# INSPiRE – An Approach to Mission Quality Management using Network Slicing for Space Applications

Apurva N. Mody, Junaid Islam, Bryan Crompton, Dap Tran, David Simpson, Dylan J. Gormley
AiRANACULUS
Lowell, MA, USA
apurva.mody@AiRANACULUS.com

Aaron Smith
NASA Glenn Research Center
Cleveland, OH,
aaron.smith@nasa.gov

Mieczyslaw M. Kokar, Jakub J. Moskal
VIStology, Inc.
Framingham, MA, USA
{mkokar, jmoskal}@vistology.com

Tommaso Melodia
Northeastern University
Boston, MA, USA

***Abstract*— Managing traffic between the Earth-Moon and Earth-Mars is a complex process requiring significant investment in resources and expertise at NASA. INSPiRE improves the performance of space networks by enabling a dynamic re-configuration process that works for any mixed topology over a heterogeneous and multi-vendor network. To achieve the desired functionality, INSPiRE incorporates a set of algorithms, machine learning processes, and policy inference to handle unpredictable, disruptive events. INSPiRE draws parallels from the current notion of the 3GPP (5G and beyond) Network Slicing approach, where the same physical network divides into several virtual networks, and for each of these virtual networks, there is a guaranteed Quality of Service for the missions that they serve.**

***Keywords*— *Cognitive Communications, Space Communications, Satellite Communications, Deep Space, Decision Engine***

## I. INTRODUCTION

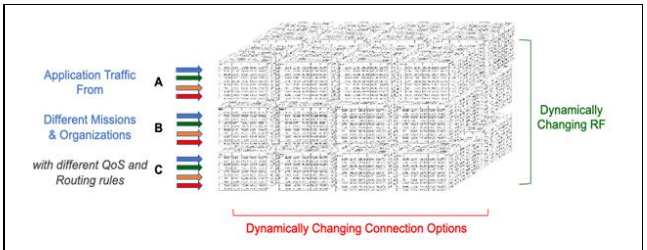


***Figure 1***. Current NASA network optimization matrix is complex and dynamically changing.

Routing mission-critical traffic across an ever-changing wireless network is a complex task that requires significant expertise and resources. Some variables include the location of fixed and mobile radios, their channel capacity, and understanding the relative priority of mission-critical traffic and application (e.g., voice, video, data). Unfortunately, due to variables such as RF noise and equipment outages, accurate application utilization is impossible to predict, thus making even the best plans a best-effort guess. Fig. 1 helps us understand the complexity of managing mission-critical wireless systems. We have different mission-critical applications from organizations traveling over a changing mobile network. This complexity creates an unsolvable problem, as traditional optimization algorithms would never be able to compute the solution fast enough. For example, a smart transportation system that utilizes multiple wireless networks (4G, 5G, WiFi6) to communicate with autonomous vehicles, traffic sensors, and mobile apps would never be able to provision and prioritize mission-critical traffic with any level of reliability; at best current mission-critical wireless networks are "best efforts."

## II. LUNANET INTRODUCTION

NASA's Artemis missions will incrementally work towards a permanent human presence on the Lunar South Pole, and over the next decade, a there will be a significant increase in the number of science instruments, vehicles, and devices on the Lunar surface. Additionally, NASA would like to provide enhanced communication services to these lunar assets, as the technology that we have grown accustomed to on Earth has improved since the Apollo era. To address this, NASA has proposed a Lunar network architecture called LunaNet [3]. This network will encompass the Lunar surface-to-surface links, proximity links from the Lunar surface to orbital relays, direct-to-Earth (DTE) links that connect the Lunar surface assets to ground stations on Earth, as well as the links that connect satellites to other satellites. Each of these communications links will operate in many frequency bands, ranging from UHF to optical, and the Lunar surface network will likely include both 4G/5G and Wi-Fi™ technologies. Hence, a network optimization technology needs to be protocol and waveform agnostic and be able to operate in a heterogeneous multi-vendor ecosystem. The LunaNet specification describes the following services:

**Networked Communications (COM)**: The networked communication services in LunaNet will enable users to transfer data to other nodes using addressable and routable data units. The primary COM services are real-time critical data transmission, data aggregation and transmission in a store-andforward mode, and messaging. LunaNet user applications will be networked-based, using either Delay/Disruption Tolerant Networking (DTN) Bundle Protocol (BP) or Internet Protocol (IP). The standardized messaging services are expected to be utilized by applications such as service acquisition, PNT, and alerts. As shown in [3] Table 4, the implementation details of this messaging service are still "To Be Determined." As such, a functional messaging implementation is utilized and described in this work to demonstrate the cognitive functionality that such services will enable.

