# Human and Robotic Space Mission Use Cases for High-Performance Spaceflight Computing

**Raphael Some, Richard Doyle,**
**Larry Bergman, William Whitaker**
***Jet Propulsion Laboratory***
***California Institute of Technology***

**Wesley Powell, Michael Johnson**
***NASA Goddard Space Flight Center***

**Montgomery Goforth**
***NASA Johnson Space Center***

**Michael Lowry**
***NASA Ames Research Center***

## Introduction

Spaceflight computing is a key resource in NASA space missions and a core determining factor of spacecraft capability, with ripple effects throughout the spacecraft, end-to-end system, and mission. Onboard computing can be aptly viewed as a "technology multiplier" in that advances provide direct dramatic improvements in flight functions and capabilities across the NASA mission classes, and enable new flight capabilities and mission scenarios, increasing science and exploration return.

Space-qualified computing technology, however, has not advanced significantly in well over ten years and the current state of the practice fails to meet the near- to mid-term needs of NASA missions. Recognizing this gap, the NASA Game Changing Development Program (GCDP), under the auspices of the NASA Space Technology Mission Directorate, commissioned a study on space-based computing needs, looking out 15-20 years. [7]

The study resulted in a recommendation to pursue high-performance spaceflight computing (HPSC) for next-generation missions, and a decision to partner with the Air Force Research Lab (AFRL) in this development.

## Formulation

To sharpen understanding of the processing gap, a multi-center NASA team conducted a study to address the following questions:

- What are the paradigm shifting NASA space-based applications that drive flight computing?
- What are the requirements imposed on flight computing by these applications?

This paper reports on the development of use cases and on the derivation and summarization of future flight computing requirements, performed during the NASA HPSC formulation activity.

## Use Cases

To identify the use cases, a series of workshops was held with mission designers, scientists and engineers from NASA Johnson Space Center (JSC), NASA Goddard Space Flight Center (GSFC) and the Jet Propulsion Laboratory (JPL). NASA Ames Reseach Center (ARC) and NASA Kennedy Space Center (KSC) also provided support. Both robotic and human spaceflight mission applications were examined.

These workshops defined the types of applications desired for future spacecraft and missions through the 2025 timeframe. Nineteen generic applications were identified that required significantly higher performance computing than currently available with current or planned space-qualified computers.

Table 1 lists the identified use cases, organized by the two primary NASA mission directorates: the Human Exploration Mission Operations Diretorate (HEOMD) and the Science Mission Directorate (SMD).

**Table 1** — Future NASA Mission Application Use Cases Requiring High Performance Spaceflight Computing

| HEOMD Use Cases | SMD Use Cases |
|---|---|
| Cloud Services | **Extreme Terrain Landing** |
| **Advanced Vehicle Health Management** | Close Proximity Operations/Formation Flying |
| Crew Knowledge Augmentation | Fast Traverse |
| Improved Displays and Controls | New Surface Mobility Systems |
| **Augmented Reality for Recognition and Cataloging** | Imaging Spectrometers |
| Tele-Presence | Radar |
| Automated Guidance, Navigation and Control | **Low Latency Products for Disaster Response** |
| Human Movement Assist | Space Weather |
| Autonomous & Tele-Robotic Construction | Autonomous Mission Planning |
| | Immersive Environments - Science Ops/Outreach |

Space limitations preclude offering a description of all 19 NASA use cases listed, but to provide a sampling of the range of space-based capabilities enabled by HPSC, the use cases accented in **bold** above are expanded on here.

Advanced Vehicle Health Management

Integrated Vehicle Health Management (IVHM) or Integrated System Health Management (ISHM) is meant to enable better management of vehicle health. This is achieved through correct use of reliable sensing to monitor component health combined with vehicle/system usage, knowledge of likely future events, and diagnostic and prognostic reasoning to optimize vehicle safety, reliability, availability, and maintenance efforts.

Attempts have been made by NASA to address IVHM, but they have not met with full success; the amount of sensor data and processing power available has never allowed an adequate implementation.

Providing this processing power will allow the monitoring of more sensors at higher frequencies and the comparison of this data over time and between areas of the vehicle, both of which will significantly enhance failure detection. Additional processing power will also allow the implementation of more sophisticated isolation and recovery functions, and the implementation of algorithms which predict impending failures, thus allowing preventative maintenance

A comprehensive implementation of IVHM will be extremely important to the success of future long duration exploration missions, including those such as crew-tended Cis-Lunar missions, which may involve a significant amount of un-crewed operations. IVHM will allow optimum use of crew time for maintenance by allowing the risk of failure to be balanced against the resources required for scheduled maintenance. Ultimately, IVHM capabilities will evolve to the implementation of a "Digital Twin", per the OCT's Modeling Simulation, Information Technology and Processing Technology Area Strategic Roadmap.

