Research Statement

James Cutler
Stanford University
jwc@stanford.edu

I. Research Overview

The goal of my past research has been the development of space system infrastructure to make space science data more accessible and to enable experiments to have a wider impact on society. I have been inspired by the advent of web-based services that have sprung from the success of the terrestrial Internet and its ability to connect globally-distributed, heterogeneous computing resources. My research vision is shaped by this guiding question: "Can we extend our networking capabilities to include space systems so that scientists and other end users can access space data as simply and reliably as we access terrestrial web sites today?"

As an electrical engineer, I bridged the research programs of a computer science laboratory and an aerospace laboratory to develop the depth needed to understand the requirements and design of aerospace information systems. We found the complexity of space communication limits our access to space systems and elevates mission operational cost. Ultimately, this results in reduced mission capabilities and yields. In particular, ground stations, the access point between space and terrestrial networks, suffer from monolithic designs, narrow interfaces, and unreliability that raise significant financial barriers for low-cost, experimental satellite missions. With the breadth of my multidisciplinary background, I extended computer science techniques to develop technology for recovery-oriented, flexible access networks built from commercial-off-the-shelf (COTS) components. The result has been a reduction in these barriers where we reduced recovery time of typical ground station software failures by a factor of four and enabled low-cost ground support for experimental space missions. The contributions I have made are described below.

A. Small Satellite Technologies

Through my small satellite projects, I have revolutionized the capabilities of university satellites and demonstrated their ability to perform low-cost, high-risk science missions. Comparable to commodity Internet technologies such as web servers and databases, our technology simplifies the design and deployment of space missions. As the project manager, systems engineer, and payload designer for the Opal satellite project,[6] I extended the capabilities of deployable spacecraft. My colleagues and I developed novel storage and deployment mechanisms to survive the harsh environments of launch and space. We used them to successfully launch six picosatellites from the Opal mothership into low-Earth orbit. The Aerospace Corporation has extended our concept to develop "on-board, on-call" inspector satellites that improve long-term satellite operations. The Opal project was funded by DARPA, JPL, and NASA.

Partnering with the California Polytechnic State University, we expanded the picosatellite concept into a standard that describes 1 kg, 10 cm cubic satellites called CubeSats and their deployment mechanism, the P-POD.[8] As a systems engineer and communication system designer, I helped develop the world's first orbiting CubeSat, Stanford's QuakeSat-1, whose primary mission is to monitor low-frequency electromagnetic signals associated with pre- and post-earthquake activities. With COTS components, QuakeSat-1 has been operational since June 2003, far surpassing its six month expected life. With the introduction of CubeSats, the university satellite community has grown from 20 to over 50 universities. Our efforts have enabled low-cost launch opportunities and the development of commercial commodity solutions. CubeSats have dramatically enhanced university space programs and greatly expanded a university's ability to reach space.

B. A Framework for Robust and Reliable Space Communication Systems

Based on my extensive small satellite experiences, I decomposed ground station services and captured them in an extensible framework that has simplified reuse of ground station services and improved portability across heterogeneous installations.[3,4] This capability, combined with selective customization through virtual machine technology, allowed us to deliver a "just in time" ground station for QuakeSat-1 at a fraction of the price of current commodity solutions.

My decomposition work was informed by principles of robust system design[5] and governed by the recovery-oriented computing (ROC) philosophy. ROC's perspective is that hardware faults, software bugs, and operator errors are facts to be coped with, not problems to be solved. Emphasis is on recovery, not solely on fault avoidance. Hence, our ground station reference implementation called Mercury was a candidate for a ROC mechanism called recursive recovery (RR),[2] a high availability technique whose effectiveness in reducing recovery time has been demonstrated on research prototypes of Internet server systems.[1] Augmenting Mercury to implement RR and extending RR concepts reduced recovery time of typical ground station software failures by a factor of four, dropping recovery time to within the "window of recovery" and effectively eliminating the adverse effects of these failures. Since the time of failures cannot be predicted, RR allowed us to mitigate the effects of transient software failures and greatly reduce their potential impact on ground station operations.

Crucial to the successful application of RR and other similar techniques is the concept of a "window of recovery", the time period in which a failure can be recovered without affecting system performance. In the ground station, radio bandwidths and antenna beamwidths provide margins of error in which the system can recover before communication is disrupted. An aspect of my future work will determine if these techniques can be applied to systems where the window is extremely small or nonexistent such as those found in real-time, aerospace control systems.

C. Global Ground Station Network

My decomposition work abstracted core services of a ground station network and captured them in an XML-based command control language called the Ground Station Markup Language (GSML). GSML provides a layered interface to a *virtual ground station* that abstracts away device heterogeneity in ground stations. It removes the constraints of current narrow ground station interface standards by providing layered access to low-level ground station hardware and high-level automation services. GSML dissolves monolithic architectures by enabling user-level customization of ground station services to best meet mission requirements.

