The Distributed Spacecraft Autonomy (DSA)

Caleb Adams, Project Manager | Jeremy Frank, Deputy Project Manager

Talk Roadmap

- Distributed Space Missions
- The Distributed Spacecraft Autonomy Project
- DSA-Starling
- DSA-LPNT
- Lessons Learned
- Bibliography

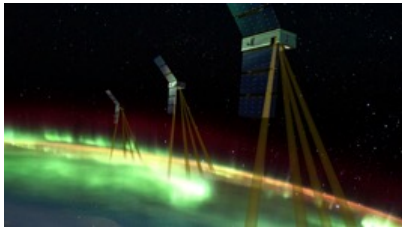

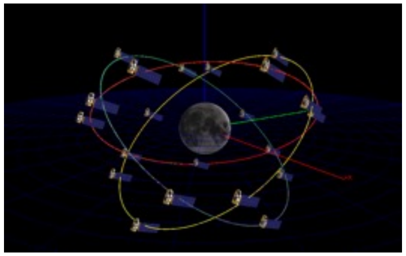

Talk Roadmap

- Distributed Space Missions
- The Distributed Spacecraft Autonomy Project
- DSA-Starling
- DSA-LPNT
- Lessons Learned
- Bibliography

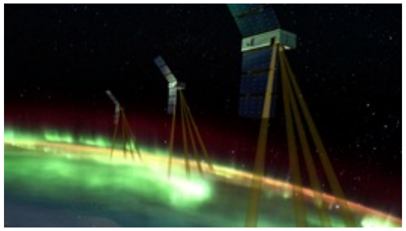

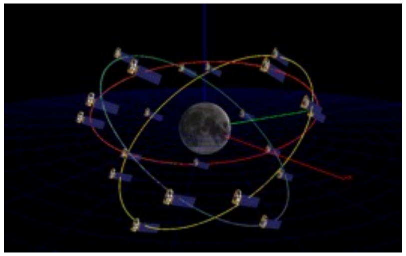

Distributed Space Missions

- **What are Distributed Space Missions?**
  - Distributed Space Missions (DSMs) are systems that utilize multiple spacecraft to accomplish mission objectives.
  - DSMs can involve heterogeneous spacecraft, consisting of different types, or homogeneous spacecraft, all the same type.
- **Why Distributed Space Missions?**
  - These missions have become more prevalent due to the decreased launch costs and development costs per vehicle.
  - Multiple spacecraft can increase spatiotemporal coverage of space missions. This is invaluable for many science needs, both on Earth and in deep space.
  - DSMs with space-to-space communications enable decreased data delivery latency (when an event is detected to delivery to ground). This is invaluable for disaster response, space weather, and other applications.

Distributed Space Missions

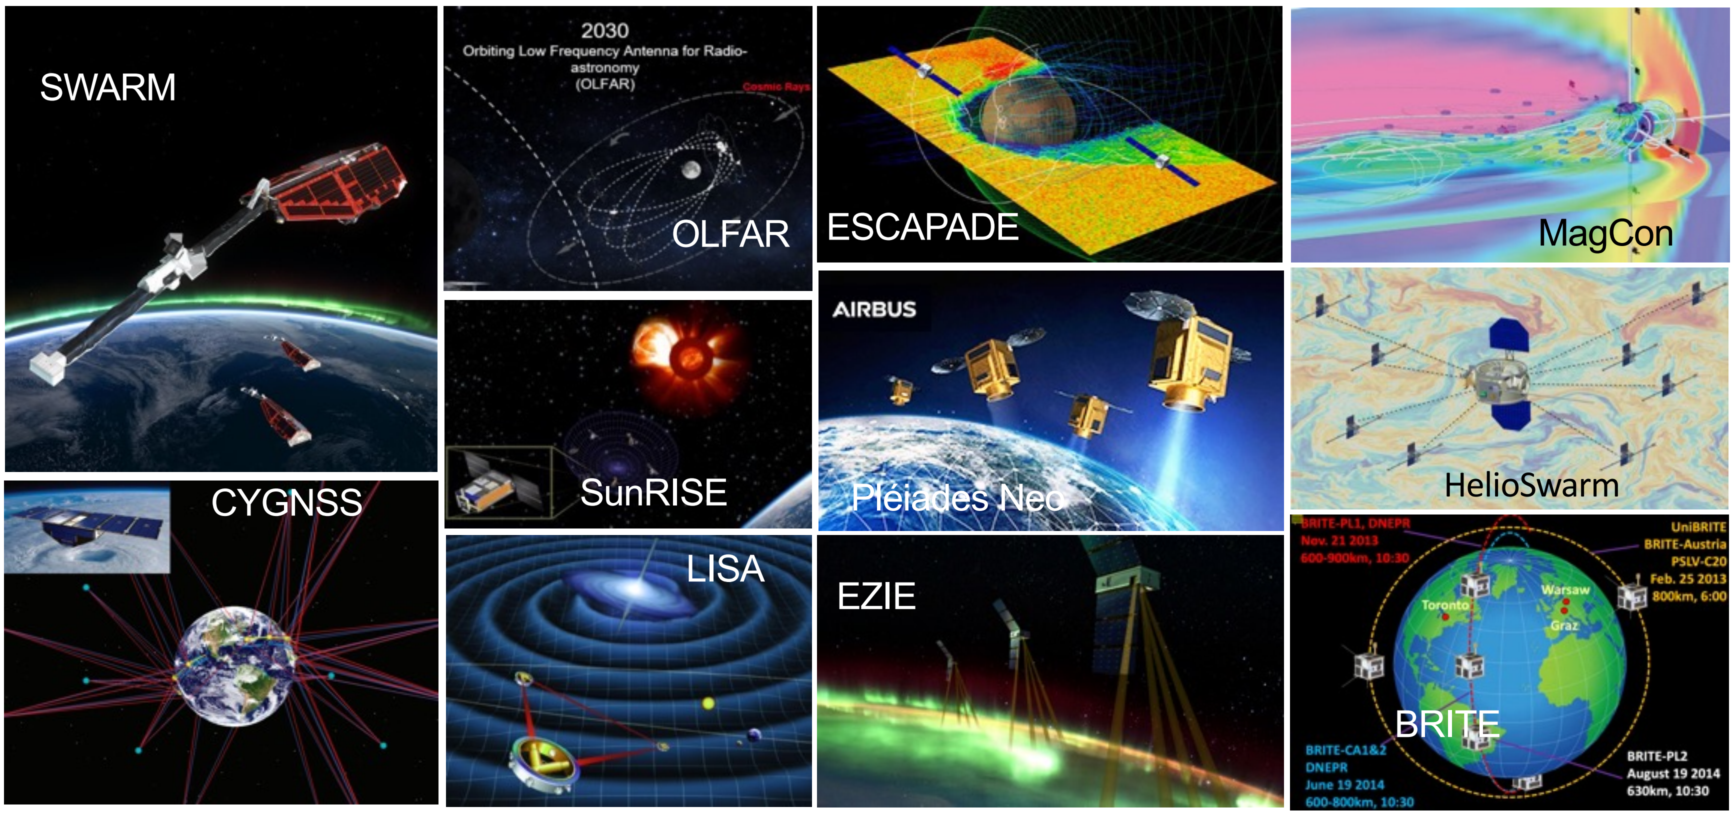

Distributed Space Missions

- **Heliophysics**

	# S/C
The Sun Radio Interferometer Space Experiment (SunRISE) is an array of six CubeSats that will work together to study solar activity.	6
Electrojet Zeeman Imaging Explorer (EZIE) is a six-spacecraft mission to explore the Sun and the system that drives space weather near Earth.	6
Helioswarm is a constellation or “swarm” of nine spacecraft that will capture the first multiscale in-space measurements of fluctuations in the magnetic field and motions of the solar wind known as solar wind turbulence.	9

Notional
Funded
Flight

Distributed Space Missions

- **Astronomy** — **# S/C**
  - The Bright Target Explorer (BRITE) Constellation provides milli-magnitude (0.1% error) differential photometry of bright stars using a six satellite constellation. — 6
  - The Laser Interferometer Antenna (LISA) mission uses three spacecraft to determine how and when the massive black holes have formed and grown over cosmic time. — 3
  - The Orbiting Low Frequency Antenna for Radio Astronomy (OLFAR) mission concept uses distributed sensors in space to explore frequency bands for radio astronomy. — 25-50
- **Planetary Science**
  - The Escape and Plasma Acceleration and Dynamics Explorers (ESCAPADE) mission is a twin-spacecraft science mission that will orbit Mars to investigate the structure, composition, variability, and dynamics of Mars' unique hybrid magnetosphere. — 2

Notional
Funded
Flight

Distributed Space Missions

- **Earth Science**

 # S/C

- Swarm will create a detailed survey of Earth's geomagnetic field and its temporal evolution as well as the electric field in the atmosphere, by using a three-satellite constellation that carries sophisticated magnetometers and other instruments. — 3
- Cyclone Global Navigation Satellite System Mission (CYGNSS) is a NASA Earth System Science Mission to collect the first frequent space-based measurements of surface wind speeds in the inner core of tropical cyclones using six spacecraft. — 6
- Pléiades Neo satellites comprises four identical 30-cm-resolution spacecraft when fully deployed, the constellation will offer reactive tasking, along with the ability to visit any point on the globe at least twice daily. — 4
- The Magnetospheric Constellation (MagCon) mission is designed to understand the transport of mass and energy across the boundaries of, and within, Earth's magnetosphere using a constellation of up to 36 small satellites — 36

Notional
Funded
Flight

Distributed Space Missions

- **Other Distributed Space Mission drivers**
  - Position Navigation and Timing
  - Commercial Imaging
  - Surveillance
  - In-Space Inspection, Assembly and Servicing
  - Planetary surface: CADRE
  - Planetary Surface: Lunar/Mars Geophysics

Distributed Space Missions

- **Technical Challenges of Distributed Space Missions**
  - Knowledge – How do we know the positions and movements of the spacecraft in the DSM?
  - Communications – How do we get information to and from, and between, the spacecraft in the DSM?
  - Control – How do we maintain the configuration of spacecraft in the DSM?
  - Operations – How do we command the DSM configuration and return data from it?
  - Access – How do we get the DSM into space and deploy it?
- **We know how to do all of these with small numbers of spacecraft and traditional operations, but not how to do it cost effectively for large numbers of spacecraft.**
  - (Some counterexamples: PlanetLabs, surveillance assets...)

- **Technical Challenges of Distributed Space Missions**
  - Multi-spacecraft coordination.
    - DSMs can improve spatiotemporal coverage without coordination overhead (e.g. CYGNSS).
    - However, coordinating activities of multiple spacecraft (e.g. stereo images, heterogeneous data collection from multiple instruments, etc.), even from ground, is more complex.
  - Operations complexity.
    - We want one operations team to manage many spacecraft. Unfortunately, each spacecraft may be complex enough that each requires a significant operations team in its own right.
    - Coordinated operations (see above) significantly increases operational complexity of the system as a whole, requiring more staff.

Distributed Space Missions

- **Technical Challenges of Distributed Space Missions**
  - Space to space communications.
    - The technology is still in its infancy.
    - Space-to-space radio communications options include omni-directional and directional antennas.
    - Space-to-space Laser-communications. (Look up that JPL cat video...)
  - Flight dynamics and formation flying. Ensuring precise spacecraft formations is critical for some complex science needs (e.g. Helioswarm).
    - This is one of several forms of coordinated operations that increases operational complexity.
    - At best, failure to maintain formation jeopardizes science, or space-to-space communications; at worst, it risks spacecraft damaging each other due to thruster plumes, or collisions.

(Autonomous) Distributed Space Missions

- **Autonomous Distributed Space Missions**
  - The 2018 Workshop on Autonomy for Future NASA Science Missions explored how autonomy can impact many upcoming NASA missions.
    - Workshop discussed Autonomous Distributed Space Missions for Astrophysics, Earth Science, Heliophysics, Moon, Mars, and Venus.
  - NASA's Earth Science Technology Office (ESTO) is arguably at the forefront of thinking about and advocating for autonomous distributed space missions.
  - ESTO defines an Intelligent and Collaborative Constellation as follows:
    - Uses onboard intelligence to perceive its environment, takes actions that maximizes its chances of success in creating coordinated observations; it can potentially learn from its experiences.
    - May include Real-time data understanding, Situational awareness, Problem solving, Planning and learning from experience, Communications and cooperation between multiple spacecraft.
- **Note: little to no autonomy in any of the previously described DSMs.**

(Autonomous) Distributed Space Missions

- **Advantages of Distributed Autonomous Space Missions**
  - Fault tolerance and redundancy: The system becomes more resilient to individual failures.
  - Increased computational capacity: Distributing decision-making among multiple spacecraft allows for parallel processing and sharing computational load.
  - Enhanced adaptability and flexibility: Enables spacecraft to make autonomous and localized decisions based on local perception and information.
  - Efficient task allocation and coordination: Spacecraft can autonomously allocate tasks among themselves based on their capabilities, proximity, and availability.
  - Increased robustness to communication delays and failures: Spacecraft can continue operating even in scenarios where communication between nodes is intermittent or disrupted.
  - Improved scalability: Distributed decision-making enables the system to scale up or down easily.
    - Perhaps not for space-based...

(Autonomous) Distributed Space Missions

- **Technical Challenges of (Autonomous) Distributed Space Missions**
  - Knowledge – How does each spacecraft know its own positions and movements, and those of the spacecraft in the DSM?
  - Communications – How do spacecraft get information to and from ground and each other?
  - Control – How do the spacecraft maintain the configuration the DSM?
  - Operations – How do we inform the DSM of the desired configuration? How does the swarm achieve the desired configuration on its own?
  - Access – How do we get the DSM into space and deploy it?
- **We don't know yet how do to do most of these tasks autonomously, with DSMs of any size.**
- **Enter DSA!**

Talk Roadmap

- Distributed Space Missions
- **The Distributed Spacecraft Autonomy Project**
- DSA-Starling
- DSA-LPNT
- Lessons Learned
- Bibliography

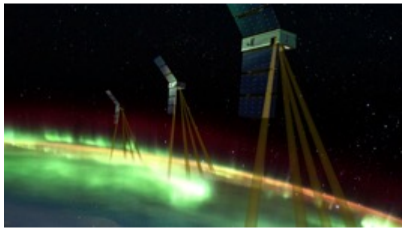

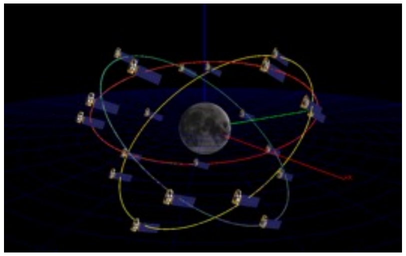

Distributed Spacecraft Autonomy

- **NASA's Distributed Spacecraft Autonomy (DSA) Project focuses on the following technical areas:**
  - Enhanced adaptability and flexibility: Distributed decision-making enables spacecraft to make autonomous and localized decisions based on local perception and information.
  - Efficient task allocation and coordination: By distributing decision-making, spacecraft can autonomously allocate tasks among themselves based on their capabilities, proximity, and availability.
  - Ad hoc Network Communications: communication infrastructure that is scalable, robust, and automatically self-configuring.
  - Human-Swarm Interaction: ground control software that enables the ability to command and interact with a distributed space mission as a collective, rather than managing individual spacecraft.

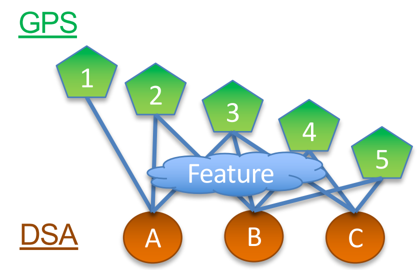

$$\max \sum_{s,g} \alpha\, r^{e}_{s,g} o_{sg} + (1-\alpha)\, r^{x}_{s,g}\, o_{sg} + \sum_{s,g} \alpha\, g^{e}\, c_g + (1-\alpha)\, g^{x}\, c_g$$

$$s.t. \forall s \sum_{s} o_{sg} < s_c$$

$$\forall g \sum_{g} o_{sg} > n$$

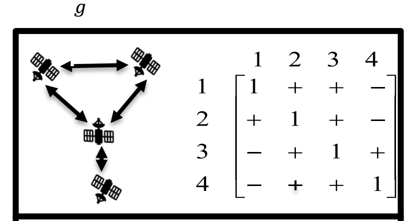

Example Swarm Topology

Distributed Spacecraft Autonomy

- **NASA's Distributed Spacecraft Autonomy (DSA) Project**
  - DSA is organized into two major sub-projects.
  - DSA-Starling is a software payload on the Starling space mission.
  - DSA-LPNT is a processor-in-the-loop scalability study building on DSA-Starling.