Augmented Reality for Recognition and Cataloging

Augmented reality (AR) is a real-time view of a physical, real-world environment whose elements are augmented by computer-generated sensory input such as sound, video, graphics or GPS data. These augmentations may be based on background data known to the system, or they may be based on non-visual sensor inputs. Augmented reality may be presented using a heads-up display, a hand-held device, or other display capability.

The use of augmented reality to support crew activities will significantly improve their efficiency and effectiveness. A "heads-up" display overlaying multi-spectral images of the surrounding environment, along with context-sensitive information

about items within the field of view, could be used to locate interesting materials/samples on the surface of a body and avoid potentially hazardous areas. Similarly, augmented reality views of vehicles and equipment could be used to locate and troubleshoot problems more quickly, as well as to provide general situational awareness.

A vehicle support version of this capability could be implemented for, and upgraded during, crewed Cis-Lunar missions. An environment support version of this capability could be first implemented for a Near-Earth Asteroid, Lunar surface, or Mars mission.

Extreme Terrain Landing

Landing site selection is restricted by the need to find a safe site that is also scientifically interesting. Reducing the landing site uncertainty, i.e., size of the "landing ellipse", and increasing the mechanical hazard tolerance of the lander, are two methods employed to increase the number of selectable landing sites. However, even with these improvements, many of the sites desired for future spacecraft landing, e.g., for collection and return of Mars samples from diverse geological contexts, are not safely accessible. This is due, in part to the fact that geological diversity is correlated with hazardous landing sites due to the desire to have stratigraphy and outcrops in the sampled terrain. This desire for landing near "interesting" terrain drives future planetary and primitive body missions. Similar requirements to deal with uncertainty in close proximity operations about comets and asteroids, and rendezvous/docking with potentially uncontrolled spacecraft crop up in numerous future mission scenarios. The highest degree of real time criticality, coupled with high rate sensor, guidance and navigation processing, however, is evidenced in planetary landing, especially on a body with significant but thin atmosphere, such as Mars.

Robotic planetary lander missions to date cannot "see" hazards or landmarks for navigation. In this sense they have all landed blindly. If the lander is equipped with hazard detection (HD) capability then hazards in the vicinity of the landing target can be identified and the vehicle diverted to avoid them. Futhermore, if the lander is equipped with a Terrain Relative Navigation (TRN) capability, then it can recognize landmarks and compute map-relative position, which can then be used in two different ways. First, if there is enough fuel, the lander can be guided to a pinpoint landing (landing ellipse < 100-m radius). Second, if the vehicle is limited in fuel, then a multiplicity of safe landing sites can be found within the landing area and the lander guided to the closest safe site.

Both TRN and HD require processing beyond the capabilities of current flight processors. Benchmarks of TRN software on a RAD 750 showed that an update from a single camera took ~10 seconds. To achieve 100m TRN accuracy, an update every 1 second is required. To achieve 1m pinpoint landing TRN estimates to the surface

are required at 10 Hz. A two order-of-magnitude improvement in flight processor performance is required. [4]

Low-Latency Products for Disaster Response

Successes with science event detection and response on Earth-observing spacecraft, notably the Autonomous Sciencecraft capability that has been operational on the Earth-Observing One (EO-1) spacecraft since 2004, led to further thinking about whether event detection and response could be generalized across multiple platforms: flight, ground and possibly even air-based. An event of interest might be detected initially on any given platform, but the optimal follow-on observational response might best occur on a different platform with a better view or sensor.

Improved onboard computing capacity can greatly assist such objectives by: 1) enabling rapid turnaround of products for disaster response (hours/minutes instead of weeks), 2) automatic cueing of additional assets such as other space platforms or UAVs, not currently feasible due to time delays, 3) direct broadcast of evolving data products in the field for situational awareness and decision support, and 4) delivery of real-time high-definition video.

