


Enabling and Enhancing Astrophysical Observations with Autonomous Systems

Rashied Amini^{1,a}, Steve Chien¹, Lorraine Fesq¹, Jeremy Frank², Ksenia Kolcio³,
Bertrand Mennesson¹, Sara Seager⁴, Rachel Street⁵

July 10, 2019



Endorsements

John Day¹, Russell Genet⁶, Jason Glenn⁷, Ryan Mackey¹,
Marco Quadrelli¹, Rebecca Ringuette⁸, Daniel Stern¹, Tiago Vaquero¹

¹*NASA Jet Propulsion Laboratory*

²*NASA Ames Research Center*

³*Okean Solutions*

⁴*Massachusetts Institute of Technology*

⁵*Las Cumbres Observatory*

⁶*California Polytechnic State University*

⁷*University of Colorado at Boulder*

⁸*University of Iowa*

^a rashied.amini@jpl.nasa.gov

1 Executive Summary

Autonomy is the ability of a system to achieve goals while operating independently of external control [1]. The revolutionary advantages of autonomous systems are recognized in numerous markets, *e.g.* automotive, aeronautics. Acknowledging the revolutionary impact of autonomous systems, demand is increasing from consumers and businesses alike and investments have grown year-over-year to meet demand. In self-driving cars alone, \$76B has been invested from 2014 to 2017 [2]. In the previous Planetary Science Decadal, increased autonomy was identified as one of eight core multi-mission technologies required for future missions [3].

The impact of autonomous systems on our ability to observe the universe can be just as revolutionary [4]. However, relevant autonomy work to date has been limited in scope and too disjoint to confidently deliver anticipated capabilities, like in-space assembly (ISA), in a low risk and repeatable manner in the 2020s or even the 2030s. This paper includes the following so that the astrophysics community can realize the benefits of autonomous systems:

- A description of autonomous systems with relevant examples
- Enabled and enhanced observations with autonomous systems
- Gaps in adopting autonomous systems
- Suggested recommendations for adoption by the Astro2020 Decadal

As we consider the observations necessary to answer new science questions formed in the 2010s, the need for autonomy is clear. Concept studies for the Astro2020 Decadal require operations that are more complex than ever before. Increasingly complex space- and ground-based observatories have more systems, components, and software. More engineering complexity invariably means that there are more paths for anomalies to disrupt a system's ability to perform its mission. This can reduce observational efficiency and potentially negate the advantages of larger apertures and more sensitive detectors.

Servicing is a legal requirement for WFIRST and the Flagship mission of the 2030s [5], yet past and planned demonstrations may not provide sufficient future heritage to confidently meet this requirement. In-space assembly (ISA) is currently being evaluated to construct large aperture space telescopes [6]. For both servicing and ISA, there are questions about how nominal operations will be assured, the feasibility of teleoperation in deep space, and response to anomalies during robotic operation.

The past decade has seen a revolution in the access to space, with low cost launch vehicles, commercial off-the-shelf technology, and programs that have enabled numerous cubesat missions. NASA and academic institutions will be operating more small satellites and operations centers will need to adapt. The need will be greater if future human exploration goals to launch dozens of cubesats per SLS launch is met [7]. Operating autonomous observatories provides one solution to this impending problem. Notably, several ground-based observatories, like Las Cumbres and ALMA observatories, have begun using autonomous operations to command large arrays of telescopes, identifying advantages for observatories that follow their example. Planet and presumably SpaceX's Starlink, private space mission operators, have reached a break point where traditional commanding is inadequate to command their large constellations and are operating spacecraft with automated scheduling [8].

Gehrels/Swift is an inspiring example of the time-domain observations that autonomous systems enable. The multi-messenger approach for characterizing the physics leading to and resulting from gravity wave events will require similar missions to Gehrels/Swift. Gehrels/Swift relies on prescriptive state machines, statically-programmed conditions and routines also used in spacecraft fault protection, to execute autonomous Gamma-ray burst (GRB) follow-up observations. The system autonomy approach detailed in this paper offers several advantages over state machines in terms of dynamic decision-making and scalability. One major advantage is the ability to make decisions using on-




Figure 1: Effective, reliable autonomous systems must coordinate between the resources utilized by a system’s lower level functions to achieve system-level goals. Appendix A offers an illustrated example of a system autonomy framework.

board analysis of data to change an observation program.

Dynamic decision-making also enables the restoration of functionality in the event of an anomaly. This type of decision-making is enabled by on-board health monitoring software, which monitors and diagnoses hardware anomalies to support autonomous systems. This results in greater observational efficiency and universally benefits all observatories. For observatories with competed time, this means more PIs can be supported. For mapping missions, like the Galaxy Evolution Probe, Probe of Inflation and Cosmic Origins, and Cosmic Dawn Intensity Mapper Probe, greater depths can be reached per unit time [9, 10, 11]. For time-domain surveys, this results in less gaps in data.

As evidenced by private investments and developments in ground-based observatories, the adoption of autonomous systems in space is inevitable. There are two questions to the field: “*When will we start using it?*” and “*How will we start using it?*” Given the ambitions of the community, the time to begin is now. In order to use it in a repeatable, low risk, and cost-effective way, NASA, spacecraft vendors, and the astrophysics community need to cooperatively develop a coherent technical path forward. To do so, **our primary recommendation is for NASA to incentivize the use of autonomous systems for competed space**

missions, for instance through a cost cap credit. Adoption in the 2020s will reduce the risk of future Flagship servicing missions.

2 Understanding Autonomous Systems

Observing the proceedings of the Space Astrophysics Landscape in 2020 and Beyond meeting, it is clear that a gap exists between the expectations of the astrophysics community and the technical readiness of autonomy technologies required to meet these expectations. To understand this gap, we need to first define autonomy in a relevant context.

