

Quantum Computation for Artificial Intelligence

Pedro Araújo Rosa da Costa
pedro.a.r.costa@tecnico.ulisboa.pt

Instituto Superior Técnico, Lisboa, Portugal

March 2017

Abstract

Quantum computation methods (*e.g.*, Grover’s algorithm [20]) have been proven to be capable of quadratically improving the time-performance of classical search procedures. For Artificial Intelligence (AI) problems, this is roughly equivalent to halving the depth necessary for guaranteeing a solution. Unfortunately, study of these models is frequently left to those in the mathematics/physics communities, despite the aforementioned advantages. However, quantum computers will one day become a reality. Thus, we propose exploring how their methodologies may aid us in solving problems in AI, assuming the latter’s symbolic outlook. In this document, we provide scientists interested in quantum computing models with a brief introduction to operating principles, upon which we describe how a hybrid quantum/classical algorithm, Quantum Iterative Deepening Search (QIDS) [46], can be employed for solving interesting AI tasks within the standard-circuit quantum computation paradigm. We present two different applications, one possessing a more concrete representation, another that is more abstract, namely: a solver for the n -blocks generalization of the Blocks World planning problem; and a hybrid system of inference, able to prove statements in Propositional Calculus given a knowledge-base of assertions in that language. We also investigate how space-efficiency may be improved in the former, and suggest a new heuristic perspective for quantum computing frameworks. Our n -blocks solver can find solutions within an N -sized state-space in $O(\sqrt{N})$ time. Similarly, for large datasets, our inference system is able to prove a propositional statement in $O(\sqrt{N})$ time, with N the number of possible deduction sequences. Both exhibit a quadratic speed-up with regards to blind classical approaches, which require $O(N)$ evaluations.

Keywords: quantum computing, artificial intelligence, problem-solving, blocks world, automated inference

1. Introduction

Classical computation models are still at the heart most research techniques in the field of AI. However, some unconventional computation paradigms possess unique benefits. In particular, quantum computing algorithms which are able to quadratically improve the time performance of an equivalent classical procedure have been identified for two decades (*e.g.*, Grover’s algorithm [20]).

Additionally, the inevitable advent of a universal quantum computer has been generally accepted, with several private corporations and governments invested in realizing this vision. Thus, it is imperative that we in the AI community continue to accompany such efforts with our own research [55].

In this work, we propose investigating the benefits inherent to the application of quantum methods, such as QIDS [46], to symbolic AI tasks. In particular, we present quantum/classical solving systems for two distinct problems, one more abstract than the other: the solution of generalized n -

blocks Blocks World (BW) instances; and the proof of a given propositional expression from a program specified at the knowledge level, *i.e.*, a knowledge-base (KB). For both systems, we wish to demonstrate that the use of quantum techniques can provide a quadratic time-performance advantage. Additionally, we will suggest how space requirements may be reduced when employing QIDS, and offer a new perspective on the utilisation of heuristics alongside quantum computing methods.

The next section introduces the reader to quantum computation, with a particular emphasis on its standard-circuit model. There we also describe the two problem domains on which this work is focused, namely those of Blocks World planning and theorem proving. A brief survey of past research efforts on each is conjointly provided.

Section [Blocks World Solver](#) presents a quantum/classical approach to solving generalized Blocks World instances. We additionally propose a mechanism for reducing the solver’s (and QIDS’s)

space requirements, and advance an alternative viewpoint on the use of heuristics within quantum computation settings.

In section [Quantum Propositional Inference](#) we show how QIDS may be applied to tackle a more abstract problem, that of proving entailed Propositional Calculus expressions from a knowledge-base specified in that language.

Lastly, we convey our conclusions in the corresponding section.

2. Background

We begin with an [Introduction to Quantum Computation](#) and the main algorithms from that field which we will utilise in solving AI problems. This is followed by an overview of investigative work within the two specific domains we intend to tackle using symbolic strategies and the above computation models, in sub-section [Two Example Problems in Symbolic AI](#).

2.1. Introduction to Quantum Computation

In this sub-section we first review those concepts from [Reversible Computation](#) which provide the basis for quantum computing in the standard circuit paradigm. Subsequently, we describe the principal methods from [Quantum Computation](#) which we will employ, such as Grover’s algorithm, QIDS and the latter’s limited-depth counterpart.

2.1.1 Reversible Computation

To be able to effect a quantum computation, we must perform some manner of transformation over a physical system. The transitions that these quantum systems are able to carry out are restricted by nature to those which adhere to certain properties. The laws enforcing these properties reveal that reversible processes are precisely those which are permitted in nature [\[16\]](#).

However, standard Turing machines do not act in a reversible fashion. We cannot always obtain unique results from calculating the inverse of their transition functions, *i.e.*, the information which would allow us to retrace the process has been lost. The commonplace computer also operates in an irreversible manner: its *NAND* and *NOR* gates, the cornerstones of digital circuitry, are irreversible.

Fortunately, in his study of *Logical Reversibility of Computation*, Bennett proved that every irreversible computation has a classical reversible analogue, which obtains identical calculations at similar cost [\[1\]](#). The method he presents is able to simulate the operation of a single-tape Turing machine by way of a three-taped one, where each tape respectively records the *input*, the *history* of the computations, and the program *output*. Initially, the latter two are blank.

In such a machine, an irreversible computation

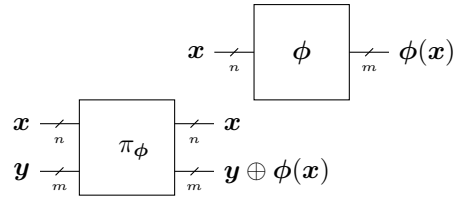


Figure 1: Similarly to how we interpret gates in digital networks, each box calculates a function based on its inputs. The rightmost gate calculates $\phi(x)$ in an irreversible process. However, its reversible analogue π_ϕ , computes $\phi(x)$ in a reversible fashion, by using the ancillary inputs y . Note: $\text{---}/\text{---}$ represents a k bit wire. (Source: [\[48\]](#)).

is divided into three phases: the first stage executes the same calculation as its single-tape analogue, storing the result in the input tape. But instead of losing intermediate calculations, these are recorded in the history tape; the second phase consists in copying the contents of the input to the output tape; finally, the inverse of the original computation in the input tape is calculated, using data from the history tape. When all three stages have been executed, the intermediate steps have been retraced, so as to obtain the program’s input in the input tape, a blank history tape, and the output of the irreversible operation on the output tape.

The technique was subsequently expanded by Toffoli [\[48\]](#), who proposed an approach for calculating any finite function with a reversible logic network, making use of ancillary bits of data for storing the history of intermediate computations.

Toffoli demonstrated that a reversible analogue $\pi_\phi : \mathbb{Z}_l \rightarrow \mathbb{Z}_l$ of a finite function $\phi : \mathbb{Z}_n \rightarrow \mathbb{Z}_m$ could be determined in the following manner:

$$\pi_\phi : (\mathbf{x}, \mathbf{y}) \mapsto (\mathbf{x}, \mathbf{y} \oplus \phi(\mathbf{x})) \quad (1)$$

where \mathbb{Z}_l is the set of integers from 0 to $2^l - 1$, $l = n + m$, \mathbf{y} are the ancillary bits and \oplus is the bitwise *XOR* operation.

In figure 1 the reader may see both ϕ and π_ϕ depicted from a black box perspective, as if these were simple gates in a logic circuit. Note how passing $\phi(x)$ as ancillary input to the reversible gate results in $\pi_\phi(\mathbf{x}, \phi(\mathbf{x})) = (\mathbf{x}, \phi(\mathbf{x}) \oplus \phi(\mathbf{x})) = (\mathbf{x}, \mathbf{0})$.

The author also described how any function under composition could be computed by combining *AND/NAND* (or Toffoli) gates alone. Theoretically, such networks were shown to operate without loss of information, and with an efficiency comparable to the corresponding irreversible computation. Storage requirements were also demonstrated to increase in proportion with the number of function arguments, instead of how many gates were used.

The operations in Toffoli’s reversible circuits may

be mathematically expressed by *unitary* transforms of linear algebra.¹ These will reproduce the action of each gate in the circuitry, according to every input/output combination defined in equation (1). In particular, unitarity of transformations is a main characteristic of quantum mechanical processes.

Note that replacing each gate in Toffoli’s classical reversible circuits with its quantum analogue, we obtain a quantum circuit that is equivalent in function, and that is comparably efficient to the irreversible calculation [34, p. 101].

2.1.2 Quantum Computation

Various models for quantum computation could be applied for problem solving in AI, although we will employ the *standard quantum-circuit model*. Every operation that is carried out within this paradigm is reversible, and any randomness restricted to the act of measurement.

Quantum algorithms in this model typically apply one of two methods: using the *quantum Fourier transform (QFT)* to find the period of some problem describing function; or effecting *amplitude amplification* to discover a marked item in a database.

Peter Shor designed a QFT-centred strategy for performing fast integer factorization in his eponymous algorithm [39, 40], proving that quantum computation is capable of tractably breaking the classical cryptographic techniques we currently rely upon. Furthermore, it is widely considered that no efficient classical alternative to the technique exists.

However, such QFT-based strategies are effective when problems bear unique structural regularities. Amplitude amplification methods, on the other hand, are applicable to a wider range of domains. Grover’s algorithm [20] is paradigmatic of such strategies. It is able to ascertain a solution to a black box problem within an unstructured N -element dataset in $O(\sqrt{N})$ time.

Since [Grover’s Algorithm](#) is central to the QIDS approach we employ, we describe it under its own heading. This is followed by a broad specification of the [Quantum Iterative Deepening Search](#) method, as well as a discussion on how it can be applied as a problem-solving approach in AI. Finally, a restricted version of QIDS, [Limited Quantum Iterative Deepening Search](#), which we will utilise in our inference method, is defined and analysed.

Grover’s Algorithm Grover’s search algorithm is able to discover a marked item from within an unsorted dataset bearing N constituents [20]. To do so, his procedure presumes access to a black-box oracle function that can evaluate whether a given element is marked. In its initial formulation, the

¹ U is a unitary transform provided its conjugate transpose U^\dagger is its inverse. *I.e.*, $U^{-1}U = U^\dagger U = I$ [27, p. 19].

database held a single tagged element, which the method could find in $O(\sqrt{N})$ oracle queries. However, the algorithm can be generalized for multiple marked elements, without suffering an impact on its overall complexity.

Compared to classical procedures, and assuming the dataset is unstructured, Grover’s method yields a quadratic speed-up, since the former techniques require $O(N)$ oracle interrogations. The advantage stems from quantum parallelism and interference.

Assuming an efficiently computable black-box predicate p exists, then we can construct a transform U_p from its reversible analogue π_p , as presented in [Reversible Computation](#). Applying U_p to the dataset’s equivalent to the even distribution superposition of its basis states (calculated using the *Walsh-Hadamard* transform, W), we can obtain a mapping akin to that of equation (1):

$$U_p : (W |00\dots 0\rangle) \otimes |a\rangle = \frac{1}{\sqrt{N}} \sum_{\mathbf{x}} |\mathbf{x}, a\rangle \mapsto \frac{1}{\sqrt{N}} \sum_{\mathbf{x}} |\mathbf{x}, a \oplus p(\mathbf{x})\rangle \quad (2)$$

where $p(\mathbf{x})$ is an n -ary boolean function, such that $p(\mathbf{x}) = 1$ for marked items and $p(\mathbf{x}) = 0$ otherwise, and a is an auxiliary qubit, prepared in the superposition $H|1\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle) := |-\rangle$. H is the *Hadamard* matrix, and can obtain an even distribution of a single qubit’s bases. The consequence is that each input is now intertwined (*i.e.*, *entangled*) with its corresponding predicate result.

At each step, the algorithm first inverts the sign of those elements which are marked. The remaining items are left unchanged. This is followed by a diffusion transformation, which essentially reflects each element about the superposition’s average.

The process progressively increases the probability of obtaining one of the marked elements upon measurement of the superimposed qubits. Successful computations are thus reinforced, and after $O(\sqrt{N})$ steps the marked item can be deterministically retrieved [5].

It was later shown [2] that any quantum algorithm making use of an oracle function cannot employ less than $\Omega(\sqrt{N})$ black box queries, thus demonstrating that Grover’s method is optimal.

Quantum Iterative Deepening Search This hybrid strategy uses quantum techniques, such as superpositioning and Grover’s algorithm, along with classical iterative-deepening [26], to perform a quantum hierarchical search of a tree, *i.e.*, a *Quantum Iterative Deepening Search* [46] procedure.

The methodology inherent to the approach stipulates that we encode the state space of a problem into a qubit register (*i.e.*, a Hilbert space), which we manipulate using linear transformations. Ancil-

lary qubits record each operator application, so that every transformation is reversible. For a particular goal or solution test $g(\mathbf{s})$ (with \mathbf{s} a state encoding) of the problem, its respective solution test matrix U_g (we will refer to it as G) can be obtained according to equation (1), and implemented as a logic gate corresponding to figure 2, left. The gate will distinguish whether or not a given state is a solution by flipping its ancillary input qubit.

