

Cognitive Anti-jamming Satellite-to-Ground Communications on NASA's ScaN Testbed



Sudharman K. Jayaweera, Shuang Feng,
Abriel Holland, and Christos Christodoulou.



BLUECOM Systems and Consulting LLC, Albuquerque, NM.
ECE Department, University of New Mexico, Albuquerque, NM.

Work performed under NASA STTR contract NNX17CC01C

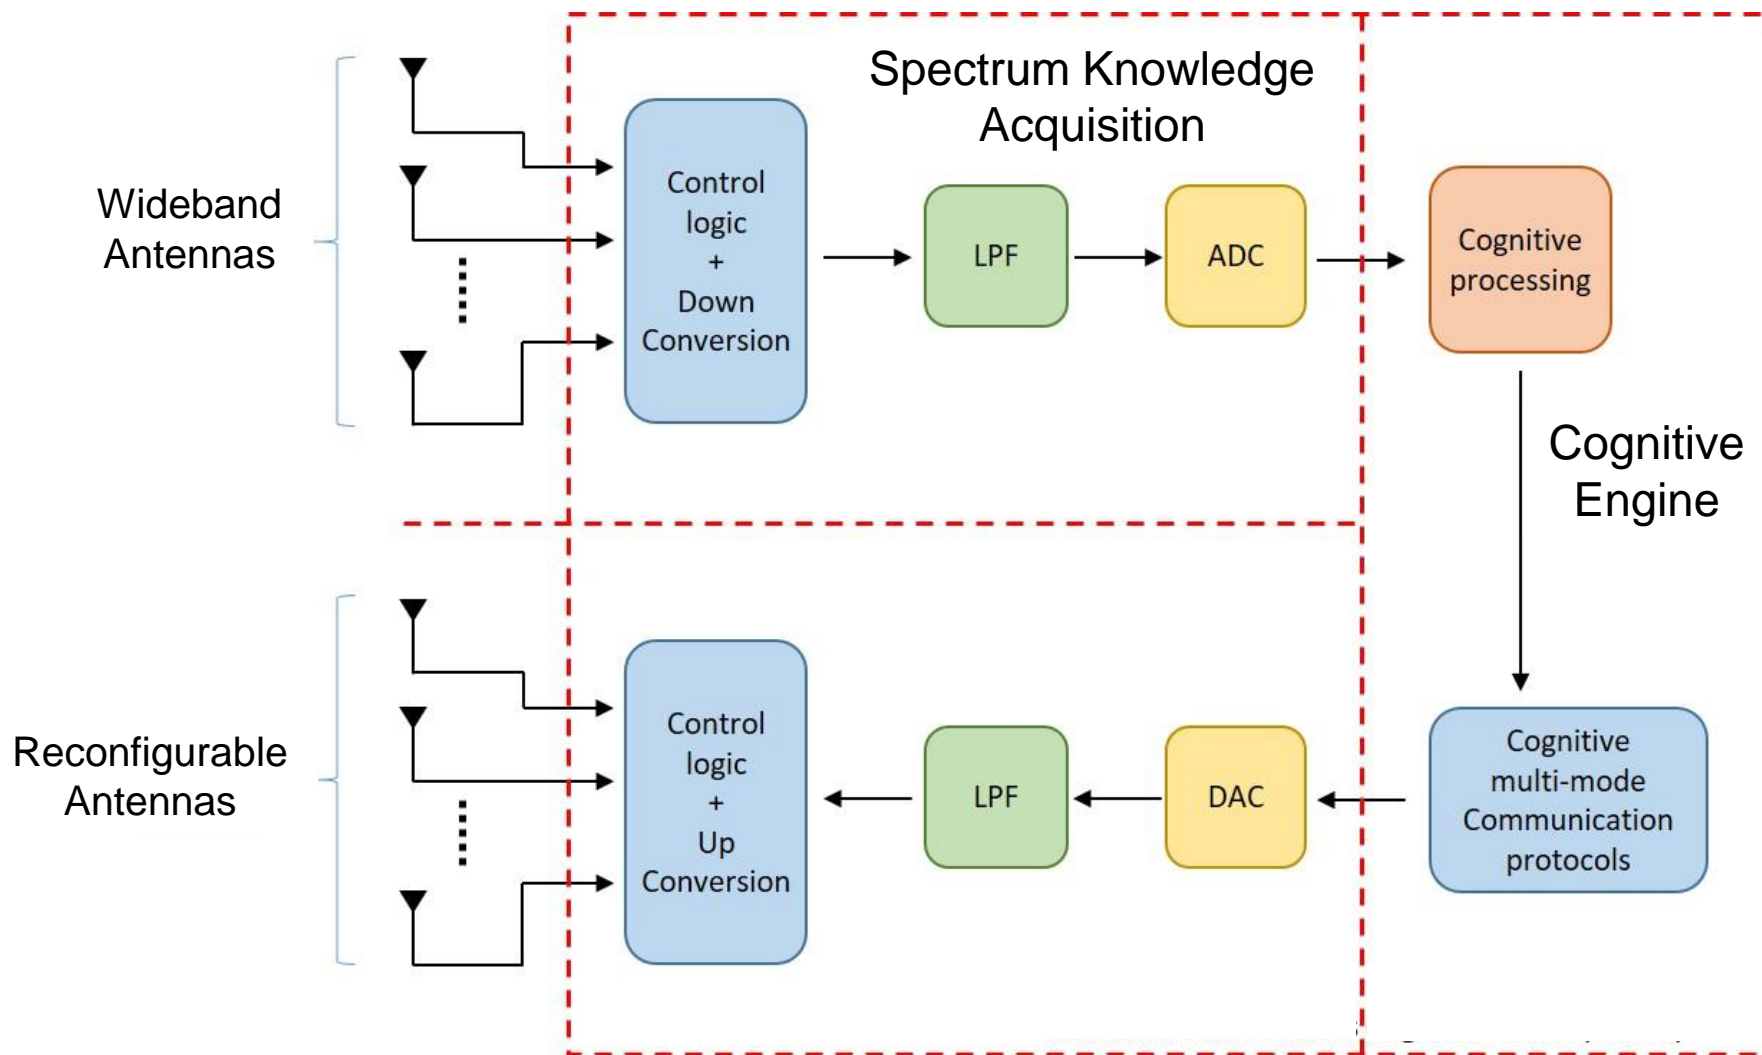
Dale Mortensen, Marie Piasecki, and Mike Evans