The specific onboard data processing capabilities required by these scenarios include fast spatial/spectral sub-setting, geo-rectification, atmospheric correction, rapid broadcasts in the event of library matches, along with the ongoing challenges of processing data streams approaching 1 Tb/sec (as is the case with the HyspIRI VSWIR/TIR hyperspectral instruments). Such intense onboard data processing would likely impact other aspects of the flight system; there would need to be sufficient memory and storage for both processing and archiving. Downlink capacity, although greatly ameliorated by onboard processing, would still need to be sufficient to support direct broadcast objectives.

Outcome

Based on these and the other use cases [1,2], a straightforward conclusion is that a high performance spaceflight computer will indeed be game changing because the capability will be needed for many planned space missions across the agency, and will enable new and dramatic mission applications that are strongly desired by advance mission planners and future users.

## Derived Requirements

To derive requirements for future NASA flight-based computing, follow-on discussions were held with various NASA personnel beyond the engineers, scientists and mission designers involved in the development of the use cases shown in Table 1. The HPSC team derived computing system and processor requirements for each

use case by characterizing the required computing, the environment, the criticality of the application, and the system constraints [3,5,6].

The requirements template developed for this purpose is shown here:

- Space Environment(s)

— Radiation environment at the time of application; e.g., geosynchronous (GEO), low-Earth orbit (LEO), deep space.

- Spacecraft Power Environment(s) / Constraint(s)

— Available power for avionics and computing, e.g., small spacecraft or rover with limited power availability (6 Watts processor power, 10-15 Watts computer power), medium sized spacecraft (7-12 Watts processor power, 15-30 Watts computer power), or large spacecraft with large power budget (>12-25 Watts processor power, >30-100 Watts computer power)

- Criticality/Fault Tolerance

— Is this application life or mission critical, must it operate through faults, can it tolerate errors if detected?

- Real-Time

— Does the application have a hard real-time deadline; if so, what is the required timing?

- Type(s) of Processing

— Algorithm kernel(s). Is it primarily e.g., digital signal processing (DSP), data base query, is it parallelizable, is it amenable to a data flow approach?

- Memory Access Pattern

— What is the primary data movement pattern, e.g., does it fetch data from memory once and then flow through a processing chain, or does it access data in a continuous random access pattern, or does it access sequential data in a continuous and regular pattern ?

- Duty Cycle

— What is the pattern of execution, e.g., is the application called continuously over a long period of time, is it called once and operate for only a short duration, is the application execution profile spiky and/or unpredictable?

- Data Rate

— What are the I/O and memory access data rates

Table 2 summarizes the application requirements set. As shown in the table, the application list from the study above was augmented with known applications from missions currently under development or in an early study phase.

The table is primarily organized vertically, by application throughput requirement. The first section of the table lists applications requiring one to ten giga-operations per second, the second section lists applications requiring ten to fifty giga-operations per second, and the third section lists applications requiring fifty to one hundred giga-operations per second.

Each application is then characterized with respect to its primary processing requirements by the following parameters:

Type of processing: digital signal processing (DSP), general purpose processing (GP), and parallelizability of algorithms (P). Applications requiring at least 20% of any of these types of computing were so noted.

Mission criticality: mission critical applications were assumed to require extremely high levels of reliability, and therefore, fault tolerance.

Power: available power for spacecraft computing (LP=low power, MP=medium power, HP=high power) per the requirements template above.

Radiation environment: radiation environment drives both the radiation tolerance and fault tolerance requirements of the processing system.

