SSC24-WV-03

DORRE: Autonomous Event Response Using an 8-Spacecraft Constellation

Michael Swartwout
Saint Louis University
3450 Lindell Blvd, St. Louis, MO 63103; 314/977-8214
mswartwo@slu.edu

Keith J. Bennett
Bennett Research Technologies
bennett@bennettresearchtech.com

ABSTRACT

Autonomy - the ability of system to retask itself in response to unexpected/unmodeled events - is an enabling technology for space missions. However, autonomy carries risk: what if the system makes unexpected or dangerous decisions? How can human operators trust that the system will not exceed certain bounds of behavior? Are we sure that an autonomous vehicle will outperform a human-directed system?

To address these questions, the Space Systems Research Laboratory (SSRL) of Saint Louis University is developing the Distributed Observation Reasoning and Reaction Experiment (DORRE). DORRE consists of 8 spacecraft operating together in low-Earth orbit to conduct a series of experiments based around event detection and response. Using on-board RF sensors and imagers, the performance of a human-directed team of operators will be measured against the autonomous network: quantifying cost of operations, timeliness of response, accuracy of performance and introduction of risk. The DORRE experiment will quantify the design parameters where autonomy can improve over human performance (e.g., mission geometry, communications architecture, timeliness of response).

The DORRE spacecraft are identical low-cost 3U vehicles in development in SSRL. At present, DORRE consists of one pathfinder spacecraft scheduled to launch in the latter half of 2024 (DARLA), a second pathfinder in assembly (DARLA-02) and an extensive flight-software-in-the-loop multi-spacecraft simulator used to test experiments and operations.

As the project was refined, it became clear that the process to assemble and test of 8 spacecraft in a university environment was an entire mission unto its own. SSRL lacks the facilities to truly automate the assembly & test process, but many of the principles of automation can and have been applied to this student-labor-intensive project. The lessons learned thus far and a proposed approach for (more) rapid assembly of multiple flight vehicles will be presented.

This paper will outline the DORRE mission concept and spacecraft design, focusing on the simulation and hardware-in-the-loop tests of the autonomous flight experiments. The near-term flight plans for DARLA and DARLA-02 will be presented, as well as the process to accelerate the assembly & test campaign. The paper will conclude with an overview of the pathway to flight in late 2025.

INTRODUCTION

Spacecraft autonomy – the ability of a system to retask itself in response to changes to the internal system or external environment – is a feature that is claimed by many systems but fully implemented by few (if any). In this paper, we will discuss a mission in development to demonstrate and evaluate the performance of autonomous systems.

First, we must contrast ***autonomy*** with ***automation***. It is helpful to begin with the SAE standard for levels of driving automation. In that taxonomy, Levels 0-2 involve increasingly-capable automated features like blind spot warnings, adaptive cruise control and automated emergency braking. At Level 3, the car takes complete control of operations but will abort to driver control when it encounters situations beyond its control. Most modern spacecraft still operate at an equivalent to Level 3 automation; the spacecraft runs on its own until it either runs out of schedule or one of its telemetry sensors exceeds its limits, and then in both cases the spacecraft waits in standby or safe mode for a human team to decide what to do next.

For this paper, we are interested in the spaceflight equivalent of Level 4 or Level 5 automation, where the vehicle operates on its own in a wide range of operating conditions, up to and including the point where it can manage unexpected events.

Second, our mission will center on ***event response***. An ***event*** is defined as an unexpected change to the system. The event can be external or desirable, such as the natural occurrence of something scientifically interesting, such as an aurora or meteor. Events can be internal or undesirable such as the battery overheating or the memory buffer overflowing. A highly autonomous system must be capable of identifying events and generating useful responses.

Among the many challenges in developing and testing highly autonomous systems is creating a test and evaluation framework to assess the systems' performance. In what circumstances would an autonomous system outperform a human-directed one? At what point do the consequences of undesirable decisions outweigh the cost and performance benefits of an autonomous system?

The Space Systems Research Laboratory at Saint Louis University, in collaboration with Bennett Research Technologies, is developing the Distributed Observation Reasoning and Reaction Experiment (DORRE). The DORRE mission will create an 8-spacecraft autonomous event-response network and perform several dozen experiments to quantitatively evaluate the relative performance of human-operated systems compared to autonomous agents operating the same system.

In this paper, we will provide a short review of related work, identify the driving characteristics of the experiment we want to conduct, and introduce the DORRE mission. The automated reasoning system will be introduced, the experimental setup presented, and expected results discussed. An overview of the DORRE flight vehicles will be provided, as well as details for the first two precursor missions, DARLA and DARLA-02, scheduled to fly in 2024 and 2025, respectively. As the student teams developed DORRE, we discovered that the process of getting an 8-spacecraft, student-led project from concept to flight could be an experiment in itself. To that end, we will outline our plans to automate the assembly and test process, and provide preliminary work for getting the full DORRE mission to flight.