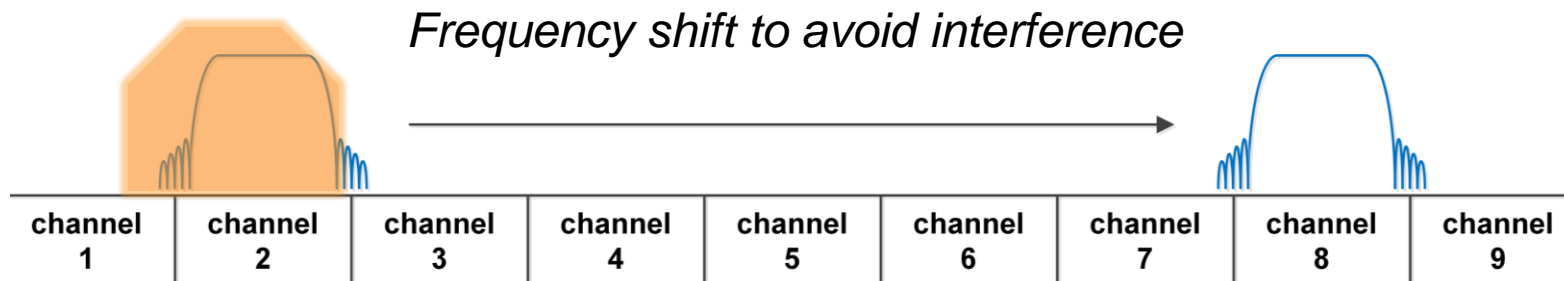
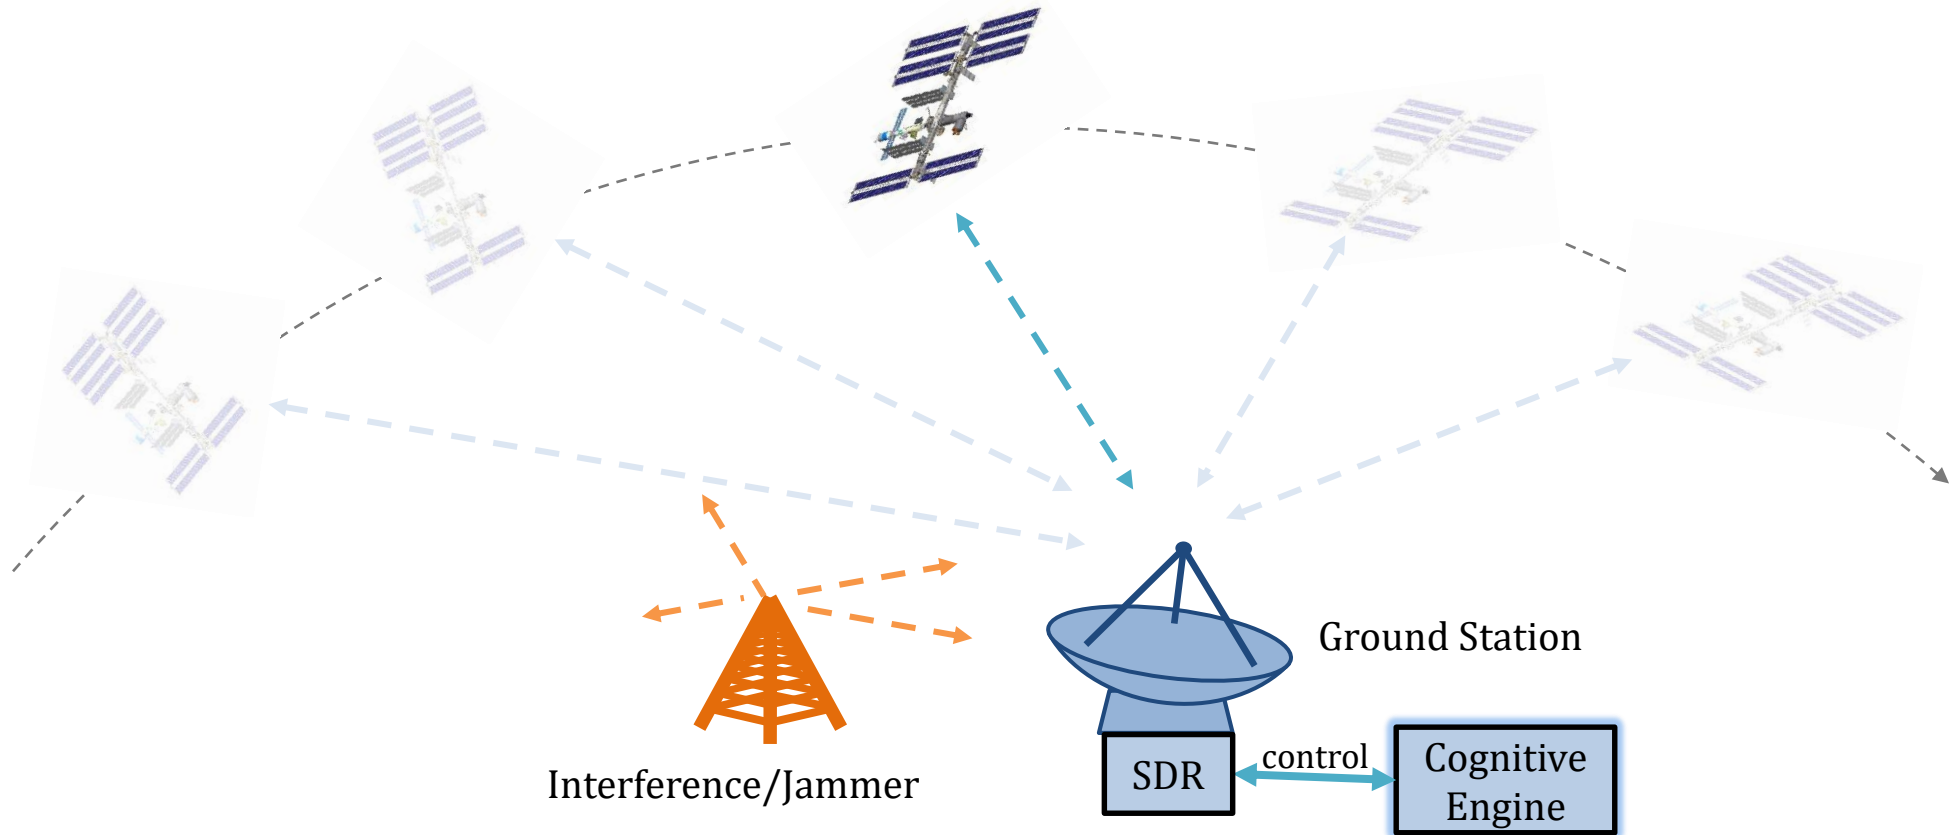
NASA Glenn Research Center, Cleveland, OH.

