

Artificial Intelligence enabled Dynamic Coalition Architecture for Space Traffic Management

**W. Thomas Vestrand
Przemek Wozniak
Sean Brennan
Troy McVay
Lucas Parker
Rebecca Holmes Sandoval
Yancey Sechrest**

*Los Alamos National Laboratory
P.O. Box 1663
Los Alamos, NM 87545*

ABSTRACT

We have now entered a new space age—where technical barriers and cost of access to space are rapidly decreasing. The easier accessibility is fostering a new commercial entrepreneurialism that is excited about using space-based systems to transform a broad range of ground-based enterprises. This rapidly increasing pace of international space activity is also generating new challenges for flight safety and Space Traffic Management (STM) in near-Earth orbit. One of the key STM challenges is developing a reliable sensor ecosystem that can persistently monitor space traffic to provide accurate measurements and actionable information on traffic anomalies in real time. A consequence of this persistence requirement is the need to globally distribute the sensors, creating a strong demand for new approaches to international STM cooperation. We present a concept for a collaborative STM architecture based on the Dynamic Coalition Architecture (DCA) designed to enable coordinated space surveillance by independent organizations. The DCA employs Artificial Intelligence (AI) concepts drawn from distributed problem solving and multi-agent system research to organize sensors into an efficient sensing ecosystem, and allow entities with different priorities to pool assets on case-by-case basis, while ensuring that the autonomous collection provides reliable measurements. We describe a simulation environment that we are developing to explore the DCA concept employing Case-Based Reasoning to coordinate real-time follow-up of emergent STM anomalies. We present the first results from our simulations and briefly discuss issues associated with sensor coalition formation and management such as the real-time negotiation, modeling sensor effectiveness, and autonomous collaborative space surveillance.

1. INTRODUCTION

Economic opportunities and new space technology, which is lowering the barrier of entry for access to near-Earth orbit, are generating an accelerating pace of international space activity and a rapidly expanding, global, commercial space industry. The task of tracking and managing this burgeoning international space traffic is now starting to challenge traditional approaches and legacy systems. Modernization of real-time Space Situational Awareness (SSA) capabilities and the establishment of new Space Traffic Management (STM) norms and protocols are therefore becoming essential for both flight safety and the future economic vitality of the space enterprise.

The United States Congress tasked the National Academy of Public Administration to conduct a study of STM requirements, independently review the transfer of STM functions from the Department of Defense to a civilian agency, and make recommendations on how the US government should address the emerging STM challenges. The National Academy study panel recently released their report which included some important findings [1]. First, that the collection of reliable, timely, space surveillance data is foundational for the STM task and, as such, “SSA and

STM should be combined and conceptualized as an ecosystem”. Second, that while it makes sense to have the Department of Commerce lead the US civilian effort, multiple government “agencies must continue to work collaboratively now, and in the future, to achieve a safer space domain” [1]. They also concluded that “international coordination and cooperation are essential” for enabling a successful SSA/STM mission. However, if the pace of Space activity continues to accelerate as forecasted, effective coordination will require a globally distributed ecosystem of instruments capable of providing reliable, real-time, space surveillance measurements.

Most of us have experienced the new capability for real-time, ground-based, automobile traffic management provided by our smart phones. This technology helps us efficiently get from point A to point B by providing an autonomous sensing ecosystem that collects real-time GPS tracking data from the phones, relays the data to a cloud-based traffic estimation model, and broadcasts back to each sensor (cell phone) in real time a suggested optimal course of action. The success of that system and the Internet of Things (IoT) concept is stimulating a substantial interest from high-tech companies in autonomous sensor ecosystems.

For real-time space traffic management, the large data volume associated with, e.g., raw imagery from space surveillance sensors and the typically very limited communication bandwidth recommends an *Edge Computing* approach. In this framework, the bulk of the complex sensor data are analyzed near the sensors (at the edge) to distill knowledge before it is transmitted to another location. We refer to this locally extracted knowledge as *sensed information*, which is distinct from the raw sensor data that would otherwise have to be transferred to the cloud for analysis. In the context of the autonomous sensing ecosystems operating on a global scale, this approach allows complex data-intensive sensors to be included in the ecosystem, promotes faster global situational awareness when the movement of information is throttled by the network band-width, and facilitates faster real-time follow-up of important events and anomalies. It also allows knowledge of the raw data, sensor capabilities and analysis techniques to remain proprietary and be controlled by the sensor owner. We assume that the contributing sensors in the STM ecosystem are all capable of generating and sharing sensed information.

In this paper we describe our work to apply concepts from Artificial Intelligence research to a simulated global ecosystem of sensors that coordinates international assets conducting real-time sky surveillance to enable effective STM. Section 2 of the paper describes system components, the simulation environment, and the relevant AI concepts. In section 3, we describe the first results from the simulation tools for quantifying the utility of our approach for STM. Finally, in section 4 we briefly discuss directions for future work.