Related Work

Autonomy is a broad field; this conference alone has more than 1000 prior papers that use the words "autonomy" or "autonomous". For the purposes of this paper, we focus on three aspects: autonomy for decision-making (as opposed to data extraction or event detection), distributed autonomy (as opposed to centralized) and autonomy on board space systems.

Autonomy for space systems has proceeded in fits and starts; one of the very first papers presented at this conference in 1987 discussed the trades between on-board autonomy and human-directed operations [1]. However, despite the many papers covering the topic in the 37 years since then, the predominant mode of operations for space systems is to be human-directed.

In the late 1990s, the Deep Space 1 mission began with a fully-autonomous flight software system (the Remote Agent); by the time the mission flew, the Remote Agent was relegated to a software experiment that was operated for only a short period of time [2]. The Remote Agent's experience is typical for today: the perceived risks of handing control authority to the spacecraft outweigh the perceived benefits.

Still, work in autonomy progresses. The four-spacecraft NASA Starling mission, currently in orbit, pushes autonomy to the vehicles to demonstrate autonomous navigation and channel selection for a distributed GPS-based measurement experiment. At the same time, the complete dataset for all measurements is stored and downloaded [3] – so while autonomous operations are actively pursued, there is still a "safe" backup in case the autonomy makes inadvisable decisions.

There is a real tension between the desire to field autonomous systems that improve performance and save costs against the justifiable fear that an autonomous system will make costly or irrevocable decisions. This tension has given rise to the study of trusted autonomy. In this context, "trust" is defined by Lee and See as, "the attitude that an agent will help achieve an individual's goals in a situation characterized by uncertainty and vulnerability." [4] Recent work highlights the needs to set rigorous bounds or guardrails around autonomous system so that "incorrect" decisions do not have unacceptable consequences [5].

AUTONOMY:
IT SHOULD SCARE YOU A LITTLE BIT

Two ways for trust in autonomous systems to improve are to fly missions with autonomous elements, or to develop experiments that quantitatively evaluate the performance of autonomous systems. The DORRE mission aims to do both through a series of experiments in event response.

As we structured our autonomy-experiment mission, we identified five key characteristics to embody.

Comparative

Adding autonomy does not automatically improve a mission. We want to quantitatively and qualitatively compare the performance of human-directed systems compared to autonomous systems. Without clear metrics that identify the situations where autonomy improves the mission and the situations where autonomy is likely to worsen the outcomes, how can mission designers choose wisely?

For the comparisons to be meaningful, the human team and the autonomous system should be using identical spacecraft in identical configurations. Ideally, as with a ground experiment, one would run the experiment with a human-directed team, then reset the initial conditions and repeat with the autonomous system. On orbit, this approach is not practical, as one cannot reset the spacecraft geometry, nor can one perfectly repeat the environmental effects experienced during the first run.

As discussed in the later section, this need to have a good comparison between systems drove a lot of the spacecraft and experimental designs.

Consequential

Autonomy is scary: by definition, an autonomous system is making operational decisions without human intervention. A natural consequence is that the autonomous system might make the wrong decision (or at least an unexpected one). Any autonomy experiment that doesn't account for the risk of "wrong" decisions is incomplete; effectively, one has put guardrails around the mission and the results are less trustworthy.

For the DORRE mission, this means that we need to allow spacecraft to make wrong decisions. More to the point, both the human operators and autonomous systems must be able to make consequential decisions, decisions that could result in blown schedules, missed contacts or even dropping into safe mode.

We need to be able to observe what happens when bad decisions are made, and to develop methods that protect the system when those decisions happen as well as developing policies for what kinds of information need to be shared between autonomous systems and human supervisors. Of course, we still want to introduce a form of guardrails around the vehicles so that we don't prematurely end the mission through vehicle mismanagement. It is enough that bad decisions cause the spacecraft to drop out of nominal operations.

Constrained

A corollary to ***consequential*** operations is ***constrained*** operations. As we experimentally compare the relative performance of human-directed and autonomous systems, we are most interested in the edge cases. Our time on orbit is limited, and if all of our experiments result in all systems behaving exactly as expected, then we could have simply relied on simulations and not flown at all. Instead, we need to introduce situations where system behavior will bump against the limitations of the hardware: filling up memory, overscheduling instruments and/or reducing the battery charge below a desired threshold.

Having spacecraft with limited performance is not necessarily a limitation to the experiment, because we can evaluate how human and autonomous systems manage their operations in the presence of constraints.

Messy

Given a ***consequential***, ***constrained*** student-build autonomy experiment, it is inevitable that parts of the experiment will get ***messy***: spacecraft will reset or lock up, components will overheat or freeze up, human operators and autonomous agents will interact with one another in unexpected ways. As with the previous goals, we embrace the messiness of an autonomy experiment; if anything, we want these effects to occur at an accelerated rate compared to a "normal" system, so we can observe those events during our experimental run time.