Transitions within the problem domain (*e.g.*, moving a disk in Tower of Hanoi, a robot’s arm, etc) can be modelled according to equation (1) and a state transition function \mathbf{f} , producing a unitary gate T , as in figure 2. As inputs, these transition operators receive both a state and some choice on how to proceed, and are “chained” together, so that each following operator accepts the state resulting from the previous transition.

Because Grover’s method requires that our input be entangled with the result of the predicate function, each gate’s behaviour is “undone” by applying the corresponding operator inverses in opposite order (see figure 3 for a depth 2 chain, figure 4 for a depth d one). This would leave us with our input on the registers again, a useless computation. However, the approach prescribes the use of the C_{not} gate² as a predicate for Grover’s algorithm, thereby “preserving” the circuit’s outcome without compromising permutability.

Every possible sequence of choices (*i.e.*, a *path*) is explored by placing the qubits recording these (the *path register*) in an even distribution superposition of its basis states, using the method described in the previous heading. In this manner, the technique effectively implements a uniform branching factor tree search over the problem space with depth equivalent to the number of transitions in the chain.

Although most problems possess a non-constant branching factor domain, we can always map their

² the gate flips an auxiliary qubit contingent on a control qubit’s value. Its behaviour is defined by the mapping $C_{not} : (c, x) \mapsto (c, c \oplus x)$. In QIDS, the result of the goal test is taken as the control qubit, so that the C_{not} ’s ancillary qubit may be used to predicate Grover’s algorithm, as in equation (2).

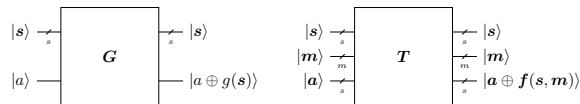


Figure 2: Unitary operator gates for general domains. The $|\mathbf{s}\rangle$ represent states, the $|\mathbf{m}\rangle$ are transition choices, and the qubits $|\mathbf{a}\rangle$ and $|\mathbf{a}'\rangle$ are auxiliary inputs, for storing the output of computations. The left gate is the reversible analogue to the solution test function $g(\mathbf{s})$. The right gate computes the state resulting from a transition \mathbf{m} over the state \mathbf{s} , according to a transition function $\mathbf{f}(\mathbf{s}, \mathbf{m})$.

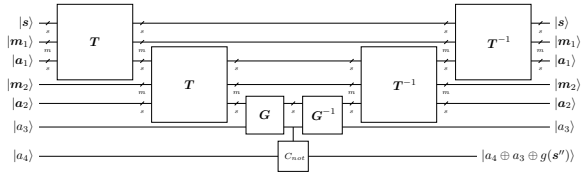


Figure 3: Reversible circuit for a tree search over two transition operations and a goal test, for a generic problem domain. Note that $\mathbf{s}'' = \mathbf{a}_2 \oplus \mathbf{f}(\mathbf{m}_2, \mathbf{s}')$, and $\mathbf{s}' = \mathbf{a}_1 \oplus \mathbf{f}(\mathbf{m}_1, \mathbf{s})$.

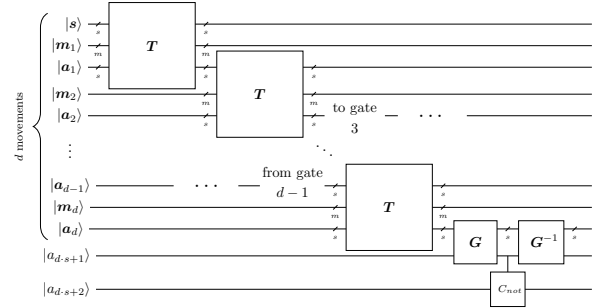


Figure 4: First section of the reversible chain for depth d . The remainder of the circuit is analogous with that of figure 3. Note that the auxiliary qubits differ in their indexing, according to whether or not they represent multiple-qubit systems. Multiple-qubit kets are highlighted in bold, indexed according to transition order. Single qubit kets are not shown in bold face, and are indexed according to the number of auxiliary qubits for the entire circuit.

non-uniform tree structure into a uniform analogue, by considering any excess transitions as having no effect over a state [45].

To execute search over all depths of the tree, the technique begins by building a depth 0 chain (0 transitions and a goal test). When a solution is found, the algorithm terminates, returning the path register. Otherwise, it increments depth, repeating the above. Naturally, if no solution exists, it will proceed indefinitely.

QIDS approaches have been successfully designed to solve such problems as the sliding block puzzle [44] or the simulation of Turing machines [46].

Limited Quantum Iterative Deepening Limiting the depth at which QIDS should stop the search process allows us to define its depth-limited counterpart, aptly dubbed Limited Quantum Iterative Deepening Search (LQIDS).

The procedure accepts the following as input:

- A starting state \mathbf{s}_0 for the search;
- A function $\mathbf{f}(\mathbf{s}, \mathbf{m})$, which calculates the state resulting from transition \mathbf{m} in state \mathbf{s} ;

- A goal test function $g(\mathbf{s})$, evaluating whether its argument is a goal state;
- A depth limit ζ for the search policy.

From depths 0 to ζ , LQIDS is identical to QIDS. Nonetheless, for clarity, we will describe it here in detail. Firstly, a quantum circuit akin to that of figure 4 for depth d is constructed, which entails using equation (1) and functions \mathbf{f} , g to obtain quantum gates for transitions and goal test (respectively T and G). Formally, the two gates may be portrayed as the two linear transforms:

$$\begin{aligned} T &: \alpha |\mathbf{s}\rangle |\mathbf{m}\rangle |\mathbf{a}\rangle \mapsto \alpha |\mathbf{s}\rangle |\mathbf{m}\rangle |\mathbf{a} \oplus \mathbf{f}(\mathbf{s}, \mathbf{m})\rangle \\ G &: \alpha |\mathbf{s}\rangle |a\rangle \mapsto \alpha |\mathbf{s}\rangle |a \oplus g(\mathbf{s})\rangle \end{aligned} \quad (3)$$

where \mathbf{a} is T 's ancillary qubit vector (with length equal to that of \mathbf{s}), a is G 's single auxiliary qubit and α are the coefficients for each ket.

Secondly, the starting state \mathbf{s}_0 is encoded into an initial state register $|\mathbf{s}\rangle$, and the even distribution superposition of every d -length path within the tree, $|\mathbf{w}\rangle = |\mathbf{m}_1\rangle \dots |\mathbf{m}_d\rangle$, obtained.

For every depth d , an equivalent number of transitions T are effected, upon which a goal test gate and a C_{not} operator are applied, so our circuit's input is $|\mathbf{s}\rangle |\mathbf{m}_1\rangle |\mathbf{a}_1\rangle \dots |\mathbf{m}_d\rangle |\mathbf{a}_d\rangle |a_{d.s+1}\rangle |a_{d.s+2}\rangle$. Further, the entire circuit operation $U_{c,d}$, can be mathematically expressed over this input as follows:

$$\begin{aligned} U_{c,d} &= C_d^{-1} (I^{\otimes s+d(m+s)} \otimes C_{not}) C_d \\ C_d &= (I^{\otimes d(m+s)} \otimes G) \prod_{i=1}^d (I^{\otimes (d-i)(m+s)} \otimes T) \end{aligned} \quad (4)$$

where $I^{\otimes n} = \bigotimes_{i=1}^n I$ does not alter the top n qubits. We also consider each qubit beyond G and T 's inputs in C_d 's definition to merely suffer an identity I transform, otherwise the notation would become unnecessarily cumbersome.

Provided every auxiliary input qubit is set to $|0\rangle$ (save for $|a_{d.s+2}\rangle$, which must be set to $|-\rangle$), running Grover's algorithm over the circuit will deterministically extract a solution (should one exist at that depth). Measuring the path register, we can readily verify if its choices lead us to a goal state via some verification function v , usually accepting \mathbf{s} , \mathbf{w} , \mathbf{f} , g and the current depth as arguments. For our purposes, v involves applying the d transitions in \mathbf{w} to \mathbf{s} in conformity with \mathbf{f} , and testing the final state via g . Success should be reported accordingly, along with the problem's solution (*i.e.*, the sequence of transitions \mathbf{w}).

Otherwise, the algorithm should repeat the above for a higher depth, until such a solution is determined, or the depth limit surpassed. In the latter case, LQIDS will indicate that none was found.

Note that the gates T and G must be pre-calculated to implement the circuit, a one-off

$O(2^{2s+m})$ time- and space-complexity process. Here, we establish a bound on the number of calculations by considering T as the largest of the gates.

When the shallowest solution to a problem is $l \leq \zeta$, then LQIDS has time-complexity identical to QIDS, that is $O(2^{lm/2})$, where m is the number of qubits needed to represent a transition choice. Otherwise, it runs in $O(2^{\zeta m/2})$ oracle queries.

Similarly, in terms of space-complexity, if $l \leq \zeta$, both algorithms require $O(s+l(m+s))$ qubits, with s the amount of qubits needed for state characterization. Naturally, when $l > \zeta$, LQIDS employs $O(s + \zeta(m+s))$ qubits in its circuits.

2.2. Two Example Problems in Symbolic AI

This sub-section examines the two symbolic AI problems which we attempt to solve with methods from quantum computation.

In [Blocks World](#) we describe that popular toy-problem and impart a comprehensive review of relevant previous work pertaining to it. Subsequently, we discuss research relating to our second problem domain, that of automated inference via [Knowledge-based Systems](#).

2.2.1 Blocks World

We begin with an introduction to the specific task with which we will illustrate the QIDS technique, under the heading [Problem Description](#). Finally, in the next of these, [From Blocks World to General Planning](#), we discuss which features of this micro-world are often relevant to AI problems involving other domains, and how these have historically influenced the development of general-purpose planning methods.

Problem Description BW has been a popular planning domain in AI research [6, 37, 41, 43, 51], first appearing in Winograd's [53] work as a micro-world for the SHRDLU simulated robot. While this early version contemplated coloured blocks of varying size and shape arranged over a surface, a simpler variant assumes a 2D representation, where blocks are colourless squares of identical size.

A Blocks World problem is traditionally exemplified using three blocks, although an n -block world is analogous. In that micro-world, three blocks lie arranged in some initial configuration over a table (as in figure 5, left). Each may have at most one block over it, and may be either over the table or another block.

Operators consist of moving one of the clear boxes³ to the table, or over another block that is also clear. The table is infinite, *i.e.*, it is always said to be clear. Our goal is to find a sequence of

³ that is, one without a block over it.

operations enabling us to transform the initial block configuration into the final one (figure 5, right).

On large-sized problems, the optimal solution for a BW instance is generally hard to find. Additionally, similar problems can have drastically different solutions. These characteristics make BW a relevant toy-world, particularly suited for showcasing search and planning efficiency over large instances.

Even when multiples blocks of the same label (*i.e.*, *interchangeable blocks*) can be present in the same state, finding a trivial solution is an easy task. For instance, we may simply place every block that is stacked over another in the start state on the table. The goal state can then be easily reached by moving each to its final position in succession. This approach yields a plan of length no worse than double that of an optimal solution.

However, Chenoweth [8] concluded that finding optimal solutions for a BW in which interchangeable blocks are permitted is NP-hard, similar in complexity to the travelling salesman problem [18, 31], optimally solving sliding block puzzles [33] or the bin-packing problem [19].

Chenoweth’s earlier work [7] had already shown that if no interchangeable blocks were present, “good” quality solution plans could be found in time $O(n^2)$, where n is the number of blocks. We refer to such a version of BW as an Elementary Blocks World (EBW)⁴. In keeping with Chenoweth’s results for BW with interchangeable blocks, Gupta and Nau [21, 22] demonstrated that while non-optimal solutions to EBW problems are similarly simple to obtain, optimal ones still involve NP-hard complexity (but no worse).

For years, the difficulty inherent to this task was ascribed to the existence of deleted-condition interactions, such as creative destruction [37] or Sussman’s Anomaly [43], where the side-effect of attaining one sub-goal deletes a condition needed for achieving the overall goal. Surprisingly though, Gupta and Nau found that optimal planning difficulty in EBW cannot be attributed to these interactions. Instead, they prove it to be the byproduct of a situation they term *deadlock*, wherein several

⁴ its description is well-known, and can be found in several AI textbooks, such as [30, 36].

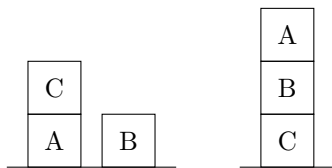


Figure 5: Example of a Blocks World problem: the left arrangement corresponds to the initial state, the final state to the right.

goals have yet to be achieved, none of which can be accomplished directly.