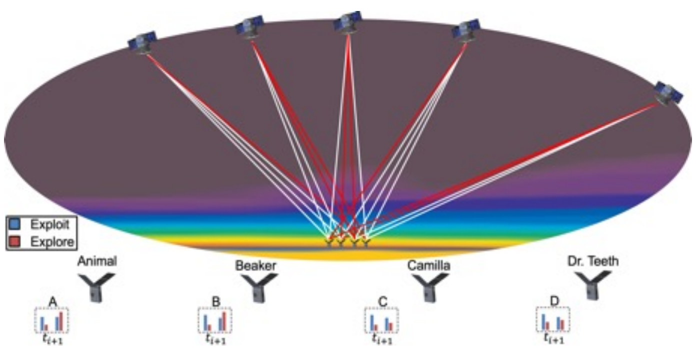

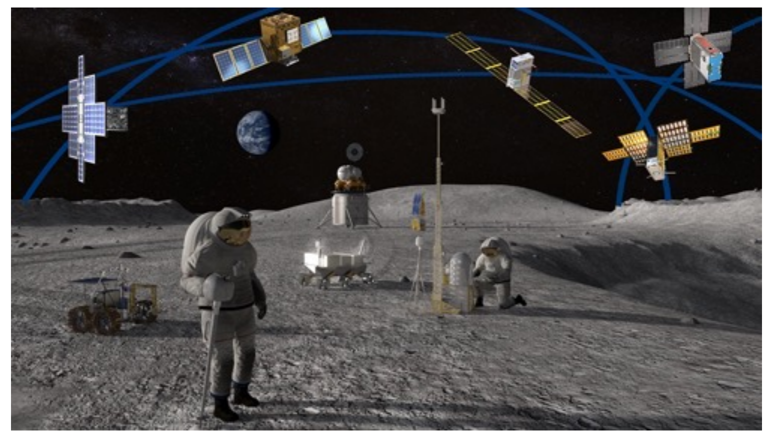

Distributed Spacecraft Autonomy

- **DSA –Starling Overview**
  - Starling consists of 4 Cubesats in Low Earth Orbit (LEO). Spacecraft communicate with each other via radios in a broadcast network. DSA-Starling is a software payload on the Starling space mission.
  - As a science proxy, we will use well-known upper-atmospheric anomalies, which are detected and reported via the GPS receivers aboard each spacecraft.
  - As these signals are detected, spacecraft need to decide which GPS channels to monitor, and which to ignore, in order to characterize the anomalies.
  - DSA-Starling will demonstrate how an Autonomous Distributed Space Mission solves this problem.

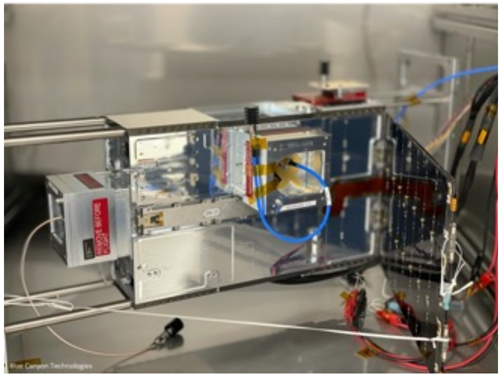

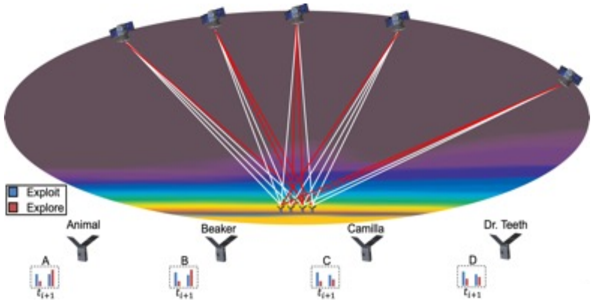

Distributed Spacecraft Autonomy

- **DSA –LPNT Overview**
  - In the upcoming decades there is expected to be a substantial increase in Lunar missions supporting and inspired by NASA's Artemis Program. These missions require Position, Navigation and Timing (PNT) services.
  - Creating a dedicated constellation (where the PNT antenna would always be on) is too costly. An alternative is to use small-sat Lunar assets to create an ad-hoc, non-dedicated PNT constellation.
  - We focus in this talk on the problem of distributed position estimation of the LPNT spacecraft.
  - LPNT provides a novel problem to demonstrate scalability of the DSA technology from 4 spacecraft to dozens.

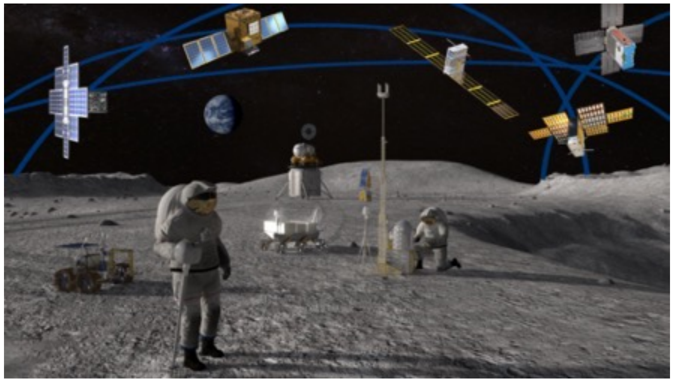

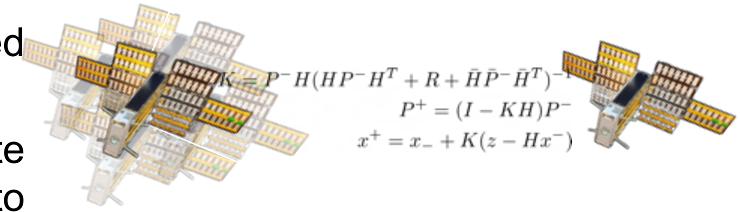

Distributed Spacecraft Autonomy

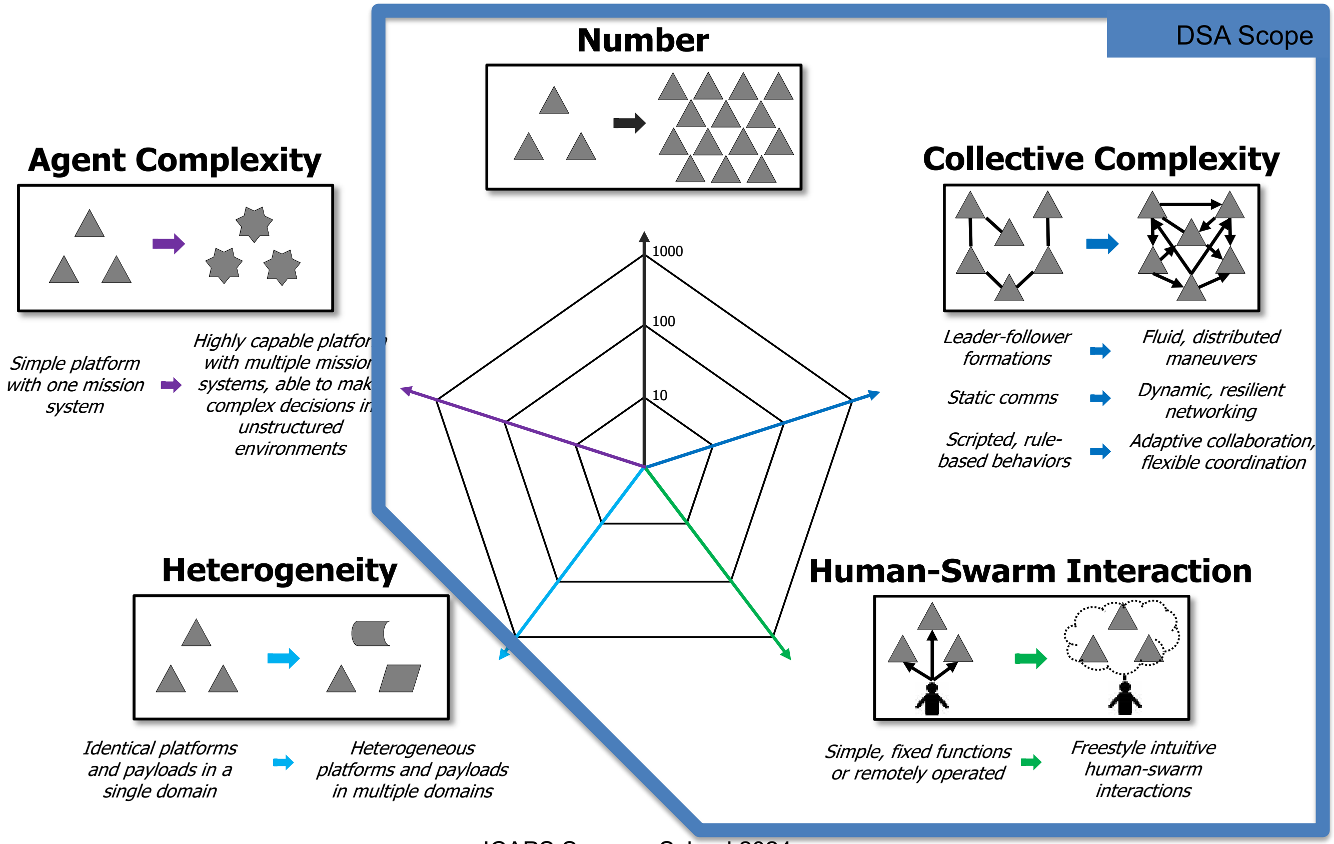

Talk Roadmap

- Distributed Space Missions
- The Distributed Spacecraft Autonomy Project
- DSA-Starling
- DSA-LPNT
- Lessons Learned
- Bibliography

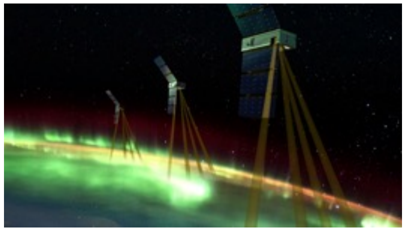

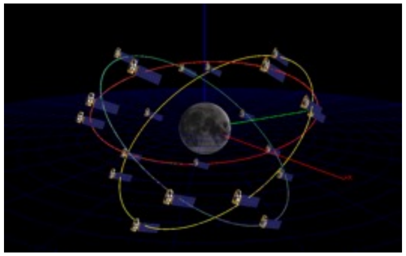

Talk Roadmap

- Distributed Space Missions
- The Distributed Spacecraft Autonomy Project
- DSA-Starling
  - The Starling Mission
- DSA-LPNT
- Lessons Learned
- Bibliography

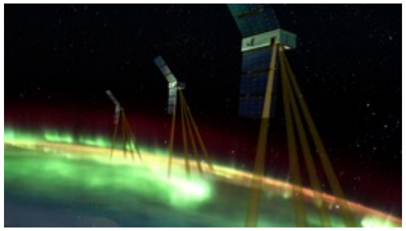

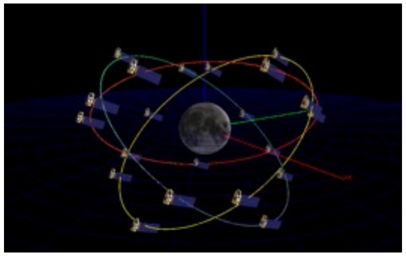

DSA-Starling

DSA-Starling

- **The Starling Mission**
  - Starling consists of 4 Cubesats in a Sun-synchronous orbit more than 480 km above Earth and 270 km apart.
    - Sun-synchronous orbits are nearly polar orbits that allow a spacecraft to consistently see the same amount of sunlight each orbit and generate the same amount of power with its solar panels.
  - Spacecraft communicate with each other via radios in a broadcast network.
    - Mobile Ad-hoc Network (MANET): Starling spacecraft communicate with each other via two-way S-band crosslink radios/antennas. If one spacecraft communications node fails, the communications route automatically reconfigures to maintain full communication capabilities for the remaining spacecraft.

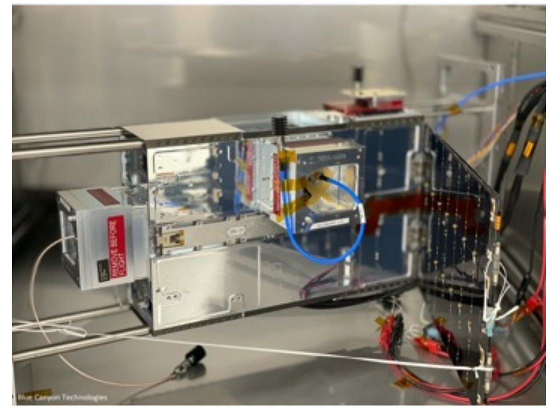

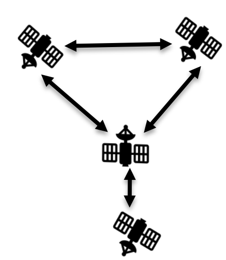

	1	2	3	4
1	1	+	+	−
2	+	1	+	−
3	−	+	1	+
4	−	+	+	1

Example Swarm Topology

DSA-Starling

DSA-Starling

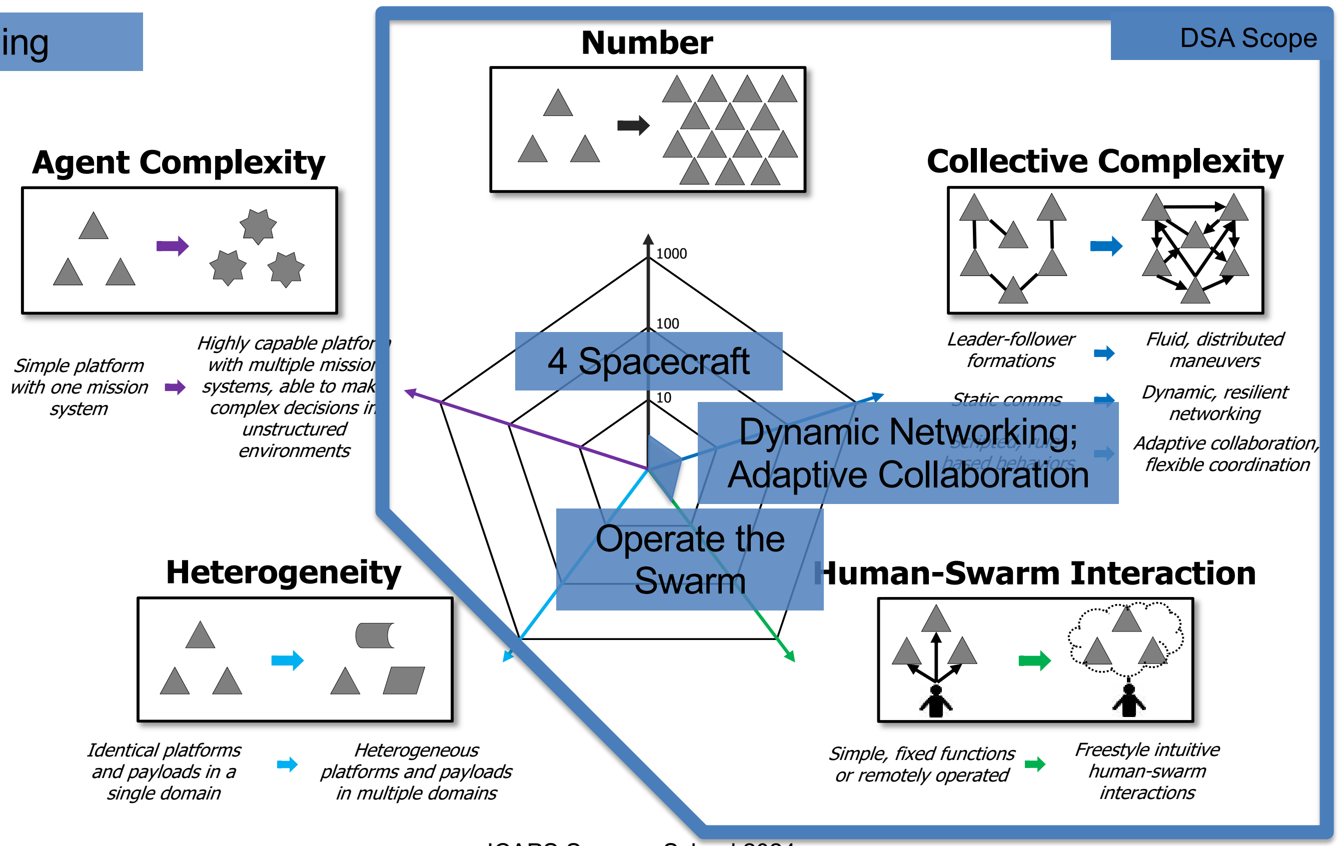

Talk Roadmap

- Distributed Space Missions
- The Distributed Spacecraft Autonomy Project
- DSA-Starling
  - DSA-Starling Science Proxy
- DSA-LPNT
- Lessons Learned
- Bibliography

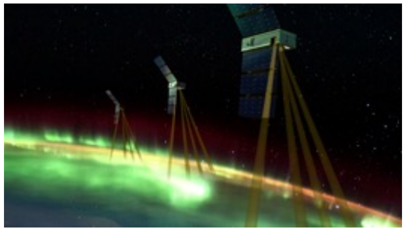

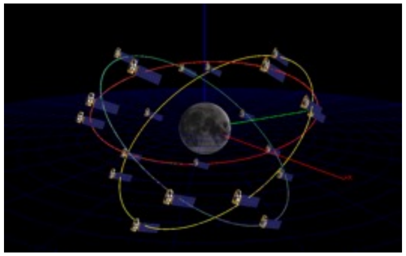

DSA-Starling

- **DSA-Starling Science Proxy**
  - Previous missions (e.g. SWARM, CHAMP) have established the upper atmospheric anomaly science that will drive DSA-Starling-1's autonomy algorithms.
  - Starling's spacecraft will collect Total Electron Content (TEC) information using an 8-channel dual-frequency GPS receiver also used for precise orbit determination.
  - SWARM data products provide testing and training sets for DSA-Starling-1:
    - Same Data rate (1 Hz), Similar altitude (450 km vs. ~500 km), Similar formation (pushbroom)

DSA-Starling is an autonomy demonstration with a science **use case**

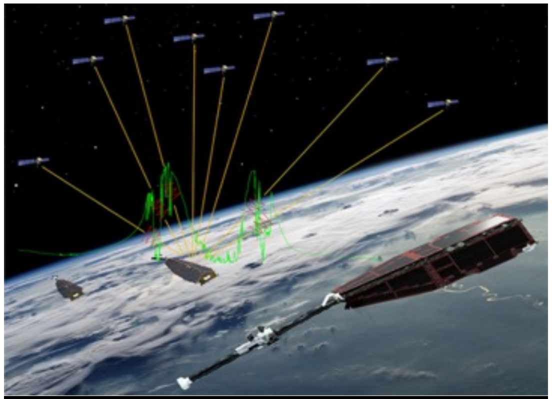

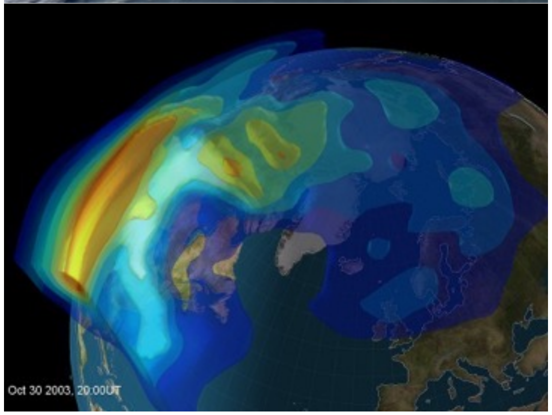

DSA-Starling

- **DSA-Starling Science Proxy**
  - GPS satellites are geostationary satellites orbiting at 20,000 km.
  - The upper atmospheric anomalies occur between the GPS satellite orbital altitude and the DSA-Starling orbital altitude.
  - Each satellite's GPS receivers will therefore detect the anomaly when Starling spacecraft pass through a region where the anomaly occurs.
  - For our purposes, the 'channel' refers to the signal a Starling spacecraft receives from a single GPS satellite.
  - As we see on the next slides, the anomalies manifest different ways relative to DSA-Starling's orbit, leading to different autonomy objectives.