Experimental

The mission is structured as a series of experiments to test our performance hypotheses across a range of initial conditions and objectives. But more than that, as we gain experience with the operational system, we expect to identify new kinds of experiments that could be flown. We intentionally have structured the system to allow for other experiments to fly using the same hardware. In particular, we are open and eager to partnering with researchers interested in evaluating their own autonomy algorithms.

DORRE MISSION OVERVIEW

With that in mind, the DORRE mission is structured to be a series of event-response experiments to compare the performance of a human-directed network (***operators***) with an autonomous system (***agents***). The experiment simulates the need to collect time-sensitive information about a natural event such as a sudden volcanic eruption or destructive hurricane. We have selected visible-light images and RF measurements as representative of a range of typical sensors.

As shown in Figure 1, a DORRE space vehicle is divided into shared resources and team-specific components. The autonomous agent "side" and the human-directed operator “side” each have exclusive control over a set of receiving radios. Both sides share a single imager, power system and communications system. Because the human-directed payloads and the autonomous payloads are on the same vehicle, most of the driving experimental variables are the same for each: orbit/viewing geometry, space environmental effects, communications windows. This enables us to isolate the performance variables.

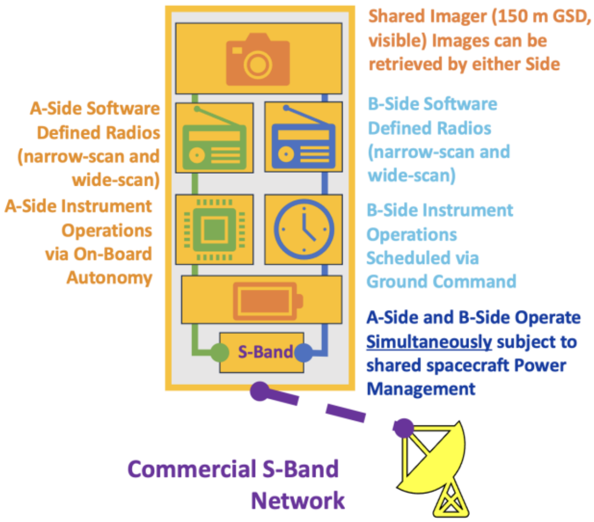

Figure 1: DORRE Payload Schematic

The DORRE spacecraft will use a commercial S-Band network to provide a range of communications windows to operate the network. However, DORRE does not operate in real-time nor does it have direct crosslink between vehicles. Rather, DORRE operates asynchronously via a messaging system that routes commands and data between spacecraft and ground nodes. As discussed previously, we believe that it is important to assess the performance of autonomous systems with constrained resources.

Aside: as structured, the DORRE mission implicitly tests the relative performance of centralized control systems vs. distributed architectures. A single team of human operators will develop and a manage a schedule for all 8 spacecraft; this is inherently centralized control. By contrast, we decided to push autonomy out to the individual spacecraft; we believe high communication latency lends itself to autonomous operations.

Experimental Runs

The DORRE mission consists of several dozen experimental runs. Each run has a fixed start and end time and a set of data-collection objectives. All experiment parameters are selected by an Experiment Coordinator and are not shared with the teams prior to the start of the experiment.

After the experiment begins, at a time of the Experiment Coordinator’s choosing, both teams are notified of an Event and directed to collect a series of measurements about the event (the “Shotlist”). An Event is tagged at one or more locations on the surface of the Earth, and the Shotlist is a list of images and RF measurements required by the Coordinator. These measurements will have characteristics such as: sunlit or darkness; directly overhead or at the horizon; and, in the case of the RF sensors, one or more frequency bands of interest.

Once notified of an Event, team of student operators will use commercial orbit-prediction tools to determine which vehicles will be in view of the Event location(s). They will create command sequences for each vehicle and queue them for delivery to the spacecraft when each is within one of the communication windows. The student team will continue to monitor the system performance, updating their plans as Shots are collected (or not collected) and responding to anomalies.

On the autonomous side, when the Coordinator announces the Event, he will queue the announcement into the messaging system. As an individual spacecraft agent passes through its next communications window, it will receive notification of the event and generate its own plan. On-board orbit and attitude propagation software allows the agent to predict when it will be over the Event location(s) and what sensors will be visible.

The autonomous system manages its response via a global Shotlist (defined by the Coordinator) and a local Agenda that indicates which Shots each agent has collected or will collect. An agent’s Shotlist is augmented with a confidence metric predicting how the quality of the agent’s Shot, and then assesses the quality of the data of the collected Shots. The planned Shots that an agent will collect are called Bids; these Bids are distributed via the communications network to the other agents. As an agent receives the other agents’ Bids, it may opt to add or modify its own Bids.