Presenter: Dale Mortensen

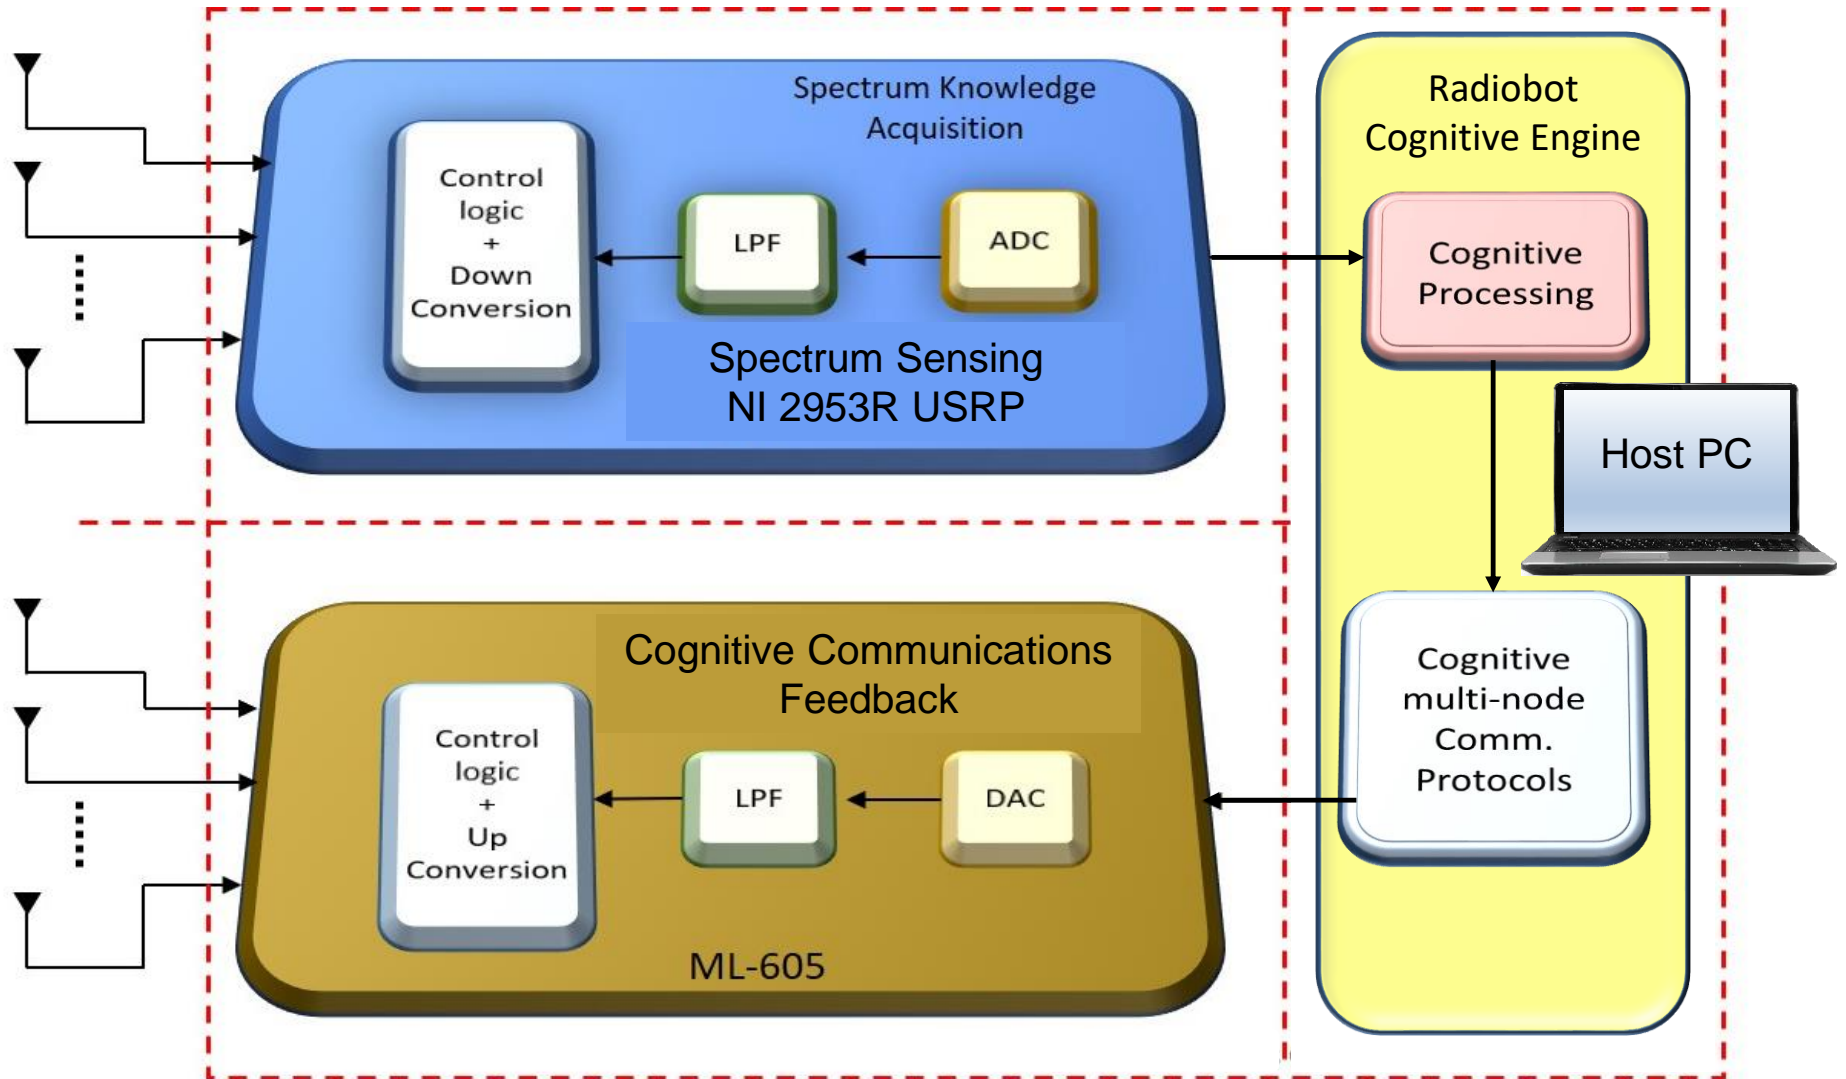
Wideband Autonomous Cognitive Radio (WARC) Architecture



