National Aeronautics and Space Administration

Automating DSN Scheduling Problems Using Quantum Computing and Deep Reinforcement Learning

PI: Brian D. Wilson (398M)

Co-I's: Alexandre Guillaume (398K), Mark Hoffmann (1762), Shahrouz Ryan Alimo(393K), Jonathon Sabol (393), Mark Johnston (3970), and Chris Mattmann (1761)

Program: Strategic Initiative

Project Objective:

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 over-subscribed and leading to scheduling "conflicts" that must be resolved while maximizing science return and accounting for critical events.

We are pursuing a multi-pronged approach to solving these large and challenging DSN antenna scheduling problems:

- **Formulation as a Quadratic Unconstrained Binary Optimization (QUBO) problem for solution on the D-Wave quantum annealer** (see Figure 1a below for a description).
- **Formulation as an "action space" and reward function for a learning agent to de-conflict the schedule via Deep Reinforcement Learning (DeepRL).**
- **Formulation as a Mixed Integer Linear Programming (MILP) problem for solution by classical optimization solvers or quantum-inspired solvers.**

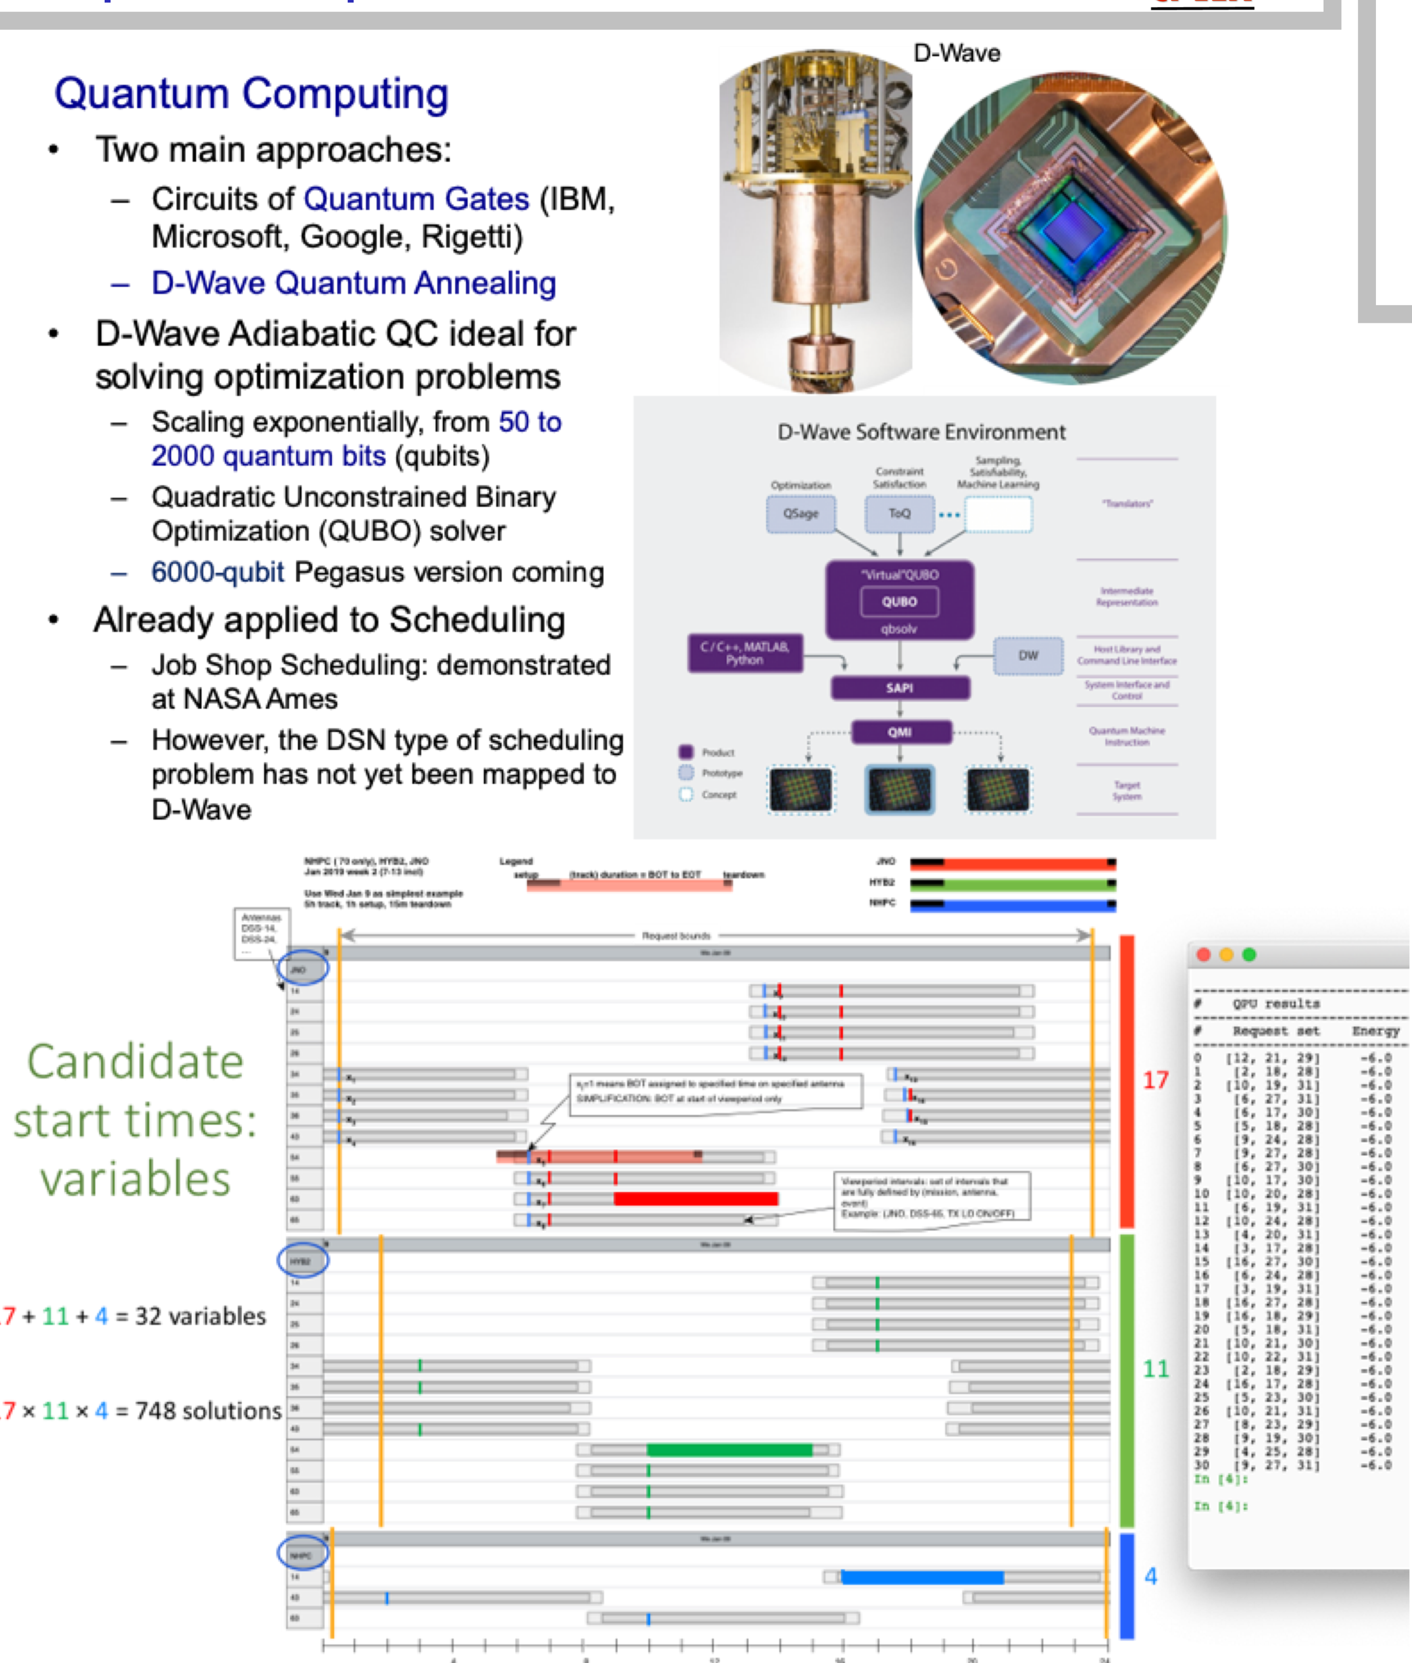

Figure 1. (a) the D-Wave quantum annealer. (b) Schematically depicts a small "synthetic" problem of slotting tracking passes for three spacecraft (red, green, blue) onto antennas at the three complexes. A QUBO problem with 32 variables, 2-hour time granularity, and 748 resulting solutions. On the right, the actual "instantaneous" output from the D-Wave shows the first 30 lowest-energy (unconflicted) solutions.