A snapshot of this process is shown in Figure 2. Note that as the experiment progresses, it is likely that every Agent will be operating with an outdated Shotlist. Because communications are asynchronous, an agent might receive the Shotlist from the Experiment Coordinator and collect multiple Shots before other agents are even aware that the Event has occurred. And those agents may in turn collect duplicate Shots before they received the first agent’s Bids. As with the other complications of this mission, this is a feature. We want to measure how the distributed asynchronous system behaves, because latency and data dropout are major risk factors for using autonomous agents.

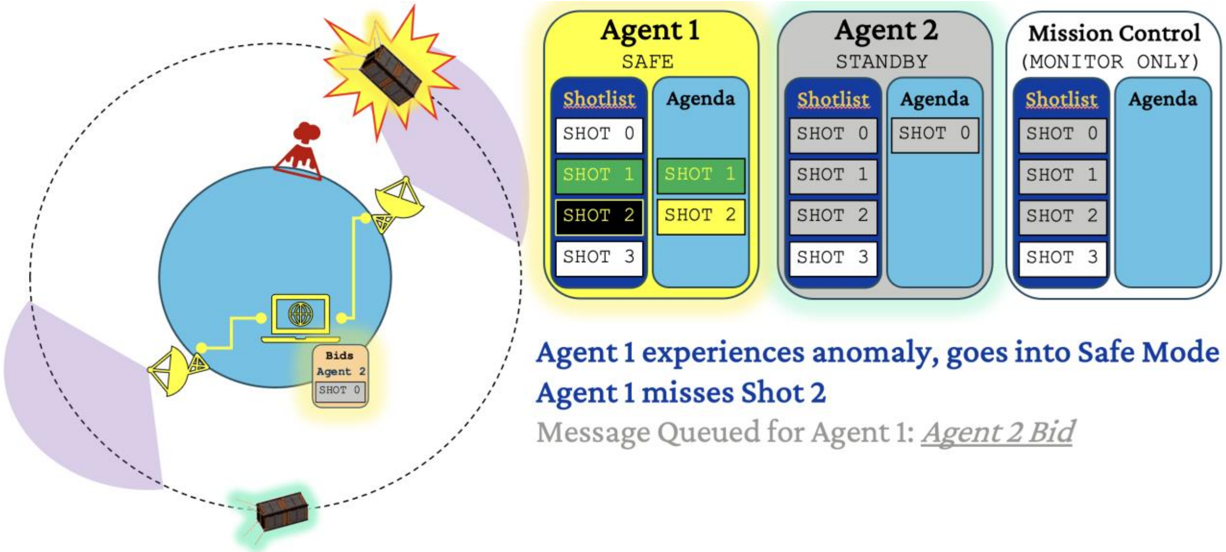

Figure 2: Mid-Experiment Snapshot of the DORRE Automated Reasoning System

An experimental run continues until both sides collect all Shots or the experiment duration is completed, whichever comes first. We do not plan to inject faults into our network, but we will not interrupt an experiment if (when) such faults occur. The DORRE spacecraft have low power margins, and we expect that the agents and operators may inadvertently drop a vehicle into safe mode. This is a feature, not a bug: as noted in the previous section, we want our autonomy experiment to have consequences, so that we can develop the tools for trusted autonomy.

Each experimental run is set at 72 hours, and nominally we would not conduct more than one experiment in a 7-day period. This duration is a compromise between making the data collection process long enough to be interested and allow for complications, and keeping it short enough so that we can make dozens of runs and not wear out our student team. It also gives us time to evaluate performance between runs.

A nominal set of experimental runs is provided in Table 1. As mission life/vehicle performance allows, we will add runs. Also, the DORRE network is capable of adding algorithms after launch, and we remain open to working with others to fly their autonomy experiments.

Performance metrics and expected outcomes

The performance of the autonomous agent side and the student operator side will be quantitatively evaluated after each run a set of metrics.

- **Cost**. The effort required for each side will be converted to a cost (dollar) metric, such as the hourly rates for the operator team (including overtime / late-night shifts), leasing of the communications networks, and engineering effort for troubleshooting / anomaly resolution. The consumption of limited resources on the spacecraft will also be converted to a dollar cost, including on-board data storage and spacecraft down time. It is expected that on balance the autonomous side will complete an experiment at a lower recurring cost.
- **Timeliness**. The time that it takes for a side to complete the Shotlist will be measured from the moment that the Event Coordinator announces the Event. The timeliness metric will be heavily influenced by orbit geometry / communications windows and whether the operator side chooses to have operators on call. Given the nature of event response, we assume that earlier collection of useful data is better, and so we will pro-rate that value of Shot based on the data return. We do not know which side will perform better in timeliness, and varying the timing and geometry of an experiment is one of the primary factors in having so many experimental runs.
- **Quality**. After an experimental run is concluded, all the Shots (measurements) that were collected by both sides will be downlinked and evaluated by the Experiment Coordinator's team. How well each side satisfies the Shotlist requirements will be measured and converted to a dollar value reflective of the quality of the response. (Redundant Shots are factored into the cost metric described above.) It is anticipated that that autonomous side will perform better in the quality metric, because it can perform real-time evaluation of orbit, attitude and on-board capabilities, and thus can make real-time adjustments to when and how data is collected.