A hierarchy of systems is represented in Figure 1. At the bottom of the hierarchy is the functional-level, where control and autonomy is exercised in a limited domain. Functional control is the commanding actuators and sensors, *e.g.* a command is sent and a motor turns at a commanded rate. Functional autonomy is decision-making within the boundaries of the functional element. A simple example is a state machine that (dis)engages a heater based on thermometer input. A more complicated example is an attitude controller that takes inputs of attitude knowledge (*e.g.* star trackers). Its output is control system actuation to maintain a desired attitude. Pre-programmed routines filter inputs and evaluate conflicting knowledge, resulting in predictable behavior.

More complex forms of functional autonomy have already been demonstrated and are currently being developed. For instance, autonomous optical navigation determines deviation from desired orbit ephemeris and has been used on Deep Space-1, Deep Impact/EPOXI, other planetary missions, and soon Arcsecond Space Telescope Enabling Research in Astrophysics (ASTERIA) [12, 13]. On-going work on servicing and ISA utilizes computer vision as a knowledge source to control robotic actuation [6]. On-orbit robotic servicing was first demonstrated on DARPA’s OrbitalExpress in 2007 [14]. In the next few years, RESTORE-L will be used to service Landsat-7 in low earth orbit using teleoperation after autonomous docking [15].

However, functional elements utilize system

resources, *e.g.* time, power, attitude, data storage, *etc.* Spacecraft are resource limited and efficient use is critical to mission success. Different activities may utilize resources in a mutually exclusive way; for instance, a space telescope may not be able to point its telescope at a target while simultaneously pointing its antenna toward Earth for communications. Some resources are zero-sum but accommodating of multiple spacecraft goals; for instance, all powered equipment require power but not all subsystem power modes can be supported simultaneously. Thus, there is a state of competition between different system goals. In the current state of practice, this competition is resolved by human planning during operations. Tools are used to define system activities, like observing and transmitting data, based on commands that are tied to certain resources. The goals of the scientists to observe the sky and goals of the engineers to preserve the spacecraft are merged using these tools to develop time-ordered sequences of commands that are uplinked to the spacecraft, *e.g.* [16]. An extension of time-ordered sequences is conditional sequencing, where sequences use conditional statements as a state model. This approach has the capability of storing pre-defined routines that can later be executed [17].

Autonomy poses a challenge to operational planning: how can you command a system that makes its own decisions? State machine-based autonomy is predictable in well-defined environments, and so resource budgets can be allotted because the input domain is well characterized. Spacecraft health is further ensured by fault protection state machines, adding another layer of protection. The use of state machines enables *Gehrels*/*Swift* to detect GRBs with the wide-field Burst Alert Telescope and slew to observe with its two other payloads [18]. However, machine learning-based decision-making and variable environments, exemplified by computer vision-guided robotic control, means that resource utilization cannot be effectively bounded in advance and so reliable, safe operation cannot be readily assured with traditional commanding.

Coordination of resources used by functional elements, prescriptive or not, can be accom-

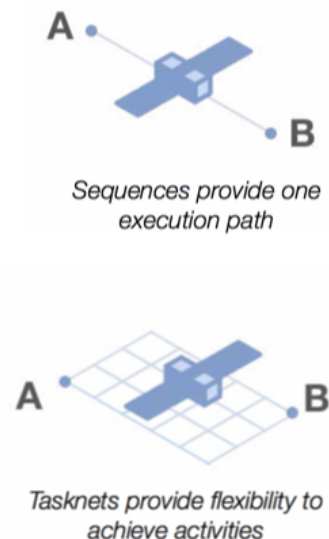


Figure 2: Task networks offer numerous pathways in time and state-space to achieve goals requested from ground operators. Implementing tasknet-based commanding enables “push button-get science” missions.

pished at the system-level through on-board planning and execution. This approach contrasts with traditional commanding with sequences through its use of task networks (tasknets, though goal and constraint networks are also used in the literature), described in Figure 2. Tasks are defined as commands that are associated with metadata defining the state conditions required for their execution and state impacts that result from their execution. Thus, graph networks of tasks can be constructed with tasks as nodes and edges connecting tasks whose state impacts are the state requirements of another task. Moreover, tasks can have temporal constraints to be sequence-like. In this manner an autonomous system can be commanded like a traditional system if desired.

Sets of tasks can be defined as independent, uniquely prioritized system goals. Some goals identify system state transitions, such as the acquisition of new science data. Other goals identify states that need to be maintained and restored if lost, such as those related to spacecraft health. The role of on-board planning and execution is to negotiate between the constraints of all goals so that they can be executed without

conflict or in violation of safe resources limits.

The final level of hierarchy at the top of Figure 1 supports multiple autonomous systems in a multi-agent architecture.

3 Examples of Relevant Autonomous Systems

Spitzer Space Telescope, Dawn, Juno, and many other planetary missions have made use of conditional sequencing with the Virtual Machine Language (VML). Spitzer reported several advantages over traditional sequencing using VML. In particular, it made observations contingent on telescope settling state rather than sequenced time, which added one or two extra observations in an $11\frac{1}{2}$ hour observing window. It also had the advantage of reducing spacecraft safing due to on-board memory overflow [19]. However, as reported in [19], the limiting factor in implementing more of these autonomous behaviors was that there was “no fast and effective way of modeling the flight system behavior on the ground.”

Autonomous systems relying on on-board planning and execution are beginning to see increased use on space- and ground-based observatories. A prominent example is the use of ASPEN/CASPER on Earth Observer-1, which used on-board science planning and execution to detect novel terrestrial scenes, like disasters, to autonomously perform follow-up observations [20, 21]. Extending the work of CASPER, the Intelligent Payload EXperiment (IPEX) cubesat executed one year of autonomous payload operations using its on-board planner [22, 4]. PPlan Execution Interchange Language (PLEXIL) is funded to be used to for a technology demonstration mission of multi-agent autonomy. Later in 2019, ASTERIA will demonstrate the use of the Multi-mission EXECutive (MEXEC). Next year, Mars 2020 will use the Onboard Scheduler to maximize science return by using excess time and power at the end of each Martian sol to plan additional measurements [23, 24, 25]. Temporal planning and scheduling systems also include Ix-TeT [26], used for robotic control, and EUROPA [27]. Other systems have been developed based on similar principles since then, notably IDEA and T-REX, used for autonomous underwater

vehicles [28]. Appendix A offers an example of how these software are implemented in practice.