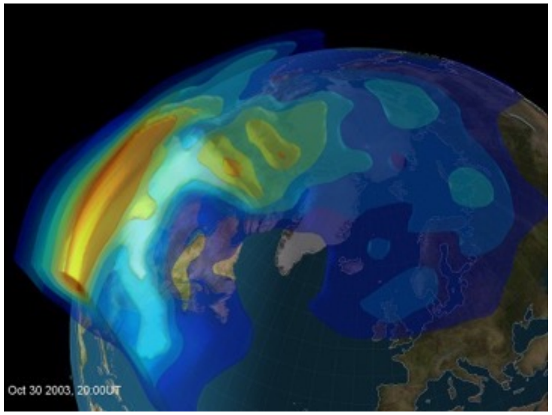

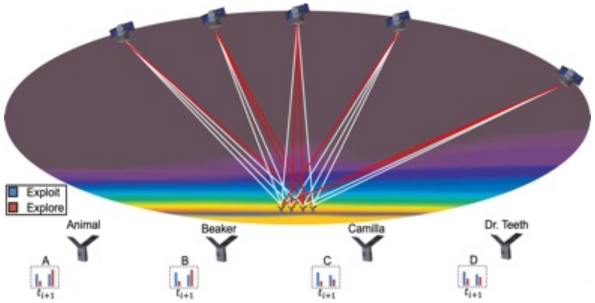

DSA-Starling

- **DSA-Starling Science Proxy**
  - The Equatorial Ionization Anomaly is a feature in the ionospheric F region (150-800km altitude). Spacecraft passing the magnetic equator will see an increase in Total Electron Count (TEC) from their GPS receivers. The Equatorial Ionization Anomaly has a few notable characteristics:
    - Relatively large: most GPS channels will be saturated, and spacecraft will see approximately the same feature.
    - Spatially Constrained: spacecraft can expect to see the feature within +/- 15^{o} latitude.
    - Reliable: it is generally present regardless of solar activity or time of year, and spans more than 12 hours of local time (meaning it is likely to appear regardless of the local time of the ascending node (LTAN) of the spacecraft).

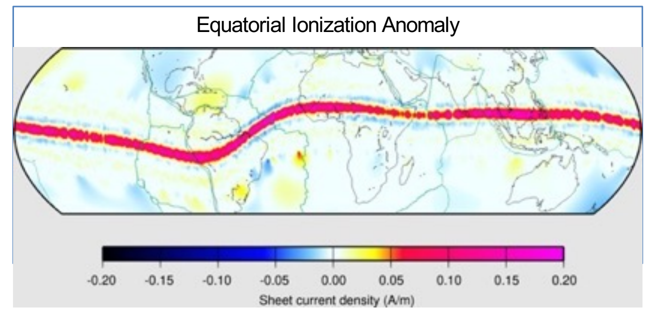

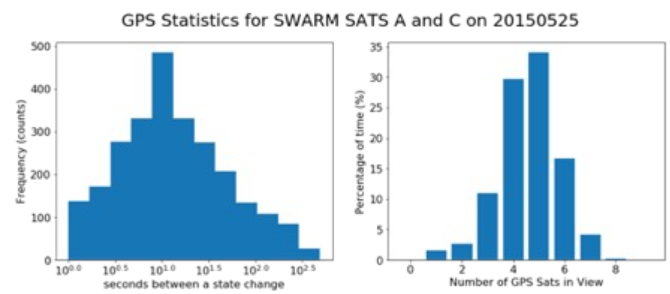

DSA-Starling

- **DSA-Starling Science Proxy**
  - Polar patches are "islands" of high plasma density in a low-density environment in the ionospheric F region (150-800km altitude). Their extent ranges from 100 to 1000 km in horizontal size. They occur nowhere with a geomagnetic latitude with magnitude less than 55° and are most likely above magnitude 80° latitude.
  - Patches occur in an annual cycle with maxima at the equinoxes and December solstice and a June solstice minimum which reflects the global ionospheric seasonal asymmetry in electron density. This indicates that the ideal launch window for DSA to perform its demonstration lays between day 285 of one year and day 70 of the next year.

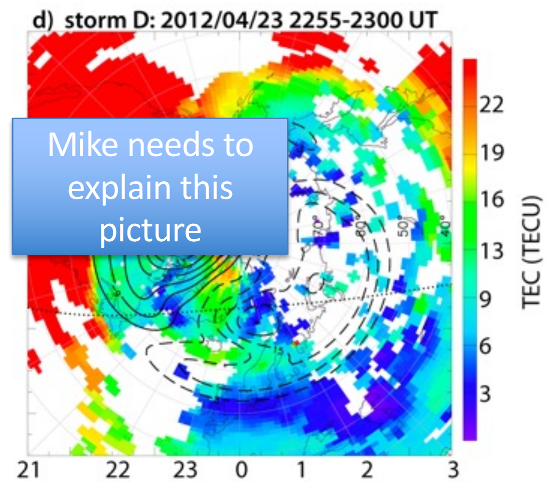

Mike needs to explain this picture

Occurrence of polar patch phenomena over time for the GPS instrument of the SWARM satellites.

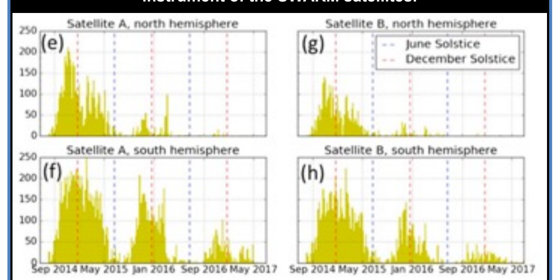

DSA-Starling

- **DSA-Starling Science Proxy**
  - Typically, 10-20 GPS satellites are in view of each Starling spacecraft.
  - The GPS receivers can typically receive data from all in-view GPS satellites.
  - For our experiments, we artificially limit the number of GPS channels any Starling spacecraft can receive to 4. This limitation drives choices in which channels to receive in the presence of the equatorial anomaly or the polar patch phenomena.

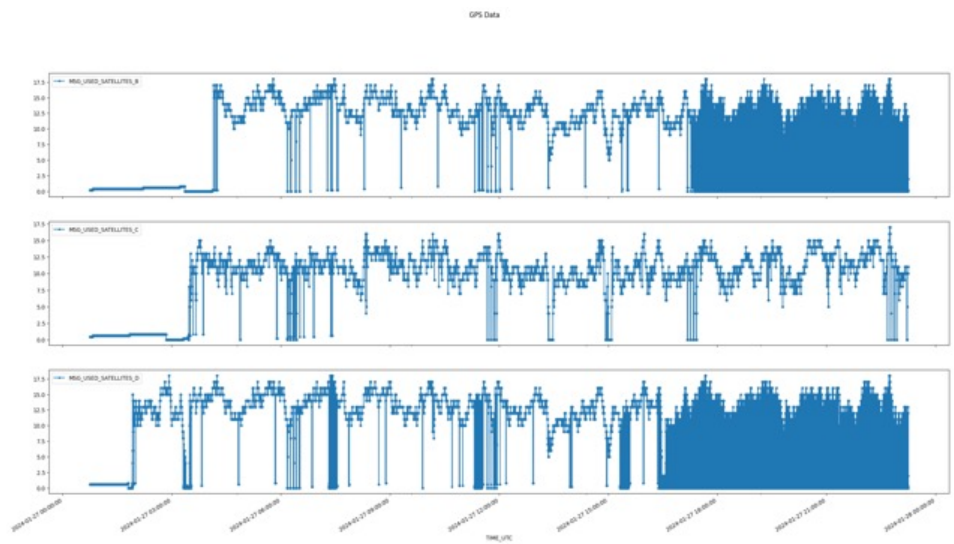

DSA-Starling

- **DSA-Starling Science Proxy**
  - The behavior of the high TEC signal requires distinct GPS channel allocation strategies.
  - During equatorial passes, we want to ensure that the maximum number of GPS channels are covered by DSA a whole, with each asset listening to as few channels as possible. Efficiently distributing GPS channel allocations across DSA assets requires two steps: The assets need to communicate to each other the value of monitoring each GPS satellites that are currently in their field of view. The assets then need to reach consensus on which asset is listening to which GPS satellite. We '*Explore' TEC over the Equator.*
  - Use Polar Patches to characterize the ability of DSA to react to transient, localized phenomena. Some GPS TEC rays will pass through the patch and some will not. The satellite will not know if it is passing through a patch until it is well inside the patch. A satellite that determines that it is passing through a patch will communicate which GPS satellites it is watching to the other satellites so that they may begin observing before they also enter the patch. *We 'Exploit' TEC over the Poles.*

Talk Roadmap

- Distributed Space Missions
- The Distributed Spacecraft Autonomy Project
- DSA-Starling
  - DSA-Starling Software Architecture
- DSA-LPNT
- Lessons Learned
- Bibliography

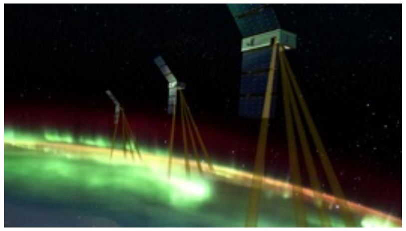

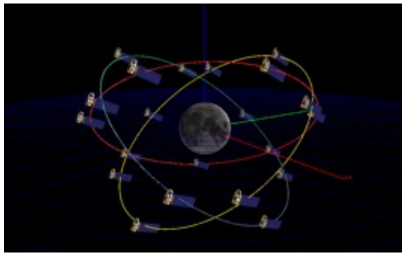

DSA-Starling

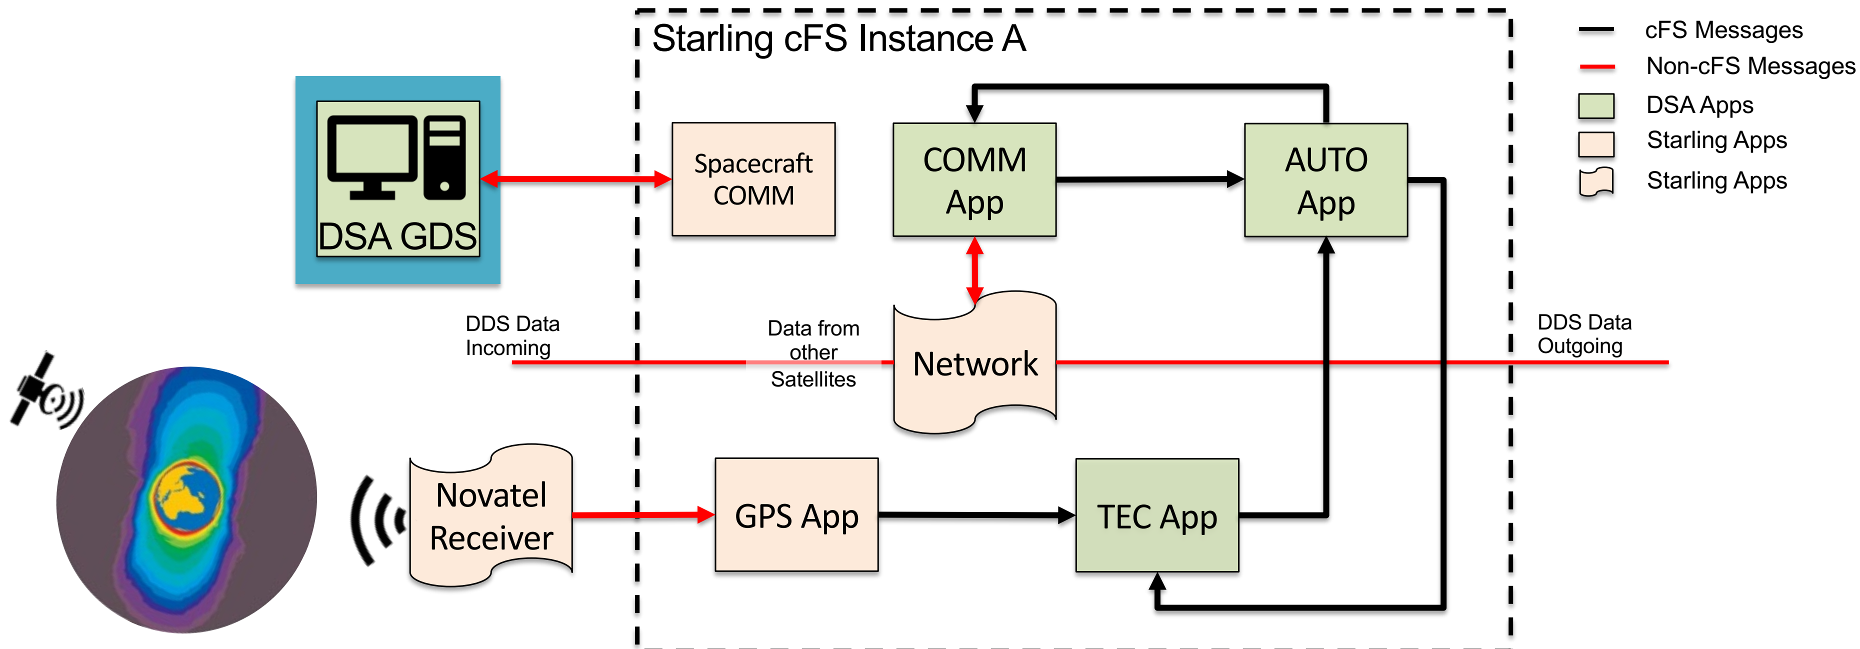

DSA-Starling

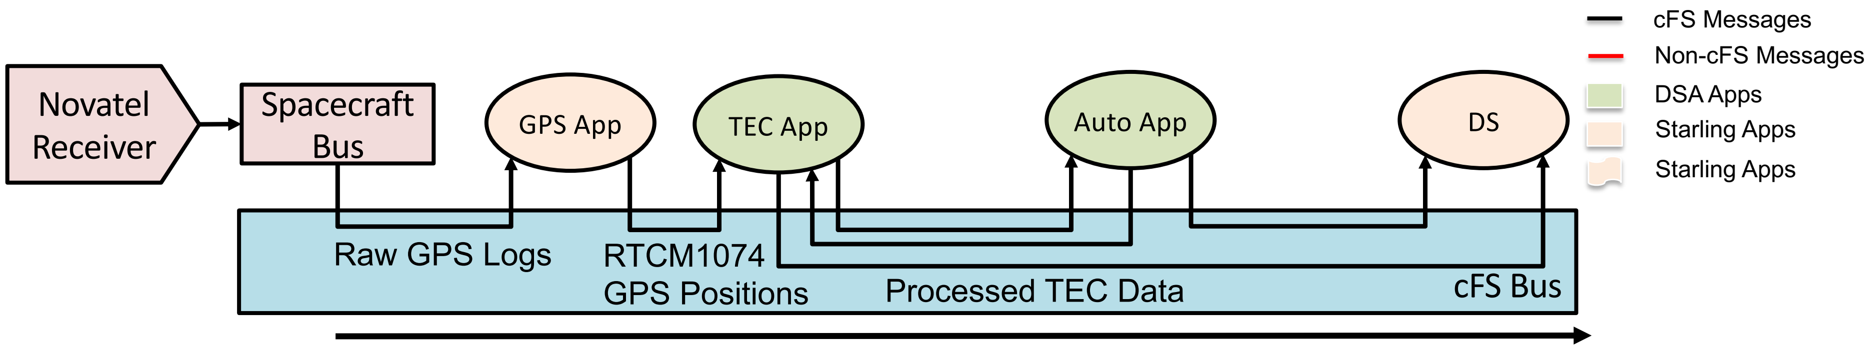

Spacecraft Bus streams Raw GPS Logs from Novatel Receiver:
RTCM1074
GPSEphem

TEC App

1. Organizes GPS logs into tracks $(T_1, T_2, \ldots, T_n)$, n=ID of GPS sat
2. Converts tracks into TEC timeseries
3. Outputs **reward values** for each track based on **GPS position** and the **presence of TEC features** in the timeseries

Auto App

Fuses local reward values with other satellites' **reward values** to calculate the subset of GPS channels to monitor

Talk Roadmap

- Distributed Space Missions
- The Distributed Spacecraft Autonomy Project
- DSA-Starling
  - The AUTO App
- DSA-LPNT
- Lessons Learned
- Bibliography

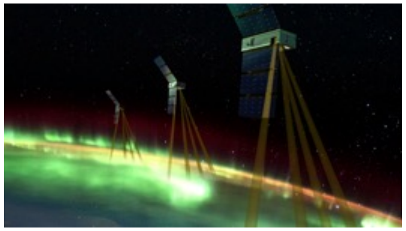

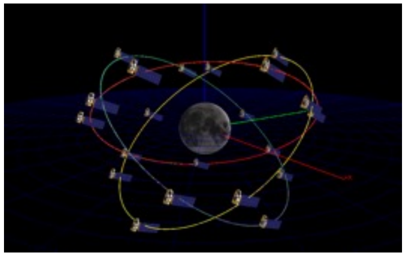

DSA-Starling

- **The Math Behind The AUTO App**
  - The 'Explore' problem is to monitor as many channels as possible. More distant GPS satellites are less useful.
  - The 'Exploit' problem is to listen to channels with high TEC signals.
- **Setting up the problem**
  - Requires combing local and global information from other DSA satellites, which requires communication.
- **The problem isn't just a matching problem**
  - Additional Constraints: Channel capacity, minimum coverage requirements.
  - Balancing/tuning the Explore/Exploit objectives is valuable; the 'blended' problem isn't quite a matching problem.