- **Risk**. The fourth metric covers the unmodeled behaviors of each side. For example, as each side makes operational decisions that put the spacecraft into Safe mode, that will be counted against them as risk. On the other hand, if the human team is monitoring the weather over the Event and knows that the next spacecraft pass will be clouded over, they can use that knowledge to capture images at a later pass.

The performance of each side will be assessed using each individual metric and then an aggregate score will be assigned. Overall, we anticipate that the human-operator side will be most successful when the communications windows are limited and the team has sufficient time to plan the entire response from the start. The human team will also perform better when the Shots are collected early and they have time to evaluate the data quality to schedule more. We expect that the autonomous side will perform better when agents can share multiple sets of messages between themselves in order to adjust to real-time changes.

Table 1: Example Set of DORRE Experimental Runs

Run	Event Start (hrs)	Time of Event (shift)	Event Location(s)	Shots / Event	Shot Constraints	SV Constraints	Timing Constraints on Operators	Geographic Constraints
0	< 1	normal	Himalayas	20	Elevation only	None	3 hours to plan	Spacecraft can respond following first pass
1	< 1	swing	Tonga	20	Elevation only	None	3 hours to plan	Spacecraft can respond following first pass
2	< 1	overnight	Seoul	20	Elevation, night only	None	3 hours to plan	Spacecraft can respond following first pass
3	< 24	normal	Great Barrier Reef	20	Elevation, day only	None	3 hours to plan	Limited observation times
4	< 24	swing	Madagascar	20	Elevation, day/night	None	3 hours to plan	Limited observation times
5	< 24	overnight	Cape Horn	20	Elevation, day/night	None	3 hours to plan	Limited observation times
6	24 - 48	normal	Panama Canal	20	Elevation, day/night	None	< 1 hour to plan	Observations delayed
7	24 - 48	swing	Galapagos	20	Elevation, day/night	None	< 1 hour to plan	Observations delayed
8	24 - 48	overnight	Hawaii	20	Elevation, day/night	None	< 1 hour to plan	Observations delayed
9	< 12	normal	Himalayas	40	Elevation, day only	None	Event in middle of constellation flyover	
10	< 12	swing	Tonga	40	Elevation, day/night	None	Event in middle of constellation flyover	
11	< 12	overnight	Seoul	40	Elevation, night only	None	Event in middle of constellation flyover	
12	< 12	normal	St Louis Great Barrier Reef	20	Elevation, day/night	None	3 hours to plan	Distributed sensing
13	< 12	swing	San Francisco Panama Canal	20	Elevation, day/night	None	< 1 hour to plan	Competed sensing
14	< 12	overnight	Madrid Galapagos	20	Elevation, day/night	None	< 1 hour to plan	Serial sensing
15	< 12	normal	St Louis	20	Elevation, day/night	Only 6 participate	< 1 hour to plan	
16	< 12	normal	San Francisco	20	Elevation, day/night	4 imagers, 4 SDRs	< 1 hour to plan	
17	< 12	normal	Madrid St Louis	20	Elevation, day/night	4 imagers, 4 SDRs	3 hours to plan	
18	< 12	normal	St Louis	20	Elevation, day/night	Disable 2 spacecraft ~12 hrs	3 hours to plan	
19	< 12	normal	San Francisco	80	Elevation, day/night		Strain the comm throughput	
20	< 12	normal	Madrid	80	Elevation, day/night		Strain the comm throughput	

DORRE SPACECRAFT

The DORRE spacecraft are intended to be as low-cost as possible while providing enough performance to meet the experiment needs. At the forefront of the DORRE design is the recognition that this is an undergraduate student-led design, fabrication and analysis program. While the authors have designed the experiments and guide the mission, the students have the lead in all the engineering activities. They have scoped the spacecraft accordingly.

Each DARLA space vehicle is a 3U CubeSat with aggregate parts cost on the order of $100,000 (Figure 3). The body-mounted arrays on the frame are augmented by two deployable solar panel wings to provide 7 W daily on-average power. In nominal mode, the spacecraft consumes 4 W and 11.5 W when all payload instruments are active. As intended, both the on-board agents and human operators must carefully manage instrument operations to maintain battery charge.

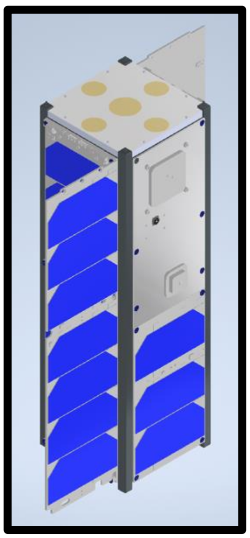

Figure 3: DORRE External View

The vehicles are passively magnetically stabilized such that the S-Band antenna will be downward facing over the Northern hemisphere; attitude knowledge is provided by sun sensors, rate gyros and magnetometers. Orbit knowledge is provided by an on-board GPS receiver.