Satellite-to-Ground Cognitive Anti-Jamming (CAJ) Communications: Concept of Operations



Implemented WACR System



WACR operation with two separate SDR modules instead of a single SDR module.

Radiobot Cognitive Engine: CAJ Policy Options

1. Load a pre-learned policy from a file and keep updating the policy during the communications phase.
2. Learn a policy during a training period and keep updating the policy during the communications phase.
3. Learn a policy during a training period and keep it fixed during the communications phase.

CAJ Policy with Reinforcement Learning: Watkin's Q-Learning Algorithm

$$Q(s, a) = (1 - \alpha)Q(s, a) + \alpha \left(r + \gamma \max_{a'} Q(a, a') \right)$$

NOTE: Learning rate (α) and Forgetting factor (γ) both held constant for this experiment.

Exploration vs Exploitation

$$\pi_t(s) = \mathop{\text{arg max}}_a Q_t(s, a)$$

$$a_t(s) = \begin{cases} \pi_t(s) & \text{with probability } 1 - \varepsilon \\ U(A \setminus \{s\}) & \text{with probability } \varepsilon \end{cases}$$

- Learned policies can be used with an exploration rate (ε) during the communications phase.
 - Allows discovery of possible new optimal (state, action) pairs.
 - Must be balanced with exploitation of the already learned policy.
- Complete exploitation of previously learned policy is obtained setting ε to zero.