2. SPACE TRAFFIC MANAGEMENT ECOSYSTEM: SOLUTION COMPONENTS

2.1 Multi-Agent AI-Enabled Architecture

A core component of our STM concept is a new architecture employing a distributed network of *sensor agents* that we call a Dynamic Coalition Architecture (DCA) [2]. Figure 1 shows a schematic view of the proposed system. Here we use the term *agent* to denote a computer system, a software machine, capable of independent (autonomous) action on behalf of its user or owner. The DCA concept employs ideas originally developed for *spatially distributed* Artificial Intelligence (AI) systems; where most solutions fit on a continuum between the extremes of distributed problem solving and multi-agent systems. The distributed problem solving approach [3] divides a complex task into smaller units and distributes the load among identical workers, which are spatially distributed, but centrally configured and controlled (command-and-control nodes). At the other end of the spectrum is the multi-agent system approach [4], which maps the problem to a set of autonomous intelligent agents with heterogeneous capabilities working on various parts of the same problem, but with no central control. Here the autonomous agents are self-motivated and act only according to their own success criteria—which can generate inefficiencies in overall system performance. A collaborative behavior tends to emerge, when none of the agents can accomplish its goals without some help from other agents. Our DCA concept blends these well-established AI methodologies to create an architecture that maintains multi-agent autonomy, while adding elements of distributed problem solving and partially centralized control to efficiently collect knowledge and enable effective real-time space traffic management.

The key feature of DCA is the creation of temporary coordinated coalitions for each STM anomaly or data collection event in real time that promotes much more efficient use of resources when and where they are needed. Data collection coalitions are formed and managed through real-time negotiation between agents in a globally distributed sensor

ecosystem and the central authority known as the *Coalition Manager*. The role of the Coalition Manager is to monitor activity, assemble and coordinate effective coalitions, and facilitate voluntary information sharing in order to maximize the overall system utility. The negotiation acceptance criteria for each agent are private and owned, controlled, and configured by the sensor owner. The sensor agent can then opt in or opt out of a potential coalition in real time based on the nature of the event. This feature of the DCA approach allows the agent owners/schedulers to accomplish their primary mission goals and only join new STM coalitions when time is available or when the event of interest is so important that they are willing to temporarily drop their current task. The sensor agent approach therefore enables the rapid formation of coalitions that include sensors in the ecosystem that are unwilling to be fully dedicated command-and-control sensors, but are willing to contribute under the terms that they control. While designed for autonomous sensors, in this architecture the coalition agents are not limited to autonomous robotic sensors. If the natural timescale of STM collection event is long, the coalition can also include agents which encapsulate human knowledge and query experts to provide additional and/or more refined information to further optimize response.

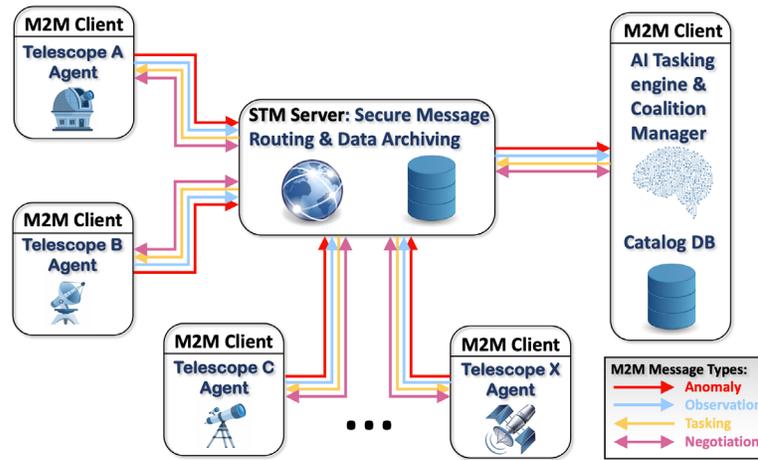


Fig. 1. The DCA concept employs a multi-agent system approach, machine-to-machine communications, and AI-based sensor tasking to coordinate STM space surveillance on a global scale. The role of the Coalition Manager is to assemble and coordinate temporary data collections involving multiple autonomous sensors. Each emerging anomaly or event of interest offers an opportunity for sensors with relevant capabilities to participate in a corresponding coalition, but there is no obligation. The messaging server provides secure message delivery for interacting agents and archives all message traffic in the STM sensor ecosystem.

The details of the coalition formation mechanism will likely evolve in the future to match specific needs. However, the basic logical flow of this step in the coalition management cycle is relatively independent of such details and is captured in Fig. 2 showing the initiation of a coordinated coalition for a hypothetical STM anomaly of Type N. Each invited sensor agent autonomously determines whether or not this new Type N event meets their currently criteria for participation and sends back a *yes* or *no* reply indicating the intent to join the coalition. The coalition manager then collects responses from potential participants and determines if there is sufficient interest to merit a coordinated follow-up. This decision must be made before the end of the opportunity window, which depends on the type of the triggering event. In case there is not enough interest to commit resources, the potential coalition is canceled and agents that did express interest are notified. But if sufficient interest is present, the start of the coalition is announced and participating agents are assigned roles, i.e., what to observe, how, and when. The role assignments reflect corresponding sensor capabilities and data collection needs for the characteristic physics signature of a Type N STM event. This is determined by the Case-Based Reasoning (CBR) engine by a comparison with similar follow-up cases (see the discussion of CBR in section 2.4).