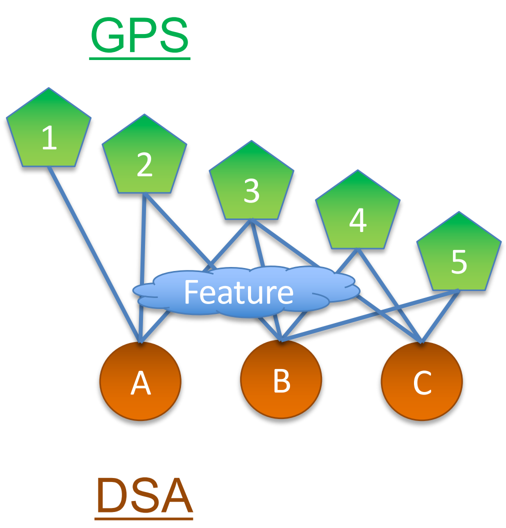

DSA-Starling

The Math Behind The AUTO App

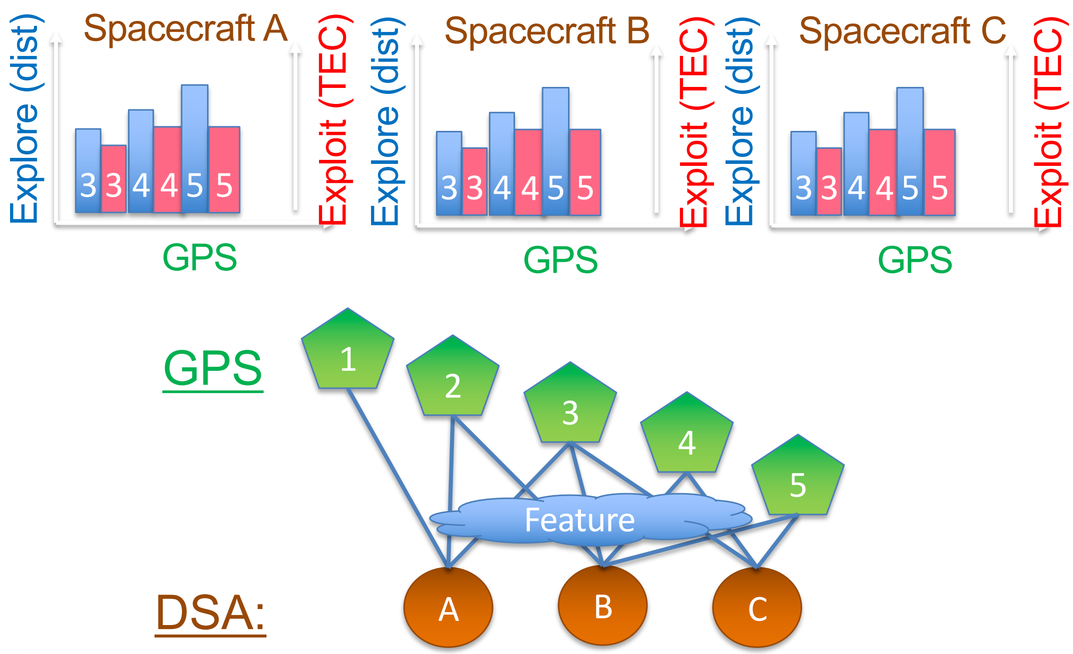

Shared Information

- Visibility: A → {1,2,3}, B → {2,3,4,5}, …
- Observations: <gpsId, exploreReward, exploitReward>
- Explore reward ∝ distance(sat s, gps g)
- Exploit reward ∝ TEC(sat s, gps g)

DSA-Starling

The Math Behind The AUTO App

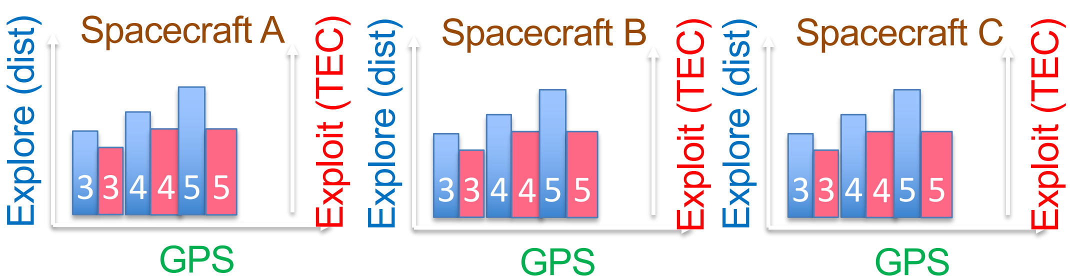

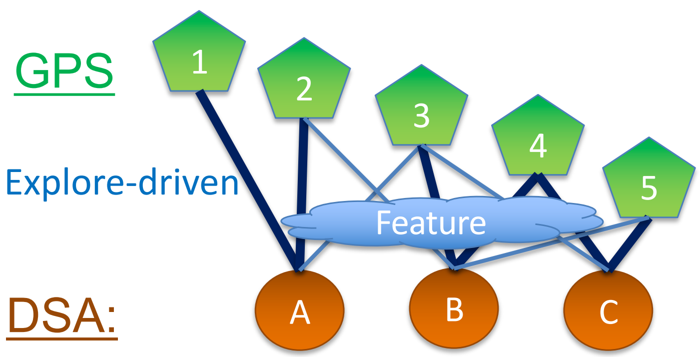

Shared Information

- Visibility: A → {1,2,3}, B → {2,3,4,5}, …
- Observations: <gpsId, exploreReward, exploitReward>
- Explore reward ∝ distance(sat s, gps g)
- Exploit reward ∝ TEC(sat s, gps g)

Exploration Scenario

- Objective driven by **explore rewards**
- GPS assignments based on **distance**

DSA-Starling

The Math Behind The AUTO App

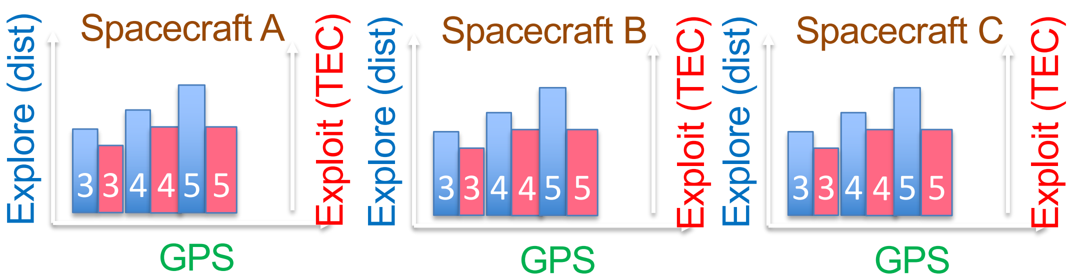

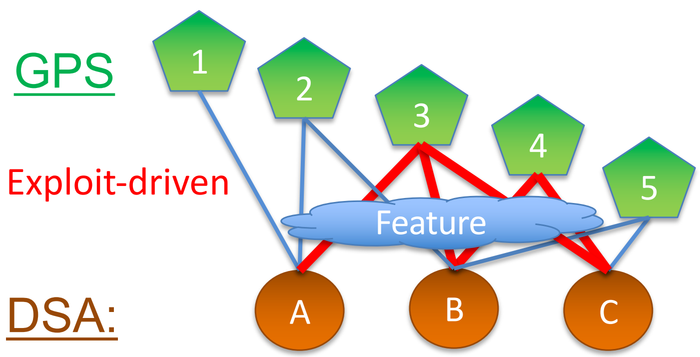

Shared Information

- Visibility: A → {1,2,3}, B → {2,3,4,5}, …
- Observations: <gpsId, exploreReward, exploitReward>
- Explore reward ∝ distance(sat s, gps g)
- Exploit reward ∝ TEC(sat s, gps g)

Exploitation Scenario

- Objective driven by **exploit rewards**
- GPS assignments based on **prior TEC observations**

DSA-Starling

The Math Behind the AUTO App

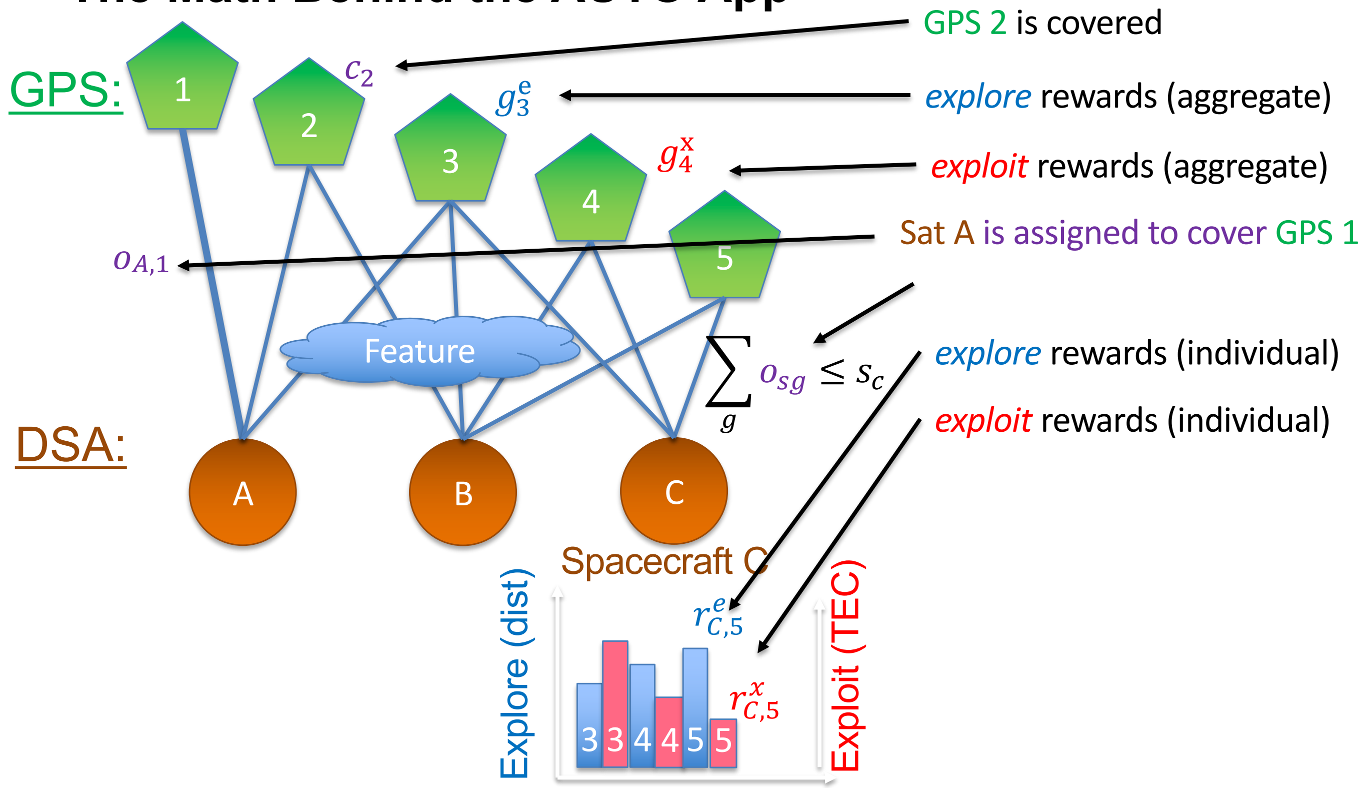

DSA-Starling

The AUTO App is posed and solved as an MILP:

Decision Variables:

$o_{s,g} = 1 \Leftrightarrow DSA\ s$ is assigned to observe GPS g, $o_{s,g} \in \{0,1\}$, For all GPS visible to DSA s

$c_g = 1 \Leftrightarrow GPS$ is covered by at least n DSA s $c_g \in \{0,1\}$, For all visible GPS

Individual rewards:

$r^e_{s,g}$ = *explore* rewards for DSA *s* observing GPS *g*

$r^x_{s,g}$ = *exploit* rewards for DSA *s* observing GPS *g*

Aggregate rewards:

g^e = *explore* rewards for all DSA *s* observing GPS g (avg. or max.)

g^x = *exploit* rewards for all DSA *s* observing GPS g (avg. or max.)

Objective

$$max \sum_{s,g} \alpha\ r^e_{s,g}\ o_{sg} + (1-\alpha)\ r^x_{s,g}\ o_{sg} \quad + \quad \sum_{s,g} \alpha\ g^e\ c_g + (1-\alpha)\ g^x\ c_g$$

α = parameter to control balance between *explore* and *exploit*, $0 \le \alpha \le 1$

Constraints:

- $\forall s$: $\sum_g o_{sg} \le s_c$ Each sat's assignments cannot exceed it's channel **capacity**
- $\forall g$: $\sum_s o_{sg} \ge n$ Each GPS must be **covered** by minimum # of satellites (optional)

DSA-Starling

- **The AUTO App design**
  - The MILP solves a ‘blended’ Explore and Exploit problem.
  - The α parameter can be tuned to experiment with different balances of Explore/Exploit.
  - We chose lp_solve as the foundation of AUTO.
- **The DSA ‘Consensus’ Algorithm**
  - DSA satellites share their local TEC observations to achieve ‘distributed consensus’.
  - The MILP is *constructed* from this information.
  - Each DSA sat solves the same problem; each satellite obtains identical solutions.
  - Each satellite execute the same plans (without the need to communicate).

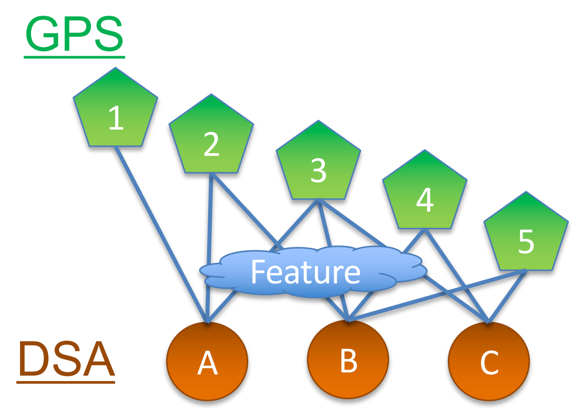

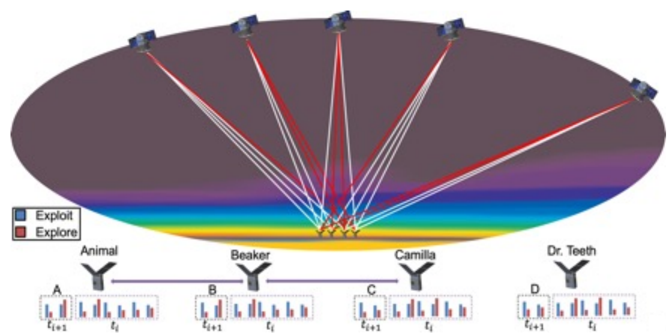

Talk Roadmap

- Distributed Space Missions
- The Distributed Spacecraft Autonomy Project
- DSA-Starling
- DSA-LPNT
  Lunar Position Navigation and Timing Services
- Lessons Learned
- Bibliography

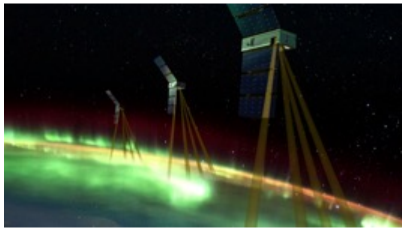

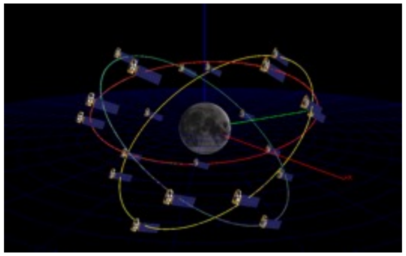

DSA-LPNT

- **A New Use Case: Position Navigation and Timing**
  - Deep Space Network (DSN) based localization is available when missions are in view of the Earth; low-cost surface missions may not be able to support the large power, mass, and weight requirements that these navigation solutions would entail. Additionally, DSN is heavily oversubscribed.
  - High Altitude / Weak GPS has positioning errors of 100 m, which will not meet Lunar mission localization requirements.
  - We explore the potential of ad-hoc PNT services using proposed Lunar small-sat science missions in low Lunar orbit, at 100km. As an alternative, one could use dedicated spacecraft in 'frozen' lunar orbits (requiring low/no propellant for orbit maintenance) at 4000km.