FY18/19 Results:

Objective 0. Formulated a proposed series of seminars to educate the community about the current state-of-the-art for Quantum Computing (QC) and potential applications at JPL. First invites from OCSCT are out.

Objective 1. Mark Hoffmann developed an API and Python library for preparing and aggregating scheduling datasets (by week): the mission requests (User Loading Profiles or ULP's), the spacecraft view periods from each antenna, and the final (actual) de-conflicted schedule.

Objective 2. Developed first version of a code library to "generate" synthetic but representative scheduling problems, based on an actual week of mission requests and view periods (from the corpus), but with a simplified set of antenna & resource constraints.

Objective 3. Alex Guillaume, the quantum lead, mapped several small synthetic problems to QUBO and successfully computed on the D-Wave QC a set of "optimal" (lowest energy) scheduling solutions. Figure 1b (below left) depicts a "synthetic" problem of slotting tracking passes for three spacecraft (red, green, blue) onto antennas at the three complexes.

Objective 4. Mark Hoffmann developed a DeepRL system, applying the Deep Q-Network (DQN) method to learn action sequences (the policy) for small scheduling problems, now being extended to more realistic action spaces. Figure 2a schematically shows how the Agent learns about the Environment (the scheduling space) as it "acts" to fill and de-conflict the schedule. Figure 2b below shows a snapshot of a partially filled weekly schedule for 22 spacecraft (different colors) onto antennas from the three DSN complexes. Animations of this visualization enable us to examine the sequence of placing and de-conflicting actions taken by the DeepRL agent as the schedule is filled by the set of mission requests.

Objective 5 (added in mid-year). Ryan Alimo and Jonathon Sabol developed a first formulation of "synthetic" scheduling problems as MILP problems and solved 1 and 2-day problems using the AMPL/CPLEX solver. To scale to week-long problems, a hybrid approach combining optimization solvers and deep learning will be investigated. Figure 3 describes the first MILP formulation which uses a "collision distance" concept.