2.2 STM Agent Communication

Fast and effective machine-to-machine (M2M) communication between autonomous agents is essential for the operation of the STM space surveillance system proposed here. Agents must understand each other and at the same time their activity should be transparent to humans. To this end LANL has developed a text markup language called SatChat ML, a dialect of XML (extensible markup language). This language essentially annotates a document to describe its structure in a way that allows the text to be parse and understood by both STM agents (software machines) and humans, as shown in Fig. 4. The SatChat standard defines an *XML schema*, which describes the grammar and the vocabulary of all possible constructs that are allowed in a valid SatChat message. In other words, the SatChat schema provides a simple Ontology for M2M messaging. The communication layer includes state-of-the art encryption that is transparent to the users. Each registered agent is issued a permanent universally unique identifier (UUID), essentially an address, to receive tailored content and track the origin of SatChat messages. The messages themselves also carry their own UUIDs to enable message cross-referencing and sophisticated patterns of communication. The message types and ontology enable a broad range of content (e.g. anomaly reporting, agent tasking and negotiation, sharing observations) that is tailored SSA/STM applications.

```
<AstrometricDeviation>
  <Object uuid="90d1009c-c2ec-58bc-b6cc-c9a85d2b217d">
    <Name>CZ-4C R/B</Name>
    <Type>space object</Type>
    <Catalog>USSPACECOM</Catalog>
    <CatalogNumber>40341</CatalogNumber>
    <Description>#40341 CZ-4C R/B</Description>
  </Object>
  <Orbit>
    <TLE title="CZ-4C R/B" >
      <Line1>1 40341U 140800 17312.65815385 -.00000092 +00000-0 +39421-5 0 9996</Line1>
      <Line2>2 40341 063.4602 359.6087 0233963 005.4015 354.9465 13.68990302145551</Line2>
    </TLE>
  </Orbit>
  <ExpectedPosition>
    <RA>116.08213355866314</RA>
    <Dec>8.703089660800005</Dec>
  </ExpectedPosition>
  <MeasuredPosition>
    <RA>116.15483333333333</RA>
    <Dec>8.57546</Dec>
  </MeasuredPosition>
  <Probability>0.01</Probability>
</AstrometricDeviation>
```

Fig. 4. A snippet from a SatChat ML message. While designed for efficient machine-to-machine (M2M) communication between STM agents, SatChat ML messages are also human readable. In this example, the message communicates a measured deviation from the expected position of a rocket body.

The SatChat Message Passing Layer package is written in object-oriented Python and provides a server/client communication architecture that utilizes the ZeroMQ Application Programming Interface (API) for the management of communication sockets. This provides the capability to pass XML message traffic amongst the ecosystem agents with minimal latency as well as network resiliency through self-healing features such as automatic socket reconnection. The software also provides cybersecurity functionality, including IP whitelisting and passing of encrypted traffic.

The logical structure of the SatChat communication network utilizes a star topology (see Fig. 1.). The central hub is the STM Server and all STM agents participating in the ecosystem network send and receive XML messages via the STM Server. At connection and registration with the STM Server, an agent subscribes to specific message types, with the default being to receive all messages. The agent also specifies which Author UUIDs (i.e. other agents in the STM network) it wishes to receive messages from, with the default being to receive messages from all other Author UUIDs. All STM agents also must follow the same cybersecurity posture as the STM Server (enabled/disabled elliptic curve encryption).

2.3 Agent Negotiation

The central idea in the Dynamic Coalition Architecture is to use autonomous real-time negotiation to form voluntary partnerships (that are easy to opt in and out of) between assets controlled by international government organizations, independent commercial providers, and academic institutions. When applied to an STM ecosystem with a global

footprint, the DCA can optimize the collection of space surveillance data for routine STM monitoring. But the DCA approach really excels in its capability to resolve time critical STM anomalies such as dangerous satellite conjunctions. As they emerge, each STM anomaly is spawned as a potential coalition and is distributed by a coalition manager in real time to a collection of potential participants with an invitation to join the anomaly resolution coalition. Some of the key information included in the coalition participation request is a characterization of the anomaly with a measure of importance for the corresponding collection request, estimated confidence level for the existence of the anomaly, and the spatial and temporal windows of opportunity for a useful collection. The system then negotiates with the agents to collect the needed anomaly resolution assets. At any given time, the current negotiation priorities for each agent are private and owned by the agent.

The autonomous process for real-time negotiation is shown in the protocol timeline displayed in Fig. 5. The time progresses along the vertical axis and horizontal lines represent messages passed between agents via the server and offset vertical lines depict different computation threads. The STM *negotiation sub-package* handles all but one of the protocol function calls which are on the outer edges of the timeline, along with the negotiation phase messages shown in italics. Messages in bold are standard agent message types. The negotiation protocol always begins with a triggering STM anomaly event or collection request. A coalition announcement is then sent to a subset of agents, to which they may respond with *Accept* or *Reject*. Agents that accept are sent a *Task* message and begin observing. However, agents can short-circuit this step by responding to the announcement with the rapid attribute.

Observations are collected in one of two modes: either immediately sending *Observation* XMLs as they are generated then calling *retask* to continue at the end of the user's process method (which will be called again immediately), or collecting Observation XMLs in the user's process method and sending all of them upon task completion (see *Long-running Event Response* in Figure 5). The first mode (*Punctuated Event Response*) is declared to the coalition manager by responding to the announcement with the punctuated attribute. The second mode is normally not employed due to inherent scaling problems, and the larger coordination subsystem will enforce its rarity.

The coalition manager periodically requests status to which the responder autonomously replies. This status is that which is reported by the operating system for the local process doing the work (e.g. *sleeping, running, dead, etc.*). Reconfiguration is accomplished by the coalition manager with a new but related announcement (the *Reconfig* message) with an altered task. Agents are free to refuse further participation in their responses. Otherwise, running tasks are abandoned and the new task begun.

At the responder agent, tasks timeout automatically. At the coalition manager, coalition formation itself has a time limit, and the subsequent coalition also has a set lifetime. A coalition that completes its natural life-cycle ceases when its constituent agents finish their tasks, i.e. there is no explicit halting command. However, the coalition manager can cancel the coalition early through the *Stop* message.

