimistically assume that quantum operations can be implemented in a Brownian fashion.

We now proceed to analyze the asymptotic complexity of classical and quantum algorithms for collision and preimage search. We will generally ignore logarithmic factors. As we are engaged in asymptotic analysis, units are strictly speaking irrelevant,

but assuming natural units (e.g.  $c = \hbar = k = 1$ ) may be desirable to keep constant factors small.

### 3 Collision Search

The best known classical algorithm for finding collisions in a random function is the parallel collision search algorithm of Van Oorschot and Wiener [6]. If the range of the function is of order  $N$ , then, given  $M$  parallel processes each with memory  $O(1)$  the algorithm can find a collision in expected serial depth  $O(\frac{\sqrt{N}}{M})$ . The communication cost between threads is negligible compared to overall computational costs as long as  $M$  is smaller than  $\sqrt{N}$  by at least a logarithmic factor.

An improvement over classical collision search has been claimed by Brassard Høyer and Tapp (BHT [7]). Their algorithm is a serial process consisting of  $O(N^{\frac{1}{3}})$  operations and requires a memory of size  $N^{\frac{1}{3}}$ . This can be generalized to arbitrary memory size,  $M < O(N^{\frac{1}{3}})$ , giving a serial complexity of  $O(\sqrt{\frac{N}{M}})$ . The BHT algorithm may be further generalized to a parallel algorithm involving  $p$  parallel processors and a shared memory  $M$ , where  $p < M < O((Np)^{\frac{1}{3}})$ .<sup>1</sup> In this case, the serial complexity is  $O(\sqrt{\frac{N}{Mp}})$ .

Bernstein [8] has observed that the BHT algorithm, even if parallelized, does not improve upon the Van Oorschot - Wiener algorithm, when measured in terms of memory and serial depth. Since the BHT algorithm also requires  $O(\sqrt{\frac{N}{M}})$  random access queries to a memory of size  $M$ , each requiring  $O(M)$  gates, it also does not improve upon Van Oorschot Wiener algorithm when evaluated in terms of circuit size and depth (See Beals et al. [9] for a more thorough analysis.) However, BHT does represent an improvement over all classical algorithms in terms of query complexity. Furthermore, the Quantum RAM model of Giovanetti et al. [10] gives a theoretical argument that despite their large gate complexity, quantum memory access operations can be performed at logarithmic energy cost. A question therefore remains whether there exists a physically realistic model of computation where BHT is actually cheaper than the classical algorithms for the same problem. However, if there is such a model, it is not the Brownian model of computation, as we proceed to show:

We first analyze the quantum algorithm, calculating the total energy required to perform a collision search, given a maximum time limit  $t$  and a maximum memory size  $M$ . (Here, we assume, following the quantum RAM model, that the energy complexity of the BHT is dominated by oracle queries rather than memory access): The per operation energy  $\epsilon$  scales with the serial complexity divided by  $t$ , i.e.:

$$\epsilon_{quant} = O\left(\frac{\sqrt{\frac{N}{Mp}}}{t}\right). \quad (1)$$

The total energy  $E$  is then the product of the parallelism, the serial complexity, and the per operation energy, i.e.:

---

<sup>1</sup> Note this also implies that  $M < O(\sqrt{N})$ . The constraint arises from the requirement that the serial complexity,  $O(\frac{M}{p})$ , of filling a table of size  $M$  with oracle values does not exceed the serial complexity  $O(\sqrt{\frac{N}{Mp}})$  of Grover search.

$$E_{quant} = O\left(p \cdot \sqrt{\frac{N}{Mp}} \cdot \frac{\sqrt{\frac{N}{Mp}}}{t}\right) = O\left(\frac{N}{Mt}\right). \quad (2)$$

Now, we analyze the classical algorithm: The per operation energy again scales with the serial complexity, i.e.:

$$\epsilon_{cl} = O\left(\frac{\sqrt{N}}{Mt}\right). \quad (3)$$

The total energy  $E$  is again the product of the parallelism (In this case  $p = O(M)$ ), the serial complexity, and the per operation energy, i.e.:

$$E_{cl} = O\left(M \cdot \frac{\sqrt{N}}{M} \cdot \frac{\sqrt{N}}{Mt}\right) = O\left(\frac{N}{Mt}\right). \quad (4)$$

Thus, even under optimistic assumptions within the Brownian model of computation, we find that quantum computers provide no advantage in terms of energy, memory, or time, for solving the collision search problem.