CAJ with a Random Policy

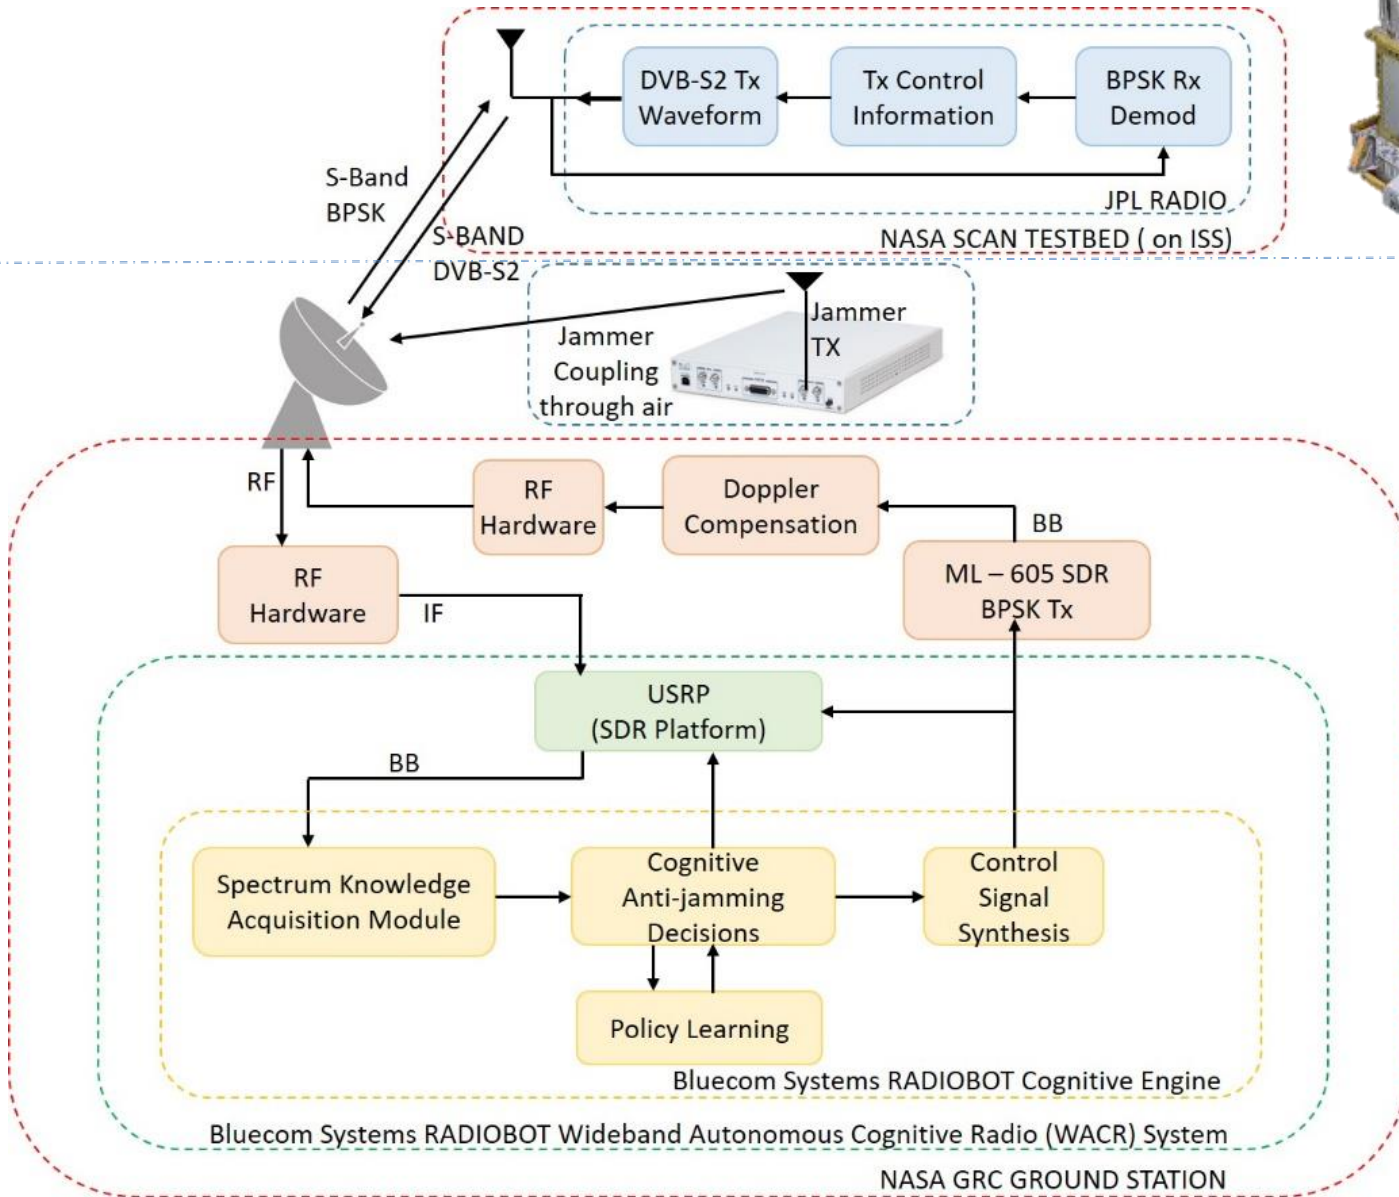
- Set the exploration rate to unity during communications phase to achieve a random channel selection policy.
- Random channel selection policy does not mean it is a traditional radio.
 - Even when the policy is to select channels randomly, the radio is still a WACR.
- Random policy is used to evaluate the effectiveness of the learning process, not the effectiveness of cognitive communications.
 - To perform anti-jamming communications, even with a random policy, the radio still needs the spectrum knowledge of the cognitive radio.
 - Hence, it is still autonomously mitigating the jammer.

Flight Testing System Configuration




FLIGHT SYSTEMS

GROUND SYSTEMS




Flight Testing Ground Station Antenna Setup



Over-the-air jammer antenna setup on same rooftop as main ground station.