2.4 Artificial Intelligence directed Follow-Up

An experimental aspect of our effort is the use of Artificial Intelligence to help shape the surveillance data collection and continuously improve the collection capability of the autonomous STM system. The AI technique that we are building into the DCA simulation environment is called Case-Based Reasoning (CBR) [5-7] with Learning. CBR is a well-established AI technique that was inspired by human reasoning and learning. Humans solve problems by applying previous experience (our own or learned) adapted to the current situation. The CBR process [7] tries to mimic that approach with four main steps: retrieve, reuse, revise and retain as shown graphically in Fig. 6. In our application of CBR, each type of STM anomaly (or routine collection scenario) is mapped to a case which contains a follow-up execution plan describing what measurements need to be made and when they need to be made to collect the required space surveillance data. These case "recipes" for STM collection are generated based on: (1) historical, successful, human in-the-loop responses, (2) previously successful automated collections and (3) putative new cases that are built using subject matter experts. The STM system then learns and refines the follow-up by evaluating the quality of the response after each executed coalition and saves successful responses as new cases. This provides a natural way for the system to bootstrap and to learn from experience even when the events of interest are poorly understood. CBR has been used very successfully in the Health Sciences [6] and can work well in autonomous systems that can outpace systems with human operators in the response loop.

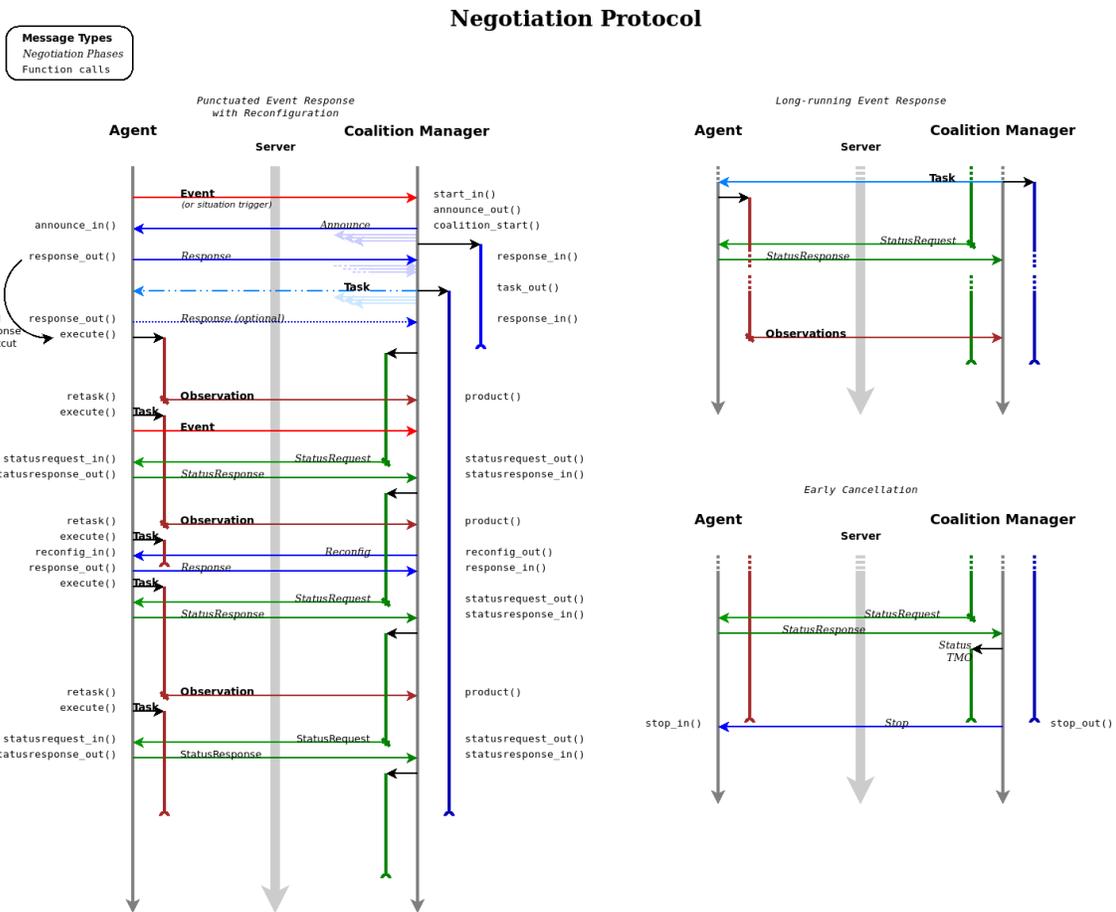


Fig. 5. The local flow for autonomous, real-time, negotiation between the coalition manager and STM sensor agents. Each sensor agent accepts or rejects membership in STM collection coalitions based on negotiation rules configured and privately held by the sensor owner.

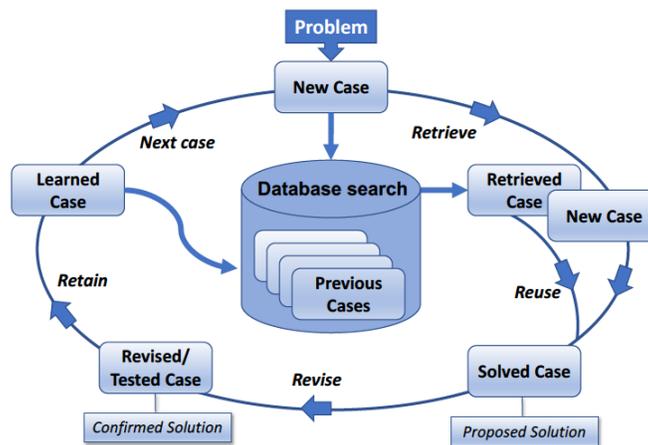


Fig. 6. The CBR learning cycle has four steps: (1) retrieve the case most like the current problem, (2) reuse that case as a proposed solution to address the current problem, (3) assess the outcome and revise if necessary, and (4) retain the revised case as a learned case for future application. This figure is based on the illustration presented in Aamodt and Plaza [7].