## 4 Preimage Search

Grover's algorithm finds preimages in a function with domain size  $N$  in serial complexity  $O(\sqrt{N})$ . Grover's algorithm can be generalized to take advantage of  $M$  parallel processes each with memory  $O(1)$ , in which case the serial complexity is reduced to  $O\left(\sqrt{\frac{N}{M}}\right)$ . If we implement Grover's algorithm in a Brownian fashion, we find that

$$\epsilon_{quant} = O\left(\frac{\sqrt{\frac{N}{M}}}{t}\right), \quad (5)$$

and,

$$E_{quant} = O\left(M \cdot \sqrt{\frac{N}{M}} \cdot \frac{\sqrt{\frac{N}{M}}}{t}\right) = O\left(\frac{N}{t}\right). \quad (6)$$

A naïve Brownian implementation for classical search would divide the key space among  $M$  parallel processes, each of which would deterministically step through  $\frac{N}{M}$  keys searching for the correct one. Such a deterministic classical algorithm would require,

$$\epsilon_{det} = O\left(\frac{N}{Mt}\right),$$

and,

$$E_{det} = O\left(M \cdot \frac{N}{M} \cdot \frac{N}{Mt}\right) = O\left(\frac{N^2}{Mt}\right).$$

This already allows us to compete with Grover's algorithm if we allow ourselves a memory of size  $O(N)$ . However, we can exploit the structure, or rather the lack of structure, of the search problem to improve upon this figure. In particular, rather than deterministically stepping through the keys, dissipating a driving energy each time, we can simply allow Brownian motion to drive the system on a random walk through

the keyspace. We will still require a latching energy to end the computation, once the correct key has been found, and an initialization energy to create the necessary energy barriers to prevent unwanted transitions from occurring.

If the search is implemented by  $M$  parallel processes, each of size  $O(1)$ , then each process must reach  $\frac{N}{M}$  keys. This requires the processes to operate at a temperature:

$$kT = O\left(\frac{N}{Mt}\right). \quad (7)$$

The initialization energy should be of order  $MkT$  i.e.:

$$E_{init} = O\left(M \cdot \frac{N}{Mt}\right) = O\left(\frac{N}{t}\right).$$

This is identical to the energy required by a Brownian implementation of Grover's algorithm. All that remains is to show that the latching energy is negligible. Indeed, we find that the energy required to suppress backwards transitions from the final state for a time of order  $t$  is  $O(kT \ln(tkT)) = O\left(\frac{N}{Mt} \ln\left(\frac{N}{M}\right)\right)$ . This is negligible as long as  $M$  is at least logarithmic in  $N$ .

Thus, as with collision search, the quantum and classical preimage search algorithms appear to offer the same tradeoffs between time, energy, and memory:

$$E_{cl} = E_{quant} = O\left(\frac{N}{t}\right). \quad (8)$$

## 5 Preimage Search at constant Power and Temperature / Energy Scale

In contrast to the collision search case, matching the time/ memory/ energy tradeoffs of Grover's algorithm with a classical search requires a somewhat unrealistic assumption. We assume that if a computational process can be accomplished at a temperature  $T$  in a time  $t$ , then an isomorphic computation can also be accomplished at a temperature  $\alpha T$  in a time  $\frac{t}{\alpha}$ . This would be true if physics were scale invariant, but the physics of the real world is almost certainly not scale invariant. A more realistic model would therefore restrict the range of temperatures where a given computation is considered feasible. We will therefore repeat the analysis of the previous section assuming a fixed temperature  $T$ . For added realism, in addition to memory  $M$ , and time  $t$ , we will express the resources required for search in terms of power,  $P = \frac{E}{t}$ , rather than energy, since a fixed power budget is a more common limitation than a fixed energy budget.