Flight Testing: Relative Powers of Satellite and Jammer Signals



Flight Testing Event Data

Test #	Jammer Type	Policy Type	Exploration Rate (ϵ)	Total Number of Sensing Periods During the Complete Event-pass	Total Number of Sensing Periods with Sufficient Signal Quality Between Channel Transitions	Number of Channel Transitions
1	sweep	random	1.0	214710	29380	21
2	sweep	random	1.0	218545	96337	77
3	sweep	pre-learned	0.3	235192	132380	81
4	sweep	pre-learned	0.3	120370	298	4
5	Markov	random	1.0	192751	51412	67
6	Markov	pre-learned	0.0	229520	72908	79
7	Markov	pre-learned	0.3	266661	112660	115

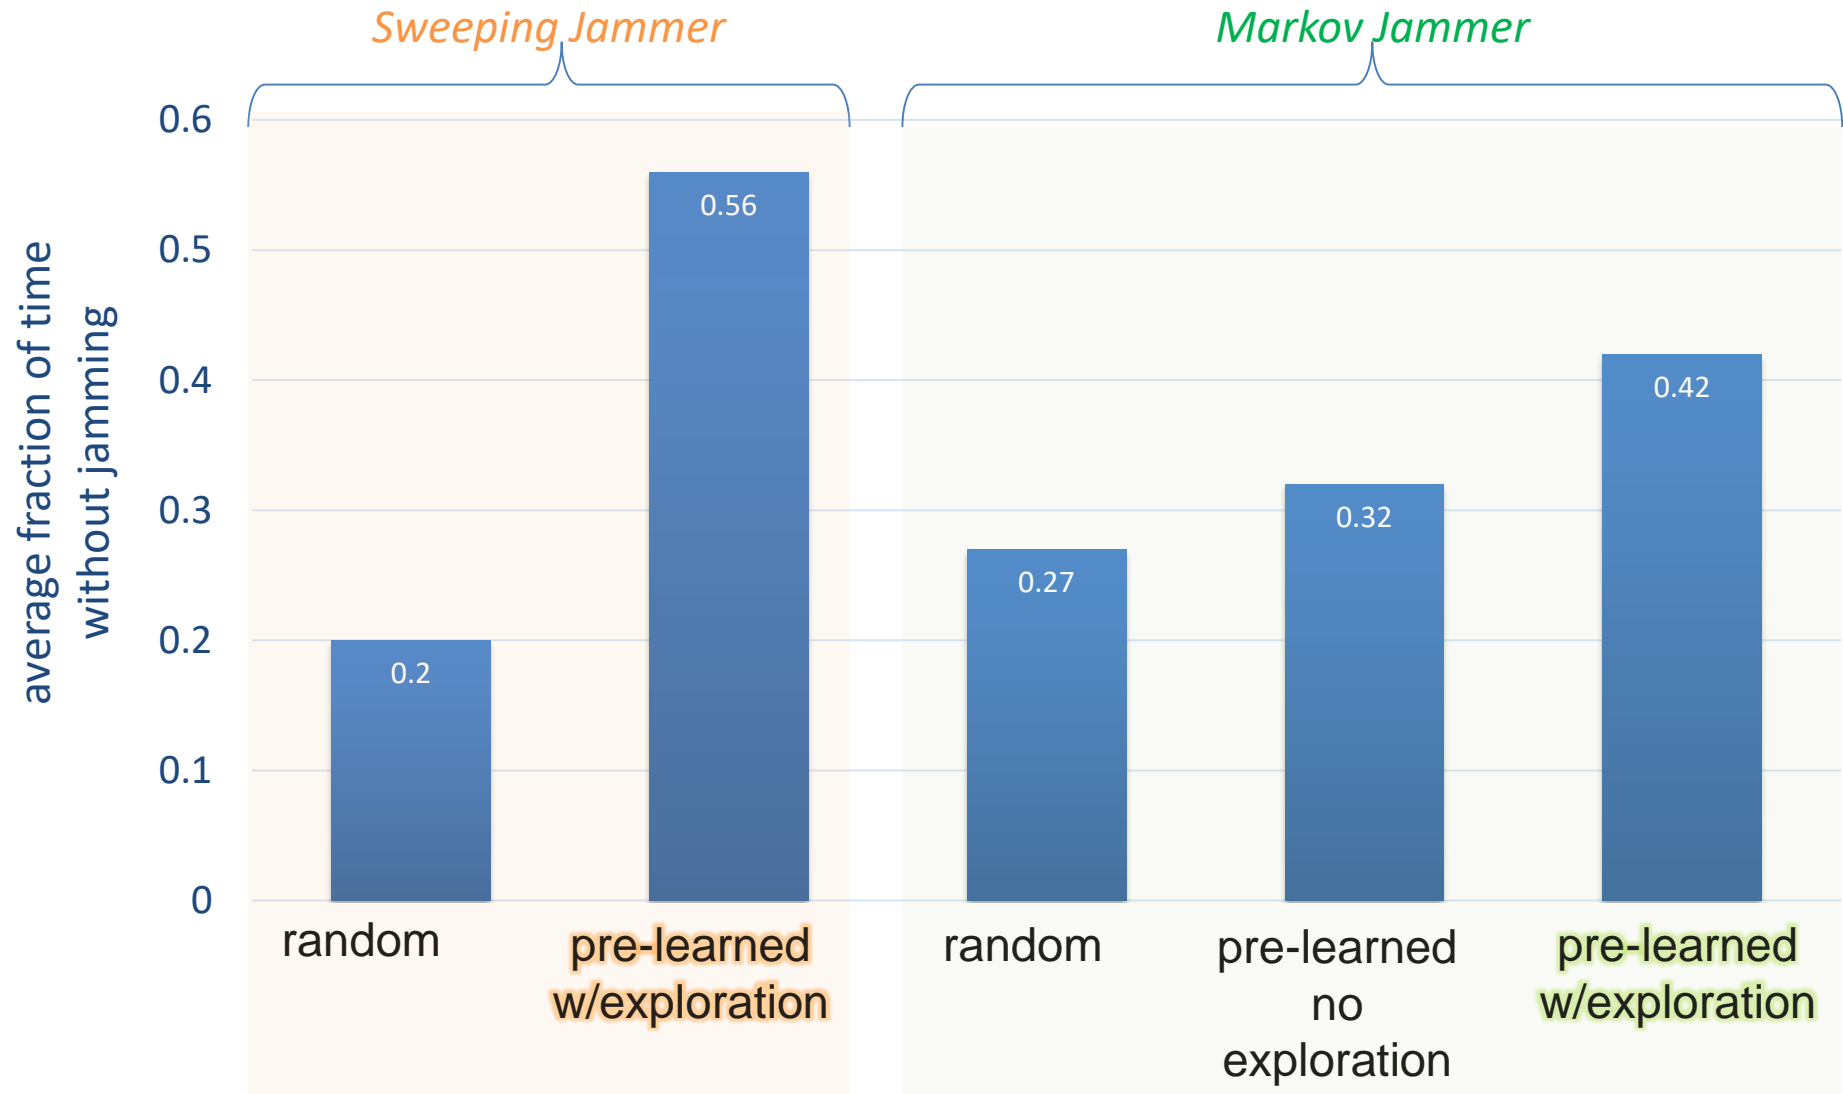
Learning rate (α) set to 0.3, and Forgetting factor (γ) set to 0.8 for all tests.