Mission	**Altitudes (km)**	**Incl. (deg)**
Luna-Hmap (completed)	5 x 3000	90
LRO (completed)	20 x 165	90
Lunar Flashlight (failed)	20 x 1000	90
LADEE (completed)	25 x 60	157
KPLO/Danuri (ongoing)	100 x 100	90
Lunar Ice Cube (failed)	100 x 100	90
SELENE (completed)	100 x 100	90
Lunar Trailblazer (planned)	100km	90
Chandrayaan-1 (completed)	200 x 200	90

DSA-LPNT is an autonomy demonstration with **a new use case**

DSA-LPNT

- **A New Use Case: Position Navigation and Timing**
  - Ideally, localization and PNT service provision should be done with as little interaction with Earth-based mission control as possible.
  - Communication between satellites is not necessarily pervasive; the orbits of the scientific missions permit pairs of spacecraft to communicate directly periodically, or not at all. Information exchange, therefore, must rely on establishing relays.
  - If LPNT uses spacecraft whose primary mission is science, they will perform other tasks, or have constraints on how often they can perform LPNT-related duties.
  - Limited time and resources require satellites to schedule communication activities to provide the best possible quality of PNT service.
  - PNT only works if the service providing assets know exactly where *they* are so they can tell *you* exactly where *you* are. (Consider the time and energy needed to keep terrestrial GPS accurate.)

DSA-LPNT

DSA-Starling

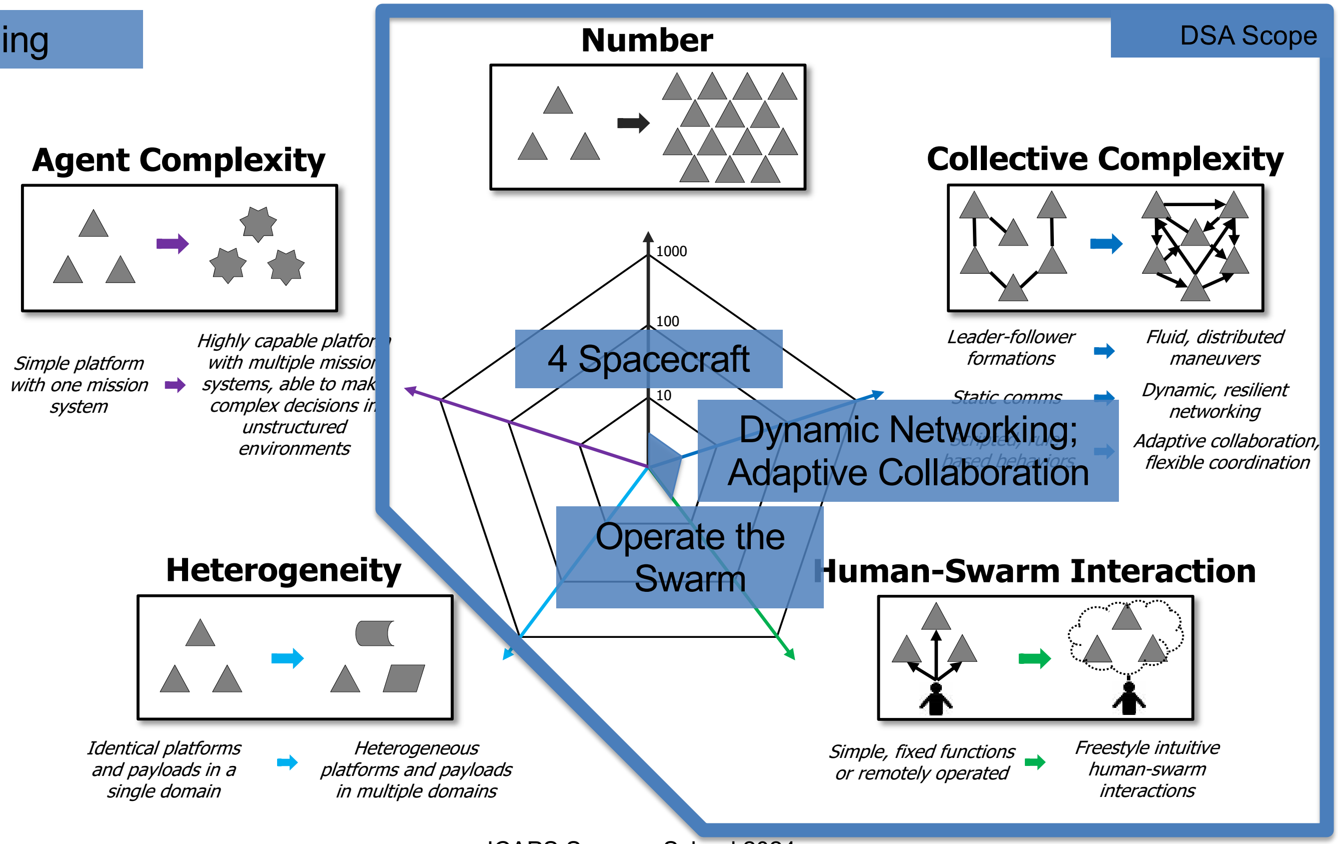

DSA-LPNT

DSA-Starling

DSA-LPNT

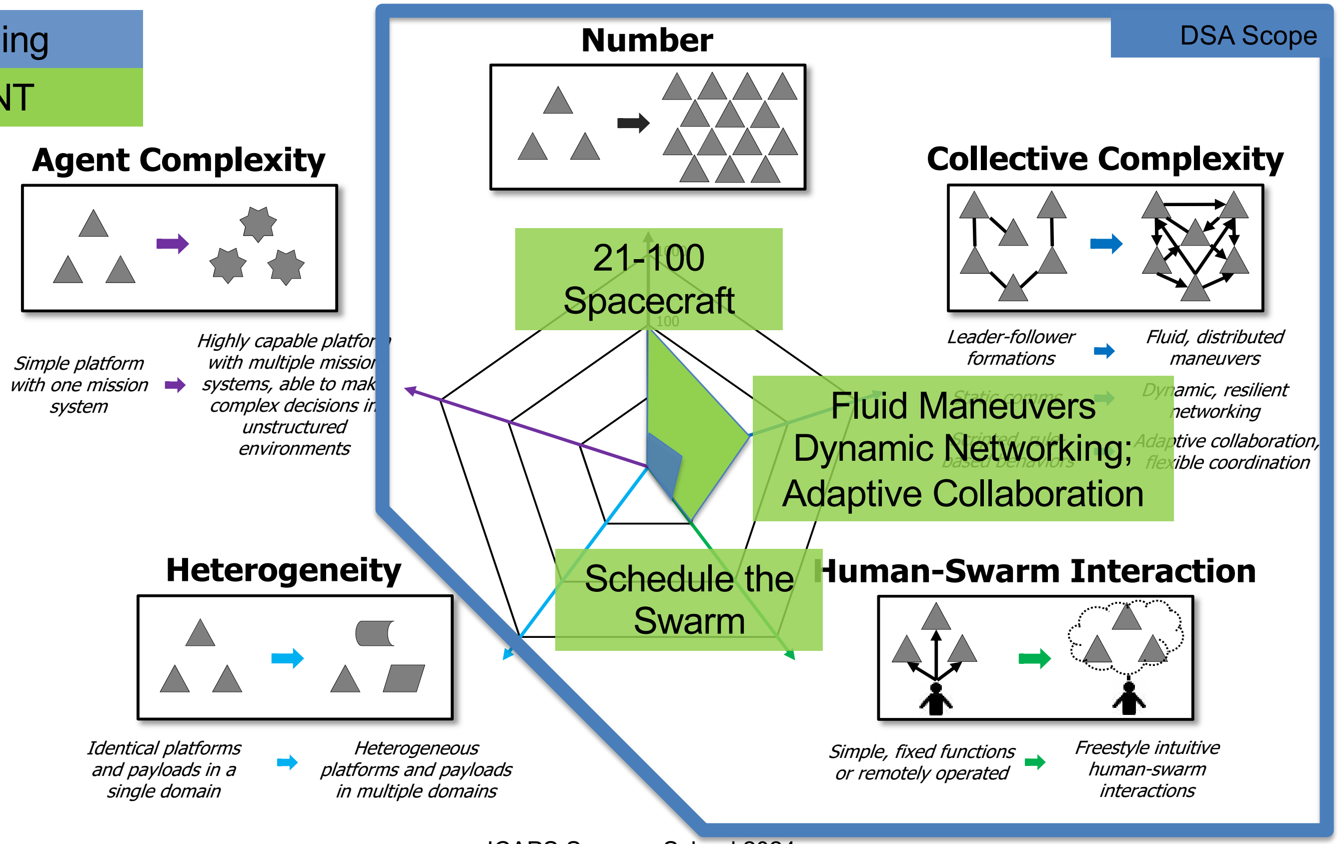

Talk Roadmap

- Distributed Space Missions
- The Distributed Spacecraft Autonomy Project
- DSA-Starling
- DSA-LPNT
  - A Brief Diversion: Kalman Filters
- Lessons Learned
- Bibliography

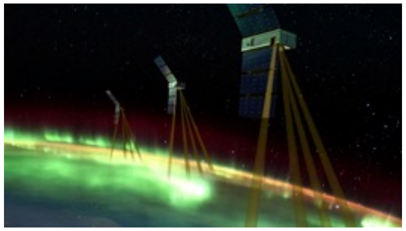

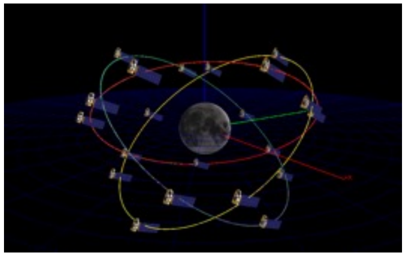

A Brief Diversion: Kalman Filters

- Kalman Filters are used for guidance, navigation and control of spacecraft (among other things.). For our purposes, the Kalman Filter is used to refine the position and velocity estimates of spacecraft in their orbit around the Moon.
- DSA really uses a Distributed Extended Kalman Filter (DEKF), which splits the state amongst spacecraft, and drives the need to communicate information to update the navigation state.

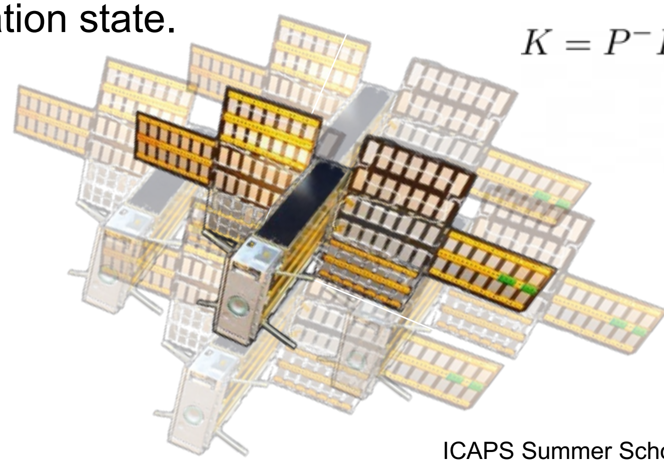

$$K = P^{-}H(HP^{-}H^{T} + R + \bar{H}\bar{P}^{-}\bar{H}^{T})^{-1}$$

$$P^{+} = (I - KH)P^{-}$$

$$x^{+} = x_{-} + K(z - Hx^{-})$$

DSA-LPNT

- **The parts of the (DE)KF we really need:**
  - Covariance matrix P_i: each spacecraft's uncertainty in its position and velocity. (6 x 6 matrix). (This is the only matrix specific to a single S/C on this slide.)
  - Measurements z: pseudorange, velocity between pairs of spacecraft. (2(n-1) dimensional vector if there are n spacecraft.)
  - Measurement sensitivity matrix H describes how the measurements (pseudoranges and velocities) inform the receiving spacecraft's estimate of its position and velocity. (2(n - 1) x 6 dimensional matrix.)
  - R: Measurement noise. (2(n-1) x 2(n-1) dimensional matrix.)
  - The Kalman Gain K is a function of P, H and R. It represents the bias between between predicted state and the measurement. (6 x 2(n-1) dimensional matrix.)

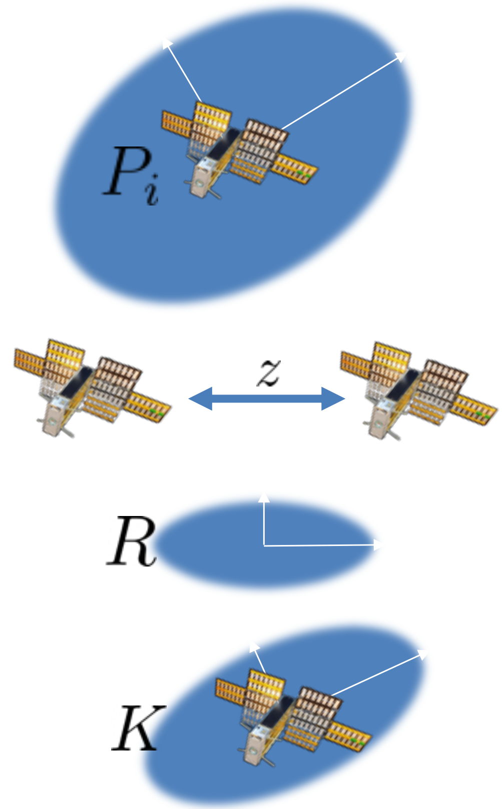

Talk Roadmap

- Distributed Space Missions
- The Distributed Spacecraft Autonomy Project
- DSA-Starling
- DSA-LPNT
  - The Measurement Selection Problem
- Lessons Learned
- Bibliography

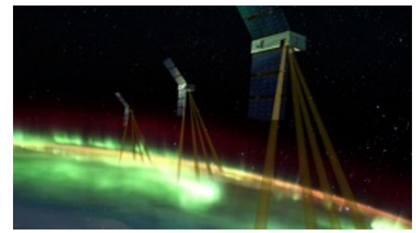

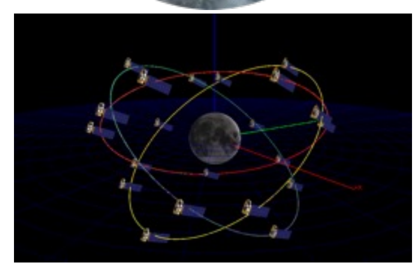

DSA-LPNT

- **The Measurement Selection Problem**
  - We have assumed that the PNT service providing spacecraft are non-dedicated, i.e. they have other things to do.
  - Spacecraft may have limited communications capability (limited numbers of antennas, ranges, frequency and other hardware limitations).
  - Measurement requires two spacecraft to simultaneously point antenna at each-other and activate radios. Both obtain measurements useful for localization. This means they need to maneuver or orient their antennas and simultaneously operate their radios.
  - This is why we can't have nice things, like all the measurements from all spacecraft all the time to drive the DEKF updates.

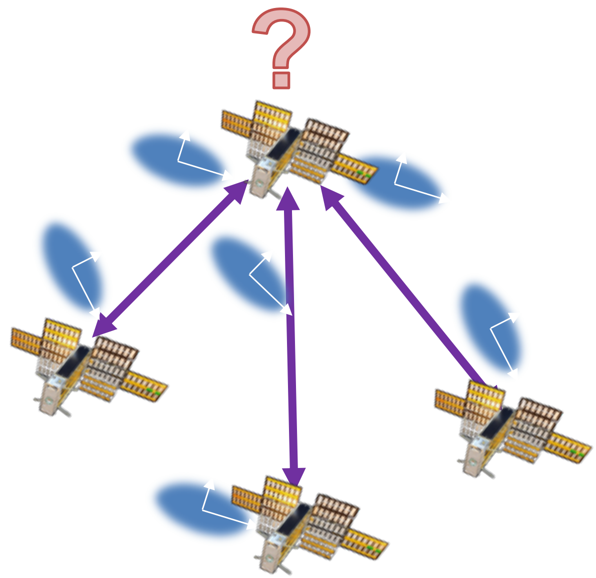

- **The Measurement Selection Problem**
  - As (part of) the DEKF Measurement Update Step, we can compute $P_{i\leftarrow j}$, which evaluates how spacecraft i's state uncertainty changes if we only use measurements from spacecraft j.
  - We only need the rows corresponding to measurement j from the Measurement Sensitivity Matrix *H*, *not the actual measurement*, to compute this quantity!
  - We can then use $P_{i\leftarrow j}$ to *choose the best measurement* to make.

$$K_{i\leftarrow j} = P_i h_j^T \left(h_j P_i h_j^T + r_{jj} + \bar{h}_j \bar{P}_i \bar{h}_j^T \right)^{-1}$$

$$P_{i\leftarrow j} = (I - K_{i\leftarrow j} h_j) P_i$$

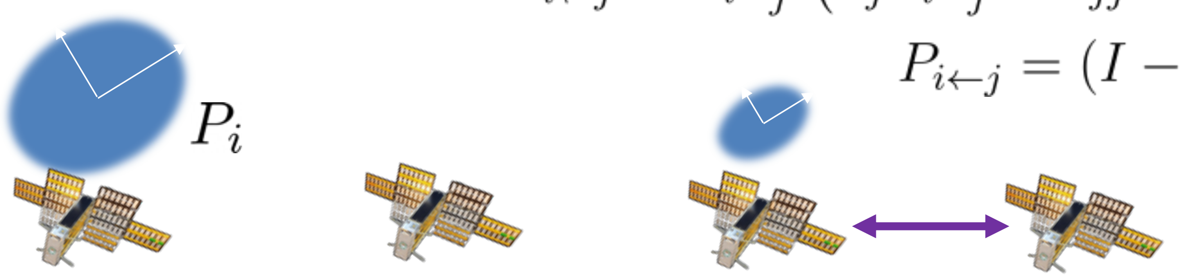

DSA-LPNT

- **The parts of the DEKF we really need:**
  - The square of the Frobenius Norm of the matrix $P_{i \leftarrow j}$ is:

$$||P_{i \leftarrow j}||_F^2 = \sum_{ij} p_{ij}^2$$

  - This is the scalar estimate of the resulting uncertainty of spacecraft i updating its position by communicating with spacecraft j. A large measure indicates high covariance, and a low measure indicates low covariance.
  - We can select the measurement that results in the lowest covariance.