| QCI | Bearer Type | Priority | Packet Delay | Packet Loss | Space Communications Domain Example |
|---|---|---|---|---|---|
| 1 | GBR | 2 | 100 ms | $10^{-2}$ | Crew conversational voice |
| 2 | | 4 | 150 ms | $10^{-3}$ | Crew conversational video (live streaming) |
| 3 | | 3 | 50 ms | $10^{-3}$ | Telerobotics |
| 4 | | 5 | 300 ms | $10^{-6}$ | Non-conversational video (buffered streaming); science data |
| 65 | GBR | 0.7 | 75 ms | $10^{-2}$ | Mission Critical user plane Push To Talk voice (e.g., MCPTT) |
| 66 | GBR | 2 | 100 ms | $10^{-2}$ | Non-Mission-Critical user plane Push To Talk voice |
| 75 | GBR | 2.5 | 50 ms | $10^{-2}$ | V2X messages |
| 5 | Non-GBR | 1 | 100 ms | $10^{-6}$ | IMS Signaling |
| 6 | | 6 | 300 ms | | Video (buffered streaming) TCP-based (e.g., science data, www, e-mail, chat, ftp, p2p file sharing, progressive video, etc.) |
| 7 | | 7 | 100 ms | $10^{-3}$ | Voice, Video (live streaming), Telerobotics |
| 8 | | 8 | 300 ms | $10^{-6}$ | Vehicle-to-surface data and video (buffered streaming) |
| 9 | | 9 | | | |
| 69 | Non-GBR | 0.5 | 60 ms | $10^{-6}$ | Mission Critical delay sensitive signaling (e.g., MC-PTT signaling) |
| 70 | Non-GBR | 5.5 | 200 ms | $10^{-6}$ | Mission Critical Data (e.g., example services are the same as QCI 6/8/9) |
| 79 | Non-GBR | 6.5 | 50 ms | $10^{-2}$ | V2X messages |
| 80 | Non-GBR | 6.8 | 10 ms | $10^{-6}$ | Low latency eMBB applications (TCP/UDP-based); Augmented Reality |
| 82 | GBR | 1.9 | 10 ms | $10^{-4}$ | Discrete Automation (small packets) |
| 83 | GBR | 2.2 | 10 ms | $10^{-4}$ | Discrete Automation (big packets) |
| 84 | GBR | 2.4 | 30 ms | $10^{-5}$ | Intelligent Transport Systems |
| 85 | GBR | 2.1 | 5 ms | $10^{-5}$ | Electricity Distribution - high voltage |

***Figure 2***. CCSDS Red, Green, and Blue books define the QoS that NASA's networks will require to support.

**Position, Navigation, and Timing (PNT):** LunaNet is likely to provide PNT services for users on the Moon as well as for the proximity links. The PNT services will enable the users to determine the position and velocity of an orbiting or the lunar surface-based asset using reference signals.

**Detection and Information (DET)**: LunaNet detection and information services provide alerts and other critical information to users. This is used to enhance situational awareness of the users which may include astronauts, rovers, and other assets on the lunar surface. DET service will also alert the users of potentially dangerous solar activity. These alerts may be enabled using smartphones that use Wi-Fi™ and 4G network that is planned to be deployed on the lunar surface. LunaNet detection and information services will also include a lunar search and rescue capability, or LunaSAR.

**Science Services (SCI)**: The SCI services will enable various researchers to conduct measurements and experiments. Some other uses of the SCI service include radio astronomy enabled by the radio telescope on the lunar surface.

**NASA Missions**: Fig. 2 shows a list of various mission requirements that NASA will need to support [10, 11, 12]. The mission requirements are defined based on the ***Quality-of-Service Class Identifier (QCI).*** Each QCI value represents a particular application that needs to be supported with a corresponding Priority (0.5 – highest, 10 – lowest), Packet Delay and Packet Loss. The table does not provide the Throughput that is required to be supported. ***Hence, Throughput, Routing and allocation of network resources is something that the system needs to derive***. As an example, QCI of 65 means that mission critical push to talk application needs to be supported with a Priority of 0.7.