Las Cumbres Observatory and Atacama Large Millimeter/submillimeter Array (ALMA) are examples of ground-based observatories whose operations are autonomously planned and executed. Las Cumbres Observatory is a network of 18 telescopes at six sites that operate as a single observatory, enabling persistent observation. Scientist request observations, which are assigned and scheduled through a global scheduler [29, 30]. General-purpose software has been developed for autonomous telescope operations that can be adopted by future observatories operating on these principles [31]. ALMA dynamically schedules and executes 30 minute “scheduling blocks” based on weather, science priority, project completion, and other parameters [32].

Automated scheduling has traditionally been used for operational planning. Most relevantly, Space Telescope Institute uses SPIKE for planning Hubble observations [33]. Planet uses automated scheduling to operate its fleet of earth observing cubesats. In human spaceflight Timeliner has seen significant use on-board the International Space Station (ISS) and is being considered as a candidate for the Lunar Gateway [34]. While automated planning streamlines operations, it still has drawbacks when the plan cannot succeed due to operational conditions.

4 Astrophysics with Autonomous Systems

4.1 Enabled Missions and New Science

Autonomous systems have already enabled new astrophysics. Both ground- and space-based transient event observatories are fundamentally enabled by autonomous systems. Autonomous transient event detection and follow-up observation capability has been demonstrated with Gehrels/SWIFT and the Zwicky Transient Facility [35]. As exemplified by LCO, time-domain astronomy observations, *e.g.* supernovae, microlensing, near earth asteroids, tidal disruption events, gravitational wave events, *etc.* require real-time, highly reactive telescope scheduling. In these cases, observations cannot be planned

in advance and the configuration of the observations may need to evolve over time according to the characteristics of event.

With the projected improvements to ground-based detection and localization of gravity wave (GW) events, there is a need for observatories that can rapidly observe potential multi-messenger signals. Ground-based observatories will require the ability to respond to external signal, verify observability of the GW ellipse given current observatory conditions, and re-task to perform GW follow-up observation while maintaining knowledge of the past observation. Space-based observatories will be required to do the same while also maintaining spacecraft health. *Gehrels/SWIFT* itself was launched in 2004 and may need replacement in the 2020s to retain the community’s ability to perform GRB detection and localization over large areas of the sky. If ESA’s *Theseus* is selected for M5, system autonomy software and expertise can serve as a potential NASA contribution to that mission.

Given the past priorities of the Astrophysics Decadal and NASA funding, it is expected that space-based time-domain observatories will be competed and are subjected to cost cap. For instance, the Gravitational-Wave Ultraviolet Counterpart Imager (GUCI) has already been proposed for the SmallSat call [36]. These missions can be architected using state machine autonomy, following the *Gehrels/SWIFT*. Given the bounded nature of time-domain observations and the additional advantages afforded by on-board planning/execution, we note that these missions can alternatively use on-board planning/execution as a relatively low risk means of demonstrating and increasing the community’s confidence in the technology.

As discussed in [6], system-level autonomy is required for ISA and servicing in order to coordinate robotic autonomy with the rest of the spacecraft. One example of how critical system-level autonomy is to ISA and servicing is the coordination of a mass model as robotic operation is performed. At a high-level, a servicer spacecraft has the goals of performing robotic operation and assuring attitude control in the presence of disturbance (gravity gradient, solar pres-

sure). To accomplish the former goal, a robotic arm moves, changing the spacecraft’s moment of inertia. To accomplish the second goal, the attitude control system maintains attitude based on a model of the spacecraft’s moment of inertia. If robotic action is not coordinated, the attitude controller’s moment of inertia model will not be consistent with reality. This may lead to over/under actuation of reaction wheels, potentially leading to collision risk and mission failure for both spacecraft.

Multi-agent autonomy also enables new observations. The AEON Network is ground-based facility currently under development operating numerous telescopes that will allow astronomers to submit requests for observation in real-time. Through multi-agent autonomy, a large network of ground- and space-based observatories, like AEON, can coordinate their observing programs across multiple facilities and wavelengths, serving as a powerful tool for characterizing new discoveries. Multi-agent autonomy can also be used on a constellations of low cost satellites as distributed transient event, namely GRB, observatories [37, 38]. By using low cost scintillating detectors on low cost smallsat/cubesat platforms, localization can be performed through time-of-arrival similar to the Interplanetary Network. One advantage of this approach is the timeliness of observation. For instance, in simulations of flooding event observations by an earth observing constellation, a multi-agent architecture measured flood area to 96% accuracy over time as opposed to a centrally planned architecture observing with 70% accuracy, owing to the timeliness of observation [39]. Additionally, multi-agent coordination of more than two assets may be required, or would greatly facilitate, interferometry missions such as LISA.

A unique class of missions that would benefit from ISA and multi-agent autonomy are radio and possibly NIR/optical/UV interferometry missions that may require ISA of large apertures and coordination.

Autonomous systems also complement the increased access to space afforded by small satellites and low-cost launch vehicles. As the total number of missions increases, let alone missions

that may utilize more than one spacecraft such as GUCI, the ground stations and operations facilities become a bottleneck for commanding and monitoring spacecraft. At some point, large numbers of traditional spacecraft cannot be efficiently commanded through traditional means. Autonomous systems reduce the human effort required to command as the burden can be off-loaded to an on-board planner.