We have also extended our concept of a virtual ground station to include components that span physically-separate ground station installations. Tight coupling at the physical layer of large, homogeneous antenna arrays such as NASA's Deep Space Network increases signal-to-noise ratios. Our interest is in loosely-coupled stations under different administrative domains that coordinate to provide additional benefits. Globally-distributed teams increase communication contact windows with satellites. Locally clustered teams leverage the strengths of independent stations and harness redundancy across the team during failures. We have demonstrated application-level coordination between stations with stations in Alaska and California communicating with QuakeSat-1. Others have demonstrated network-level coordination through MobileIP.[9] We are currently exploring low-level handoffs and multiplexing mechanisms that work independently of application and network-level implementations to enable more ad hoc handoffs.

We have developed a GSML reference implementation for ground station network control that harnesses the benefits of COTS components, including rapid prototyping and deployment, while overcoming the challenges of COTS reliability and mission critical usage. Just as our CubeSat development work led to a drastic increase in university satellite research programs, our ground station research has laid a foundation for increased reliable and flexible access to satellites. Our early ground station network prototype, the Mercury Ground Station Network (MGSN), is harnessing the idle resources of heterogeneous ground stations at the University of Würzburg in Germany, the Norwegian University of Science and Technology (NTNU) in Norway, and Aalborg University in Denmark. The MGSN is a fertile testbed for exploring automated space system technologies and an excellent pedagogical tool for training future space system and networking engineers.

II. Research Philosophy

The research problems I face are multifaceted and multidimensional, requiring in-depth, cross-disciplinary knowledge to obtain solutions. The ground station decomposition and modeling work required extensive knowledge of high-level satellite operational tasks and low-level ground station implementation details. Application of ROC techniques to Mercury required application-level knowledge of ground stations to understand their ability to tolerate restarts at multiple levels as well as lower-level systems knowledge of ROC principles. By far, the QuakeSat-1 ULF work has been the most multifaceted. Our work requires a fundamental grasp of electromagnetic wave propagation, state-of-the-art electronic design for low-noise systems, and geological understanding of the wave sources. I expect future systems research problems to be similarly multifaceted. To meet the wide range of challenges in my future research, I anticipate working with a multidisciplinary team to leverage their diverse specialties and backgrounds.

Another pillar of my research is building real systems deployed in functional and working environments. Starting with theoretical foundations and extending them when needed, I have played a significant role in nine orbiting satellites, and my ground station technology is currently providing mission critical support for several missions. These real world systems provide first-hand access to raw data, which is crucial to validate the theoretical underpinning of the research projects.

I am also a strong proponent of open-source research projects. Open-source promotes collaboration that can lead to novel research directions and contributions. Open-source projects allow others to more easily extend software systems. To this end, I have released all of my ground station control software, and it is actively used by a growing team of international developers.

III. Future Research Directions

The focus of my research will continue to be reliable distributed systems with a particular emphasis on aerospace informational systems. Two potential future projects are described below.

A. ROC-Hardened Space Computing

Failures are to be expected in the harsh radiation environment of space. Traditionally, systems are built with special radiation-hardened components that are more expensive and several generations behind current terrestrial components. I propose to apply similar high-availability techniques as used in my prior research to develop a ROC-hardened spacecraft computing platform. ROC-hardening is complementary to radiation-hardening: we detect common transient failures such as those induced by radiation, we rapidly coerce all such failures to crashes (to maximize the likelihood of fail-stop behavior), and we design the system to recover safely and rapidly from this single failure type. In essence, we make unexpected crashes the expected case. Our group has shown the effectiveness of this approach in a state storage system.[7]

My future research will develop a ROC-hardened spacecraft computing platform based on commodity Internet middleware for hosting onboard science applications. The three primary functions of the platform will be fault detection, containment, and recovery. Rapid fault detection is an integral part of recovery. Programmer-inserted assertions, periodic consistency checks, and run-time audits on data structures can be deployed during development time. Mechanisms orthogonal to applications such as watchdog timers and heartbeats provide external monitoring by independent agents. Fault containment mechanisms reduce fault propagation and enable localized recovery. Physical isolation boundaries such as clustering limit many failures to single computer systems. Software mechanisms such as virtual machine monitors isolate multiple operating systems running on the same platform. Microkernels provide a mechanism for intra-operating system containment. Fault recovery mechanisms rely on some form of redundancy in time, data, or functionality. Functional redundancy such as Tandem processor pairs and clustering enable recovery of failed CPU components. Log-based recovery, checkpointing, and transactions provide time redundancy to roll back failed systems and retry calculations. System level architectures such as ARMORs[10] and ROC's J2EE extensions provide middleware recovery primitives and automated services to application developers. We will investigate the extent to which ROC-hardening can mitigate radiation-induced failures in non rad-hardened systems.