The challenge for efficient and reliable collection by an integrated autonomous STM ecosystem is to make the right measurements, at the right times, in the right places. So, after successfully negotiating for access to agent resources, the coalition manager compares of the optimal response, as determined by CBR, with the set of available ecosystem resources. This allows the coalition management agent to: (1) determine if the response can capture the essential measurements and (2) assign roles to the correct agents for execution when the sensor resources are available. The ranking of agents for role assignment is based on geographic location, basic capabilities (e.g. telescope sensitivity, field-of-view, field-of-regard, slewing speed and astrometric accuracy) and potential real-time reconfiguration capabilities (e.g. frame rates, color filters).

Even if the sensors have identical capabilities on paper, factors such as weather at the sensor location can affect success rate and quality of the collection results. We are therefore developing and experimenting with ideas for the measurement of agent effectiveness. Some of the agent effectiveness metrics that we will be exploring include: Measurement accuracy—Are the measurements consistent with test calibrators?; Connected—How often is the agent on-line and responding in a timely manner?; Dependable—How often does the agent successfully collect when assigned a role in the coalition?; and Capability—Is the agent capable of delivering what is needed?

2.5 Visualization Dashboards

A suite of interactive visualization tools have been developed to understand behavior of the ecosystem, both in "real time" as the system is operating, and "play back" visualizations to analyze activity over past time spans. The visualizations are browser based: the front-end runs in a web browser and can be viewed from any machine with a connection to the visualization server. The visualization back-end consists of both a Bokeh server and a Flask server, which access system information through queries to a Postgres database maintained by the central STM server. The browser-based visualizations are interactive, and enable a user to "drill down" to understand system behavior. For example, graphical elements that represent a message can be clicked, resulting in a display of the message's xml contents.

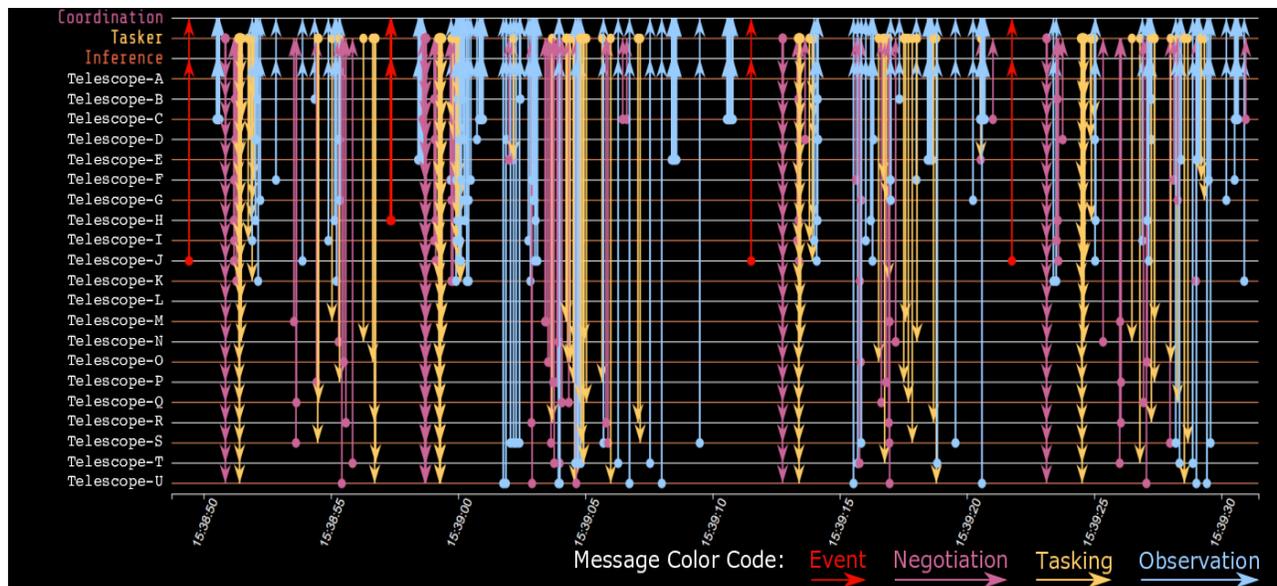


Fig. 7. The live dashboard for the STM simulation environment allows one to track the passage of messages in the ecosystem and visualize patterns of agent interaction. Each arrow is color coded to denote the type of message and can be “clicked on” to bring up a new screen that displaying the message content.

