Federated Scheduling of Model-Driven Observations for Earth Science

A. Branch[1], S. Chien[1], Y. Marchetti[1], H. Su[1], L. Wu[1], J. Montgomery[1], M. Johnson[1], B. Smith[1], L. Mandrake[1], P. Tavallali[1]

[1]Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive Pasadena, CA 91109
firstname.lastname@jpl.nasa.gov

Introduction

Comprehensive sensing of many Earth science phenomena, such as hurricanes and extreme weather events, is challenging due to the large spatial scales and complex dynamics. Adaptive sensing is a technique that assists in studying these events by utilizing online analysis of the phenomena to determine the locations most advantageous for sensing. Focused observations are made such that the information gain is maximized. This new data is incorporated into the model and the cycle can repeat. This technique is illustrated in Figure 1.

We are currently undergoing an effort to apply this adaptive sensing concept to study hurricanes/typhoons (Tavallai et al 2020). In this effort, we simulate a storm and use these simulations to formulate a utility function that estimates the relative knowledge gain from a set of observations (Torn and Hakim 2008). An observation scheduler can then be used to request observations from assets such that utility is maximized. These observations are taken and the cycle repeats.

When studying these complex Earth science phenomena, it is usually the case that a set of assets are required, with different sensing capabilities. Assets that are often under the control of external organizations. These organizations all have unique systems and requirements in place to produce observation schedules for their assets. An interface may be made available to external researchers to request specific observations. Depending on a number of factors including available resources and contention from other requests, these observations may or may not be fulfilled. Additionally, requested observations may have an associated cost to consider. This naturally leads to the problem of requesting observations from a set of assets where the goal is to maximize the utility of the fulfilled observations while minimizing the cost of those observations. Here we describe this, the federated scheduling problem.
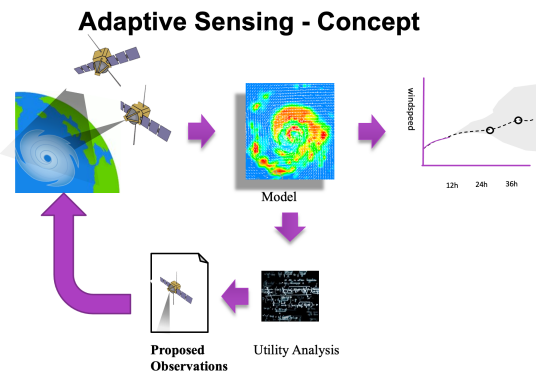

Figure 1: Adaptive Sensing Concept

Problem Statement

In the federated scheduling problem, inputs are the utility function that we want to maximize, black boxes that represent each asset, and the orbits, specifications, and observation cost models of those assets. Each black box accepts as inputs a set of requested observations and outputs the observations which are planned, and eventually the observations that were fulfilled. As these processes are not instantaneous, we assume that each black box will have a response time on the order of hours, preventing real-time search. In application these black boxes would be other organizations. Our goal is to

request observations that maximize the utility of the fulfilled observations while minimizing the cost of those observations. This problem is illustrated in Figure 2.

The form that requested observations take is dependent on specific assets. Most commonly they will be either a point or polygon, sometimes with an associated static priority. In our initial testing we have seen that regions of semi-homogeneous utility are common. In these scenarios it may be more advantageous to request any observations in a given region rather than a specific observation. This approach would likely result in more fulfilled observations by providing flexibility to the black box scheduler.
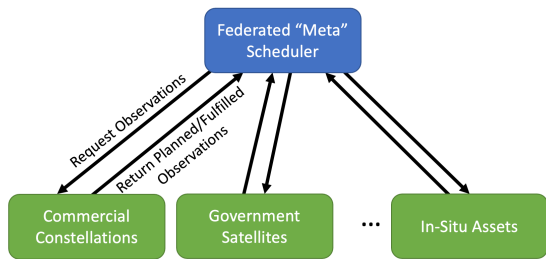

Figure 2: Federated Scheduling Problem Overview

Proposed Approaches

This project is in the early prototyping stages and is starting development of methods to address the federated scheduling problem presented here.

A simple baseline approach is to request observations at all locations in the study area, filtered by the visibility of a given asset. The visibility of an asset, or the set of all potential observations, can be calculated via the asset orbit and specification information. This approach does not consider the utility of the requested observations and will potentially result in repeat observations. For these reasons we would expect that the performance would decrease as the probability of request fulfillment decreases.

For more sophisticated approaches probabilistic models of observation fulfillment for each asset would be beneficial. Using these models one would be able to calculate the expected utility for a set of requests. This expected utility could then be maximized and cost minimized. Asset visibility would also be utilized here to prevent requests of completely infeasible observations. There are a number of complications to resolve in using this general approach, such as the construction of probabilistic models and efficiently calculating expected utility with non-linear utility functions.

Through this project we will be exploring these approaches and others. In our testing we will simulate operations of a variety of assets, including those from commercial providers, government agencies, and in-situ assets such as autonomous aerial vehicles. Enabling adaptive sampling techniques using a wide variety of assets and interfacing to existing systems would significantly increase our capabilities of studying these large scale Earth Science phenomena.

Related Work

This project builds on prior work in which satellite observations were directed based on alerts from other assets to track flooding (Chien et al 2020) and volcanic activity (Chien et al 2019). However, that prior work focused on direct scheduling of observations, not this federated problem. Additionally, they did not estimate the utility of observations or attempt to maximize the utility gain, instead using event → observation triggers. Fioretto (2018) provides a survey of the general multi-agent constraint optimization problem, but does not address the specific problem proposed here. Clement and Barrett (2003) develops a method of coordination of activities between self-interested

assets/entities. This work differs in that it does not consider the utility maximization component and does not use a single entity requesting observations from assets with no inter-asset communication. Chien et al (2000) outlines three general frameworks for planning across multiple assets. This lays the groundwork for the type of distributed planning required in the federated scheduling problem, but does not directly address the problem of distributing goals (requested observations) to maximize a non-linear utility function.

Acknowledgements

References

Chien, S. et al. A comparison of coordinated planning methods for cooperating rovers. Proceedings of the fourth international conference on Autonomous agents. 100-101. 2000.

Chien, S. et al. Automated Volcano Monitoring Using Multiple Space and Ground Sensors. Jnl Aerospace Information Systems, 17:4: 214-228. 2020.

Chien, S. et al. Using Taskable Remote Sensing in a Sensor Web for Thailand Flood Monitoring. Jnl Aerospace Information Systems, 16(3): 107-119. 2019.

Clement, B. J. and Barrett, A. C.. Continual coordination through shared activities. In Proceedings of the second international joint conference on Autonomous agents and multiagent systems, 57-64. 2003.

Fioretto, F., Pontelli, E., and Yeoh, W. Distributed constraint optimization problems and applications: A survey. Journal of Artificial Intelligence Research 61, 623-698. 2018

Tavallai, P. et al. Adaptive, Model-Driven Observations for Earth Science. International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS 2020). 2020.

Torn, R. D. and Hakim, G. J. Ensemble-based sensitivity analysis. Monthly Weather Review, 136(2), 663-677. 2008.