Current plans for human exploration offer new platforms for astrophysics missions, creating new opportunities for the development of observatories. Most imminently, lunar exploration may create dozens of opportunities for new measurements. With cubesats piggybacking launches and opportunities to use the Lunar Gateway as a platform for payloads, managing multiple missions and scheduling observations that may have conflicting pointing and thermal requirements on Lunar Gateway becomes increasingly difficult to coordinate across multiple teams [40]. Again, an operational bottleneck results that can be resolved through automated planning. Additionally, returning to the moon creates new opportunities for lunar surface-based observatories. This offers unique opportunities for some radio bands, cosmic ray, MeV γ -ray, X-ray, and UV measurements that cannot be made from Earth’s surface. A Probe mission concept, FARSIDE, is a ~ 10 MHz radio observatory on the farside of the Moon. As it requires a rover for deployment, autonomous mobility and robotic assembly capability is critical to mission feasibility. [41]

4.2 Efficient Observing Programs

The traditional paradigm of commanding reduces the overall efficiency of targeted observing programs as observation length is pre-determined in advance. Later, data is downlinked and analyzed on the ground. However, the efficiency of observing programs can be improved by analyzing data on-board to inform system decision-making.

One example is exoplanet direct imaging, exemplified by HabEx and LUVOIR, that requires a level of 10^{-10} raw contrast to perform direct imaging of exo-Earths. This raw contrast can only be effectively achieved in cases where

exozodiacal light is not so bright that it reduces the effective raw contrast at the exoplanet’s location. Even if exozodiacal light is previously characterized in mid-IR wavelengths [42], these observations may not predict the exozodiacal light at HabEx/LUVOIR near UV to near IR wavelengths. Additionally, not all systems will have constrained inclination that impacts the apparent brightness of the exozodiacal dust. Currently, HabEx and LUVOIR will schedule their observations in advance and use a pre-determined observing program based on little or no knowledge of the actual level of exozodiacal optical brightness around individual targets.

In an autonomous system, coronagraphic imaging can be analyzed on-board the spacecraft to evaluate the contribution of exozodiacal light and the determine the value of continuing observation. In this case, excessive exozodiacal light can be detected on-board within a fraction of the planned observation time. On-board data processing software can then alert the on-board planner to truncate the observation so the observatory can perform the next scheduled observation. Data from the truncated observation is later downlinked for future analysis. In this example, more targets are observed more quickly, resulting in more observing time for other targets of interest and greater exo-Earth yield during the primary mission [43].

Recommendation: NASA should use ROSES as a means of funding software development for on-board data processing.

The advantage of on-board data processing in union with a system planner is not limited to space-based observatories. Subsystems that evaluate weather and seeing conditions can aid to autonomously reschedule planned observations that may not be possible when scheduled, improving their net efficiency.

Recommendation: NASA and NSF should incentivize the development of future ground-based observatories with automated scheduling/execution, following the example of ALMA and LCO.

As discussed above, on-board data processing

and multi-agent autonomy can be used in a coordinated network of ground- and earth-based to perform GW follow-up observatories. In such an architecture, localization that is currently performed *post-hoc* as data is released can be performed on-board within the constellation, resulting in localization while the source is still emitting brightly.

4.3 Adaptive Fault Protection Enables More Observations

Traditional space systems have fault protection schemes that enter safe modes, requiring human diagnosis and commanding to restore nominal operation. As a result, 4% of nominal space-flight operations are blocked by spacecraft safings [44]. Notably, [44] presents a lower bound on blocked operational time, as other anomalies can occur that restrict nominal operations and do not cause safing.

Figure 3 indicates that on-board planning/execution with on-board health diagnosis may mitigate the impact of about 50% of safings. This adds an additional week of nominal operations per year. There are two reasons. Rather than relying on state machines for executing fault protection, health maintenance tasknets can restore the minimum functionality required to perform science operations while not endangering spacecraft health [45, 46]. Second, this architecture permits integration of on-board health diagnosis to monitor the health of hardware and local models for attitude knowledge and control, a major cause of safing events, to inform these health maintenance tasknets [47, 48].

While an additional week of data per space observatory may seem marginal, if applied to NASA’s fleet of space-based observatories it would result in 11 additional weeks of science per year for the community. The benefit is most useful to observatories with PI-directed observations, like Hubble and Spitzer, where additional PIs can be supported. For mapping missions, like GEP, additional mapping depth per unit time is achieved. For time-domain surveys, coverage is more complete in time.

Even without on-board planning/execution, health diagnosis models and software can im-

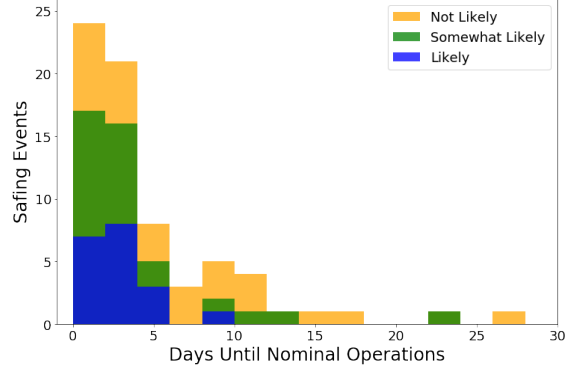


Figure 3: Histogram of safing events binned on the number of days between suspension and restoration of nominal operation. With on-board planning and execution and on-board health diagnosis, about 50% of anomalies resulting safing may be averted. Result based on analysis of the [44] safing dataset.

prove ground- and balloon-based observatories. Ballooning in particular suffers from a high failure rate, owing from *ad hoc* integration of multiple payloads on-site and schedule constraints forcing limited testing. Recently, NASA JPL evaluated technologies for a self-reliant rover during which on-board health diagnosis was found to be effective in the build, integration, and testing environments in discovering and diagnosing hardware issues previously undetected [49, 50, 51]. Health diagnosis software can be used for ballooning systems that are used repeatedly, such as mirror motor control, pressure vessels, and power generation, to detect hardware issues. This can reduce complexity and stress during the balloon integration phase and improve success rate of balloon missions.