Resolving deadlocks requires an *enabling-condition interaction* [22, 29]. These achieve several goals simultaneously and can make it easier to accomplish others. However, when multiple deadlocks exist, it is frequently unclear which course of action obtains the best plan for the problem.

Problems in which no deadlocks are present can be optimally solved in time $O(n^3)$, using a hill-climbing strategy proposed by the authors, which moves blocks into positions consistent with the goal whenever doing so is possible. This is true regardless of the existence of deleted-condition interactions. In particular, traditional examples of situations which illustrate these interactions, such as creative destruction and Sussman’s Anomaly, do not contain deadlocks, and are thus easily solved by such a method.

When deadlocks *are* present, a solution with length no more than double that of an optimal one can be obtained in $O(n^3)$ with the same approach. Furthermore, the authors demonstrate that deciding whether or not a given EBW problem is solvable can be done in $O(n \log n)$ time-complexity, and that provided the decision is affirmative, a solution which moves no block more than twice can be extracted in time $O(n)$.

Similar complexity results were found for other generalizations of EBW, where the above technique is still applicable. In worlds where the blocks may vary in size and shape, the problem complexity is equivalent. If the table capacity is limited, the results are comparable: querying for solution existence, and finding one of length $O(n \log n)$ if it exists, is a low-order polynomial time problem. However, when both features are present, blocks of variable shape/size and limited table capacity, planning is more difficult. Because the shortest plans for some instances of this problem have exponential length, no deterministic polynomial-time procedure exists for finding them.

Nevertheless, for all of these EBW worlds, producing an optimal plan remains an NP-hard problem.

From Blocks World to General Planning

Several features of the BW domain are present in more realistic planning environments. Those traits have influenced the strategies employed for solving the latter problems. For instance, deleted-condition interactions exhibit some of the limitations of linear (progression/regression) planning methodologies. Such totally-ordered approaches cannot take advantage of goal decomposition and often compromise solution quality.

This led to the development of Partial-order

Planners (POPs), which allow interleaved actions to be included in a plan producing the overall objective. The distinction between the two strategies is pronounced. Whereas total-order planners are forced to make decisions earlier, POPs postpone these until the last possible moment, following a least-commitment heuristic. Plans generated from partially-ordered methods are often advantageous when adaptation to external factors is a necessity. Because some operators can be considered for execution in a non-deterministic fashion, more flexible scenarios may be contemplated.

On the other hand, the enabling-condition interactions observed in BW are particularly relevant when tractably obtaining an optimal plan is crucial. As observed in the previous heading, although creative destruction and Sussman’s Anomaly do influence plan quality, enabling-condition interactions are actually responsible for the complexity of optimal planning in that space. These effects remain pertinent for several domains where such plans are required [29], including logistic and transportation problems [50], scheduling of tasks [14], assuring resource redistribution [47], etc.

The QIDS approach, along with the decomposition heuristic we will introduce in section 3, are particularly suited to these types of optimal, parallelization-susceptible planning problems.

2.2.2 Knowledge-based Systems

The traditional search algorithm’s ability to make predictions and deductions from data they already possess is limited in scope by implementation-related details. Information-centric architectures arose from a need to endow our procedures with the capacity to operate over the knowledge level, reasoning about things that they know in a domain-independent fashion.

At the heart of these frameworks lies a *knowledge-base*, the collection of sentences currently held as being true. When combined with an *inference engine*, algorithms can deduce new facts from these, and add them to the database. Typically, a particular form of logic (*e.g.*, Propositional, Predicate, First-order, etc) is enforced when representing and manipulating phrases to discover new information.

Although elementary, the particular form of calculus upon which our method is focused, Propositional Logic [3], is still capable of drawing far-reaching conclusions. The *well formed formulae (wff)*, phrases or sentences allowed in its context are defined by a familiar *syntax* [36, p. 244], so we do not describe it here.

We next explore the two principal approaches for fact deduction: **Model Checking** and **Theorem Proving**.

Model Checking It is usual to consider the KB as a single sentence, constructed from its fact collection using some logical connective (*e.g.*, conjunction). We may then ask our knowledge-based system whether a particular sentence logically follows from the dataset,⁵ also known as the *propositional entailment problem*.

One way to go about answering such a query is dubbed *model checking*. It consists in enumerating every possible combination of true/false attributions to the atomic sentences in the database, similarly to how we solve Boolean Constraint Satisfaction Problems. The combinations themselves are referred to as *models*, and we say that a model satisfies a formula α whenever α is true for that configuration of assignments. Provided our query sentence is satisfied for every model where the KB is also evaluated truthfully, we may say that it follows logically from it [36, p. 240].

Propositional entailment is a co-NP-complete problem, so all known inference procedures for it are exponential on input size. The above process, for example, is $O(2^n)$, where n is the number of atomic propositions in the KB. Of course, it quickly becomes intractable.

However, an alternative model checking approach consists in proving that there are no models which satisfy both the KB and our query’s negation. Assuming our inquiry phrase is α , this equates to solving the SAT problem [9, 49], *i.e.*, finding a model which satisfies the Boolean *formulae* in $KB \wedge \neg\alpha$.

SAT is an NP-complete problem and has been the subject of much research over the years. The design of procedures able to solve these problems is a vast area of investigation, so we do not explore it further. Suffice it to say that modern SAT algorithms are capable of dealing with tens of thousands of propositional variables in an efficient manner.

Theorem Proving The model checking approach examined attribution possibilities in order to prove propositional entailment. *Theorem proving* methods, on the other hand, successively apply inference rules to formulas in a KB, until a query sentence is obtained,⁶ which may be advantageous in domains where the number of possible models for the KB is quite large, while the length of a proof is short. This is true in many task worlds, since a proof can ignore irrelevant variables during search.

A variety of inference rules may be used in deduction, and as long as these are *sound*⁷ they shall preserve entailment. One of the most familiar rules,

⁵ for a logical sentence α , this can be expressed as $KB \models \alpha$, *i.e.*, α *semantically* follows from fact set KB.

⁶ if a sentence α can be derived from KB we write $KB \vdash \alpha$, *i.e.*, α is *syntactically* derivable from KB.

⁷ an inference rule is sound when $\alpha \vdash \beta$ implies $\alpha \models \beta$, where α, β are formulas of a particular logic.

Modus Ponens (MP), is an example of a sound inference mechanism. The *resolution* rule is another.

Theorem provers can typically be divided into two categories: those which use the general form of the resolution rule for inference; and forward or backward chaining methods, which restrict that mechanism to a subset of propositional *formulae*.

As long as the problem of finding a proof is adequately defined, any classical search algorithm which is complete can be used to solve it. However, without suitable inference rules, even these techniques cannot produce a semantically complete⁸ deduction algorithm.

When used in conjunction with a complete search procedure, the resolution rule is sufficient for construction of a complete inference system. Because the resolution mechanism accepts disjunctions of literals, these processes must first convert the KB and query formula to a specific canonical format, conjunctive normal form (CNF).⁹

A proof by contradiction strategy is then employed: the technique attempts to prove the query formula α by finding an inconsistency,¹⁰ demonstrating that $\text{KB} \wedge \neg\alpha$ cannot be satisfied. The resolution rule is tried for each pair of clauses within $\text{KB} \wedge \neg\alpha$, until either all options have been exhausted (in which case α is false in KB), or an empty clause is obtained, and α has been proven.

Since any sentence in Propositional Calculus can be converted to its CNF analogue, the methodology is complete for that language.

In many situations, however, resolution performs less efficiently than simpler alternatives. Provided our KB is restricted to a certain subset of propositional expressions, forward or backward-chaining approaches may outperform more general resolution techniques. The language of Horn clauses is one such subset, and allows chaining algorithms to construct reasoning processes which are typically easier for us to follow than the general mechanisms.

Chaining algorithms make use of a KB consisting of definite clauses to prove a particular Horn formula α , using either MP or a restricted form of resolution as their inference rule. In forward variants, such as that of [13], deduction begins with the KB's set of facts, and every implication whose antecedent is satisfied will assert its consequent literal into the database. This is done (with or without heuristic assistance) until such time as α is obtained, and has thus been proved, or every implication rule has been evaluated. Backward chaining, on the other hand, starts from query formula α , which the method at-

⁸ a system is semantically complete iff $\alpha \models \beta \rightarrow \alpha \vdash \beta$.

⁹ advanced in [38], the notation represents any expression in Propositional Logic as a conjunction of disjunctive clauses.

¹⁰ a set of formulas Γ is inconsistent iff some formula γ can be found s.t. $\Gamma \vdash \gamma$ and $\Gamma \vdash \neg\gamma$.

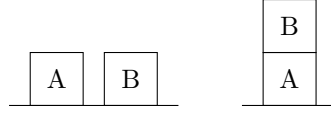


Figure 6: A 2-EBW problem: To the left is the initial state, with the final state on the right.

tempts to infer by recursively finding supporting evidence in the KB for every literal in α , essentially performing the dual procedure with respect to forward chaining approaches.

Both chaining strategies are complete for Horn clauses, and more importantly, run in time linear with the size of the knowledge-base.

3. Blocks World Solver

To illustrate the technique, we begin with its description and analysis for 2-blocks world problems in sub-section [2-EBW Solver](#). In [n-EBW Solver](#), the approach is extended for problems of n blocks, and the generalization is adapted to environments where qubit availability might be reduced, in [Reducing Space Requirements](#). We conclude the section with [Decomposition Heuristic](#), where we propose adopting an alternative heuristic perspective when solving problems using these computers.

3.1. 2-EBW Solver

We assume the initial and final states of such problems are of a form similar to figure 6. This will also serve as our example problem instance.

Since each block may be either on the table or over the other cube, a single qubit is sufficient for keeping track of its position. Two boxes mean a total of two qubits, s_1 and s_2 , can characterize a state. Similarly, one qubit is enough to reflect a single movement decision: we can either move block A or B. We will assume the encoding scheme in table 1 throughout our exposition.

Bit \ Value	0	1
s_1	A on table	A on B
s_2	B on table	B on A
m	move A	move B

Table 1: A 2-EBW problem encoding scheme.

Given a goal test of the form:

$$g(s_1, s_2) = \begin{cases} 1, & \text{if } (s_1, s_2) \text{ is a goal state} \\ 0, & \text{otherwise} \end{cases} \quad (5)$$

we are able to construct the goal test transform in accordance with equation (1), which we will designate as G in our calculations. Assuming registers of the form $|s_1 s_2, a\rangle$, where a is the auxiliary

qubit used to calculate the result of $g(s_1, s_2)$, then G maps every state s.t. $(s_1, s_2) \neq (0, 1)$ into itself. For $(s_1, s_2) = (0, 1)$, we have that $|s_1 s_2, a\rangle \mapsto |s_1 s_2, a \oplus g(s_1, s_2)\rangle$. The entire matrix can be consulted in appendix A.

We can extract a block movement operator similarly. Moving one of the boxes results in a new state description, so our output now involves two auxiliary qubits, as well as the input state (s_1, s_2) and a movement choice m . Thus, given a function $\mathbf{f}(s_1, s_2, m)$ which computes the new state:

$$\mathbf{f}(s_1, s_2, m) = \begin{cases} (\neg m, m), & \text{if } (s_1, s_2) = (0, 0) \\ (0, 0), & \text{if } (s_1, s_2, m) = (0, 1, 1) \\ (0, 0), & \text{if } (s_1, s_2, m) = (1, 0, 0) \\ (s_1, s_2), & \text{otherwise} \end{cases} \quad (6)$$

its unitary operator T can be determined. We leave its specification for appendix A.

The two operations can be portrayed as black box logic gates in a quantum circuit, as in figure 2.

If we consider that the entire depth-one circuit is modelled by an operator U_c , its application to the problem depicted in figure 6 results in the following superposition:

$$\begin{aligned} U_c |00\rangle \otimes H |0\rangle \otimes |000\rangle \otimes |-\rangle &= \\ &= \frac{1}{\sqrt{2}}(|000000\rangle - |001000\rangle) \otimes |-\rangle \end{aligned} \quad (7)$$

where we assume the layout of the qubits follows the pattern $|s_1 s_2 m a_1 a_2 a_3 a_4\rangle$. Note the application of a Hadamard transform to those qubits representing the different path choices m , and the auxiliary superposition $|-\rangle$ prescribed by Grover's technique. The full computation is exhibited in appendix B.

In this case, the resulting superposition is composed of merely two basis elements, so the solution $|001000\rangle$ is easily discerned, where moving block B leads to the goal state. In superpositions consisting of a higher number of bases, running Grover's algorithm over them will deterministically recover the solution.

Although the circuit of equation (7) is able to solve any 2-EBW problem whose initial state is a single movement removed from its final one, it is incapable of finding solutions for other instances.

Because we do not always know how many actions must be performed to obtain a solution, we adopt the QIDS approach when designing the quantum circuits, in keeping with their portrayal of figure 4.