Memory access: R=random, S=sequential,

Duty cycle: C=continuous, S=spikey

Memory access rate: in Bytes per second

I/O rate: in bits per second

Table 2 Applications Requirements Summary

| Apps Summary | Processing | | | Mission | Power | | | Raiation | Memory | Duty | Memory | I/O |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | DSP | GP | P | Critical | LP | MP | HP | Environ's | Access | Cycle | Rate | Rate |
| Throughput = 1-10 GOPS | | | | | | | | | | | | |
| Autonomous Mission Planning | | X | X | X | X | X | X | All | R | C,S | 1GB/S | 100Mb/S |
| Disaster Response | X | X | | | X | | | LEO | R,S | C | 200,B/S | 1Gb/S |
| Hyspiri | X | X | X | | | X | | LEO | R,S | C | 200MB/S | |
| | | | | | | | | | | | | |
| Throughput = 10-50 GOPS | | | | | | | | | | | | |
| Fast Traverse | X | X | X | X | X | | | Mars | C | C | 500MB/S | |
| Extreme Terrain Landing | X | X | X | X | X | | | Mars, GEO | R, S | C | 50MB/S | |
| Adept | | X | X | | | X | | LEO, MEO | R, S | C | 375MB | |
| Optimum Observation | X | X | X | | X | X | X | All | R | C | 1GB/S | 1Gb/S |
| Space Weather | X | | X | | X | X | X | All | R | C | 125MB/S | 500Mb/S |
| Robotic Servicing | X | X | X | X | X | X | X | GEO | R | C,S | 125MB/S | |
| Cloud Service | X | X | X | | | | X | All | R,S | C,S | 6.25GB/S | 10Gb/S |
| Advanced ISHM | | X | X | | X | | X | All | R | C | 6.25GB/S | 10Gb/S |
| Autonomous and | | X | X | X | X | | X | GEO, Mars, | R | C,S | 12.5GB/S | 50Gb/S |
| Telerobotic Construction | | | | | | | | | | | | |
| Througput = 50-100s GOPS | | | | | | | | | | | | |
| Hyperspectral Imaging | X | X | X | | X | X | X | All | R,S | C,S | nGB/S | |
| RADAR Science | X | X | X | | | | X | All | R,S | C | 1GB/S | |
| RADAR EDL | X | | X | X | X | X | X | All | R,S | S | 1GB/S | |
| Automated GN&C | X | X | X | X | | | X | All | R,S | S | 12.5GB/S | 10GbS |
| Human Movement Assist | X | X | X | | X | | | All | R,S | S | 12.5GB/S | 50GbS |
| Crew Knowledge Augmentation | | X | X | | | | | All | R,S | S | 12.5GB/S | 10Gb/S |
| Improved Displays and Controls | | X | X | X | X | | X | All | R,S | C | 12.5BG/S | 50Gb/S |
| Augmented Reality | | X | X | | X | | X | All | R,S | S | 12.5GB/S | 50Gb/S |
| Telepresence | | X | X | | X | | X | All | R,S | S | 12/5GB/S | 50Gb/S |

As the table shows, each of the generic applications will be utilized in multiple missions and mission scenarios. On the first line, for example, it is seen that Autonomous Mission Planning is expected to be used on small, medium and large spacecraft, in a multiplicity of radiation environments. It is also shown that in at least some cases, this application will be mission critical. This leads to a situation that can be appropriately characterized as "requirement explosion", a situation which is unwieldy and difficult to handle with straightforward system analysis tools.

## Eigen-Applications

In performing the requirements analysis, over 60 variations of the application and use cases shown were examined. To reduce this to a manageable number for subsequent analysis, the most stressing and cross-cutting of these applications were condensed into ten representative applications, which we termed '"Eigen-Applications." Each Eigen-Application represents a broad class of mission applications with similar requirements. The Eigen-Apps collectively define the required advanced flight computing capability for future NASA applications. Table 3 summarizes the Eigen-Applications.

Table 3 Eigen-Applications

Eigen-Application Definitions

| Eigen-App | Throughput | DSP | GP | P | LP | MC |
|---|---|---|---|---|---|---|
| 1 | 1-10 GOPS | X | X | X | X | |
| 2 | 1-10 GOPS | | X | X | X | X |
| 3 | 10-50 GOPS | X | X | X | X | X |
| 4 | 10-50 GOPS | X | X | X | X | |
| 5 | 10-50 GOPS | | X | X | X | X |
| 6 | 10-50 GOPS | | X | X | X | |
| 7 | 50-100 GOPS | X | X | X | X | X |
| 8 | 50-100 GOPS | X | X | X | X | |
| 9 | 50-100 GOPS | | X | X | X | X |
| 19 | 50-100 GOPS | | X | X | X | |

Following a similar pattern to Table 2, Table 3 is divided into three primary sections by required throughput: one to ten giga-operations (eigen-apps 1-2), ten to fifty giga-operations (eigen-apps 3-6) and fifty to one hundred giga-operations (eigen-apps 7-10). Each eigen-app is then characterized by processing type – again, the 20% threshold was used to define processing type requirement (digital signal processing (DSP), general purpose processing (GP), and parallelizability of algorithm (P)). Only low power applications were included in the eigen-application specification as this was defined to be a key stressing requirement. Finally, mission criticality (MC) was noted. All applications were assumed to require a high degree of radiation tolerance. Memory and I/O requirements were not considered as these were deemed to be secondary considerations for the subsequent analysis.

While the Eigen-Applications do not represent all possible future spacecraft applications, they do cover the majority of advanced computing requirements for future NASA missions and specifically, the more stressing and more important mission applications.

