Quantum Computing for Earth Observation (QC4EO) Study

Summary of UC1 Mission Planning for Earth Observation Acquisitions

Mission planning is a pivotal aspect in EO, targeting the optimal scheduling of acquisition requests from end-users. State-of-the-art classical methods leverage deterministic and meta-heuristic algorithms to generate optimal solutions for smaller satellite constellations with thousands of optimization variables. However, as larger constellations featuring numerous satellites ($N \geq 100$) increases, the task of finding optimal solutions becomes more challenging. The difficulty is further intensified by the necessity to account for more complex mission constraints, impacting both the required time and the quality of the solution. Consequently, the concept of quantum advantage should consider not only the computational time but also the quality of the solution obtained. The QC4EO study proposes two distinct approaches to attack this problem: in the first one, the mission planning is cast into an optimization problem, while in the second approach relies on machine learning.

1.1.1 Instance for the quantum optimization approach

The objective of the mission planning is finding the optimal scheduling of satellite observations for a given list of user requests. For each satellite and for each "acquisition request", there are several "data take opportunities" and "downlink opportunities". The solution must return, for each satellite, a time-ordered shortlist of the possible observations and downlinks opportunities taking in consideration several constraints, i.e., on-board available memory, minimum preparation time between two subsequent acquisitions, time ordering of the acquisitions with respect to the downlink opportunities, available battery. We have formulated the problem in terms of a QUBO, where the original constraints now appear as penalty terms in a quadratic problem. We studied the feasibility of a quantum solution with the aim to improve the quality of the optimal solution. Our QUBO formulation is amenable for an implementation both on digital (or general purpose) quantum computer, using the Quantum Approximate Optimization Algorithm (QAOA), and analog specific-purpose quantum computers (or simulators) that are not able to run generic quantum algorithms, but are optimized to solve specific problems, i.e., QUBO problems using quantum annealing algorithm in the case of QA. We underline the importance of proceeding in both directions: digital quantum computers have been recently applied to similar scheduling problems with a reduced set of variables and constraints; on the other hand, quantum annealers (and quantum simulators) can already allocate up to thousands of qubits and have shown promising results in the solution of similar optimization problems.
However, D-Wave states to be able to support real-world size applications with up to 1 million variables and 100,000 constraints via their quantum-classical hybrid solver: obviously, these numbers must be tailored on the specific problems based on the form of the interactions that appear in the cost function (including constraints).

	Problem size	**Hardware requirements**	**Timeline**		
			Up to 5 years	**Up to 10 years**	**Up to 15 years**
Minimum-size problem	2 satellites, 2.000 requests	Analog hardware (superconducting, neutral atoms): ~10^6 qubits (polynomial scaling of qubits). Digital hardware (superconducting): ~10^6 qubits (polynomial scaling of gates), error correction required	Problem implemented on NISQ devices	Problem implemented on NISQ devices	Problem implemented on NISQ devices
Full-size problem	10-100 satellites, 10.000-100.000 requests	Analog hardware (superconducting, neutral atoms): ~10^9 qubits (polynomial scaling of qubits), error correction required. Digital hardware (superconducting): ~10^9 qubits (polynomial scaling of gates), error correction required	No feasible implementation envisioned	No feasible implementation envisioned	Problem implemented if fully scalable error correction on superconducting devices become available

The QUBO approach to a medium real size mission planning problem involves 10^6 binary variables. Although photonic platforms could potentially accommodate such a great number of qubits, we focused on forthcoming platforms such as superconducting qubits (digital quantum computers and quantum annealers), ion-trap digital quantum computers, and arrays of Rydberg atoms confined with optical tweezers. The main bottleneck for an implementation on quantum annealers (such as D-Wave Quantum Systems) is the required number of qubits, which depends both on the number of binary variables and the number of auxiliary qubits to account for the limited connectivity in the hardware. Our study anticipates that, while existing hardware can already manage very small size problems, it is feasible to solve at least the minimum-size problem (2 satellites, 2000 requests) within the 3 to 5 years. Our estimates are based on the D-Wave roadmaps, where they plan to reach more than 7000 qubits by 2024 with improved connectivity. Further, Rydberg simulators appear promising due to their intrinsic flexibility which allows for the adjustment of the connectivity to meet the specific requirements of each problem. On the other hand, the application of digital quantum computers appears feasible in a 5+ years horizon due to the number of qubits and gates required (see table below) even for a small-to-medium size problem. Since QAOA is a hybrid-variational algorithm we expect that the implementation of error mitigation techniques will help the convergence towards optimal results. However, the timeline for solving the full-size problem for both approaches, general purpose quantum computers and quantum annealer-simulators extends beyond 15 years.

1.1.2 Instance for the quantum machine learning approach

In the second approach, the mission planning problem is addressed through a hybrid classical-quantum method based on a Quantum Neural Network (QNN) and a Reinforcement Learning (RL) framework. In this algorithm, an agent interacts with an environment, using a reward function to evaluate and assign values to the agent's actions. The RL agent employs a policy model to determine which actions based on the state of the environment, thus transforming the state itself. The policy model is trained to select actions that maximize the probability of achieving a positive reward. We considered the policy model as a parameterized quantum circuit to establish the input-output relation of the model and apply it within the mission planning context. This quantum subroutine follows a classical pre-processing step based on classical neural network to reduce the number of relevant features, which serve as input for the QNN. This approach has demonstrated promising results for small-case instance and is viable for implementation on a general-purpose digital quantum computer. Specifically, the proposed quantum circuit requires few qubits (the exact count depends on the number of input data features) and few layers, yielding a number of logical gates that is linear with the input size. In previous work, a circuit with 4 qubits (thus encoding 32 features) and 8 layers, resulting in the implementation of more than 100 logical gates, was proposed to solve a small-medium size problem. This narrow quantum circuit can be already implemented on current state-of-the-art quantum hardware technology without full error-correction, such as IBM quantum computer based on superconducting qubits. Furthermore, IBM roadmap outlines the potential for hardware with over 4000 qubits by 2025, with plans for scaling to 10k-100k error-corrected qubits beyond 2026. This makes the implementation of the QML approach for the mission planning problem feasible within the next 15 years, even for full-size scenarios. The primary challenges in this algorithm stem from the error rates of quantum gates, which are device-dependent and significantly impact the quality of the obtained solution. In this context, it is worth to highlight another promising platform: trapped-ion based quantum computers. Despite having a smaller number of currently available qubits, these machines feature longer coherence times and notably lower error rates, often orders of magnitude below those of IBM processors. Furthermore, this technology provides enhanced connectivity compared to superconducting architectures. A notable player in this domain is the American public company IonQ, which has set ambitious plans to achieve 1024 qubits by 2028.

	Problem size	Hardware requirements	Timeline		
			Up to 5 years	Up to 10 years	Up to 15 years
Minimum-size problem	2 satellites, 2.000 requests	Accuracy related to number of features. Experimental demonstration with 32 features. Hardware requirements scale linearly: 2^f features encoded using 2^q qubits and 2^l layers with $f = q + l$.	Problem implemented on NISQ devices	Problem implemented on NISQ devices	Problem implemented on NISQ devices
Full-size problem	10-100 satellites, 10.000-100.000 requests	Accuracy related to number of features. Hardware requirements scale linearly: 2^f features encoded using 2^q qubits and 2^l layers with $f = q + l$.	No feasible implementation envisioned	No feasible implementation envisioned	Problem implemented if fully scalable error correction on superconducting devices becomes available