Recommendation: Integrate the use of health diagnosis software for elements that are repeatedly used on ballooning platforms.

5 Addressing the System Autonomy Gap

For astrophysics, autonomous systems can enable and enhance missions that deliver revolutionary data sets, reduce the cost of missions, and reduce the burden on scientists in developing and maintaining observing programs. A future where “press button – get science” missions is on the horizon, but work remains that requires

the community’s awareness and support.

The primary gap is cultural. Autonomous systems imply a different paradigm of design and operations compared to traditionally-commanded systems. Compared to autonomy in the private sector, a small proportion of our space science and spaceflight communities have relevant expertise to review the opportunities and risks associated with autonomous systems. This is compounded by traditional engineering preference for heritage designs and expectations of predictability. On point, how can scientists, engineers, and proposal reviewers be confident in a mission concept that operates itself? Is it possible to design and deploy autonomous systems that are partially autonomous to placate the concerns of the community? These questions need to be formally addressed if NASA is to meet its legal requirement to perform servicing requirement for future large, space-based observatories, let alone to reap the benefits of autonomous systems for observation.

Limited institutional capacity to adopt autonomous systems is exemplified by the examples of autonomous systems above: most of these missions were or will be designed and built by NASA. Given the high cost and risk associated with changing the process by which spacecraft are designed, built, and tested, spacecraft vendors have till now relied on conditional sequencing and not autonomous planning/execution. Thus, government-industry cooperation is required to make use of autonomous systems reliably and repeatably for all NASA missions.

Recommendation: NASA should incentivize the use of autonomous systems for competed space missions. Specifically, small sat missions, missions of opportunity, SMEX, MIDEX, and Probe missions can include a credit for using the technology. *We note that transient event observatories offer a low risk path to maturing this critical technology.*

Another aspect of the cultural gap is NASA’s definition of technology readiness and its reference to “operational environment” that is overly

restrictive to software technologies that can be effectively validated outside of the operational environment, *e.g.* on-board science data processing software.

Recommendation: NASA should evaluate the applicability of the Technology Readiness Level as a means of evaluating the maturity of autonomy and on-board data processing software.

Other gaps are technical. Autonomy frameworks, described in Appendix A, define rules for how system-level planning and execution interface with traditional components and systems and functional autonomy. Community acceptance of these frameworks can reduce adoption risk and promote repeatability by permitting traditional design and operations approaches. By defining a convention for how autonomous missions should be designed and built, frameworks would also improve reviewability of autonomous missions and portability of testing methodology. Remaining work includes improving the verifiability of tasknets, which is critical to reaping the benefits of integrated fault protection. Relatedly, telemetry that permits reconstruction of on-board decision-making requires further study and definition. Ground systems and tools for commanding of autonomous spacecraft require further maturation.

Finally, some observations will benefit from on-board data analysis. For these observations, new software will be required to perform this function, which will be the responsibility of science community. While not the subject of this white paper, processing-intensive data processing may require high performance computing. High performance computing does not necessarily enable autonomous systems, but is enhancing by permitting intensive processing of science data and on-board scheduling over larger search spaces.

Recommendation: NASA should fund technology demonstrations of high-performance space-based computing for on-board data processing.

A Example of a System Autonomy Framework

In order to create autonomous systems repeatedly and reliably, a framework must be defined. Similar to a legal constitution, a system autonomy framework defines responsibilities and capabilities of system components and how they interface with one another in governing system behavior. For instance, the manner in which on-board science data processing is interfaced to inform decisions at a system-level should be identical across all missions regardless of which decisions it informs. This enables reliable multi-mission adoption of on-board planning/execution. Under a unified framework, engineers, scientists, and managers can work toward the same set of requirements that assure mission success. Reviewers can also use the same framework to verify compliance. Without a unified framework, development, mission assurance, and review can become intractable given the complexity in designing autonomous systems.

There are several requirements that define an effective autonomy framework. It should make guarantees about acceptable behavior, enable confident operator oversight and insight, readily accommodate new information, and not require extensive tailoring or *ad hoc* modification to support multiple missions. Pragmatically, such a framework must afford a practical path toward adoption. To do so, system autonomy must integrate with existing components. Human workflows involved across mission phases should deviate minimally from existing practice. Also, the framework must support varying degrees of autonomy – permitting sequence-like commanding to highly autonomous operation. Without this practical path, NASA and industry partners will have to invest in brand new software and processes and accept significant risk in implementing a major leap toward systems autonomy at once.

There are several examples of such a framework. The Framework for Robust Execution and Scheduling of Commands On-Board (FRESCO) is under development at NASA JPL; the NASA Platform for Autonomous Systems (NPAS) is

under development at NASA SSC; and, a vehicle management system for autonomous spacecraft habitat operations is under development at NASA ARC and JSFC [13, 52, 53]. Finally, the European Robotic Goal-Oriented Autonomous Controller (ERGO) has been under developed by an EU-funded consortium of industry and academia [54]. Below, we use FRESCO as an example to illustrate how systems autonomy is implemented.

Tasknet Tasknets are data structures that encapsulate the potential envelop of spacecraft behavior. They are graph networks where nodes are tasks and the edges are the state and temporal dependencies between tasks. Tasknets can be defined as goals for the spacecraft to achieve, to transition states (*e.g.* an imaging survey goal results in a set of images being taken) and to maintain states (*e.g.* a pointing knowledge maintenance goal restores pointing knowledge through optical navigation if a knowledge uncertainty threshold is violated). Tasknets have been described in literature since the 1970s [55, 56].

Planner and Executive Planners create and maintain schedules of tasknets based on their prioritization and projected timelines of future states. Scheduling tasks is performed by a search function, whose search space can be constrained based on how tasks are defined. This permits traditional, sequence-like behavior or highly autonomous behavior within the same framework. At a certain time before scheduled execution, the planner passes tasks to the executive for execution. Executives are responsible for intelligent execution and monitoring the impact of executed tasks. Under nominal operation, they receive receipt of successful task execution and proceed to dispatch the next scheduled tasks. If a task fails, they can exercise contingency behaviors specified by the task, which can include replanning requests to the planner.