3.2. n -EBW Solver

For an arbitrary n -block EBW problem, we can follow a similar encoding strategy. We may represent a state in such a world by registering the position of each block, *i.e.*, which object it is on. Each of these may be over one of $n - 1$ blocks, or the ta-

ble. Thus, n^n different state configurations exist, for which $s = \lceil \log_2 n^n \rceil$ qubits are required.

In addition, a single block movement can be defined as a $(source, destination)$ pair, in which case n blocks may be chosen as a *source*, with $n - 1$ other blocks as a *destination*. This amounts to $n(n - 1)$ eventualities, and $m = \lceil \log_2 n(n - 1) \rceil$ qubits for each action.

Assuming we have access to goal test $g(\mathbf{s})$ and transition result $\mathbf{f}(\mathbf{s}, \mathbf{m})$ functions analogous to those in equations (5) and (6), each of our operators can be reversibly defined as per equation (1), where \mathbf{s} is a state description vector in the complex Hilbert space \mathcal{H}^{2^s} and \mathbf{m} encodes our movement in a space \mathcal{H}^{2^m} . Consequently, our goal test gate G now accepts a state representation consisting of s qubits, and a single auxiliary qubit in which to calculate the test. Similarly, the movement gate T takes s qubits describing the current state, m qubits defining the movement decision, and a further s auxiliary qubits in which to compute the resulting state. Notice that the largest of these gates is the movement transition transform, requiring $O(2^{2s+m})$ pre-calculations.

Using the QIDS methodology reflected in figure 4's design for n -EBW problems, we see that a total of $n_{c,d} = s + d(m + s) + 2$ qubits are needed for construction of a depth d search (*i.e.*, performing d movements), $n_{c,d} - (s + dm) = ds + 2$ of which are auxiliary qubits. Applying Grover's algorithm means that each iteration requires $O(\sqrt{N})$ steps, with $N = 2^{dm}$, *i.e.*, the number of different paths. Naturally, if the shallowest depth for a solution is l , the entire process takes $O(\sqrt{2^{lm}})$ steps, with $O(l(m + s))$ space-complexity.

We may structure the d^{th} such circuit's input $|\mathbf{x}\rangle$ into several functionally distinct segments, like so $|\mathbf{x}\rangle = |\mathbf{s}\rangle (\otimes_{j=1}^d |\mathbf{m}_j\rangle |\mathbf{a}_j\rangle) |a_{d.s+1}\rangle |a_{d.s+2}\rangle$, where $|\mathbf{s}\rangle$ contains the initial state description, $|\mathbf{m}_j\rangle$ and $|\mathbf{a}_j\rangle$ correspond (respectively) to the action superposition and ancillary qubits for the j^{th} application of the movement operator, $|a_{d.s+1}\rangle$ is the auxiliary qubit for the goal test, and $|a_{d.s+2}\rangle$ for the C_{not} .

Accordingly, for circuit d , we can define the unitary operator $U_{c,d}$, as per equation (4). This would act over members of the above register space, performing the n -block equivalent of the computation suggested by figure 4.

3.3. Reducing Space Requirements

The reader will note that the n -EBW Solver we have presented requires a number of qubits that is linearly depth-dependent. However, in all likelihood, employing qubits will continue to be an expensive affair for a long time to come. Therefore, we may wish to decrease the amount of qubits the solver requires, at the expense of further one-off, pre-runtime calculation steps.

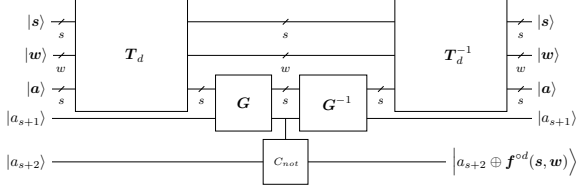


Figure 7: Circuit using the T_d operation, computing the result of a sequence of d movements and a goal test in n -EBW in a single step. s , s and m are as in section 3.2. $w = dm$, and $\mathbf{w} = (\mathbf{m}_1, \dots, \mathbf{m}_d)$ is the path register, where each $\mathbf{m}_i \in \mathcal{H}^{2^m}$ corresponds to an action over the i^{th} state obtained from executing actions \mathbf{m}_1 to \mathbf{m}_{i-1} sequentially, beginning at the initial state s .

Each iteration d of the method produces a circuit requiring $n_{c,d} = s + d(m + s) + 2$ qubits. Of these, $n_{c,d} - (s + dm) = ds + 2$ are auxiliary, *i.e.*, correspond to intermediate calculations. Because $s \gg m$, eliminating these will conserve more memory than any alternate encoding for our actions. This can be done, for instance, by calculating the results from performing the d actions in iteration d at once, instead of saving the intermediate states.

Assuming we have the function \mathbf{f} of n -EBW Solver, then we may define the construct described in the previous paragraph:

$$\mathbf{f}^{od}(s, \mathbf{w}) = \begin{cases} \underbrace{\mathbf{f}(\dots \mathbf{f}(s, \mathbf{m}_1) \dots, \mathbf{m}_d)}_{d \text{ times}}, & \text{if action sequence} \\ & \mathbf{w} = (\mathbf{m}_1, \dots, \mathbf{m}_d) \\ & \text{is valid in state } s \\ s, & \text{otherwise} \end{cases}$$

where our d actions are sequentially ordered into the vector $\mathbf{w} = (\mathbf{m}_1, \dots, \mathbf{m}_d)$, with \mathbf{m}_d corresponding to the last chronological movement, \mathbf{m}_1 equating to the first, and each \mathbf{m}_{i+1} immediately following \mathbf{m}_i for all $i \in \mathbb{N}$ s.t. $i < d$. Note that an action sequence \mathbf{w} is said to be valid when each of its individual actions are applicable in the state on which they are to be performed. Given these definitions, $\mathbf{f}^{od}(s, \mathbf{w})$'s reversible counterpart, which we refer to as T_d , can be extracted using equation (1), producing the circuit in figure 7.

While we may still use the state encoding described in n -EBW Solver, our actions must be mapped differently. Because we wish to perform the d actions in a single operation and each requires $m = \lceil \log_2 n(n-1) \rceil$ qubits, a total of $w = dm$ qubits can be used for path specification. Consequently, an action sequence vector \mathbf{w} lies in a $\mathcal{H}^{2^{dm}}$ Hilbert space, with separate qubits for different actions, organized hierarchically so that those referring to \mathbf{m}_{i+1} are contiguous to those of \mathbf{m}_i , and possess a lower order of magnitude (for instance).

The T_d operation involves a total cost of $2s + w = 2s + dm$ qubits. Whenever $d \geq 1$, $n \geq 2$, this

is less than or equal to the $s + d(m + s)$ qubits needed solely for movement computation in the n -EBW extension to figure 4's circuit. For the entire circuit of figure 7, the tally is $n_{c,d} = 2s + dm + 2$ qubits, and the amount of auxiliary space, $n_{c,d} - (s + dm) = s + 2$, is no longer depth-dependent. Assuming the shallowest solution has depth l , the solver will still extract a solution in $O(\sqrt{2^{lm}})$ steps. However, its circuits involve merely $O(s + lm)$ space-complexity. The absence of a depth coefficient over s is of note, considering $s \gg m$.

While the reduction in space is significant, the new movement operator has a drawback: for every iteration d , the T_d transform must have been pre-calculated. The cost of this computation is bounded from above by that of determining the largest such matrix, equivalent to $O(2^{2s+lm})$ steps. Although this is a one-off calculation, it rapidly becomes intractable, so would only be practical for less complicated instances.

An additional consequence is that the technique is only applicable up to a certain depth-limit, unlike our n -EBW Solver. Fortunately, in any n -EBW problem, the worst case length of an optimal solution is bounded from above by $2n$, equivalent to placing every block over the table, and then stacking each into their final positions in a bottom-up fashion [22]. Thus, we need only determine the T_d transforms up to that depth.

3.4. Decomposition Heuristic

While the quadratic speed-up obtained from employing Grover's algorithm is desirable, one may naturally consider if it would be possible to yield further gains using quantum analogues to classical heuristics. The latter usually provide a way to compare between different states, so as to select the best option among them.

In order to simulate such behaviour within a quantum framework, we would need the ability to exchange information between the various states of our superposition. Unfortunately, quantum mechanical laws such as the No-cloning theorem impose severe limitations on the communication mechanisms that may be applied.

We argue that a different heuristic perspective must be adopted within quantum computation, to circumvent the communication restrictions preventing path comparison. In particular, we propose that divide-and-conquer heuristics are more useful in such a scenario. These methods are capable of simplifying search processes whenever certain problem regularities are found, decomposing them into less complicated sub-problems.

We demonstrate an example of such a heuristic for an n -EBW problem. Suppose we can extract the directed acyclic graphs $\mathcal{G}_I = (V_I, E_I)$ and $\mathcal{G}_F =$

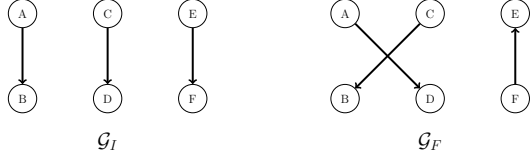


Figure 8: Directed acyclic graphs for representing initial (\mathcal{G}_I , left) and final (\mathcal{G}_F , right) states of the 6-EBW problem portrayed in figure 9.

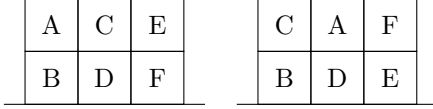


Figure 9: 6-EBW instance used to describe the decomposition approach. As usual, to the left lies the initial state and to the right our goal.

(V_F, E_F) from whichever representation we hold of the initial and final states of the problem, where V_i is graph \mathcal{G}_i 's set of *vertices*, containing a labelled vertex for each block in state i , and E_i is \mathcal{G}_i 's set of *edges*. For every block x over another y in state i s.t. y is not the table, a directed edge (x, y) is included in E_i . Figure 8 gives examples of these graphs for the 6-EBW problem in figure 9.

From these structures, we are able to devise an undirected graph $\mathcal{G} = (V, E)$, where $V = V_I = V_F$ and every edge in $E_I \cup E_F$ is present in E in an undirected manner. Figure 10 shows \mathcal{G} for the cited 6-EBW instance. A trivial, uninformed search procedure over it can obtain the graph's *connected components* [10, p. 88], *i.e.*, sets of vertices reachable to and from each other along a sequence of edges. These components capture sets of blocks whose movements from the initial to the final state *may* depend on each other's interim positions.

For each such component $V^{(i)}$, we then extract from \mathcal{G}_I and \mathcal{G}_F the sub-graphs [10, p. 88] $\mathcal{G}_I^{(i)}$ and $\mathcal{G}_F^{(i)}$ induced by $V^{(i)}$. These may be understood as representing partial-states, *i.e.*, reduced

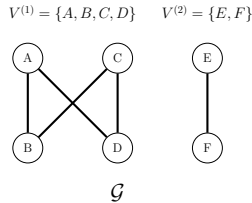


Figure 10: The undirected graph \mathcal{G} , crafted from \mathcal{G}_I and \mathcal{G}_F of figure 8. Two connected components $V^{(1)}$ and $V^{(2)}$ can be extracted from it, which reflect sets of blocks whose movements from initial to final configuration may be interdependent.

sections of the more general states described by \mathcal{G}_I and \mathcal{G}_F . In practice, we have decomposed our problem $\langle \mathcal{G}_I, \mathcal{G}_F \rangle$ into (hopefully) smaller, $|V^{(i)}|$ -block EBW sub-problems $\langle \mathcal{G}_I^{(i)}, \mathcal{G}_F^{(i)} \rangle$, where each part-state of component i can be encoded by our solver approach using $s^{(i)} = \lceil \log_2 |V^{(i)}| \rceil$ qubits.

Assume the initial and final part-state encodings are $\mathbf{s}_I^{(i)}$ and $\mathbf{s}_F^{(i)}$. Then we may define the following goal test function, where $\mathbf{s}^{(i)}$ represents a part-state within component i 's state-space:

$$g^{(i)}(\mathbf{s}^{(i)}) = \begin{cases} 1, & \text{if } \mathbf{s}^{(i)} = \mathbf{s}_F^{(i)} \\ 0, & \text{otherwise} \end{cases}$$

and craft its respective goal test operator as we have previously described.

If individual movement operations are considered, we know that $m^{(i)} = \lceil \log_2 |V^{(i)}| \rceil - 1$ qubits are sufficient for capturing a decision. Thus, given the transition function $\mathbf{f}(\mathbf{s}, \mathbf{m})$ for the generalized states, we can obtain an $\mathbf{f}^{(i)}(\mathbf{s}^{(i)}, \mathbf{m}^{(i)})$ for sub-problem i . To do so, we must define a relevant general state from which the action is to be performed. Let $\mathbf{s}_I^{(i)}$ be the generalized state in which every block in $V^{(i)}$ is over the same block specified in part-state encoding $\mathbf{s}^{(i)}$, while every other block is on its position in \mathbf{s}_I . Likewise, let $\mathbf{m}^{(i)}$ define the generalized encoding of the action equivalent to $\mathbf{m}^{(i)}$. We may then extract those values of $\mathbf{f}(\mathbf{s}_I^{(i)}, \mathbf{m}^{(i)})$ which are relevant to our $V^{(i)}$ blocks, and map them from the generalized encoding to our partial state representation. This gives us $\mathbf{f}^{(i)}(\mathbf{s}^{(i)}, \mathbf{m}^{(i)})$, which we may use to compute the unitary movement operator for sub-problem i .