B. Magnetospheric/Ionospheric Wave Event Research

Building off of my QuakeSat-1 development and science studies, I am working with Quakefinder, LLC, and a variety of international researchers to study magnetospheric/ionospheric wave events and perturbations. We have recently developed an algorithm to study coherent wave events in three axis data. We have demonstrated the technique on a class of geomagnetic micropulsations called Pc1s using triaxial search-coil magnetometers in California. The algorithm identifies specific wave events and calculates wave parameters such as polarization and wave normal angles. We are currently applying this algorthim to study the long term behavior of Pc1 events in mid latitude regions.

Through ground and space-based measurements, I am studying whether ionospheric perturbations and electromagnetic anomalies preceed earthquakes. Past research has shown various anomalous signals, such as ultra low-frequency (ULF) signal increases, correlated in time with earthquakes. I am analyzing data from a large ULF receiver network deployed in California. I am working with with Dr. Michel Parrot (CNES) to analyze data from the DEMETER satellite mission. I am working with Dr. Eric Calais, Dr. Jennifer Haase, and Dr. Jim Garrison at Purdue University to study ionospheric electron content measurements derived from GPS reciever data. I have recently submitted a proposal to NSF to fund this work.

I am also program manager at Quakfinder for our follow-on space mission, QuakeSat-2. I am developing new technologies for measurement of low-frequency electromagnetic signals in the near-Earth space environment. We are working with Dr. Umran Inan at Stanford's VLF research group and Dr. Vitaly Chmyrev at Geoscan to develop an extremely low-noise search coil receiver system.

Also integral to our work is data fusion, the ability to analyze multiple data streams for correlations to provide greater insight into regional wave events. I am partnering with NASA Ames on a proposal to develop an intelligent sensor web framework to support this data fusion. Our project technologies will enable two-way interaction between the modeling / assimilation system and the sensing system to verify the correctness of sensor web data streams, assess earthquake risk and feedback an alert status to the sensor web. We will ingest data from several sources and process it for our needs using the controller for our study area, California. We will create a server for unified earthquake data access and produce preliminary visual products, which will enable timely production of earthquake risk maps.

IV. Conclusions

Due to the Challenger and Columbia tragedies and the insufficient use of the International Space Station, our generation has lost sight of a potential space age. However, recent advances, such as the X-Prize winner SpaceShipOne, are rekindling this dream. Space exploration is no longer confined to the realm of governmental institutions. My work has helped to throw open the doors for university research programs to perform end-to-end mission experimentation and technology development in space. Complimenting my future research will be a training program for future space system engineers. It will contain core electrical and aerospace material and a system-centric project component to provide experience with the multidisciplinary nature of building space systems. My experience with the design and launch of nine orbiting satellite missions and the building a global ground station network uniquely positions me to train future space system engineers. Given the creative potential of universities, I look forward to what university space programs will produce in the future for space exploration.

References

[1] George Candea, James Cutler, and Armando Fox. Improving availability with recursive micro-reboots: A soft-state system case study. In *Performance Evaluation Journal*, Summer 2003.

[2] George Candea and Armando Fox. Recursive restartability: Turning the reboot sledgehammer into a scalpel. In *Proc. 8th Workshop on Hot Topics in Operating Systems*, pages 110–115, Elmau/Oberbayern, Germany, May 2001.

[3] James Cutler. Ground station virtualization. In *Proceedings of the The Fifth International Symposium on Reducing the Cost of Spacecraft Ground Systems and Operations (RCSGSO)*, Pasadena, CA, July 2003.

[4] James W. Cutler. The ground station markup language. In *Proc. IEEE Aerospace Conference*, Big Sky, Montana, March 2004.

[5] James W. Cutler, Armando Fox, and Kul Bhasin. Applying the lessons of internet services to space systems. In *Proc. IEEE Aerospace Conference*, Big Sky, Montana, March 2001.

[6] James W. Cutler and Gregory Hutchins. Opal: Smaller, simpler, luckier. In *Proc. AIAA Small Satellite Conference*, Logan, Utah, September 2000.

[7] Benjamin C. Ling and Armando Fox. A self-tuning, self-protecting, self-healing session state management layer. In *5th Annual Workshop On Active Middleware Services*, Seattle, WA, 2003.

[8] Isaac Nason, Jordi Puig-Sair, and Robert J. Twiggs. Development of a family of picosatellite deployers based on the cubesat standard. In *Proceedings of IEEE Aerospace Conference*, Big Sky, Montana, 2002.

[9] OMNI. Operating missions as nodes on the internet, 2002. `http://ipinspace.gsfc.nasa.gov`.

[10] K. Whisnant, S. Bagchi, B. Srinivasan, Z. Kalbarczyk, and R.K. Iyer. Incorporating reconfigurability, error detection, and recovery into the Chameleon ARMOR architecture. Technical Report CRHC-98-13, University of Illinois at Urbana-Champaign, 1998.