Quantum algorithms for optimised planning – Technical & Business Feasibility Study

Premise – Existing artificial intelligence (AI) planning techniques can be enhanced by quantum algorithms to deliver optimised plans/schedules in real-time for complex tasks, considered too demanding for classical computers.

Quantum computing algorithms promise to increase the processing power by several orders of magnitude, depending on problem characteristics. Could they enhance existing AI planning algorithms to achieve anything like these improvements? When could this happen? This project will perform initial experiments with early quantum computing systems (D-Wave quantum annealer) and explore how/where universal quantum computing algorithms, often called circuit-model approaches could help realise more optimised planning.

Background – Generating optimal solutions to complex planning problems, best represented by the traveling salesman problem, has been designated in complexity theory as NP-hard. As a result, performing optimised planning on large complex planning problems is often not considered to be a good use of processing power.

For smaller planning/scheduling problems that limit the search for optimal solutions within a well-constrained solution space (e.g. route optimisation on roads for individual vehicles, airline gate allocation at airports, and local factory scheduling), optimised solutions can be generated dynamically with large conventional processors. Often these approaches seek to minimise key plan factors (e.g. time, cost) or maximise other measures (e.g. throughput), potentially using simulated annealing techniques.

As the solution space grows (e.g. traffic optimisation on wider transport networks, global telecom network optimisation, and just-in-time manufacturing across large industry networks) the processing power becomes not only prohibitively expensive, but also the solutions take too long to compute.

Consequently today, most (commercial) approaches to solving large planning and scheduling problems fall in two camps: (1) Seeking 'feasible' solutions, and selecting optimal solutions by hand (prone to human error); (2) Developing heuristics to find near optimal-solutions (often missing good solutions). Both approaches avoid generating optimal solutions, but still require considerable processing power to solve constraint satisfaction (rather than constraint optimisation) problems, which are a core part of automated planning systems.

Over the past 20 years, a range of constraint satisfaction algorithms, often described as SAT solvers, has been developed to check plan consistency (confirming that the plan is valid by checking all plan constraints are met). These have helped to find feasible plans within complex planning tasks, rather than optimal plans.

What current problem does it address – Enhancing AI/Hierarchical Task Network (HTN) planning techniques with quantum technology could permit optimal plans to be developed and refined dynamically (every second) in response to changing situation. Rather than choosing between the top 10-20 plans, the top 100, 10,000, 1M plans could be maintained in parallel and/or prioritised, based on known facts, with plan confidence levels changing as the situation/facts change. In addition to optimising plans for operational planning tasks, these techniques could be applied to plan/intent recognition and prediction tasks, such as threat assessment and financial fraud, permitting not only the recognition of threats/crimes as they occur, but also predicting them from early precursors.

Who will benefit from this? – Such enhanced capabilities could help operations teams to develop optimised plans/schedules for telecommunications/transportation networks, oil exploration and military/law enforcement operations, enabling them to manage resources more efficiently and maximise operational benefits.

Plantagenet SYSTEMS · BT · UCL · University of BRISTOL · D-Wave The Quantum Computing Company™

In addition to operational planning teams, they could help law enforcement teams not only recognise threats/crimes as they occur, but also predict them from early precursor information. They could help investigators with complex financial fraud cases that are often too convoluted, involving multinational corporations and organised crime groups.

What does your Experiment consist of? – Our proposal is to perform initial experiments with early quantum computing systems, such as the D-Wave quantum annealer, and then explore how circuit-model-based quantum algorithms could help. Although the D-Wave system is a special-purpose quantum annealer (with limited functionality), it is well suited to solving constraint satisfaction (CS) and quadratic unconstrained binary optimisation (QUBO) problems. Note that more than 99% of processing power required for running AI planners involves solving CS tasks.

We plan to leverage UCL's expertise with the D-Wave system to explore how to map quantum annealing approaches onto existing HTN planning algorithms, and then exploit University of Bristol's expertise in circuit-model bases quantum algorithms to enhance these algorithms further beyond quantum annealing methods.