Finally, for each $V^{(i)}$ with more than one block, we prepare a superposition of actions $\mathbf{w}^{(i)}$ and a quantum circuit according to the method shown in sub-section *n-EBW Solver*, where the initial state qubits should correspond to part-state $\mathbf{s}_I^{(i)}$. Running the described iterative-deepening approach alongside Grover's algorithm over these will recover the solution.

If there are k such connected components with $|V^{(i)}| > 1$, we obtain as many movement sequences, *i.e.*, solution paths $\mathbf{p}^{(i)}$. Since blocks in $V^{(i)}$ are never affected by those in $V^{(j)}$, where $i \neq j$, (*i.e.*, there are no deleted or enabling-condition interactions between the components) these can be independently ordered or interleaved in any plan taking us from the initial state to the final one. In a sense, this result is similar to what a POP framework would obtain.

The entire approach has total time-complexity of $O(\sum_{i=1}^k \sqrt{2^{l^{(i)} m^{(i)}}})$ oracle queries, where $l^{(i)}$ is the shallowest solution depth of the i^{th} connected component, *i.e.*, the length of $\mathbf{p}^{(i)}$. Naturally, the technique is dependent on how difficult the

most complex of the sub-problems is. However, a less restricted but easier to determine upper bound can be defined. Since the worst case optimal plan length for every $|V^{(i)}|$ -block world is double the number of blocks, we know the process is $O(\sum_{i=1}^k \sqrt{2^{|V^{(i)}|m^{(i)}}}) = O(\sum_{i=1}^k 2^{|V^{(i)}|m^{(i)}})$. This is dominated by the component with the most blocks, and so is $O(2^{|V^{(j)}|m^{(j)}})$, assuming $\forall_{i,i \neq j} |V^{(i)}| \leq |V^{(j)}|$.

The process is able to reduce time-complexity whenever an instance is decomposable. We posit that this is true of many configurations, where certain parts of the problem are not interdependent.

In remaining situations, the method does not increase time-complexity or compromise answer optimality. If every block is interdependent, a single connected component is extracted with $|V^{(1)}| = |V|$, yielding time-complexity $O(\sqrt{2^{lm}})$, equivalent to solving the more general problem without using the heuristic, where l is the shallowest solution depth and m is as defined in *n-EBW Solver*.

On the other hand, if we assume more than one component is found, we admit that the largest exponent in the terms of sum $\sum_{i=1}^k \sqrt{2^{l^{(i)}m^{(i)}}}$ occurs for some connected component j . Then decomposition allows us to solve the instance in $O(\sqrt{2^{l^{(j)}m^{(j)}}})$ steps, with $l^{(j)}$ the shallowest solution depth for the j^{th} sub-problem. The method we have described always finds the shortest paths for each of the sub-problems, and resolving the problem in this manner will never conceive plans which include movements between blocks in disconnected components. Consequently, we know $l^{(j)} \leq l$. Since $s^{(j)} < s$, $m^{(j)} < m$, time and space-complexity are improved.

However, such a technique is not always applicable. In the n -puzzle case [44], our divide-and-conquer approach is not so easily achieved, particularly when optimal solutions are desired. Some parts of the puzzle may not be independently solvable, and in many situations, a player must move tiles already placed in their final position to complete other regions of the board.

4. Quantum Propositional Inference

By accomplishing propositional inference using a combination of quantum and classical processes, we propose a knowledge-based mechanism which is able to prove Propositional Calculus expressions. Our strategy is detailed in this section.

We begin by presenting the *Inference Principles* upon which our system is grounded, along with a proof by cases approach to deducing expressions. Here, we also introduce the principal algorithm achieving that purpose. We then characterize the process of transforming each fact of a KB consisting of sentences from Propositional Logic into a structured dataset S suitable for our goals, in sub-section

Preparing the Knowledge-base. In *Fact Representation* we discuss how we have chosen to portray the state of inferences within our framework, as well as the alternatives considered in that regard. Sub-section *Dividing the Knowledge-base* contains a detailed exposition of how we decompose KB S into several sub-KBs σ of S , where proofs may be crafted in a more direct fashion. The use of quantum search techniques in our *Deductive Process* is subsequently elaborated upon. Finally, our *Analysis of the Method* is conveyed.

4.1. Inference Principles

Suppose we are given a certain knowledge-base \mathcal{KB} , a collection of assertions in Propositional Logic.

We designate $S = \langle \mathcal{F}, \mathcal{R} \rangle$ as the structured fact database resulting from the particular set of transformations we describe in the next sub-section. These will partition the original KB into two sets: whereas the *set of facts* \mathcal{F} will contain those expressions in \mathcal{KB} that are classified as *simple sentences* (i.e., they do not explicitly follow a pattern $\alpha \rightarrow \beta$, where an implication is the operator with the *lowest precedence*), the *set of rules* \mathcal{R} shall encompass the remaining *formulae*.

Elements of \mathcal{F} will then correspond to a collection of states of inference (with each state consisting of a semantic valuation of every literal), and rules in \mathcal{R} to propositions upon which MP may be applied to produce new assessments. The modification procedure induces a simpler characterization of a deductive state, as well as of the inference process itself, while still affording the knowledge-base designer with specification flexibility.

Often, we will refer to S as $\mathcal{F} \cup \mathcal{R}$. The two definitions are analogous, and represent the same KB in different ways, relying on context to disambiguate.

Our inference engine is predicated upon the use of two traditional inference rules, namely: *Modus Ponens* and *disjunction elimination*. As stated, we wish to determine if a query expression η ,¹¹ given in propositional form, follows from S . Algorithm 1 describes our approach, though we also provide an overview of our strategy here.

We first transform the facts in \mathcal{F} so that they possess only expressions containing \neg, \vee, \wedge . This is done in accordance with definitions for the remaining operators of Propositional Calculus. In addition, rules in \mathcal{R} are modified similarly, so that their antecedents and consequents are expressed merely with the connectives above.

Because it is more direct to carry out deduction in a KB σ whose facts in \mathcal{F} and consequents of \mathcal{R} are sentences composed only with literals and \wedge , a proof by cases strategy is employed in deriving

¹¹ where it is assumed that the propositional symbols in η are present in $\mathcal{F} \cup \mathcal{R}$.

Algorithm 1 Top level routine of our method, which attempts to prove the query η from the assumptions in S .

```

procedure PROVE( $S, \eta$ )
  ALL $\sigma$ CONTRADICTIONARY  $\leftarrow$  True
  PROOF  $\leftarrow$   $\emptyset$ 
   $\sigma$ PROOF  $\leftarrow$   $\emptyset$ 
  W  $\leftarrow$   $\emptyset$ 

  SYSLIST  $\leftarrow$  DIVIDE-KB( $S$ )

  for each  $\sigma \in$  SYSLIST do
    if CONTRADICTIONARY? $(\sigma) = False$  then
      ALL $\sigma$ CONTRADICTIONARY  $\leftarrow$  False
       $\sigma$ PROOF  $\leftarrow$  PROVE-QUERY( $\sigma, \eta$ )
      W  $\leftarrow$  SECOND( $\sigma$ PROOF)
      if FIRST( $\sigma$ PROOF) = False then
        return  $\langle False, \emptyset \rangle$ 
      else
        PROOF  $\leftarrow$  PROOF  $\cup$   $\{ \langle \sigma, W \rangle \}$ 
      end if
    end if
  end for each

  if ALL $\sigma$ CONTRADICTIONARY = True then
    return  $\langle Impossible, \emptyset \rangle$ 
  else
    return  $\langle True, PROOF \rangle$ 
  end if
end procedure

```

query η . This requires that we first eliminate the disjunctions present in these expressions. We do so by dividing S into the unique sub-cases σ of the KB, each portraying the disjunction alternatives, and such that $\bigvee_i \sigma_i = S$.

Each of the resulting independent knowledge-bases σ consist of a collection of working hypotheses, the cases over which disjunction elimination (or *proof by cases*) may be performed. Obtaining a proof of η in S can be accomplished by deducing a set of proofs for all of the σ . If two systems σ_1 and σ_2 were generated from the process of KB division, we would extract a proof using this technique. Frequently, its application is depicted as follows:

$$\begin{array}{ccc}
 [\sigma_1] & & [\sigma_2] \\
 \vdots & & \vdots \\
 (\sigma_1 \vee \sigma_2) & \eta & \eta \\
 \hline
 & \eta & E_\vee
 \end{array}$$

where the vertically bracketed σ_i are hypotheses, and the vertical dots indicate intermediate deductions via MP from which η can be inferred.

Thus, in addition to indicating whether or not a proof attempt has succeeded, Algorithm 1 returns

the set PROOF, a collection of $\langle \sigma_i, \alpha_{i,1} \dots \alpha_{i,l_i} = \eta \rangle$, where $\alpha_{i,1} \dots \alpha_{i,l_i} = \eta$ is the query's proof in each σ_i , with the $\alpha_{i,j}$ wff, and l_i the length of such a sequence in σ_i .

4.2. Preparing the Knowledge-base

Before inclusion into \mathcal{F} , simple statements must undergo the transformation process described below. Every \leftrightarrow and \rightarrow operator in the expression should be expanded in accordance with the definitions for implication and equivalence. Subsequently, every negation is successively administered using De Morgan's laws until these all occur in immediate precedence to a positive literal.

Likewise, those phrases in \mathcal{KB} not classified as simple sentences, thus being in the form $\alpha \rightarrow \beta$, will suffer the same transformations, but only for α and β . This leaves us with an expression retaining a single implication, where α and β contain no \leftrightarrow or \rightarrow operators.

Finally, an index over each sentence in the KB should be constructed, recording whether the phrase exhibits a disjunctive connective and if so, cataloguing the set of its disjunctive cases. For example, formula $A \wedge (B \vee C)$ would be stored as bearing a disjunction, along with its set of individual cases $\{A \wedge B \wedge C\}$, $\{A \wedge \neg B \wedge C\}$ and $\{A \wedge B \wedge \neg C\}$. Although the expansion is exponential, it is limited to the propositional variables within the sentence, and need only be done once, upon which multiple queries to the KB may be performed.

We assume Algorithm 1 to accept a KB $S = \langle \mathcal{F}, \mathcal{R} \rangle$ that has already been processed in this manner. Thus, when S is divided into its disjunctive cases $\sigma = \langle \mathcal{F}_i, \mathcal{R}_j \rangle$, definite semantic values for each disjunction in \mathcal{F} will be obtained from σ 's fact set \mathcal{F}_i . Similarly, because for each rule $\alpha \rightarrow \beta \in \mathcal{R}_j$, its consequent is now a conjunction of literals, it is easy to evaluate whether α is satisfied, and to modify literal values to conclude β when this is so.

Note that although some facts of \mathcal{F} could have been transformed in a different manner and included in \mathcal{R} as rules (*i.e.*, those with an \vee connective as the least precedent operator), we have chosen to provide the KB designer with some flexibility in defining assertions. Such a choice essentially translates into his/hers ability to prescribe a set of models in a more straightforward fashion.

4.3. Fact Representation

Our first task lies in assuring that states of inference have a manageable, but faithful representation.

If we were to consider capturing the truth or falsehood of every possible statement resulting from a combination of literals under the \vee and \wedge connectives, our memory requirements would sky-rocket.

However, there are only a total of $2^{2^2} = 16$ [11, p. 104] unique Boolean functions (*i.e.*, truth tables)

$H : \mathbb{B}^2 \rightarrow \mathbb{B}$ under the set of positive literals $\Sigma = \{A, B\}$, where $\mathbb{B} = \{F, T\}$, F represents a false value and T a true one. These tables uniquely determine (and are singularly determined by) a sentence in a canonical format, namely *full disjunctive normal form (DNF)*.¹² It is not surprising then that any propositional phrase can be portrayed by a unique one in full DNF [4, p. 73].

Since 16 unique functions are possible, the same number of qubits would be required to track these essential statements during inference with two symbols. For a set with 10 propositional variables, however, $2^{2^{10}} = 1024$ qubits would be needed. We would like to have a more condensed account of truthfulness between deductions, without sacrificing propositional expressiveness.