Flight Testing: Performance

Evaluation of CAJ Communications

Test #	Jammer type	Policy Type	Exploration Rate (ϵ)	Average time in a Channel Without Being Jammed	Average Fraction of time in a Channel Without Being Jammed
1	Sweep	Random	1.0	1399	0.14
2	Sweep	Random	1.0	1251	0.44
1 & 2	Sweep	Random	1.0	1325	0.20
3	Sweep	Pre-learned, continuously updated through exploration	0.3	1634	0.56
5	Markov	Random	1.0	767	0.27
6	Markov	Pre-learned and fixed.	0.0	922	0.32
7	Markov	Pre-learned, continuously updated through exploration	0.3	980	0.42

Flight Testing: Policy vs Random Performance



Conclusions

- Results show that the developed WACR approach is an effective anti-jamming tool, regardless of learning type and channel selection algorithms are used.
- Reinforcement learning aided cognitive anti-jamming communications policy significantly outperforms the random channel selection policy, both in terms of the average unjammed time in a channel as well as the fraction of time in a channel without being jammed.
- Performance is consistent regardless of the type of the jammer: Sweep or Markov.
- Allowing learning-based policy update and policy exploration during actual RF environment will lead to better performance with cognitive anti-jamming communications.
- Best possible performance improvements with the CAJ communications policy can be expected to be higher than what is observed in these tests since these tests only allowed a very short learning period length, and parameters of the algorithms (i.e. learning rate, and forgetting factor, etc.) were unoptimized arbitrary values.

Contact Info:

jayaweera@bluecomsystems.com

505.615.1807

dale.mortensen@nasa.gov

216.433.6823