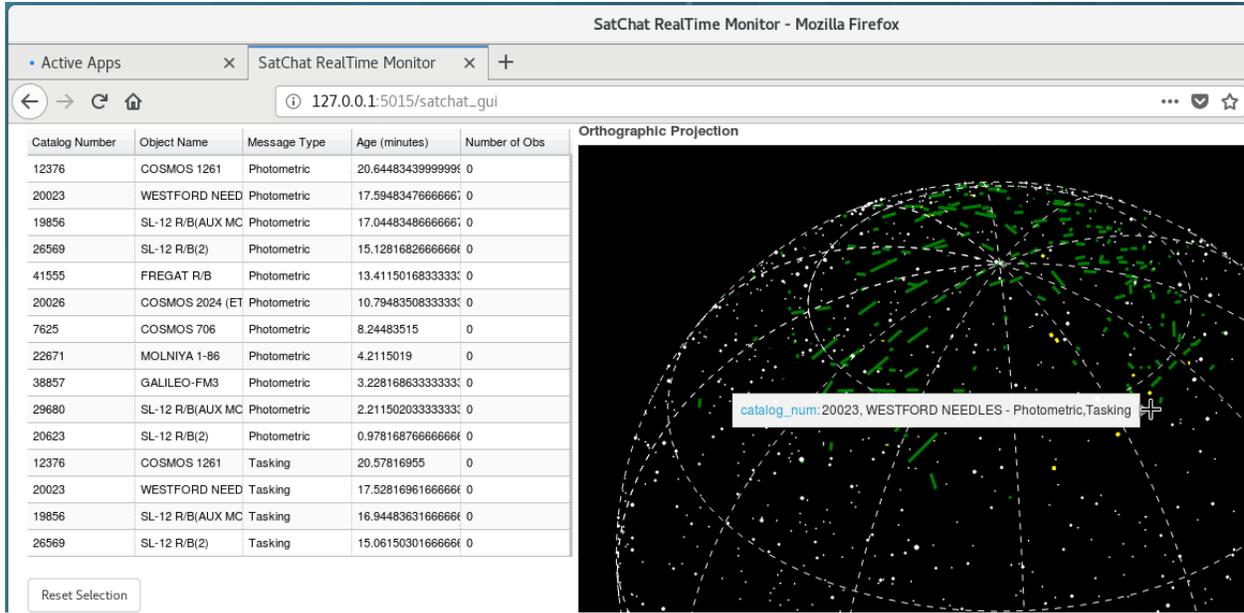


Fig. 8. The real-time dashboard for STM sensor agent collections. This screen allows one to monitor STM anomalies and associate them with named objects.

3 SIMULATIONS AND FIRST RESULTS

3.1 Space Traffic Management Test Ecosystem

Our first instantiation of the DCA-enabled STM testbed assumes the existence of a network of sensor agents deployed at twelve geographic locations around the globe (see fig. 8). For simplicity we have also assumed that all of the STM sensors are optical telescopes, but we acknowledge that ground-based radio frequency systems and space-based sensors are likely to play an important role in space traffic management. All of the selected sensor locations are established astronomical sites that currently host operating optical telescopes that routinely share observations with the international astronomical community. At each site, we assume three classes of optical instruments are potentially operating: (1) a full-sky persistent monitor; (2) a more sensitive, wide-field, patrol instrument; and (3) a sensitive, narrow-field, telescope on a fast slewing mount that can quickly follow up STM anomaly alerts. Finally, we assume that the instruments are running real-time astrometric and photometric analysis pipelines that extract sensed information there at the site or, as we described in section 2, at the “edge”.

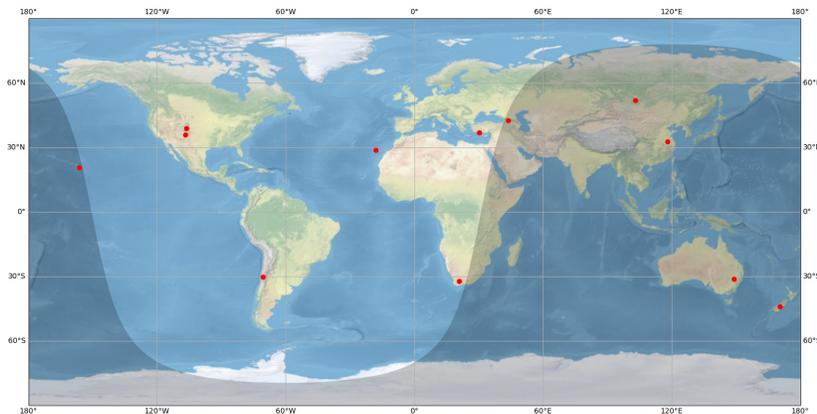


Fig. 9. The geographic locations used for observatory sites in our STM simulation. All of the sites have operating robotic telescopes that conduct follow-up observations of astrophysical transients.

To provide realistic representations of capability for each of the telescope classes employed in the simulation environment, we use actual values achieved by our RAPTOR (*RAPid Telescopes for Optical Response*) telescopes deployed at the Fenton Hill Observatory Site in Northern New Mexico as part of the Thinking Telescopes project [8,9] operated by Los Alamos National Laboratory. The wide-field patrol system, RAPTOR-P (fig. 10., left panel), employs eight Canon 200mm F1.8 lenses with Apogee U-10 cameras. The cameras employ a 2Kx2K Thomson front illuminated CCD (Charge Coupled Device) with 14 micron pixels. Altogether the array provides a Field-of-View (FoV) that covers ~500 sq-degrees in the night sky. That wide FoV and the array's deployment on a fast-slewing mount means that it can patrol the full sky--using back-to-back 30 second exposures for sensitive imaging differencing--in less than an hour. The typical sensitivity this system achieves is a 3σ -sensitivity of $R\sim 16^{\text{th}}$ magnitude for 30 second exposures of sidereal tracked stars. The narrow-field, fast follow-up telescope, RAPTOR-S (fig. 10., right panel), has a 0.41-meter mirror and is also deployed on a fast slewing mount that allows it to slew anywhere in the visible sky in less than ten seconds and begin follow-up imaging. It employs a camera with a 1Kx1K E2V back-illuminated CCD chip that achieves 3σ -sensitivity limiting magnitude of $R\sim 18.5$ in 10 second exposures of tracked stars.