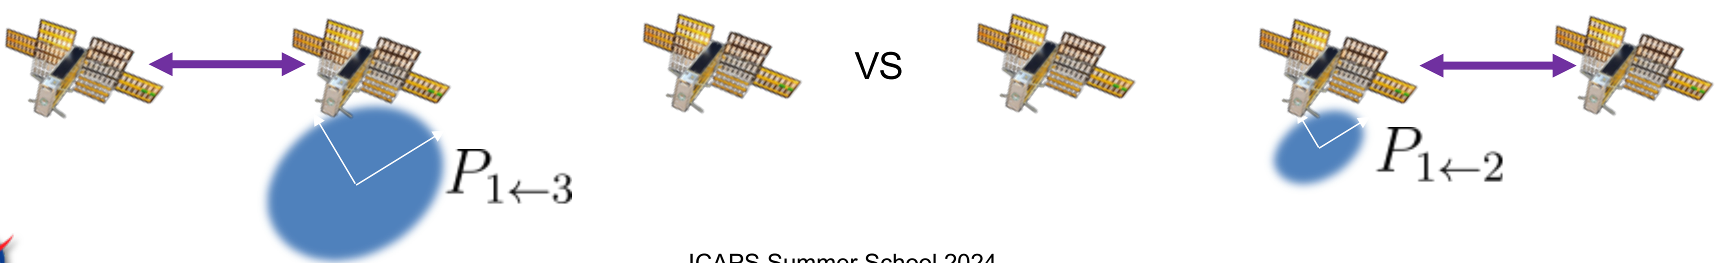

DSA-LPNT

- **The Distributed Measurement Selection Problem is a Matching Problem**
  - Suppose each spacecraft only has a *single* antenna.
  - Define $f_{ij} \equiv ||P_{i \leftarrow j}||^2_{\mathrm{F}}$.
  - The ‘global’ quality for both spacecraft i and j if they communicate is thus $f_{ij} + f_{ji}$ (each edge includes contributions from both spacecraft that communicate.)
  - If we have all of these qualities, we can solve the problem of finding the *best set of simultaneous communication activities to perform,* subject to the constraint that each spacecraft can perform only one communication act.
  - The resulting problem is a *minimum matching problem* (because we want to minimize the sum of covariances across the LPNT network.)

Talk Roadmap

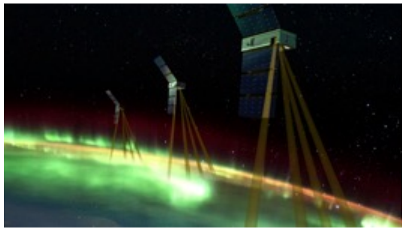

- Distributed Space Missions
- The Distributed Spacecraft Autonomy Project
- DSA-Starling
- DSA-LPNT
  From AUTO to MPLAN
- Lessons Learned
- Bibliography

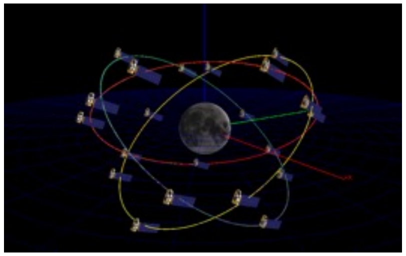

DSA-LPNT

- ~~The Distributed Measurement Selection Problem is a Matching Problem~~

DSA-LPNT

- **From AUTO to MPLAN**
  - Suppose spacecraft have *two* antennas. We can compute $P_{i \leftarrow j,k}$, the covariance matrix resulting from communicating with 2 spacecraft j,k. (We explain why we need 2 antennas later.)
  - Define $f_{ijk} \equiv ||P_{i \leftarrow j,k}||_F^2$. This is the scalar estimate of the resulting uncertainty of spacecraft i updating its position by communicating with j and k (from computing the Frob. Norm of the resulting covariance matrix).
  - We similarly define $f_i \equiv ||P_i||_F^2$ which is provably *larger* than $P_{i \leftarrow j}$ and $P_{i \leftarrow j,k}$
  - This requires generalization of the matching problem to a *multi-matching problem*: each spacecraft can now choose to communicate with 0-2 neighboring spacecraft.
  - The consequence of this generalization is that we now have something like an MILP, and we can reuse AUTO’s lp_solve technology to solve the resulting problem.

DSA-LPNT

Ingredients of MPLAN App Formal Problem Formulation

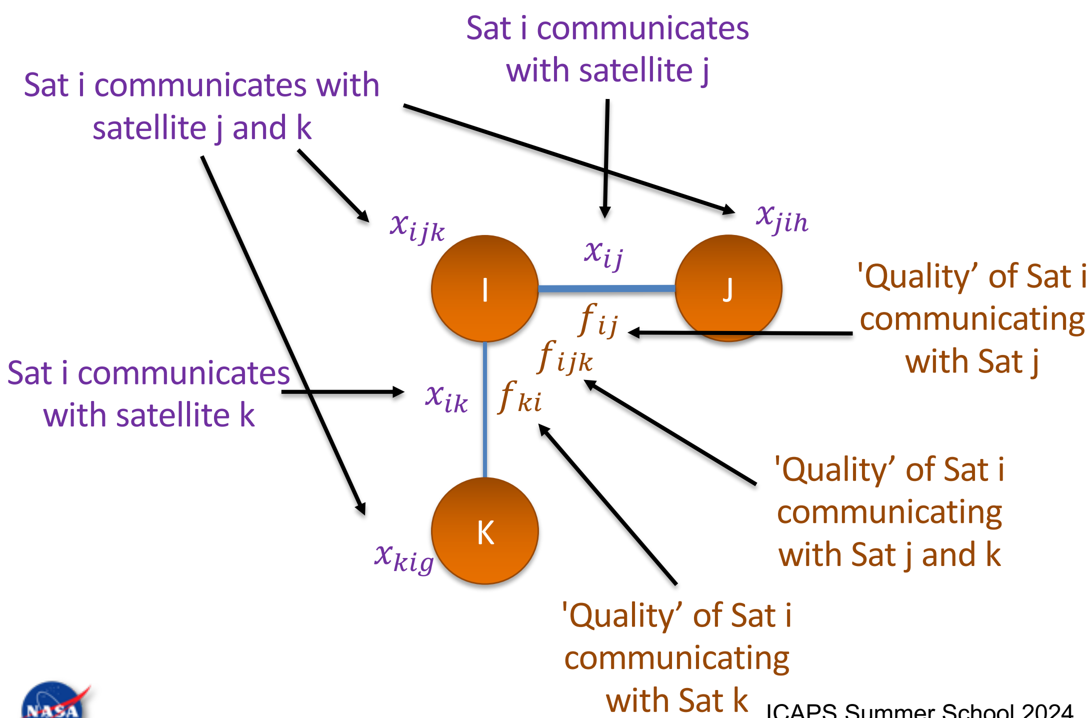

Decision Variables: x_{ijk}, x_{ij}, x_i

Individual costs: f_{ijk}, f_{ij}, f_i

Objective

$$min \sum_{i,j,k} f_{ijk}\, x_{ijk} + \sum_{i,j} f_{ij}\, x_{ij} + \sum_{i} f_i\, x_i$$

Constraints:

$$\sum_{i,j,k} x_{ijk} + \sum_{i,j} x_{ij} + \sum_{i} x_i = 1$$

$$\sum_{k} x_{ijk} + x_{ij} = \sum_{k} x_{jik} + x_{ji}$$

DSA-LPNT

- **The DSA 'Consensus' Algorithm Revisited**
  - All spacecraft broadcast P_i and x_i . Each computes H and K, and subsequently, all f_{ijk} from P, K, H and R. Spacecraft also need to then broadcast f_{ijk} to achieve global consensus.
  - Each spacecraft solves the same multi-matching problem; each obtains identical solutions. Each spacecraft executes the same plans to communicate (which changes the communication topology) to provide measurements for the DEKF update.
  - Thus, each step of localization requires 2 communication topologies, a 'broadcast network topology' to send the state, covariance matrices, and matching weights, and the 'measurement network' topology to provide measurements for the DEKF update.
- **We must 'Communicate to communicate!' How?**
  - Each spacecraft has 2 antennas; a fully connected broadcast network is, therefore, a Hamiltonian Path (HP).
  - The HP Broadcast network topologies are computed offline. Doing so requires offline analysis of orbits and committing to a fixed update schedule.

DSA-LPNT

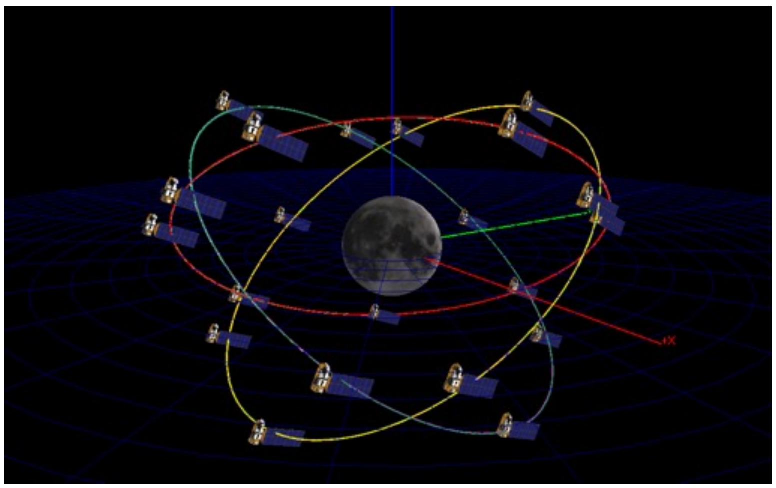

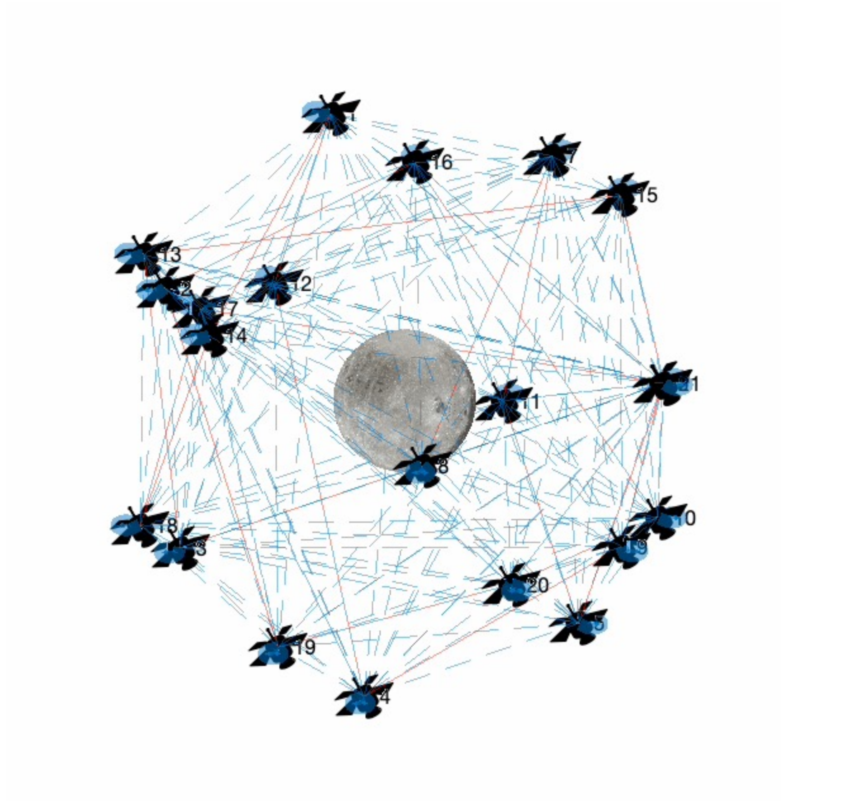

Talk Roadmap

- Distributed Space Missions
- The Distributed Spacecraft Autonomy Project
- DSA-Starling
- DSA-LPNT
  From MPLAN to DSA LPNT Software Stack
- Lessons Learned
- Bibliography

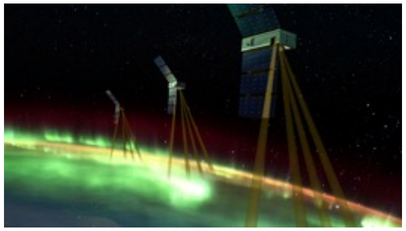

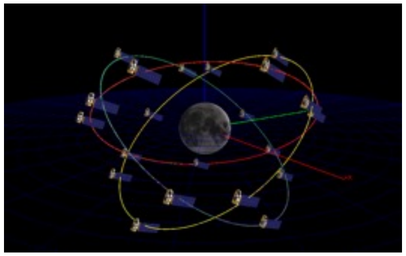

DSA-LPNT

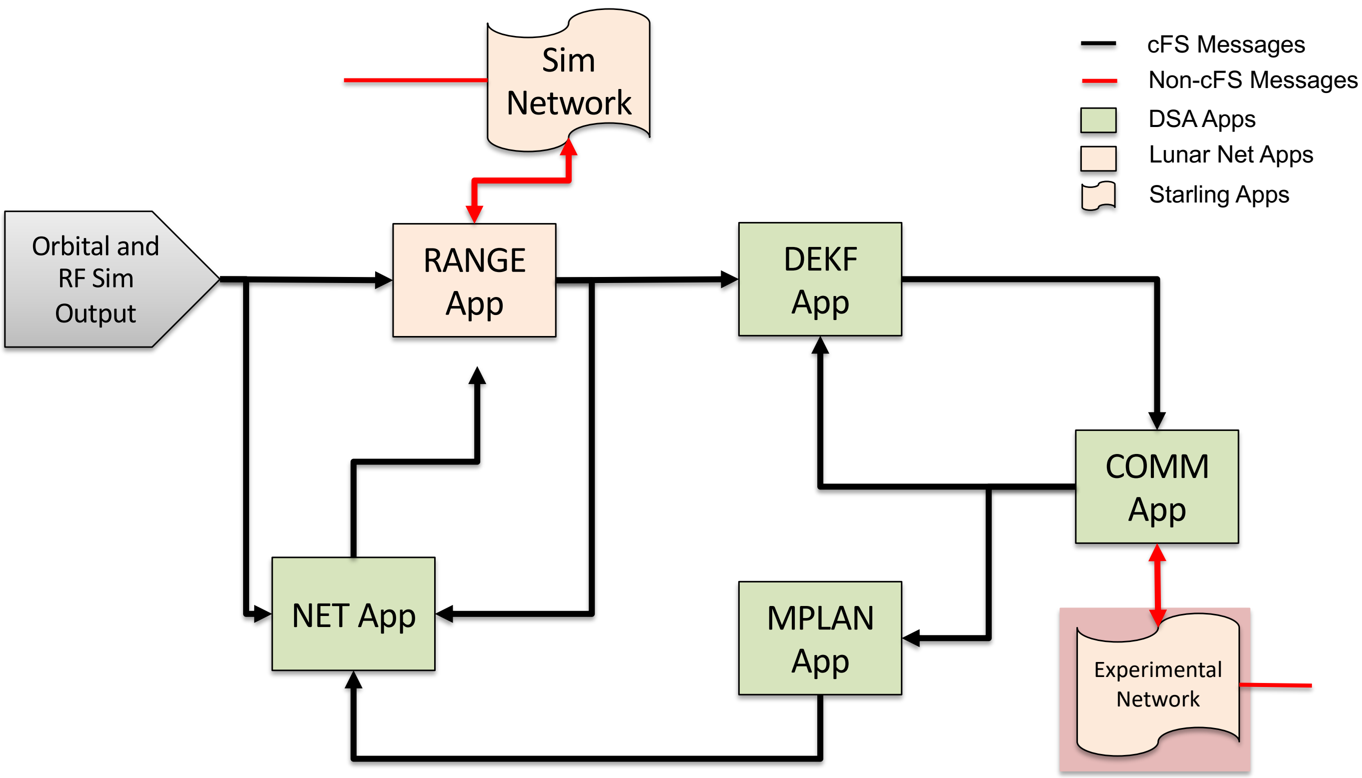
cFS Messages
Non-cFS Messages
DSA Apps
Lunar Net Apps
Starling Apps
Sim Network
Orbital and RF Sim Output
RANGE App
DEKF App
COMM App
NET App
MPLAN App
Experimental Network