## III. INSPiRE INTRODUCTION

As discussed in the previous section, ***NASA will need to support a wide variety of missions with various communications needs to support the Lunar operations whereas the resources may vary.*** The NASA enterprise network spans the Lunar surface, the Lunar proximity and the DTE links. This mission critical communications needs to be carried out while the space assets (e.g., Orbital Relays) are in motion and may or may not be available. Also, any of the links may face interference from local sources or from natural phenomena such as a Solar flare. Providing Quality of Service (QoS) in such a scenario is extremely challenging.

Of late, Network Slicing concept has gained significant traction for 3GPP 5G networks [Rel. 17, 14]. Network slicing enables the multiplexing of virtualized and independent logical

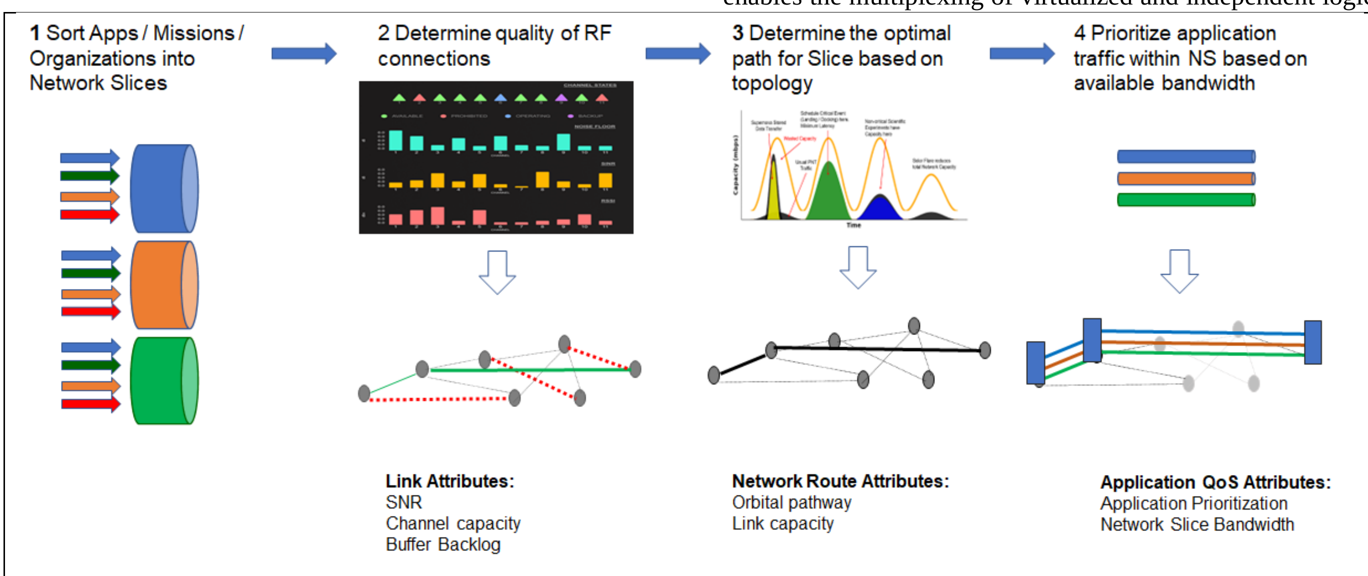


***Figure 3***. INSPiRE Concept of Operations (CONOPS)

networks on the same physical network infrastructure where each network slice is an isolated end-to-end network tailored to fulfil diverse requirements requested by a particular application. Hence, a network slice is a virtual logical network that also

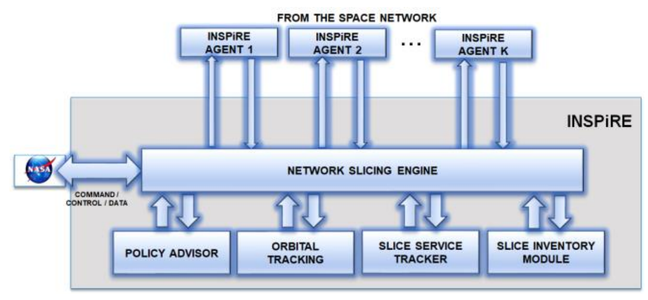


***Figure 4**.* INSPiRE architecture

defines the QoS. The 3gpp defines three types of network slices – enhanced Mobile Broadband (eMBB), Ultra-reliable Low Latency Communications (URLLC) and massive Machine-type Communications (mMTC). All applications (e.g., high-definition video, voice, Augmented Reality, Virtual Reality - AR/VR etc.) are mapped to one of these three categories. The three broad classes of network slices map the applications to the 5G Orthogonal Frequency Division Multiplexing (OFDM) waveform resources to fulfil the QoS requirements. While the network slicing concept is highly applicable to managing the needs of NASA's Lunar missions, simply applying the same 3gpp construct is not going to work since NASA's entire network will be heterogeneous with many wireless communications modalities and it will be dynamically changing over time and location.