## A Joint Investment

Based on the flight computing use cases and requirements developed during the HPSC formulation activity, with some additional supporting gap analysis and assessment of extant and emerging computing architectures, NASA made the decision in December 2012 to proceed into an implementation phase for the HPSC project. Another important factor was a series of discussions with AFRL/Kirtland that revealed similar interests and objectives concerning future flight computing capability, and first-order alignment of requirements. As part of the same decision meeting, AFRL and NASA entered into partnership and issued a joint BAA in April

2013 to solicit, under a study phase, architectural designs for a rad-hard general-purpose multicore flight computer [8,9] that addresses the derived requirements from the NASA study and additional capabilities defined by AFRL . At the time of this writing, the proposal period has closed and selections are pending.

## Next Steps

As the HPSC project continues, the NASA team is embarking on an effort to develop application and architectural benchmarks to be utilized to evaluate the hardware architecture designs that will emerge from the AFRL/NASA-funded study. These benchmarks can be seen as a logical extension of the work to develop use cases and derive requirements. They will provide a more thorough basis to evaluate the efficacy of designs, going beyond the letter of derived requirements to the details of algorithms and computational kernels and the subtleties of system-level support for energy management and fault tolerance, all of great importance to NASA applications. We plan to report on the development and use of these benchmarks in a future paper. We hope they will prove to be of high utility for the space flight computing community.

## Conclusions

We have reported here on the development of use cases and the derivation of requirements to assess and capture the flight computing needs of future NASA missions. Some of the mission applications examined in the HPSC formulation activity had already been identified within NASA programs. In these cases, our effort in the HPSC study served to articulate the relevant flight computing requirements with greater fidelity. Some of the identified mission applications, however, were new or significantly extended, and serve to broaden the base of NASA mission-pull for high performance spaceflight computing.

It has been more than 15 years since the previous investment at NASA in a flight computer. Our study provides a clear and broad basis for the decision that a next-generation flight computing capability will be a necessary and timely ingredient for future mission success. Partnered with AFRL, we are taking the first steps toward realizing that investment.

## Acknowledgments

The work described in this publication was performed at the Jet Propulsion Laboratory, California Institute of Technology, under contract from the National Aeronautics and Space Administration. The authors would like to thank Gabriel Mounce and Keith Avery of AFRL/Kirtland for their support of the agency-level partnership described herein.

## References


1. Tara Estlin, Benjamin Bornstein and Larry Matthies, "Software on Mars," *Communications of the ACM*, November 2012.

2. David R. Thompson, Daniel Q. Tran, David McLaren, Steve A. Chien, Larry Bergman, Rebecca Castano, Richard Doyle, Tara Estlin and Matthew Lenda, "Autonomous Onboard Science Data Analysis for Comet Missions," *11th International Symposium on Artificial Intelligence and Robotics Applications for Space,* Torino, Italy, September 2012.

3. James Alexander, Bradley J. Clement, Kim P. Gostelow, John Y. Lai, "Fault Mitigation Schemes for Future Spaceflight Multicore Processors, *AIAA Infotech*, Garden Grove, CA, June 2012.

4. James Alexander, Yang Cheng, William Zheng, Nikolas Trawny and Andrew Johnson, "A Terrain Relative Navigation Sensor Enabled by Multi-Core Processing," *IEEE Aerospace*, Big Sky, MT, March 2012.

5. Peter M. Kogge, Benjamin J. Bornstein and Tara A. Estlin, "Energy Usage in an Embedded Space Vision Application on a Tiled Architecture," *AIAA Infotech*, St. Louis, March 2011.

6. Kim P. Gostelow, "The Design of a Fault-Tolerant, Real-Time, Multi-Core Computer System," *AIAA Infotech*, St. Louis, March 2011.

7. Wesley A. Powell, Michael A. Johnson, Jonathan Wilmot, Raphael Some, Kim P. Gostelow, Glenn Reeves and Richard J. Doyle, "Enabling Future Robotic Missions with Multicore Processors," *AIAA Infotech*, St. Louis, March 2011.

8. Villalpando, Some, et al., "Reliable Multicore Processors for NASA Space Missions", Proceedings of the IEEE Aerospace Conference, March 2011.

9. Villalpando, Some, et al., "Investigation of the Tilera Processor for Real Time Hazard Detection and Avoidance on the Altair Lunar Lander", Proceedings of the IEEE Aerospace Conference, March 2010.