DSA-LPNT

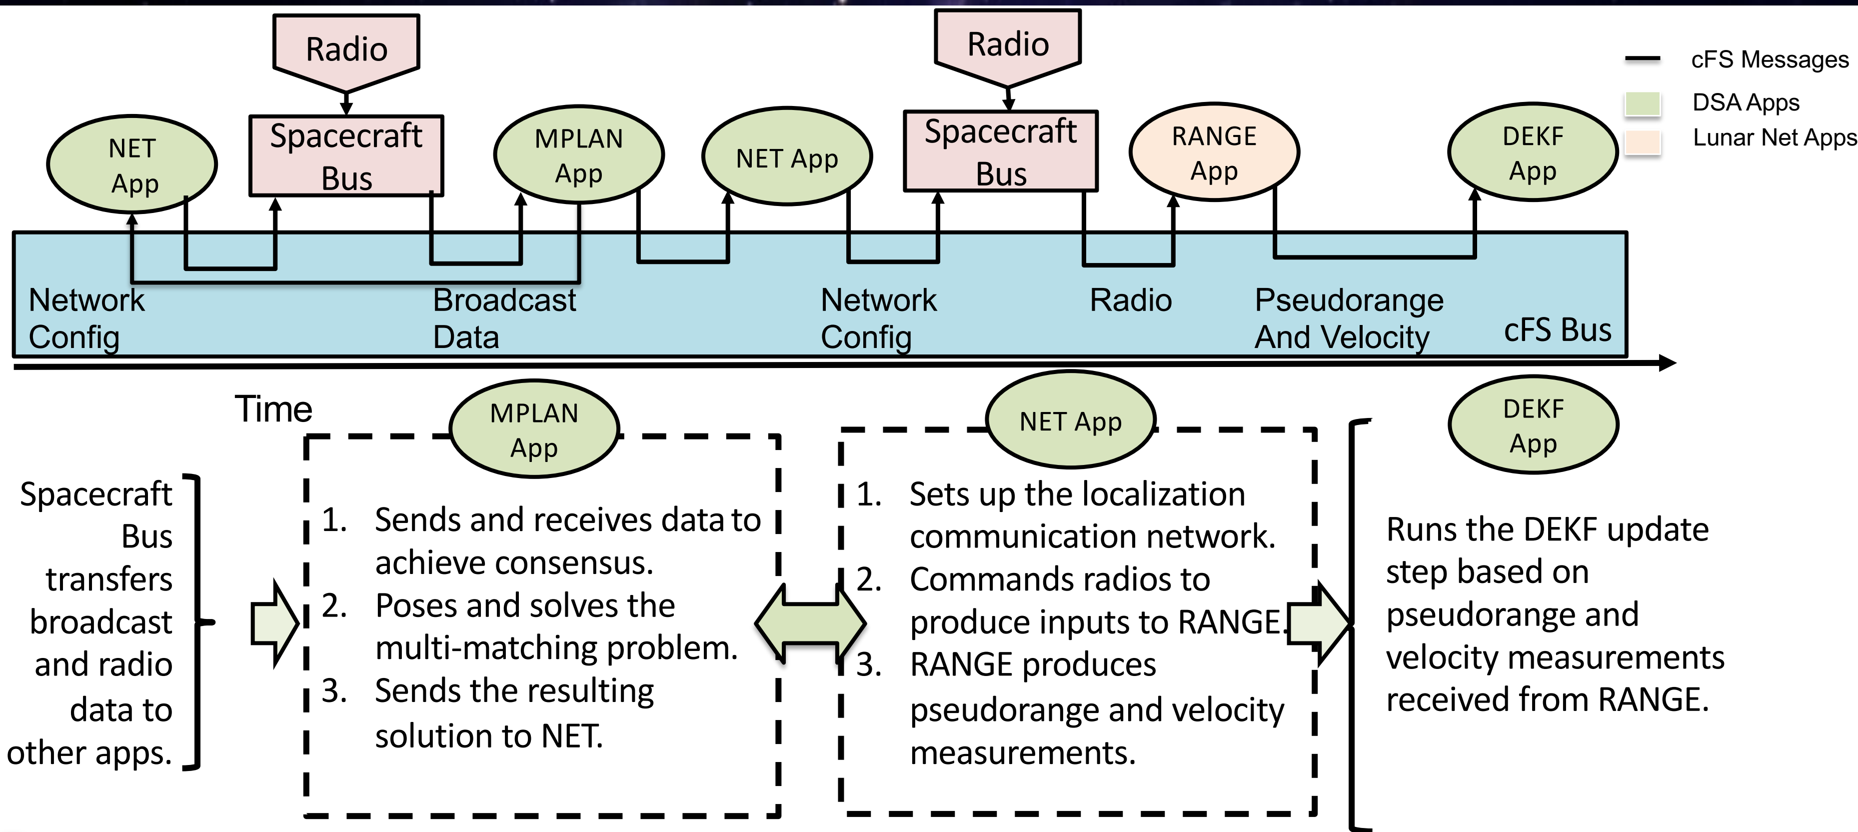
Radio
Spacecraft Bus
Radio
Spacecraft Bus
NET App
MPLAN App
NET App
RANGE App
DEKF App
cFS Messages
DSA Apps
Lunar Net Apps
Network Config
Broadcast Data
Network Config
Radio
Pseudorange And Velocity
cFS Bus
Time
Spacecraft Bus transfers broadcast and radio data to other apps.
MPLAN App
1. Sends and receives data to achieve consensus.
2. Poses and solves the multi-matching problem.
3. Sends the resulting solution to NET.
NET App
1. Sets up the localization communication network.
2. Commands radios to produce inputs to RANGE.
3. RANGE produces pseudorange and velocity measurements.
DEKF App
Runs the DEKF update step based on pseudorange and velocity measurements received from RANGE.

DSA-LPNT

- **Analyzing Performance:**
  - We evaluate localization performance on a 21-node, 3 plane (7 S/C each) Lunar constellation in a 4000 km frozen orbit.
  - The constellation was simulated for 7 days.
  - Fully Sampled DEKF: every spacecraft receives measurements from every S/C in view.
  - 1-Matching: Every spacecraft is assumed to have a single antenna and thus can receive only 1 measurement. (And yet, somehow, they can magically coordinate...☺)
  - 2-Matching: Every spacecraft is assumed to have 2 antennas, and thus can receive 2 measurements.
  - The measurement update cadence is 10 minutes. This accounts for antenna slew rates, network latency, setup/teardown, and the need to establish the broadcast topology as well as the localization topology.
  - The results shown were performed in MATLab: we have plans to replicate them on our Processor in the Loop testbed, as well as scale from 21 to 100 nodes. (Maybe at 100 km...maybe not...because it's harder.)

DSA-LPNT

- **Results:**
  - 1-matching results are quite poor. (This is why we assume spacecraft have 2 antennas!)
  - 2-matching results: median error is only 1.6x that of the (physically unachievable) fully sampled DEKF.

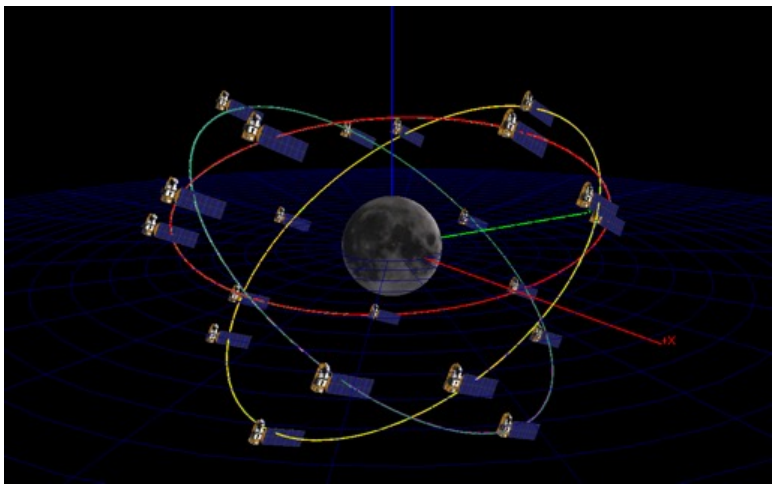

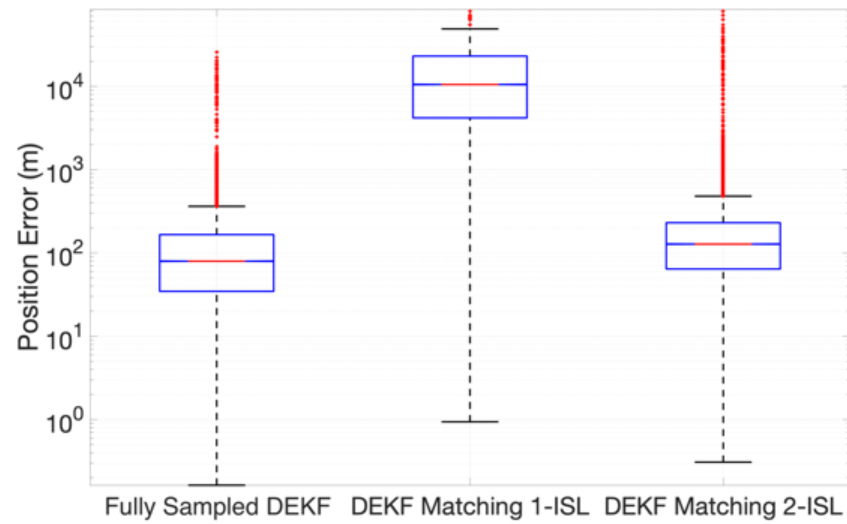

Talk Roadmap

- Distributed Space Missions
- The Distributed Spacecraft Autonomy Project
- DSA-Starling
- DSA-LPNT
- Lessons Learned
- Bibliography

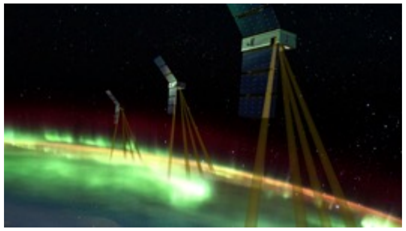

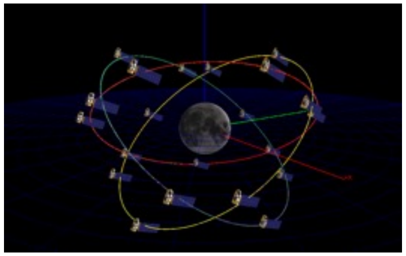

Lessons Learned

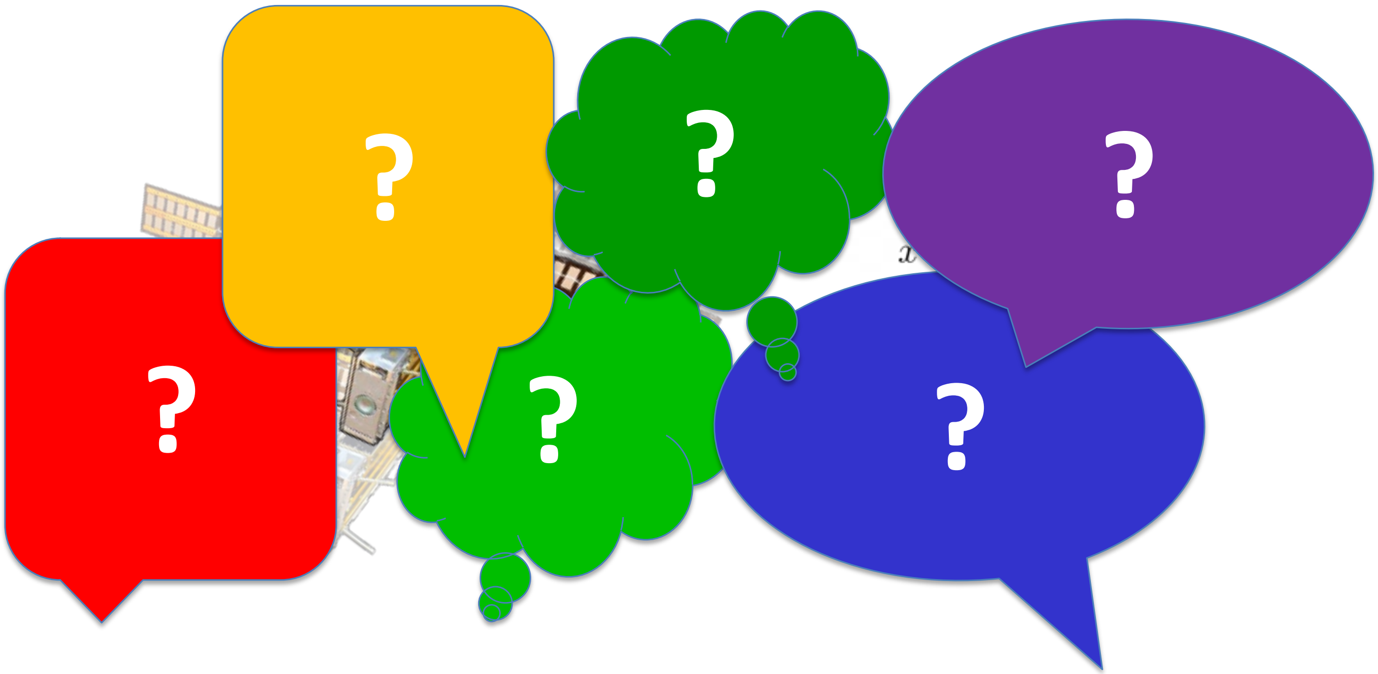

Lessons Learned

- **Solving the Right Problem, Part I: Applications**
  - It is easy to describe simple, contrived multi-agent autonomy problems in isolation, and solve them. When confronted by the real problem, many proposed problem formulations and their solutions are quickly shown to be inapplicable.
  - Examples: the simple looking matching problems for DSA-Starling and DSA-LPNT give way to more complex problem formulations, requiring a more general-purpose approach to each (i.e. using MILPs instead of specialized algorithms)

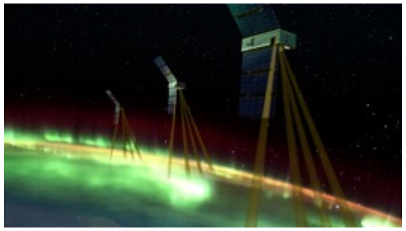

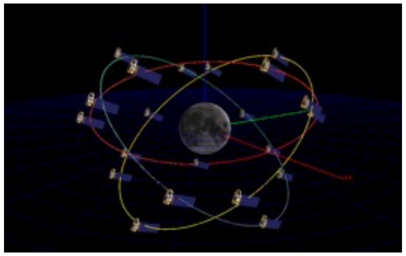

Lessons Learned

- **Solving the Right Problem, Part II: Real Problems are...Complicated**
  - It took considerable time, effort and research to formulate the planning problems we solved for both DSA-Starling and DSA-LPNT.
  - Examples: developing the matching intuition for DSA-LPNT.
  - It is also important to ensure that the problem you choose can actually be solved.
  - Example: It’s not yet clear that even the ‘fully sampled’ DEKF can localize a Lunar constellation at 100km!
- **Solving the Right Problem, Part III: Planning Problems We solved**
  - The AUTO App and MPLAN App.
  - HP (for the broadcast network).
  - Service Scheduling problem (which we didn’t talk about, but which turns out to be interesting, see Biblography).
  - STNs for DSA-Starling ops (to build command sequences, which we didn’t talk about).

Lessons Learned

- **The Drive for General Purpose Technology**
  - Because DSA was conceived of as a two-tier project, we wanted to reuse approaches as much as possible.
  - In particular, we reused the lp_solve and MILP formulation between DSA-Starling and DSA-LPNT, even though specialized approaches might have been better choices (but see previous slide).
  - This establishes a technology maturation path from DSA-Starling (4 spacecraft) to the much larger and more complex LPNT problem (21-100 spacecraft). Note the emphasis on technology, as opposed to theory: as computer scientists, it's easy to assert scalability, but showing the technology scales is quite a different matter!

$$max \sum_{s,g} \alpha\, r^e_{s,g} o_{sg} + (1-\alpha)\, r^x_{s,g}\, o_{sg}$$
$$+ \sum_{s,g} \alpha\ g^e\ c_g + (1-\alpha)\ g^x\, c_g$$
$$s.t. \forall s \sum_{s} o_{sg} < s_c$$
$$\forall g \sum_{g} o_{sg} > n$$

$$min \sum_{i,j,k} f_{ijk}\, x_{ijk} + \sum_{i,j} f_{ij}\, x_{ij} + \sum_{i} f_i\, x_i$$
$$s.t. \sum_{i,j,k} x_{ijk} + \sum_{i,j} x_{ij} \sum_{i} x_i = 1$$
$$\sum_{k} x_{ijk} + x_{ij} = \sum_{i,j} x_{jik} + x_{ji}$$

Lessons Learned

- **Can we talk?**
  - DSA-Starling and DSA-LPNT may appear to have similar assumptions about the availability of communication between spacecraft. In DSA-Starling, we (hopefully) rely on a broadcast network that allows spacecraft to communicate freely by abstracting away the details of the underlaying relay network (MANET). In DSA-LPNT, we established that HP broadcast relay networks existed, so we made the same assumption.
  - The reality of space-to-space communication is that spacecraft can't communicate freely. The Starling radios cannot be turned on constantly (nor do they need to be). For DSA-LPNT, the requirement to communicate in order to localize, combined with assumptions of limited directional antennas and the non-dedicated availability of the spacecraft, limits communication opportunities.
  - This drove our specific 'distributed consensus' approach, which is to ensure each agent solves the same problem by establishing global state.