Fig. 10. The left hand panel shows a wide-field patrol telescope called RAPTOR-P that has a field-of-view of ~500 sq-degrees and can patrol the full visible sky in less than an hour. The right-hand panel shows RAPTOR-S, a 0.41-meter telescope that can rapidly slew and follow-up anomalies in less than 10 seconds after receipt of an anomaly message.

To measure the astrometric accuracy of STM observations by the RAPTOR-P and RAPTOR-S telescopes, we collected metric measurements of the GPS (Global Positioning System) satellites. The GPS satellites were selected as calibration targets because high-precision orbits were easily available to us. During the observations the GPS satellites themselves were not tracked; the telescope mounts tracked at the sidereal rate so the satellites appeared as streaks in the imagery. Our real-time streak detection and astrometric calibration software was used to determine the streak endpoints that correspond to shutter open and close times. Those position measurements were then compared with the high precision orbit predictions to determine the measurement errors. Fig. 11. shows the measurement deviation distributions both along the streak (track) and across the satellite track as well as the Median Absolute Deviation values for the two systems. Those values are used in the STM sensor agent software to simulate realistic agent performance.

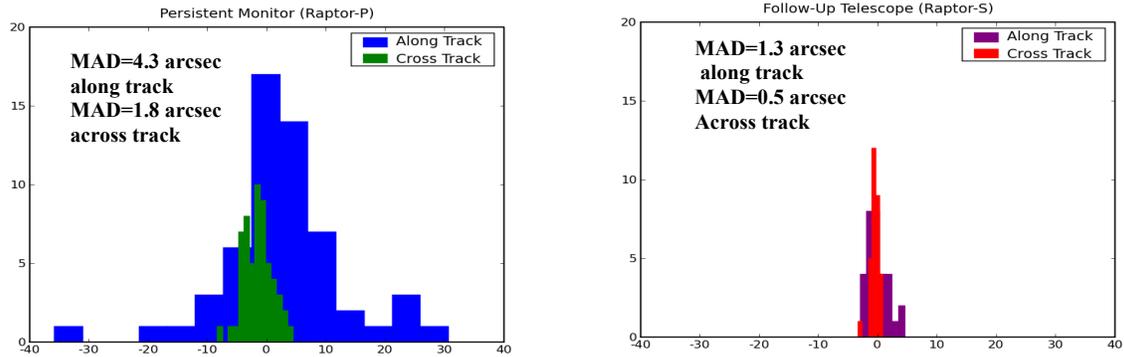


Fig. 11. Histograms displaying the accuracy of endpoint extraction for GPS satellites with the wide-field patrolling telescope (RAPTOR-P) and the narrow-field follow-up telescope (RAPTOR-S). The MAD (Median Absolute Deviation) of the distribution for along track and across track measured endpoints are displayed.

To model the full sky persistent monitor we used the measured performance of our RQD2 (*RAPTOR-Q Demonstrator 2*) system shown in fig. 12. The RQD2 system [10] uses Canon 24mm F1.4 EF camera lenses that have a focal ratio of 1.4 and a 17 mm clear aperture. When used with the Apogee U10 cameras, these lenses provide a 60 degree field of view. Four of the RQD2 cameras are mounted on the enclosure so that they point in the four cardinal directions at a 45 degree elevation above the horizon. The fifth camera is pointed at the zenith. Together, the cameras cover about 95% of the sky above an elevation of 12 degrees. The system takes 10 second long exposures of the sky at a rate of three times all night long. The use of 5 rectilinear lenses with custom-built baffles isolates problems associated with the bright moon that compromise fisheye lens systems and enables higher astrometric accuracy. A detailed description of the approach to streak detection and the achieved astrometric precision of this system is presented in [11].



Fig. 11. The RQD2 full-sky persistent monitor. This fully autonomous, portable, observatory extracts real-time photometric and astrometric measurements and is able to help recognize STM anomalies as they emerge.

The photometric precision of satellite measurements by RQD2, due to the fixed orientation of the lenses, are very much dependent on the rate that it appears to move across the sky. Slower movement generates a longer dwell in the image pixels, a better signal-to-noise ratio for photometric measurements and easier detectability. Fig. 13 displays 45,768 photometric measurements of objects measured during fourteen nights from a RQD2 system deployed at the Fenton Hill Observatory site. The measurements are associated with 1,565 unique objects in the spacetrack.org catalog. Notice that higher range objects are detectable to fainter magnitude because (due to Kepler's law) they have slower apparent angular motion across the sky. To provide a better feel for the sensitivity of the RQD2 system, we plotted on the figure (red line) the magnitude that would be measured for a Lambertian Sphere with diameter of 1 meter

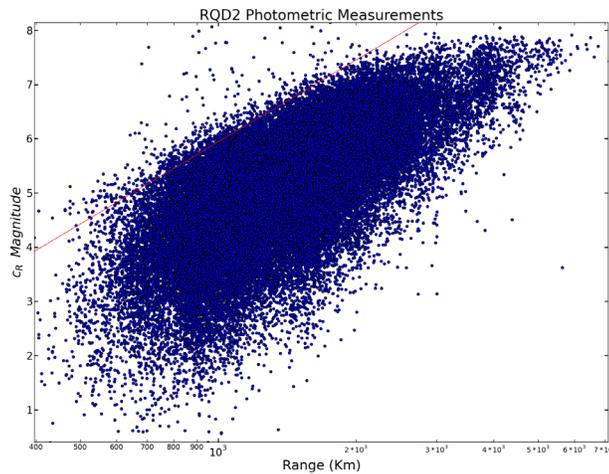


Fig. 12. Photometric measurements taken by RQD2 of satellites in spacetrack.org. The cameras are unfiltered, so the measured instrumental magnitude values are transformed to an R-band equivalent using the catalog R-band magnitudes of calibration stars in the images. The red line denotes the predicted magnitude for a Lambertian sphere with one meter diameter and 20% albedo.