Currently, MEXEC and PLEXIL are maintained by NASA JPL and ARC, respectively [24, 57]. A wider survey of command execution systems is presented in [58].

State Database A state database serves as a “single source of truth” for the system, maintaining component status and abstracted system

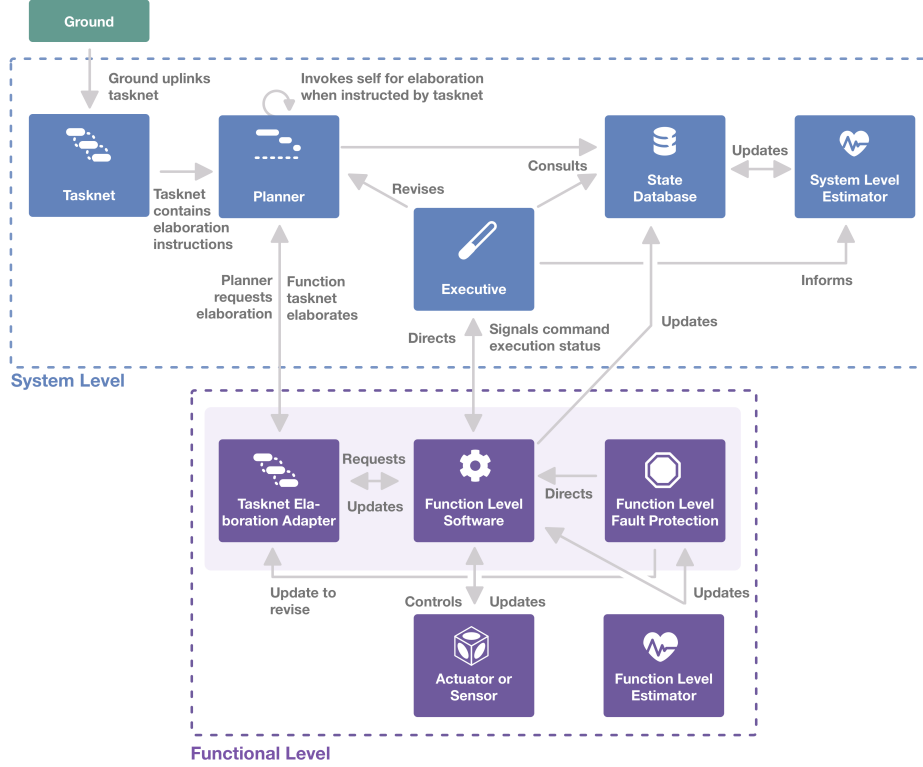


Figure 4: An example of an autonomous system framework, the Framework for Robust Execution and Scheduling of Commands On-Board (FRESKO), defines capabilities and interfaces that will resulting in repeatable and predictable implementations of systems autonomy for complex systems, such as spacecraft. This figure offers a simplified description of FRESKO components and interfaces.

states used in decision-making.

System-Level Estimator Estimators that inform decision-making use system telemetry as input to models of system behavior. System health monitoring software is one such estimator. System health monitoring serves two purposes. First, it can be used to identify potentially faulty components to alert operators to potential future risks. In rover testing, it was able to successfully identify undiagnosed hardware problems[51]. Second, it permits the creation of tasknets that operate only if healthy component states are reported, reducing the risks of autonomous operation. MONSID, developed by Okean Solutions, uses linked models of hardware behavior to monitor the health status of components [48, 59, 51].

Function-Level Software and Components Function-level software can perform a multitude of functions, ranging from hardware control to data processing and functional autonomy. Depending on its function, its internal in-

terfaces will vary. For instance, if a hardware controller includes local fault protection, an interface for a signal to interrupt system-level execution over that controller’s domain is required. Traditional hardware controllers and on-board data processing can also be used to inform the scheduling of tasknets. The on-board processing for exozodiacal light in exoplanet coronagraphy serves as an example.

B Acknowledgements

Thank you to Ellen Van Wyk, NASA JPL, for illustrations included in this white paper and SpaceX for the cover photograph.