Lessons Learned

- **The Existential Pleasures (and Dreads) of Engineering**
  - Considerable time and energy was spent not on algorithm development, but design trades, engineering, analysis, and quality assurance.
  - Example: Why did we fly lp_solve when Gurobi and CPLEX are available? Licensing, ease of integration and deployment, the list goes on...
- **The Existential Pleasures (and Dreads) of Operations**
  - Anything?
- **The Needs of the Team**

LESS TEXT

  - C/C++ skills are priority for S/C software. The ground systems could use other technology but we needed UI/UX expertise. Testing and quality assurance expertise is critical to S/C flight projects.
  - Being able to work across disciplines as part of a team is critical.
  - Example: DSA-LPNT required expertise in spacecraft hardware design, state estimation, and theoretical computer science.

Lessons Learned

- **What problems remain to be solved? Plenty!**
- **General:**
  - Think about other multi-agent approaches like DECPOMDPs, or consensus adjustment algorithms: would they work? Why or why not?
  - Constructing MILPs or other problems on the fly. Problem? Solution? Generalization?
- **DSA-Starling:**
  - What if we could turn radios off and on?
  - Incorporate power estimate into decision making?
- **DSA-LPNT:**
  - Can we trade out directional antennas for omni directional? What changes?
  - What if we were on Mars? Would MPNT work? Why or why not?
  - The existence of HPs is crucial to DEKF performance. What if there isn't one? What changes?

Bibliography

- **DSA papers**

  - D. Cellucci, N. Cramer, and J. Frank, “Distributed Spacecraft Autonomy,” in 2020 AIAA Ascend Conference. AIAA, 2020.
  - Hagenau, B.; Peters, B.; Burton, R.; Hashemi, K.; and Cramer, N. 2021. Introducing the Lunar Autonomous PNT System (LAPS) Simulator. In 2021 IEEE Aerospace Conference, 1–11.
  - N. Cramer, D. Cellucci, C. Adams, A. Sweet, M. Hejase, J. Frank, R. Levinson, S. Gridnev, L. Brown. Design and Testing of Autonomous Distributed Space Systems. Proceedings of 35th Annual Small Satellite Symposium, 2021.
  - B. Kempa, J. Frank, and N. Cramer. Swarm Mentality: Toward Automatic Swarm State Awareness with Runtime Verification. Proceedings of the AAAI Spring Symposium Series ‘Can We Talk’, 2022.
  - C. Adams, B. Kempa, M. Iatauro, J. Frank, W. Vaughan. Design and Testing of Autonomous Distributed Space Systems. Proceedings of 37th Annual Small Satellite Symposium, 2023.
  - C. Adams, B. Kempa, W. Vaughan and N. Cramer, "Development of a High-Performance, Heterogenous, Scalable Test-Bed for Distributed Spacecraft," 2023 IEEE Aerospace Conference, Big Sky, MT, USA, 2023, pp. 1-8

Bibliography

- **DSA Background and Motivation**
  - Schreiter, L. F., Arnold, D., Sterken, V. J., Jäggi, A., & Stolle, C. Imaging the topside ionosphere and the plasmasphere using Swarm GPS observations. 20th EGU General Assembly, EGU2018, Proceedings from the conference held 4-13 April, p.7023
  - J. L. Moigne, M. Little, and M. Cole, "New Observing Strategies (NOS) for Future NASA Earth Science Missions. Proc. IGARSS 2019.
    - https://esto.nasa.gov/
  - Science Missions," in AGU Fall Meeting, 2020.Cubesat Thinking Inside Box?
    - https://nap.nationalacademies.org/catalog/23503/achieving-science-with-cubesats-thinking-inside-the-box
  - CADRE
    - https://www.nasa.gov/cooperative-autonomous-distributed-robotic-exploration-cadre/#:~:text=NASA's%20Cooperative%20Autonomous%20Distributed%20Robotic,of%20NASA's%20A%2DPUFFER%20technology.
  - 2018 NASA Science Mission Directorate Autonomy Workshop
    - https://science.nasa.gov/technology/2018-autonomy-workshop/
  - AAAI 2022 Spring Symposium Series: Can We Talk?
    - https://sites.google.com/view/aaai-2022-spring-symposium/home

Bibliography

- **Related DSA papers**
  - S. Niemoeller and R. Burton and J. Frank and R. Levinson and N. Cramer. Scheduling Position, Navigation and Time Service Requests from Non-dedicated Lunar Constellations. Proceedings of the 43rd IEEE Aerospace Conference, 2022.
  - J. Frank, R. Levinson, E. Hillsberg, N. Cramer and R. Burton. Distributed Scheduling of Position Estimation Updates in Ad-Hoc Lunar Constellations. Proceedings of the AAAI Spring Symposium Series 'Can We Talk', 2022
  - Walter Vaughan, Alan George, Brian Kempa, Daniel Cellucci, Nicholas Cramer. Evaluating Network Performance of Containerized Test Framework for Distributed Space Systems. Proceedings of 36th Annual Small Satellite Symposium, 2022.
- **Everything I Know About Kalman Filters I Learned From This Cartoon**
  - https://www.bzarg.com/p/how-a-kalman-filter-works-in-pictures/

The DSA TEAM

- **Core Team:**
  - Caleb Adams (PM)
  - Jeremy Frank (DPM)
  - Caroline Lassiter
  - Michael Iatauro
  - Sergei Gridnev
  - Rich Levinson
  - Nate Benz
  - Steven Hu
  - Michael Ho

- **Special Thanks To**
  - Nick Cramer (previous PM)
  - Roland Burton (PNT expert)

Thank you!


20km
20km
20km
Directive: Satellites maintain spacing at 20km.

The parts of the DEKF we need to explain distributed consensus

The general EKF update formula requires computing the measurement sensitivity matrix and the Kalman Gain, which requires all states from all nodes and all covariance matrixes. Our matching formulation, in turn, uses the full Measurement Sensitivity and Kalman Gain matrixes. As a result, we must communicate state and covariance information globally across the broadcast network.

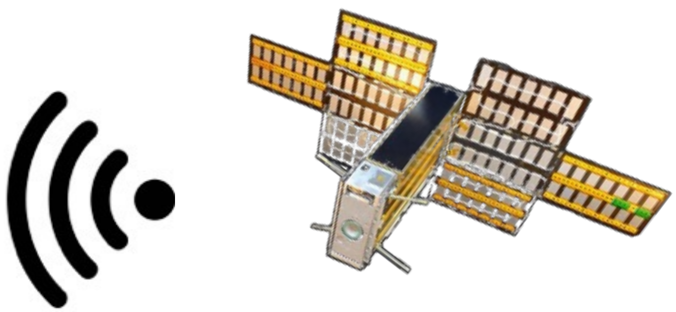

Computing the Kalman Gain Matrix

$$K = P^- H (H P^- H^T + R + \bar{H} \bar{P}^- \bar{H}^T)^{-1}$$

$$P^+ = (I - KH) P^-$$

The parts of the DEKF we need to explain distributed consensus

States

Sender $x_0 = \begin{bmatrix} r_0 \\ v_0 \end{bmatrix}$

Target $x_1 = \begin{bmatrix} r_1 \\ v_1 \end{bmatrix}$

Measurement

$$z = \begin{bmatrix} \|r_1 - r_0\| \\ \frac{(v_1 - v_0)^{\mathrm{T}}(r_1 - r_0)}{\|r_1 - r_0\|} \end{bmatrix}$$

Measurement Covariance assumed known and constant

$$R = \begin{bmatrix} \sigma_r^2 & 0 \\ 0 & \sigma_v^2 \end{bmatrix}$$

Computing the Measurement Sensitivity Matrix

$$H = \frac{\partial z}{\partial x} = \begin{bmatrix} -\frac{(r_1 - r_0)^{\mathrm{T}}}{\|r_1 - r_0\|} & 0 \\ -(v_1 - v_0)^{\mathrm{T}} \left(\frac{I}{\|r_1 - r_0\|} - \frac{(r_1 - r_0)(r_1 - r_0)^{\mathrm{T}}}{\|r_1 - r_0\|^3} \right) & -\frac{(r_1 - r_0)^{\mathrm{T}}}{\|r_1 - r_0\|} \end{bmatrix}$$

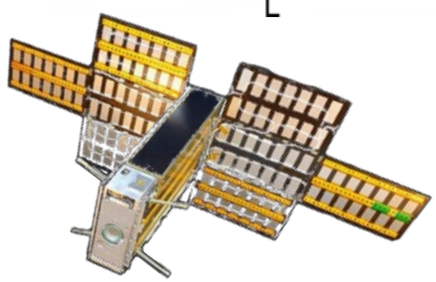

To compute *H* (above) we need the state.
To compute *K* (below) we need *P*, which means we need all the covariances.

Computing the Kalman Gain Matrix

$$K = P^- H (H P^- H^T + R + \bar{H} \bar{P}^- \bar{H}^T)^{-1}$$
$$P^+ = (I - KH) P^-$$

The parts of the DEKF we need to explain distributed consensus

We need both *H* and *K* to compute $P_{i \leftarrow j,k}$ and $f_{ijk} \equiv ||P_{i \leftarrow j,k}||_F^2$

$$K = P^- H(HP^- H^T + R + \bar{H}\bar{P}^- \bar{H}^T)^{-1}$$
$$P^+ = (I - KH)P^-$$

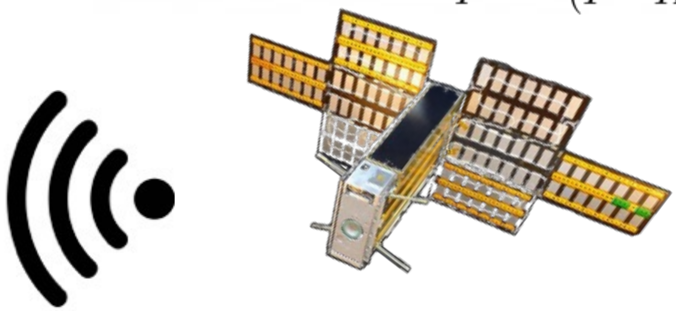

$$K_{i \leftarrow j} = P_i h_j^T \left(h_j P_i h_j^T + r_{jj} + \bar{h}_j \bar{P}_i \bar{h}_j^T\right)^{-1}$$
$$P_{i \leftarrow j} = (I - K_{i \leftarrow j} h_j) P_i$$
$$x_{i \leftarrow z} = x_i + K_{i \leftarrow j}(z_i - h_j x_i)$$

The parts of the DEKF we need to explain distributed consensus

Recall we define $f_i \equiv ||P_i||_F^2$ which is provably *larger* than $P_{i \leftarrow j}$ and $P_{i \leftarrow j,k}$

- Why? Fundamentally, because a measurement decreases uncertainty as long as assumptions underlying the KF are met.
- In the EKF and thus DEKF, this proof does not hold (but we use it anyway)

https://dsp.stackexchange.com/questions/75465/kalman-filter-why-do-we-decrease-the-state-uncertainty-regardless-of-the-curren
https://dsp.stackexchange.com/questions/79032/proving-that-the-uncertainty-can-not-increase-during-the-update-step-of-a-kalman
https://en.wikipedia.org/wiki/Kalman_filter
https://en.wikipedia.org/wiki/Extended_Kalman_filter

The parts of the DEKF we need to explain distributed consensus

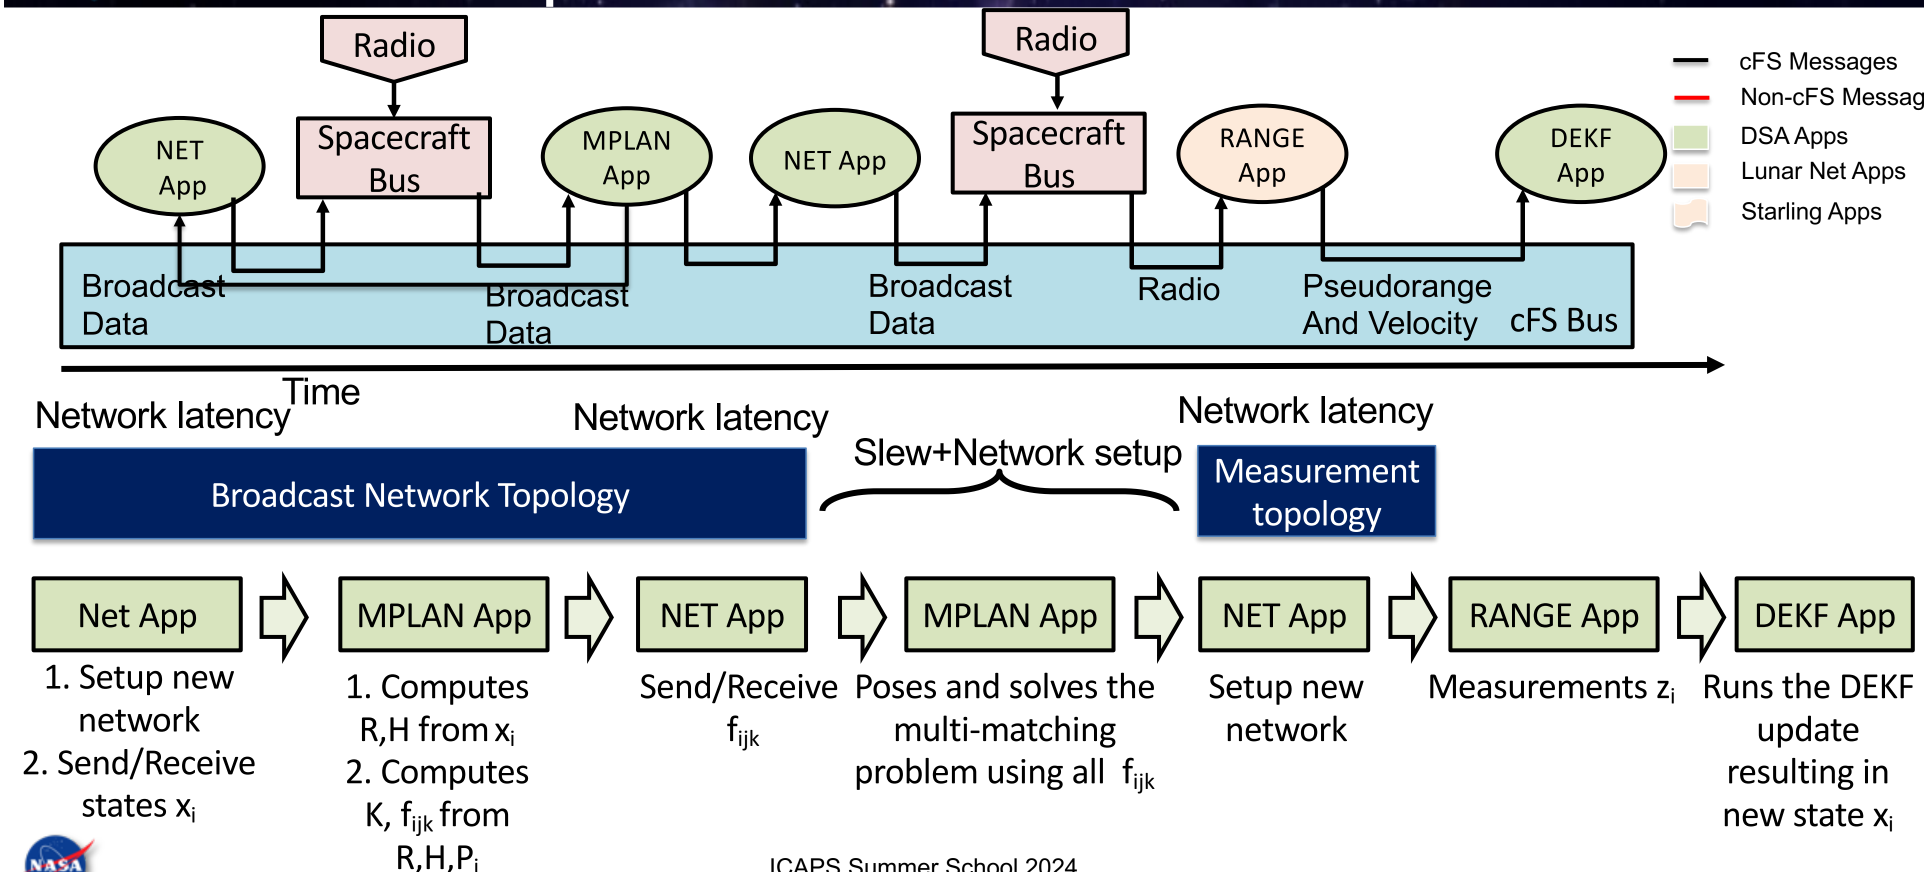

DSA-Starling

Application	Purpose	Feature
Total Electron Count (TEC) Application	Provide a platform to calculate total electron content from local GPS data and provide explorative & exploitative reward values to the autonomy app.	•Process GPS data from inputs: GPS_Positions, RANGE to capture: •Position GPS Satellites in view •Total Electron Content in view •Range of GPS satellites •Produce explorative reward values, which define the range of observability for the satellites, as well as exploitative reward values, which focus on maximizing the observance potential of regions with prolonged or elevated TEC density.
Autonomy (AUTO) Application	Manage and allocate resources autonomously to maximize the observation of TEC features in the ionosphere	•Manage swarm membership and collect observation of GPS satellites in view •Generate a set of GPS channel assignments each node in the swarm based on explorative and exploitative reward values from DSA TEC Feature App •Optimize GPS channel assignments based on other nodes' observations
Communication (COMM) Application	The DSA swarm autonomy demonstration is involving swarm communication of the entire swarm	Data Distribution Service (DDS) implementation by RTI for pub/sub framework • Middleware has flight heritage • DDS fulfills all functional requirements and supports most of the non-functional requirements. • Publish/Subscribe reliable communication schema for this configuration.