Notice that sentences in $\mathcal{F} \cup \mathcal{R}$ may contain disjunctive connectives. Suppose this were not the case for facts in \mathcal{F} , as well as consequents of rules in \mathcal{R} . These would then be conjunctions of literals. \mathcal{F} would be a single, large, disjunction-free conjunction, the *fact conjunct*. In that case, the deductive process would be simple: for each rule $\alpha \rightarrow \beta \in \mathcal{R}$ check that α 's propositional symbols are appropriately asserted in the fact conjunct so as to make α true. If so, and upon inference with *Modus Ponens*, every factor of the conjunction β can be independently asserted. Thus, we need only keep a record of each literal's semantic value. For n propositional symbols, this would involve linear memory requirements, namely a $2n$ qubit register would suffice.

We thus define our KB state in the following manner: if $\{a_1, a_2, \dots, a_n\}$ is the set of propositional symbols appearing in $\mathcal{F} \cup \mathcal{R}$, then qubit register $|\kappa\rangle = |a_1 a_2 \dots a_n \bar{a}_1 \bar{a}_2 \dots \bar{a}_n\rangle$ will portray the state of our fact database at each step. Initially, each a_i in the register will have its value set to 1 for truth, provided a_i is affirmed in the fact conjunct \mathcal{F} . Otherwise, its value will be 0. The same shall apply to the \bar{a}_i qubits, each reflecting whether or not the negative literal $\neg a_i$ is asserted in \mathcal{F} . This entails an initial stage of $O(2^{\mu_\omega} \cdot |\mathcal{F}|)$ time-complexity, where $|\mathcal{F}|$ is the fact set's cardinality, $\omega \in \mathcal{F}$ is its most complex formula¹³ and μ_ω the depth within the phrase of the conjunction with the highest precedence.

At this point, the reader may make the following objection: why should we keep both positive and negative literal values? Will not the information captured in the former be mirrored in the latter? We justify such a choice with our intention of

¹² that is, *formulae* which are disjunctions of conjunctive clauses, where the latter contain exactly one of the two literals corresponding to each propositional variable [11, p. 103]. Originally proposed by Boole in [4].

¹³ this is measured according to how deep in the formula the conjunction with the highest precedence is. *E.g.*, for a phrase $\omega = a \wedge ((b \wedge (c \wedge d)) \wedge e)$, $\mu_\omega = 4$ and we may establish an $O(2^{\mu_\omega})$ upper bound on the number of literals to set.

maintaining the inference system *monotonic*,¹⁴ and as simple as possible, without sacrificing recognition of propositional *formulae*. If we only tracked variables a_i , a *non-monotonic* logic would result.

While such mechanisms can be appropriate for situations in which beliefs may change over time, in general they are also more complex, both in time and space terms [17, p. 220]. For this reason, we opt to model the monotonic variety, maintaining the record described by $|\kappa\rangle$.

4.4. Dividing the Knowledge-base

We now turn our attention to the obvious: the hypothesis laid out in the previous sub-section, that every fact in \mathcal{F} and the consequents of rules in \mathcal{R} are all conjuncts of literals, does not hold for every KB expressed in Propositional Logic.

However, any knowledge-base S for which the above conjecture is unsupported can be divided into several modules in which it *is* satisfied. This is what is carried out at the beginning of Algorithm 1, when the procedure DIVIDE-KB is executed. That function obtains a collection of sub-KBs $\sigma = \mathcal{F}_i \cup \mathcal{R}_j$ of $S = \mathcal{F} \cup \mathcal{R}$, in which each fact of \mathcal{F}_i /consequent of a rule in \mathcal{R}_j corresponding to a fact in \mathcal{F} /consequent of a rule in \mathcal{R} bearing a disjunction, has been replaced by one of the disjunctive's individual cases.

DIVIDE-KB returns a list of $O(k^{2\xi})$ independent knowledge bases, where disjunctions can only occur in the antecedents of elements of their rule sets. ξ is the maximum of the number of *facts in \mathcal{F} /consequents of \mathcal{R}* possessing disjunctions, and k the maximum number of unique operands of the full DNF form of any of those expressions, restricted to the symbols they contain. Space-performance is also $O(k^{2\xi})$, whereas with respect to time, the entire method runs in $O(|\mathcal{B}| + k^{2\xi})$ operations, with $\mathcal{B} = \mathcal{F}$ or $\mathcal{B} = \mathcal{R}$. Naturally, we take \mathcal{B} so as to make expression $|\mathcal{B}| + k^{2\xi}$ the largest, as a worse case analysis.

4.5. Deductive Process

Having divided our original S , we now describe our approach for proof using quantum techniques. We ask the reader to refer back to Algorithm 1 whenever necessary.

More formally, a proof of query formula η from $S = \langle \mathcal{F}, \mathcal{R} \rangle$ is the sequence of wff s.t.:

$$\begin{array}{ll}
 S = \mathcal{F} \cup \mathcal{R} = \bigvee_{i,j} \mathcal{F}_i \cup \mathcal{R}_j & \text{Premiss} \\
 \sigma_1 = \mathcal{F}_1 \cup \mathcal{R}_1 & \text{Hypothesis} \\
 \alpha_{1,1} & \text{Modus Ponens} \\
 \vdots & \vdots
 \end{array}$$

¹⁴ in a monotonic logical system, an inference will never remove a literal from the set of true statements. Non-monotonic logics, however, allow this to occur [28, p. 335].

Algorithm 2 Detects whether a given knowledge-base $\sigma = \mathcal{F}_i \cup \mathcal{R}_j$ is contradictory.

procedure CONTRADICTORY?($\sigma = \mathcal{F}_i \cup \mathcal{R}_j$)

- Prepare the initial KB state vector \mathbf{s}_0 according to the facts in \mathcal{F}_i
- Let g be the goal test function and \mathbf{x} a KB state, such that

$$g(\mathbf{x}) = \begin{cases} 1, & \text{if } \mathbf{x}'\text{s contradiction bit is set} \\ 0, & \text{otherwise} \end{cases}$$

return FIRST(LQIDS($\mathbf{s}_0, \mathbf{f}, g, |\mathcal{R}_j|$))

end procedure

$\alpha_{1,l_1} = \eta$	<i>Modus Ponens</i>
\vdots	\vdots
$\sigma_{u \cdot v} = \mathcal{F}_u \cup \mathcal{R}_v$	Hypothesis
$\alpha_{u \cdot v, 1}$	<i>Modus Ponens</i>
\vdots	\vdots
$\alpha_{u \cdot v, l_{u \cdot v}} = \eta$	<i>Modus Ponens</i>
η	$E_{\vee} (S = \bigvee_i \sigma_i)$

where we consider $\mathcal{F} \cup \mathcal{R}$ as the conjunction of all expressions in that set. We also admit that u, v are bounded by $O(k^\xi)$ sets of $\mathcal{F}_i, \mathcal{R}_i$. In addition, an ordering of $\mathcal{F}_i \cup \mathcal{R}_j$ pairs is assumed, so that they are indexable by each σ 's subscript. Thus, there are $u \cdot v = O(k^{2\xi})$ independent sets σ_i .

The $\alpha_{i,j}$ are the wff obtained by inference from fact set σ_i using MP, and which lead up to the proof of η in that knowledge set. Their index j denotes the j^{th} wff in that proof, and l_i its length in σ_i . These $\alpha_{i,j}$ are deduced exclusively using MP, a fact which we will address in concluding this sub-section.

At the end of the proof structure, and upon obtaining η from each σ , we may effect disjunction elimination over the sub-KBs to derive η in S .

Returning to our analysis of Algorithm 1, having our list of σ_i we must now show η will be deduced from each. However, in proceeding, we require the notion of a *contradiction* within a set of facts σ : let a be a propositional symbol, should both a and $\neg a$ follow by deduction from σ when applying any combination of its rules and MP in sequence, then we may classify σ as contradictory, *i.e.*, there are no models satisfying σ 's collection of facts and σ is inconsistent. In particular, no model of σ will satisfy S , and in that regard, a proof of η from S ignores that σ in the proof structure above. This is what is first ascertained by the CONTRADICTORY? procedure, presented in Algorithm 2.

The CONTRADICTORY? function will accept a set of assertions $\sigma = \mathcal{F}_i \cup \mathcal{R}_j$, and make use of the limited form of the Quantum Iterative Deepening Search methodology to discern if these derive

a contradiction. This requires that we redefine our notion of a KB state to track such a situation.

Henceforth, our KB state will be $|\mathbf{s}\rangle = |\boldsymbol{\kappa}\rangle |c\rangle = |a_1 a_2 \dots a_n \bar{a}_1 \bar{a}_2 \dots \bar{a}_n\rangle |c\rangle$, where the $|c\rangle$ qubit will record whether or not a state has become contradictory. Although such information is already contained within $|\boldsymbol{\kappa}\rangle$, we would like to test and record such an occurrence in $O(1)$ time after its starting value is set. Having $|c\rangle$ included in the state description provides a simple means of achieving this.

As we discussed in 2.1.2, when we have defined a domain (state space, transition function, goal test), an iterative tree search procedure can be performed over it using QIDS. In this case, our domain consists of a state space given by every possible $|\mathbf{s}\rangle$, where we may trigger the rules in \mathcal{R}_j to transition between these. The depth of the search can be limited to the number of rules $\zeta = b = |\mathcal{R}_j|$ (*i.e.*, the branching factor), since repetition of a rule will never infer additional facts in a monotonic system of inference. The solution/goal test g will control the termination behaviour of CONTRADICTORY?'s call to LQIDS. For any KB state \mathbf{x} , it is able to test the $|c\rangle$ qubit and assume a value 0 or 1 according to its definition in Algorithm 2.

In this case, LQIDS takes \mathbf{s}_0 for its initial state (ascertained from \mathcal{F}_i , according to our description in section 4.3), a transition function \mathbf{f} (defined below) and the solution test g . Its depth limit is determined from \mathcal{R}_j , as $|\mathcal{R}_j|$. We describe the algorithm's execution for each step, from $d = 0$ to $|\mathcal{R}_j|$.

To begin with, the quantum circuit for the current depth is realized. The transition operator T will behave according to the rules in \mathcal{R}_j . This requires we formalize its classical counterpart $\mathbf{f}(\mathbf{s}, \mathbf{m})$, which will output the state resulting from transition \mathbf{m} when the KB state is a vector $\mathbf{s} = [a_1 a_2 \dots a_n \bar{a}_1 \bar{a}_2 \dots \bar{a}_n c]^T$:

$$\mathbf{f}(\mathbf{s}, \mathbf{m}) = \begin{cases} \mathbf{s}, & \text{if } \mathbf{m}'\text{s antecedent is not satisfied in } \mathbf{s} \\ \mathbf{s}', & \text{if } \mathbf{m}'\text{s antecedent is satisfied in } \mathbf{s} \text{ but} \\ & \text{a contradiction ensues¹⁵ from it} \\ \mathbf{s}'', & \text{otherwise} \end{cases}$$

where \mathbf{s}'' has bits corresponding to literals asserted in \mathbf{m} 's consequent conjunct set to 1, while every other bit is inherited from \mathbf{s} . \mathbf{s}' bears the same properties as \mathbf{s}'' , but its contradiction bit is set.

In this way, using equation (3) and \mathbf{f}, g , the gates T and G are extracted and the circuit constructed.

The reader will notice \mathbf{f} is not defined within Algorithm 2. It will be reused in a further procedure, so we prefer to characterize it here with the description above.

\mathbf{f} 's evaluation can be made $O(1)$, provided we pre-calculate and index its output for every situation. Assuming $s = 2n + 1$ is the number of bits re-

¹⁵ *i.e.*, at least one of the a_i and \bar{a}_i in \mathbf{s}' are both set.

quired for state data, and $m = \lceil \log_2 b \rceil$ the amount of bits needed to encode transition decisions, this entails an $O(2^{4n+\log_2 b})$ calculation.¹⁶ Evaluation of g can also be carried out in $O(1)$ time by adopting a similar indexing strategy, which requires a one-off $O(2^{2n})$ pre-computation. Multiple queries can then be performed without recalculation of these gates.

We now delineate the circuit’s initial input values. The state register $|s\rangle = |\kappa\rangle |c\rangle$ is defined by s_0 (and thus, by \mathcal{F}_i), as per our previous instructions. $|c\rangle$ ’s starting value can be set in $O(n)$ time, by checking whether $\forall_i a_i$ and \bar{a}_i are simultaneously asserted in κ . Every possible d -length rule application sequence in \mathcal{R}_j is constructed by placing the path qubits $|w\rangle$ in an even distribution superposition. Finally, each ancillary qubit is set according to the description laid out in 2.1.2.

We subsequently run Grover’s procedure over the circuit’s superposition. Should there be a solution at depth d , then Grover’s algorithm will deterministically find it, *i.e.*, measuring the path register $|w\rangle$ and verifying the sequence of rules from \mathcal{R}_j as leading to a solution will uncover whether the search succeeded. Should this be the case, LQIDS returns $\langle True, w \rangle$. Otherwise, it repeats the above for the next depth, until a solution is found or depth $|\mathcal{R}_j|$ is surpassed. If g is unsatisfied past this point, *False* is reported.