From Equation 8 we find:

$$N = O(Pt^2).$$

Plugging this into Equation 7 gives us:

$$M = O\left(\frac{Pt}{T}\right).$$

We can now calculate time and memory requirements in terms of  $T$ ,  $P$ , and  $N$ :

$$t_{cl} = O\left(\sqrt{\frac{N}{P}}\right) \quad (9)$$

$$M_{cl} = O\left(\frac{\sqrt{NP}}{T}\right) \quad (10)$$

A similar analysis may be done in the quantum case. Here we use Equation 5 as a lower bound for  $T$ . If the per gate energy  $\epsilon$  exceeds  $kT$ , we enter the thermodynamic regime of irreversible computing, as opposed to Brownian computing, at which point the time per gate not only fails to further decrease with increasing  $\epsilon$ , but must in fact increase to prevent the waste heat from heating the computing system to a temperature higher than  $T$ . Combining this bound with equation 6 then yields the following time and memory requirements for Grover search at fixed power and temperature:

$$t_{quant} = O\left(\sqrt{\frac{N}{P}}\right) \quad (11)$$

$$M_{quant} = O\left(\frac{P}{T^2}\right) \quad (12)$$

Thus we find, fixing power and temperature, that our classical search strategy recovers the squareroot time scaling of Grover's algorithm. However, unlike Grover's algorithm, whose space requirement is determined only by the power budget and maximum operating temperature, the classical algorithm also requires memory that scales, like the time, with the squareroot of the size of the search space.

## 6 Factors associated with the cost of oracle queries

The asymptotic complexities given in previous sections ignore the computational complexity of individual oracle queries. Most of the results of previous sections remain substantively similar if these factors are included. We will model each oracle query as a circuit with depth  $d_0$ , width  $m_0$ , and total gates  $g_0$ .

In the case of powered Brownian computation, the effect of these factors is fairly straightforward. The memory imposed limit on parallelism (and number of table entries in the case of BHT) is now  $p_{max} = O\left(\frac{M}{m_0}\right)$ . Likewise, if  $t_0$  is the time per query required to complete the computation in time  $t$ , we will now require an energy per gate of  $\epsilon = O\left(\frac{d_0}{t_0}\right)$ . We must also ensure that all the bits or qubits in the circuit advance through it roughly synchronously. This can be done, for example, by associating a clock state of size  $O(\log(d_0))$  to each bit or qubit in the oracle circuit, and imposing a restoring potential proportional to the squared difference of the clock states of neighboring qubits. This will tend to couple the clock states of nearby qubits, but will not dissipate any net energy. As with other energy barriers ensuring correct computation, this potential need only extend logarithmically far from the equilibrium point, relative to the total size of the computation. We will generally ignore the logarithmic memory cost of the clock state and the logarithmic computational costs associated with creating interactions between the clock state, but in more detailed models, they may be subsumed into  $m_0$  and  $g_0$  respectively. Finally, we must take into account the number of gates required to perform an oracle query,  $g_0$ . Making these substitutions into equations 2, and 4 gives the following energy costs for quantum and classical collision search:

$$E_{quant} = O\left(p \cdot g_0 \sqrt{\frac{m_0 N}{Mp}} \cdot d_0 \sqrt{\frac{m_0 N}{Mp}}\right) = O\left(\frac{g_0 m_0 d_0 N}{Mt}\right); \quad (13)$$

$$E_{cl} = O\left(\frac{M}{m_0} \cdot g_0 \frac{m_0 \sqrt{N}}{M} \cdot \frac{m_0 d_0 \sqrt{N}}{Mt}\right) = O\left(\frac{g_0 m_0 d_0 N}{Mt}\right). \quad (14)$$

Again, we find the classical and quantum complexities to be identical. In both cases, the useful memory size is bounded above by  $O(m_0\sqrt{N})$ .

Similarly, we may make the same substitutions in equations 5 and 6 to include these factors in the per gate and total energy cost Grover search:

$$\epsilon_{quant} = O\left(\frac{d_0\sqrt{\frac{m_0N}{M}}}{t}\right); \quad (15)$$

$$E_{quant} = O\left(\frac{M}{m_0} \cdot g_0\sqrt{\frac{m_0N}{M}} \cdot \frac{d_0\sqrt{\frac{m_0N}{M}}}{t}\right) = O\left(\frac{g_0d_0N}{t}\right). \quad (16)$$