***INSPiRE provides the network slicing construct for such a dynamically changing, heterogeneous multivendor network***. INSPiRE improves the performance of mission-critical networks by enabling a dynamic re-configuration process that works for any mixed topology over such a heterogeneous and multi-vendor network. To achieve the desired functionality, INSPiRE incorporates a set of algorithms, machine learning processes, and policy inference to handle unpredictable, disruptive events. INSPiRE draws parallels from the current notion of the 3gpp network slicing approach, where the same physical network divides into several virtual networks. For each of these virtual networks, there is a guaranteed Quality of Service for the missions that they serve.

Fig. 3 shows the CONOPS of the INSPiRE System. It consists of the following parts. 1. Prioritize and sort the Applications into slices, 2. Determine the quality of the RF connection, 3. Predict the network topology over the period of time that the network slice has to be orchestrated and pre-plan the routes that the packets should take, 4. Dynamically prioritize the applications traffic within the network slice based on the available resources while maintaining the minimum QoS that is desired for a particular application.

INSPiRE helps the mission manager not only in the planning and execution phase but also when the network needs to be re-configured on the fly due to catastrophic events such as interference, congestion, or malfunctioning node. The high-level architecture of INSPiRE has two primary inputs: 1. Real-time spectrum and channel data provided by a cross-layer signal processing system referred to as CLAIRE, and 2. Mission-critical application data that arrives at its input buffer from functions/applications such as life support systems, vehicle control, and astronaut communications. Based on the available channel capacity, INSPiRE arranges packets in its output buffer based on policy.

The INSPiRE architecture consist of two parts. The Mission Control Module and the Space Module. Fig. 4 shows the INSPiRE Mission Control Module, which consists of the Network Slicing Engine, the Policy Advisor, Orbital Tracking, Sice Service Tracker, and the Slice Inventory Module. Fig. 7 shows the INSPiRE Space Module. The Space Module is designated by the INSPiRE Agent, but it also consists of the Packet Inspection and Sorting, the Policy-based Packet Scheduler and the DTN Cache.

Based on the available channel capacity, INSPiRE arranges packets in its output buffer based on the QCI priorities, including never falling below the minimum data level for functionalities. In the extreme case when the available channel capacity falls below the minimum data level, which might happen during a catastrophic event like of natural disaster, INSPiRE can connect to the Policy Advisor module to help solve resource contention resolution by looking at higher-level concepts such as which traffic type is vital for human life.

The Policy Advisor implements a Policy-Based Control paradigm in which the network operator or mission manager provides a set of policies to analyze the state of the network and to provide guidance on what actions INSPiRE could, should, or should not implement in a specific situation.

The policies are written in formal, human, and computer-interpretable languages. The Web Ontology Language (OWL) defines a vocabulary for the policies. SPARQL used to specify policies, expressed as rules, as well as to query the high-level state of the network. A logical inference engine matches the preconditions of the policy rules against the state of the network and derives the actions. The operator can dynamically add ontology concepts and policies during the system run time.

## IV. INSPiRE OPERATION

Fig. 5 provides the INSPiRE Operations Flow. The flow consists of the following steps:

1. Every 24 hours, NASA Mission Manager provides a list of all the Applications (along with their QCIs) that the network needs to supports that day. The mission manager provides the source and destination IP addresses,
2. INSPiRE Network Slicing Engine asks the Orbital Tracker to map the routes that are available over the course of this time duration, including capacities and anticipated latencies,
3. INSPiRE Network Slicing Engine updates the Slice Inventory Module. It then queries the Slice Inventory Module if all these new Network Slices can be supported by the network,

4. If the answer is No, then the Policy Advisor is invoked to reprioritize the Network Slices based on Organization, Missions, and Applications,
5. INSPiRE Network Slicing Engine develops the Slice definitions and sends these to the Slice Service Tracker Module,