The DORRE vehicle will slowly rotate around its long axis such that the side-facing radio antennae and imaging payload sensors will have periods of facing the horizon and nadir.

Communications is maintained via a full-duplex S-band transceiver operating in the commercial bands with one of the global S-band communications networks. S-band is preferred over UHF, not because of the required data rates but because of the relative availability of S-band stations. During an experiment, each DORRE vehicle requires on the order of 1 MB per day downlink and 0.5 MB per day uplink; the messaging system by which the agents pass information can involve a lot of uplink to the vehicles.

As noted in the previous section, DORRE vehicles have no active crosslink; however, due to the asynchronous nature of the communications systems, the vehicles effectively have an asynchronous crosslink via messages routed through the ground network.

The primary payload instruments are four software-defined radios; two controlled by the on-board agent, two controlled by operator command. Nominally, during the experiment one SDR will be configured to scan over a wide field of view while the other will have a narrower antenna beam and focus on specific frequencies. In both cases, the SDRs are measuring signal strength only; no attempts are made to decode the signals.

The secondary instrument is a commercial Pi camera which is jointly operated by both agent and operator. There is no conflict where both sides want to operate the camera at the same time; the same file can be tagged and shared with both sides. In the case where one team wants to operate the camera slightly before the other, the commands will be executed sequentially, as close to the requested time as possible. Timing delays are an expected part of the ***Messy*** feature of this experiment.

The flight software was developed by BRT in support of prior SSRL missions. The Python-based Artificial Reasoning for Exploration and Space (ARES) framework is both a common operating system for the ground nodes and spacecraft and the data protocol used to manage the transmission and reception of asynchronous messages across the system.

These spacecraft are designed to operate autonomously as long as on-board resources allow, particularly battery charge. As shown in Figure 4, DORRE spacecraft have as few operating modes as possible while meeting deployment, testing and safety constraints. The main means of transitioning between modes is the voltage of the on-board battery; a customizable threshold determines when the payload is active (Payload Mode) and when the spacecraft charges in Nominal Mode. When anomalies occur, the spacecraft reverts to Safe Mode which requires additional review before returning to payload operations.

