SPACE SYSTEMS SYMPOSIUM (D1)
Space Systems Architectures (4)

Author: Mr. Edward Birrane
The John Hopkins University Applied Physics Laboratory, United States, Edward.Birrane@jhuapl.edu

Mr. Christopher Krupiarz
The John Hopkins University Applied Physics Laboratory, United States,
Christopher.Krupiarz@jhuapl.edu
Mr. Alan Mick
The John Hopkins University Applied Physics Laboratory, United States, Alan.Mick@jhuapl.edu

FLYING THE CLOUD:
THE CASE FOR ROUTING IN SPACE

Abstract

Cloud computing offers a new paradigm for robust, high-throughput data transfer as emergent behavior in highly connected networks. Clouds provide robustness by offering transparent redundancy to applications within the system – when an application service experiences an error a backup service can take over without requiring manual intervention. Data throughput is increased by enabling multiple paths in the network for concurrent data exchange and migrating services away from high-load and high-congestion areas. There are several significant architectural hurdles in recreating a cloud computing concept into space, including longer propagation delays, lower available transmission bandwidth, and low-powered processors. However, applying these concepts to flight systems will result in an overall increase in science data return and application fault tolerance, especially in the context of envisioned space mesh networks.

Spacecraft and planetary landers typically power multiple processors, including dedicated and redundant Command and Data Handling (CDH), Guidance and Control (GNC), and processors for science instruments, bus controllers, and other special-purpose hardware. Connections between processors are typically hard-wired and processors are typically special-purpose to conserve power. Mesh configurations consisting of orbiters, spacecraft, and landers exist today and, at the logical level, demonstrate dozens of processors in flight communicating over highly-specialized links. Flight architectures support memory-protected execution and multiple application executables rather than monolithic, single-binary-file code uploads.

This paper proposes an infrastructure to transform the model of dozens of processors running individual applications into a service-oriented, cloud architecture. The critical enabler of this capability is an application-level router resident on spacecraft and landers that not only routes between space assets, but within them as well. Such a router allows applications to be individually addressed both internal and external to the spacecraft. The focus of communications moves from platforms as network endpoints to applications as network endpoints, with very little additional hardware necessary over what is being flown today.

An analysis of the advantages of this approach for data throughput, streamlined mission design, and reduction in software implementation and testing costs is performed. Additionally, standards for software bus publish and subscribe mechanisms and supporting protocols are presented. Ultimately, this paper concludes that the core hardware resources for clustered computing exist in space systems today and that the introduction of inter-asset routing and software publish and subscribe mechanisms presents a compelling extension to this architecture.