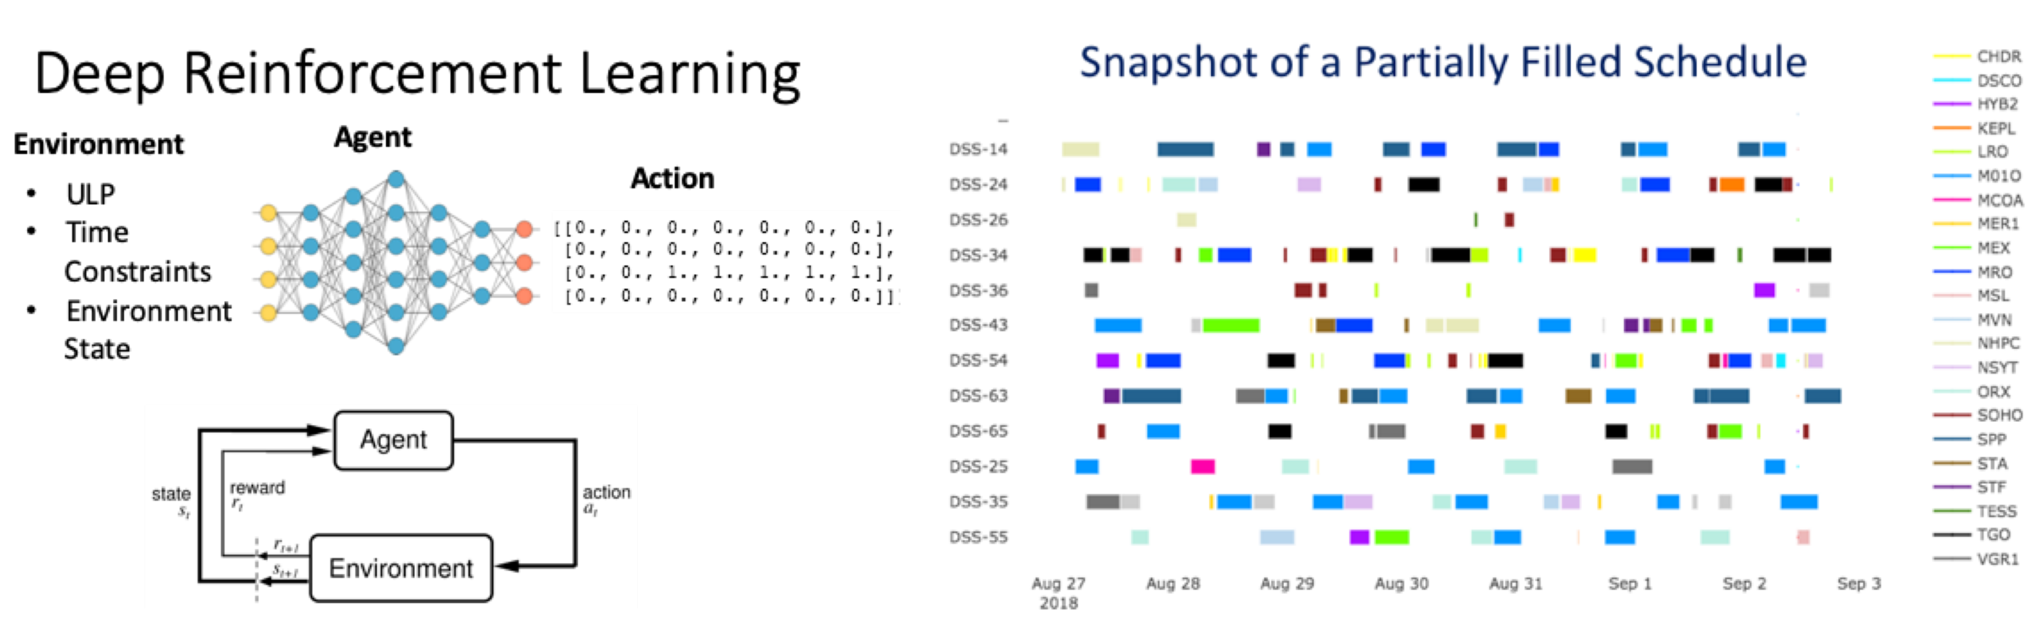

Figure 2. (a) Schematic of how the DeepRL Agent learns about its Environment (the scheduling space). **(b)** A snapshot of a partially filled weekly schedule for 22 spacecraft (different colors) onto antennas from the three DSN complexes. Animations of this visualization enable us to examine the sequence of placing and de-conflicting actions taken by the DeepRL agent as the schedule is filled. The visualizations enable all team members to understand the scheduling outcomes.

