National Aeronautics and Space Administration

Automating DSN Scheduling Problems Using Quantum Computing and Deep Reinforcement Learning

Principal Investigator: Brian Wilson (398); Co-Investigators: Alexandre Guillaume (398), Shahrouz Alimo (393), Edwin Goh (174), Hamsa Shwetha Venkataram (174), Mark Johnston (927), Thomas Claudet (393)

Program: FY21 R&TD Strategic Initiative   Strategic Focus Area: Quantum Scheduling

Objectives

Scheduling problems have traditionally been solved by formulating them as Mixed Integer Linear Programming (MILP) problems and then applying one of the many MILP solvers (e.g. CPLEX, Gurobi). However, for large scheduling problems with many integer variables the long solution time renders this approach problematic. The D-Wave quantum annealer solves optimization problems that have been formulated as Quadratic Unconstrained Binary Optimization (QUBO) problems. There has been increasing interest by the classical algorithm developers in solving QUBO formulations with improved performance by exploiting quantum-inspired algorithms, leading to a performance "race" between the classical, quantum and hybrid algorithms. DeepRL has also been shown to have unexpected power; AlphaGoZero achieved better than human performance at the game of Go after three weeks of training (playing itself). Without actually applying these techniques, it is difficult to predict (or characterize) which will be most effective.

For these reasons, we pursued a multi-pronged approach to solving DSN antenna scheduling problems:

- Formulation as **QUBO** for solution on the D-Wave quantum annealer or hybrid solvers.
- Formulation as **MILP** for solution by modern optimization solvers,
- Formulation as an "action space" and reward function for an agent to de-conflict the schedule via Deep Reinforcement Learning (**DeepRL**).

Background

The goal of this initiative is to improve the scheduling of spacecraft tracking passes onto the 34-m and 70-m antennas at the Deep Space Network (DSN) complexes. There are many months when multiple missions are requesting tracking passes from each DSN complex, resulting in the antennas being oversubscribed and leading to scheduling "conflicts" that must be resolved while maximizing science return. Experience with the DSN's mid-range Service Scheduling Software (S3) indicates that antennas can be as much as 60% over-subscribed during some weeks. Currently, priority levels in the scheduling tools are under-utilized and human intervention is needed, including negotiation between mission schedulers, to relax requirements appropriately. This leads to slower cycle times for "de-conflicting" interim and final schedules. Moreover, faster iteration with the DSN will enable JPL mission designers to consider many more "what if" scenarios in their planning processes.

This effort aims to: ***further automate the solution of antenna scheduling problems, learn from the corpus of existing schedules, generate "less conflicted" candidate schedules for human consideration, investigate greater use of existing priority levels in automatically generating schedules, and decrease the cycle time***. To accomplish this, we need to automatically solve these large, challenging optimization problems quickly.

Approach

Antenna scheduling is a Multi-Objective Optimization problem, where optimal decisions need to be taken in the presence of trade-offs between two or more conflicting objectives. The goal is to find a weekly tracking schedule that satisfies the constraints (e.g. antenna arraying or splitting tracks), has no conflicts, and maximizes the number of scheduled hours (RMS over the missions) but not at the expense of any one mission (percentage of satisfied hours for worst mission). To support the exploration, we generated synthetic but realistic problems (sets of mission requests) for multiple weeks that closely resemble the actual 'request set' that was used for the final scheduling outcome.

Results

QUBO. Alex Guillaume developed a sophisticated QUBO formulation for the problems and then solved them using the hybrid solver from D-Wave, yielding some of the best schedules in terms of total tracking hours scheduled. Schedules were also generated using a custom Quantum Inspired Optimization (QIO) solver from Microsoft and AlphaQUBO from MetaAnalytics. Figure 1 shows an example schedule for the oversubscribe week 40 of 2018. A detailed discussion of the QUBO hybrid solution compared to MILP and the Microsoft Custom Solver appears in the report appendix and reference. The quantum hybrid solutions were competitive in number of requests and hours scheduled, with a faster runtime than MILP.

MILP. Ryan Alimo and Thomas Claudet developed an improved version of the MILP formulation and solver, and an additional iterative algorithm to optimize the tradeoff between overall scheduled hours and a minimum percentage of hours for each mission. The delta-MILP algorithm iteratively increases a percentage threshold, adjusts weights on the mission requests, and reruns the MILP solver until all missions exceed that percentage of scheduled hours (actual / requested). Schedules for 6 weeks were generated with the threshold set to at least 40% for each mission, while only reducing the total scheduled hours by a small percentage. Figure 2 shows an example of a delta-MILP schedule for the very oversubscribed week 44 of 2016; Figure 3 shows the fraction of requested hours scheduled for each mission.