In the case of unpowered Brownian computation, we must calculate the temperature  $T$  required for random Brownian motion to power the traversal of an oracle circuit of depth  $d_0$  and containing  $g_0$  gates in time  $t_0$ . To do this, we create a random variable,  $x$  indicating the total number of gates that have been completed at a time  $t$ . We expect that  $x$  will obey the usual formula for Brownian motion,  $\langle x^2 \rangle = Dt$ , for some  $D$ , which will depend on  $T$ ,  $g_0$ , and  $d_0$ . We will then require  $Dt_0 = O(g_0^2)$ . It remains to determine the scaling of  $D$ : Note that at any given time, on average  $O\left(\frac{g_0}{d_0}\right)$  gates will be exposed to activation by thermal noise. (The remaining gates will be disallowed by the clock states associated with their input/output bits.) Each of these gates is expected to contribute  $O(Tdt)$  to  $d\langle x^2 \rangle$ . The coupling potential between neighboring clock states will also drive the activation of individual gates, but it should have no net effect on  $x$ , since every gate driven forward by the coupling potential will be counterbalanced by another gate driven backwards. Thus we find that  $D = O\left(\frac{Tg_0}{d_0}\right)$  and therefore  $T = O\left(\frac{g_0d_0}{t_0}\right)$ .

We may now apply this analysis to equations 7 and 8. Since, in order to complete a preimage search of size  $N$  in time,  $t$  with memory  $M$ , we need  $t_0 = \frac{m_0N}{Mt}$ , we find that:

$$T_{cl} = O\left(\frac{g_0m_0d_0N}{Mt}\right), \quad (17)$$

and,

$$E_{cl} = O\left(M \cdot \frac{g_0m_0d_0N}{Mt}\right) = O\left(\frac{g_0m_0d_0N}{t}\right) = O(m_0E_{quant}). \quad (18)$$

Note that, when we include cost factors associated with the size and computational complexity of oracle queries, the mostly unpowered randomized preimage search is more energy intensive than Grover's algorithm by a factor of  $O(m_0)$ . Nonetheless, this factor is generally expected to be logarithmic in  $N$  and may easily be overwhelmed by the various costs associated with implementing fault tolerant quantum computation. We may also consider the fixed power and temperature scenario discussed in section 5. In this case, equations 9, 10, 11, and 12 become:

$$t_{cl} = O\left(\sqrt{\frac{g_0m_0d_0N}{P}}\right) \quad (19)$$

$$M_{cl} = O\left(\frac{\sqrt{g_0m_0d_0NP}}{T}\right) \quad (20)$$

$$t_{quant} = O\left(\sqrt{\frac{g_0 d_0 N}{P}}\right) \quad (21)$$

$$M_{quant} = O\left(\frac{m_0 d_0 P}{g_0 T^2}\right) \quad (22)$$

## 7 Conclusion

## References

1. Fredkin, E., Toffoli, T.: Conservative logic. *International Journal of Theoretical Physics* **21** (1982) 219–253
2. Bennett, C.H.: Logical reversibility of computation. *IBM J. Res. Dev.* **17** (1973) 525–532
3. Bennett, C.H.: The thermodynamics of computation: a review. *International Journal of Theoretical Physics* **21** (1982) 905–940
4. Landauer, R.: Irreversibility and heat generation in the computing process. *IBM J. Res. Dev.* **5** (1961) 183–191
5. Margolus, N., Levitin, L.B.: The maximum speed of dynamical evolution. *Physica D: Nonlinear Phenomena* **120** (1998) 188 – 195 *Proceedings of the Fourth Workshop on Physics and Consumption.*
6. van Oorschot, P.C., Wiener, M.J.: Parallel collision search with cryptanalytic applications. *Journal of Cryptology* **12** (1999) 1–28
7. Brassard, G., Høyer, P., Tapp, A. In: *Quantum cryptanalysis of hash and claw-free functions.* Springer Berlin Heidelberg, Berlin, Heidelberg (1998) 163–169
8. Bernstein, D.J.: Cost analysis of hash collisions: Will quantum computers make shares obsolete. (SHARCS09 Special-purpose Hardware for Attacking Cryptographic Systems) 105
9. Beals, R., Brierley, S., Gray, O., Harrow, A.W., Kutin, S., Linden, N., Shepherd, D., Stather, M.: Efficient distributed quantum computing. In: *Proc. R. Soc. A. Volume 469.*, The Royal Society (2013) 20120686
10. Giovannetti, V., Lloyd, S., Maccone, L.: Quantum random access memory. *Physical review letters* **100** (2008) 160501