CONTRADICTION? will use the first element (*True* or *False*) of the LQIDS algorithm’s result to determine if indeed some contradiction has been found in the facts of σ . This would preclude its inclusion in a proof of η , since as we have stated, no model of σ is a model of S . Note however, that should every σ be contradictory, then of course, there are no models for S . In such a context, anything and everything is provable, so we must detect this situation, and classify the proof attempt as *Impossible* within our knowledge-base.

On the other hand, if at least one σ is not contradictory, then perhaps a proof for η can be found. Algorithm 1 attempts to extract one for each of these, via PROVE-QUERY, defined in Algorithm 3.

The procedure’s structure is almost identical to that of CONTRADICTION?’s. The differences lie in the goal test used to invoke LQIDS and built into the circuit, as well as in how it returns its result. While in CONTRADICTION? we tested the contradiction qubit, in this case we look for satisfiability of η . By proceeding in the manner prescribed by LQIDS, we will obtain the shortest proof of η under σ , should one exist. If it does not, PROVE-QUERY shall report η to be unprovable in σ , which is enough to signal to PROVE that it was unable to derive η from S .

Otherwise, Algorithm 3 will return the tuple

¹⁶ $\lceil \log_2 b \rceil \leq \log_2 b + 1$.

Algorithm 3 Finds a sequence of inferences from a knowledge-base $\sigma = \mathcal{F}_i \cup \mathcal{R}_j$, which when applied in succession will satisfy the query expression η .

```

procedure PROVE-QUERY( $\sigma = \mathcal{F}_i \cup \mathcal{R}_j, \eta$ )
  ◦ Prepare the initial KB state vector  $s_0$  according to the facts in  $\mathcal{F}_i$ 
  ◦ Let  $g$  be the goal test function and  $x$  a KB state, such that
    
$$g(x) = \begin{cases} 1, & \text{if } x \text{ satisfies expression } \eta \\ 0, & \text{otherwise} \end{cases}$$

  return LQIDS( $s_0, f, g, |\mathcal{R}_j|$ )
end procedure

```

$\langle True, w \rangle$ from LQIDS. Algorithm 1 verifies that the deduction is successful and adds the pair $\langle \sigma, w \rangle$ to the set PROOF. Should η be inferred for every non-contradictory σ , then a tuple $\langle True, PROOF \rangle$ is returned from PROVE, where set PROOF reflects the proof by cases strategy delineated at the beginning of this section.

Notice how LQIDS’s use of our f effectively models MP. This is why for each σ_i , the $\alpha_{i,j}$ of the aforementioned strategy are all deduced merely with this inference rule.

4.6. Analysis of the Method

The entire PROVE method has $O(|\mathcal{B}| + k^{2\xi} + k^{2\xi} \cdot 2 \cdot (2^{b \lceil \log_2 b \rceil / 2} + b \cdot (n + 2^{\mu_\omega} \cdot |\mathcal{F}| + b))) = O(|\mathcal{B}| + k^{2\xi} \cdot (2^{b \lceil \log_2 b \rceil / 2} + b \cdot (n + 2^{\mu_\omega} \cdot |\mathcal{F}| + b)))$ time-performance, where \mathcal{B} , k , ξ , b , n and μ_ω retain their previously defined meanings. Term $2 \cdot (2^{b \lceil \log_2 b \rceil / 2} + b \cdot (n + 2^{\mu_\omega} \cdot |\mathcal{F}| + b))$ arises from running LQIDS twice for each σ : once in CONTRADICTION? and again in PROVE-QUERY. Within these procedures, $2^{\mu_\omega} \cdot |\mathcal{F}|$ operations are needed to initialize the literal qubits of the starting state of the KB, n are required to set the first contradiction qubit and our verification process v is $O(b)$. Each of these steps is executed up to a maximum depth of b . From footnote¹⁶ and provided $b \geq 2$, we can also assert the bound to be $O(|\mathcal{B}| + k^{2\xi} \cdot (2^{(b \log_2 b)/2} + b \cdot (n + 2^{\mu_\omega} \cdot |\mathcal{F}| + b)))$.

The equivalent classical technique performs in $O(|\mathcal{B}| + k^{2\xi} \cdot (2^{b \log_2 b} + b \cdot (n + 2^{\mu_\omega} \cdot |\mathcal{F}| + b)))$ time. Since this is always greater than or equal to the hybrid method’s complexity, we know that there is an upper bound improvement.

As b grows quite large, $|\mathcal{B}| = |\mathcal{R}| = b$. Terms $|\mathcal{B}|$ and $b \cdot (n + 2^{\mu_\omega} \cdot |\mathcal{F}| + b)$ may be disregarded, as they are executed in both procedures. In that case, we let $N \simeq k^{2\xi} \cdot 2^{b \log_2 b}$ represent the number of deduction paths which are classically examined. When there are $\xi = b$ rules of \mathcal{R} with consequents possessing a disjunction and $k \ll b$, the algorithm runs at $O(\sqrt{N})$ time, a quadratic speed-up in comparison with classical $O(N)$ exhaustive search approaches.

In space-complexity terms, the technique requires

that we store $O(k^{2\xi})$ set elements. Additionally, since $s = 2n + 1$ qubits are required for state characterization, the quantum circuits need up to $O(n(b + 1) + b \log_2 b)$ qubits, which can be reused for inference in each of the σ . These two bounds constitute the method’s space-performance.

The methodology is semantically incomplete for wff of Propositional Calculus built using its logical connectives and positive literals present in $S = \mathcal{F} \cup \mathcal{R}$. On the other hand it does produce sound deductions. To see why this is so, note that both MP and E_\vee are sound derivation mechanisms.

PROVE is a complete search algorithm. For any input query η and KB S , it returns a solution, should one exist according to g . However, although the procedure LQIDS provides an optimal deduction of η (if it exists) for each σ , PROVE itself cannot ensure optimal proofs in S .

5. Conclusions

Our introduction of hybrid quantum/classical mechanisms for solving BW instances and performing knowledge-based propositional inference shows how quantum computation models may be utilised for solving non-trivial symbolic AI problems.

In particular, our BW Solver can find a goal state from an N -sized state-space in $O(\sqrt{N})$ steps, a quadratic improvement over blind classical methods, which require $O(N)$ operations. For n -block problems, the procedure has a complexity of $O(\sqrt{2^{lm}})$ steps, with $O(l(m + s))$ space-complexity, where l is the shallowest solution depth, and $m = \lceil \log_2 n(n - 1) \rceil$, $s = \lceil \log_2 n^n \rceil$ the number of qubits needed for action and state encoding (respectively).

Additionally, we have suggested a modification that is able to reduce the solver’s qubit requirements, so that it may be applied in more conservative memory environments. Because most of the necessary space emanates from auxiliary computations, omitting these was our primary concern. Consequently, instead of storing the results of d movement operations, the d^{th} iteration of the procedure prescribes an operation T_d , which calculates the outcome of the d movements at once, without keeping intermediate states within memory.

While the technique’s time-complexity is identical to that of the conventional solver’s, qubit requirements are now $O(s + lm)$, where the depth-dependent amount of auxiliary state storage has been eliminated. However, the process entails pre-calculating each of the T_d operators, *i.e.*, a one-off computation of at least $O(2^{2s+lm})$ steps, corresponding to the largest such transform. Since this quickly becomes intractable, application would be limited to less complicated instances, although the strategy is similarly amenable to more general tasks.

Within this domain, a novel heuristic perspective has also been proposed for use in quantum computation settings, as a contrast to the informed, comparative-appraisal heuristics often seen in traditional classical approaches. Because of the no-communication constraints intrinsic to quantum mechanical theory, the latter provide limited benefits in this context. Alternatively, we show how a divide-and-conquer heuristic can produce better results than our conventional QIDS solver for the same problem, provided it is decomposable.

If such is the case, then a classical method determines sets of blocks which are interdependent based on the initial and final configurations. These are separately solved using our n -EBW approach, resulting in $O(\sqrt{2^{l^{(j)}m^{(j)}}})$ ¹⁷ time-complexity, where j is the most complex of the sub-problems found, and $l^{(j)}$, $m^{(j)}$ have the same meaning in j as l , m had for the non-decomposed instance. Since the general problem is divisible, we have that $l^{(j)} \leq l$ and $m^{(j)} < m$. In particular, the speed-up is significant for situations where $l^{(j)} \ll l$ and $m^{(j)} \ll m$, *i.e.*, when the largest set of interdependent blocks is much smaller than considering all n blocks to be coupled. In addition, because $s^{(j)} < s$, space-complexity is also lowered.

On the other hand, when the problem instance is not decomposable, we have that $l^{(j)} = l$, $m^{(j)} = m$ and $s^{(j)} = s$, so the approach has complexity identical to that of solving without the heuristic.

Regarding the problem of propositional inference, our technique can provide a quadratic time-performance benefit with respect to its classical analogue, given a large KB of propositions. For an N sized state space (N different deduction paths), the speed-up obtained is equivalent to executing an $O(\sqrt{N})$ search, instead of an exhaustive $O(N)$ procedure, halving necessary verification depth within the domain.

The approach is a first attempt at employing the QIDS framework to address the problem of proving Propositional Logic *formulae*. Its characteristics make it a middle ground alternative to classical model checkers and theorem provers, lying within the complexity interval defined by the co-NP-complete propositional entailment problem and the linear time bounds encountered in language restricted, resolution-based procedures.

Acknowledgements

The author would like to thank Andreas Wichert for his openness, sense of humour, guidance and support over the previous year. Additionally, special

¹⁷ j is the most complex of the decomposed problems, $l^{(j)}$ its shallowest solution depth, $m^{(j)} = \lceil \log_2 n^{(j)}(n^{(j)} - 1) \rceil$ the number of qubits coding j ’s actions, and $n^{(j)}$ the j^{th} sub-problem’s number of blocks.

thanks go to Luísa Coheur and Catarina Moreira for their invaluable suggestions and input on improving this work. Responsibility for any errors or mistakes remains solely with the author.

A. 2-EBW Unitary Transforms

Table 2: The goal test unitary operator G . It can be inferred from $g : \{0, 1\}^2 \rightarrow \{0, 1\}$ in equation (5) over its $2^{(2+1)} \times 2^{(2+1)}$ tuples.

$ s_1 s_2 a\rangle$	000	001	010	011	100	101	110	111
000	1	0	0	0	0	0	0	0
001	0	1	0	0	0	0	0	0
010	0	0	0	1	0	0	0	0
011	0	0	1	0	0	0	0	0
100	0	0	0	0	1	0	0	0
101	0	0	0	0	0	1	0	0
110	0	0	0	0	0	0	1	0
111	0	0	0	0	0	0	0	1

Table 3: The movement unitary operator T can be determined from $f : \{0, 1\}^3 \rightarrow \{0, 1\}^2$ in equation (6). Each of its $2^{(3+2)} \times 2^{(3+2)}$ cells has a 1 if its row is (s_1, s_2, m, a_1, a_2) and its column $(s_1, s_2, m, a_1 \oplus f_1(s_1, s_2, m), a_2 \oplus f_2(s_1, s_2, m))$. Otherwise, the cell is valued as 0. Its truth table is partially represented here, where we assume $y_i = a_i \oplus f_i(s_1, s_2, m)$ and only non-zero valued row/column pairs are depicted.

Input					Output				
s_1	s_2	m	a_1	a_2	s_1	s_2	m	y_1	y_2
0	0	0	0	0	0	0	0	1	0
0	0	0	0	1	0	0	0	1	1
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
0	1	0	1	1	0	1	0	1	0
0	1	1	0	0	0	1	1	0	0
0	1	1	0	1	0	1	1	0	1
0	1	1	1	0	0	1	1	1	0
0	1	1	1	1	0	1	1	1	1
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
1	0	1	0	0	1	0	1	1	0
1	0	1	0	1	1	0	1	1	1
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
1	1	1	1	0	1	1	1	0	1
1	1	1	1	1	1	1	1	0	0

B. 2-EBW Calculations

We detail the steps involved in computing the result of equation (7) here. Throughout the exposition we assume that the layout of the qubits corresponds to $|s_1 s_2 m a_1 a_2 a_3 a_4\rangle$. Registers presenting merely n qubits, with $n < 7$, refer to the first n left qubits of this arrangement.

The first gate in the circuit acts upon the encoding of the initial state in the following manner:

$$\begin{aligned} T|00\rangle \otimes H|0\rangle \otimes |00\rangle &= T|00\rangle \otimes \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \otimes |00\rangle \\ &= \frac{1}{\sqrt{2}}(|00010\rangle + |00101\rangle) \end{aligned}$$

where the Hadamard transformation has been applied to qubit m so that we may explore all avenues. The remaining auxiliary qubits are appropriately assumed to be 0.