DeepRL. Edwin Goh and Hamsa Venkataram developed a DeepRL approach that involves an 'agent' learning a 'policy' to schedule the next tracking request given a 'state' of the environment (schedule so far). A simulation model of the DSN scheduling process was developed according to OpenAI's Gym API specifications. To improve generalization from week to week, a streaming formulation was adopted where requests were sequentially presented to the agent. This enabled the design of a three-dimensional, visual representation of the state space whose optimal latent embeddings could be learned by a convolutional policy network. Several variations were investigated: varying the agents' action space, changing the state space resolution, and experimenting with different heuristics. An investigation into curriculum learning (progressively harder problems) demonstrated the potential for reduced policy variance and improved time to convergence. The final generated schedules were competitive with the other two techniques in terms of number of hours scheduled but differed qualitatively. Figure 4 shows a snapshot of the generated schedule for week 44 of 2016.

Fig 1

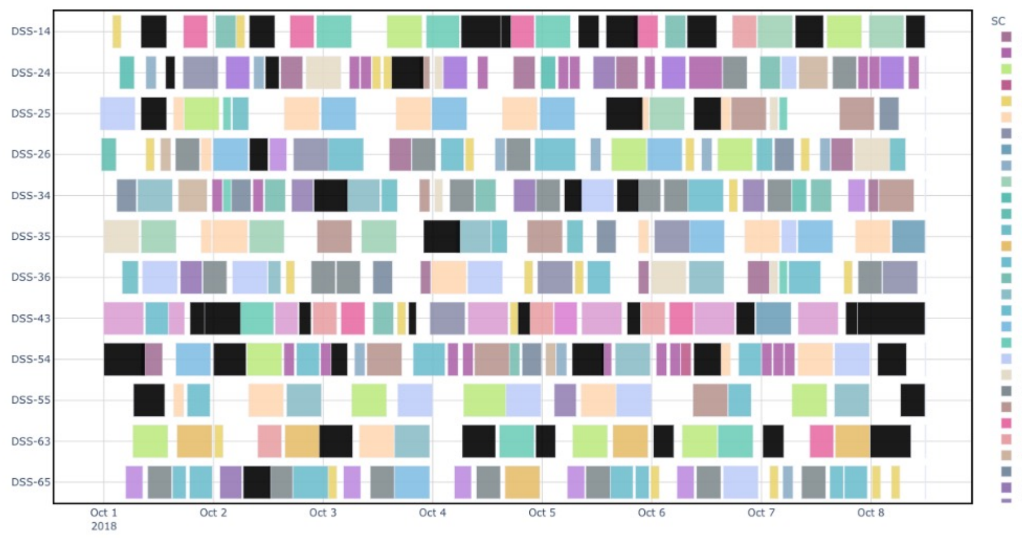

Fig 2

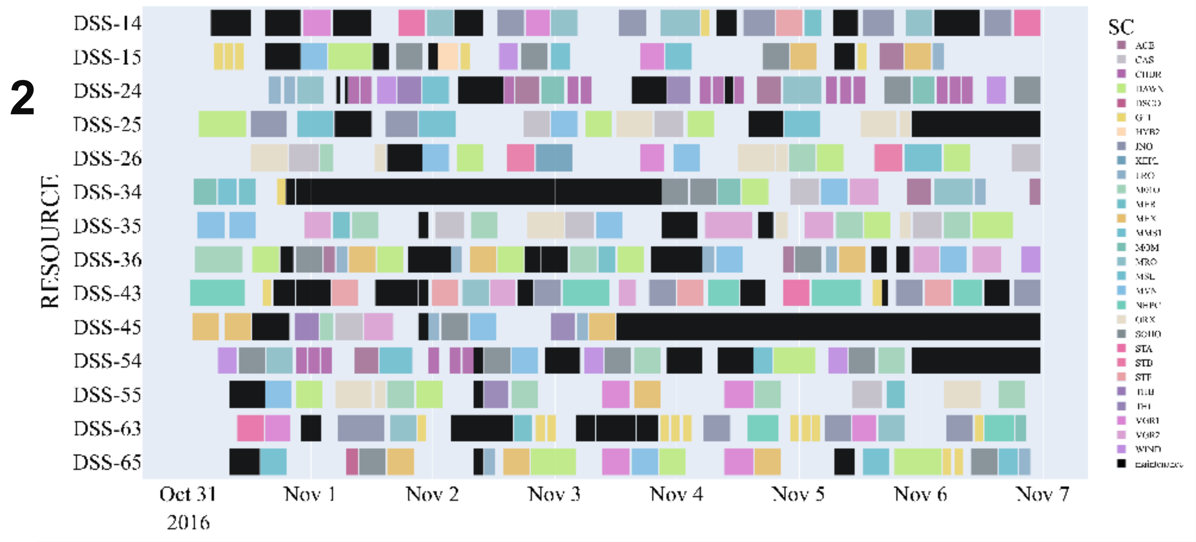

Fig 3

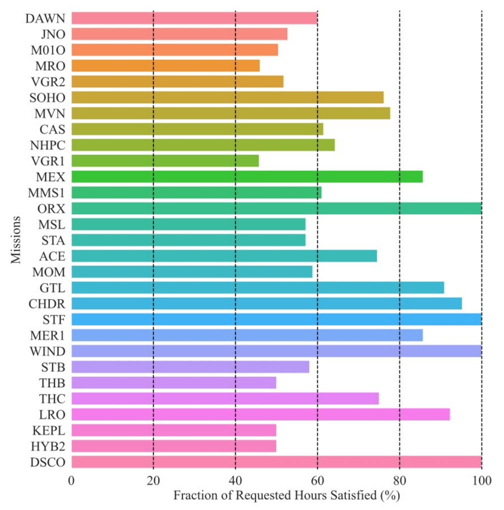

Fig 4

Significance/Benefits to JPL and NASA

There is a revolution going on in optimization prompted by the advent of solving large QUBO problems using the D-Wave quantum annealer and hybrid solver, or other quantum-inspired classical solvers. Deep reinforcement learning has also shown tremendous efficacy for selected optimization problems and games. The performance race between advancing QUBO, QIO, DeepRL and MILP solvers is quite competitive, with papers comparing multiple techniques on the same problems to explore the landscape of their proper use. We found that all three techniques could reasonably solve a synthetic week-long DSN antenna scheduling problem (for the set of constraints that we handled). The team has accumulated a cutting-edge toolbox for solving scheduling and optimization problems: quantum annealing, quantum-inspired solvers or hybrids, deep RL, and advanced MILP solvers.