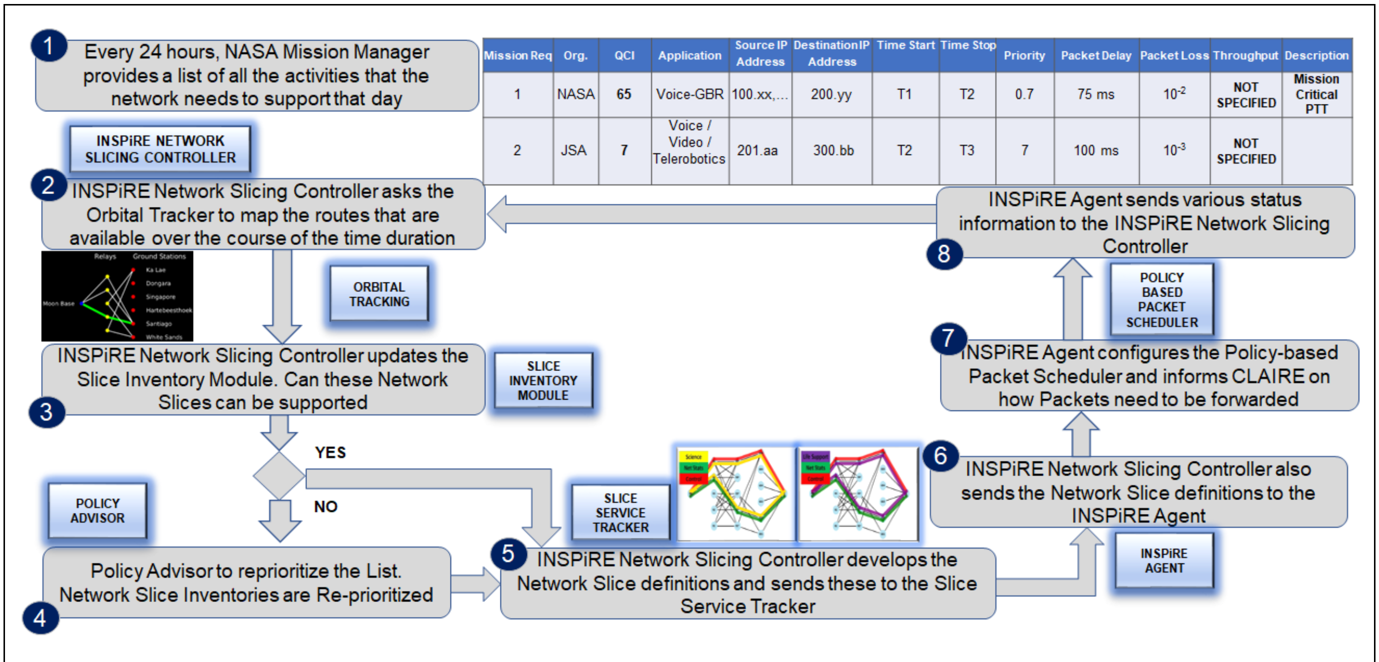


***Figure 5***. INSPiRE Operation Flow

6. INSPiRE Network Slicing Engine also sends the Network Slice definitions to the INSPiRE Agents,
7. INSPiRE Agents configure the Policy-based Packet Scheduler and inform CLAIRE of how the packets need to be forwarded.
8. INSPiRE Agents send various status information to the INSPiRE Network Slicing Engine so that it can track the performance and make configuration changes if needed.

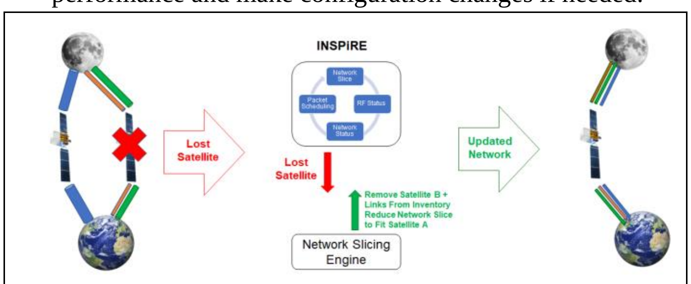


***Figure 6***. Network Slicing Engine operation when a satellite is lost

To illustrate the value of INSPiRE in managing a complex mission-critical network, we can examine how it would handle a shared space network as shown in Fig. 6. In this simple example, we have traffic from three organizations (NASA, ESA, JAXA) connecting over two satellite relays. Each organization has a mix of high-priority and low-priority traffic. INSPiRE sorts the traffic into three network slices (for simplicity). While there may be 100s of applications, they are treated as a set of network slices. Next, INSPiRE has the NASA slice going to one satellite and the partner agencies to the second. If one satellite fails, the traditional approach to rebuilding routing would not be efficient as it requires all applications to reconnect – losing critical data. Instead, INSPIRE will ask the Network Slicing Engine what to do. Network Slicing Engine instructs the INSPiRE Agents to 1. Remove the routing entry related to the lost satellite and 2. Re-route the Network Slice while maintaining the min application value (i.e., only two network slices instead of 100s of applications).