3.2 Example of first STM test results

The integration of the full AI-enabled DCA simulation package for STM applications is still a work in progress. But we have started to test and quantify the utility of a STM system with a global footprint that negotiates with International partners that have hidden priorities. For example, to test the time to detection of a TLE change for the STM network described in section 3.1, we randomly selected 100 TLEs from the RQD2 list of detected objects with orbital periods less than 3 hours. We then gave them a slight delta-v (on the order of 10 meters/sec) at random times distributed around UTC (Coordinated Universal Time) 2021-08-26T00:00:00.00. A total of 3,000 events were then generated (30 events per modified TLE) at a random time in the following six hours. These event times were then used to simulate the start of a negotiated STM collections and the time to successful collection was calculated.

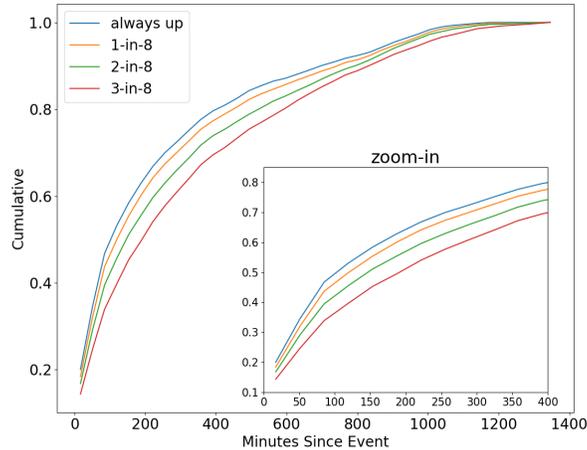


Fig. 13. Cumulative fraction of all events versus difference of observation time and event time. Reported fraction (y-value) gives the fraction of all events for which the first reported observation time following the event is less than or equal to the reported lag time (x-value). Higher negotiated availability (blue - 100%, orange - 87.5%, green - 75%, red - 62.5%) yields higher fraction of events at lower lag times.

To mimic unsuccessful negotiation (or an unsuccessful collection) with a STM sensor agent, we compiled statistics for random agent collection failures (for any possible reason) of 0%, 12.5%, 25% and 37.5%. A plot of the results is shown in fig. 13. The various curves are ordered with increasing cumulative percentage of TLE recovery at lower lag times with an increasing probability of successful negotiation. Of course, the lag times plotted there are very dependent on the specific geographic locations of the sensors; with shorter lags being achieved with greater geographic diversity. Hence the need for international collaboration whenever possible and the motivation to keep the barrier of entry low for autonomous, coordinated, STM measurements.

4.0 DIRECTIONS FOR FUTURE WORK

In this paper, we have described the concept for a new architecture for STM that enables real-time coordination between international government and commercial partners. One of the nice features of this DCA approach is that it can easily incorporate modern AI-driven cognitive sensing. Our next steps will be to use our realistic STM simulation environment to explore ecosystem responsiveness to important STM anomalies such as dangerous conjunctions, lost object search and new launches. Another line of exploration that will be explored is optimization of the composition of the sensor ecosystem by sensor placement and telescope capability.

5.0 REFERENCES

[1] M. Dominguez et al., "Space Traffic Management", National Academy of Public Administration Report, August, 2020. Available at: https://www.napawash.org/uploads/NAPA_OSC_Final_Report.pdf

- [2] R. Seaman, W. T. Vestrand and F. Hessman, "Reengineering observatory operations for the time domain," *Proc. SPIE9149*, Observatory Operations: Strategies, Processes, and Systems V, (2014).
- [3] R. Davis and R.G. Smith, "Negotiation as a metaphor for distributed problem solving", *Artificial Intelligence*, vol. 20, 63-109, (1983).
- [4] N. R. Jennings, "On agent-based software engineering", *Artificial Intelligence*, vol. 117, 277-296 (2000).
- [5] M. Richter and R. Weber, "Case-Based Reasoning - a textbook", Springer ISBN: 978-3-642-40166-4 (2013).
- [6] S. Begum et al., " Case-Based Reasoning Systems in the Health Sciences: A Survey of Recent Trends and Developments", *IEEE Trans. on Systems, Man and Cybernetics*, vol. 41, no. 4. (2011).
- [7] A. Aamodt and E. Plaza, "Case-based reasoning: Foundational issues methodological variations and system approaches", *Artificial Intelligence Communications*, vol. 7, pp. 39-59, (1994).
- [8] W. T. Vestrand et al., "RAPTOR experiment: a system for monitoring the optical sky in real time," *Proc. SPIE 4845*, Advanced Global Communications Technologies for Astronomy II, (2002).
- [9] W. T. Vestrand et al., "Autonomous global sky surveillance with real-time robotic follow-up: Night Sky Awareness through Thinking Telescopes Technology", *Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) 2008*, (2008).
- [10] J. Wren et al., "A portable observatory for persistent monitoring of the night sky," *Proc. SPIE 7737*, Observatory Operations: Strategies, Processes, and Systems III, (2010).
- [11] P. Wozniak, L. Prasad and B. Wohlberg, "Moving point source detection and localization in wide-field images", *Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) 2018*, (2018).