Quantum Computing is slowly becoming a reality, with potential huge speedups and capabilities to solve otherwise intractable problems. It is time to test its near-term capabilities and map JPL's challenging problems to this potentially revolutionary paradigm. In 2020, the D-Wave quantum annealer scaled from 2000 qubits to 5640 qubits, and the hybrid optimization solver scaled to one million variables. Gate QC's are currently limited to ~100 qubits, but with plans to scale by the end of the decade. We don't want to be surprised as their size and capabilities of QC continue to scale. Potential applications include quantum optimization, quantum chemistry, quantum machine learning, fast eigensolvers, sampling, etc.

National Aeronautics and Space Administration
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California
www.nasa.gov

PI/Task Mgr Contact
Email: Brian.Wilson@jpl.nasa.gov

Publications:
Sabol, A., Alimo, R., Kamangar, F., & Madani, R. (2021). Deep Space Network Scheduling via Mixed-Integer Linear Programming. *IEEE Access*, *9*, 39985-39994.
Sabol, J. A., Alimo, R., Hoffmann, M., Goh, E., Wilson, B., & Johnston, M. (2021). Towards Automated Scheduling of NASA's Deep Space Network: A Mixed Integer Linear Programming Approach. In *AIAA Scitech 2021 Forum* (p. 0667).
Claudet, T., Alimo, R., Goh, E., Johnston, M., Madani R. & Wilson, B. (2021). Scheduling NASA's Deep Space Network with Mixed Integer Linear Programming. Needs to be revised to be accepted in *IEEE Access*.
Guillaume, Alexander, Edwin Y. Goh, Mark D. Johnston, Brian D. Wilson, Anita Ramanan, Frances Tibble, and Brad Lackey, "Deep Space Network scheduling using Quantum Annealing", To be submitted.

Clearance Number:
RPC/JPL Task Number: R19033