## V. Modeling & Simulations

We illustrate the operation of INSPiRE through some simple use cases. Based on the pre-defined QCI, INSPiRE ensures that the Policy-Based Packet Scheduler prioritizes the packets / bundles belonging to those services. INSPiRE adopts the same terminology as NASA's HDTN protocol stack [15] which consist of Ingress, Egress, and the (Policy-based Packet) Scheduler. There are multiple approaches to implementing the policy-based packet scheduler. In this case, the information required for such an implementation is packet (service) type, policies (service requirements), network connections, etc.

Fig. 7 shows the architecture of the INSPiRE Node in space and how priority & capacity aware network slice scheduling takes place. It is important to note that while for the 3gpp systems, the resource block allocation takes place based on the applications that are mapped to eMBB, URLLC and mMTC network slices, ***for a heterogeneous network which may or may not use 3gpp technologies, such a resource block allocation is not possible***. Hence, we have developed a generic architecture that performs an equivalent function at each of the LunaNet Nodes. Such an operation consists of Packet Type Inspection and Sorting, Policy-based Packet Scheduler connected to the INSPiRE Agent and the DTN Cache.

In the figure, the Packet Type Inspection and Sorting function sorts of various types of packets with different QCI designations into SCI, DET, COM, or PNT buffers. This allows us to then perform dynamic buffer management of the queues. In this case, the CLAIRE agent which interfaces to the Radios informs the INSPiRE Agent the total Capacity / Throughput that is available.

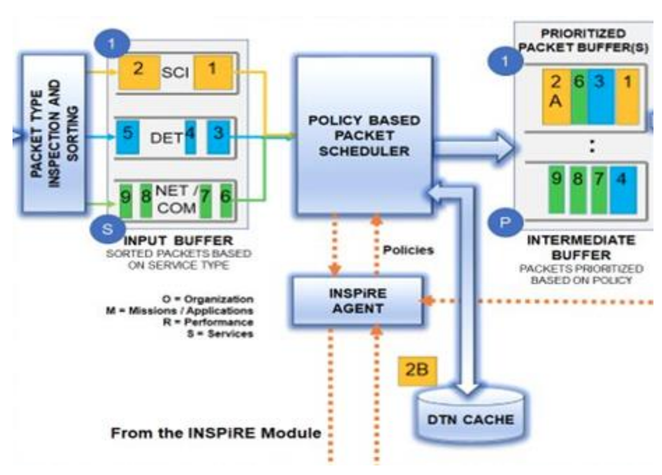


***Figure 7***. Architecture of the INSPiRE Node in Space - Priority & capacity aware network slice scheduling

CLAIRE also informs the INSPiRE Agent the Channel Quality Indicator (CQI) and differential buffer backlog that is present on a particular route. INSPiRE Agent also interfaces to the Network Slicing Engine and the Policy Advisor which determines which application needs to be prioritized. Based on all this information, the INSPiRE Agent sets the packet scheduling decision which is conveyed to the Policy-based Packet Scheduler. In this example, bundle 2 gets divided into two parts, where 2A is forwarded to the Prioritized Packet Buffers, whereas 2B is sent to the DTN Cache.

Fig. 8 is an example of [TOP] Policy-based Packet Scheduler during the normal operation and [BOTTOM] Interference / congestion event. In this case, traffic belonging to Slice A (Red) has Priority 1 (High), Slice B (Blue) has Priority 2, Slice C (Green) has Priority 3 and Slice 4 (Purple) has Priority 4 (Low). During the normal operation, the Capacity (Yellow) at the Egress is higher than the rate at which packets arrive at the Ingress. Hence the DTN Cache growth rates average to zero, and all Slices get the desired throughputs. Fig. 8 [BOTTOM] shows the INSPiRE operation during an interference event. Here the total capacity (Yellow) at the egress is impacted due to interference. The INSPiRE Agent learns about the capacity reduction from CLAIRE. Under this circumstance, INSPiRE maintains the throughput of high priority Slice A, while substantially reducing the throughput that is allocated to Slices B and C. Slice D has the lowest priority, and hence the corresponding packets are populated in the DTN Cache.