Appending the auxiliary a_3 qubit to this result and executing the goal test, we obtain:

$$\begin{aligned} I^{\otimes 3} \otimes G \left(\frac{1}{\sqrt{2}}(|00010\rangle + |00101\rangle) \otimes |0\rangle \right) &= \\ &= \frac{1}{\sqrt{2}}(|000100\rangle + |001011\rangle) \end{aligned}$$

where $I^{\otimes n} = \underbrace{I \otimes \dots \otimes I}_{n \text{ times}}$ leaves the top n qubits.

This is subjected to the C_{not} gate, after concatenating the results with the a_4 auxiliary qubit, suitably prepared in the superposition $|-\rangle$, as disclosed in Grover's method:

$$\begin{aligned} I^{\otimes 5} \otimes C_{not} \left(\frac{1}{\sqrt{2}}(|000100\rangle + |001011\rangle) \otimes |-\rangle \right) &= \\ &= \frac{1}{\sqrt{2}}(|000100\rangle - |001011\rangle) \otimes |-\rangle \end{aligned}$$

Finally, we wish to recover the auxiliary values present at the outset of the computation, in keeping with Bennett's stages of reversible computing. Thus, the inverse of the goal test is first applied over the superposition above:

$$\begin{aligned} I^{\otimes 3} \otimes G^{-1} \otimes I \left(\frac{1}{\sqrt{2}}(|000100\rangle - |001011\rangle) \otimes |-\rangle \right) &= \\ &= \frac{1}{\sqrt{2}}(|000100\rangle - |001010\rangle) \otimes |-\rangle \end{aligned}$$

followed by the inverse of the movement operator:

$$\begin{aligned} T^{-1} \otimes I^{\otimes 2} \left(\frac{1}{\sqrt{2}}(|000100\rangle - |001010\rangle) \otimes |-\rangle \right) &= \\ &= \frac{1}{\sqrt{2}}(|000000\rangle - |001000\rangle) \otimes |-\rangle \end{aligned}$$

Since in this case our superposition consists of two basis elements, the solution $|001000\rangle$ is easy to perceive, so we know that moving block B will lead to our goal. In more complex superpositions, running Grover's amplification over these will recover the solution with certainty.

References

- [1] C. H. Bennett. Logical reversibility of computation. *IBM Journal of Research and Development*, 17(6):525–532, Nov. 1973. 2
- [2] C. H. Bennett, E. Bernstein, G. Brassard, and U. Vazirani. Strengths and weaknesses of quantum computing. *SIAM Journal on Computing*, 26(5):1510–1523, Oct. 1997. 3
- [3] G. Boole. *The Mathematical Analysis of Logic*. Philosophical Library, New York, NY, USA, 1847. 7

- [4] G. Boole. *An Investigation of the Laws of Thought: On which are Founded the Mathematical Theories of Logic and Probabilities*. George Boole's collected logical works. Walton and Maberly, London, UK, 1854. [14](#)
- [5] G. Brassard, P. Hoyer, M. Mosca, and A. Tapp. Quantum amplitude amplification and estimation. *arXiv eprint: quant-ph/0005055*, May 2000. [3](#)
- [6] D. Chapman. Planning for conjunctive goals. Technical Report 802, AI Laboratory, Massachusetts Institute of Technology, Cambridge, MA, USA, May 1985. [5](#)
- [7] S. V. Chenoweth. Mathematical foundations for blocks world including a proof of np-completeness. Master's thesis, Wright State University, Dayton, OH, USA, 1986. [6](#)
- [8] S. V. Chenoweth. On the np-hardness of blocks world. In T. L. Dean and K. McKeown, editors, *Proceedings of the Ninth National Conference on Artificial Intelligence, Volume 2, AAAI '91*, pages 623–627, Menlo Park, CA, USA, July 1991. AAAI Press / MIT Press. [6](#)
- [9] S. A. Cook. The complexity of theorem-proving procedures. In M. A. Harrison, R. B. Banerji, and J. D. Ullman, editors, *Proceedings of the Third Annual ACM Symposium on Theory of Computing, STOC '71*, pages 151–158, New York, NY, USA, May 1971. ACM. [7](#)
- [10] T. H. Cormen, C. E. Leiserson, and R. L. Rivest. *Introduction to Algorithms*. MIT Press, Cambridge, MA, USA, 1st edition, 1990. [11](#)
- [11] B. Davey and H. Priestley. *Introduction to Lattices and Order*. Cambridge mathematical text books. Cambridge University Press, Cambridge, UK, 2nd edition, June 2002. [13](#), [14](#)
- [12] P. A. M. Dirac. A new notation for quantum mechanics. *Mathematical Proceedings of the Cambridge Philosophical Society*, 35(3):416–418, July 1939. No citations.
- [13] W. F. Dowling and J. H. Gallier. Linear-time algorithms for testing the satisfiability of propositional horn formulae. *The Journal of Logic Programming*, 1(3):267–284, Oct. 1984. [8](#)
- [14] R. Englert. Planning to optimize the umts call setup for the execution of mobile applications. *Applied Artificial Intelligence*, 19(2):99–117, Sept. 2005. [7](#)
- [15] S. E. Fahlman. A planning system for robot construction tasks. *Artificial Intelligence*, 5(1):1–49, Jan. 1974. No citations.
- [16] R. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6-7):467–488, June 1982. [2](#)
- [17] D. Gabbay and K. Schlechta. *Conditionals and Modularity in General Logics*. Cognitive Technologies. Springer, Berlin, Deutschland, 2011. [14](#)
- [18] M. R. Garey, R. L. Graham, and D. S. Johnson. Some np-complete geometric problems. In *Proceedings of the Eighth Annual ACM Symposium on Theory of Computing, STOC '76*, pages 10–22, New York, NY, USA, 1976. ACM. [6](#)
- [19] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., New York, NY, USA, 1979. [6](#)
- [20] L. K. Grover. A fast quantum mechanical algorithm for database search. In G. L. Miller, editor, *Proceedings of the Twenty-eighth Annual ACM Symposium on Theory of Computing, STOC '96*, pages 212–219, New York, NY, USA, 1996. ACM. [1](#), [3](#)
- [21] N. Gupta and D. S. Nau. Complexity results for blocks-world planning. In T. L. Dean and K. McKeown, editors, *Proceedings of the Ninth National Conference on Artificial Intelligence, Volume 2, AAAI '91*, pages 629–633, Menlo Park, CA, USA, July 1991. AAAI Press / The MIT Press. [6](#)
- [22] N. Gupta and D. S. Nau. On the complexity of blocks-world planning. *Artificial Intelligence*, 56(2-3):223–254, Aug. 1992. [6](#), [10](#)
- [23] W. Heisenberg. *The Physical Principles of the Quantum Theory*. University of Chicago Press, Chicago, IL, USA, 1930. No citations.
- [24] A. Horn. On sentences which are true of direct unions of algebras. *The Journal of Symbolic Logic*, 16(1):14–21, Mar. 1951. No citations.
- [25] G. Kirchhoff. Ueber das verhältniss zwischen dem emissionsvermögen und dem absorptionsvermögen der körper für wärme und licht. *Annalen der Physik*, 185(2):275–301, 1860. No citations.
- [26] R. E. Korf. Depth-first iterative-deepening: An optimal admissible tree search. *Artificial Intelligence*, 27(1):97–109, Sept. 1985. [3](#)

- [27] R. J. Lipton and K. W. Regan. *Quantum Algorithms via Linear Algebra: A Primer*. MIT Press, Cambridge, MA, USA, 2014. [3](#)
- [28] G. Luger. *Artificial Intelligence: Structures and Strategies for Complex Problem Solving*. Pearson Addison Wesley, Boston, MA, USA, 5th edition, Oct. 2005. [14](#)
- [29] D. S. Nau. Enabling-condition interactions and finding good plans. In A. Lansky, editor, *Proceedings of the 1993 AAAI Spring Symposium on Foundations of Automatic Planning: The Classical Approach & Beyond*, pages 93–97, Menlo Park, CA, USA, Apr. 1993. AAAI Press. [6](#), [7](#)
- [30] N. J. Nilsson. *Artificial Intelligence: A New Synthesis*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, Apr. 1998. [6](#)
- [31] C. H. Papadimitriou. The euclidean traveling salesman problem is np-complete. *Theoretical Computer Science*, 4(3):237–244, June 1977. [6](#)
- [32] E. L. Post. Introduction to a general theory of elementary propositions. *American Journal of Mathematics*, 43(3):163–185, July 1921. No citations.
- [33] D. Ratner and M. Warmuth. Finding a shortest solution for the $n \times n$ extension of the 15-puzzle is intractable. In T. Kehler and S. J. Rosenschein, editors, *Proceedings of the Fifth National Conference on Artificial Intelligence*, AAAI '86, pages 168–172, Menlo Park, CA, USA, Aug. 1986. AAAI Press. [6](#)
- [34] E. Rieffel and W. Polak. *Quantum Computing: A Gentle Introduction*. Scientific and Engineering Computation. MIT Press, Cambridge, MA, USA, 2011. [3](#)
- [35] J. A. Robinson. A machine-oriented logic based on the resolution principle. *Journal of the ACM*, 12(1):23–41, Jan. 1965. No citations.
- [36] S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall Press, Upper Saddle River, NJ, USA, 3rd edition, 2009. [6](#), [7](#)
- [37] E. D. Sacerdoti. *A Structure for Plans and Behavior*. PhD thesis, Stanford University, Stanford, CA, USA, Aug. 1975. [5](#), [6](#)
- [38] E. Schröder. *Der Operationskreis des Logikkalküls*. B. G. Teubner, Leipzig, Deutschland, 1877. [8](#)
- [39] P. W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. In S. Goldwasser, editor, *Proceedings of the Thirty-fifth Annual Symposium on Foundations of Computer Science*, SFCS '94, pages 124–134, Washington, DC, USA, Nov. 1994. IEEE Computer Society. [3](#)
- [40] P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, Oct. 1997. [3](#)
- [41] E. M. Soloway and E. M. Riseman. Levels of pattern description in learning. In R. Reddy, editor, *Proceedings of the Fifth International Joint Conference on Artificial Intelligence - Volume 2*, IJCAI '77, pages 801–811, San Francisco, CA, USA, Aug. 1977. Morgan Kaufmann Publishers Inc. [5](#)
- [42] H. P. Stapp. The copenhagen interpretation. *American Journal of Physics*, 40(8):1098–1116, 1972. No citations.
- [43] G. J. Sussman. A computational model of skill acquisition. Technical Report 297, AI Laboratory, Massachusetts Institute of Technology, Cambridge, MA, USA, Aug. 1973. [5](#), [6](#)
- [44] L. Tarrataca and A. Wichert. Problem-solving and quantum computation. *Cognitive Computation*, 3(3):510–524, Dec. 2011. [4](#), [12](#)
- [45] L. Tarrataca and A. Wichert. Tree search and quantum computation. *Quantum Information Processing*, 10(4):475–500, Nov. 2011. [4](#)
- [46] L. Tarrataca and A. Wichert. Quantum iterative deepening with an application to the halting problem. *PLoS ONE*, 8(3):1–9, Mar. 2013. [1](#), [3](#), [4](#)
- [47] S. Thiébaux, M. Cordier, O. Jehl, and J. Krivine. Supply restoration in power distribution systems: A case study in integrating model-based diagnosis and repair planning. In E. Horvitz and F. Jensen, editors, *Proceedings of the Twelfth International Conference on Uncertainty in Artificial Intelligence*, UAI '96, pages 525–532, San Francisco, CA, USA, Aug. 1996. Morgan Kaufmann Publishers Inc. [7](#)
- [48] T. Toffoli. Reversible computing. In J. de Bakker and J. van Leeuwen, editors, *Proceedings of the Seventh Colloquium on Automata, Languages and Programming*, volume 85, pages 632–644, Berlin, Deutschland, July 1980. Springer. [2](#)

- [49] B. A. Trakhtenbrot. A survey of russian approaches to perebor (brute-force searches) algorithms. *IEEE Annals of the History of Computing*, 6(4):384–400, Oct. 1984. [7](#)
- [50] M. Veloso. *Learning by Analogical Reasoning in General Problem Solving*. PhD thesis, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA, Aug. 1992. [7](#)
- [51] R. Waldinger. Achieving several goals simultaneously. Technical Report 107, AI Center, SRI International, Stanford, CA, USA, July 1975. [5](#)
- [52] A. N. Whitehead and B. Russell. *Principia Mathematica*, volume 1. Cambridge University Press, Cambridge, UK, 1910. No citations.
- [53] T. Winograd. *Procedures as a Representation for Data in a Computer Program for Understanding Natural Language*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, USA, Feb. 1971. [5](#)
- [54] L. Wittgenstein. *Tractatus Logico-philosophicus*. International Library of Psychology, Philosophy, and Scientific Method. Harcourt, Brace and Company, New York, NY, USA, 1922. No citations.
- [55] M. Ying. Quantum computation, quantum theory and ai. *Artificial Intelligence*, 174(2):162–176, 2010. [1](#)