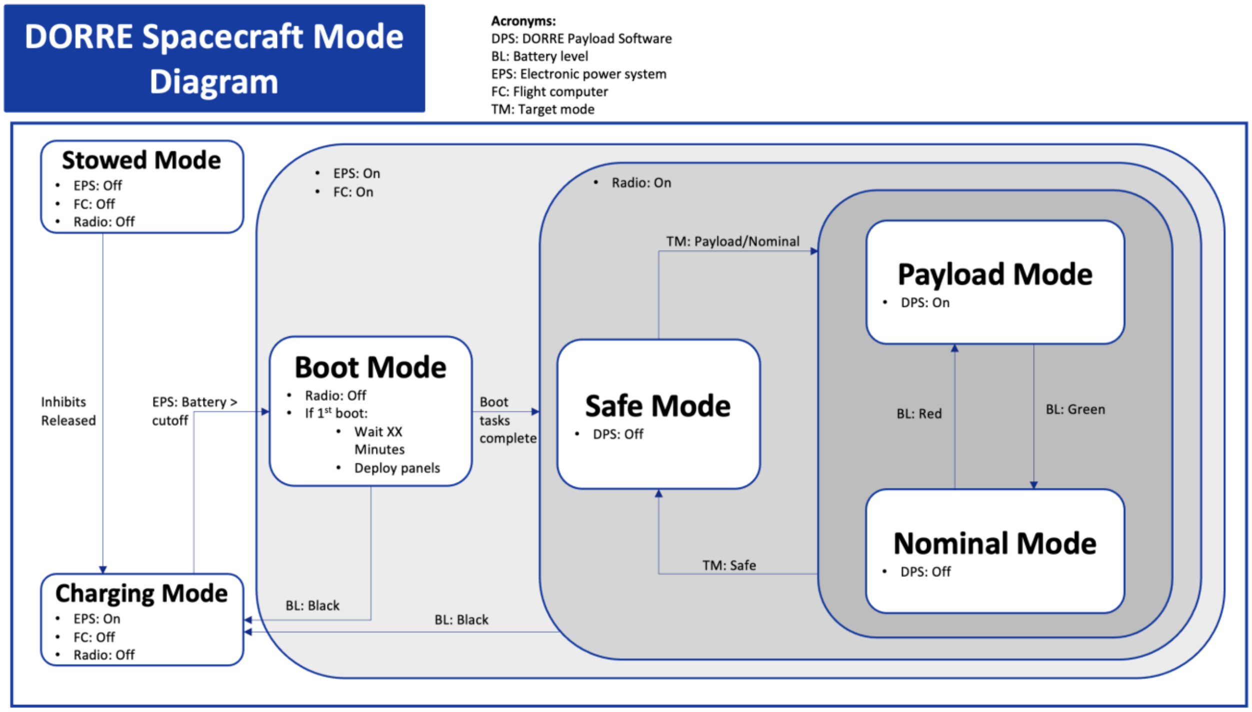

Figure 4: DORRE On-Board Modes

A series of flight tests are planned before the full DORRE experiment is flown. Called the Demonstration of Artificial Reasoning, Learning, and Analysis (DARLA) these missions have three related goals: technical risk-reduction before the DORRE flight, experimentation in support of the autonomy mission, and improving the technical capabilities of the student team while preserving institutional knowledge. (And, of course, because it's great to build and fly your own spacecraft.)

First Flight Test: DARLA

As of this writing (June 2024), the first DARLA flight is scheduled for October 2024 on a Firefly Alpha launch. The NASA CubeSat Launch Initiative (CSLI) is sponsoring DARLA. This flight will validate the flight software, assembly & test techniques and baseline performance of a SLU-built DORRE-class spacecraft. The DORRE flight spacecraft is shown in Figure 5.

Second Flight Test: DARLA-02

The second DARLA mission has been selected for flight by NASA CSLI. Assembly and test of DARLA-02 will commence in late 2024 after the delivery, launch and initial operations of the first DARLA mission. Any necessary changes based on those flight results will be folded into the development process. In addition to advancing the readiness of the technologies for the full

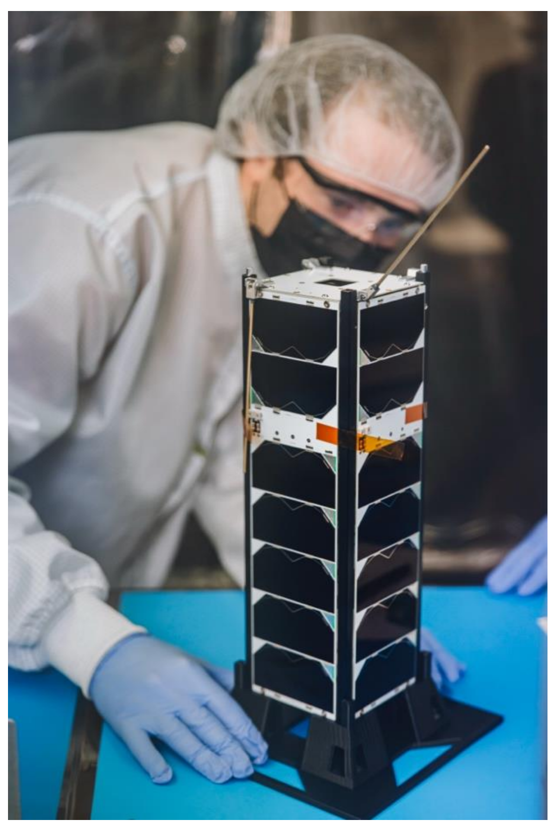

Figure 5: DARLA Flight Vehicle

DORRE mission, DARLA-02 will serve as a pathfinder for SSRL's multi-spacecraft assembly and test process. Using the same processes that were developed for the first DARLA – but with a new group of students – DARLA-02's integration and test campaign will be closely measured to scope out our prospects for building 8 spacecraft in a reasonably short period of time.

EIGHT SPACECRAFT: ARE YOU SERIOUS?!?

In addition to questions about the mission itself, the DORRE team has faced skepticism about our ability to build eight spacecraft. Given how often universities delay (or cancel) single-spacecraft missions, it is understandable to doubt how a school would manage to build 8.

We are trying to be only slightly flippant when we reply that we aren't building 8 spacecraft; we're building one spacecraft 8 times. And, in truth, our baseline is to build the same spacecraft 10 times – DARLA and DARLA-02 are pathfinders for this process. Under this approach, the first DORRE spacecraft will take a longer time to build, as we resolve the expected integration snafus. Beginning with the second vehicle, we will cement the assembly procedures and apply an apprentice model: the assembly team for the second spacecraft will have apprenticed under the assembly team for the first. In turn, the apprentices on the second spacecraft will lead development of the third. At this point we will rotate back through with three experienced teams until all 8 vehicles are completed.

We expect that the first vehicle will take 3 months to integrate in our assembly room (Figure 6), and the remaining 7 will require 2 weeks each. The testing campaign will follow a qualification / acceptance model, wherein the first vehicle is tested at higher levels and the remainder at an accelerated pace.

Each vehicle will be "owned" by a dedicated two-student team that will follow it through its entire assembly / test / delivery / flight operations phase. This team has primary responsibility for the health and safety of the vehicle.

PATHWAYS TO FLIGHT
(AKA BROTHER, CAN YOU SPARE $2M?)