References

- [1] NASA Autonomous Systems – Systems Capability Leadership Team. Autonomous systems taxonomy, 2019.
- [2] CF Kerry and J Karsten. Gauging investment in self-driving cars.
- [3] Space Studies Board, National Research Council, et al. *Vision and voyages for planetary science in the decade 2013-2022*. National Academies Press, 2012.
- [4] Steve Chien and Kiri L Wagstaff. Robotic space exploration agents. *Science Robotics*, 2(7):eaan4831, 2017.
- [5] Science Committee on Commerce and Transportation. *National Aeronautics and Space Administration Authorization Act of 2010: report of the Committee on Commerce, Science, and Transportation on S. 3729*. U.S. G.P.O., 2010.
- [6] Rudranarayan Mukherjee et al. When is it worth assembling observatories in space? July 2019.
- [7] Kimberly F Robinson and Andrew Schorr. Nasa’s space launch system: Exceptional opportunities for secondary payloads to deep space. In *2018 AIAA SPACE and Astronautics Forum and Exposition*, page 5138, 2018.
- [8] V Vitaldev. Automated scheduling and operation of a heterogeneous fleet of satellites. IWPSS 2019, 2019.
- [9] Jason Glenn, Charles M Bradford, Rashied Amini, Katey Alatalo, Lee Armus, Andrew Benson, Duncan Farrah, Adalyn Fyhrie, Sarah Lipsy, Bradley Moore, et al. The galaxy evolution probe: a concept for a mid and far-infrared space observatory. In *Space Telescopes and Instrumentation 2018: Optical, Infrared, and Millimeter Wave*, volume 10698, page 106980L. International Society for Optics and Photonics, 2018.
- [10] S Hanany. Pico probe of inflation and cosmic origins. The Space Astrophysics Landscape for 2020 and Beyond, April 2019.
- [11] T-C Chang for the CDIM Team. Cdim cosmic dawn intensity mapper. The Space Astrophysics Landscape for 2020 and Beyond, April 2019.
- [12] *Deep Space-1. Navigation: Primary Mission*. JPL: DESCANSO Design and Performance Summary Series, April 2004.
- [13] Lorraine Fesq and Rashied Amini. A new paradigm for autonomous spacecraft: From research to deployment. Low Cost Planetary Missions (LCPM-13), June 2019.
- [14] T Weismuller and M Leinz. Gn&c technology demonstrated by the orbital express autonomous rendezvous and capture sensor system. In *29th annual AAS guidance and control conference*, pages 06–016. American Astronautical Society, 2006.
- [15] R Ticker. Restore-l mission information.
- [16] P Maldague, A Ko, D Page, and T Starbird. Apgen: A multi-mission semi-automated planning tool. In *First international NASA workshop on planning and scheduling*, pages 363–365, 1998.
- [17] Christopher Grasso and Patricia Lock. Vml sequencing: Growing capabilities over multiple missions. In *SpaceOps 2008 Conference*, page 3295, 2008.
- [18] D Palmer. Grb follow up observation autonomy. In correspondence, June 2019.
- [19] David S Mittman and Robert Hawkins. Scheduling spitzer: the sirpass story. 2013.
- [20] S. Chien, G. Rabideau, R. Knight, R. Sherwood, B. Engelhardt, D. Mutz, T. Estlin, B. Smith, F. Fisher, T. Barrett, G. Stebbins, and D. Tran. Aspen - automating space mission operations using automated planning and scheduling. In *International Conference on Space Operations (SpaceOps 2000)*, Toulouse, France, June 2000.

- [21] Daniel Tran, Steve Chien, Rob Sherwood, Rebecca Castano, Benjamin Cichy, Ashley Davies, and Gregg Rabideau. The autonomous sciencecraft experiment onboard the eo-1 spacecraft. In *Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, pages 163–164. ACM, 2005.
- [22] Steve Chien, Joshua Doubleday, D Thompson, K Wagstaff, John Bellardo, Craig Francis, Eric Baumgarten, Austin Williams, Edmund Yee, Daniel Fluit, et al. Onboard autonomy on the intelligent payload experiment (ipex) cubesat mission: A pathfinder for the proposed hypsiri mission intelligent payload module. In *Proc 12th International Symposium in Artificial Intelligence, Robotics and Automation in Space, Montreal, Canada*, 2014.
- [23] Lorraine Fesq, Patricia Beauchamp, Amanda Donner, Rob Bocchino, Brian Kennedy, Faiz Mirza, Swati Mohan, David Sternberg, Matthew W Smith, Martina Troesch, et al. Extended mission technology demonstrations using the asteria spacecraft. In *2019 IEEE Aerospace Conference*, pages 1–11. IEEE, 2019.
- [24] Vandt Verma, Dan Gaines, Gregg Rabideau, Steve Schaffer, and Rajeev Joshi. Autonomous science restart for the planned europa mission with lightweight planning and execution. In *International Workshop on Planning and Scheduling for Space (IWPSS 2017)*, Pittsburgh, PA, 2017.
- [25] Gregg Rabideau and Ed Benowitz. Prototyping an onboard scheduler for the mars 2020 rover. In *Proceeding of International Workshop on Planning and Scheduling for Space, Pittsburgh, PA*, 2017.
- [26] Philippe Labone and Malik Ghallab. Planning with sharable resource constraints. In *Proceedings of the 14th international joint conference on Artificial intelligence*, volume 2, pages 1643–1649. Citeseer, 1995.
- [27] Javier Barreiro, Matthew Boyce, Minh Do, Jeremy Frank, Michael Iatauro, Tatiana Kichkaylo, Paul Morris, James Ong, Emilio Remolina, Tristan Smith, et al. Europa: A platform for ai planning, scheduling, constraint programming, and optimization. *4th International Competition on Knowledge Engineering for Planning and Scheduling (ICKEPS)*, 2012.
- [28] Conor McGann, Frederic Py, Kanna Rajan, Hans Thomas, Richard Henthorn, and Rob McEwen. T-rex: A model-based architecture for auv control. In *3rd Workshop on Planning and Plan Execution for Real-World Systems*, volume 2007, 2007.
- [29] Nikolaus Volgenau and Todd Boroson. Two years of lcogt operations: the challenges of a global observatory. In *Observatory Operations: Strategies, Processes, and Systems VI*, volume 9910, page 99101C. International Society for Optics and Photonics, 2016.
- [30] T Boroson, T Brown, A Hjelstrom, DA Howell, T Lister, A Pickles, W Rosing, E Saunders, R Street, and Z Walker. Science operations for lcogt: a global telescope network. In *Observatory Operations: Strategies, Processes, and Systems V*, volume 9149, page 91491E. International Society for Optics and Photonics, 2014.
- [31] RA Street, M Bowman, ES Saunders, and T Boroson. General-purpose software for managing astronomical observing programs in the lsst era. In *Software and Cyberinfrastructure for Astronomy V*, volume 10707, page 1070711. International Society for Optics and Photonics, 2018.
- [32] A Bridger. The alma observing tool: Proposal preparation submission. In *ALMA Early Science Workshop, IRAM*. ALMA, 2010.
- [33] Mark D Johnston. Spike: Ai scheduling for nasa’s hubble space telescope. In *Sixth Con-*