We will also engage with NQIT (Oxford) Quantum Computing Hub to explore how their Q2020 universal quantum processor might enhance HTN planning algorithms (outside the scope of this task). Whereas today's quantum annealers might permit 100s or 1000s of competing plans to be maintained and tracked, future annealing or universal processors (like Q2020) may be able to target 100,000 to 1M+ competing plans eventually.

Experimental summary

1. Explore how quantum annealing techniques from D-Wave can enhance simple planning techniques to solve optimised planning tasks and refine these algorithmically (UCL).
2. Refine experiments to show how they will support optimised planning for network optimisation (BT).
3. Explore strengths and weaknesses of both quantum algorithmic approaches and propose integration of circuit-model approaches (Bristol) to address larger scale and more diverse optimised planning tasks.
4. Confirm business case for exploiting quantum-enhanced optimisation algorithms, including innovation workshops with other industry players to explore optimised plan recognition for financial fraud and other law enforcement tasks.

These activities would be performed over 12 months. Outputs include the following:

- Reports on experiments and enhancements from circuit-model approaches.
- Case studies on telecommunications network optimisation.
- Outputs from other commercialisation workshops on optimised plan recognition.

Who will the team comprise? – Leading academics from Bristol, UCL, plus BT and D-Wave, led by Plantagenet Systems[1], will explore not only how existing AI/HTN planning techniques can be enhanced with quantum technology to support optimised planning tasks, but also which planning tasks would benefit from improved optimisation. We will also engage with several other industry players from wider markets (e.g. distribution logistics, finance, law enforcement).

[1] Note that Plantagenet Systems Ltd is an independent consulting company that specialises in early R&D into quantum technologies. During 2016, Plantagenet Systems held two innovation workshops on the commercialisation of quantum technologies in finance and cyber security, funded by the Knowledge Transfer Network for Quantum (part of InnovateUK).

Overview of workplan

This task will be completed within 12 months involving 4 main work packages (WP)

- **WP1 Quantum Annealing experiments** – *Explore how quantum annealing techniques from D-Wave can enhance AI planning techniques to solve optimised planning tasks and refine these algorithmically (UCL):*
  - WP1.1 Review existing planning algorithms – *Confirm specific AI planning algorithms to benchmark (steady-state STRIPS, HTN planner).*
  - WP1.2 Explore quantum annealing versions – *Explore quantum annealing versions for optimised planning related to selected AI planning techniques.*
  - WP1.3 Map onto D-Wave processor – *Map quantum annealing planning algorithms onto D-Wave processor (Washington 2K qubits).*
  - WP1.4 Complete experiments – *Perform experiments on D-Wave processor.*
  - WP1.5 Review results and explore options – *Review results and next steps.*
- **WP2 Use case development/experiments** – *Refine experiments to show how quantum annealing versions will support optimised planning for network optimisation (BT):*
  - WP2.1 Existing optimised planning use cases (benchmark) – *Review existing optimised planning use cases for network optimisation and oil exploration to establish benchmarks.*
  - WP2.2 Refined quantum annealing versions – *Apply quantum annealing versions (WP1) to both use cases and refine.*
  - WP2.3 Review results and explore options – *Review results and next steps.*
- **WP3 Quantum annealing vs circuit-model** – *Explore strengths and weaknesses of both quantum algorithmic approaches and circuit-model approaches (Bristol) to address larger scale, more diverse optimised planning tasks:*
  - WP3.1 Strengths & weakness of both quantum approaches – *Explore strengths and weaknesses of both quantum approaches.*
  - WP3.2 Explore circuit-model versions – *Propose circuit model versions for optimised planning related to selected AI planning techniques.*
  - WP3.3 Review results and explore options – *Review results and next steps.*
- **WP4 Wider market use cases/workshops** – *Hold innovation workshops with other industry players to explore optimised planning/plan recognition for distribution logistics, financial fraud and other law enforcement tasks and perform detailed market assessment:*
  - WP4.1 Organise market use cases workshops – *Confirm invite list and organise series of one day workshops (2-3).*
  - WP4.2 Run and complete workshops – *Hold workshops to present progress and seek additional market use cases.*
  - WP4.3 Review results and final report – *Review results and combine outputs from all WPs into final report, together with detailed market assessment for relevant markets and business case/roadmap for next steps.*