The baseline plan is to fly the full DORRE after the completion of the DARLA-02 mission. However, this timeline is not under our complete control. With parts costs at $100k/vehicle, plus testing and launch costs and paying for time on a commercial communications network, the price of the full DORRE mission could approach $2M. The authors are actively seeking partners to manage such an undertaking. In the absence of progress towards that goal, SSRL may need to field a DARLA-03 and a DARLA-04 before the full DORRE could be realized.

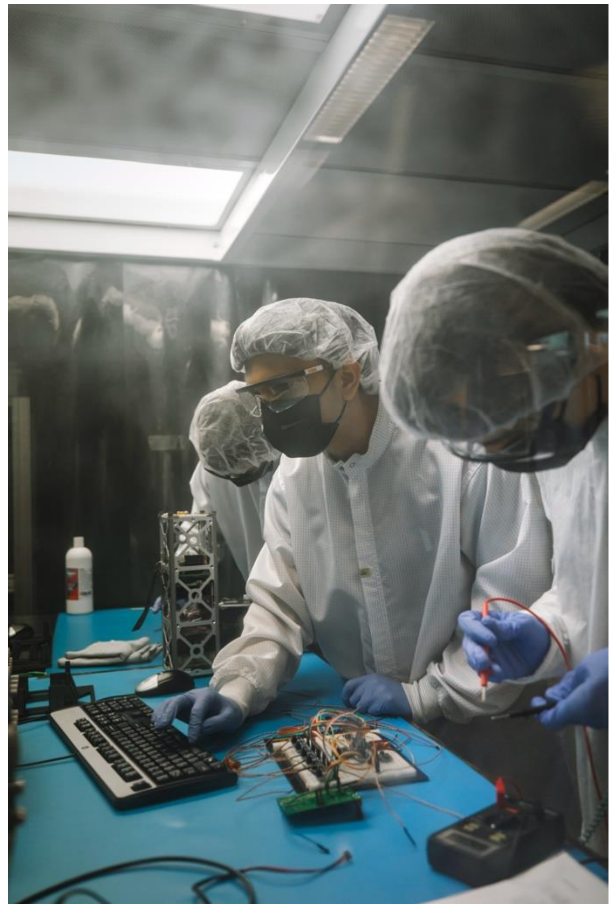

Figure 6: SSRL Spacecraft Integration Lab

In parallel, the SSRL team is evaluating dramatically less expensive bus concepts (e.g., the PyCubed system) that might enable 8 or 12 or 16 1U spacecraft that could approximate the DORRE performance, but at an order-of-magnitude lower cost.

CONCLUSION

The DORRE mission has been structured to provide quantitative assessment of the relative performance of a human-operated system compared to a distributed autonomous system. The chosen experimental framework is a response to an unplanned science event. We believe that this framework addresses the needs for an autonomous demonstration to be comparative, consequential, constrained, messy and experimental. We look forward to demonstrating key parts of the DORRE mission in two precursor flights in the next 18 months (DARLA and DARLA-02) and are actively seeking mission partners to complete the full mission.

ACKNOWLEDGEMENTS

More than 100 SLU students have contributed to the DORRE, DARLA and DARLA-02 missions. We are particularly indebted to the program managers and lead systems engineers: Connor Morris, Celia Taylor-Puckett, Damian Oven, Declan Byrne, Michael Dompke, Eric Suter, Benjamin Schneiderheinze, Owen Binkley, Nathan Brubaker and Melanie Reym.

The DORRE mission and the DARLA demonstration vehicles were developed with the support of the Space Force University Nanosat Program via a subcontract to the Space Dynamics Laboratory (CP0071892 and CP0053704). The NASA Missouri Space Grant Consortium provided additional support.

REFERENCES

1. Smith, D. "Consolidated Space Test Center Capability to Support Small Satellites," Proceedings of the 1st AIAA/USU Conference on Small Satellites, Logan, UT, 1987. https://digitalcommons.usu.edu/smallsat/1987/all1987/25/

2. Aljabri A.S, et. al., "Infusion of Autonomy Technology into Space Missions: DS1 Lessons Learned," 1998 IEEE Aerospace Conference Proceedings (Cat. No.98TH8339), Snowmass, CO, USA, 1998, pp. 315-329 vol.2, doi: 10.1109/AERO.1998.687918.

3. Adams, C., et. al, "An Overview of Distributed Spacecraft Autonomy at NASA Ames," Proceedings of the 37th Annual Conference on Small Satellites, Logan, UT, 2023. https://digitalcommons.usu.edu/smallsat/2023/all2023/78/

4. J. D. Lee and K. A. See, "Trust in automation: Designing for appropriate reliance," *Human Factors*, vol. 46, no. 1, pp. 50–80, 2004. doi: 10.1518/hfes.46.1.50_30392

5. K. L. Hobbs et al., "Space Trusted Autonomy Readiness Levels," 2023 IEEE Aerospace Conference, Big Sky, MT, USA, 2023, pp. 1-17, doi: 10.1109/AERO55745.2023.10115976.