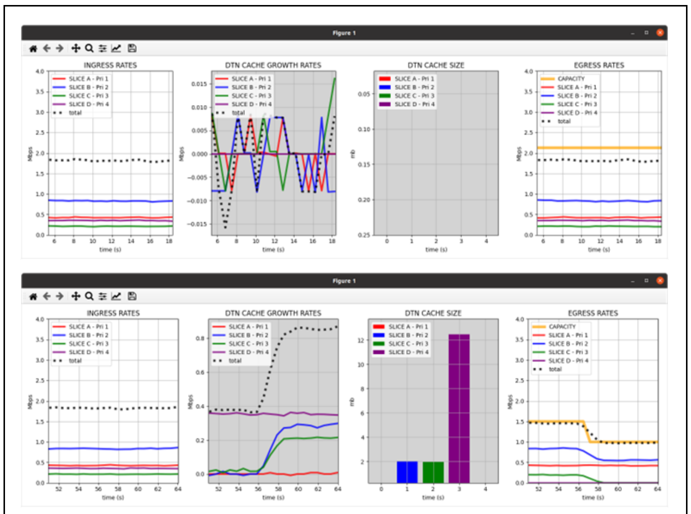


***Figure 8***. Policy-based Packet Scheduler during [TOP] Normal Operation. Capacity at the egress designated as Yellow [BOTTOM] Interference / Congestion event

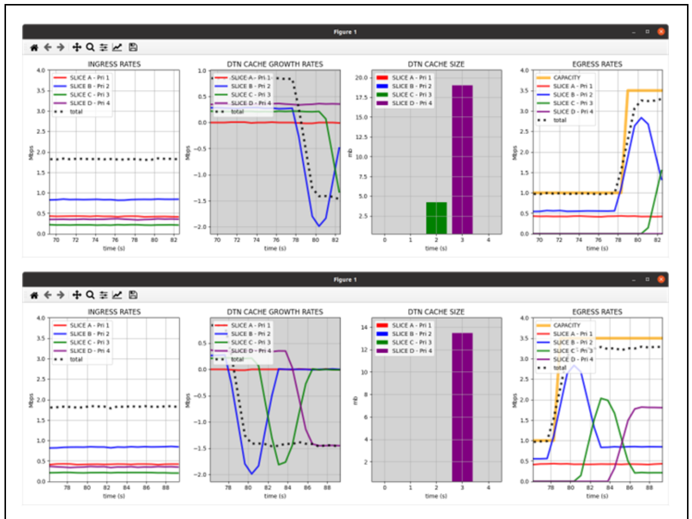


***Figure 9***. Policy-based Packet Scheduler during after an interference event where the DTN Cache packets are forwarded as soon as the capacity (Yellow) is restored.

Fig. 9 shows the operation of the Policy-based Packet Scheduler after the interference event where the capacity is restored. Once the capacity (Yellow) is restored, it maintains the throughput of the highest priority Slice A while allocating resources to Slice B and C such that the packets that are stored in the DTN Cache get transmitted.

There are several approaches to mapping the Radios to the Network Slices. The first approach is to use the highest-priority-gets-forwarded-first systems. In this approach, packets with the higher priority have the first choice and it selects the Radio with highest capacity and most robust connection.

Another approach is a variant of preemptive scheduling. Links can provide their available bandwidth as their state, and services can "subscribe" to links they wish to send their data through. Whenever a link has the available bandwidth, it accepts the service and updates its available bandwidth. Also, it is essential to note that these links should describe the bandwidth of the entire network and not only adjacent links if the packets are forwarded to a dead-end connection. The base station INSPiRE app should be able to handle these cases.