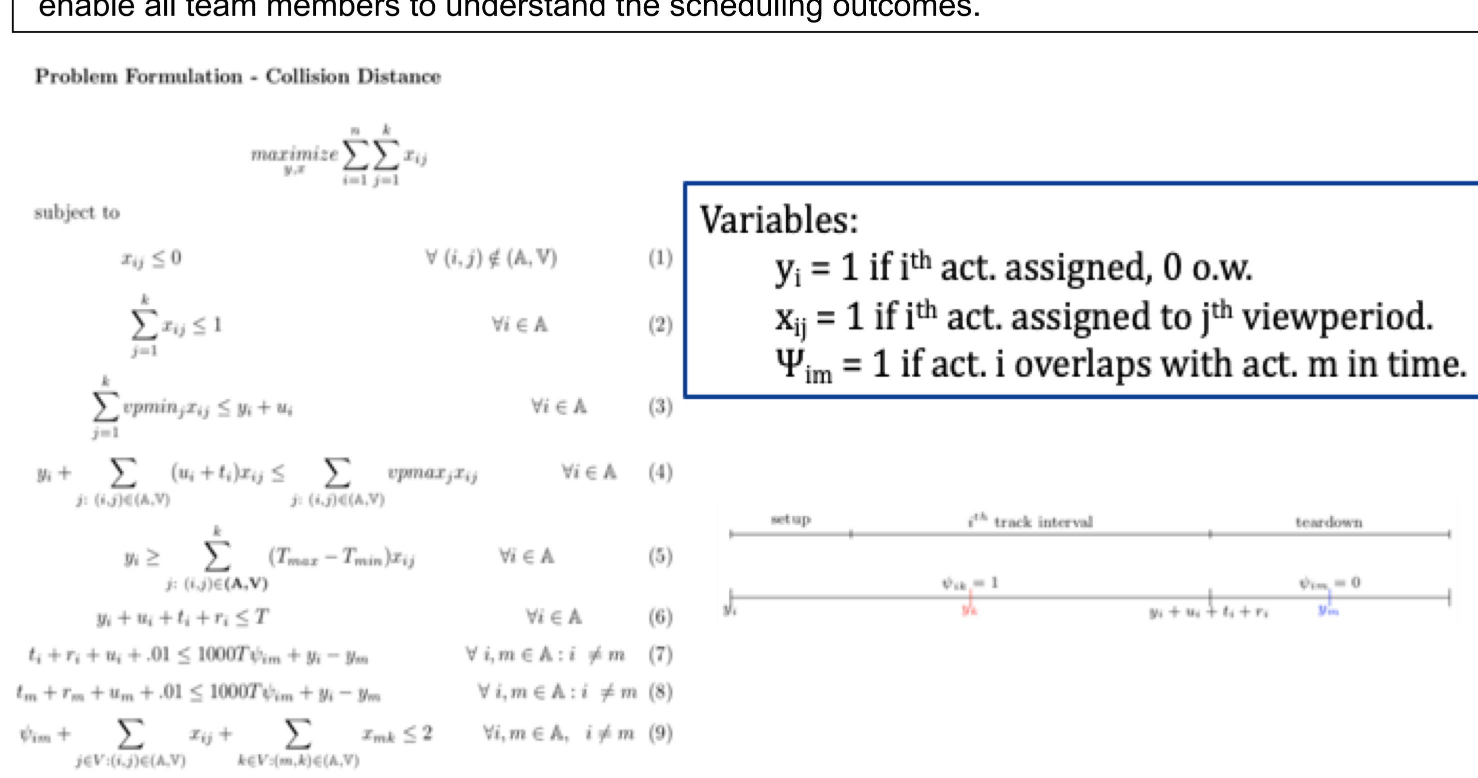

Figure 3 (left). MILP Formulation using a "collision distance" concept, many binary variables and many linear constraints. Small "synthetic" problems were solved using the modern AMPL/CPLEX solver.

Scaling up the problem size leads to thousands of variables that challenge the solvers. To overcome this difficulty, we will investigate a hybrid approach combining optimization solvers and deep learning.

Benefits to NASA and JPL:

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. Currently, priority levels in the existing scheduling tools are under-utilized and human intervention is needed, including negotiation between mission schedulers, to "relax" mission requirements appropriately. This leads to slower cycle times for "de-conflicting" interim and final schedules. Faster iterations and the capability to automatically generate "candidate" DSN schedules will enable JPL mission designers to consider many more "what if" scenarios in their planning processes. There are many other scheduling, planning and optimization problems that are vital to the DSN and JPL, including science activity scheduling, resource scheduling on spacecraft or ground computers, or even project/personnel scheduling. Each of these scheduling problems is unique, and so they are not solvable with a single algorithm. However, a breakthrough in solving one of these challenging optimization problems using QC or DeepRL can then be reapplied, after more algorithmic work, to improve capabilities in many areas.

Quantum Computing is slowly becoming a reality. It is time to test its near-term capabilities and map JPL's challenging problems to this potentially revolutionary paradigm. We don't want to be surprised as the size of QC's in qubits continue to scale. QC speedups, relative to classical computers, have not been demonstrated for any "real" problems yet, but this may change in the next few years. If QC is not yet feasible, then we can still improve DSN scheduling in the near-term using Deep Learning networks and Reinforcement Learning techniques, or classical linear or quadratic solvers.

National Aeronautics and Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

www.nasa.gov

Publications:
No publications submitted in the first year; there will be several in the 2nd year.

PI Brian D. Wilson:
Brian.Wilson@jpl.nasa.gov
818-726-4383

Poster RPC-134