- ference on Artificial Intelligence for Applications*, pages 184–190. IEEE, 1990.
- [34] PK Barnes, AT Haddock, and CA Cruzen. Autonomous science operations technologies for deep space gateway. In *Deep Space Gateway Concept Science Workshop*, volume 2063, 2018.
 - [35] Roger M Smith, Richard G Dekany, Christopher Bebek, Eric Bellm, Khanh Bui, John Cromer, Paul Gardner, Matthew Hoff, Stephen Kaye, Shrinivas Kulkarni, et al. The zwicky transient facility observing system. In *Ground-based and Airborne Instrumentation for Astronomy V*, volume 9147, page 914779. International Society for Optics and Photonics, 2014.
 - [36] Stephen B Cenko. The gravitational-wave ultraviolet counterpart imager (guci) network. In *American Astronomical Society Meeting Abstracts*, volume 234, 2019.
 - [37] Judith Racusin. Current and future large missions as science motivation for grb detecting smallsats. In *Towards a Network of GRB Detecting Nanosatellites*, 2018.
 - [38] Rashied Amini. Utilizing fractionated space mission design and small satellites for a next generation gamma ray burst observatory. In *22nd SmallSat Conference*, 2008.
 - [39] Sreeja Nag, Alan S Li, Vinay Ravindra, Marc Sanchez Net, Kar-Ming Cheung, Rod Lammers, and Brian Bledsoe. Autonomous scheduling of agile spacecraft constellations with delay tolerant networking for reactive imaging. In *11th International Workshop on Planning and Scheduling for Space*, 2019.
 - [40] William Gerstenmaier and Jason Crusan. Cislunar and gateway overview.
 - [41] Jack Burns, Gregg Hallinan, et al. Cdim cosmic dawn intensity mapper. The Space Astrophysics Landscape for 2020 and Beyond, April 2019.
 - [42] B Mennesson, R Millan-Gabet, E Serabyn, MM Colavita, Olivier Absil, G Bryden, M Wyatt, W Danchi, Denis Defrere, O Doré, et al. Constraining the exozodiacal luminosity function of main-sequence stars: complete results from the keck nuller mid-infrared surveys. *The Astrophysical Journal*, 797(2):119, 2014.
 - [43] Christopher C Stark, Aki Roberge, Avi Mandell, Mark Clampin, Shawn D Domagal-Goldman, Michael W McElwain, and Karl R Stapelfeldt. Lower limits on aperture size for an exoearth detecting coronagraphic mission. *The Astrophysical Journal*, 808(2):149, 2015.
 - [44] Travis Imken, Thomas Randolph, Michael DiNicola, and Austin Nicholas. Modeling spacecraft safe mode events. In *2018 IEEE Aerospace Conference*, pages 1–13. IEEE, 2018.
 - [45] Gordon B Aaseng, Adam Sweet, and John Ossenfort. Performance analysis of an autonomous fault management system. In *2018 AIAA SPACE and Astronautics Forum and Exposition*, page 5149, 2018.
 - [46] Jeremy D Frank and Gordon B Aaseng. Transitioning autonomous systems technology research to a flight software environment. In *AIAA SPACE 2016*, page 5530. 2016.
 - [47] Ksenia Kolcio, Louis Breger, and Paul Zetocho. Model-based fault management for spacecraft autonomy. In *2014 IEEE Aerospace Conference*, pages 1–14. IEEE, 2014.
 - [48] Ksenia Kolcio and Lorraine Fesq. Model-based off-nominal state isolation and detection system for autonomous fault management. In *2016 IEEE Aerospace Conference*, pages 1–13. IEEE, 2016.
 - [49] Daniel Gaines, Robert Anderson, Gary Doran, William Huffman, Heather Justice, Ryan Mackey, Gregg Rabideau, Ashwin

- Vasavada, Vandana Verma, Tara Estlin, et al. Productivity challenges for mars rover operations. In *The 26th International Conference on Automated Planning and Scheduling*, 2016.
- [50] Daniel Gaines, Robert Anderson, Gary Doran, William Huffman, Heather Justice, Ryan Mackey, Gregg Rabideau, Ashwin Vasavada, Vandana Verma, Tara Estlin, et al. Productivity challenges for mars rover operations. In *The 26th International Conference on Automated Planning and Scheduling*, 2016.
- [51] Ksenia Kolcio, Ryan Mackey, and Lorraine Fesq. Model-based approach to rover health assessment-mars yard discoveries. In *2019 IEEE Aerospace Conference*, pages 1–12. IEEE, 2019.
- [52] Fernando Figueroa, Mark Walker, and Lauren W Underwood. Nasa platform for autonomous systems (npas). In *AIAA Scitech 2019 Forum*, page 1963, 2019.
- [53] Richard Levinson, Jeremy D Frank, Michael Iatauro, Adam Sweet, Gordon B Aaseng, Michael Scott, John Ossenfort, James Soeder, Tam Ngo, Zachary Greenwood, et al. Development and testing of a vehicle management system for autonomous spacecraft habitat operations. In *2018 AIAA SPACE and Astronautics Forum and Exposition*, page 5148, 2018.
- [54] J Ocón, FJ Colmenero, J Estremera, K Buckley, M Alonso, E Heredia, J Garcia, AI Coles, AJ Coles, M Martinez, et al. The ergo framework and its use in planetary/orbital scenarios. In *Proc. 69th International Astronautical Congress (IAC), IAF, Bremen, Germany*, 2018.
- [55] Richard E. Fikes and Nils J. Nilsson. Strips: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2(3):189 – 208, 1971.
- [56] SA Vere. Deviser-an ai planner for spacecraft operations. 1985.
- [57] Vandi Verma, Ari Jónsson, Corina Pasareanu, and Michael Iatauro. Universal-executive and plexil: engine and language for robust spacecraft control and operations. In *Space 2006*, page 7449. 2006.
- [58] Vandi Verma, Ari Jónsson, Reid Simmons, Tara Estlin, and Rich Levinson. Survey of command execution systems for nasa spacecraft and robots. 2005.
- [59] Ksenia Kolcio, Lorraine Fesq, and Ryan Mackey. Model-based approach to rover health assessment for increased productivity. In *2017 IEEE Aerospace Conference*, pages 1–13. IEEE, 2017.