## VI. POLICY ADVISOR

Policy Advisor is invoked by the Network Slicing Engine by sending to it requests to derive specific decision to address issues that the network has experienced and is not able to achieve the desired QoS as given in the Mission Requirements provided by the Mission Manager. For instance, the NSC could request to "*deemphasize*" one of the mission requirements. The PA would need to infer which one to choose. Such a decision is derived by the PA based on the policies that are applicable to a specific request (in this case *deemphasize*). The NSC needs to send the PA the current set of mission requirements for the applications, e.g., Mission-Critical Voice, Voice/Video/Telerobotics, V2X-Messages, Telerobotics, Crew Conversational Video, Electricity Distribution, Mission Critical Delay Sensitive Signaling. In the simplest case, the policy may be "*deemphasize the application with the lowest priority*" as given by their QCIs. In this specific case, it would return the reply "*deemphasize Voice/Video/Telerobotics*" application. In case there are more than one such applications executed on behalf of different organizations; those would have to be analyzed with respect to other organization-specific policies. For instance, one of the policies might assign higher preference level to NASA vs. other countries applications, or just the opposite. The important aspect of this process is that policies can be modified by the Mission Manager on the fly, and this would not require any modifications of the code that controls the communications.

## CONCLUSIONS

In this paper, we presented the architectural construct and capabilities of the INSPiRE System, which provides an ability to manage wide variety of missions and applications over a dynamically changing heterogeneous multi-vendor network. We leveraged the 3gpp concept of Network Slicing to address this problem. While the 3gpp defined Network Slicing concept is rigid, we provided an architectural construct on how this may be implemented for LunaNet, as it applies to Lunar Surface Links, Lunar Proximity Links and Direct to Earth Links. We provided examples of the INSPiRE operation. The work on this project is still on-going and we are expecting to illustrate more results during future publications.

## ACKNOWLEDGMENT

We thank the Cognitive Communication Project team members Dr. Rachel Dudukovich, Dr. Janette Briones, Joseph Downey, and Dr. Adam Gannon for their valuable input the design of INSPiRE.

## REFERENCES

[1] Space Communications and Navigation (SCaN) Network Architecture Definition Document (ADD)

[2] Lunar Communications Relay and Navigation Systems (LCRNS), Preliminary Lunar Relay Services Requirements Document," NASA Technical Document, May 17th, 2022.

[3] Draft LunaNet Interoperability Specification," NASA Technical Document, March 2022.

[4] Pavel Puricer, Pavel Kovar, and Miroslav Barta, "Modernized Solar Radio Spectrograph in the L Band Based on Software Defined Radio" www.mdpi.com/journal/electronics, 2019

[5] Report of the Interagency Operations Advisory Group Lunar Communications Architecture Working Group The Future Lunar Communications Architecture September 25th, 2019

[6] David Israel, "LunaNet: A Flexible and Extensible Lunar Exploration Communication and Navigation Infrastructure," March 9th, 2020

[7] Report of the Interagency Operations Advisory Group - Lunar Communications Architecture Working Group The Future Lunar Communications Architecture September 25th, 2019

[8] CCSDS 734.2-B-1 CCSDS Bundle Protocol Specification. Blue Book. Issue 1. September 2015.

[9] CCSDS 734.1-B-1 Licklider Transmission Protocol (LTP) for CCSDS. Blue Book. Issue 1. May 2015.

[10] CCSDS Green Book, Report Concerning Space Data System Standards "WIRELESS NETWORK COMMUNICATIONS OVERVIEW FOR SPACE MISSION OPERATIONS," May 2017

[11] CCSDS Blue Book, Report Concerning "SPACECRAFT ONBOARD INTERFACE SERVICES—HIGH DATA RATE 3GPP AND WI-FI LOCAL AREA COMMUNICATIONS," February 2022

[12] CCSDS Red Book, Draft Recommendation for Space Data Systems Standard "SPACECRAFT ONBOARD INTERFACE SERVICES—HIGH DATA RATE WIRELESS PROXIMITY NETWORK COMMUNICATIONS," March 2021

[13] LunaNet Concept of Operations and Architecture, September 2020

[14] 3gpp Network Slicing construct - Link

[15] NASA High Rate Delay Tolerant Networking - Link

[16] A. N. Mody et al., "Machine Learning Based Cognitive Communications in White as well as the Gray Space," Proceedings of the IEEE MILCOM, 2007.

[17] A. N. Mody et al. "Recent Advances in Cognitive Communications," IEEE Communications Magazine, Special Issue on Network Centric Military Communication, Oct. 2007

[18] A. N. Mody, M. Sherman, R. Martinez, R. Reddy, and T. Kiernan, "A Survey of IEEE Standards Supporting Cognitive Radio and Dynamic Spectrum Access," Proceedings of MILCOM, Nov 2008.

[19] A.N. Mody, M. Sherman, "On Making the Current Military Radios Cognitive without Hardware or Firmware Modifications," Proceedings of MILCOM 2010.